Eye movements in dynamic environments

Acknowledgments

I would like to thank my supervisor Prof. Constantin Rothkopf for his scientific support over the course of my PhD. I am greatful for his competent feedback and suggestions during all stages of my dissertation and for him always pushing my research towards high quality. Also, I am very thankful that I was given the freedom to choose interesting research questions to work on. A special thanks to Inge Galinski, our secretary, for her great assistance in all formal matters, for her crocheting stories, and for her endless supply of gummy bears. Of course I also thank my research assistant Stefan Helfmann for his reliability in helping me run my experiments.

Next, I want to thank my colleague and friend Nils Neupärtl for making work so much fun. I enjoyed the exchange of research ideas, discussing the events of the day, and of course the occasional video game during lunch time. Thanks to my family for providing me with the opportunity to receive great education by means of their financial and moral support over all my years at the university. Finally, I want to thank my wife Elena for her love that gave me strength in difficult times.

Abstract

The capabilities of the visual system and the biological mechanisms controlling its active nature are still unequaled by modern technology. Despite the spatial and temporal complexity of our environment, we succeed in tasks that demand extracting relevant information from complex, ambiguous, and noisy sensory data. Dynamically distributing visual attention across multiple targets is an important task. In many situations, for example driving a vehicle, switching focus between several targets (e.g., looking ahead, mirrors, control panels) is needed to succeed. This is further complicated by the fact, that most information gathered during active gaze is highly dynamic (e.g., other vehicles on the street, changes of street direction). Hence, while looking at one of the targets, the uncertainty regarding the others increases. Crucially, we manage to do so despite omnipresent stochastic changes in our surroundings. The mechanisms responsible for how the brain schedules our visual system to access the information we need exactly when we need it are far from understood. In a dynamic world, humans not only have to decide where to look but also when to direct their gaze to potentially informative locations in the visual scene. Our foveated visual apparatus is only capable of gathering information with high resolution within a limited area of the visual field. As a consequence, in a changing environment, we constantly and inevitably lose information about the locations not currently brought into focus.

Little is known about how the timing of eye movements is related to environmental regularities and how gaze strategies are learned. This is due to three main reasons: First, to relate the scheduling of eye movements to stochastic environmental dynamics, we need to have access to those statistics. However, these are usually unknown. Second, to apply the powerful framework of statistical learning theory, we require knowledge of the current goals of the subject. During every-day tasks, the goal structure can be complex, multi-dimensional and is only partially accessible. Third, the computational problem is, in general, intractable. Usually, it involves learning sequences of eye movements rather than a single action from delayed rewards under temporal and spatial uncertainty that is further amplified by dynamic changes in the environment.

In the present thesis, we propose an experimental paradigm specifically designed to target these problems: First, we use simple stimuli with reduced spatial complexity and controlled stochastic behavior. Second, we give subjects explicit task instructions. Finally, the temporal and spatial statistics are designed in a way, that significantly simplifies computation and makes it possible to infer several human properties from the action sequences while still using normative models for behavior. We present results from four different studies that show how this approach can be used to gain insights into the temporal structure of human gaze selection. In a controlled setting in which crucial quantities are known, we show how environmental dynamics are learned and used to control several components of the visual apparatus by properly scheduling the time course of actions.

First, we investigated how endogenous eye blinks are controlled in the presence of non-stationary environmental demands. Eye blinks are linked to dopamine and therefore have been used as a behavioral marker for many internal cognitive processes. Also, they introduce gaps in the stream of visual information. Empirical results had suggested that 1) blinking behavior is affected by the current activity and 2) highly variable between participants. We present a

computational approach that quantifies the relationship between blinking behavior and environmental demands. We show that blinking is the result of a trade-off between task demands and the internal urge to blink in our psychophysical experiment. Crucially, we can predict the temporal dynamics of blinking (i.e., the distribution of interblink intervals) for individual blinking patterns.

Second, we present behavioral data establishing that humans learn to adjust their temporal eye movements efficiently. More time is spent at locations where meaningful events are short and therefore easily missed. Our computational model further shows how several properties of the visual system determine the timing of gaze. We present a Bayesian learner that fully explains how eye movement patterns change due to learning the event statistics. Thus, humans use temporal regularities learned from observations to adjust the scheduling of eye movements in a nearly optimal way. This is a first computational account towards understanding how eye movements are scheduled in natural behavior.

After establishing the connection of temporal eye movement dynamics, reward in the form of task performance, and physiological costs for saccades and endogenous eye blinks, we applied our paradigm to study the variability in temporal eye movement sequences within and across subjects. The experimental design facilitates analyzing the temporal structure of eye movements with full knowledge about the statistics of the environment. Hence, we can quantify the internal beliefs about task-relevant properties and can further study how they contribute to the variability in gaze sequences in combination with physiological costs. Crucially, we developed a visual monitoring task where a subject is confronted with the same stimulus dynamics multiple times while learning effects are kept to a minimum. Hence, we are not only able to compute the variability between subjects but also over trials of the same subject. We present behavioral data and results from our computational model showing how variability of eye movement sequences is related to task properties. Having access to the subjects' reward structure, we are able to show how expected rewards influence the variance in visual behavior.

Finally, we studied the computational properties underlying the control of eye movement sequences in a visual search task. In particular, we investigated whether eye movements are planned. Research from psychology has merely revealed that sequences of multiple eye movements are jointly prepared as a scanpath. Here we examine whether humans are capable of finding the optimal scanpath even if it requires incorporating more than just the next eye movement into the decision. For a visual search task, we derive an ideal observer as well as an ideal planner based on the framework of partially observable Markov decision processes (POMDP). The former always takes the action associated with the maximum immediate reward while the latter maximized the total sum of rewards for the whole action sequence. We show that depending on the search shape ideal planner and ideal observer lead to different scanpaths. Following this paradigm, we found evidence that humans are indeed capable of planning scanpaths. The ideal planner explained our subjects' behavior better compared to the ideal observer. In particular, the location of the first fixation differed depending on the shape and the time available for the search, a characteristic well predicted by the ideal planner but not by the ideal observer. Overall, our results are the first evidence that our visual system is capable of taking into account future consequences beyond the immediate reward for choosing the next fixation target.

In summary, this thesis proposes an experimental paradigm that enables us to study the temporal structure of eye movements in dynamic environments. While approaching this computationally is generally intractable, we reduce the complexity of the stimuli in dimensions that do not contribute to the temporal effects. As a consequence, we can collect eye movement data in tasks with a rich temporal structure while being able to compute the internal beliefs of our

subjects in a way that is not possible for natural stimuli. We present four different studies that show how this paradigm can lead to new insights into several properties of the visual system. Our findings have several implications for future work: First, we established several factors that play a crucial role in the generation of gaze behavior and have to be accounted for when describing the temporal dynamics of eye movements. Second, future models of eye movements should take into account, that delayed rewards can affect behavior. Third, the relationship between behavioral variability and properties of the reward structure are not limited to eye movements. Instead, it is a general prediction by the computational framework. Therefore, future work can use this approach to study the variability of various other actions. Our computational models have applications in state of the art technology. For example, blink rates are already utilized in vigilance systems for drivers. Our computational model is able to describe the temporal statistics of blinking behavior beyond simple blink rates and also accounts for interindividual differences in eye physiology. Using algorithms that can deal with natural images, e.g., deep neural networks, the environmental statistics can be extracted and our models then can be used to predict eye movements in daily situations like driving a vehicle.

Zusammenfassung

Die Leistung des menschlichen visuellen Systems und der zugrunde liegenden Mechanismen sind immer noch unerreicht von moderner Technologie. Unsere Umgebung ist geprägt von komplexen zeitlichen und räumlichen Dynamiken, dennoch sind wir in der Lage aus uneindeutigen und mit Rauschen versehenen sensorischen Daten jene Informationen zu extrahieren, welche uns in die Lage versetzen, schwierige Aufgaben zu meistern. Eine Schlüsselrolle spielt dabei die dynamische Verteilung unserer visuellen Aufmerksamkeit auf mehrere Regionen im Raum. In vielen Situationen, zum Beispiel während des Fahrens eines Autos, ist es wichtig Informationen von einer Vielzahl an informationstragenden Punkten im Auge zu behalten (z. B., Abstand zum Vordermann, überholende Autos im Seitenspiegel, die Einstellungen des Entertainmentsystems). Dies wird zusätzlich dadurch erschwert, dass der Großteil der Informationen zeitlichen Veränderungen unterliegt. So verändert sich beim Fahren zwangsläufig die Umgebung, da wir uns im Auto fortbewegen. Zusätzlich bewegen sich die anderen Verkehrsteilnehmer. Da wir lediglich in einem kleinen Bereich (der Fovea) visuelle Eindrücke mit hoher Auflösung wahrnehmen können, verpassen wir mit jeder Fixation Informationen an allen Regionen, die wir in diesem Moment nicht fokussieren. Trotz der stochastischen Veränderungen in unserer Umgebung zeigen Menschen gute Leistungen in vielen visuellen Aufgaben. Dabei ist es unklar, wie das Gehirn unser visuelles System derart koordiniert, dass wir zum richtigen Zeitpunkt Zugang zu den richtigen Informationen haben. Durch die zeitliche Dynamik in unserer Umwelt müssen wir nicht nur entscheiden, wohin wir unseren Blick richten, sondern auch wann. Es ist unumgänglich, dass wir ständig Informationen verlieren, da wir nicht alles gleichzeitig fokussieren können und sich der Zustand unserer Umgebung ändert.

Viele Details der zeitlichen Steuerung von Augenbewegungen und der Verbindung zu Regelmäßigkeiten in unserer Umgebung sind ungeklärt. Auch existieren nur wenige Erkenntnisse darüber, wie Strategien zur Kontrolle von Augenbewegungen erlernt werden. Dafür gibt es drei Gründe: Erstens brauchen wir Zugang zu den Statistiken unserer Umgebung, um diese mit Augenbewegungen in Verbindung zu bringen. Diese Statistiken sind allerdings im Allgemeinen nicht zugänglich und daher unbekannt. Zweitens sind für die Modellierung des Verhaltens mittels Methoden der statistischen Lerntheorie Informationen über die Zielstruktur der Probandin notwendig. Diese latenten Strukturen sind in der Realität allerdings komplex, vielschichtig und nur teilweise abrufbar. Drittens übersteigt die Komplexität der Berechnungen, welche zur Beschreibung von natürlichem Verhalten nötig sind, die verfügbare Rechenleistung. Gewöhnlich handelt es sich nämlich nicht um das Erlernen einzelner Augenbewegungen, sondern um Sequenzen von Augenbewegungen, welche in Gegenwart zeitlicher und räumlicher Unsicherheit und dynamischer Veränderungen der Umgebung aus verzögerten Belohnungen abgeleitet werden müssen.

In der vorliegenden Arbeit stellen wir einen experimentellen Ansatz vor, welcher speziell zur Lösung dieser Problematik entwickelt wurde: Erstens nutzen wir Stimuli mit reduzierter räumlicher Komplexität und kontrollieren deren stochastisches Verhalten. Zweitens verwenden wir Aufgaben, bei denen wir Zugang zu der Belohnungsstruktur haben. Dies wird durch geeignete Instruktionen sichergestellt. Zuletzt wählen wir die zeitlichen und räumlichen Statistiken auf

eine Weise, sodass sich die Modellberechnungen signifikant vereinfachen und somit das Inferieren von Eigenschaften der menschlichen Informationsverarbeitung aus Handlungssequenzen auch mit normativen Modellen möglich wird. Wir präsentieren Ergebnisse aus vier verschiedenen Studien, welche zeigen, wie dieser Ansatz genutzt werden kann, um Einblicke in die zeitliche Koordination von Augenbewegungen zu erhalten. Für kontrollierte Umgebungen, in denen für das Verhalten relevante Größen zugänglich sind, zeigen wir wie Umgebungsdynamiken gelernt und zur Kontrolle verschiedener Komponenten des visuellen Apparats genutzt werden.

Zunächst haben wir untersucht, ob der Lidschluss von den Dynamiken und Anforderungen einer sich stetig verändernden Umgebung beeinflusst werden. Eine Verbindung zwischen der Häufigkeit des Lidschlusses und dem Neurotransmitter Dopamin gilt als erwiesen, daher stellt der Lidschluss eine beobachtbare Verhaltensweise dar, welche Rückschlüsse auf viele nicht sichtbare interne Prozesse erlaubt. Zudem führt der Lidschluss zu regelmäßigen Lücken in der visuellen Wahrnehmung. Empirische Ergebnisse legen nahe, dass 1) das Lidschlussverhalten von der aktuellen Tätigkeit beeinflusst wird und 2) dass eine hohe interindividuelle Variabilität besteht. Mittels eines computationalen Modells konnten wir die Verbindung zwischen dem Lidschlussverhalten und den Anforderungen der Umgebung quantifizieren. In einem Wahrnehmungsexperiment konnten wir zeigen, dass Häufigkeit des Lidschlusses Folge eines Trade-Off zwischen Erfordernissen der Aufgabe und dem Drang zu blinzeln ist. Erstaunlich ist dabei, dass wir in der Lage sind die Verteilung der Zeiten zwischen zwei Lidschlüssen für einzelne Personen vorherzusagen.

Im Anschluss präsentieren wir Verhaltensdaten, welche eine effiziente Anpassung der zeitlichen Abfolge von Augenbewegungen beim Menschen belegen. Regionen werden länger fokussiert, wenn bedeutungsvolle Ereignisse in diesen Regionen nur von kurzer Dauer sind und daher leicht verpasst werden können. Unser mathematisches Modell zeigt darüber hinaus wie Eigenschaften des visuellen Systems das Timing von Augenbewegungen leiten. Wir präsentieren einen Bayesianischen Learner, der die Veränderungen in den Augenbewegungsstrategien auf das Erlernen der Ereignisstatistiken zurückführt. Dadurch können wir zeigen, dass Menschen zeitliche Regelmäßigkeiten, erlernt über sensorische Beobachtungen, nutzen, um beinahe in optimaler Weise Augenbewegunsstrategien anzupassen. Diese Ergebnisse sind ein erster Schritt zu einem tieferen Verständnis von Augenbewegungen in natürlichem Verhalten.

Nachdem wir die Verbindung zwischen zeitlichen Dynamiken von Augenbewegungen, der Belohnungsstruktur der zu erledigenden Aufgabe und physiologischen Kosten für Sakkaden und Lidschlüsse nachgewiesen hatten, haben wir unser experimentelles Paradigma angewendet, um die interindividuelle Variabilität von Augenbewegungssequenzen zu erforschen. Das experimentelle Design ermöglicht die zeitliche Struktur von Augenbewegungen zu analysieren, während detaillierte Informationen bezüglich der Statistiken der Umgebung verfügbar sind. Mit dieser Grundlage können wir die internen Vorstellungen der Probanden über aufgabenrelevante Größen quantifizieren. Weiter können wir untersuchen, wie sie in Kombination mit physiologischen Kosten zur Variabilität von visuellen Verhaltensweisen beitragen. Wir entwickelten eine Aufgabe bei der die Probandin mehrere Regionen mithilfe von geeigneten Augenbewegungen im Auge behalten muss. Durch geeignete Manipulation der Stimuli konnten wir dieselbe Sequenz mehrfach präsentieren, während Lerneffekte so gering wie möglich gehalten wurden. Dies ermöglicht eine Quantifizierung sowohl der inter- wie auch der intraindividuellen Variabilität. Wir präsentieren Verhaltensdaten und Ergebnisse von unserem Modell, welche aufzeigen, wie Variabilität von Verhalten mit Eigenschaften der Aufgabe verbunden ist. Insbesondere sind wir in der Lage eine Verbindung zwischen der erwarteten Belohnung einer Entscheidung und der Variabilität in der Entscheidung herzustellen.

Abschließend untersuchten wir die computationalen Eigenschaften, welche Sequenzen von Augenbewegungen während visueller Suche unterliegen. Insbesondere untersuchten wir ob Menschen in der Lage sind Augenbewegungen zu planen. Forschungsergebnisse aus der Psychologie legen nahe, dass mehrere Augenbewegungen gemeinsam in Form eines Scanpaths vorbereitet, bzw. programmiert, werden. Unsere Fragestellung war, ob Menschen fähig sind die optimale Sequenz von Augenbewegungen auszuführen, auch wenn dies erfordert, mehr als die nächste Fixation in die Entscheidung einzubeziehen. Für eine Aufgabe aus dem Bereich der visuellen Suche leiteten wir einen Ideal Observer und einen Ideal Planner basierend auf dem Framework der belief MDPs her. Der Ideal Observer wählt jene Augenbewegung, welche zu maximaler sofortiger Belohnung direkt im Anschluss an die Augenbewegung führt. Im Gegensatz dazu maximiert der Ideal Planner die Gesamtsumme aller Belohnungen über alle Augenbewegungen hinweg. Wir konnten zeigen, dass Ideal Observer und Ideal Planner zu unterschiedlichen Verhaltenssequenzen führen, dies aber zusätzlich von der Form des Suchfeldes abhängt. Mithilfe dieses Paradigmas konnten wir die Fähigkeit einen Scanpath zu planen beim Menschen nachweisen. Der Ideal Planner lieferte eine weit bessere Erklärung für die erhobenen Daten unserer Probanden als der Ideal Observer. Insbesondere hing die Landeposition der ersten Augenbewegung innerhalb der Suchsequenz von der Zeit, die zur Suche zur Verfügung stand, ab. Dieser Effekt ist im Einklang mit den Vorhersagen des Ideal Planners, nicht jedoch des Ideal Observers. Insgesamt stellen unsere Ergebnisse die erste Evidenz dafür dar, dass unser visuelles System in der Lage ist, mehr als unmittelbare Konsequenzen in die Entscheidung für die nächste Augenbewegung mit einzubeziehen.

Zusammengefasst stellt die vorliegende Arbeit ein experimentelles Paradigma vor, welches die quantitative Erforschung der zeitlichen Struktur von Augenbewegungen in dynamischen Umgebungen ermöglicht. Während eine computationale Beschreibung für den allgemeinen Fall nicht möglich ist, haben wir die Komplexität in für die Untersuchung zweitrangigen Bereichen reduziert. Durch dieses Vorgehen konnten wir Daten über Augenbewegungen in Aufgaben mit komplexer zeitlicher Struktur sammeln und trotzdem die für die Modellierung der internen Vorstellungen der Probanden notwendigen Größen einbeziehen. Insgesamt stellen wir vier Studien vor, welche aufzeigen, wie das Paradigma zu neuen Erkenntnissen über zahlreiche Eigenschaften der visuellen Informationsverarbeitung führen kann. Unsere Ergebnisse haben klare Auswirkungen auf zukünftige Forschungsarbeiten: Erstens haben wir Faktoren ermittelt, welche bei der Generierung von visuellem Verhalten eine tragende Rolle spielen. Diese müssen für die Beschreibung der zeitlichen Folge von Augenbewegungen in die Betrachtung mit einbezogen werden. Zweitens sollten zukünftige Modelle für Augenbewegungen berücksichtigen, dass auch Belohnungen über die unmittelbare Belohnung einer Handlung hinaus das Verhalten beeinflussen können. Drittens sind die Ergebnisse über den Zusammenhang zwischen Variabilität und den Eigenschaften der Belohnungsstruktur nicht auf Augenbewegungen beschränkt. Vielmehr handelt es sich um eine allgemeine Vorhersage des Modells, welche auf andere Bereiche übertragen werden kann. Zukünftige Arbeiten können demnach den Ansatz nutzen, um Variabilität in anderen Verhaltensmodalitäten zu untersuchen.

Unsere Modelle sind außerdem relevant für zahlreiche technologische Anwendungen. Der Lidschluss, zum Beispiel, wird bereits in Systemen zur Erfassung von Aufmerksamkeit und Wachheit im Rahmen des Straßenverkehrs verwendet. Das von uns entwickelte Modell ist in der Lage die zeitlichen statistischen Kennwerte des Blinzelns zu beschreiben und dabei insbesondere physiologische Unterschiede zwischen Personen zu berücksichtigen. In Verbindung mit modernen Algorithmen für komplexes hochdimensionales Datenmaterial wie Bilder und Videos, zum Beispiel tiefe neuronale Netze, können die Statistiken der Umgebung abgeleitet werden.

Auf diese Weise ist eine Anwendung der im Rahmen dieser Arbeit entwickelten Modelle auf alltägliche Problemstellungen möglich.

Contents

List of Figures

List of Tables

Introduction

Humans have adapted to their environment over the course of evolution (Darwin, 1859). Crucially, we have learned to perform suitable actions depending on the needs of the situation - for example, gathering food, fleeing from predators, and avoiding dangerous or life-threatening situations. In order to do so, the human brain was designed to process sensory information and to choose suitable actions based on these percepts. While reasoning about human behavior has a long tradition (see Hergenhahn & Henley, 2013, for an overview), we are far from understanding the mechanisms that generate behavior. In the beginning of quantitative experimental psychology in the late 19th century, human behavior was described by a set of laws like the Gestalt laws. However, over the last hundred years few laws have endured empirical and theoretical advances in the field (Halvor Teigen, 2002). Where does the complexity originate that makes finding general explanation for why we act the way we do so difficult?

1.1 Eye movements in dynamic environments

We live in a rich and complex environment and eye movements are deeply intertwined with goals and actions (e.g., Rothkopf, Ballard, & Hayhoe, 2007; Henderson, 2003; Hayhoe & Ballard, 2014). We heavily rely on visual input and in many situations visual information is vital for survival or success in a task, for example in the mundane example of navigating through a crowded city in order to safely and quickly reach a destination (see Figure 1.1A). Visual information plays a crucial role in many of the associated subtasks, which have been investigated for static environments: To find directions to the destination, we must look for cues like street names (visual search, e.g., Najemnik and Geisler (2005) or Wolfe (1998)). Once we have determined our goal, we walk along a path while avoiding obstacles (locomotion, e.g., Rothkopf et al. (2007) or Warren Jr, Kay, Zosh, Duchon, and Sahuc (2001)). In order to cross the street, we might have to activate the traffic lights by pressing a button (visuo-motor hand control, Trommershäuser, Landy, and Maloney (2006) or Körding and Wolpert (2006)). When the traffic light turns green (event detection), we can safely cross the intersection. It is advantageous to monitor other road users to avoid collisions as well as to be able to react to unpredictable events (monitoring).

In an ever-changing environment, we need to constantly use new sensory information to monitor our surroundings to avoid missing crucial events. However, the fraction of the visual environment that can be perceived at a given moment is limited by the placement of the eyes and the arrangement of the receptor cells within the eyes (Land & Nilsson, 2002). Thus, continuously monitoring environmental locations, even when we know which regions in space contain relevant information, is unfeasible. Instead, we actively explore by targeting the visual apparatus towards regions of interest using proper movements of the eyes, head, and body (Yarbus, 1967; Findlay & Gilchrist, 2003; Hayhoe & Ballard, 2005; Land & Tatler, 2009). This constitutes a fundamental
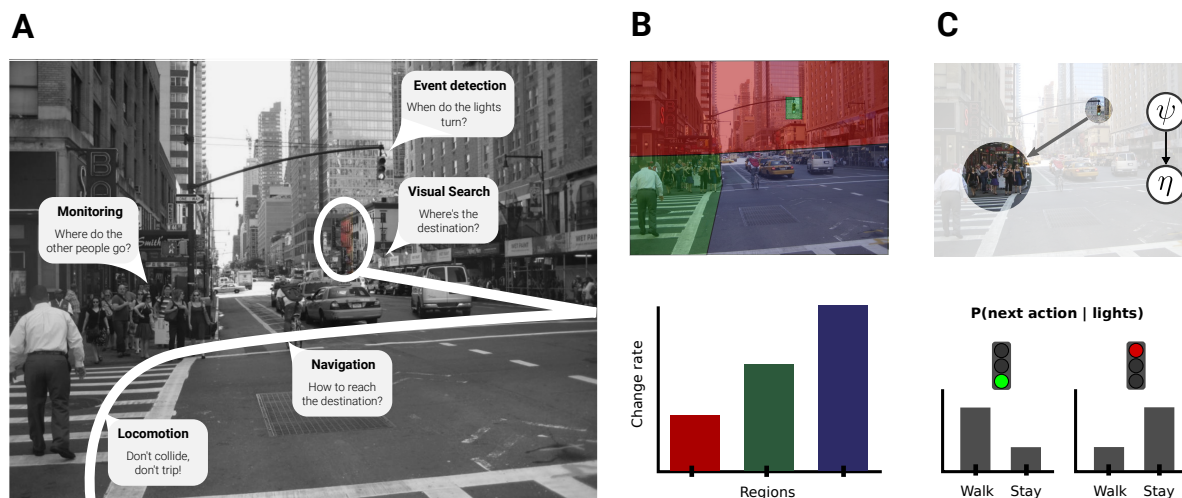
Figure 1.1: Properties of visual tasks in natural environments. (A) Subtasks involved in the everyday scenario of walking a crowded street to get to a destination. (B) Locations in the scene have different environmental dynamics (top). The dynamics can differ with respect to how quickly the information in the scene changes (bottom). (C) Locations can be dependant, i.e., the future state of one location can depend on the state of another location. Clearly, the state of the traffic light informs us about the probability of the next actions of the pedestrians.

computational problem requiring humans to decide sequentially when to look where. Solving this problem arguably has been crucial to our survival from the early human hunters pursuing a herd of prey and avoiding predators to the modern human navigating a crowded sidewalk and crossing a busy road.

Extensive research has been conducted to study eye movements for tasks in static environments. However, event detection and temporal monitoring only arise in dynamic environments. In the present thesis, we primarily focus on how eye movements are scheduled in these environments and the various complicating factors that accompany dynamic changes: In a dynamic world humans not only have to decide where to look, but the timing of gaze to potentially informative locations in the visual scene is crucial. For example in the context of a moving pray, the location of relevant visual targets is not constant, instead, the temporal course of the scanpath determines the received information. Thus, the "when" and the "where" become entangled. This is further complicated by the fact that locations in our environment change at different rates (Figure 1.1B) and are highly correlated (Figure 1.1C).

Little is known about how the timing of eye movements is related to environmental regularities and how gaze strategies are learned. This is due to the complexity of natural behavior. The true environmental state is only accessible through high dimensional and ambiguous visual input, which itself is an ongoing research area. Also, our receptor cells are not distributed uniformly but are most dense at the fovea, further complicating the extraction of information from the raw visual data. In addition, the spatial and crucially, the temporal statistics are unknown and mutually depend on each other.

Here, we investigate the following questions:

- How can we construct stimuli suited to study temporal eye movement patterns while preserving the temporal complexity of real-world problems?

- Do environmental regularities guide visual behavior (saccades, fixations, blinks)?

- How does uncertainty about the state of the environment, action variability, and physiological costs affect the temporal course of eye movements?

- What role do delayed rewards play in tasks involving eye movement sequences?

- Can we quantify the factors causing variability in sequences of eye movements?

1.2   Overview of the thesis

The present thesis is structured as follows: First, influential empirical and computational results related to eye movement strategies are outlined (Chapter 2). The relationship between eye movements and attention is discussed (Chapter 2.1). A long line of research has reported a close connection between where we target our visual apparatus to and what visual information is processed and used for action selection. Gaze behavior has been shown to be a valid behavioral marker for internal processes thus justifying using eye tracking to infer cognitive states. When describing sequences of eye movements rather than isolated gaze targets in an environment that itself has a rich dynamic the perception of time is omnipresent. Therefore, we summarize the findings concerning how humans deal with time, in particular, how the brain is able to estimate and reproduce time intervals (Chapter 2.2). The uncertainty introduced by variability in time interval estimation has severe implications on gaze strategies in a dynamic environment. In Chapter 2.3 we outline findings that suggest that humans are capable of finding near-optimal action sequences despite this variability in action execution by accounting for it when selecting appropriate actions. This is important as fixation durations, a critical factor for vision in dynamic environments, have been shown to be variable. Subsequently, we outline how physiological and psychological costs influence behavioral choices (Chapter 2.4) and give a brief discussion on how the brain could implement the various computations involved in properly controlling the visual apparatus and finding suitable action sequences (Chapter 2.5). Finally, we present related work showing how eye movements are used beyond the purpose of gathering visual information (Chapter 2.6).

Next, we describe the framework of partially observable Markov decision processes (POMDPs), a computational approach which all of the modeling in this thesis is built on (Chapter 3). We show how different models for cognitive processes and visual behavior can be derived from POMDPs. Our presentation is restricted to the mathematical principles underlying the computational methods used later in this thesis.

In Chapter 4, computational approaches that have been used to describe human eye movement behavior are reviewed. Also, we describe the experimental designs and task structures used to study where humans look. We outline the differences between tasks and show their limits with respect to the conclusions that can be drawn from them. Finally, we discuss which properties of experimental designs are advantageous for studying the temporal dynamics of eye movements.

In Chapter 5 we use a controlled detection experiment with parametrically generated event statistics to investigate human blinking control. Subjects were able to learn environmental regularities and adapted their blinking behavior to detect future events better. Crucially, our design enables us to develop a computational model that allows quantifying the consequence of blinking in terms of task performance. The model is based on optimal control of blinking by trading off intrinsic costs for blink suppression with task-related costs for missing an event under perceptual uncertainty. Remarkably, this model is not only sufficient to reproduce key

characteristics of the observed blinking behavior but also predicts the well known and diverse distributions of time intervals between blinks, for which an explanation has long been elusive.

In Chapter 6 we present behavioral data establishing that humans learn to adjust their temporal eye movements efficiently. Our computational model shows how established properties of the visual system determine the timing of gaze. A Bayesian learner only incorporating the scalar law of biological timing can fully explain the course of learning these strategies. Thus, humans use temporal regularities learned from observations to adjust the scheduling of eye movements in a nearly optimal way, given their cognitive constraints.

In Chapter 7, we present data from a study investigating what factors contribute to the variability of gaze behavior within as well as between subjects. We developed a visual monitoring task where subjects switched between three locations. At each location a small dot moved within a circular boundary. The goal was to detect when a dot first moves beyond its boundary. Crucially, we generated the movement trajectories for the dots in a way that subjects were presented with the same radial trajectories multiple times while not being aware of this. This way, we were able to repeatedly probe our subjects' visual system using a constant stimulus while avoiding learning effects. Our results show that variability of temporal eye movement sequences is linearly related to the expected reward in the task.

Finally, in Chapter 8 we investigate whether humans are capable of finding the optimal scanpath even if it requires incorporating more than just the next eye movement into the decision. For a visual search task, we derived an ideal observer as well as an ideal planner based on the formalization of the POMDP as a belief MDP. We show that depending on the search shape the ideal planner and ideal observer lead to different scanpaths. This allows us to investigate whether humans are capable of planning eye movements. Following this paradigm, we found evidence that humans are indeed capable of planning scanpaths. Our subjects' behavior was better explained by the ideal planner compared to the ideal observer. In particular, the location of the first fixation differed depending on the shape and the time available for the search, a characteristic well predicted by the ideal planner but not by the ideal observer. Overall, our results are the first evidence that our visual system is capable of planning in computational terms.

1.3   Contributions

This thesis extends our understanding of active perception and the fundamental relationship between sensing and acting through state-of-the-art models for human eye movement behavior by using novel experimental paradigms to study the temporal course of action sequences related to the visual system. All chapters of the present thesis are based on Hoppe and Rothkopf (2016), Hoppe, Helfmann, and Rothkopf (2018), and Hoppe and Rothkopf (2019), and may contain previously published content:

- **Gaze behavior in dynamic environments can only be described if temporal uncertainty, acting variability and the biological law of timing is taken into account (Chapter 6)**.
  In this work, we designed an experiment using stimuli with a simplistic spatial but a rich temporal structure to gain insights into how humans schedule their eye movements in a temporally uncertain environment. Crucially, we developed a computational model that described not only how the temporal regularities present in the environment guide the visual behavior, but also how behavior changes over the course of learning these statistics.

This work was published in:

Hoppe, D., & Rothkopf, C. A. (2016). Learning rational temporal eye movement strategies. *Proceedings of the National Academy of Sciences*, 201601305.

- **Blinking behavior is affected by environmental statistics (Chapter 5).** By using a controlled experiment with known temporal statistics we are able to link blinking behavior to task-related rewards and physiological costs. In particular, this is the first work that derives key properties of the distribution of time intervals between consecutive blinks.

  This work was published in:

  Hoppe, D., Helfmann, S., & Rothkopf, C. A. (2018). Humans quickly learn to blink strategically in response to environmental task demands. *Proceedings of the National Academy of Sciences*, 201714220.

- **Eye movements are planned (Chapter 8).** Here, we tested whether sequences of eye movements are planned, i.e. whether the sequence is chosen to maximize the overall reward, as opposed to performing the eye movement that maximized the immediate reward. Based on the methodology of partially observable Markov decision processes we derived a computational model for a visual search task. Using this model, we derived predictions for an ideal planner (maximizing the total reward) and for an ideal observer (maximizing the immediate reward). Our results question the assumption of greedy behavior that is present in all state of the art models for eye movements.

  This work was published in:

  Hoppe, D., & Rothkopf, C. A. (2019). Multi-step planning of eye movements in visual search. *Scientific reports*, 9(1), 144.

- **Variability in temporal eye movement sequences (Chapter 7).** By using a temporal monitoring task we investigated variability in eye movement sequences. We found a connection between human action variability and properties in the reward structure of the task.

  This work is currently in preparation for submission.

Human eye movement strategies

The work presented in this thesis builds on several influential findings that we review briefly in the following chapter. First, to justify using eye tracking as a method to study how the visual system acquires sensory information, it is necessary to show that the measured gaze traces are related to the information that is processed. Second, we review how the human brain processes temporal information, as the main part of the current thesis investigates vision in the presence of temporally changing environments. Next, we describe related work presenting evidence for the capability of humans to account for sensory uncertainties and variability in the execution of actions. In particular, we list studies showing humans are able to perform nearly optimal even if this requires incorporating behavioral variability. Next, we motivate the approach of inverse reinforcement learning (see Chapter 3, for details on the mathematical methods) used in all experiments presented in the thesis by showing related work pointing to the many-layered reward structure of human actions. Finally, we review current proposals of how the brain may implement Bayesian computations.

2.1   Eye movements and attention

What information the visual system processes at a given time is not directly accessible. Instead, we can measure the current alignment of the visual apparatus using eye tracking. By measuring the orientation of the eyes, we can detect when gaze is targeted at a location (fixation), when gaze is redirected to a new location (saccade), when gaze is following a moving target (smooth pursuit), and when gaze is interrupted (eye blinks). While eye movements are closely related to internal cognitive processes, visual attention can also be enhanced at other locations in the visual field. In an early study, Posner (1980) showed that subjects' reaction times in a detection experiment were smaller at cued locations in the visual field compared to trials without cued locations and trials where a different location was cued (see Figure 2.1A-B). Hence, while the visual apparatus was directed towards the same region (the center of the screen), visual processing was different depending on which location was cued.

Various studies have investigated how visual attention influences performance in a spatial detection task (see Smith & Ratcliff, 2009, for a review of the most important results as well as a computational model). As with reaction time in the detection task, performance in a letter discrimination task is also greatly affected by which location is cued (Hoffman & Subramaniam, 1995). Thereby, the letter discrimination task serves as a behavioral marker for the spatial distribution of visual attention. Studies investigating visual attention have extensively used discrimination tasks for arrays of letters in order to measure performance at various spatial locations (e.g., Irwin, 2011; Hoffman & Subramaniam, 1995; Baldauf & Deubel, 2008b). The found effects remain when using symbolic cues instead of arrows (Hommel, Pratt, Colzato, &
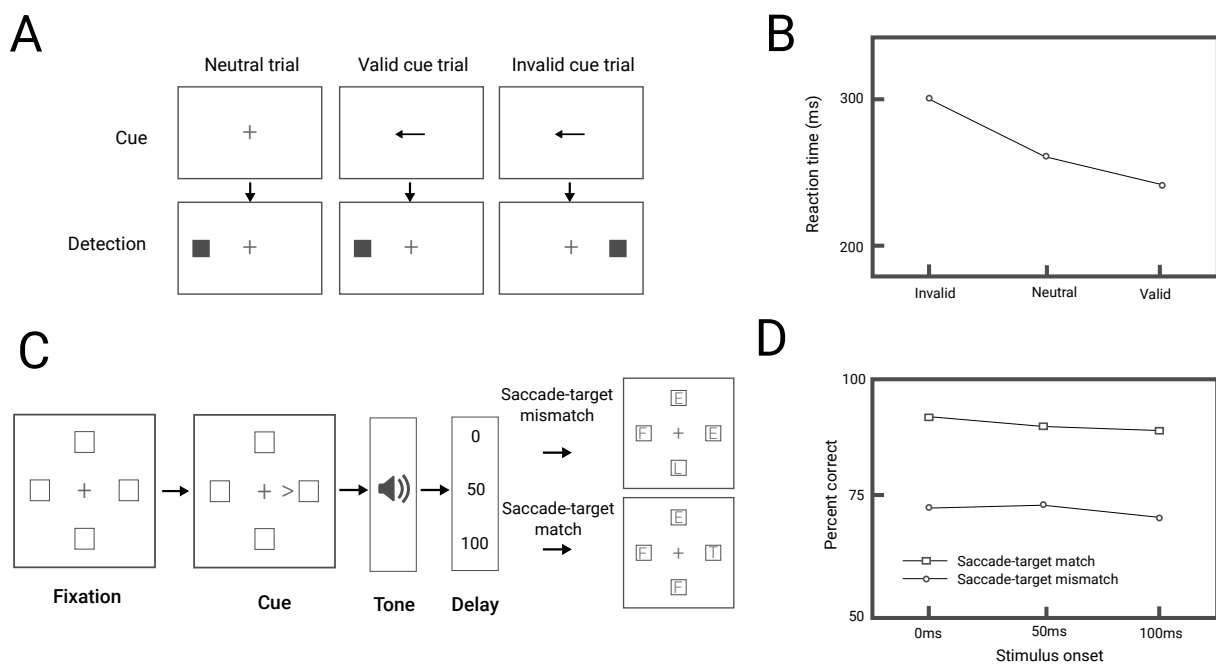
Figure 2.1: (A) Experimental design adapted from Posner (1980). Subjects reacted to an increase of luminance after receiving either a valid cue, an invalid cue, or without any clue. (B) Results from the experimental design presented in A. Subjects responded quicker if the location was cued. (C) Dual task experimental paradigm adapted from Hoffman and Subramaniam (1995). (D) Results illustrating the shift of attention prior to making a saccade.

Godijn, 2001).

While these results point to the complexity of attention, nevertheless, the current target of gaze is an indicator of the internal cognitive state. When making an eye movement, attention has shown to be necessarily targeted at the target location of the movement (Deubel, Shimojo, & Paprotta, 1997). Various studies have reported that saccades are preceded by a shift in covert attention towards the saccadic target (Hoffman & Subramaniam, 1995; Peterson, Kramer, & Irwin, 2004). Using a dual-task paradigm Hoffman and Subramaniam (1995, Figure 2.1C-D) found that discrimination performance was improved if the location of the target stimulus was also the target of a saccade. Four letters were grouped around a central fixation point. Subjects performed a saccade towards one of the four locations cued by an arrow. After the presentation of the cue but before the execution of the saccade three distractors (letters Es and Fs) and one target (T or L) were displayed. Subsequent to the saccade, subjects indicated whether the target letter comprised a T or an L. Performance was improved if the target letter was at the location of the saccadic target. This suggests that visual attention is shifted towards the target of a saccade prior to execution. In a similar study, Godijn and Theeuwes (2003) found that attention precedes locations of future saccades, and visual attention is higher for closer saccadic goals.

Studies suggested that spatial locations in our visual field are divided into movement relevant and movement irrelevant locations. Before executing a sequence of two or three saccades, visual attention at the respective saccadic targets is enhanced (Baldauf & Deubel, 2008a). This enhancement is restricted to the distinct location and does not extend to locations inbetween. Further, visual attention is stronger for locations targeted early in the programmed sequence than for later locations. While they are not functionally equivalent (see Smith & Schenk, 2012,

for a recent review), there is evidence for a coupling between motor preparation and spatial attention (saccades: Hoffman and Subramaniam (1995); blinking: Irwin (2011); hand movements: Baldauf and Deubel (2008b); head direction relative to the body: Nakashima and Kumada (2017)). Finally, in their study, Henderson, Shinkareva, Wang, Luke, and Olejarczyk (2013) could classify what task a person was performing using features derived from gaze patterns.

2.2 Time perception

Time plays a crucial role on many scales in our lives (see Figure 2.2A), especially when dealing with dynamic environments. Time intervals involved in visual tasks range from 30 ms (e.g., deciding where to look next: Stanford, Shankar, Massoglia, Costello, and Salinas (2010)) over 200-300 ms (e.g., fixations: Rayner (2009)) to a couple of seconds (car in traffic, Figure 2.2C). In many situations, it is of great significance to keep an internal belief about the environmental state. This is a direct consequence of the fact that only a small proportion of relevant information can be monitored using our sensory systems. Hence, overwhelming research has demonstrated that we rely on predictions and extrapolation of dynamic events to succeed in many tasks (Miall, Christensen, Cain, & Stanley, 2007; Hayhoe, McKinney, Chajka, & Pelz, 2012; Zago, McIntyre, Senot, & Lacquaniti, 2008). For example, while driving, we need to keep information about the location of other cars as accurate as possible even if not looking at them all the time. In order to do that we need an estimate of the velocity, the last location of the car and, crucially, the time elapsed since the previous measurement. Prior work has shown that accurate models for sensory systems are needed to explain behavior (Schmitt, Bieg, Herman, & Rothkopf, 2017). Hence, the internal estimate about time intervals plays a major role in the internal belief about the environmental state and therefore in building computational models for eye movements.

How humans perceive and estimate time has been a question from the beginning of the field of psychology (e.g., James, 1890). Classical experimental designs for the study of time perception comprise verbal estimation (e.g., -*How long is the following time interval? - Five seconds*), production (e.g., *Please press this button, wait five seconds, and press it again.*), reproduction (e.g., *Please reproduce the following time interval using the buttons*), and the method of comparison (e.g., *Is the first time interval longer than the second?*) (Allan, 1979; Grondin, 2010). It is unclear, which of the paradigms is closest to the role time estimation plays in natural behavior.

On a neural level, the perception of time has been associated with the cerebellum and the basal ganglia (see Ivry, 1996, and the references therein). Two dominant theories about the implementation of temporal information are the global internal clock and a distributed timing mechanism. Recent results suggest that the representation of temporal information is distributed. In their study, Motala, Heron, McGraw, Roach, and Whitaker (2018) found that adaption effects in a temporal reproduction task were only present within a modality but not across modalities. This suggests that processing of temporal information is done separately for different modalities. One of the most influential findings in the field of time perception and particularly relevant to developing computational models for behavior involving temporal dynamics is the scalar law of timing that follows from Weber's law (see Figure 2.2B). Weber's law relates the size of change of a stimulus to the perceived change by

$$dp = k\frac{dS}{S},$$

where dp is the perceived change, dS is the actual change, S is the stimulus size, and k is the
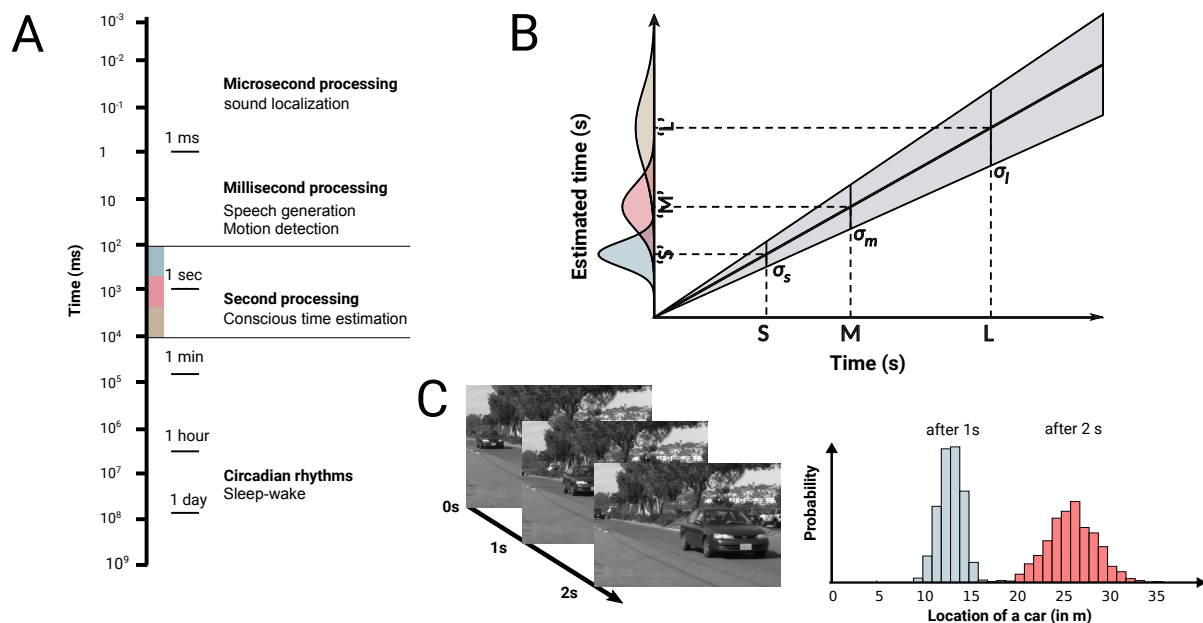
Figure 2.2: (A) Different scales of time that are relevant in our life (adapted from Mauk & Buonomano, 2004). (B) The scalar law of biological timing (adapted from Shi, Church, & Meck, 2013). (C) Effect of the scalar law on estimating the location of a car driving at 50 km/h. Uncertainty about the location is shown after an interval of 1 s (blue) and 2 s (red). The velocity is assumed to be known, hence the uncertainty is only caused by the uncertainty of the time interval.

Weber fraction. The Weber fraction for time estimation has been reported to range between 0.05 and 0.35 and can be assumed to be constant for intervals greater 200 ms (Mauk & Buonomano, 2004). The differential equation can be solved by integrating both sides leading to Fechner's law

$$p = k \ln S + C \,.$$

The assumption that time estimation is unbiased has been questioned despite its use in computational models (Hudson, Maloney, Landy, & Friston, 2008). For example, Verordt's law suggests that short intervals are overestimated while long intervals are underestimated (Eisler, Eisler, & Hellström, 2008).

However, in contrast to psychophysical tasks probing the subject with time intervals, time is often intertwined with complex environmental dynamics and seldomly isolated in more natural behavior (e.g., crossing a street; Figure 2.2C). In their study, Brown and Merchant (2007) compared the quality of produced time intervals in a single task design with a dual task design, where participant concurrently performed a sequencing task. The results showed an increase in temporal variability if an additional task was present. The fact that the variability in temporal judgment is affected by resources used simultaneously for other requirements indicates that uncertainty about time-related dynamics could be higher in complex behavior compared to simple psychophysical experiments.

The estimation of time intervals has been shown to be affected by various properties of the environment leading to numerous illusions (Eagleman, 2008). Crucial to models for temporal gaze allocation, the judgment of time intervals is affected by how fast information changes (see Eagleman, 2008, for an overview) as well as how predictable a stimulus is (Pariyadath & Eagleman, 2007). Even the quality of static visual input can influence judgments (Oliveri et al., 2008). Besides, subjects have been shown to use temporal cues and combine them with other
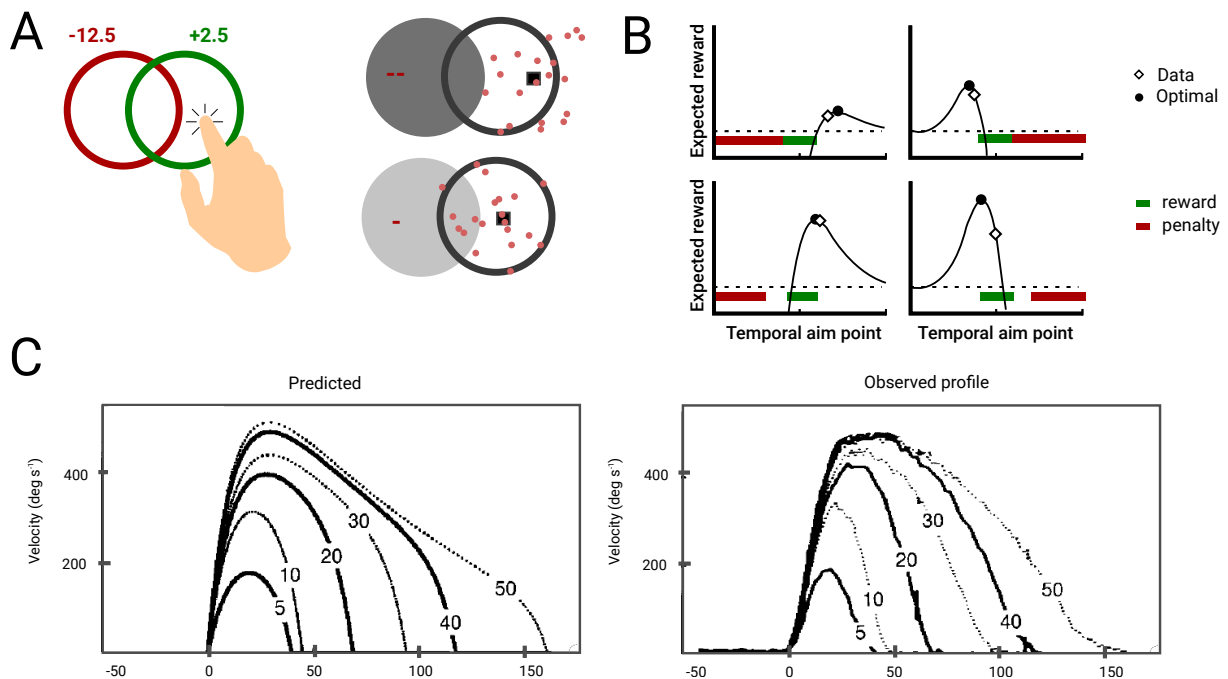
Figure 2.3: (A) Experimental design and results adapted from Trommershäuser, Landy, and Maloney (2006) and Trommershäuser, Maloney, and Landy (2008). The dark gray area is associated with higher penalty than the light gray area (left panel). (B) Results from Hudson, Maloney, Landy, and Friston (2008). Subjects reached with their finger and were rewarded if the duration of the reaching was correct. The reward (green area) was always centered at 650 ms, while the penalized durations were manipulated. (C) Results for the velocity profiles of saccades from Harris and Wolpert (1998).

noisy percepts to improve an internal model of the dynamic environment (Chang & Jazayeri, 2018).

2.3 Action variability

We can execute actions only with limited precision introducing variability to all behavior. Importantly, humans have shown to take into account variability in their actions. In the influential study of Trommershäuser et al. (2006), subjects made ballistic hand movements and tapped on a display with a finger (Figure 2.3A). The display contained two circular areas, one yielding a positive reward, the other a negative reward leading to four different regions. The results showed that dependent on the specific values for the rewards subjects targeted different locations in the areas. If the negative reward was high, the target location chosen by subjects was further away from the intersection yielding more taps to fall outside of any reward shape. Crucially, subjects' choice of endpoints was in accordance with a model based on decision theory that maximizes the expected reward while considering the reaching variability. Hence, the data suggest, that humans are capable of including motor variability when deciding for future actions.

Besides including the endpoint variability in the decisions involved in movement planning, studies have investigated whether humans account for temporal uncertainty when timing movements. Indeed, humans have shown to be able to consider temporal uncertainty (Figure 2.3B; Hudson et al., 2008). In their study, subjects performed nearly optimal in a reaching task that required them to make hand movements with a specific duration.

As for ballistic movements of the hands, humans have shown to account for behavioral variability also when controlling movements of the eyes. In their paramount study, Harris and Wolpert (1998, Figure 2.3C) showed that the velocity profile of human eye movements can be described by a computational model minimizing endpoint variability while adjusting for signal dependent noise.

2.4 Behavioral costs

All actions are associated with costs from contraction of big muscles during walking (e.g., Hall, Figueroa, Fernhall, & Kanaley, 2004) to individual spikes of a single neuron performing a computation (Lennie, 2003). These rewards and costs guide behavior and complex organisms must consider various dimensions when choosing appropriate actions. The sources of costs influencing behavior are manifold and effects can be found on many levels. These respective costs have to be traded off with the expected reward that is gained by the actions, for example, the amount of metabolic energy consumed by a specific behavior. For many species, a connection between behavior and resting metabolic rate has been established (see Biro & Stamps, 2010, for a review). Besides physiological and anatomical costs, cognitive demand has been found to influence what actions humans take. The *law of least mental effort* suggests that actions with minimal cognitive demand are chosen if everything else is constant. Kool, McGuire, Rosen, and Botvinick (2010) let subjects repeatedly choose between two task that were associated with different mental efforts. The overall results showed a preference for the task with the smaller effort.

Also, for hand movements, prior studies have used inverse computational control theory to recover the compounds that contribute to the overall costs of an action (Berret, Chiovetto, Nori, & Pozzo, 2011). The results show that the total costs of a movement can be attributed to unique cost functions like energy, effort, hand jerk, and torque, among others. The multi-dimensionality of reward and cost structures has been investigated extensively in economic choice behavior. For example, a product's value and therefore the expected reward of deciding in favor of that respective product can depend on multiple characteristics (see Gupta & Kim, 2010, and the references therein). Also, there exist various theoretical accounts for the connection between costs and behavior. Inverse reinforcement learning methods have been developed to infer latent cost structures from behavioral data (Ng, Russell, et al., 2000; Dimitrakakis & Rothkopf, 2011). For example, Rothkopf and Ballard (2013) showed how different reward distributions affect the trajectories in a navigation task. Overall, behavior is deeply intertwined with costs and rewards and investigating the detailed composition of these structures is a huge ongoing research area.

2.5 The Bayesian brain

The models developed in the current thesis describe human behavior using partially observable Markov decision processes (POMDPs, for details, see Chapter 3). This corresponds to the computational level of the three levels of analysis proposed by Marr (1982). According to Marr, the *computational level* describes what problem is to be solved, what sensory information is available, and what computations have to be performed to solve the problem. The *algorithmic level* states how the computations are carried out and what representations are used. Finally, the *implementational level* describes, how the algorithm and representations are realized in the real system.

POMDPs are a suitable framework for modeling human actions as many studies have shown that behavior is reward driven (e.g., Navalpakkam, Koch, Rangel, & Perona, 2010) and perceptual processes can be understood in terms of Bayesian inference. For example, human data on how multiple sensory cues are combined into a single percept are in accordance with Bayesian inference (cue combination; Ernst & Banks, 2002; Chang & Jazayeri, 2018). Also, humans have shown to be able to exploit stochastic dependencies in our environment to improve sensory measurements (explaining away; Battaglia, Kersten, & Schrater, 2011). Many complex visual tasks have been successfully understood in terms of Bayesian perception, i.e., object perception (Kersten, Mamassian, & Yuille, 2004) or color constancy (Brainard & Freeman, 1997). Also, Bayesian models are regarded as a suitable approach to unfolding the mechanisms behind behavior (Zednik & Jäkel, 2016).

While our work does not claim to explain how our subjects perform any of the computations (i.e., on the algorithmic and implementational level), prior research has proposed how populations of neurons could implement the computations involved in Bayesian inference (see Pouget, Beck, Ma, & Latham, 2013; Knill & Pouget, 2004, and the references, therein). Two approaches have been proposed to incorporate uncertainty in neural computations: representing the distribution parameters and sampling procedures (see Fiser, Berkes, Orbán, & Lengyel, 2010, for a discussion). The former suggests, that neurons represent not only an estimate for a perceptual quantity but a full probability distribution over potential states of an uncertain quantity (population codes Ma, Beck, Latham, & Pouget, 2006). The product of two probability distributions, as necessary for Bayesian inference, can be realized by adding two population codes. In the latter approach, Bayesian computations are realized using sampling. Thereby, the posterior distribution is approximated by repeatedly drawing samples from it. The sampling procedure can explain idiosyncracies of human behavior such as the unpacking effect and the base-rate neglect (Sanborn & Chater, 2016).

2.6 Eye movements as actions

Throughout the entire thesis, eye movements are investigated as actions tied to gathering information while not altering the state of the world. While the key purpose of the visual system is perception, often in the specific context of a task, there are exceptions where visual behavior is not merely a way to sense but has an active impact on the environment itself (Foulsham, 2015). As an example, we can tell another person to close an open door by looking at it, hence we can actively deliver information using eye movements. Many aspects of social interactions are influenced by gaze behavior (see Kleinke, 1986, for example). Also, when working together, eye movements carry information about future actions that are used by collaborators (Khoramshahi, Shukla, Raffard, Bardy, & Billard, 2016). The ability to voluntarily use our eyes as effector is utilized by modern technical applications like gaze-contingent displays as well as eye controlled systems.

In addition, we have learned to use the information that is immanent to another person's gaze. In their study, Kuhn, Tatler, and Cole (2009) recorded eye movements while subjects watched video clips of two versions of a magic trick. The videos differed with respect to the magician's gaze during the critical part of the trick. They found that subjects were more likely to detect the misdirection if gaze was hinting towards essential locations for the misdirection. Eye movements also play a vital role as a communicative device (Senju & Csibra, 2008). Infants were shown videos of an adult gazing at one of two objects. If the adult either made eye contact or spoke to the camera before looking at one the objects, infants were more likely to follow the gaze.

Hence, in the presence of ostensive signals infants' eye movements are guided by adults' gaze targets. Finally, gaze has been suggested to play a role during language acquisition. Yu, Ballard, and Aslin (2005) found better performance in speech-segmentation as well as word-learning if participants were presented with another person's gaze location in addition to a visual scene.

Computational modeling of human behavior

Here, we will summarize the theoretical underpinnings used for our computational models. First, we will give a brief introduction to probability theory. Without any claim to comprehensiveness, we selected the fundamental concepts underlying the computational models presented in subsequent chapters. For a detailed treatment of filtering and acting in partially observable domains see Thrun, Burgard, and Fox (2005). For a comprehensive presentation of Bayesian inference and graphical models see Gelman et al. (2014) and Jordan (2003), respectively. For an introduction to the application of Bayesian inference to computational modeling of cognition and perception see Lee and Wagenmakers (2014). Using probabilistic reasoning to make sound decisions is treated in DeGroot (2005), and for details on Markov Decision Processes, reinforcement learning and dynamic programming see Sutton and Barto (1998). For an initial treatment of inverse reinforcement learning, see Ng, Russell, et al. (2000).

3.1 Computational foundations

3.1.1 Probability theory

In order to build computational models for human behavior, we first need an abstract mathematical representation of the problems that humans are designed to solve. For survival we need to make the right decisions and actions based on sensory information. However, our sensors are not perfect and the environment is stochastic and ambiguous. Therefore, our brain has to be equipped to manage stochasticity and uncertainty. Quantitative models describing human action selection (e.g., decisions in a two alternative choice task), or internal states (e.g., fMRI data) heavily rely on probability theory.

In probability theory, a random variable x captures the uncertainty about a quantity. For example, the outcome of a coin flip x_i is a priori unknown, but we can assign a probability $p(x_i)$ to seeing the event x_i. The probability of an outcome x_i is between 0 and 1, an impossible outcome has probability zero and the probability of either x_k or x_j is the sum of the probabilities of the individual outcomes $p(x_k) + p(x_j)$. A probability distribution $p(x)$ assigns a probability to all outcomes $x_i \in \mathcal{X}$ in the sample space $\mathcal{X}$.

The probability of two random variables x and y is denoted the joint probability $p(x = x_i \wedge y = y_i)$. The joint distribution $p(x, y)$ assigns a probability to each combination of outcomes of x and y. The distribution of the individual random variable $p(x)$ can be obtained from the joint distribution of x and y by marginalization

$$p(x) = \int p(x, y) dy \,. \tag{3.1}$$
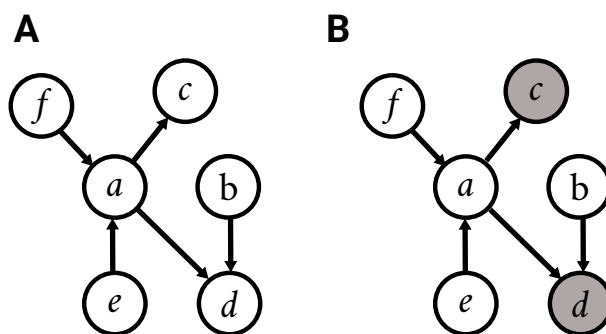
Figure 3.1: Examples for graphical models. (A) Graphical model showing the dependencies for the random variables a, b, c, d, e, and f. (B) Same graphical model as in A after having observed the outcome for c and d.

The probability distribution of a random variable x after we have observed the outcome of y is the conditional distribution $p(x|y)$. We can express the joint distribution as the product of a conditional distribution and a marginal distribution as

$$p(x, y) = p(x|y)p(y) \,. \tag{3.2}$$

Bayes' rule follows by simply rewriting Equation 3.2:

$$p(x|y) = \frac{p(x, y)}{p(y)} = \frac{p(y|x)p(x)}{p(y)} \tag{3.3}$$

Two random variables are independent if knowing the outcome of one of the variables does not give us any information about the outcome of the other. Hence, if two random variables x and y are independent, it follows that $p(x|y) = p(x)$ and $p(y|x) = p(y)$, respectively, and therefore the joint distribution can be written as the product of the marginal distributions

$$p(x, y) = p(x)p(y) \,. \tag{3.4}$$

Two random variables are conditionally independent if they become independent in the presence of a third variable. If x is conditionally independent of y given z it follows that

$$p(x, y|z) = p(x|z)p(y|z) \,. \tag{3.5}$$

Computations are simplified if variables are independent or conditionally independent.

3.1.2 Graphical models

A graphical model captures the dependency structure between random variables. It consists of a set of random variables (nodes) as well as a set of connections (edges) between the nodes. An edge between two random variables represents a dependency. For direct edges in a graphical model we can derive a factorization of the joint distribution that considers all independence assumptions and conditional independence assumptions made by the model using

$$p(\text{nodes}) = \prod_{\text{node} \in \text{nodes}} p(\text{node}|pa(\text{node})) \tag{3.6}$$
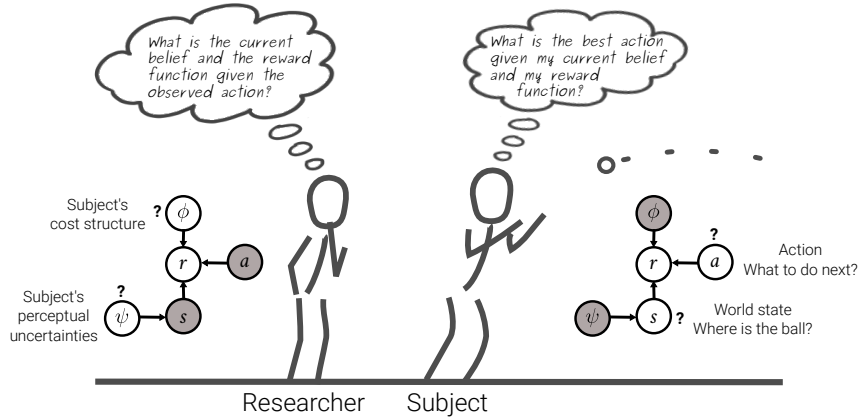
Figure 3.2: Different levels of computational models during ball catching. The figure depicts an agent (right) trying to catch the ball reasoning about the current position of the ball (s) and what action (a) to perform next. The researcher is observing the agent's action (a) while trying to infer the uncertainties of the agent's sensory system (ψ) and the structure of the agent's cost function (ϕ).

where $pa(\cdot)$ denotes the parents of node $\cdot$. For example, the joint distribution shown in the graphical model in Figure 3.1A can be factorized as

$$p(a, b, c, d, e, f) = p(a|e, f)p(d|a, b)p(c|a)p(b)p(e)p(f) . \tag{3.7}$$

In practice, we are interested in computing $p(\mathcal{A}|\mathcal{B})$, i.e., the probability distribution of a set of random variables $\mathcal{A}$ after having observed a set of random variables $\mathcal{B}$ from a joint distribution $p(\mathcal{A}, \mathcal{B}, \mathcal{C})$. For example, when assessing a person's IQ, we want to infer the latent concept of intelligence (intelligence $\in \mathcal{A}$) through the measurement of the test score (score $\in \mathcal{B}$) while integrating out the influence of confounding effects (e.g., fatigue $\in \mathcal{C}$). Using the rules for marginalization and conditional distributions we can compute any distribution of the random variables in a graphical model using the joint distribution

$$p(\mathcal{A}|\mathcal{B}) = \int_{\mathcal{C}} \frac{p(\mathcal{A}, \mathcal{B}, \mathcal{C})}{p(\mathcal{B})} d\mathcal{C} . \tag{3.8}$$

For example, if we want to compute $p(a, b|c, d)$ from the joint distribution $p(a, b, c, d, e, f)$ in Figure 3.1B we can use the same approach and set $\mathcal{A} := a, b$, $\mathcal{B} := c, d$, and $\mathcal{C} := e, f$ yielding

$$p(a, b|c, d) = \int_{e} \int_{f} \frac{p(a|e, f)p(d|a, b)p(c|a)p(b)p(e)p(f)}{p(c, d)} df \, de . \tag{3.9}$$

3.1.3   Levels of modeling

A computational model for behavior is a quantitative description for the action selection of an agent. It specifies how information in the form of sensory percepts is processed and used to choose appropriate actions. However, it is important to distinguish two different levels of computational models that arise when studying active agents (see Figure 3.2): (A) The cognitive model of the human or the animal solving the problem. This model describes how the agent processes sensory observation, makes inferences about latent environmental states, and finally takes actions that optimizes a mixture of intrinsic and task-related rewards and costs (reinforcement learning). (B)
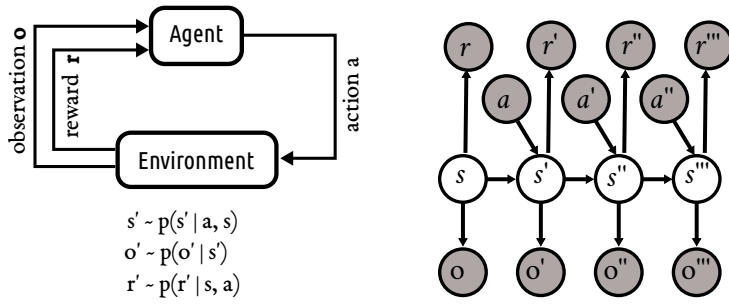
Figure 3.3: Agent-environment relationship. The left panel shows the procedure of the agent interacting with the environment. The right panel shows the statistical relationship.

The model of the researcher, who observes data in the form of the subjects' actions and makes inferences about latent characteristics (inverse reinforcement learning). While the subjects' reasoning is targeted at what the best action is given the current belief and the reward function, the researcher is reasoning about what the current belief and the reward function is given that the observed action of the subject is close to optimal.

3.2 Partially observable Markov decision processes

Next, we describe a general formalization of how an agent interacts with its environment. The procedure is as follows (see Figure 3.3): The agent, currently in a state s, takes an action a and transitions to the next environmental state s'. It then receives an observation from the sensory systems o' as well as a reward r'. This formulation as well as all computational models presented in the present thesis build on the concept of partially observable Markov decision processes (POMDP). A Markov Decision Process (MDP Bellman, 1957; Sutton & Barto, 1998) is a mathematical formalization of an agent acting in an environment. Formally, it denotes a tuple $(\mathcal{S}, \mathcal{A}, T, R, \gamma)$, where $\mathcal{S}$ is a set of states, $\mathcal{A}$ is a set of actions, the transition function $T = p(s'|s, a)$ is a conditional probability distribution of the next state given the current state and the current action, R represents the reward function, and finally, γ denotes the discount factor.

The state space $\mathcal{S}$ is a formal description of what states of the environment are relevant to the problem. Usually, this depends on what problem we want to solve. For example, a potential state-space is the location of a target in visual a search task. Another example is the location and velocity of a ball in a ball catching task. A state formulation is denoted *complete* if it is Markovian. The Markov property is satisfied, if $p(s'|s^-, a^-, o^-) = p(s'|s, a)$, where s^-, a^-, and o^- denote all states, actions, and observations prior to s', for all states. Hence, the state s captures all relevant information from past actions and observations.

In POMDPs (Kaelbling, Littman, & Cassandra, 1998; Murphy, 2000), the states are not directly observable but can only be inferred through observations. For example, in a visual search task, the location of the target is not known. With each eye movement, the distribution of potential target locations is altered. A POMDP is formalized as a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{Z}, \mathcal{O}, \mathcal{T}, \mathcal{R}, z_0)$, where $\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}$ define the underlying MDP. In addition, $\mathcal{Z}$ is a set of belief states, $\mathcal{O}$ is a set of observations and z_0 is the initial belief state. A belief state is a probability distribution over states and therefore summarizes the agent's knowledge about the current state. In our visual search example, the belief state could be the probability distribution over potential target locations. The true state s is not observable directly, instead, it must be inferred from observations.
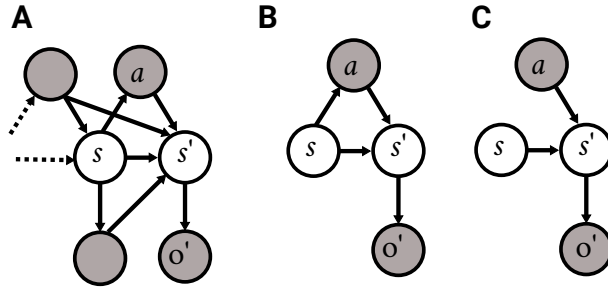
Figure 3.4: Perception as graphical models for inferring the current state. (A) The current state s' depends on the complete history. (B) Assuming the Markov property, the current state is independent of past actions and observations given the last state. (C) The action does not depend on the state.

3.2.1 Perception

Perception uses sensory systems to collect measurements of latent environmental quantities. In the framework of POMDPs the task-relevant quantities are formalized as a state. Our environment is complex, ambiguous, and our sensory systems are imperfect leading to additional uncertainty in our perceptions. As a consequence, the environmental state is not fully observable but must be inferred from perceptual observations. Probability theory provides consistent methods both to quantify the uncertainty we have about the state using probability distributions as well as to update these when novel measurements are received.

The graphical model shown in Figure 3.4 depicts the random variables of interest and the dependencies. Assume that what we currently know about our environment is summarized in the random variable s. Instead of knowing the exact value for s, we keep a probability distribution $p(s)$ assigning each potential state a value indicating the degree of belief of that particular state being the true one. Next, we perform an action a and as a consequence collect a new measurement o' and transition to the next state s'. How can we incorporate this new measurement into what we know about the state s?

Formally, we want to compute $p(s'|a^-, o^-)$, the probability of s' given all past actions a^- and all past observations o^-. From the graphical model, we can derive this probability distribution using the rules for probability distributions introduced in the last section

$$p(s'|a^-, o^-) = \int\limits_{s^-} \int\limits_{a^-} \int\limits_{o^-} \frac{p(s', s^-, a^-, o^-, a, o')}{p(a, o')} do^- \, da^- \, ds^- \, .$$ (3.10)

We can simplify this problem by making several independence assumptions leading to the Bayes' filter and finally to Bayesian inference.

Assumption 1: Markovian state If we assume our state formulation to be complete, i.e., each state summarizes all relevant information about the latent quantities of interest, the sequence of states becomes a Markov chain. Hence, the probability distribution of each state is independent of the past given the last state. In the graphical model, this corresponds to the absence of edges between the state and past observations, actions, and states (Figure 3.4B)

$$p(s'|a, o') = \int\limits_{s} \frac{p(s', s, a, o')}{p(a, o')} ds \, .$$ (3.11)

From the graphical model, we can derive the factorization as

$$p(s'|a, o') = \frac{p(o'|s') \int_s p(s'|s, a)p(a|s)p(s)ds}{p(a, o')} \tag{3.12}$$

where $p(o'|s')$ is the likelihood of the observation given the new state, $\int_s p(s'|s, a)p(a|s)p(s)ds$ is the prior probability of the new state before observing o', and $p(a, o')$ is the marginal probability of the observations.

Assumption 2: Action does not depend on the state If we make the assumption that the action is independent of the state (Figure 3.4C), i.e., $p(a|s) = p(a)$, computations simplify and we get the equation for the Bayes filter

$$p(s'|a, o') = \frac{p(o'|s')p(a) \int_s p(s'|s, a)p(s)ds}{p(a, o')} \propto p(o'|s') \int_s p(s'|s, a)p(s)ds. \tag{3.13}$$

This assumption is easily violated, for example, when we design an experiment a based on what we already know s about a quantity of interest s'. It is apparent that in some situations the prior knowledge has an impact on the experimental design. For example, during visual search, our prior belief about the location of a target influences where we direct our gaze to next (Torralba, Oliva, Castelhano, & Henderson, 2006).

Assumption 3: Action does not change the state If we assume that the action does not change the state, $p(s'|a, s)$ becomes a point mass at $s' = s$, computations simplify and we obtain

$$p(s'|a, o') = \frac{p(o'|s')p(s)}{p(o')}, \tag{3.14}$$

which is Bayes' rule as it is applied in data analysis. This assumption is valid for purely perceptual tasks, for example, eye movements usually do not influence the course of the environment. Although, under some circumstances, they have been shown to do so, e.g., in social settings (see Chapter 2.6).

3.2.2 Action

In the previous section, we have shown how observations can be used to update our belief about potential states. Here, we describe an approach for finding actions that lead to good performance, i.e., yield high rewards. In particular, our goal is to find an action sequence that maximizes the expected sum of future rewards, i.e., the optimal policy π^*. A policy is a mapping between states and actions and the optimal policy connects each state to the optimal action. In partially observable environments, complete states comprise probability distributions over states s, the belief states $b = p(s)$. Hence, we account for the uncertainty about the current state through the belief state b.

The expected sum of long-term rewards associated with performing action a in belief state b and following policy π after that is denoted the state-action value $Q^\pi(b, a)$. We can compute $Q^\pi(b, a)$ by

$$Q^\pi(b, a) = \int_{b'} p(b' \mid b, a)\big[R(b', a) + \gamma V^\pi(b')\big]db' \tag{3.15}$$

where $V^\pi(b')$ is the expected future reward in the next belief state $b' = p(s')$ and $p(b'|b, a)$ is the probability of transitioning from belief state b to belief state b' when performing action a. $R(b', a)$ is the expected immediate reward and γ is the discount factor. Essentially, this means that the value of an action based on the current belief is a combination of the immediate reward and the expected long-term reward, weighted by how likely the next belief is under the action. The discount factor controls the trade-off between immediate rewards and long-term rewards.

The transition probability $p(b'|b, a)$ can be computed as

$$p(b'|b, a) = \int_{o'} p(b'|b, a, o')p(o'|b)do' \,. \tag{3.16}$$

Crucially, under the assumption that the agent uses the Bayes' filter to update its belief $p(b'|b, a, o')$ is a point mass at $b' = p(s'|b, a, o')$. This is due to the fact, that given a belief b, an action a, and an observation o', there is a unique next belief b' under the Bayes' filter. Plugging this into Equation 3.15 yields

$$Q^\pi(b, a) = \int_{o'} p(s'|b, a, o')p(o'|b)\big[R(p(s'|b, a, o'), a) + \gamma V^\pi(p(s'|b, a, o'))\big]do' \tag{3.17}$$

We can compute the Q values using dynamic programming. Once we have computed the Q values, we can derive the optimal policy π^*, i.e., a mapping between belief states and actions that maximize the reward. Given the optimal Q values $Q^*(s, a) = \max_\pi Q^\pi(b, a)$ the action according to the optimal policy a_{π^*} can be derived as

$$a_{\pi^*} = \arg\max_a Q^*(s, a) \tag{3.18}$$

3.3 Applications

The presented framework of partially observable Markov decision processes is powerful but computationally expensive. Due to its generality, it subsumes various prominent methods, including Bayesian decision theory, ideal observer theory, Bayesian experimental design, and reinforcement learning.

3.3.1 Observable states and delayed rewards

If the state space is observable, i.e., the current state s is fully known, our belief state $b = p(s)$ is one at s and zero everywhere else. Hence, we can simplify the computation of the state-action values by using s instead of b in Equation 3.15

$$Q(s, a) = \int_{s'} p(s' \mid s, a)\big[R(s', a) + \gamma V^*(s')\big]ds' \,. \tag{3.19}$$

This yields the general formalism of reinforcement learning, that has been applied frequently to experimental data (see Dayan & Daw, 2008, and the references, therein).

3.3.2 Observable states and immediate rewards

If delayed rewards do not play a role, computations can be simplified further by setting the discount factor γ to zero yielding

$$Q(s, a) = \int_{s'} p(s' \mid s, a) R(s', a) ds' \,. \tag{3.20}$$

For example, this approach was used by Trommershäuser et al. (2006, see Chapter 2). In their study, perceptual uncertainty was negligible, thus the state was assumed to be known. As the task only comprised single ballistic action, delayed rewards did not play a role. In their reaching model, subjects chose the location with the maximum expected reward while considering their reaching variability $p(s'|s, a)$.

3.3.3 Partially observable states and immediate rewards

The sensory information we perceive is noisy, and due to the limited resources of our visual system, we can only perceive a small proportion of our environment at any given time. Hence, in most situations, we have uncertainty about the current state. Nevertheless, under some circumstances, it is sufficient to compute the optimal next action rather than the action that maximizes the long-term return. A first example is a scenario in which the task only comprises a single action and only a single reward is received. Clearly, the sum of all returns is equal to the single reward in this scenario. A second scenario is if the action does not change the state. For example, when tossing a coin several times the state before each toss does not change, hence finding the optimal action for this state is sufficient (this is exploited in the model for the temporal event detection task in Chapter 6). Third, maximizing the immediate reward instead of long-term returns can be an appropriate model for agents not capable of the computations involved in handling delayed rewards (a characteristic often implicitly assumed for human eye movements, which we question in Chapter 8).

Removing delayed rewards from our full POMDP model yields

$$Q^\pi(b, a) = \int_b p(b' \mid b, a) R(b, a) db' \,, \tag{3.21}$$

alternatively, Equation 3.17 leads to

$$Q^\pi(b, a) = \int_{o'} p(s'|b, a, o) R(p(s'|b, a, o), a) do' \,. \tag{3.22}$$

This corresponds to taking the action that yields the belief that leads to the maximal reward and is equivalent to the ideal searcher in the visual search task of Najemnik and Geisler (2005), which suggests fixating the location that subsequently maximizes the probability of finding the target.

If we assume that the observation has already been received the model results in Bayesian decision theory. To maximize the reward, the action that maximizes the expected reward should be taken.

$$Q^\pi(b, a, o') = \int_{s'} p(s'|a, o') R(s', a) ds' \,. \tag{3.23}$$

where $p(s'|a, o')$ is the posterior and $R(s', a)$ is the reward function.

3.3.4 Partially observable states and delayed rewards

This is the most general and most complex scenario described at the beginning of the chapter. However, for many problems, the approach is computationally not feasible (Murphy, 2000). Over the last years, approximate techniques have been developed, but guarantees are scarce and empirical evaluations currently limited (Hausknecht & Stone, 2015). Especially, when modeling human behavior, multiple characteristics further complicate the approach. The optimal action an agent can take is the action that yields the maximum expected long-term reward. The reward or utility "are numerical representations of . . . tastes and preferences" (DeGroot, 2005, p. 86). The reward function $R : (s, a) \rightarrow \mathbb{R}$ maps each state-action pair (s, a) to a scalar number r. However, seldom a single factor contributes to the reward (Rothkopf & Ballard, 2013). Instead, a state action pair is associated with multiple sources of positive (reward) and negative (cost) returns. It has been shown that many aspects contribute to the total reward associated with an action (see Chapter 2.4, for empirical work regarding the cost structure)

$$R(s, a) = \sum_i r_i(s, a) \tag{3.24}$$

where $r_i(s, a)$ denotes a component contributing to the reward associated with state action pair (s, a). The different factors contributing to the overall reward originate from the environment (e.g., task performance) or from the agent itself (e.g., intrinsic costs like mental effort). The specific structure of the reward function is unknown to the researcher, in general, and is inferred in the research process (inverse reinforcement learning), for example using a linear model and estimating how much each dimension contributes to the total reward

$$R(s, a) = \sum_i w_i x_i(s, a) \tag{3.25}$$

where x_i is an explanatory variable connected to the costs measured by the researcher and w_i quantifies how much the variable influences the total reward. Also, action selection may be influenced by uncertainties regarding consequences of actions, as humans may know which states are rewarding, but not how to reach them. This can be modeled using structured generative models of actions deviating from reward maximization in inverse reinforcement learning (Rothkopf & Dimitrakakis, 2011).

In natural behavior, the action space is high-dimensional as various end effectors are controlled simultaneously. For example, in tasks involving eye-hand coordination like making a peanut butter and jelly sandwich (Hayhoe & Ballard, 2005) the action space comprises at least tuples of actions $A = (a_{hand}, a_{gaze})$ at each time step. In reality, various joints have to be controlled using motor signals to generate proper hand movements. If we consider sports games like playing squash (Hayhoe et al., 2012) even more dimensions need to be controlled simultaneously. Finally, in addition to learning the connections between actions, rewards, observations, and states, human subjects may even have to learn what state representations to use, which action they might have at their disposal and what can be considered rewarding. This full problem is still out of scope of current machine learning.

Related work on temporal eye movements

4.1 Computational models for eye movements

Computational and algorithmic models for eye movements attempt to describe where we direct gaze and predict future fixation locations in quantitative, mathematical terms. Various factors have shown to influence where we choose to fixate including low-level image features (Borji, Sihite, & Itti, 2013), scene gist (Torralba et al., 2006) and scene semantics (Henderson, 2003), task constraints (Rothkopf et al., 2007), extrinsic rewards (Navalpakkam et al., 2010), and also combinations of factors have been investigated (Schütz, Trommershauser, & Gegenfurtner, 2012). Studies have cast doubts on the optimality of spatial gaze selection (Morvan & Maloney, 2012; Morvan & Maloney, 2009; Clarke & Hunt, 2015), but computational models of the spatial selection have in part established that humans are close to optimal in targeting locations in the visual scene, as in visual search for visible (Najemnik & Geisler, 2005) and invisible targets (Chukoskie, Snider, Mozer, Krauzlis, & Sejnowski, 2013) as well as face recognition (Peterson & Eckstein, 2012) and pattern classification (Yang, Lengyel, & Wolpert, 2016).

4.1.1 Bottom-up models

Bottom-up models predict human eye movements solely from the visual input suggesting that low-level features in the visual scene guide eye movements. In order to predict spatial targets of eye movements, low-level features like contrast and orientation are used to assign a value to each region present according to how *salient* it is. The resulting map is referred to as *saliency map* (e.g., Itti, Koch, & Niebur, 1998). "The purpose of the saliency map is to represent the conspicuity ... at every location in the visual field by a scalar quantity and to guide the selection of attended locations" (Itti et al., 1998, p. 1255).

Although many different saliency models have been proposed in the past, the approaches share a common procedure: First, the image is transformed to low-level features like color or orientation using filtering by appropriate kernels. The maps that result from filtering are then combined to yield a saliency map. Bruce and Tsotsos (2009) proposed a saliency model they term *Attention based on Information Maximization* (AIM). They computed the likelihood of the features for an input image using the distribution of feature coefficients at each location in the image. The saliency map is created from these distributions after transforming them to the information-theoretic measure of self-information. The described methods use predefined filters that are either handcrafted or based on findings of human information processing. However, more recently, machine learning has been used to learn to classify salient regions in natural images (see Zhao & Koch, 2013, for a review). Also, the computation of saliency using deep neural networks has been proposed (Zhao, Ouyang, Li, & Wang, 2015). Thereby, a deep convolutional

network is pretrained on image classification. The resulting filters are then used as initial values for predicting the saliency of image regions.

As only the visual input is considered, computational and algorithmic models based on saliency maps predict the same scanpath independent of which task a person is solving, although this is a factor known to be contributing to gaze behavior (Yarbus, 1967, Figure 4.1A). Prior studies have shown, that these computer-generated saliency maps are in accordance with the judgments of observers about the saliency of objects in scenes. In their study, Borji et al. (2013) let subjects choose which object in a scene containing two objects was considered more salient. The results showed that the computed saliency of objects selected by the subjects was significantly higher compared to the objects that were not selected. While being able to qualitatively reflect the effects found in humans regarding pop-out and conjunctive search, saliency models performed poorly for visual search in complex natural images (Itti & Koch, 2000).

A saliency map is not a direct model for eye movements, i.e., it does not imply any actions. Instead, it merely is a description of how much each location attracts gaze. Hence, an additional mechanism generating actions from the map is needed. For example, in Itti et al. (1998) the last layer of the saliency map is modeled as leaky integrate-and-fire neurons. The inputs of the neurons are the aggregated feature maps. The neurons excite a network of winner-takes-all neurons. Once a threshold of activation is reached, the model predicts an eye movement to the corresponding location.

However, additional characteristics have to be accounted for when generating realistic eye movement patterns from saliency maps. First, the saliency map is independent of the task as well as of the previous actions of the observer. Hence, a location with a high value of saliency remains a high value independent of how long it already has been fixated. In order to prevent getting stuck at a location or reattending the same location over and over again, previously attended locations are inhibited (inhibition of return). Using inhibition of return eye movements are pushed towards other high saliency regions apart from the current fixation location. However, this does not depend on the distance to the next location, hence it assumes that the amplitude of the saccade does not play a role. Clearly, the distribution of saccade amplitudes is not uniform (Tatler, Baddeley, & Vincent, 2006), instead, it is dominated by smaller amplitudes. In order to favor locations close to previously attended regions, mechanisms that promote smaller amplitudes have been suggested (proximity preference). More fundamentally, it has been disputed, whether inhibition of return actually is a biologically plausible spatial mechanism in gaze selection (Wilming, Harst, Schmidt, & König, 2013; Ludwig, Farrell, Ellis, & Gilchrist, 2009). For example, it has been shown that inhibition of return is greatly influenced by environmental statistics (Farrell, Ludwig, Ellis, & Gilchrist, 2010). Finally, saliency models do not include the temporal course of eye movements, i.e., the fixation durations. Overall, while various results indicate that eye movements are influenced by low-level features, saliency models account for only a part rather than yielding complete models for human eye movements.

4.1.2 Top down models

Besides low-level features the impact of top-down processes on the perception of visual stimuli has been investigated (see Gilbert & Li, 2013, for a review on a neural level). Eye movements have been shown to be particularly influenced by the current goals (Rothkopf et al., 2007; Yarbus, 1967; Castelhano, Mack, & Henderson, 2009, Figure 4.1B). The specific reward structure an agent attempts to maximize is usually unknown, though the performance in the current task is likely to be a factor (see Peterson & Eckstein, 2012; Najemnik & Geisler, 2005, for example).
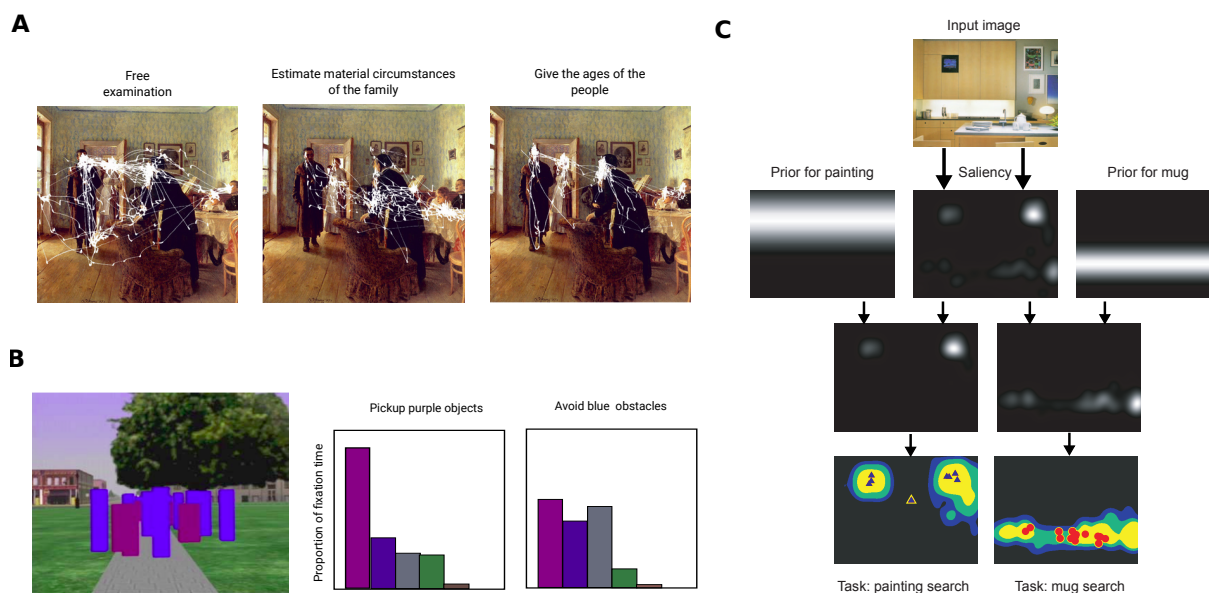
Figure 4.1: (A) Data shown form Yarbus (1967). Subjects' eye movements are shown (white traces) while looking at the painting *An Unexpected Visitor* by I.E. Repin. The graphical images overlaying the original data from Yarbus (1967) onto a color image of the painting was retrieved from Archibald (2008). (B) Experimental design and results from Rothkopf, Ballard, and Hayhoe (2007). During navigation (left), eye movements were different with respect to the instruction. If picking up the purple objects was the dominant task, purple objects were fixated more often compared to if the task was to avoid blue obstacles (right). (C) Effect of context information in visual search (adapted from Torralba, Oliva, Castelhano, & Henderson, 2006). The bottom-up features of the scene are independent of the task. However, the scene context, i.e. the prior belief about where an object is located within a scene highly depends on the object. Therefore, when searching for a painting, the upper part of the scene is more likely as paintings usually hang on walls. The results showed that eye movement patterns resembled a combination of saliency and location prior.

This indicates that we control our eye movements in a way that leads to maximal task related reward. In contrast, information theoretic measures suggest that eye movements are deployed to maximize knowledge about the environment. Itti and Baldi (2006) proposed a model based on Bayesian surprise suggesting that we move our eyes to the location that leads to the greatest difference between prior and posterior, i.e., the highest gain in information. Also, it has been suggested that actions are chosen that minimize free energy of sensory inputs (see Friston, Adams, Perrinet, & Breakspear, 2012, and the references therein). A similar methodology was used for defining goal relevance, a measure of how much an observation changes how a problem can be solved (Tanner & Itti, 2017). Another information theoretic approach was presented by Renninger, Verghese, and Coughlan (2007). Under specific conditions, information theoretic approaches lead to similar actions as maximizing task-related reward. In their study, participants studied an irregular shape for 1.2 seconds. Subsequently, they were asked to choose the previously displayed shape among a set of two similar shapes. The results showed that eye movements were well described by a computational model maximizing information. However, the proposed action selection only maximized the information in the next step, a property that we will question in Chapter 8. Finally, a combination of task related reward and information theoretic minimization of uncertainty has been presented in the past (Schwartenbeck et al., 2015).
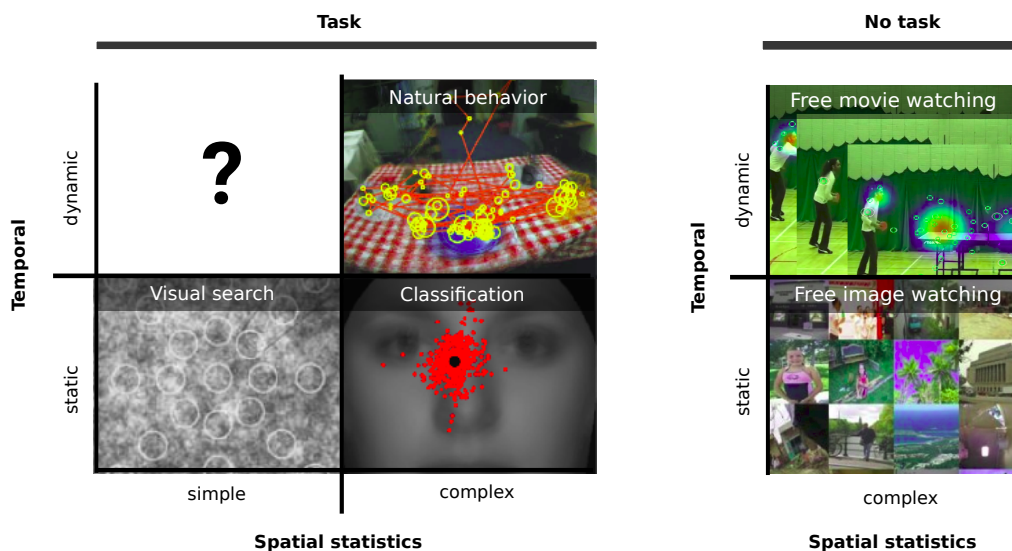
Figure 4.2: Experimental designs used to study eye movements. First, experimental paradigms can either have a clear reward structure (left) or not (right). While in the former actions are associated with different costs and benefits, for example, smart gaze behavior in visual search leads to a greater probability of finding the target, the latter has no such relationship. Second, we can categorize the stimuli as static (they remain constant during the task) or dynamic. Finally, on the spatial level stimuli can either be more complex like natural images or less complex like Gabors or Ts among Ls.

4.1.3   Combination of bottom-up models and top-down models

Humans have shown to fixate locations that are both salient as well as valuable (Navalpakkam et al., 2010). There have been attempts to combine low-level visual features with top-down processes. For example, prominent top-down factors comprise features like face detection and preference for central and horizontal locations (see Judd, Ehinger, Durand, & Torralba, 2009; Borji, 2012, for example). Using Bayes' rule Oliva, Torralba, Castelhano, and Henderson (2003) combined low-level image based features with top-down characteristics. In particular, a prior probability over locations for objects in scenes was learned from a set of images. The study gives a possible explanation, how the heuristics used in other works (e.g., horizon preference) could arise from experience. It has been shown that humans use context information when searching for a target in natural images. Results from human eye movements show that that information about the task-specific scene context is combined with bottom-up features to select suitable gaze targets (Torralba et al., 2006). The influence of scene context is shown in Figure 4.1C.

4.2   Experimental design for studying eye movements

Different experimental designs have been used to study the visual system deploying different stimuli and tasks (Figure 4.2[1]). The distribution of eye movements in our daily lifes is characterized by enormous complexity making studies about the underlying processes and predicting future behavioral actions incredibly difficult. Consider a complex visual task like navigating through a crowded city. We need to find a suitable path to our destination while avoiding colliding with traffic and pedestrians. This is done by receiving visual information through eye

---

[1]Images were taken from Mital, Smith, Hill, and Henderson (2011), Najemnik and Geisler (2005), Peterson and Eckstein (2012), Judd et al. (2009), Hayhoe and Ballard (2005).

movements and simultaneously solving the complex motor tasks involved in walking. Hence, we perform multiple actions at any given time. In addition, various aspects of this environment have their own dynamics which are stochastic and uncertain. Also, inherent in an environment containing other active agents, there are several other factors introducing uncertainty: First, extracting relevant information, e.g., the current velocity of a car driving by, is only possible with limited precision. Second, our visual apparatus can only perceive a small area at any given time. This introduces uncertainty about the state in the regions that are currently not fixated. Third, uncertainty increases over time due to memory leakage. These statistics and their uncertainties need to be considered when making predictions about the environmental state. Being able to predict the course of the environment is crucial on many levels ranging from overcoming internal delays in eye-hand coordination (Miall et al., 2007) to successfully managing complex tasks (e.g., Hayhoe et al., 2012).

In a natural task, the environmental state is high dimensional and uncertain, the dynamics (transition function) are only partially known, stochastic and depend on the current time. Also, a rich structure of dependencies connect the spatial and temporal aspects of the environment and multiple actions are often performed simultaneously. Finally, the reward structure is multi-dimensional, usually only known to the researcher to some extent, and the connection to actions is often delayed. Hence, while some authors claim that studying the visual system should always use natural images as stimuli (Yuille & Kersten, 2006), creating computational models for eye movements in these tasks is unfeasible due to the complexity of natural behavior in dynamic environments. While many of the effects found in spatial gaze selection originate from controlled experiments using simple stimuli with known statistics, research applying this paradigm to the temporal dynamics of vision is scarce.

4.3   Stimuli for complex temporal dynamics

In Chapter 3, we gave a brief introduction in partially observable Markov decision processes. While the approach is powerful in describing decisions and actions in various situations, the computations involved in finding optimal solutions are costly and for complex scenarios the analysis quickly becomes unfeasible. In particular, this applies to nonstationary environments where the state transition $p(s'|s, a, t)$ also depends on time. Here, we investigate properties of experiments and stimuli which make computational models based on POMDPs applicable. We want to probe the visual system using stimuli with temporal dynamics that reflect the dynamics found in our natural environment. The following stimulus and task characteristics are advantageous for modeling visual behavior as they reduce complexity not relevant for the research question.

(1) *Controlled temporal statistics.*
  The relevant temporal statistics should be known and controlled. This enables computing the probability distribution over the future states $p(s_{t+n}|s_t, n)$, the basis of inferring the subjects' belief about the state progression while gaze is directed elsewhere. When the transition probability is unknown, e.g., when using natural stimuli like movies, images, or real-world behavior, computing the transition in the belief space $p(b'|b, a)$ is not possible.

(2) *No peripheral vision.*
  We minimize the spatial complexity of the stimuli by eliminating the role of peripheral vision (either by adjusting the distance between stimuli or by using gaze-contingent techniques). As a consequence, information is only obtained from a single location at any given

time and we do not need to account for the decline of visual acuity with increasing eccentricity. While proper methods to account for foveation exist (see Geisler & Perry, 1998, for example) and have been used in computational models for spatial problems (Najemnik & Geisler, 2005; Peterson & Eckstein, 2012), unnecessary complexity is added to the study of the time course of eye movement strategies. By excluding peripheral vision, the action space becomes discrete.

(3) *Discrete observations.*
The detection of the environmental state should be as simple as possible. In particular, the state of a stimulus should be completely observed after an initial processing time. Hence, if fixated, the belief about the state of a stimulus is constant, unless there are changes in the stimulus. This yields an observation distribution $p(o'|s')$ that can be managed computationally. This is opposed to complex visual scenes where information about action relevant states must be integrated from multiple observations. For example, while driving, multiple observations are needed to infer the velocity from the car behind us. Also, in complex scenes, information must often be integrated across multiple fixations.

(4) *Independent locations.*
The dynamics of the discrete regions should be independent. Thus, observing a location does not reveal any information about the state of any other region.

Stimuli having these characteristics facilitate tractable computational models while still sharing similarities with the temporal dynamics found in our environment. However, in order to use the framework of partially observable Markov decision processes for studying eye movements, we need to have access to the reward structure. This excludes free viewing of static and dynamic scenes. Instead, we propose *event detection tasks* as a natural counterpart to what visual search is for spatial gaze selection in the temporal domain. There is a clear connection between internal beliefs about the current environmental state and reward. More knowledge in task-relevant domains leads to higher probability of event detection and therefore to better task performance. While eye movements are organized together with other actions in natural behavior, in an event detection task no other actions are needed and therefore the complexity is reduced further. However, the approach is not limited and can easily be extended, for example by combining it with a control task.

Humans quickly learn to blink strategically in response to environmental task demands

5.1 Introduction

5.1.1 Physiology of blinking

Blinking is an omnipresent involuntary process, which humans carry out $15 - 17$ times per minute (Bressman, Cassetta, & Bentivoglio, 1997), on average. It primarily serves the physiological purpose of cleaning the surface of the eye and providing a stable tear film (Sweeney, Millar, & Raju, 2013), of preventing optical aberrations (Montés-Micó, 2007; Koh, Maeda, Hirohara, Mihashi, & Ninomiya, 2006; Montés-Micó, 2007; Koh et al., 2008), and thus maintaining a good quality of vision (Tutt, Bradley, Begley, & Thibos, 2000). Besides these positive consequences, blinking has an immediate negative consequence on perception as while blinking the stream of visual information is interrupted (Figure 5.1A). Moreover, in addition to the break of optic signals due to the occlusion of the pupil by the upper lid, neural processing is inhibited (Bristow, Haynes, Sylvester, Frith, & Rees, 2005; Volkmann, Riggs, & Moore, 1980). This leads to perceptual gaps every two to three seconds. If not timed correctly, these gaps can lead to negative outcomes in numerous situations. In social interaction, for example, microexpressions have a duration between 160 and 500 ms and can therefore easily be missed due to a blink (Yan, Wu, Liang, Chen, & Fu, 2013). Indeed, humans have been shown to reduce these gaps by combining blinks with saccades (Fogarty & Stern, 1989; Evinger et al., 1994), during which neuronal processing is also inhibited. The same has been shown for animals (Yorzinski, 2016; Gandhi, 2012). Overall, the positive effects of blinking on the maintenance of proper vision and the negative effects of the interruption of vision constitute a fundamental trade-off. Hence, controlling actively when to blink provides a behavioral advantage compared to blinking at random points in time.

5.1.2 Blinking and visual information

Investigations of the control of blinking have revealed an intriguing multitude of additional factors influencing human blink rates so that blinking is often used as a behavioral marker for a variety of internal processes. Blinking is closely intertwined with cognitive functions connected to dopamine (see Hayes & Greenshaw, 2011, for a recent review). In particular, blink rates reflect the progress of learning (Slagter, Georgopoulou, & Frank, 2015) and the perception of time (Terhune, Sullivan, & Simola, 2016). Blinking is also affected by our current goals and actions. It depends on what task we are solving (see Garcia, Pinto, Barbosa, & Cruz, 2011;
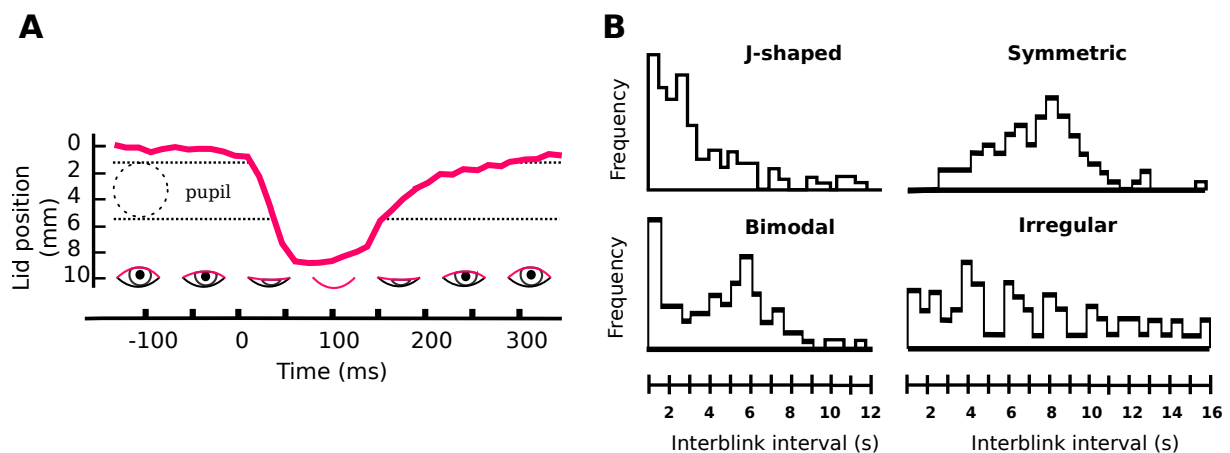
Figure 5.1: Experimental design and stimulus generation. (A) Schematic time course for a single blink (adapted from Volkmann, Riggs, & Moore, 1980). (B) Four different interblink interval distributions found by Ponder and Kennedy (1927). The shapes are not equally frequent. Most subjects showed J-shaped (62 %), Symmetric was the least often (4 %), Bimodal and Irregular comprised 12 % and 22 % of the cases, respectively.

Doughty, 2001; Cruz, Garcia, Pinto, & Cechetti, 2011, for reviews), how long we are on the task (Stern, Boyer, & Schroeder, 1994), how difficult the task is (Schieppati & Schmid, 2011), and it is an indicator of what we remember afterward (Shin et al., 2015). Even the coordination of blinks and saccades is influenced by cognitive factors (Williamson, 2004). Finally, blinks have been shown to be synchronized across people during conversation and thus might play a role in human social communication (Nakano & Kitazawa, 2010). While blinking is clearly related to these cognitive processes it is still debated whether the appropriate measure is the blinking rate during a task, the spontaneous blinking rate, or possibly the distribution of times between two consecutive blinks, the so-called interblink interval distribution (IBI), or whether these quantities are potentially inherently related (see Jongkees & Colzato, 2016, for a recent discussion). Thus, a better understanding of the process underlying blinking behavior is beneficial to understanding a wide range of perceptual and cognitive processes.

5.1.3 Related work on blinking

Although numerous empirical studies have established the importance of blinking, quantitative explanations have been scarce. Current models assume a linear increase in urge when blinking is suppressed (Berman, Horovitz, Morel, & Hallett, 2012) or an oscillating blink generator (Moraitis & Ghosh, 2014), but both models considered voluntary blink suppression and do not incorporate any task-related influences. One reason might be that environmental regularities and task related costs are usually complex and unknown. The lack of quantitative models is surprising considering the strong contingencies between environmental statistics and gaze behavior, which have been explained successfully through modeling (Najemnik & Geisler, 2005; Chukoskie et al., 2013; Peterson & Eckstein, 2012; Hoppe & Rothkopf, 2016). As with blink rates, few computational approaches exist that describe the temporal course of blinking. In particular, the IBI distribution is an active research area. In their seminal paper, Ponder and Kennedy (1927) found four different types of IBI distributions (see Figure 5.1B). Since then, many studies have presented similar results showing that subjects' IBI distributions show great variability (e.g., Zaman & Doughty, 1997). However, the origin of these different IBI

A

B
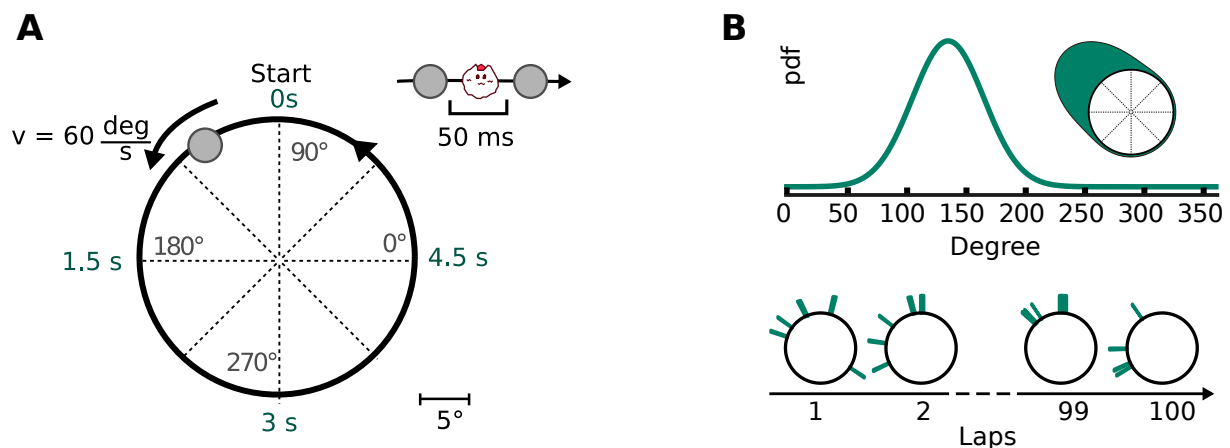
Figure 5.2: Experimental design and stimulus generation. (A) Task design. In each block a grey dot (0.3°) moved on a circular trajectory for 100 consecutive laps. The trajectory was not visible to the subject. In each lap 3 – 5 events occurred. An event comprised the circle being replaced with a comic face for 50 ms. Subjects attempted to detect as many of the events as possible by pressing a button. (B) Event generation for a single block. Three to five events per lap were drawn from a mixture of a uniform and a Gaussian distribution. Sample laps and events are given in the lower panel. Overall, four different conditions were presented to the participants. The conditions differed with respect to the mean of the Gaussian distribution.

distributions and their variability is unclear.

Here we address the question of how blinking behavior is related and adapted to the current task. We conducted a controlled blinking experiment with parametrically generated environmental statistics. Using an event detection paradigm, we created a direct connection between temporary loss of visual information due to blinking and task performance. Crucially, knowing the probabilistic structure of the task, we can investigate how blinks are linked to internal beliefs about task parameters and how participants adapted their blinking to the task. The computational model treats blinking dynamics as the result of a trade-off between the physiological need to blink and the task-related cost of blinking given subjects' beliefs. The model captures our subjects' strategic blinking behavior in the experiment and provides a quantitative explanation for changes in blinking behavior. Importantly, based on the computational model we derive the temporal structure of blinking behavior including the IBI distribution on an individual subject's basis. This provides a first explanation of the classic IBI distributions observed in numerous experiments and lays the groundwork for quantitative studies investigating blinking behavior, task structure, and subjects perceptual uncertainties.

5.2 Experimental design

5.2.1 Task design

Subjects directed their gaze to a grey dot moving on a circular trajectory (counter clock-wise) displayed on a computer monitor in order to detect events (Figure 5.2A). A circular trajectory was chosen to avoid blinks being triggered by saccades (Fogarty & Stern, 1989; Evinger et al., 1994). The velocity of the dot was chosen to lead to smooth pursuit without catch up saccades. Events were defined by replacing the dot with a stylized face for 50 ms. Hence, an average blink (Volkmann et al., 1980; Wang, Toor, Gautam, & Henson, 2011; Schiffman, 1990) could lead to

missing an event. Events were generated probabilistically and were drawn from a mixture of a Gaussian distribution (Figure 5.2B) using a mixing weight of $p = 0.8$ and a uniform distribution with mixing weight of 0.2. For each lap, the number of events was randomly chosen between three and five. Over the course of 100 consecutive laps, subjects could learn the relationship between the angle at which the dot was on its circular trajectory and the probability of an event occurring. Overall, subjects completed four blocks differing with respect to the location of the Gaussian. Blocks were presented in a random order. After each block subjects were given the percentage of detected events as feedback.

5.2.2 Blink detection

Blinks were detected using an infrared eye tracking device (Tobii EyeX eye tracker; 60 Hz). This technique has been used by past research (e.g., Kaminer, Powers, Horn, Hui, & Evinger, 2011; Terhune et al., 2016). During blinking and thus closed eyes, the eye tracking device loses track of the pupils and can not determine the gaze location. These artifacts were identified and used to detect blinks by analyzing their temporal structure. We tested different thresholds on three subjects while manually recording blinks and chose the threshold with the best agreement. For the analysis, blinks were treated as point processes. Using this procedure, we found similar statistics regarding blink rates and interblink intervals compared to studies using magnetic search coils (e.g., Garcia et al., 2011), manual video analysis (Naase, Doughty, & Button, 2005), and EEG (Shin et al., 2015).

5.2.3 Experimental setup

The test subjects were seated in front of a computer screen (Flatron W2242TE monitor; 1680 × 1050 resolution, 60 Hz refresh rate) at a distance of about 55 cm. All experimental procedures were carried out in accordance with the guidelines of the German Psychological Society and approved by the ethics committee of the Darmstadt University of Technology. Subjects gave informed consent and were aware that their eye movements were recorded, however, they were told about the purpose of the task after the experiment to prevent conscious control of the blinking behavior.

5.2.4 Procedure

The experiment was conducted in single sessions. Subjects confirmed the detection of an event by pressing a button on a keyboard. If the button press occurred within a 1s time range after the event, the event was regarded as detected. If no button was pressed, the event was marked as missed. Missing an event produced a sound to signal the miss to the subject. A different sound was played in case of a response without prior event. The sounds served as feedback to help the subject learn the event distribution. Additional feedback was given to the subject after each block in form of a score (one point per detected event, minus one for button presses without events, zero points for missed events).

5.3 Results

Overall, 25 subjects (8 male) participated in the experiment in exchange for course credit. Their age ranged from 19 to 56 years ($M = 26.52$, $SD = 10.97$). Seven subjects wore glasses, however, sufficient accuracy of the eye tracker and the detection of blinks was assured for all participants.
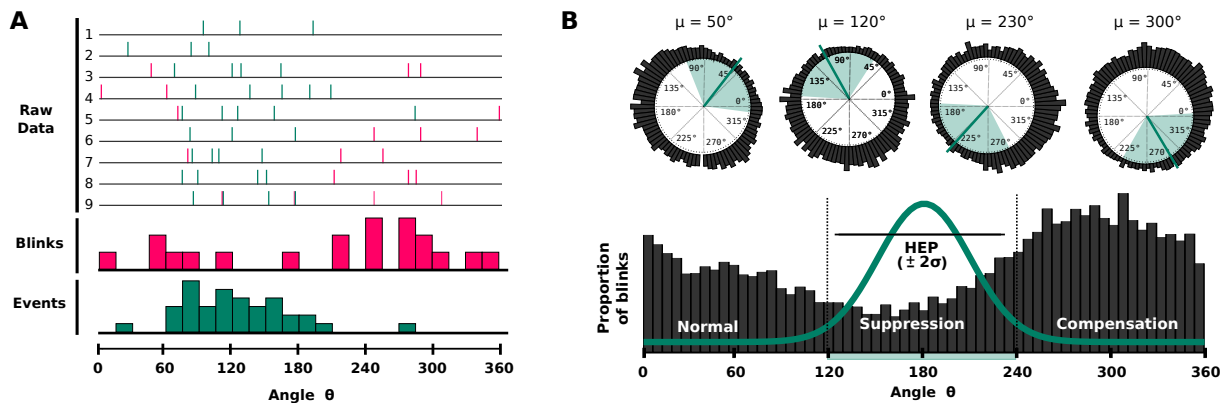
Figure 5.3: Behavioral results. (A) Raw data (blinks and events) for a single participant in nine consecutive laps. The top row depicts the temporal dynamics of blinks and events. The two bottom rows show histograms for the blinks and events aggregated over the angular locations of the laps. (B) Blinking frequencies over locations within the circular trajectory. Histograms are shown for all four conditions (upper panel). Center of the event generating mixture distribution is indicated by a green bar for each condition. Data from all conditions were normalized by rotation (lower panel). The normalized event generating distribution was centered to 180 degrees.

5.3.1 Behavioral results

Data from nine consecutive laps is shown for a single participant in Figure 5.3A. Visual inspection yields the first indications for a connection between event probability and blinking behavior. Overall, the distribution of blink locations for all participants showed a similar time course across all conditions (Figure 5.3B, upper panel). As conditions only differed with respect to the mean of the Gaussian while sharing the same variance, data was aggregated by normalizing the event distribution to peak at 180 degrees (Figure 5.3B, lower panel). The data show that blinking behavior was clearly affected by the distribution of event probabilities. Instead of being distributed uniformly over the circle two characteristics of the distribution of blinks can be observed: First, blinking is suppressed in the area of high event probability (HEP, ± 2 standard deviations from the center of the mixture distribution). Fewer blinks occurred in the HEP area (blink rate $r_{\text{HEP}} = 11.28$) compared to the remaining part ($r_{\text{not HEP}} = 20.46$; see Figure 5.4A). The 95 % credibility interval for the difference in blinking rates was $[8.67, 9.67]$. This indicates that blinking behavior is adapted to the event distribution in order to avoid missing events. Second, visual inspection indicates an asymmetry of blinking counts before and after the HEP region. Indeed, we found a higher number of blinks after (blink rate $r_{\text{after HEP}} = 25.16$; see Figure 5.4B) than before the HEP region (blink rate $r_{\text{before HEP}} = 15.76$; 95 % credibility interval for the difference was $[8.71, 10.01]$).

To investigate how learning of the event distribution affected blinking, we computed the proportion of blinks in the HEP area over the course of the 100 laps (Figure 5.4C). The proportion of blinks occurring in the HEP region declined over the course of the first laps and was constant afterward. This indicates that in the beginning subjects learned the hidden event distribution by observing the event locations. Bayesian change point analysis (see also Perreault, Bernier, Bobée, & Parent, 2000) revealed that steady state was reached on average after about 13 laps (95 % credibility interval of the change point was $[9.15, 17.47]$).
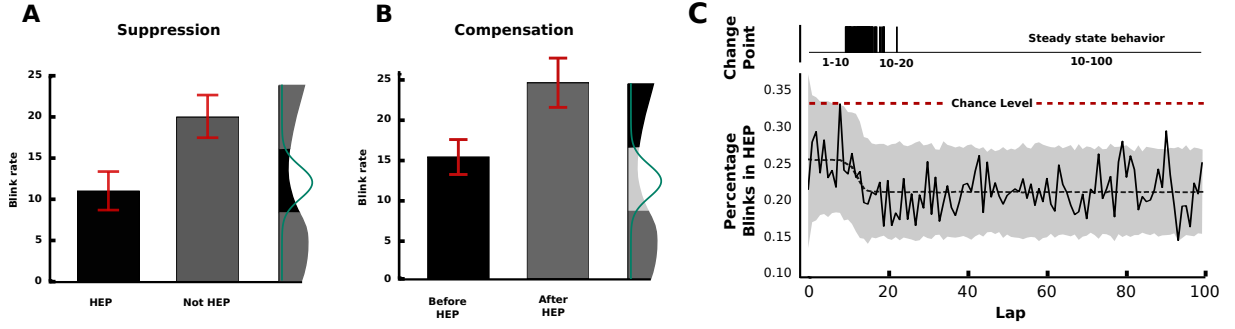
Figure 5.4: Behavioral results. (A) Comparison of the mean number of blinks aggregated for all conditions between the area ±2 standard deviations (HEP) from the center and the remaining part of the circle. Error bars correspond to the standard error of the mean. (B) Differences in number of blinks between the 60° area before and after the HEP area. Again, error bars correspond to the standard error of the mean. (C) Percentage of blinks occurring in the HEP area over the course of 100 laps. Chance level assumes that blinks are uniformly distributed over the circle (red dashed line).

5.3.2 A computational model for blinks

The design of our experiment gives access to the statistical structure of the task in terms of a probabilistic generative model. Crucially, we are able to determine the loss of information due to a blink by quantifying the influence on task performance at each point in time during the experiment. This is a distinct advantage compared to investigations studying human blinking behavior during reading or free viewing, where it is difficult to capture the consequences of blinking quantitatively.

The computational model is motivated by capturing the fundamental trade-off between costs and benefits of blinking and comprises two distinct components that control blinking behavior: 1) internal costs for blink suppression and 2) external task-related costs for blink execution. Internal costs for blink suppression arise due to the subjects physiological need to blink from time to time to maintain healthy vision. Formally, we denote the costs for blink suppression as c_s. It is important to note that c_s does not depend on the current location within the lap since it is independent of the task. Further, our data suggest that c_s does not depend on the time since the last blink. The second factor is the cost associated with blink execution $c_e(\theta)$. It points in the direction opposite to c_s, as more blinks lead to higher costs. This component denotes the amount of task performance that is lost in case of a blink. In our experimental procedure, this is the probability of detecting events. In contrast to the costs for blink suppression c_s, the task-related costs $c_e(\theta)$ depend on the current angle and therefore are not constant over the course of a lap (higher blink related costs in HEP region). Whether to blink or not to blink at any point in time therefore can be described as a trade-off between c_s and $c_e(\theta)$.

Here we assume that the probability of a blink occurring at an angle θ can be modeled as being inversely proportional to the sum of the costs of blink suppression c_s and blink execution c_e:

$$P(\text{blink at } \theta \mid \alpha, \psi) \propto \frac{1}{(1-\alpha)c_s + \alpha c_e(\theta, \psi)} . \tag{5.1}$$

where c_s is the cost for blink suppression, $c_e(\theta, \psi)$ is the cost for blink execution in terms of reduced task performance, α is between 0 and 1 and regulates how much weight is put on the task and $\psi = [\text{mixing proportion, perceptual uncertainty}]$ are parameters describing the subject's belief about the experiment given previous observations. Intuitively, a weight of $\alpha =$
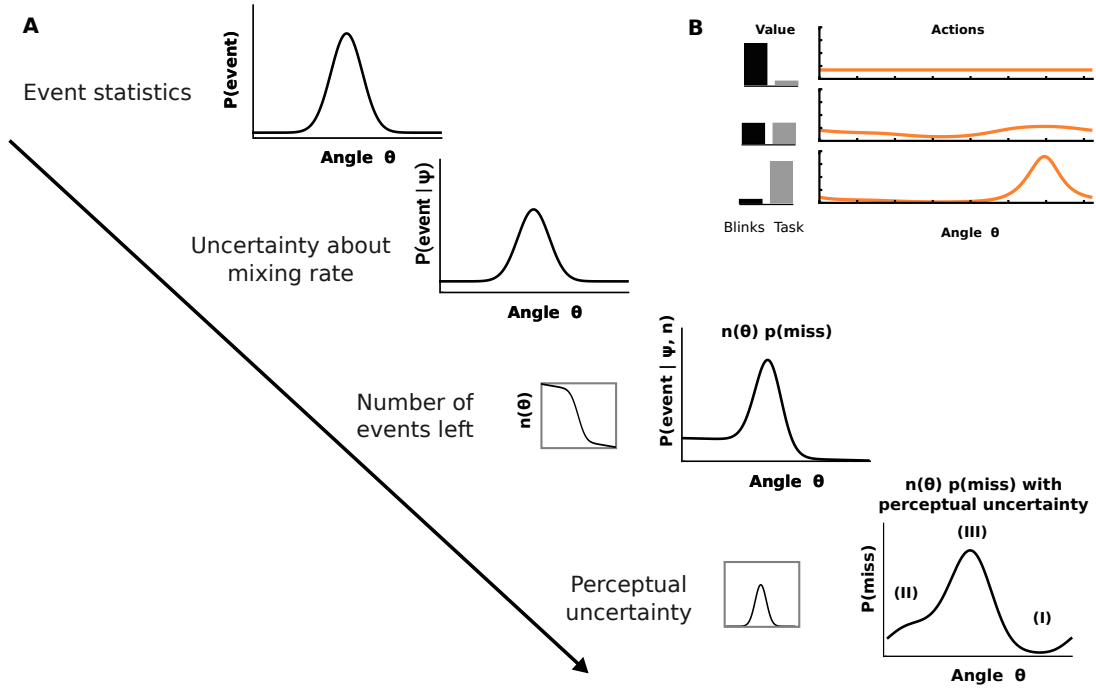
Figure 5.5: Schematic of the computational model for blinking. (A) Steps of the computational model describing subjects' belief about the event probabilities. (B) The belief structure is used for action selection while balancing multiple costs.

1 corresponds to blinking only because of the external task and therefore suppressing blinks for maintaining vision while a weight of $\alpha = 0$ corresponds to putting complete priority on maintaining healthy vision and thereby blinking only to lubricate the cornea (Figure 5.5).

The cost for blinking $c_e(\theta)$ is derived as the probability of missing an event and therefore depends on the current angle θ.

$$P(\text{miss}|\theta, \mu, \sigma^2, p) = n(\theta)P(\text{event at } \theta \mid \mu, \sigma^2, p) \tag{5.2}$$

where μ, σ^2 and p are the parameters of the mixture distribution generating the events, θ is a location during the lap, and $n(\theta)$ is the average number of events left at each location during a particular lap. The belief about the probability of missing an event depends on how many events have already been seen over the course of the current lap. In order to describe our subjects' belief, we computed the average number of events observed up to a certain point within the lap. This was done for all locations in the lap, parameterized by the angle θ:

$$n(\theta) = \bar{n} - \bar{n} \int_0^\theta P(\text{event at } \tau \mid \mu, \sigma^2, p)d\tau \,, \tag{5.3}$$

where $P(\text{event at } \tau)$ is the probability of an event at an angle τ (this corresponds to our mixture of $\mathcal{U}(0, 360)$ and $\mathcal{N}(180, 30)$ at a location τ), (μ, σ^2, p) are the parameters of the event generating distribution and $\bar{n}$ is the mean number of events per lap.

The subjects' belief about the costs for blinking is furthermore influenced by their uncertainty about the true underlying event distribution during the experiment. We accounted for the imperfect knowledge by assuming that subjects do not have access to the exact parameters

of the mixture distribution. In particular, we allowed the mixing proportion to be different from the true value that was used for event generation. In addition we hypothesized that subjects have perceptual uncertainty given previous observations about the exact parameters describing the Gaussian distribution from which the events are drawn. Overall, four parameters (mixing proportion p, cost trade off α, average blinking rate r, and the spread of the temporal uncertainty) were estimated from data across subjects. For parameter estimation, only data from the subjects' steady state behavior were used, as the changes in blinking behavior at the beginning of a block are related to learning the new event distributions. Thus, changes in the subjects' belief, which occur during learning the experimental statistics at the beginning of a new block, would affect these estimates. Changepoint analysis revealed that steady-state behavior was reached after 20 laps.

The starting point for modeling our subjects' belief is the probability distribution of the events (see Figure 5.5). Subjects do not have access to the true distribution, therefore we allowed for uncertainty about the mixing weight of the mixture distribution. Next, we incorporated recent observations by computing the average number of events seen during a lap. We used this to weight the probability of an event by the number of events left in a lap, $n(\theta)$. Finally, we accounted for uncertainty regarding the current location. Computationally, this was done by convolving the belief using a Gaussian distribution accounting for this position uncertainty. At the end of a lap (I) the event probability increases as it becomes more likely that the new lap has started. The probability of being in the next round saturates but the HEP region is not reached yet (II). Finally, there is a steep increase of event probability at the HEP region (III), which is followed by a decrease caused by both the decreasing probability of a single event as well as the decreasing number of events left in the lap.

5.3.3   Model results

We fitted our model to the aggregated blinking data. The result is shown in Figure 5.6A. Our model is capable of reflecting the characteristic course of blinking behavior. Moreover, we are now able to link both main effects, blink suppression as well as blink compensation, to computational quantities in our model: 1) Blink suppression follows from putting a higher value on the ongoing task. This leads to a higher proportion of blinks related to the ongoing external task. In the current experiment, this means that fewer blinks are carried out in those regions around the circular trajectory of the target where the probability of an event is high and thus the loss of task-related information is high. 2) The asymmetric shape of the curve is due to uncertainties in the subjects' belief where exactly the probability of an event is highest and the dynamic changes in this belief structure due to observing events over the course of a lap. Specifically, similar to a survival process, the fewer events have been observed during a lap, the higher the probability that an event is imminent. Finally, closer inspection of the probability of blinking in Figure 5.6A reveals a less steep decrease in blinking probability at around ninety degrees. In this region, the belief that all events from the previous lap have already occurred is small given the perceptual uncertainty and past event observations.

The results further suggest that the mixing rate of the mixture distribution is underestimated by our subjects ($\hat{p} = 0.58$ compared to $p = 0.8$). This is not surprising since learning a complex mixture distribution from noisy observations is a difficult task (e.g., Körding & Wolpert, 2004). Hence, our subjects know less about the task statistics compared to an ideal observer. Perceptual uncertainty due to observations of events as well as uncertainty from storing events in memory was estimated to $\sigma_{\text{perc}} = 29.6$ degrees.

A crucial property of the proposed model of the probability of blinking is that it allows
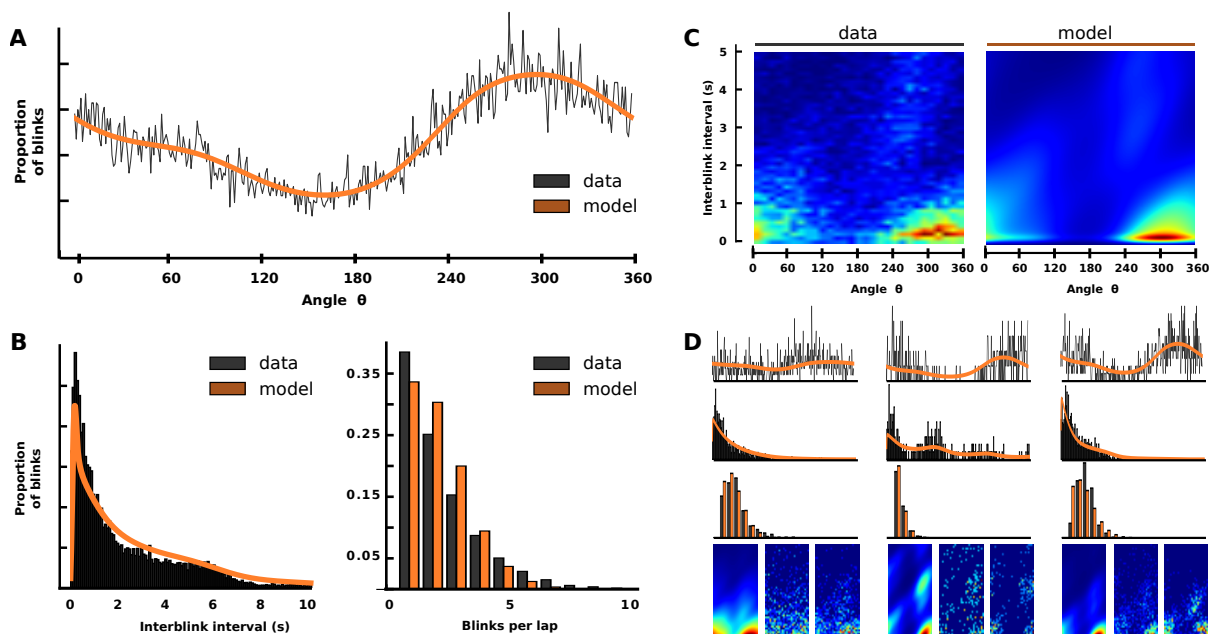
Figure 5.6: Computational model results. (A) Model results for blinking behavior. Number of blinks per degree (black line) and fitted model (orange line). (B) Model prediction of the temporal statistics of the blinking behavior. The distribution of the interblink intervals shown in the left panel. The histogram depicts all our subjects data. The right panel depicts the distribution of number of blinks per lap. Data is shown in black, model predictions are shown in orange. (C) Probability of blinking depending on the location on the circle (x-axis, from 0 to 360 degrees) and the time since the last blink (y-axis, from 0 to 5 seconds). Model predictions are shown in the right panel, subject data is shown in the left panel. We weighted very short IBIs using a cumulative Gaussian (see Hoppe & Rothkopf, 2016). This accounts for motor delays making two blinks occurring with close to zero IBI very unlikely. (D) Single participant results. From top to bottom: blinking proportion over the circle, IBI distributions, blinks per lap, and blinking probability in the angle and time domain are shown for three participants. Fitted model is depicted in orange, data is shown in black. The bottom row contains the theoretical distribution (left), samples from the theoretical distribution matched to the number of blinks of the individual subject (center), and our subjects' data (right). Results for all participants can be found in Figure 5.11 and 5.12 and Table 5.1.

deriving the temporal statistics of subjects' blinking data. The probability of a specific interblink interval d (IBI; the time between two consecutive blinks) can be computed as follows: Assume a blink occurring at a random location θ with probability $P(\text{blink at } \theta)$. Since the next blink is executed at $\theta + d$ for the next d time steps no blink occurs. Finally, d time steps after the first blink the next blink occurs. This yields

$$P(\text{blink at } \theta \mid d) = P(\text{blink at } \theta - d) \cdot$$
$$\cdot \left[\prod_{k=\theta-d+1}^{\theta-1} (1 - P(\text{blink at } \theta - k)) \right] \cdot P(\text{blink at } \theta) . \tag{5.4}$$

Figure 5.7 shows the application of Equation 5.4 to constant blink rates. For the changing blink rates we observed in our experiment, IBI distributions can be obtained by repeating this process for all values of θ and d and averaging over θ. Thus, our approach yields an analytic description of the distribution of IBIs as a product of geometric distributions, which has an intuitive link to blinking as sequential Bernoulli trials with varying probabilities. Note that previous empirical

Constant blink rate

$$p(blink) = const. = 16\frac{blinks}{min} = 0.27\frac{blinks}{s}$$

IBI follow geometric distribution

$$p(d) = (1 - p(blink))^{d-1} \cdot p(blink)$$

$p(1) = 0.73^0 \cdot 0.27 = 0.27$

$p(2) = 0.73^1 \cdot 0.27 = 0.197$

$p(3) = 0.73^2 \cdot 0.27 = 0.144$

$\qquad \vdots \qquad \vdots \qquad \vdots \qquad \vdots$

J-shaped
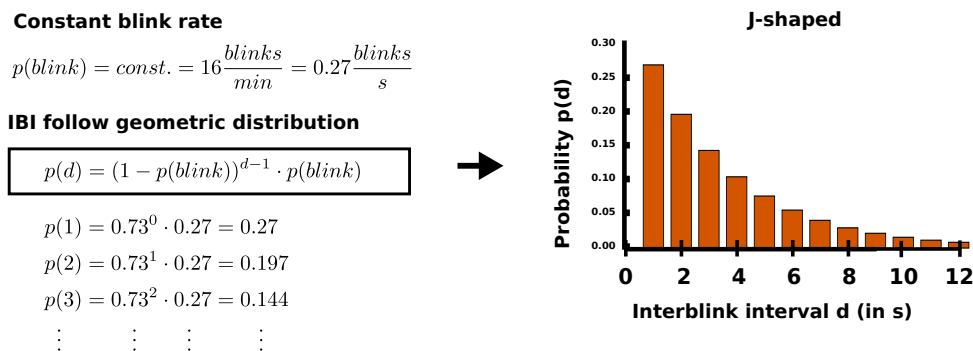
Figure 5.7: Derivation of the interblink intervals. For a constant blink of 16 blinks per minute the computations involved in reconstruction the IBI distribution are shown. The values for the first three seconds are given on in the left panel and correspond to the first three bars in the histogram shown in the right panel.

studies have proposed a power law distribution to describe the IBI distributions (Kaminer et al., 2011). However, our approach explicitly incorporates the non-stationary nature of the blinking rate in our task. By using this procedure, we are able to recover main characteristics of the temporal dynamics (Figure 5.6B–C) of subjects' blinks. Crucially, no further parameters need to be estimated.

As prior research on blinking has shown large individual differences, we fitted our model to each subject independently. We estimated the parameters for the cost trade-off as well as for the mean blinking rate on a subject-by-subject basis. The uncertainty about the temporal structure and the event generating distribution were assumed to be the same across participants. We therefore used the respective values estimated from the aggregated data. Results for three subjects are shown in Figure 5.6D. Individual blinking strategies were highly variable and IBI were qualitatively different across subjects. Remarkably, our model captures and explains this variability. As a consequence, we are able to link qualitative differences in behavior between individuals to quantitative differences in model components. The model suggests that the variability in blinking behavior can be ascribed to differences among subjects in motivational as well as physiological parameters. Specifically, three of the four IBI histogram shapes (J-shaped, bimodal, irregular) consistently found by previous studies can be explained through motivational differences across subjects combined with the temporal structure of the task. The analysis further reveals, that the least frequently found shape (symmetric) does not result from differences in the trade-off between task demands and physiological needs or differences in the general blinking rate.

Overall, participants detected 87 percent of all events. Two participants showed very low performance scores (smaller than 70 percent, $> 2\sigma$) and were excluded from the analysis. We tested whether subjects who put greater weight on the external task, i.e., a higher estimated α parameter, showed better performance (Figure 5.10). Linear regression revealed a significant relationship between the individually estimated values for α and the percentage of detected events, $r(23) = .51, p = .01$. Importantly, this means that subjects blinked more strategically instead of less often, as we found no significant connection between smaller blinking rates and higher performance, $r(23), = 0.38, p = .07$.
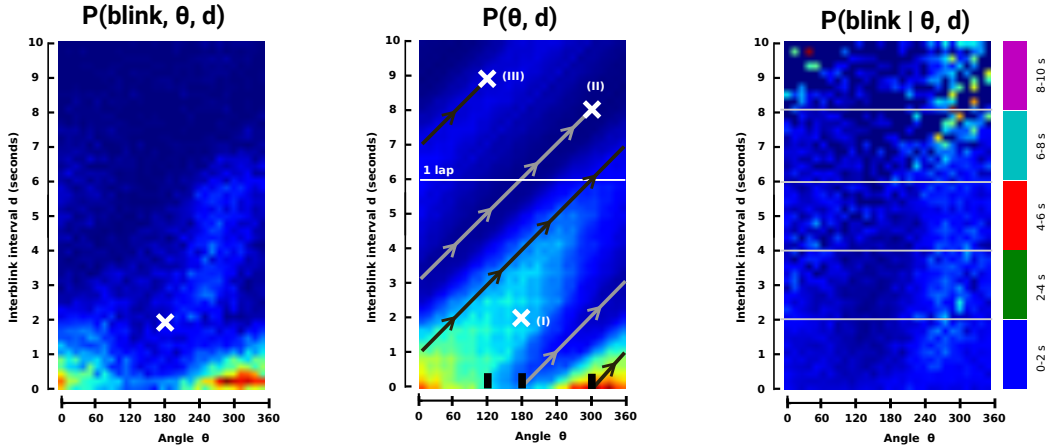
Figure 5.8: Derivation of the blink probability $p(blink|\theta, d)$. Left panel: The blinking density $P(blink, \theta, d)$ is shown depending on the location within a lap (x-axis, from 0 deg to 360 deg) and the time since the last blink (y-axis, from 0 s to 10 s). For example, the value at the white cross depicts the proportion of blinks occurring at location 180 deg in the lap and simultaneously occurring 2 s after the last blink. Importantly, this implies that the last blink occurred at location 60 deg ($180 \deg - 2s \cdot 60 \deg /s$), since the angular velocity of the visual target in the experiment was 60 deg / s. Center panel: Probability distribution of how often each combination of a position in the lap and a time since last blink is reached in our experiment. Not all combinations of (θ, d) are equally likely as they are heavily affected by the uneven distribution of blinks over the course of a lap (see Figure 5.3). For illustration purposes, three blinks (depicted by white crosses) are shown together with the corresponding two interblink intervals (gray: interval between Blink I and Blink II; black: interval between Blink II and Blink III). Blink I occurs at location 180 degrees. As can be seen by following the gray line, the same location (180 degrees) is reached again exactly after six seconds. After eight seconds Blink II is executed and following the black line leads to Blink III. Right panel: The blink probability for each combination of location in the lap and time since the last blink is shown. This can be computed as $P(blink, \theta, d)/P(\theta, d)$.

5.3.4   Constant costs for blink suppression

We tested whether the costs for blinking depend on how much time has passed since the last blink since our computational blinking model assumes that the costs for blink suppression c_s are constant and therefore neither depend on the current location θ in the lap nor on the time since the last blink d (interblink interval; IBI). The first independence follows from the definition of the cost as task independent, but the second is not apparent. While we do not claim that costs do not increase at all with time, we provide evidence that for the durations between two blinks observed in the current experiments, costs indeed did not increase significantly. The statistical independence assumption in our model can be formalized as:

$$P(blink|\theta, d) \stackrel{?}{=} P(blink|\theta). \tag{5.5}$$

This formalization provides a way to test whether costs for blink suppression depend on the IBI. If the assumption holds, we expect $P(blink|\theta)$ not to be affected by the value of d, i.e., the marginal distribution of blinks over a lap should not be different if the time since the last blink is 4 s instead of 2 s, for example. On the other hand, if costs for blink suppression increase with the time passed since the last blink, we expect a more even blink distribution for longer IBIs, as with greater costs for blink suppression less value is given to the task. We tested this empirically using our data by comparing the conditional probability of blinks at a location θ for

the following ranges for d:

$$P(blink|\theta, 0s < d < 2s) \overset{?}{=} P(blink|\theta, 2s < d < 4s) \overset{?}{=} \cdots \overset{?}{=} P(blink|\theta, 8s < d < 10s). \quad (5.6)$$

In order to test our hypothesis by applying Equation 5.5 and 5.6, we need to compute the conditional probability $P(blink|\theta, d)$. Using basic probability theory we obtain:

$$P(blink|\theta, d) = P(blink, \theta, d)/P(\theta, d) \quad (5.7)$$

where $P(blink, \theta, d)$ is the joint probability of our blinking data (Figure 5.8, left panel) and $P(\theta, d)$ is the probability of being at a location θ after d seconds have passed since the last blink. Hence, we must normalize the blinking probability at a location θ and interblink interval d by the probability of reaching location θ d seconds after the last blink. In order to compute the conditional probability $P(blink|\theta, d)$ we first need to obtain $P(\theta, d)$. We can compute this quantity from our blinking data as follows (see Figure 5.8, center panel): For each blink (tuple of location θ and IBI d) we can reconstruct all the points (θ, d) that have been visited since the last blink. In our experiment, the target stimulus moves with constant angular velocity on a circular trajectory. The velocity was 60 deg / s. Hence, we move through the parameter space along a line with slope 60 deg/s (see gray and black lines in Figure 5.8 (center panel)). The probability $P(\theta, d)$ is the normalized number of times each point Q (θ_q, d_q) is visited. $P(\theta, d)$ clearly is not uniform, instead, it is strongly influenced by where blinks occurred in the lap. For example, fewer blinks occurred in the HEP region (120 - 240 deg). Hence, the location ($\theta = 60$, $d = 4$) is not visited often (dark blue value), as the point is only reached if the last blink occurred in the HEP region (shown by the gray line). We now can compute $P(blink|\theta, d)$ according to Equation 5.7 . The results are shown in Figure 5.8 (right panel).

Our initial question, whether costs for blink suppression can be assumed to be independent of the time since the last blink, can be tested using $P(blink|\theta, d)$. We computed $P(blink|\theta)$ for five different ranges of d (see Figure 5.9). Our results indicate that blinking behavior does not change depending on the time since the last blink. This suggests that for interblink intervals between 0 and 10 s, which comprises 93 % of all blinks in our experiment, costs can be assumed to be constant.

We additionally used our behavioral results in Figure 5.9 to check whether longer times since the last blink lead to differences in the trade-off between costs for blink suppression and costs for blink execution. Therefore, we divided the lap into three equally sized areas: before HEP (0 - 120 deg), HEP (120 - 240 deg), and after HEP (240 - 360 deg). If the cost for blinking increased with time, then blinks with greater interblink intervals should less reflect the task related costs for greater IBIs (Figure 5.9, upper left panel) leading to smaller differences in the proportion of blinks across the areas. If, on the other hand, costs are constant, we would not expect a change in the relative distribution of blinks across the three areas (Figure 5.9 B, upper right panel). Our results (Figure 5.9 B, lower panel) clearly favor constant cost. Blinking was not suppressed less in the HEP region for greater IBIs.

5.4 Discussion

In our study, we investigated how blinking behavior is related to internal costs and environmental visual demands. In particular, while prior research has provided various accounts of the link between blinks and task demands, little work quantifying this connection exists. We created
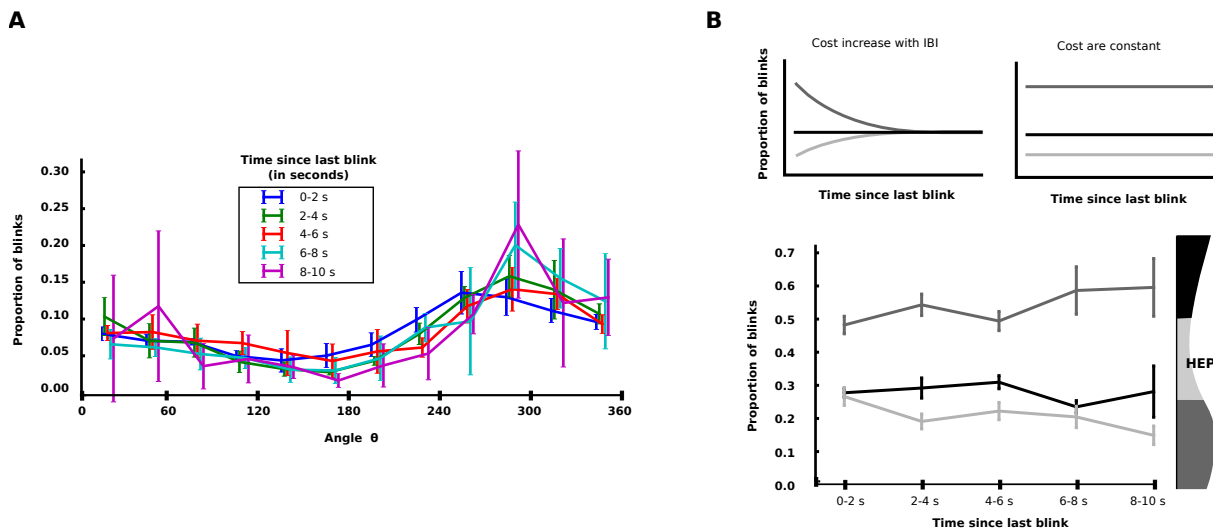
Figure 5.9: Evidence for constant costs of blink suppression for the observed IBI. (A) Proportion of blinks over the course of the laps computed separately for different interblink intervals. The colors correspond to the areas Figure 5.8 (right panel). For example, the blue line depicts the mean blinking proportion given that the time since the last blink is between 0 and 2 s. We used a bin size of 1/3 second for the time since the last blink (0 s - 0.33 s, ... , 1.66 s - 2s) and computed the mean for the six resulting values at each location. Error bars correspond to the standard deviation. (B) Schematic course of blinking proportions for the three areas (Before HEP, HEP, After HEP). If costs increased with IBI (left panel), blinking would become less effected by task related costs. For constant costs (right panel) this should not be the case. The lower panel shows the proportion of blinks dependent on the time since the last blink. Overall, this provides clear evidence, that costs for blink suppression were constant over the IBI durations in our experiments.

an event detection task tailored to studying blinking behavior quantitatively. Subjects detected events while fixating a moving dot. The event probability was linked to the spatial location of the dot, a regularity subjects could learn over the course of the experiment. Our design provides full control and knowledge regarding the temporal statistics of the visual input as well as the reward structure of the task. In particular, the consequences of blinking on the loss of information can be quantified.

The probabilistic design of the experiment allowed developing a computational model of blinking behavior. The basic assumption is that blinking is the consequence of a trade-off between an internal urge to blink in order to maintain healthy vision and external task requirements of not blinking when crucial information needs to be acquired from the environment. Given subjects' perceptual and memory uncertainties about the event generating process in the experiment, it is possible to quantify the cost of blinking in terms of task performance, i.e., the probability of detecting an event.

The behavioral data show two main effects: First, blinking was significantly suppressed in the high event probability (HEP) region, i.e. where most events occurred. Our computational model results suggest that this effect can be explained in terms of minimal loss of task-relevant information. This result is in accordance with prior research that reported a connection between blink suppression and task performance (Gregory, 1952). Also, in classical psychophysical experiments blinks have been shown to occur around the time of response (Oh, Han, Peterson, & Jeong, 2012) and toddlers watching movies showed suppressed blinking at scenes containing affective and physical events (Shultz, Klin, & Jones, 2011). Thus, subjects who weighted costs associated with the external task more tended to blink more strategically and therefore avoided
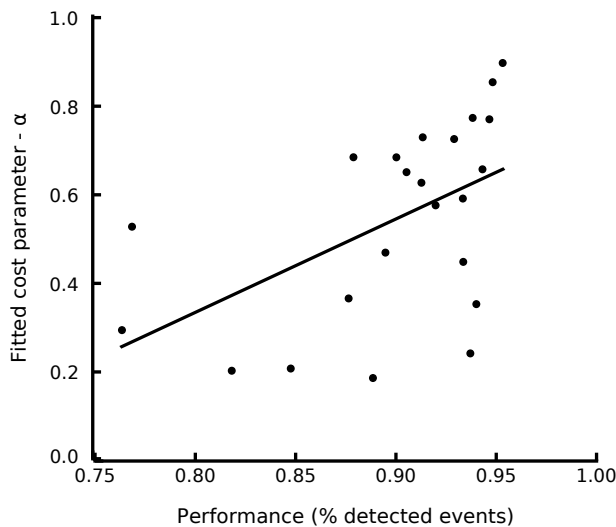
Figure 5.10: Relationship of the α parameter fitted to individual subjects and task performance. Solid line depicts the regression curve.

to blink in the HEP region.

Second, more blinks occurred after the HEP region compared to before. Although the probability of missing an event is symmetrical around the peak of the event generating distribution, the observed blinking pattern is not. This asymmetry is predicted by our model if we account for observations made over the course of a lap. The event probability and therefore the probability of missing an event during a blink is proportional to the number of events left. Hence, blinking earlier in a lap leads to greater loss of detection performance. With every observed event the number of events left in a lap decreases. Finding this asymmetry of blinking strategies in our data reveals two properties of our subjects' information processing: subjects learned the number of events per lap and they were able to dynamically incorporate recent information about observed events for deciding when to blink. Compensation between episodes of suppression has been reported in the past (Gregory, 1952). While other studies argue that blinks take on a role of breakpoints to facilitate mental processing (Nakano, Kato, Morito, Itoi, & Kitazawa, 2013), in our study, blinking was well described in terms of collecting maximal task-relevant information.

The distribution of interblink intervals (IBI) is highly variable across humans (Ponder & Kennedy, 1927). Four different shapes have been identified repeatedly: J-shaped, irregular plateau, bimodal, and symmetric (Doughty, 2002; Ponder & Kennedy, 1927; Naase et al., 2005). Here, we showed that three of these shapes arise as an immediate consequence of the trade-off between internal and external task costs. Our model is capable of capturing the characteristics of the first three types (see also individual model fits for all data in Figure 5.11 and 5.12). One participant (subject 19 in Figure 5.12), however, showed a symmetric distribution which was not well fitted by our model. One explanation could be that the symmetric shape (the least often according to Ponder & Kennedy, 1927) does not arise from interacting with the task but from physiological properties that were not captured by our model.

Few models have been developed to describe blinking behavior. In their urge model (Berman et al., 2012) assumed a linear increase in urge when blinking is suppressed. However, the model does not account for any external task related influences. In another study Moraitis and Ghosh (2014) proposed that blinks are generated by an oscillating blink generator. However, both studies used voluntary blink suppression. Here we presented a computational model that explicitly included task-related goals as well as intrinsic costs, thereby building a natural connection to the reward-related learning literature involving dopamine. Hence, the model can be applied to a broader area of investigations as long as some properties about the environmental statistics

are known. While we presented results for a psychophysical detection task, our approach is not limited to simple stimuli. Recent developments in machine learning and deep neural networks (e.g., LeCun, Bengio, & Hinton, 2015) have paved the way for retrieving statistical regularities even in complex and dynamic visual scenarios. In combination with these methods, our blinking model can readily be applied to many real-world problems. Better understanding and in particular, quantitative insights, into human blinking behavior are also relevant for building technical aid systems (see Krolak & Strumillo, 2008, for example), and detecting mental states during critical tasks to prevent accidents (Smilek, Carriere, & Cheyne, 2010).

Table 5.1: Individual parameter estimates for all subjects.

Subject	Performance	Blink rate	Blink rate estimate	Cost trade-off
1	0.879	6.37	6.54	0.68410
2	0.905	6.95	7.30	0.65063
3	0.942	4.42	4.50	0.65705
4	0.933	24.42	24.86	0.59144
5	0.847	38.95	39.03	0.21002
6	0.946	7.67	7.98	0.76929
7	0.948	5.50	5.56	0.85248
8	0.936	7.75	7.77	0.24398
9	0.919	4.20	4.36	0.57620
10	0.818	22.70	22.70	0.20509
11	0.640	18.30	18.21	0.44432
12	0.888	14.50	14.46	0.18871
13	0.894	27.35	27.27	0.47012
14	0.876	12.92	12.92	0.36731
15	0.939	26.95	26.84	0.35454
16	0.912	10.77	11.07	0.62687
17	0.938	7.25	7.29	0.77236
18	0.928	6.07	6.07	0.72495
19	0.933	16.65	16.78	0.44933
20	0.587	11.68	12.00	0.82670
21	0.769	6.90	6.90	0.52826
22	0.900	29.30	30.05	0.68391
23	0.764	30.22	30.21	0.29622
24	0.952	6.50	6.79	0.89559
25	0.913	2.92	2.93	0.72884
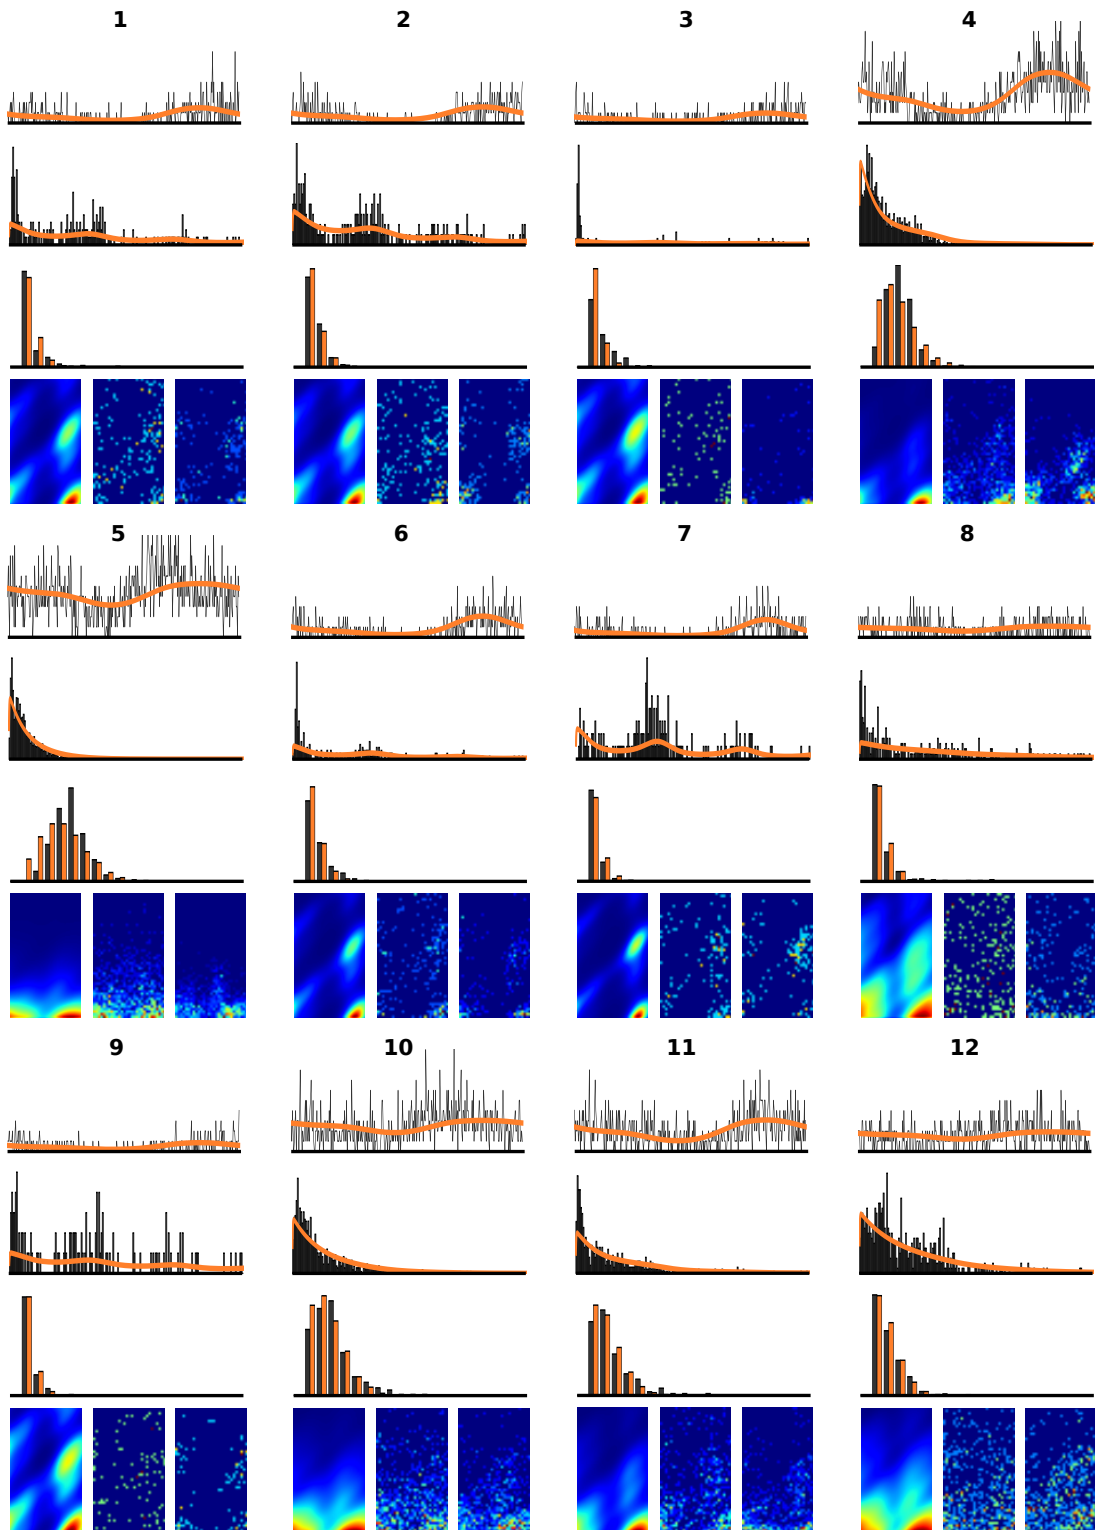
Figure 5.11: Individual data and model fits for all participants (1 – 12). For each subject the blinking rate over the circle (first row), the distribution of inter blink intervals (second row), and the distribution of blinks per lap (third row) are shown. The fourth row depicts the density of blinks in the *location on the circle* (x-axis) – *time since last blink* (y-axis) plane.
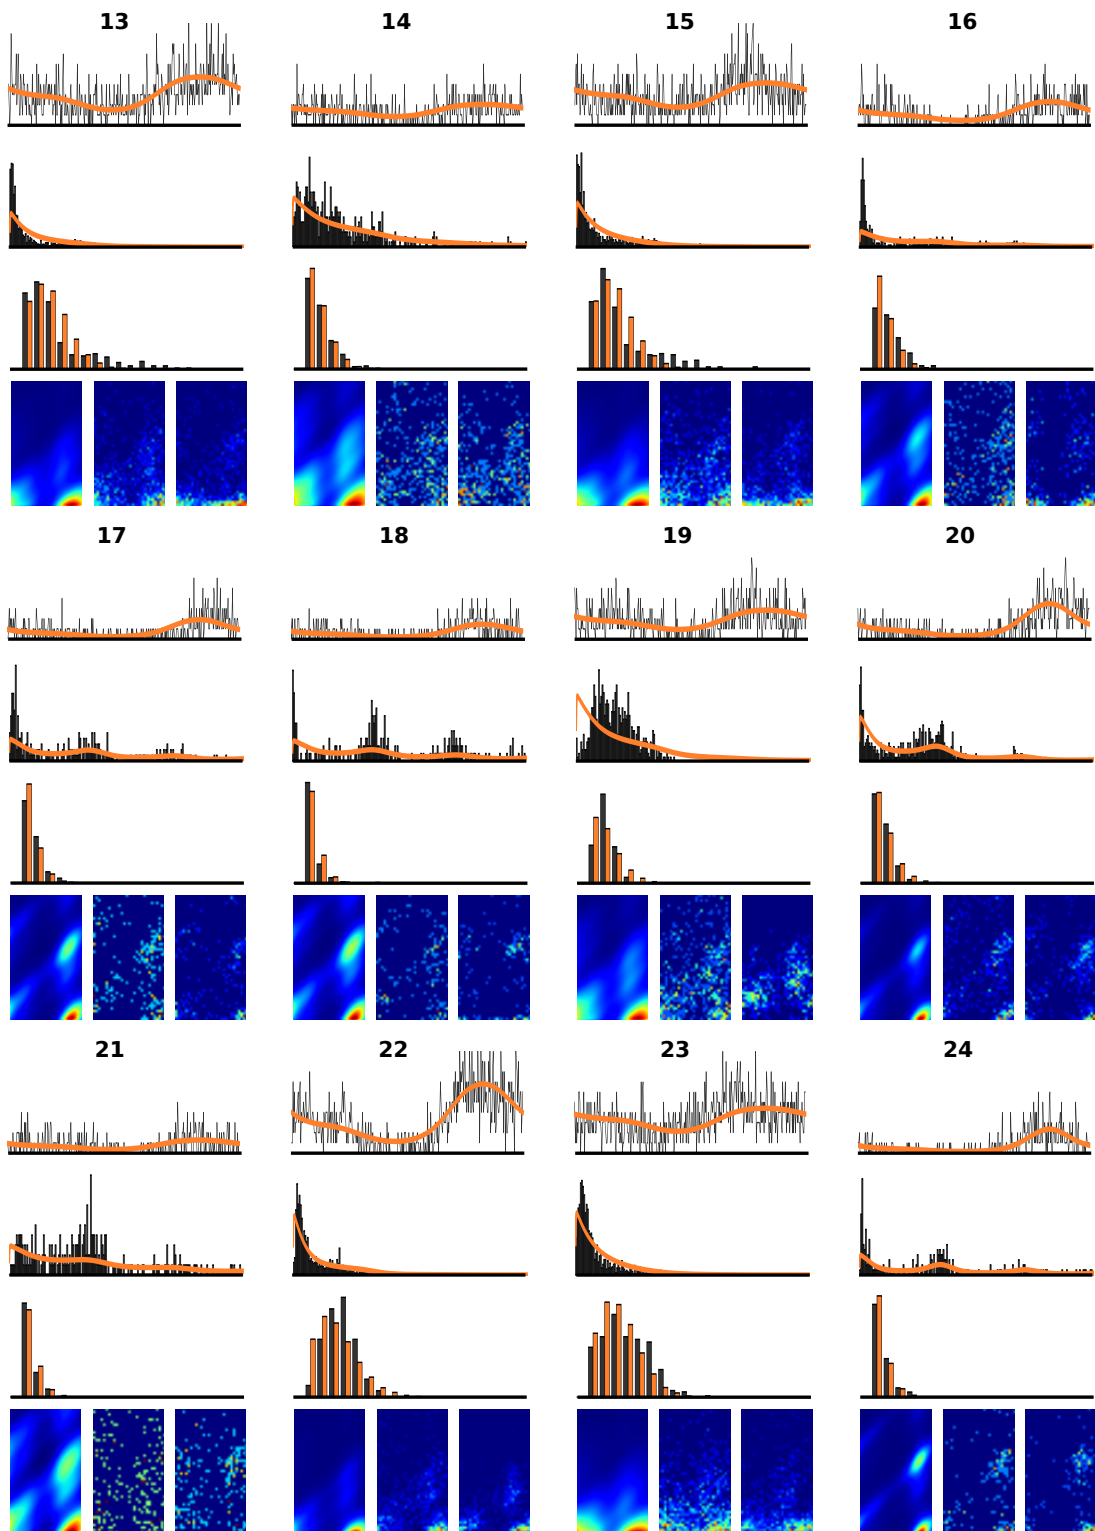
Figure 5.12: Individual data and model fits for all participants (13 – 24). Same as Figure 5.11.

Learning rational temporal eye movement strategies

6.1 Introduction

During active behavior, humans redirect their gaze several times every second within the visual environment. Where we look within static images is highly efficient as quantified by computational models of human gaze shifts in visual search (Najemnik & Geisler, 2005) and face recognition tasks (Peterson & Eckstein, 2012). But when we shift gaze is mostly unknown despite its fundamental importance for survival in a dynamic world. So far, the emphasis of most studies on eye movements has been on spatial gaze selection. While these studies provide insights into where humans look in a visual scene, the ability to manage complex tasks cannot be explained solely by optimal spatial gaze selection. Even if relevant spatial locations in the scene have been determined these regions must be monitored over time to detect behaviorally relevant changes and thus keeping information updated for action selection in accordance with the current state of the environment. Successfully dealing with the dynamic environment therefore requires intelligently distributing the limited visual resources over time. For example, where should we look, if the task is to detect an event occurring at one of several possible locations? In the beginning, if the temporal regularities in a new environment are unknown, it may be best to look quickly at all relevant locations equally long. But with more and more observations of events and their durations, we may be able to use the learned temporal regularities. To keep the overall uncertainty as low as possible, fast-changing regions (e.g., streets with cars) should be attended more frequently than slow-changing regions (e.g., walkways with pedestrians).

Although little is known about how the dynamics in the environment influence human gaze behavior, empirical evidence suggests sophisticated temporal control of eye movements in object manipulation (Ballard, Hayhoe, & Pelz, 1995; Johansson, Westling, Bäckström, & Flanagan, 2001) and natural behavior (Hayhoe & Ballard, 2005; Land & Tatler, 2009). A common observation is that gaze is often predictive both of physical events (Diaz, Cooper, Rothkopf, & Hayhoe, 2013) and events caused by other people (Flanagan & Johansson, 2003), and studies comparing novices' and experienced sportsmen's gaze coordination suggest that temporal gaze selection is learned (Land & McLeod, 2000). However, the temporal course of these eye movements and how they are influenced by the event durations in the environment has not been investigated so far. We argue that understanding human eye movements in natural dynamic settings can only be achieved by incorporating the temporal control of the visual apparatus. Also, modeling of temporal eye movement allocation is scarce with notable exceptions (Hayhoe & Ballard, 2014), but while ideal observers (Najemnik & Geisler, 2005; Peterson & Eckstein, 2012; Chukoskie et al., 2013) model behavior after learning is completed and usually exclude
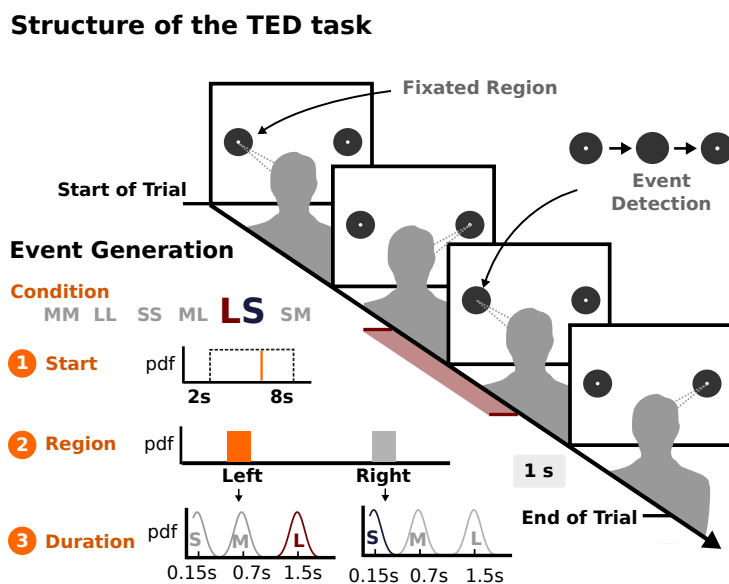
Figure 6.1: Experimental procedure for the temporal event detection (TED) task. The procedure of a single trial is shown on the right. The left panel depicts the steps involved in generating an event for a trial in the Condition LS.

intrinsic costs, reinforcement learning models based on Markov decision processes (Chukoskie et al., 2013; Hayhoe & Ballard, 2014) can be applied to situations involving perceptual uncertainty at best only approximately.

Here we present behavioral data and results from a computational model establishing several new properties of human active visual strategies specific to temporally varying stimuli. In particular, we address the question, what factors specifically need to be accounted for to predict human gaze timing behavior in response to environmental events, and how learning proceeds. To this end, we created a novel temporal event detection (TED) task to investigate how humans adjust their eye movement strategies when learning temporal regularities. Our task enables examining the unknown temporal aspects of eye movement strategies by isolating them from spatial selection and the complexity that arises from natural images. Our results are the first to show that human eye movements are consistently related to temporal regularities in the environment. We further provide a computational model explaining the observed gaze behavior as an interaction of perceptual and action uncertainties as well as intrinsic costs and biological constraints. Surprisingly, based on this model alone, the learning progress is accounted for without any additional assumptions only based on the scalar law of biological timing, allowing crucial novel insights into how observers alter their behavior due to task demands and environmental regularities.

6.2 Experimental design

6.2.1 The temporal event detection task

Participants switched their gaze between two regions marked by gray circles to detect a single event (Figure 6.1). The event consisted of the little dot within one of the circles disappearing (depicted as occurring on the left side in this trial). Participants confirmed the detection of an event by indicating the region (left or right) through a button press. The event duration
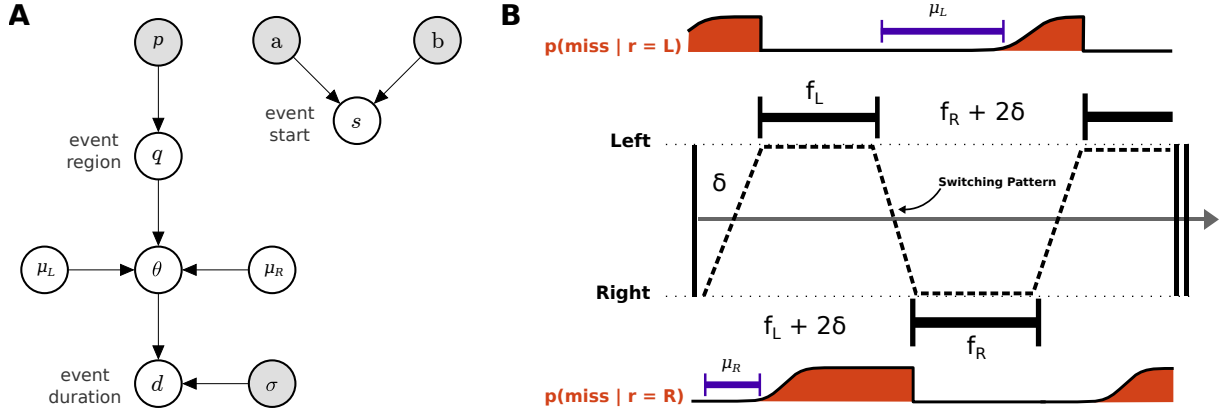
Figure 6.2: Derivation of the statistics in the TED task. (A) Probabilistic graphical model of the event generating process. (B) Probability of missing an event. The condition for the shown segment is LM, i.e., long event durations at the left region and medium event durations at the right region. The probability of missing the event (orange area) is shown for an arbitrary switching pattern (dashed line), where the left region is fixated for a duration f_L and the right region for a duration f_R. Events are missed if the event's region is not fixated for the entire duration of the event. For long event durations the offset of the probability of missing an event (purple lines) is greater than for short event durations. As a consequence, the time between two fixations can be greater if the event duration is long without missing events.

(how long the dot was invisible) was drawn from one of three probability distributions (small, medium, and long). Overall, six different conditions resulting from the combination of three event durations and two spatial locations were used in the experiment (SS, MM, LL, SM, SL, MM, LL). A single block consisted of 20 consecutive trials with event durations drawn from the same distributions. For example, in a Block LS, all events occurring at the left region were long events with a mean duration of 1.5 s, all events occurring at the right region were short events with a mean duration of 0.15 s. A single event for the Condition LS was generated as follows (lower left panel): first, the start time for the event was drawn from a uniform distribution between 2 and 8 s. Next, the region where the event occurred was either left or right with equal probability. Finally, the duration of the event was drawn dependent on where the event occurred. Given, that in the example the condition was LS and the event was determined to occur at the left region, the event duration was drawn from a Gaussian distribution with mean 1.5 s (long event).

6.2.2 Statistics of the temporal event detection task

Events for the TED task were generated according to the graphical model presented in Figure 6.2A. An event was defined as a tuple (s, d, q) with the generative process

$$s \sim \mathcal{U}(a, b) \tag{6.1}$$

$$q \sim \mathcal{B}(1, p) \tag{6.2}$$

$$\theta = \begin{cases} \mu_L & \text{if } q = 0 \\ \mu_R & \text{if } q = 1 \end{cases} \tag{6.3}$$

$$d \mid q \sim \mathcal{N}(\theta, \sigma^2), \tag{6.4}$$

where the random variable s denotes the event's starting time within a trial, which was distributed uniformly between $a = 2s$ and $b = 8s$. The region q where the event occurs (q = 0: left; q = 1: right) was drawn from a binomial distribution and the event's duration d from a Gaussian distribution. The standard deviation of the event durations σ was fixed to a small value (100 ms). Events were equally probable at both regions ($p = 0.5$). The mean of the event duration depended on the region where the event occurred. The mean event durations for the right region (μ_R) and the left region (μ_L) were fixed over the course of a block according to the respective condition. For example, in a block belonging to Condition LM (Long-Medium) μ_L and μ_R were set to 1.5 s and 0.75 s, respectively. Since the reward function for our task is known we are able to derive the relationship between switching patterns and detection probability (see the derivation of the model as well as Figure 6.2B).

6.2.3 Procedure

The TED task required participants to complete 30 experimental blocks, each consisting of 20 consecutive trials. Subjects were instructed to detect a single event on each trial and press a button to indicate the detection of the event, which could occur at one of two possible locations with equal probability. The two regions (distance 35 deg) were presented on a Flatron W2242TE monitor (1680×1050 resolution, 60 Hz refresh rate). Each region was marked with a gray circle (0.52 deg in diameter). The contrast of the target dot was adjusted to prevent extrafoveal vision. Eye movement data were collected using the SMI Eye Tracking Glasses (60 Hz sample rate). Calibration was done using a three-point calibration. For detection of saccades, fixations, and blinks dispersion-threshold identification was used.

6.3 Computational models for the temporal event detection task

All of the presented models are based on the methods described in Chapter 3, i.e., as a trade-off between task-related reward and physiological costs. In order to quantify the relationship between eye movements and reward, we derive the probability of missing an event (Figure 6.3). Although the gaze sequences in the task comprise many switches, it is sufficient to only consider a single fixation at each side to find the optimal behavior (Figure 6.2B)[1].

We first derived the probability that an event ends at time e at region q based on the generative process. This corresponds to the joint distribution $p(e, q)$ which can be factorized into $p(e, r) = p(e \mid r)p(r)$. For any event, the ending time is the sum of the starting time and the duration. The probability of an event ending at a specific time e given the region r can be computed by the sum of s and d, which results in the convolution of the respective probability density functions

$$p(e \mid q) \;\;=\;\; p(s) * p(d \mid q) = \frac{1}{b-a} \left[\psi_\theta \left(\frac{b-e}{\sigma} \right) - \psi_\theta \left(\frac{a-e}{\sigma} \right) \right]. \tag{6.5}$$

The last identity follows from the convolution of a Gaussian and a Uniform distribution, where ψ_θ denotes the cumulative distribution of a normally distributed variable centered at θ.

---

[1]After executing the two fixations (left and right) the environmental state is the same as before. Hence, the optimal pattern for these two switches is also the optimal pattern for the next two switches and so on (see also Chapter 3).

Figure 6.3: Eye movement strategies for the TED task. The relationship between switching patterns and the probability of missing the event in a single trial is shown for two conditions. For the Condition MM (top) the probability of missing an event is lower for faster switching (short fixation durations at both regions). For the Condition LS (long events at the left region, short events at the right region) the probability of missing an event can be decreased by longer fixating the region with short events, i.e. the right region.

We formalized switching patterns (SP) as a tuple (f_L, δ, f_R), where f_L and f_R are the fixation durations for the left and right region, respectively, and δ is the time needed to move focus from one position to the other (i.e., the duration of a saccade). Using our results for the distribution of event ending times, we derived the probability of missing an event when performing a switching pattern (see Figure 6.2B). For a single region (e.g., R), the time between two consecutive fixations is $f_L + 2\delta$. Events both starting and ending within this time range at region R are regarded missed. The probability of missing an event at region R while fixating L can therefore be computed as

$$p(\text{event missed} \mid f_L, \delta, f_R, q = \text{R}) \quad = \quad p(e \le f_L + 2\delta \mid q = \text{R}) = \int\limits_{0}^{f_L + 2\delta} p(e \mid q = \text{R})de \,. \ (6.6)$$

The probability of missing an event at region L can be computed analogously. The event generating process implies that events are mutually exclusive at the two regions. Therefore, we can calculate the probability of missing an event when performing the switching pattern (f_L, δ, f_R) by marginalizing over the regions

$$p(\text{event missed} \mid f_L, \delta, f_R) \quad = \quad \sum_{q \in \{R,L\}} p(\text{event missed} \mid f_L, \delta, f_R, q)p(q) \,. \qquad (6.7)$$

The derived probability of missing an event only covers the sequence of two fixations. Thus, the total time covered is $f_L + f_R + 2\delta$. To compute the probability of missing an event over the

course of the entire trial and to compare different switching patterns, we divided the probability of missing an event over the course of the switching pattern by the total time covered by the switching pattern. As a result, we get the expected probability of missing an event in a single time step Δt associated with a certain switching pattern (SP)

$$P(\text{event missed} \mid f_L, \delta, f_R)/\Delta t = \frac{P(\text{event missed} \mid f_L, \delta, f_R)}{f_L + f_R + 2\delta} \, . \tag{6.8}$$

6.3.1 Computational models of gaze switching behavior

Based on the statistical relationships derived from the event generating process we now propose different models for determining the optimal switching pattern assuming that the event durations are learned. The time needed to perform a saccade is small compared to the fixation durations and by and large determined by its amplitude as reflected by the main sequence relation. Thus, δ was treated as a fixed value and omitted from the equations. Further, the probability of detecting the event when performing a certain switching pattern can be computed from $P(\text{event missed}|f_L, \delta, f_R)$ and can be expressed as a task related reward function $r_{task}(f_L, f_R)$.

We started by deriving an ideal observer which chooses switching patterns f_L, f_R solely by maximizing r_{task}

$$SP_{\text{IO}} = \arg \max_{f_L, f_R} r_{\text{task}}(f_L, f_R) \, . \tag{6.9}$$

Ideal observer models are frequently used to investigate human behavior involving perceptual inferences. However, a perceptual problem may have multiple ideal observer models depending on which quantities are assumed to be unknown. Ideal observers usually do not consider costs but provide a general solution to an inference problem. Therefore, we extended our ideal observer in order to obtain more realistic models.

First, by using the true event durations d the ideal observer assumes perfect knowledge about the generative process. Especially, the mean duration for both regions (μ_L and μ_R) are assumed to be known. This is clearly not the case at the beginning of the blocks and even after learning perceptual uncertainty and, crucially, the scalar property prevents complete knowledge. While the former is irrelevant for the case of learned event durations, the latter is not. Building on findings regarding the scalar property we augmented the ideal observer by including signal depending Gaussian noise on the true variance of event durations σ^2. As a result the distribution $d|q$ becomes $\mathcal{N}(\theta, \sigma^2 + wf^2\theta)$, where wf is the Weber fraction.

Second, the ideal observer only considers task performance for determining optimal switching patterns. It implicitly assumes that switching patterns do not differ with respect to other relevant metrics. However, switching at a very high frequency consumes more metabolic energy and can be associated with further internal costs. Therefore, we hypothesized that human behavior is the result of trading off task performance and intrinsic costs

$$SP_{\text{IO+Cost}}(\alpha, \tau) \quad = \quad \arg \max_{f_L, f_R} [r_{\text{task}}(f_L, f_R) + \alpha \cdot r_{\text{eye}}(f_L, f_R, \tau)] \tag{6.10}$$

where α quantifies eye movement costs in units of r_{task}. To our knowledge the exact shape of these costs is unknown. We hypothesize that eye movement costs are greater for higher switching

frequencies and are therefore connected to fixation durations. We used the positive range of a bell curve as a cost function

$$c_{eye}(f_L, f_R, \tau) = \exp\left[-\frac{(f_L + f_R)^2}{\tau^2}\right].$$ (6.11)

where τ controls how quickly costs increase with switching frequency. Using c_{eye} we can compute a reward function for switching frequencies as $r_{eye} = (\max c_{eye}(f_L, f_R, \tau)) - c_{eye}(f_L, f_R, \tau)$.

So far, all models have assumed the switching patterns to be deterministic. However, humans clearly are not capable of performing an eye movement pattern without variability. In addition, they have shown to take motor variability into account when choosing actions, hence they can be described as ideal actors. The ideal actor model takes into account that there is variation when targeting a certain switching pattern f_L, f_R. This property can be considered when no longer treating fixation durations as fixed but as samples drawn from some distribution. We chose an exponential Gaussian, centered at the targeted durations f_L and f_R.

$$p(f \mid f_X) = \text{ex-Gaussian}(\zeta, \eta, \omega).$$ (6.12)

for X in (L,R), whose expected value $\mathbb{E}[f] = \zeta + \eta$ is equal to the respective target fixation duration f_X. Given f_X the shape parameters of the distribution were determined using $f_X = \zeta + \eta$ together with $\zeta = w_0 + w_1\eta$. The regression parameters w_0 and w_1 as well as ω were estimated from the data. The expected value of $r_{\text{task}}(f_L, f_R)$ is then computed by marginalizing out the variability of f_L, f_R

$$r_{\text{task}}^{act}(f_L, f_R) = \int\limits_{f_0} \int\limits_{f_1} p(f_0 \mid f_L)p(f_1 \mid f_R)r_{\text{task}}(f_0, f_1)df_1 df_0.$$ (6.13)

Finally, the abilities of the motor system and the speed of information processing constrain switching patterns. As a consequence, targeted fixations durations have a lower bound. To account for this constraint, we propose a threshold using a cumulative Gaussian

$$p(f_X) = \psi(\beta, \gamma^2)$$ (6.14)

where the mean β was estimated from the data and γ^2 was fixed to a small value. Combining all components yields the Bounded Actor

$$SP_{\text{Bounded Actor}}(\alpha, \tau, \beta, w_0, w_1, \omega) =$$ (6.15)

$$= \arg\max_{f_L, f_R}\left[p(f_L)p(f_R)\left(\int\limits_{f_0}\int\limits_{f_1} p(f_0 \mid f_L)p(f_1 \mid f_R)r_{\text{task}}(f_0, f_1)df_0 df_1 + \alpha \cdot r_{\text{eye}}(f_L, f_R, \tau)\right)\right].$$
(6.16)

6.3.2   Learning the temporal statistics

All models discussed so far have assumed perfect knowledge of the event durations despite perceptual uncertainty. They are therefore only applicable to the steady state case since event durations were unknown at the beginning of a block and had to be learned over the course of the block. Learning was formulated as updating the models' belief about the event durations

through observations. From a Bayesian perspective, this can be seen as using the observation to transform a prior distribution to a posterior distribution

$$p(\theta \mid o) = \frac{p(o \mid \theta)p(\theta)}{p(o)} \,, \tag{6.17}$$

where θ denotes the mean of the event duration and o a sequence of observations. The updated belief, the posterior distribution $P(\theta \mid o)$, is computed as the product of the prior $P(\theta)$ and the likelihood $P(o \mid \theta)$. The denominator normalizes the posterior in order to sum to one.

In the TED task there are three types of observations: First, if an event starts at a region while the respective region is fixated and is perceived for the whole event duration the observation is defined as *fully observed*. Second, if the region is not fixated at all while the event occurs the observation is defined as *missed*. Third, if only a fraction of the event is observed (event is already occurring when the region is fixated) the observation is defined as partially observed. Hereby, we assume that if a region is fixated while an event occurs, participants continue fixating the respective region until the event has finished. Our data show that this is indeed the case. The different types of observations lead to different types of censoring. The resulting likelihood can be computed by

$$p(\mathbf{o} \mid \theta) \;\; = \;\; \prod_{\text{fully observed}} p(o \mid \theta) \prod_{\text{missed}} F(t_E - t_L) \tag{6.18}$$

$$\cdot \;\; \prod_{\text{partially observed}} F(t_E - t_L) - F(t_E - t_S) \,, \tag{6.19}$$

where $P(o \mid \theta)$ is Gaussian and F is the cumulative distribution function of θ. The starting and ending time for the event are denoted as t_S and t_E, respectively. The time when the region of interest is fixated is denoted by t_L.

In order to extend our models from the steady state case to all trials, the progression of uncertainty over the course of the blocks was simulated using the proposed Bayesian Learner. The mean estimate and the uncertainty of the event duration were computed for each trial in a block as follows: We repeatedly simulated events and determined the type of the resulting observation using the mean switching pattern for each stage of a block. We then used Markov Chain Monte Carlo methods to draw samples from the posterior belief about the true event durations for each trial. For each trial, the prior distribution for the event duration is equivalent to the posterior distribution in the trial before. For the first trial, we used a Gaussian distribution fitted to the means of the three event durations. To model the initial belief prior to the first observation, we used a Gaussian distribution fitted to the true event durations for the three conditions. We controlled the subjects prior belief by having them perform a few training trials during which they became familiar with the overall range of event durations.

6.3.3 Parameter estimation

Parameters were estimated by least squares using the aggregated data from trials 10 to 20 because switching behavior was stable across subjects during this period. AIC values for the models were computed according to Burnham and Anderson (2004) as

$$AIC = n \log(\frac{SSE}{n}) + 2(p+1) + C \tag{6.20}$$
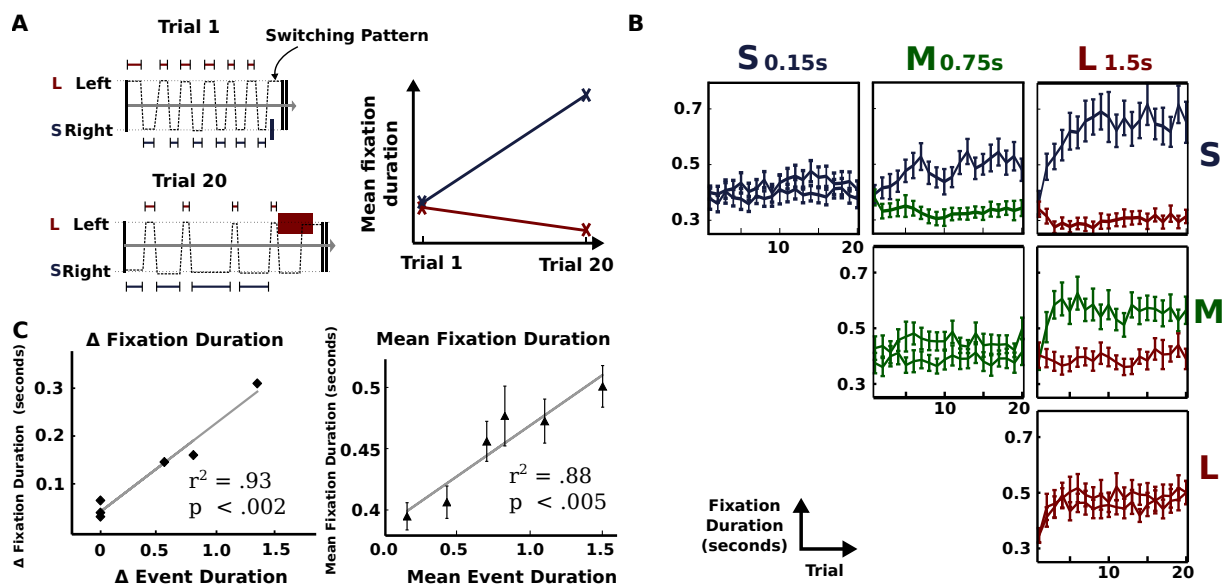
Figure 6.4: Results for fixation durations. (A) Sample switching pattern for participant MK in a single block (left) and mean fixation duration for each region (right). Switching patterns are shown for two trials, i.e., before learning the event durations (Trial 1) and after learning the event duration (Trial 20). (B) Mean fixation duration aggregated for all participants and all blocks grouped by condition (error bars correspond to standard error of the mean). Mean fixation duration is shown in seconds (y-axis) and is the same scale for all conditions. Trial number within the block is shown on the x-axis. Colors of the lines refer to the event duration used in that condition. (C) Mean fixation duration difference as a function of event duration difference (left). Each data point corresponds to a specific condition. Mean fixation duration as a function of mean event duration (right). Again, each data point is associated with one of the conditions.

where n is the number of data points, p the number of free parameters, and SSE the sum of squared errors. The constant C contains all additive components that are shared between the models. Since we use the difference of AIC values between the models for model comparison, these constant values can be ignored, as they cancel out. Akaike weights were computed according to Wagenmakers and Farrell (2004).

$$w_i(\text{AIC}) = \frac{\exp\left\{-\frac{1}{2}\Delta_i(\text{AIC})\right\}}{\sum\limits_{k=1}^{K} \exp\left\{-\frac{1}{2}\Delta_k(\text{AIC})\right\}} \tag{6.21}$$

where i refers to the model the weight is computed for, $\Delta_i(\text{AIC})$ is the difference in AIC between the ith model and the model with the minimum AIC. The Akaike weight of a model is a measure for how likely it is that the respective model is the best among the set of models. The probability ratio between two competing models i and j was computed as the quotient of their respective weights w_i/w_j.
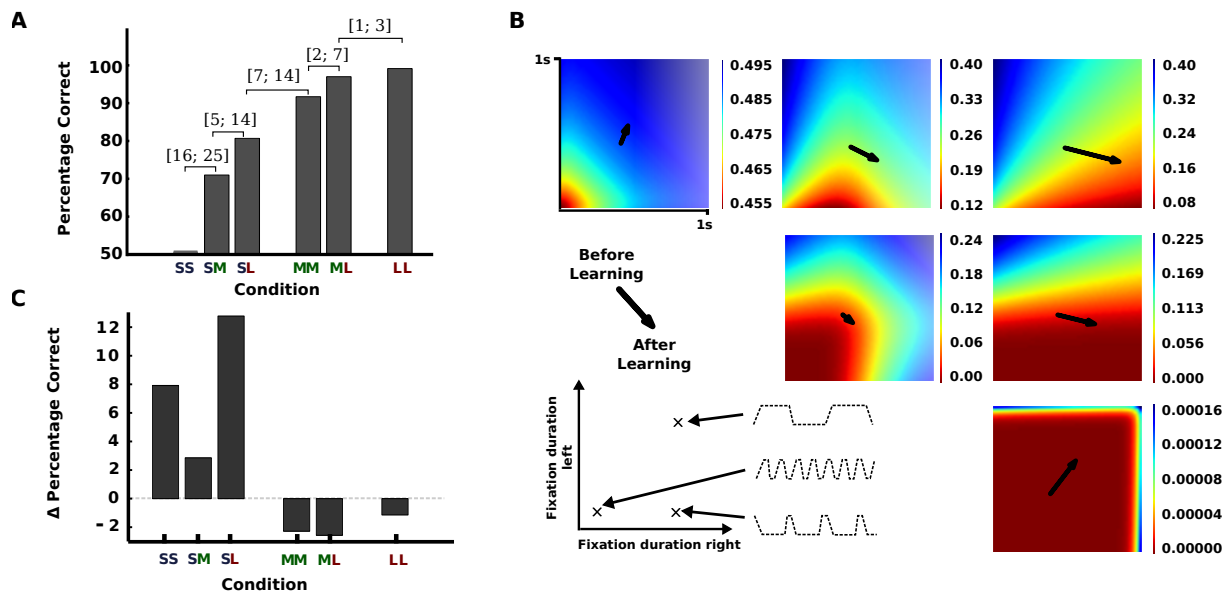
Figure 6.5: Performance results. (A) Percentage of detected events for all conditions. Values in square brackets correspond to the 95% credibility interval for the differences in proportion. Different conditions are shown on the x-axis. (B) Behavioral changes from Trial 1 to Trial 20 (arrows) for each condition. Color density depicts the probability of missing an event for every possible switching pattern . Each plot corresponds to a single condition. Conditions are arranged as in Figure 6.4 B. (C) Difference in detection performance between the first and last trial in each block. Different conditions are shown on the x-axis.

6.4 Results

6.4.1 Subjects

Ten undergraduate subjects (4 females) took part in the experiment in exchange for course credit. The participants' age ranged from 18 to 30 years. All subjects had normal or corrected to normal vision. Data from two subjects were not included in the analysis due to insufficient quality of the eye movement data. Participants completed two to four shorter training blocks (ten trials per block) prior to the experiment to get used to the setting and the task (5 minutes on average). The number of training blocks, however, did not influence task performance. All experimental procedures were carried out in accordance with the guidelines of the German Psychological Society and the university's ethics committee.

6.4.2 Behavioral results

We investigated how participants changed their switching behavior over trials within individual blocks as a consequence of learning the constant event durations within a single block. Figure 6.4 shows two example trials for participant MK, both from the same block (Block LS – long event duration at the left region, short event duration at the right region). In the first trial when event durations were unknown (upper panel), MK switched evenly between the regions in order to detect the event ($t_7 = 0.48; p = 0.64$). However, at the end of this block in Trial 20, we observed a clear change in the switching pattern (lower panel). There was a shift in the mean fixation duration over the course of the block (right panel). The region associated with the short event durations was fixated longer in the last trial of the block ($t_4 = 2.75; p = 0.04$). Hence, the
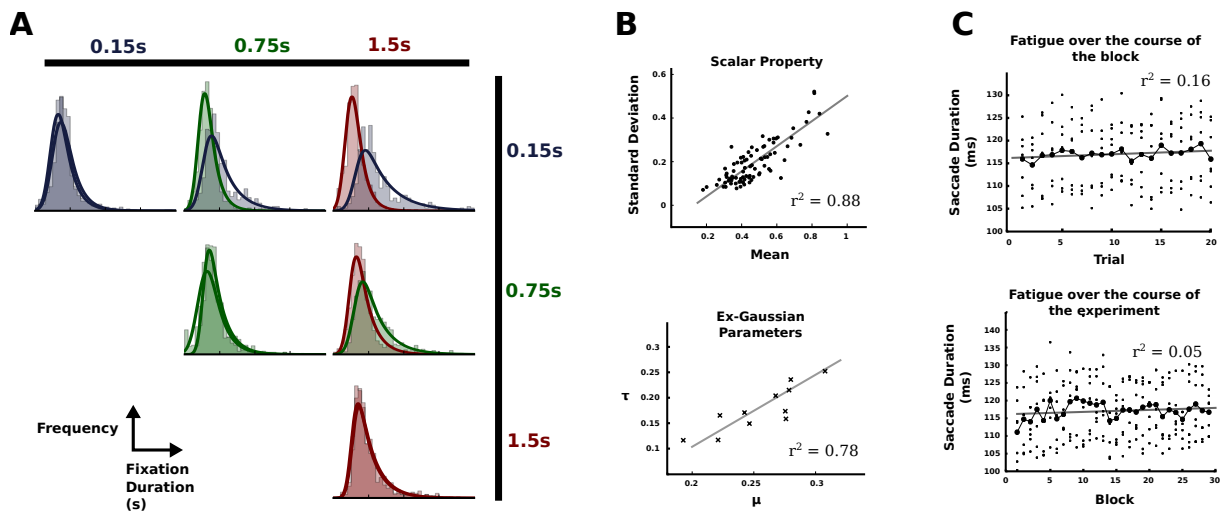
Figure 6.6: Properties of eye movement sequences. (A) Distributions of fixation durations across conditions. For each condition histograms and fitted exponential Gaussian distributions are shown for both regions. (B) Scalar law of biological timing. The upper panel depicts the relationship between the mean fixation durations in a trial and the respective standard deviation observed in our experiment. The solid line refers to the result of a fitted linear regression. The lower panel depicts the linear relationship of the parameters of the exponential Gaussian distribution, where μ denotes the mean of the Gaussian and τ denotes the rate of the exponential distribution. Each point represents the best parameter estimates for a single distribution shown in A. (C) Quantifying oculomotor fatigue. Saccade durations over the course of trials aggregate over all blocks (top). Saccade durations over the course of blocks aggregated over all trials within the respective blocks (bottom).

subject changed the gaze switching pattern over trials in this particular block in response to the event durations at the two locations.

Mean fixation durations aggregated over all participants are shown in Figure 6.4B. Variability and temporal properties of fixation durations for the different conditions are shown in Figure 6.6. Steady-state behavior across subjects was reached within the first five to ten trials within a block suggesting that participants were able to quickly learn the event durations. Behavioral changes were consistent across conditions. In conditions with different event durations (SM, SL, ML), regions with shorter event durations were fixated longer in later trials of the block (all differences were highly significant with $p < 10^{-5}$). Participants' gaze behavior was not only influenced by whether the event durations were different, but was also guided by the size of the difference (Figure 6.4C, left panel). Greater differences in event durations led to greater differences in the fixation durations (linear regression, $r^2 = 0.93, p < 0.002$). Moreover, the overall mean fixation duration aggregated over both regions was affected by the overall mean event durations (Figure 6.4C, right panel). Participants, in general, switched faster between the regions if the mean event duration was small (linear regression, $r^2 = 0.88, p < 0.005$). Blocks with long overall event durations (e.g. LL) yielded longer mean fixation durations compared to blocks with short event durations (e.g., SS). Taken together, these results establish that participants quickly adapted their eye movement strategies to the task requirements given by the timing of events.

However, further inspection of the behavioral data suggests non-trivial aspects of the gaze switching strategies. The conditions in our TED task differed in difficulty depending on the event durations (Figure 6.5A). The shorter the event durations in a specific condition, the more difficult was the task. The improvement over the course of a block differed between conditions

Table 6.1: Model comparison for all fitted models. Differences in AIC (ΔAIC), Akaike weights ($w(AIC)$) and the probability ratios (w_min/w) were computed with respect to the Full Model (FM).

Model	n	p	SSE	AIC	ΔAIC	w(AIC)	w_{min}/w
Full Model (FM)	90	6	0.15	-561.56	0	0.99	1.0
FM without Perc. Uncertainty	90	6	0.18	-543.11	18.44	9.9×10^{-5}	1.0×10^4
FM without Cost	90	4	0.65	-434.06	127.50	2.1×10^{-28}	4.8×10^{27}
FM without Processing Time	90	5	1.04	-389.14	172.42	3.6×10^{-38}	2.8×10^{37}
FM without Acting Uncertainty	90	4	1.21	-377.95	183.61	1.4×10^{-40}	7.4×10^{39}

Note: n = number of data points, p = number of fitted parameters

(Figure 6.5B-C) and participants only improved in detection accuracy in conditions associated with high task difficulty. In contrast, participants showed a slight decrease in performance in the other conditions.

This can be represented as arrows pointing from the observed switching behavior in the first trial to the behavior in the last trial within blocks in the action space (Figure 6.5B). Overall, several aspects remain unclear, i.e., the different rates of change in switching frequencies across trials in a block as well as the different directions and end points of behavioral changes observed in the performance space.

6.4.3 Bounded actor model

We hypothesized that multiple processes are involved in generating the observed behavior, based on the behavioral results in the TED task and taking into account known facts about human time constraints in visual processing, time perception, eye movement variability, and intrinsic costs. Starting from an ideal observer (Geisler, 1989) we developed a bounded actor model by augmenting it with computational components modeling these processes (Figure 6.7A).

The ideal observer chooses switching patterns that maximize performance in the task. The gaze behavior suggested by the ideal observer corresponds to the switching pattern with the lowest probability of missing the event (Figure 6.5B). We extend this model by including perceptual uncertainty, which limits the accuracy of estimating event durations from observations. An extensive literature has reported how time intervals are perceived and reproduced (Merchant, Harrington, & Meck, 2013; Ivry & Schlerf, 2008; Wittmann, 2013). The central finding is the so-called scalar law of biological timing (Gibbon, 1977; Merchant et al., 2013), a linear relationship between the mean and the standard deviation of estimates in the interval timing task (which we also found in our data, see Figure 6.6). We included perceptual uncertainty in the ideal observer model based on the values found in the literature.

The second extension of the model accounts for eye movement variability, which limits the ability of humans to accurately execute planned eye movements and has frequently been reported
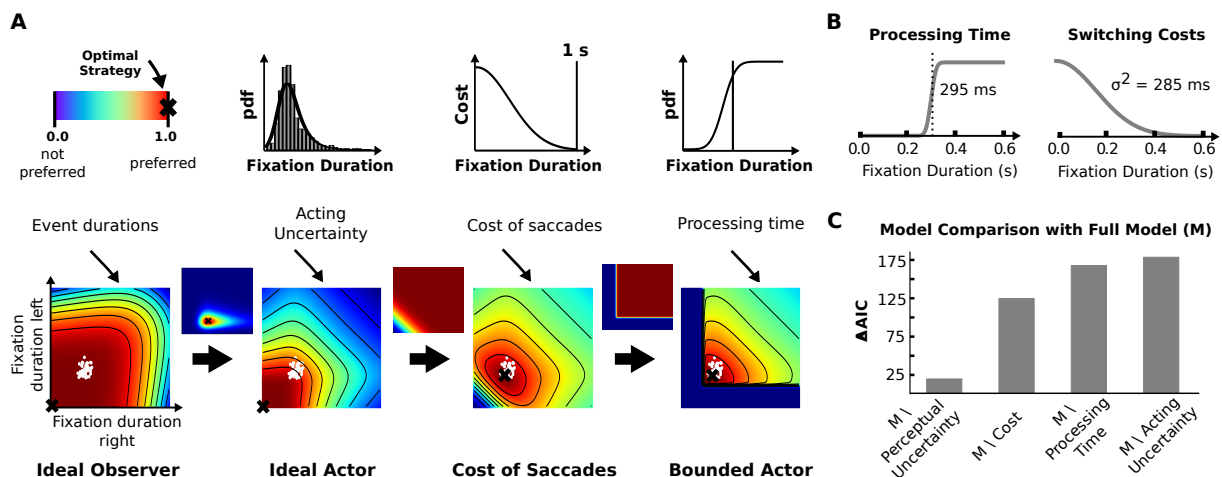
Figure 6.7: Computational model. (A) Schematic visualization describing how the different model components influence the preferred switching patterns for a single condition (Medium - Medium). Images show which regions in the action space are preferred by the respective models. Each location in the two dimensional space corresponds to a single switching pattern (see also Figure 6.5B). White dots are the mean fixation durations of our subjects aggregated over individual blocks. The optimal strategy is marked with a black cross for each model. The ideal observer only takes the probability of missing an event into account. The ideal actor also accounts for variability in the execution of eye movement plans thus suppressing long fixations (as they increase the risk of missing the event). Therefore, strategies using shorter fixation durations are favored. Next, additive costs for faster switching yield a preference for longer fixation durations. Hence, less value is given to switching patterns near the origin. Finally, the Bounded Actor model excludes very short fixation durations, which are biologically impossible, as they do not allow sufficient time for processing visual information. (B) Model components for processing time and switching costs after parameter estimation. Parameters were estimated through least squares using the last ten trials of each block. (C) Model comparison. For the absence of each component the model was fitted by least squares. Differences in AIC with respect to the full model are shown.

in the literature. While many factors may contribute to this variability in action outcomes, humans have been shown to take this variability into account (Hudson et al., 2008). Further evidence comes from experiments that have shown, that in a reproduction task the impact of prior knowledge on behavior increased according to the uncertainty in time estimation (Jazayeri & Shadlen, 2010) and studies that extended an ideal observer to an ideal actor by including behavioral timing variability (Sims, Jacobs, & Knill, 2011). In order to describe the variability of fixation durations, we used an exponential Gaussian distribution (Ratcliff, 1979) as it is a common choice for describing fixation durations (Laubrock, Cajar, & Engbert, 2013; Staub, 2011; Luke, Nuthmann, & Henderson, 2013; Palmer, Horowitz, Torralba, & Wolfe, 2011). The parameters were estimated from our data using maximum likelihood.

While behavioral costs in biological systems are a fundamental fact, very little is known about such costs for eye movements, as research has so far focused on extrinsic costs (Navalpakkam et al., 2010; Schütz et al., 2012; Hayhoe & Ballard, 2014). We included costs into our model by hypothesizing that switching more frequently is associated with increased effort. In a similar task that did not reward faster switching participants showed fixation durations of about one second (Schütz, Lossin, & Kerzel, 2013). To our knowledge, a concrete shape of these costs has not been investigated yet and in the present study, we used a sigmoid function (Figure 6.7B, right panel) for which the parameters were estimated using least squares from the measured data.
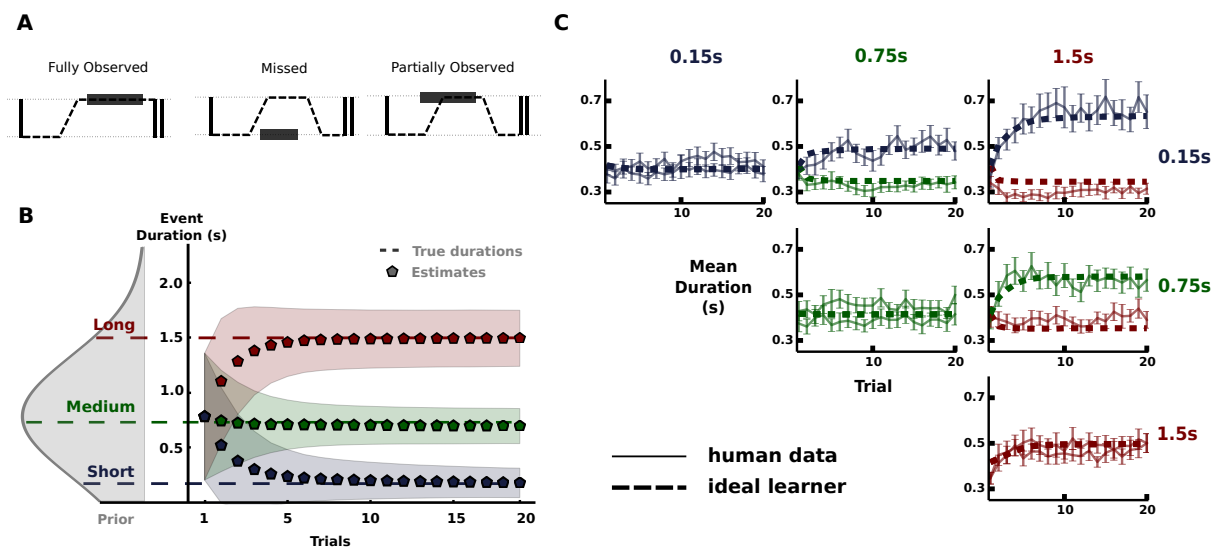
Figure 6.8: Bayesian Learner Model. (A) Different types of observations in the TED task. Dashed lines depict switching patterns, rectangles denote events. (B) Mean estimates and uncertainty (standard deviation) for the event durations over the course of a block. Prior belief over event durations before the first observation is shown on the left side. (C) Human data and model predictions for the Bayesian learner model. Fitted parameters from the full model were used. Dashed lines show the mean fixation duration of the bounded actor model over the course of the blocks. For each trial the mean estimate suggested by the ideal learner was used.

Finally, the model has so far neglected that all neural and motor processes involved take time. The time required for visual processing is highly dependent on the task at hand. It can be very fast (Stanford et al., 2010), however, in the context of an alternative choice task processing is done within the first 150 ms (Thorpe, Fize, Marlot, et al., 1996) and does not improve with familiarity of the stimulus (Fabre-Thorpe, Delorme, Marlot, & Thorpe, 2001). In our study, the demands for the decision were much lower since discriminating between the two states of a region (ongoing event, no event) is rather simple. Still, some additional constraints limit the velocity of switching, i.e., eye-brain lag, time of planning and initiating a saccade (Trukenbrod & Engbert, 2014). In addition, prior research suggests that in general saccades are triggered after processing the current fixation has reached some degree of completeness (Remington W., Wu, & Pashler, 2011). We estimated the mean processing time (Figure 6.7B, left panel) from our behavioral data using least squares.

To investigate which model should be favored in explaining the observed gaze switching behavior we fitted the model leaving out one component at a time and computed the Akaike information criterion (AIC) for each of these models. The results are shown in Figure 6.7C. The conditional probability for each model was calculated by computing the Akaike weights from the AIC values as suggested in Wagenmakers and Farrell (2004). We found that gaze switches were more than 10000 times more probable under the full model compared to the second best model (see Table 6.1). In addition, models lacking a single component deviate severely from the observed behavior. This strongly suggests that the temporal course of eye movements cannot be explained solely on the basis of an ideal observer. Instead, multiple characteristics of perceptual uncertainty, behavioral variability, intrinsic costs, and biological constraints have to be included.
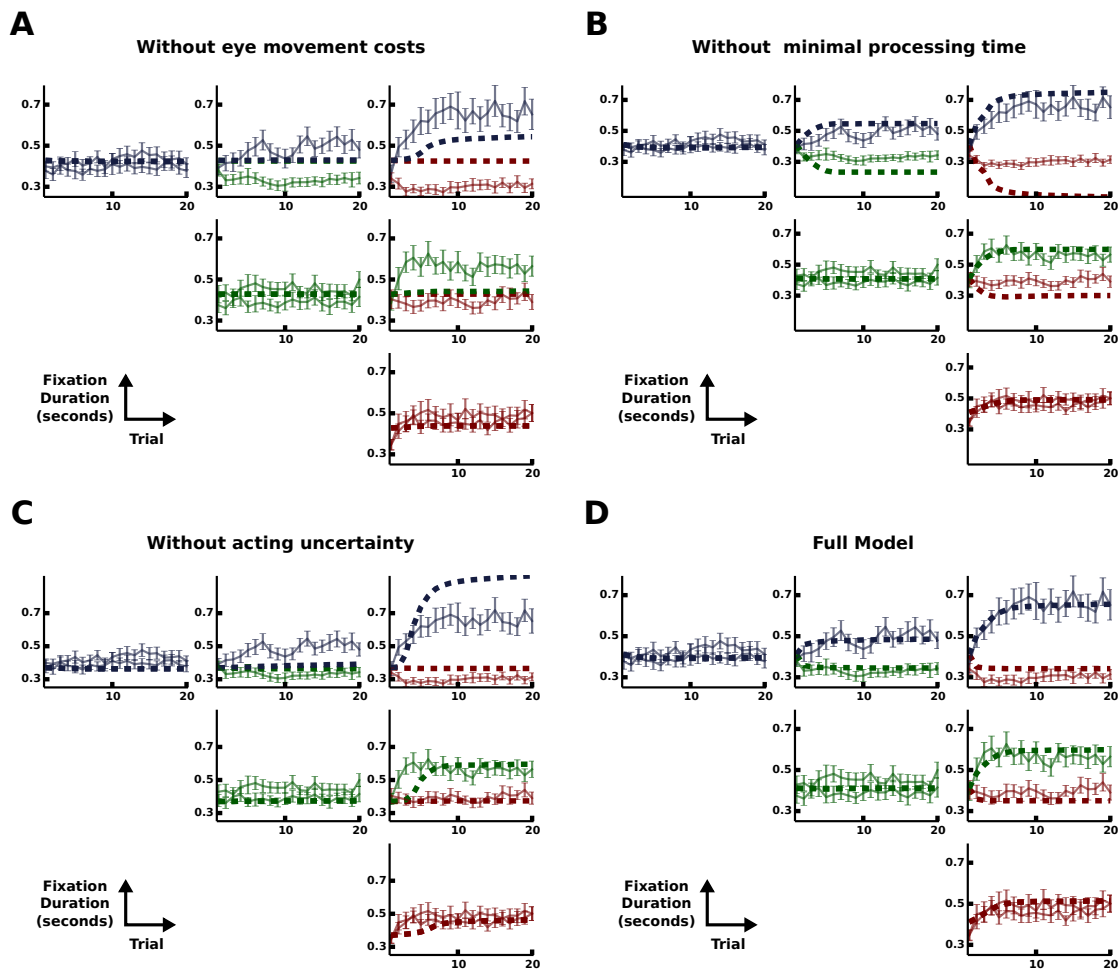
Figure 6.9: Learning of gaze strategies over blocks for all conditions according to different models. For each model (except D), a single component was removed from the full model. The remaining parameters were estimated from the data using only the steady state behavior (Trial 10-20). (A) Eye movement costs. (B) Processing time. (C) Acting uncertainty. (D) Full model with all parameters fitted.

6.4.4 Bayesian learner model

All models so far have assumed full knowledge of the event durations, as is common in ideal observer models. This neglects the fact, that event durations were a priori unknown to the participants at the beginning of each block of trials and needed to be inferred from observations. In our TED task information about the event durations cannot be accessed directly but must be inferred from different types of censored and uncensored observations (Figure 6.8A). Using a Bayesian learner, we simulated the subjects' mean estimate and uncertainty of the event duration for each trial in a block using Markov Chain Monte Carlo techniques (Figure 6.8B). We used a single Gaussian distribution fitted to the true event durations for the three conditions to describe subjects' initial belief prior to the first observation within a block. Our Bayesian learner enables us to simulate the bounded actor model for every trial of a block using the mean estimates (Figure 6.8C). The results show that our bounded actor together with the ideal Bayesian learner is sufficient to recover the main characteristics of the behavioral data. Crucially, the proposed learner does not introduce any further assumptions or parameters. This means
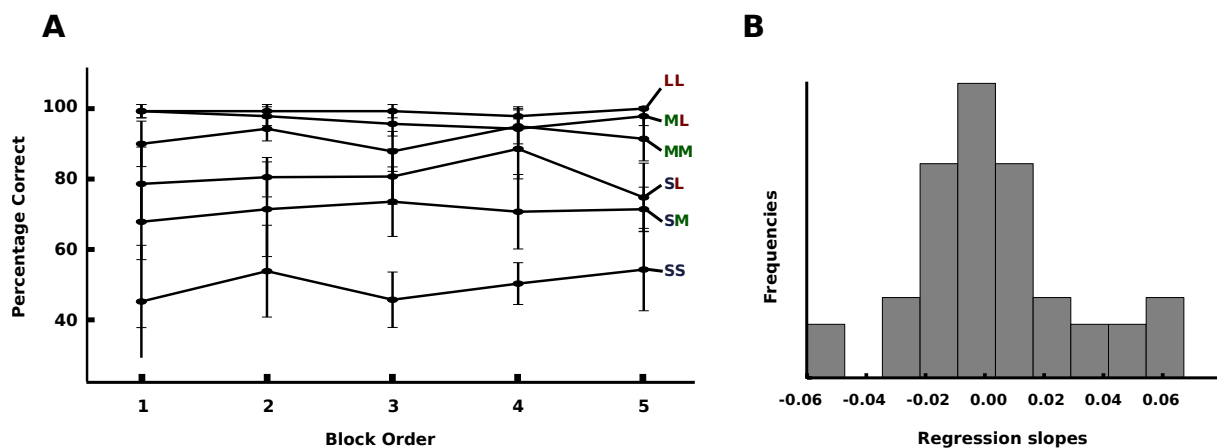
Figure 6.10: Learning effects over the course of the experiment. (A) Mean performance aggregated over all participants grouped by the position of the block within the experiment. Error bars correspond to the sample standard deviation. Different conditions are shown as separate lines. (B) Histogram of regression slopes. For each participant and each condition linear regression between the mean performance in the blocks and the position of the respected blocks in the experiment was computed. The distribution of regression slopes indicates no improvement in the task over the course of the experiment.

that the behavioral changes we observed can be explained solely by changes in the estimates of the event durations. As with the steady-state case, all included factors contributed uniquely to the model. Omitting a single component lead to severe deviation from the human data (see also Figure 6.9).

6.4.5 Lower bound model fit

Our model predicted a single optimal fixation duration for each aggregated line in Figure 6.9 (in the range of Trial 10 to 20). The sum of squared errors was computed as the variation around this predicted value and then used to obtain the AIC. Differences in AIC were used to compare different models. However, we used the AIC solely for the purpose of model comparison, not as an absolute measure of model fit, as AIC is not well suited for evaluating the absolute model fit due to the missing of a fixed lower bound. If a model were so good that it perfectly predicted the dependent variable, the error-sum-of squares would approach zero, and the natural logarithm of this value and thereby the AIC would approach negative infinity. Further, the AIC by itself is highly dependent on the units of measurement, which influence the error-sum-of squares.

Here we present an absolute lower bound on AIC for the data presented in our experiment. We computed the sum of squared errors for a model that estimates the best fitting parameter for each aggregated line in Figure 6.9. For each line, the parameter that minimizes the SSE is the mean. The lower bound for the data we collected is SSE = 0.09 and AIC = -603.15, respectively. We used n = 90 and p = 9 (one for each line separately) parameters.

Note that this is an overly optimistic bound for two reasons: First, the estimates for fixation durations underlying the lower bound are independently chosen to best represent each individual optimal strategy. For the full model, however, these predictions are not independent since each estimated parameter simultaneously affects the optimal fixation durations for all conditions. This dependency puts a constraint on the full model that is not considered by this lower bound. Second, the lower bound assumes the sample mean of the data to be equal to the optimal fixation duration and only quantifies the variation around that mean. This neglects the variability across

blocks and participants. In contrast, the full model uses all available data to estimate biologically plausible parameters. Using these parameters the optimal fixation duration is predicted by the dynamics of the full model rather than directly estimated from the data. While this is necessary to provide generalizability and indispensable to prevent overfitting, it naturally introduces a further source of variability. This is reflected in a higher sum of squared errors of the full model compared to the lower bound.

6.5   Discussion

The study investigated how humans learn to adjust their eye movement strategies to detect relevant events in their environment. The minimal computational model matching key properties of observed gaze behavior includes perceptual uncertainty about duration, action variability in generated gaze switches, a minimal visual processing time, and finally intrinsic costs for switching gaze. The bounded actor model explains seemingly suboptimal behavior such as the decrease in performance in easier task conditions as a tradeoff between detection probability and faster but more costly eye movement switching. With respect to this bounded actor model, the experimentally measured switching times at the end of the blocks were close to optimal. While ideal observer models (Najemnik & Geisler, 2005; Peterson & Eckstein, 2012; Chukoskie et al., 2013) assume, that the behaviorally relevant uncertainties are all known to the observer, here we also modeled the learning progress. The speed at which learning proceeded was remarkable, as subjects were close to their final gaze switching frequencies after only five to ten trials on average, suggesting that the observed gaze behavior was not idiosyncratic to the TED task. Based on the bounded actor model we proposed a Bayesian optimal learner, which only incorporates the inherent perceptual uncertainty about time passage on the basis of the scalar law of biological timing (Gibbon, 1977; Merchant et al., 2013; Wittmann, 2013). Remarkably, this source of uncertainty was sufficient to fully account for the learning progress. This suggests, that the human visual system is highly efficient in learning temporal regularities of events in the environment and that it can use these to direct gaze to locations in which behaviorally relevant events will occur.

Taken together, this study provides further evidence for the importance of the behavioral goal in perceptual strategies (Yarbus, 1967; Rothkopf et al., 2007; Hayhoe & Ballard, 2014), as low-level visual feature-saliency cannot account for the gaze switches (Itti & Koch, 2001). Instead, cognitive inferences on the basis of successive observations and detections of visual events lead to changes in gaze strategies. At the beginning of a block, participants spent equal observation times between targets to learn about the two event durations and progressively looked longer at the location with the shorter event duration. This implies that internal representations of event durations were likely used to adjust the perceptual strategy of looking for future events.

A further important feature of the TED task and the presented model is that, contrary to previous studies (Najemnik & Geisler, 2005; Peterson & Eckstein, 2012; Chukoskie et al., 2013), the informativeness of spatial locations changes moment by moment and not fixation by fixation, which corresponds more to naturalistic settings. Under such circumstances, optimality of behavior is not exclusively evaluated with respect to a single, next gaze shift, but instead driven by a complex sequence of behavioral decisions based on uncertain and noisy sensory measurements. In this respect, our TED task is much closer to natural vision than experiments with static displays. On the other hand, the spatial locations in our task were fixed and the visual stimuli were simple geometric shapes. Naturalistic vision involves stimuli of rich spatiotemporal statistics (Simoncelli & Olshausen, 2001) and semantic content across the entire visual field.

How the results observed in this study transfer to subjects learning to attend in more naturalistic scenarios will require further research. The challenge under such circumstances is to quantify relative uncertainties as well as costs and benefits of all contributing factors. There is nothing that in principle precludes applying the framework developed here to such situations because of its generality. Clearly, several factors indicate that participants did not depart dramatically from eye movement strategies in more normal viewing environments and thus that natural behavior transferred well to the TED task. First, subjects did not receive extensive training. Secondly, they showed only minimal changes in learning behavior across blocks over the duration of the entire experiment (see Figure 6.10), and thirdly intersaccadic interval statistics were very close to the ones observed in naturalistic settings according to previous literature (Laubrock et al., 2013; Staub, 2011; Luke et al., 2013; Palmer et al., 2011). Overall, the found concepts contribute uniquely to our understanding of human visual behavior and should inform our ideas about neuronal mechanisms underlying attentional shifts (Gottlieb, 2012), conceptualized as sequential decisions. The results of this study may similarly be relevant for the understanding of learning optimal sensory strategies in other modalities and the methodologies can fruitfully be applied there, too.

CHAPTER 7

Variability of eye movement sequences

7.1 Introduction

Variability is inherent to all actions. Even extensive training, for example in professional athletes, does not lead to deterministic action execution. Various sources have been found responsible for behavioral variability. These include fluctuations in the internal states of the subject, past experiences, the current level of attention as well as microscopic processes in the brain, among others (see Renart & Machens, 2014; Beck, Ma, Pitkow, Latham, & Pouget, 2012; Nienborg & Cumming, 2009; Faisal, Selen, & Wolpert, 2008).

However, little is known about variability in temporal human visual behavior. While spatial variability, for example, endpoints of saccades (van Beers, 2007), is straightforward to measure, investigating the connection of variability and environmental dynamics is more complicated because one has to deal with action sequences instead of single actions. Even for tasks in non-changing stationary environments such as looking at pictures quantifying the similarity of eye movement sequences is hard (see Hacisalihzade, Stark, & Allen, 1992, for an approach based on Markov matrices). Previous research indicates that the external visual input influences variability. When watching movies eye movements are similar between subjects (Goldstein, Woods, & Peli, 2007) and across species (humans and maquaces, see Shepherd, Steckenfinger, Hasson, & Ghazanfar, 2010). Also, eye movements have shown to be more consistent in Hollywood trailers than in natural movies (Dorr, Martinetz, Gegenfurtner, & Barth, 2010) and consistency of eye movements across observers increases with the observers' age (Franchak, Heeger, Hasson, & Adolph, 2016). One explanation is that the short clips are specifically designed to lead attention (Dorr, Vig, & Barth, 2012). Also, the frequency of cuts could play a role since variability between observers increases with viewing time when viewing static images (Tatler, Baddeley, & Gilchrist, 2005). Surprisingly, Wang, Freeman, Merriam, Hasson, and Heeger (2012) found eye movements to be less consistent across observers of movie junks presented in a random order the smaller the duration of the junks.

How interindividual variability relates to characteristics of a visual task is still unanswered. There is a connection between the variability of motor actions and the reward structure of the task (Pekny, Izawa, & Shadmehr, 2015). If the probability of reward was low, variability in the movements increased. Past studies concentrated on natural stimuli like movies where many aspects of the task are unknown and the statistics of the stimuli are far from understood. Also, repeatedly probing a subject with the same stimulus is impossible due to learning. For movies, it was shown that variability increased with repeated watching (Dorr et al., 2010). Further, variability serves an essential purpose during the development of skills (see Dinstein, Heeger, & Behrmann, 2015, and the references therein) and studies have shown that variability itself can be altered through reinforcement (Paeye & Madelain, 2014). During learning, variability can
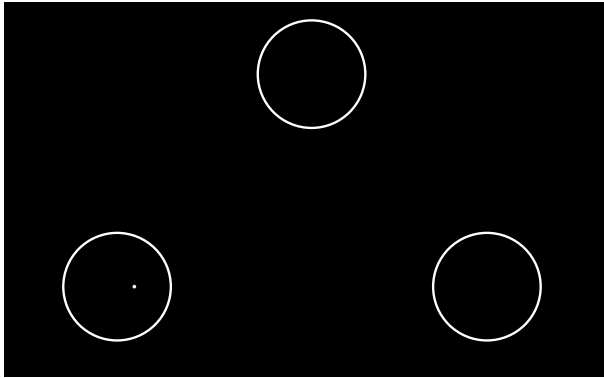
Figure 7.1: Stimulus for the temporal monitoring task. Only a single dot (at the fixated circle) is visible at any time. In this case, gaze is targeted at the lower left circle. For the other two locations, only the circular bound is shown.

serve the purpose of exploration and facilitate finding multiple equivalent solutions to a problem (see Sternad, 2018, for a review).

In this study, we fill a gap in the eye movement literature by investigating goal-directed eye movements in a changing environment with controlled dynamics. This includes simultaneous demands on spatial and temporal control of gaze. In particular, we study the influence of task characteristics on interindividual behavioral variability. Therefore, we developed a monitoring task where three spatial locations with independent dynamics are monitored over time in order to detect an event. Crucially, we designed the stimulus dynamics in a way to make repeated presentation of the same stimulus dynamics possible while preventing learning effects. All participants were shown the same movements as the primary goal of the study was to investigate variability between participants' eye movements. Using this experimental paradigm we investigated the following questions: 1) How do humans decide when to make a saccade to a new target in a dynamic environment? 2) How do humans determine the target of the next saccade? 3) How does the reward structure of the task influence variability?

7.2 Methods and materials

7.2.1 Stimulus and task design

In our visual monitoring task, three circles (diameter 8.72 deg, 300 px) were shown to the participant at different spatial locations arranged as a triangle (Figure 7.1). Within each of the circles, a small dot (diameter 0.29 deg, 10 px) moved randomly according to a random walk (see Figure 7.2). At the beginning of each trial, all dots were located at a distance of 1.74 deg (60 px) away from the center of their respective circles at a random angle. The dots moved at a rate of three times per second (step duration 333 ms). The dot locations were updated simultaneously at all three circles. The end of the trial was reached if one of the dots moved outside its circle. This was the event participants were instructed to detect by switching their gaze between the circles. A trial was successful if participants fixated the correct circle at the moment of the dot crossing the circular bound. Eye movement data were collected using the Tobii EyeX eye tracker (60 Hz).

7.2.2 Stimulus generation

The goal of our stimulus design was to investigate inter- and intra-individual variability in eye movement sequences. In our event detection task, the movement of the dots was generated using polar coordinates with different generative processes for radial and angular movement.
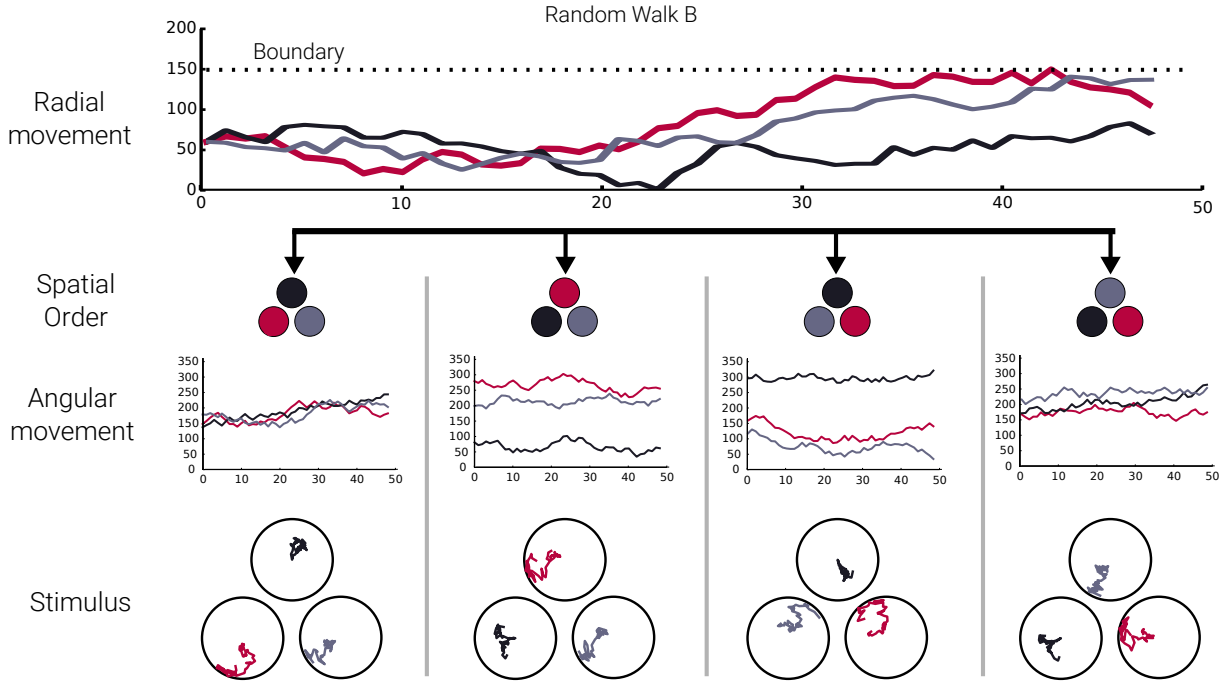
Figure 7.2: Experimental design and stimulus generation. We used the same radial stimulus movement for multiple trials (first row). To disguise this from the participants we randomly drew a spatial ordering for each trial (second row). Also, the angular movement was sampled from a uniform distribution (third row). As a result, the same radial stimulus movement lead to different presentations (fourth row).

The steps for the radial movement were drawn from a Gaussian distribution, the steps for the angular movement were drawn from a uniform distribution

$$r_{t+1} = r_t + \epsilon_r, \quad \epsilon_r \sim \mathcal{N}(0, \sigma^2) \tag{7.1}$$

$$\theta_{t+1} = \theta_t + \epsilon_\theta, \quad \epsilon_\theta \sim \mathcal{U}(-15, 15), \tag{7.2}$$

where σ is the standard deviation of the Gaussian stepsize (set to 0.29 deg / 10 px throughout this experiment), r_t is the current distance of the dot to the center and θ_t is the current angle. The increment for the angle was chosen uniformly between -15 and 15 degrees. The dot's location at time $t + 1$ in Cartesian coordinates can then be computed as

$$(x_{t+1}, y_{t+1}) = \left(r_{t+1} \cos(\theta_{t+1}), r_{t+1} \sin(\theta_{t+1})\right) \tag{7.3}$$

where x and y refer to the coordinates in the two-dimensional space.

In order to generate differently looking stimuli with the same underlying probabilistic structure, we created eight different random walks (Figure 7.3) ranging between 34 (11.3 s) and 49 (16.3 s). We prevented the subjects from noticing the repeated presentation by generating each trial as follows (Figure 7.2): First, one of the eight radial movements, each comprising of three different Gaussian random walks, was selected. Second, we randomly assigned each of the three trajectories to one of the three locations (top, left, right). Third, we sampled the angular movement uniformly for each trial. As a consequence, we were able to generate stimuli with different appearance while keeping the task-relevant dimension (radial movement) the same. Overall, each of the eight different random walks was presented 15 times, hence each participant finished a total of 120 trials.
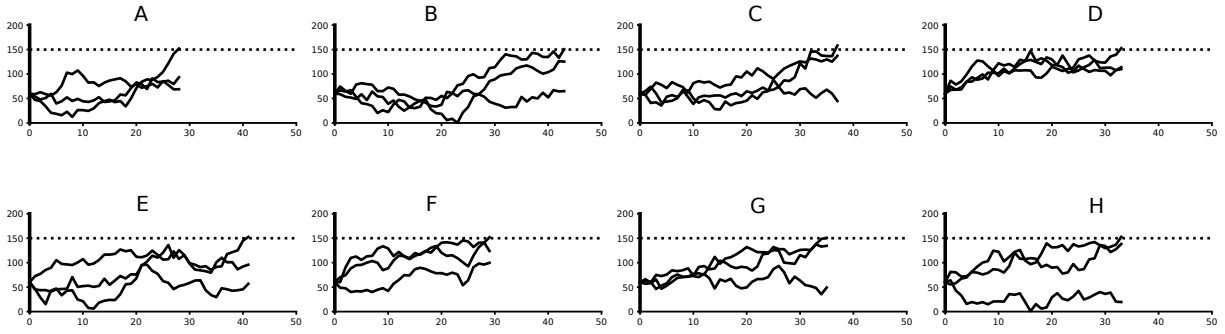
Figure 7.3: Different radial movements used in the experiment. In the beginning, all dots were located 1.74 deg (60 px) away from the center. The stimulus movement continued for 2 s (6 steps) after the circular bound was reached.

7.2.3 Computational model

In contrast to real-world situations, where we do not have access to the environmental statistics, in the presented temporal monitoring task, we have full knowledge about the process that generates the movement of the dots. Therefore, we can build a computational model based on these statistics. Also, we have access to the reward structure of the task the subject is solving, where reward is proportional to the probability of missing an event, i.e., one of the targets moving beyond the boundary. The critical event only depends on the radial position, hence we can disregard the angular movement leading to a one-dimensional problem. Using the task statistics, we can derive the optimal eye movement strategy. In each time step, the radial position of the dot is altered by additive Gaussian perturbations. For the region currently fixated, this change in location is observable up to perceptual uncertainty. For regions, which are not fixated, the change of the dot's location is not visible, hence the uncertainty increases with every time step.

   We ensure that the location of only one of the dots is perceptually accessible at any given time while the remaining two dots are hidden using a gaze-contingent display. Having full knowledge about the stimulus movement, we can compute the course of the uncertainty and thus model the belief of the subject with respect to the dots' locations. The state space $\mathcal{S}$ has three components, i.e., the location of the three dots. However, only one of the dots is observable at any given time, hence the full state is only partially observable and a belief state is used instead. We can utilize the additive property of the Gaussian distribution to calculate the belief about the location of a dot at any time as

$$p(s_n|s_0, n) \sim \mathcal{N}(s_0, n\sigma^2)\,, \tag{7.4}$$

where s_0 is the location of the dot at the time of the last observation, n is the number of steps elapsed since then, and σ^2 is the variance of the Gaussian random walk (Figure 7.4). The belief state $b(s)$ comprise the three locations as well as the duration since the last fixation for each dot.

   In our task subjects repeatedly decided which circle to fixate. A trial is considered successful if the dot which first moves beyond the boundary is fixated at the time of the event. Similar to all models presented in this thesis we assume that one component of our subjects' reward function is to solve the task in the experiment. Using our probability distribution for the dot locations $p(s_n|s_0, n)$ we can derive the optimal decision, i.e., the decision that maximizes the reward by minimizing the chance of missing the event. This corresponds to switching gaze to
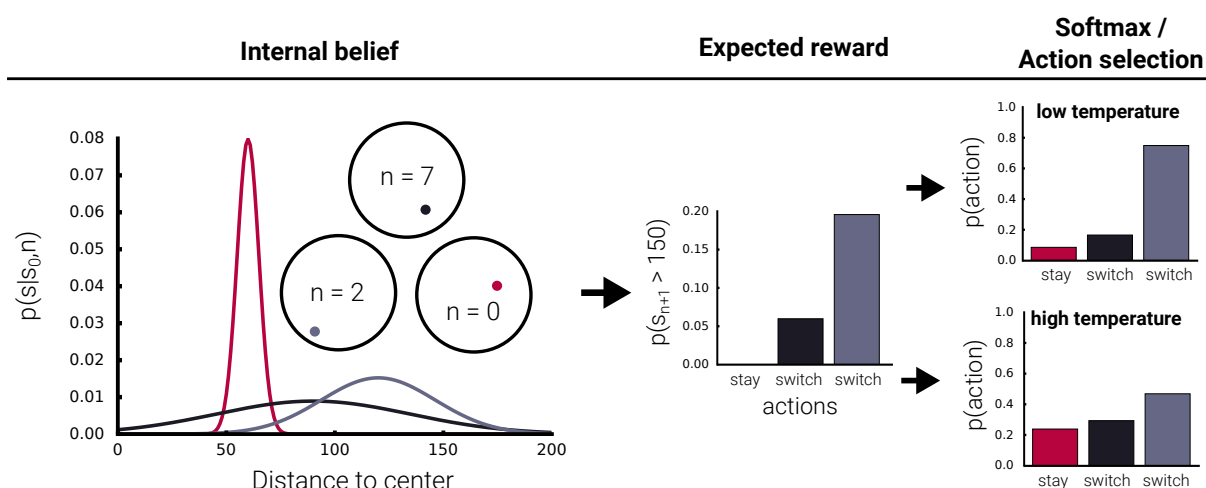
Figure 7.4: Schematic overview of the computational model. Using the observed quantites (s, n) for each location the internal belief is shown (left panel). Currently, the red location is fixated, hence the time since the last visit n is zero for that location. The belief is used to compute the expected reward for each of the three actions, which is proportional to the probability that the dot moves outside the circle (center panel). The expected reward can be transformed to a probability distribution over actions using a softmax transformation (right panel). The temperature coefficient regulates how differences in the expected reward lead to differences in action selection.

the dot that is most likely to exceed the boundary in the next time step at any given time. The reward function is given by

$$R(s,a) = \begin{cases} G & \text{if } s^a > \text{bound} \\ 0 & \text{else} \end{cases} \tag{7.5}$$

where s^a is the location of the dot at the circle targeted by action $a \in \{0, 1, 2\}$ and G is the immediate return a subject receives when successfully detecting the event. We can now compute the optimal policy

$$\pi = \arg \max_{a \in \{0,1,2\}} \mathbb{E}\big[R(s,a)\big] = \arg \max_{a \in \{0,1,2\}} \int_s p(s|o)R(s,a)ds \tag{7.6}$$

where $p(s|o)$ is the belief state summarizing the subjects' knowledge about the location of the dots and $o = (s_0, n)$ is the available information that can be observed.

Various factors are not accounted for in this model: uncertainty in the belief, inaccuracy in internal computations of the probability, among others. Hence, we used a softmax function to transform the expected reward to a probability distribution over actions.

$$p(a) = \frac{e^{Q(s,a)/\tau}}{\sum_{k=1}^{K} e^{Q(s,a_k)/\tau}} \tag{7.7}$$

where $Q(s,a) = \mathbb{E}\big[R(s,a)\big]$ and τ is the temperature coefficient. The temperature coefficient scales the differences in value and is used to relate internal rewards to actions probabilities (e.g., Neiman & Loewenstein, 2011).
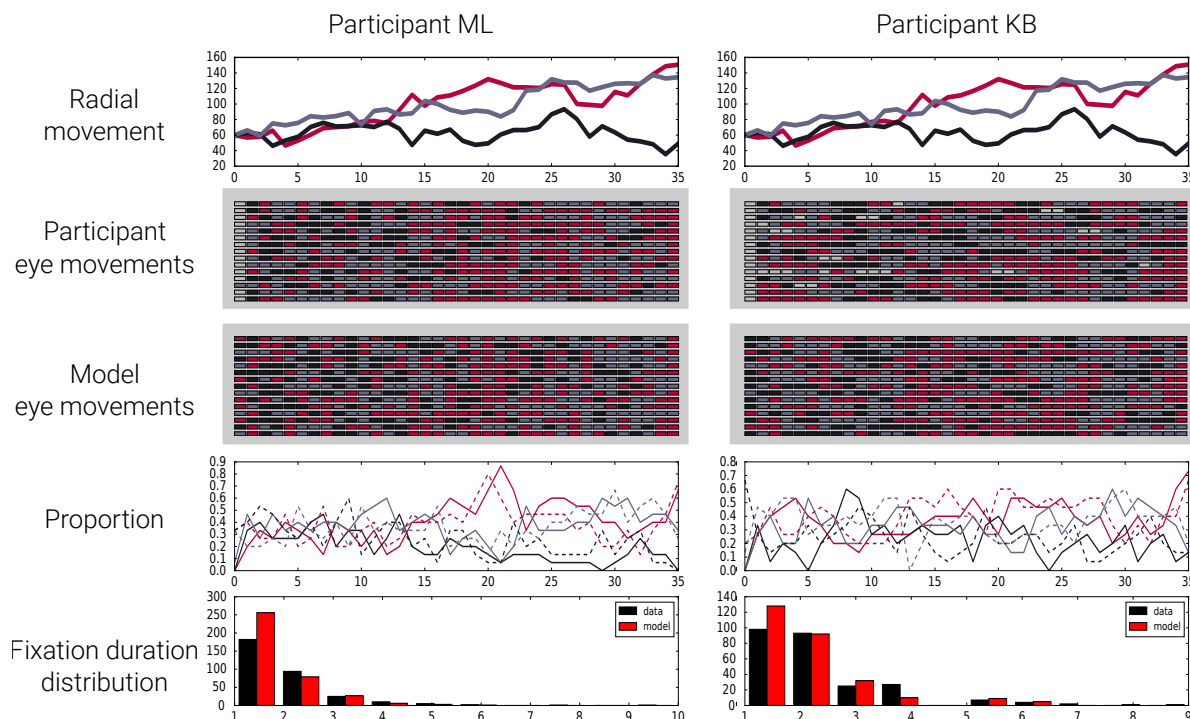
Figure 7.5: Eye movement sequences for two participants and data drawn from the model. Data is shown for the same trial (first row). Eye movement sequences of the participants are shown in the second row. Each square represents a time step in the experiment (300 ms). The colors of the squares correspond to the colors shown in the first row. Each line of squares is a single trial leading to a total of 15 lines per participants, as this stimulus was presented in 15 trials. The simulated data is shown in the third row. The proportion of how often a specific dot was fixated at each time step is shown in the fourth row. The solid lines correspond to the proportion computed from the participant data, the dashed lines to those of the computational model. The fifth row shows the fixation duration distribution.

7.2.4 Deciding when and where

So far, we have derived the probability of deciding for one of the three locations given the internal beliefs about the task statistics. However, it is unknown at what rate this probability distribution is evaluated and therefore at what rate actions are selected. For example, while in our task new information is presented every 330 ms it is not clear that subjects make decisions at the same rate. It is unlikely that we perform active decisions every time new information is perceived as information can be continuous and our information processing resources are limited.

This introduces ambiguities in our behavioral sequences, as the observation of a subject repeatedly fixating a location can originate either from (1) the subject actively deciding to stay at a location or (2) the subject not deciding. We need to account for this by explicitly modeling the different possibilities. Let us consider the eye movement sequence "1112" - a subject fixating Location 1 for three time steps and then switching to Location 2. We can incorporate the two possibilities in the computation of the likelihood by

$$p(\text{stay at } 1) = p(\text{decide}) \cdot p(\text{decide for } 1) + p(\overline{\text{decide}}) \tag{7.8}$$

where $p(\overline{\text{decide}}) = 1 - p(\text{decide})$ and $p(\text{decide for } 1)$ is the probability of Action 1 according to Equation 7.7. The total likelihood of the measured eye movement sequence "1112" can be
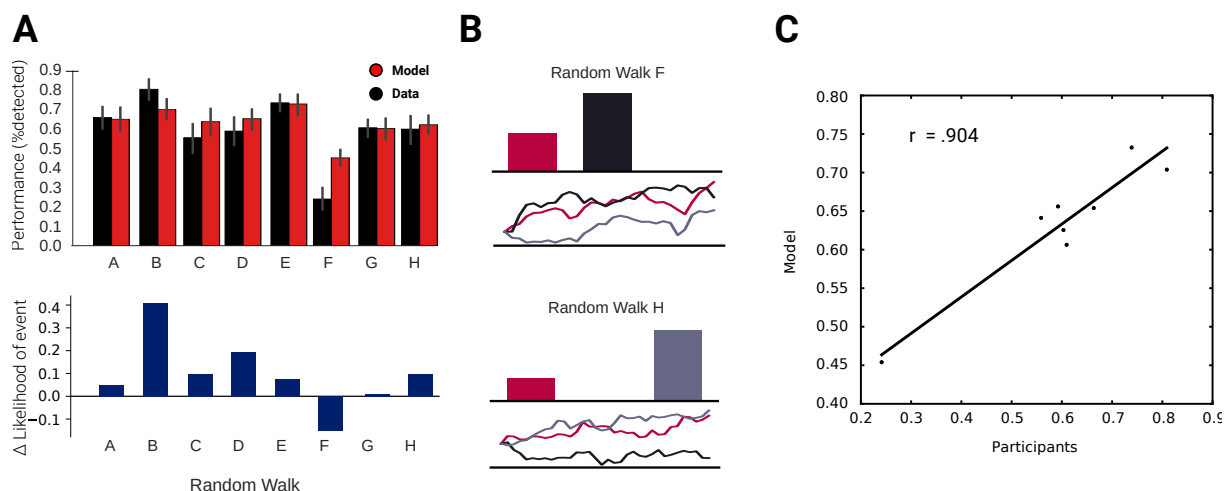
Figure 7.6: Task performance for human and model data. (A) Mean performance for the individual stimulus dynamics. Upper panel: The bars correspond to the random walks in Figure 7.3. Errorbars correspond to the 95 confidence interval of the mean. Lower panel: For each random walk we computed the likelihood for the outcome of the random walk three time steps prior to the event. Random Walk F is the only condition where another event was more likely thus explaining the drop in performance shown in the upper panel. (B) For Random Walk D and H the probability for events at each of the circles is shown three steps prior to the end of the trial. (C) Relationship between the human task performance and the model task performance.

computed as

$$p(1112) = p(\text{stay at } 1)^3 \cdot p(\text{decide}) \cdot p(\text{decide for } 2). \tag{7.9}$$

This approach captures both characteristics of the decision: which location to target next and when. Our probabilistic decision model that accounts for the rate of action selection is similar to using a stickiness factor for the softmax decision rule to address the preference for repeating the same action, that has been proposed in the past (see Gershman, 2016, for example).

7.2.5 Parameter estimation

Overall, we estimated three parameters for each subject: the probability of deciding at each step $p(\text{decide})$, an additional uncertainty about the location $sd = \sigma$ and the temperature coefficient $tau = \tau$. The parameter $t = p(\text{decide})$ accounts for differences in the temporal dynamics of the eye movement behavior, i.e., subjects differ with respect to their rate of action selection. Further, as it is unlikely that subjects have full knowledge about the random movement of the dots, we estimated our subjects additional uncertainty (variance σ^2) of the step size. The increased uncertainty can originate from various sources, i.e., uncertainty regarding the step size, uncertainty regarding the memory of the last location, as well as uncertainty regarding how much time has passed since the last observation. Finally, humans differ with respect to their action selection. We estimated the temperature coefficient τ of the softmax decision rule separately for each participant. All parameters were estimated through Maximum Likelihood using the computational model for the action selection together with our model of the temporal dynamics.

Table 7.1: Fitted parameter values for each subject. The fitted parameters comprise t (temporal dynamics), τ (action selection), and sd (additional uncertainty). In addition the likelihood for the model fit (fun) is shown for each participant as well as the fixation duration (fixdur) and the overall performance (correct).

part	fun	t	tau	sd	fixdur	correct
1	4262.02	0.846	0.172	11.878	0.565	0.650
2	4015.66	0.928	0.125	22.004	0.540	0.733
3	4428.64	0.964	0.163	4.201	0.626	0.658
4	2185.68	0.335	0.091	18.210	0.193	0.541
5	3978.10	0.631	0.196	1.964	0.446	0.525
6	4555.15	0.938	0.232	3.032	0.636	0.583
7	3880.24	0.667	0.156	8.206	0.441	0.533
8	3970.53	0.881	0.115	14.006	0.519	0.658
9	3810.49	0.611	0.195	18.784	0.428	0.600
10	4590.52	1.000	0.210	0.000	0.732	0.516
11	4195.44	0.706	0.226	2.888	0.499	0.558
12	4313.45	0.898	0.165	11.370	0.577	0.641
13	4659.51	1.000	0.299	0.000	0.834	0.425
14	3855.07	0.644	0.189	26.753	0.440	0.533
15	4290.45	0.951	0.160	14.749	0.602	0.616
16	4465.53	0.944	0.222	13.777	0.624	0.700
17	4526.29	1.000	0.221	7.764	0.748	0.616
18	4061.07	0.889	0.155	18.208	0.567	0.675
19	4171.98	0.761	0.185	11.559	0.505	0.591
20	4380.03	0.902	0.196	12.128	0.570	0.708
21	4200.85	0.779	0.173	6.642	0.504	0.575

7.3 Results

To investigate how closely our model captures the behavioral data, we used the parameter estimates for each participant to simulate eye movement sequences from our model (for the parameter estimates see Table 7.1). Thereby, we applied the model to each participant's trial yielding a simulated dataset that matches our human data in size. The results of the model fits for two different subjects as well as the raw eye movement data can be seen in Figure 7.5.

7.3.1 Subjects

Overall, 22 subjects (16 females) took part in the experiment in exchange for course credit. The participants' age ranged from 15 to 28 years ($M = 22$, $SD = 3.5$). All subjects had normal or corrected to normal vision (seven wore glasses, two contact lenses). Sufficient eye tracking quality was ensured for all subjects. Since the gaze-contingent paradigm used in the present study did not require a high amount of spatial accuracy (as it was only used to detect which of the three circles was attended), deteriorated signal quality introduced by glasses was not an issue.

7.3.2 Performance

Performance, measured as the percentage of detected events, differed between the experimental stimuli. Figure 7.6 depicts the proportion of events detected for each stimulus dynamic. Subjects showed the best results for the Condition B ($M = 0.81$, $SD = 0.14$). The least events were detected for Condition F ($M = 0.24$, $SD = 0.14$). The model results suggest a close connection between the model behavior and the participants' eye movements. Linear regression showed that the model behavior closely resembles our subjects data with respect to how they performed in the task ($p = 0.0002$, $R = .90$).

7.3.3 Temporal eye movements and stimulus dynamics

Our results show that humans are capable of scheduling eye movements in accordance with the stimulus dynamics. The proportion of eye movements targeted at a specific circle is far from uniform. Instead, it is affected by the location of the dot in that circle (see Figure 7.8). In particular, the distribution of gaze is very similar across subjects. Our simulated data reflect the eye movement patterns we observed in our subjects (see Figure 7.9). These results suggest that humans can execute temporal eye movement plans following dynamic environmental demands. They are able to monitor multiple changing locations and detect events by moving their focus to relevant spatial locations and monitoring these locations more closely.

The gaze patterns drawn from the estimated model showed similar temporal statistics as our subjects' data (Figure 7.7). In particular, the estimated rate of deciding t is sufficient to recover the distribution of fixation durations for individual participants. Surprisingly, our model is capable of explaining our data despite the implicit assumption that the probability of making a decision does not depend on the current reward distribution, a simplification made out of computational convenience.

7.3.4 Variability of eye movements

Finally, we investigated how variability in eye movement patterns was related to the stimulus dynamics. Variability in gaze behavior was computed as the entropy of the participants gaze
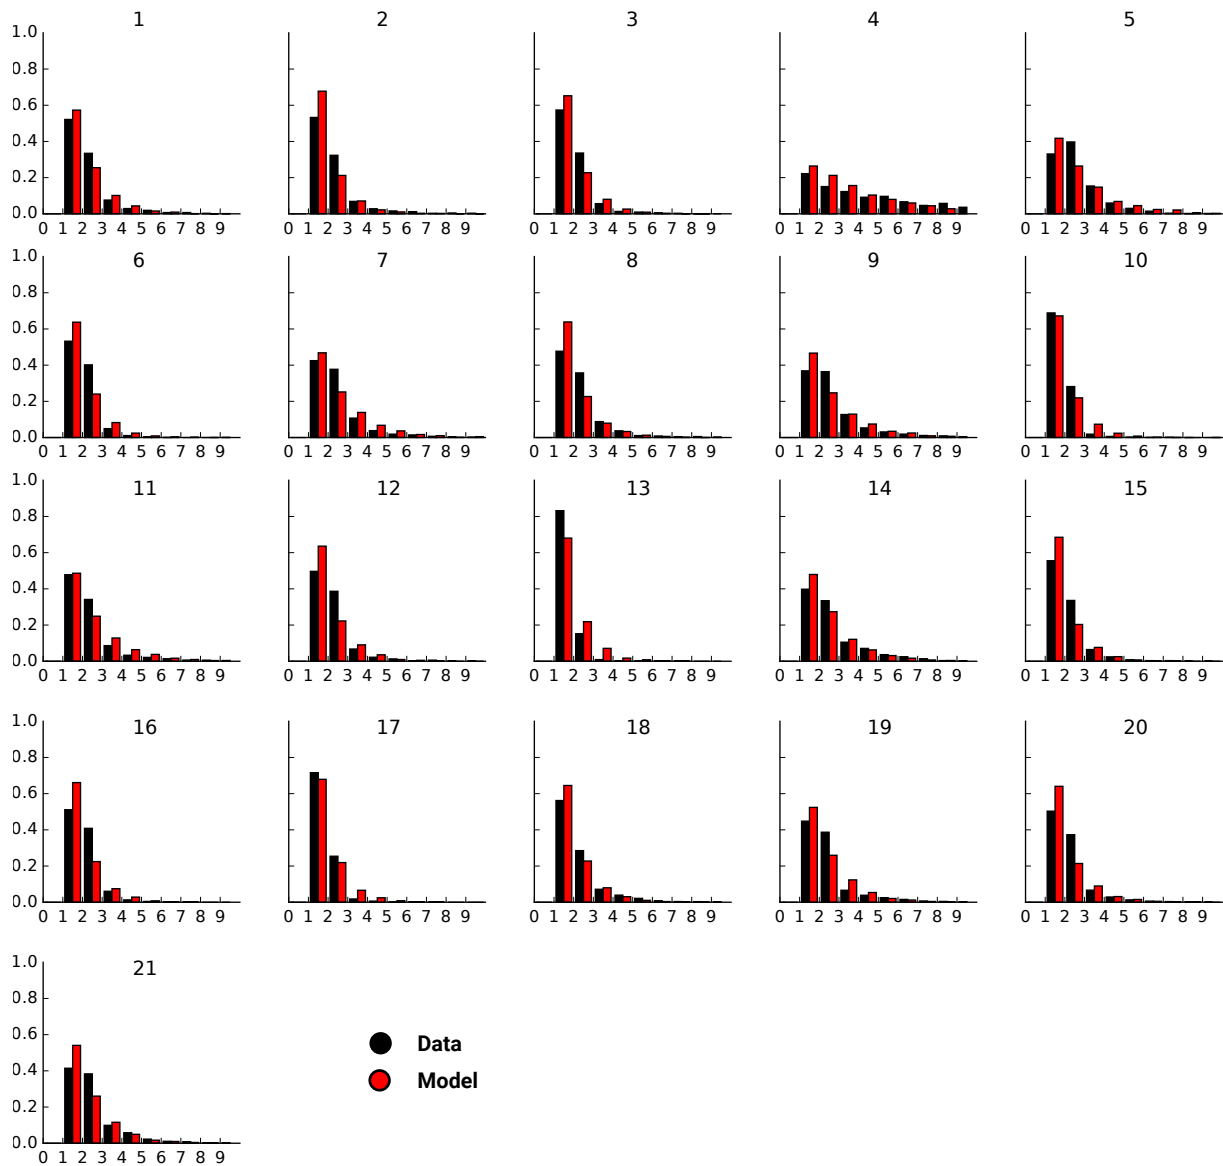
Figure 7.7: Participant fixation duration distributions and fitted model results. For each participant the fixation duration distribution is shown (black bars) as well as the results from the simulated data (red bars). The x-axis denotes the duration of the fixation (in time steps of the experiments 300 ms) and the y-axis the percentual proportion. The results for the individual paramter estimates can also be found in Table 7.1.

patterns for each time step

$$H(X) = -\sum_i P(x_i)logP(x_i) \,, \tag{7.10}$$

where X is the multinomial distribution of the aggregated eye movement patterns for each time step computed from the proportion each location is fixated at a particular time. For example, if all participants had fixated the same location at a particular time step, the distribution $p(X)$ would have been be $(1,0,0)$ and the entropy minimal. We found the entropy in the fixation patterns to decrease if the distance of the dot closest to the boundary is small (Figure 7.10A).
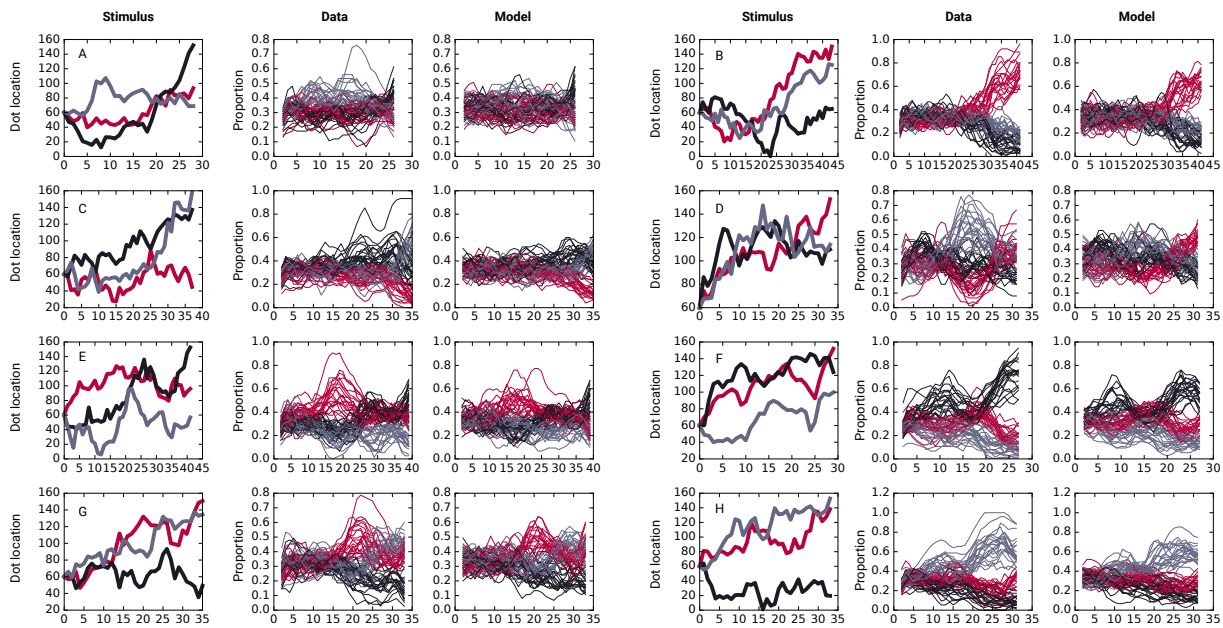
Figure 7.8: Fixation proportions for all eight different stimulus dynamics. For each stimulus dynamic the radial movement is shown (first and fourth column) as well as the fixation proportion of the participant data (second and fifth row) and the fixation proportion of the model (third and sixth row). For the fixation proportion, each line corresponds to a single participant (also compare to the fixation proportion in Figure 7.5). Lines correspond to a moving average of the fixation proportions with a window width of five time steps. For example, for the stimulus in the upper right, both the participants as well as the compuational model strongly favor the circle with the pink stimulus movement towards the end of the trial.

Our computational model well described this effect. The maximum distance is closely related to the utility of the correct action, i.e., how much reward can be expected if the correct action is chosen. We converted the maximum distance to the expected loss using the experimental dynamics. For each distance to the boundary, we can assign a probability of the dot exceeding the boundary - a quantity that is directly related to the expected loss. As a result, we reveal a nearly linear relationship between eye movement variability and expected loss (see Figure 7.10B). This relationship is also well described by our model.

7.4 Discussion

The goal of this study was to investigate variability in eye movement sequences. To this end, we developed a temporal monitoring task using stimuli with a simple spatial structure. Subjects switched their gaze between three location, at each location a dot was moving inside a circular boundary according to a stochastic process. The goal of the monitoring task was to detect when one of the dots left the circle. Crucially, participants were presented with each movement multiple times while not being aware of this. This way we were able to compute the gaze distribution for each subject separately for each time step, e.g., at time step five subject k fixated Region A 50 percent of the time. We developed a computational model for our monitoring task, that was able to reproduce key characteristics of our human data. Our results show that external environmental dynamics and task-related rewards guide eye movement sequences. We further show, that the variety of temporal statistics showed by subjects could be explained by
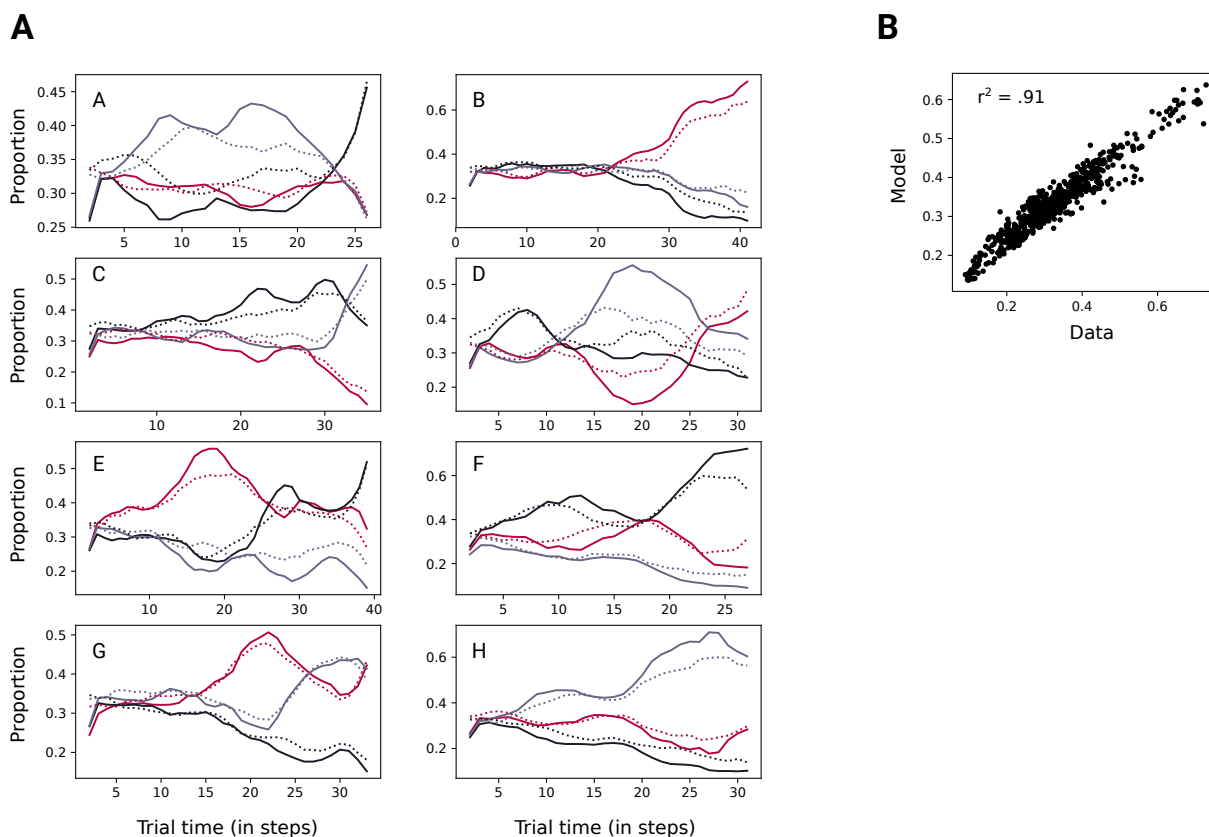
Figure 7.9: Model predictions and human data. (A) Fixation proportions for all eight different stimulus dynamics. The fixation proportions are shown for the aggregated data (dotted lines correspond to model predictions). The individual fixation proportions are shown in Figure 7.8. All data was smoothed using a moving average with a window size of three time steps. (B) Relationship between model and human data. Each dot corresponds to a single time step in (A).

a memory-less process, in which the probability for deciding what action to take at any time is independent of the current environmental state. Further, our model reflected key properties of the eye movement variability found in our data. Behavior became more similar across all participants if one of the dots was close to the boundary. In particular, this relationship became linear, when the distance was transformed to the expected loss in the task.

7.4.1   Eye movements in a monitoring task

The eye movement sequences drawn from our computational model were close to the measured human behavior. This suggests that participants were able to distribute gaze over time according to the demands of our task. In particular, they did so despite the fact that only a single location was visible at any time while the uncertainty regarding the other two increased. Due to our stimulus design, the behavior cannot be explained in terms of saliency (Itti et al., 1998), as all locations had the same visual features, or prior knowledge (Torralba et al., 2006). Also, subjects did not visit the locations in a fixed order, which would have resulted in an equal proportion of time spent at each location. Instead, the eye movements were connected to the stimulus dynamics in a complex manner ruling out a simple heuristic. Crucially, this behavior was shown by all participants indicating similar information processing across subjects.
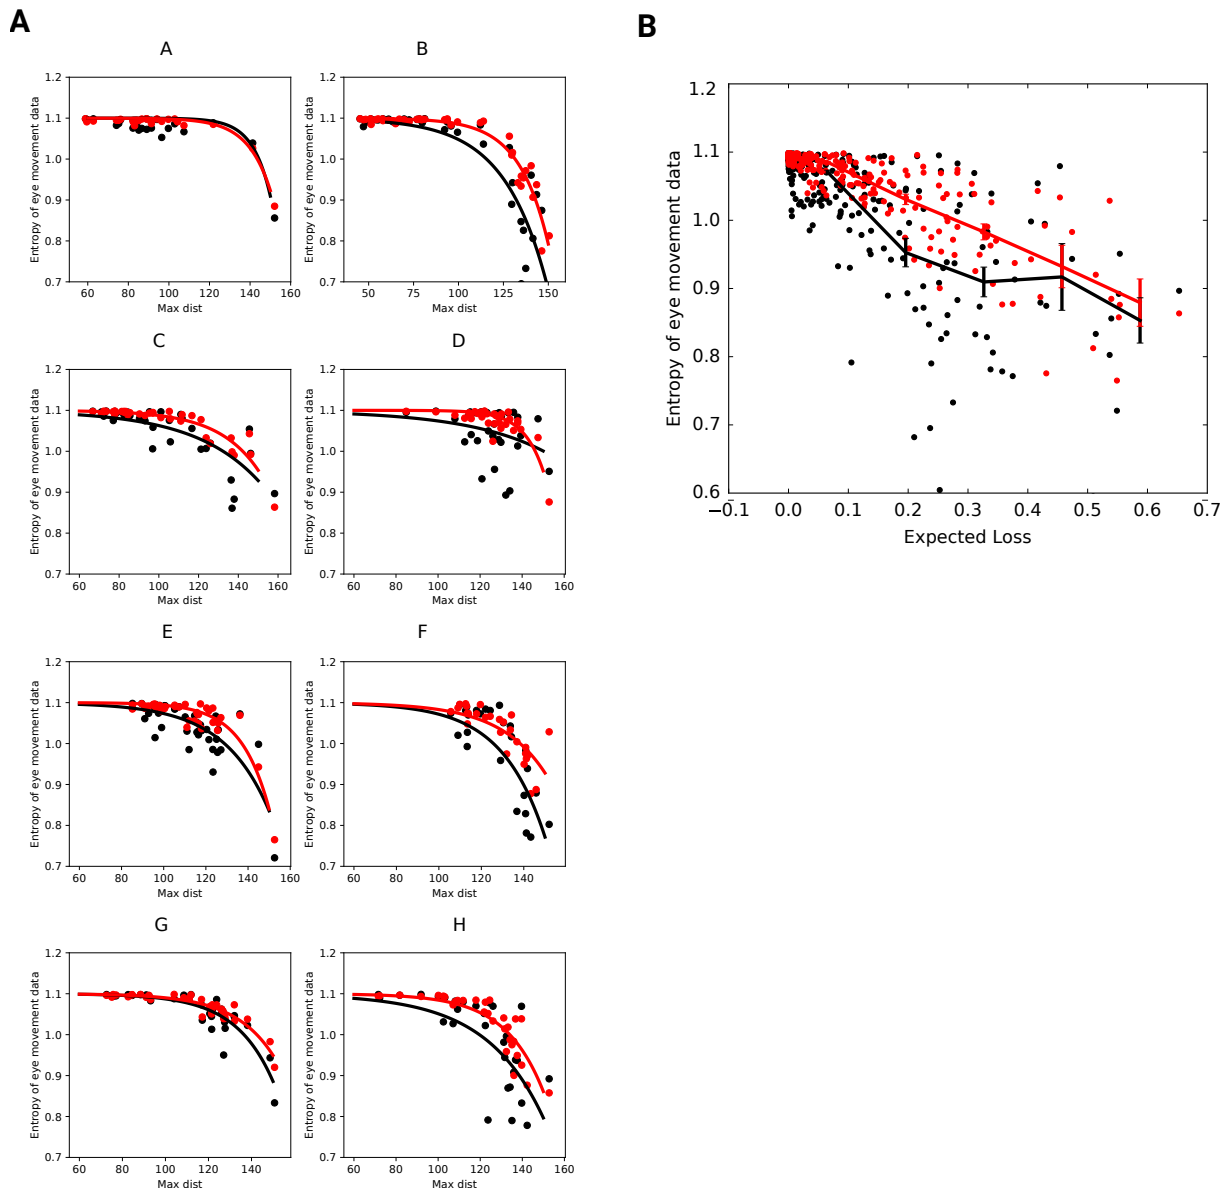
Figure 7.10: (A) Relationship between stimulus location and eye movement variability. For each trial the entropy in our eye movement data is shown dependent on how close the furthest of the dots was to the boundary. When all dots were far away from the boundary (e.g., in the beginning of each trial) entropy was maximal (1.08). However, the closer they are to the boundary, variability in terms of entropy in our eye movement data (black) as well as in our model simulations (red) decreases. (B) Relationship between entropy of eye movement sequences and expected loss. Expected loss was computed using the subject's belief that the most eccentric dot would hit the boundary. Dots correspond to the raw data for our subjects (black) and our model (red). Error bars correspond to the standard error of the mean.

7.4.2 Variability is task-dependent

The connection of variability and expected loss quantified in the present experiment is in accordance with related empirical (Valero-Cuevas, Venkadesan, & Todorov, 2009) and theoretical studies (Todorov & Jordan, 2003) reporting evidence for the minimal intervention principle. According to this principle, deviations in actions are corrected less if the consequences with respect

to reward are small. Also, fixation durations of infants have shown to be less variable when cognitive control is increased (Wass & Smith, 2014). Here, we showed that the variability of eye movement sequences is linearly related to the expected loss. We used a softmax decision rule with temperature parameters estimated separately for each participant. The softmax decision rule is frequently used to account for randomness in behavior (see Daw & Doya, 2006, and the references therein). By estimating the temperature coefficient we can represent the extent to what subjects acted according to the optimal Q values. Subjects do not have direct access to the real rewards and computations involved in computing the optimal action at every moment in the experiment are only feasible with limited precision. Prior research has shown, how using an approximation can be optimal, if the computations are costly (Vul, Goodman, Griffiths, & Tenenbaum, 2014) which they usually are (Kool et al., 2010; Lennie, 2003). The presented behavioral results and the insights gained from our computational model manifest a connection between properties of the task's reward structure and the variability of human actions. The findings are not limited to eye movements, instead they are a general prediction of the framework. Hence future work can apply the approach to test whether this relationship can also be found in other domains.

<div align="right">

8

</div>

<div align="right">

CHAPTER

</div>

Spatial planning of eye movements

8.1 Introduction

Actively deciding where to direct our eyes is an essential ability in fundamental tasks, which rely on acquiring visual information for survival such as gathering food, avoiding predators, making tools, and social interaction. As we can only perceive a small proportion of our surroundings at any moment in time due to the spatial distribution of our retinal receptor cells (Land & Nilsson, 2002), we are constantly forced to bring task-relevant parts of the visual scene into focus using eye movements (Findlay & Gilchrist, 2003). Thus, vision is a sequential process of active decisions. These decisions have been characterized in terms of optimizing performance in the ongoing task (Najemnik & Geisler, 2005; Torralba et al., 2006; Peterson & Eckstein, 2012; Hoppe & Rothkopf, 2016; Yang, Lengyel, & Wolpert, 2016), maximizing knowledge about the environment (Itti & Baldi, 2006; Renninger, Coughlan, Verghese, & Malik, 2005; Renninger et al., 2007), or targeting gaze towards locations that are most salient (Itti & Koch, 2000).

The capability of directing gaze to relevant parts of the environment is crucial for our survival. Computational models based on ideal-observer theory have provided quantitative accounts of human gaze selection in a range of visual search tasks. According to these models, gaze is directed to the position in a visual scene, at which uncertainty about task-relevant properties will be reduced maximally with the next look. However, in tasks going beyond a single action, delayed rewards can play a crucial role, thereby necessitating planning. Here, we investigate whether humans are capable of planning more than the next single eye movement.

Surprisingly, all of the reviewed computational models for eye movements are myopic, i.e., they choose actions that maximize the immediate reward (Najemnik & Geisler, 2005; Naval-pakkam et al., 2010; Schütz et al., 2012; Peterson & Eckstein, 2012; Hoppe & Rothkopf, 2016; Yang, Lengyel, & Wolpert, 2016). In practice, the problem of delayed rewards is circumvented by either investigating only single saccades or by choosing tasks where both policies lead to equivalent solutions. To our knowledge, there exist neither computational models nor empirical data examining whether humans are capable of planning eye movements. The execution of eye movement sequences has been subject to psychological research and results have shown that the latency of the first saccade was higher for longer sequences of saccades (Zingale & Kowler, 1987). Also, discrimination performance was enhanced at multiple locations within an instructed sequence of saccades (Baldauf & Deubel, 2008a). Further, if an eye movement plan was interrupted by additional information midway, the execution of the second saccade was delayed (De Vries, Hooge, & Verstraten, 2014). Although these results indicate that a scanpath of at least two saccades is internally prepared before execution, no light is shed on whether multiple future fixation locations are jointly chosen to maximize performance in a task.

8.1.1 The ideal observer

To understand the requirements of perceptual tasks, ideal-observer analysis (Geisler, 2003, 2011) has been very successful based on the idea that visual perception is inference of latent causes based on sensory signals (Knill & Richards, 1996; Kersten et al., 2004). In this framework, the goal of the visual system is to use sensory data o to infer unknown properties of the state s of the environment. For example, s could be indicating whether a predator is hiding behind a bush, and by directing gaze to the bush visual data o about the latent variable describing the true state s of the environment is obtained. This information can be incorporated into what is known about s using Bayes' theorem $p(s|o) = p(o|s)p(s)/p(o)$ (see also Chapter 3). Hence, the ideal observer combines prior knowledge $p(s)$ and sensory information $p(o|s)$ to form an updated posterior belief about environmental states relevant to the specific task. The ideal-observer paradigm has been used successfully to understand how humans choose locations for the next saccade. Specifically, human eye movements use the current posterior and target the location where they expect uncertainty about task-relevant variables to be reduced most after having acquired new data from that location in situations such as visual search (Najemnik & Geisler, 2005), face recognition (Peterson & Eckstein, 2012), and temporal event detection (Hoppe & Rothkopf, 2016).

A limitation of ideal-observer theory is that performing sensory inference by itself does not prescribe an action, i.e., information about s in the end needs to be used for action selection, e.g., whether to flee. The costs and benefits for the potential outcomes of the action can be very different, e.g., not to flee if a predator is present is more costly than an unnecessary flight. Bayesian decision theory provides such an answer by using the costs and benefits of different outcomes with the respective uncertainties of the associated outcomes. Hence, different potential outcomes of s are weighted with a reward function $R(a, s)$ to determine the action with highest expected utility: $a = \arg\max_a \int_s R(a, s)p(s|o)ds$. Thus, it may be better to flee, even when one is not certain that a predator is hiding behind a bush, because the consequences may be particularly harmful. Interestingly, if all information is equally valuable in terms of reward, the optimal action targets the location where the next fixation will reduce uncertainty the most and not the location that currently looks like the most probable target location. Indeed, both explicit monetary rewards (Schütz et al., 2012) and implicit behavioral costs (Hoppe & Rothkopf, 2016) in experimental settings have been shown to influence eye movement choices.

8.1.2 The ideal planner

However, Bayesian decision theory is limited to a particular subset of visual tasks, namely tasks that do not involve planning. Repeatedly taking the action with the maximum immediate utility may, in general, fail in tasks with longer action sequences and delayed rewards depending on the specific task structure. In these cases, an ideal planner based on the more powerful framework of belief MDPs, which contains the ideal observer and the Bayesian decision maker as special cases, is needed to find the optimal strategy. In a belief MDP only partial information about the current state s is available, therefore a probability distribution over states is kept as a belief state $b = p(s \mid D)$ (Kaelbling et al., 1998). The action-value function Q denotes the expected reward associated with performing action a in a belief state b:

$$Q(b, a) = \int_{b'} \{P(b' \mid b, a)[R(b') + \gamma V^*(b')]\}db' \tag{8.1}$$

where $V^*(b')$ is the expected future reward gained from the next belief state b'. Essentially, this
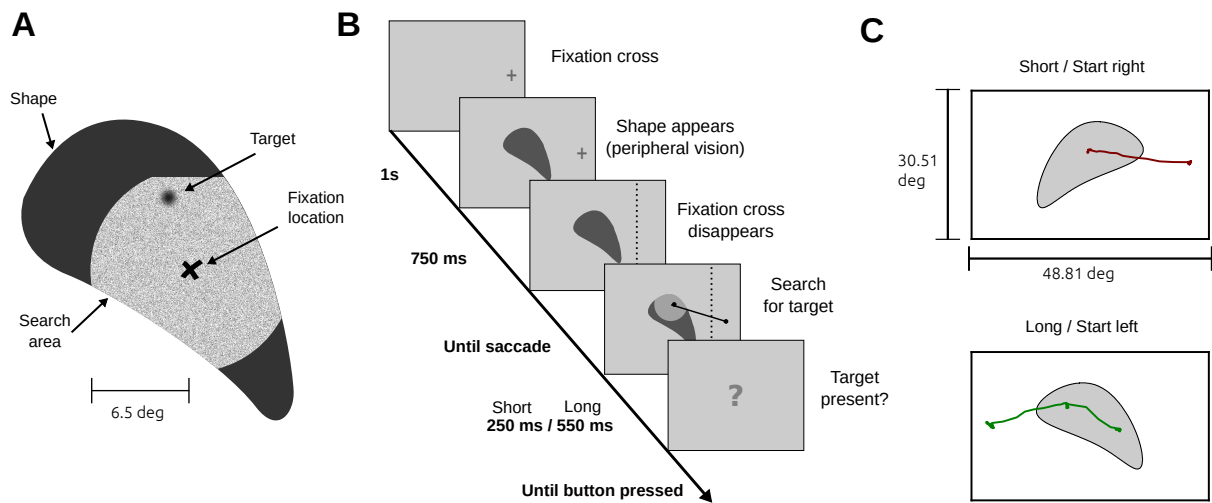
Figure 8.1: Experimental design. (A) Gaze contingent visual search paradigm. Targets were only visible in close proximity to the current fixation location (i.e., inside the search area). (B) Procedure for a single trial. Subjects fixated a fixation cross either shown on the left or the right side, respectively. The shape appeared 750 ms prior to the start of the search. The search time was initiated by the participants' gaze crossing the dotted line. The line, however, was not visible to the subjects. Depending on the condition (short or long) subjects were able to perform one or two fixations inside the shape. (C) Raw gaze data is shown for a trial with short search time and initial fixation on the right side (upper panel) and for a trial with long search time and initial fixation on the left side (lower panel). Shapes were mirrored in a counterbalanced design to ensure equal orientation with respect to the initial fixation cross.

means that the value of an action based on the current belief is a combination of the immediate reward and the expected long-term reward, weighted by how likely the next belief is under the action. Thus, as the belief about the state of task-relevant quantities depends on uncertain observations, actions are influenced both by obtaining rewards and obtaining more evidence about the state of the environment.

In the present study, we devised a task that allows probing whether ideal-observer models are sufficient to describe human eye movement strategies. For our visual search task, we derived computational models based on ideal-observer theory as well as on the framework of belief MDPs. Using these models, we specifically created our stimuli such that the two models led either to different behavioral sequences or the same. The rationale for this was not only to show the differences between ideal planer and ideal observer but to also demonstrate that the solutions of both may lead to the same action sequence, depending on the structure of the specific task. Using this experimental paradigm we are able to test whether human eye movement strategies follow the computational principles underlying ideal-observer theory and sequential Bayesian decision making or whether the strategies are planned and future rewards need to be considered (belief MDP).

8.2 Experimental design

8.2.1 Task

In our task subjects searched for a hidden target within irregularly bounded shapes (Figure 8.1A). Using a gaze-contingent paradigm the hidden target only became visible if a fixation

landed close enough ($||\mathbf{p}_{\text{Fix}} - \mathbf{p}_{\text{Tar}}|| < 6.5°$). The search area was made explicit by showing the shape's texture for all points closer than 6.5° to the fixation location. Targets within that area became visible to the participant after a delay of 130 ms. This was done to prevent participants from sliding over the image and instead encourage them to perform distinct fixations. The texture was chosen to reinforce the feeling of looking through the shape (subjects were told to imagine wearing x-ray goggles).

A single trial was as follows (Figure 8.1B): Participants fixated a fixation cross that was randomly presented either on the left or the right side of the screen. After 1 s the shape was shown in the center of the screen, thus subjects had access to the contour through peripheral vision. Shapes were mirrored if necessary yielding equal distances for left and right starting points. After 750 ms the fixation cross disappeared and participants could initiate the search for the target. Trials in which the first saccade was made while the fixation cross was still visible were dismissed and had to be repeated. It was made transparent to the participants that the search interval started once they made the first eye movement as opposed to when the fixation cross disappeared. After the search interval was over the shape disappeared and participants stated, whether it contained a target. Overall, shapes contained a target in half the trials. We used two durations as search intervals: a short interval (250 ms) providing enough time for a single saccade and a long interval (550 ms) providing enough time for two saccades. Trials were presented in blocks either containing only short intervals or long intervals, respectively.

8.2.2 Materials

Our computational models enabled us to specifically select shapes that facilitate testing our hypothesis. In particular, we identified stimuli that triggered different policies for the ideal-observer model and the ideal-planner model. First, multiple candidates shapes were generated using the following approach: Five points were drawn uniformly in a bounded area ($23.24° \times 23.24°$). Next, a B-spline was fitted to the random points. Finally, the shapes bounded by the splines using the fitted parameters were filled with a texture (white noise). We applied both models to identify shapes that lead to different policies (see Figure 8.7). Overall, four different shapes were used in the experiment (see Figure 8.2B). We chose two shapes where optimal behavior requires planning (S3 and S4) and two where it does not (S1 and S2), i.e., where the sequence of eye movements from the ideal observer and the ideal planner coincide. In each category, we selected two shapes by visual inspection ensuring that they were similar with respect to the area covered. For the display in the experiment the shapes were upscaled with a factor of 1.5 and centered on the monitor such that the center of the shapes bounding box matched the center of the screen. The target was a circular grating stimulus (0.87°in diameter). The contrast was set in a way that it was easily detected if it was within the visible search radius of the current fixation. The target's position was generated by randomly choosing a location within the shape.

8.2.3 Procedure

After signing a consent form, the eye tracker (SMI Red, 250 Hz) was calibrated using a 3 point calibration. Subsequently, subjects completed three to five short training trials, (about 1 minute) as part of the experiment instruction. During these training trials it was ensured that the search time was sufficiently long for the individual subject to execute a single saccade in the short condition and two fixations in the long condition, respectively. If necessary, the search time was adjusted (between 500 ms and 580 ms, for the long search interval). Participants
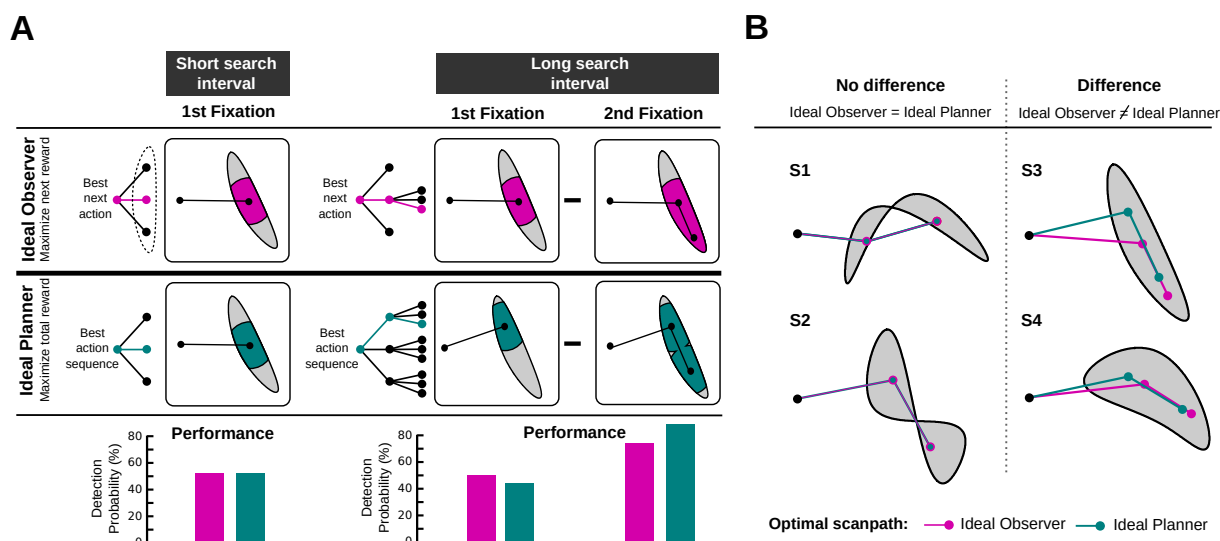
Figure 8.2: Computational models for visual search. (A) Illustration of optimal scanpaths for both models depending on the search time. For the short search interval (left side, one fixation) both models show the same behavior. For the long search interval (right side, two fixations), the ideal observer and the ideal planner differ with respect to the scanpath. While the ideal observer's next fixation is chosen to maximize the immediate reward (better performance after the first fixation, bottom row), the ideal planner's scanpath is chosen to maximize performance after two fixations. Computational complexity (depicted as decision trees) is higher for the ideal planner as in the condition with long search intervals all two-fixation sequences are evaluated in order to maximize performance. (B) Shapes used in our visual search experiment. For each shape, the optimal policy is shown for the ideal observer (pink) and the ideal planner (green). Whether these models lead to different strategies depends on the particular shape. Scanpaths are the same for Shapes S1 and S2, but differ for S3 and S4.

were encouraged to ask questions if anything was unclear. After training, participants answered ten items from a checklist to ensure that they understood the task correctly (e.g., when does the search interval start and how many targets can be found at most). Incorrect answers were documented and the solution was discussed. After successfully finishing the training, four blocks each containing 100 trials were performed. Thereby, the order of the blocks was either SSLL (two blocks with short search time followed by two blocks with long search time) or LLSS. Participants were randomly assigned to one of the two orders. Eye tracking calibration was renewed before each block.

8.3 Model

8.3.1 Visual search as planning under uncertainty

To develop a computational model of visual search as optimal planning under uncertainty it is first necessary to specify the relevant quantities describing the task, i.e., the state representation. In our visual search task (Figure 8.1), a suitable candidate for a state representation is the target location and the current location of gaze. However, in general, the exact location of the target is unknown. Therefore, we formalize the probability distribution of the target as a belief state. The action space comprises potential fixation locations and with each action, we receive information about the target, update our belief and transition to the next belief state. The reward function is an intuitive mapping between the belief state, which comprises the knowledge about the location

of a potential target, and the probability of finding the target.

How should the actor decide where to look next according to this framework? A policy π is a sequence of actions and the optimal policy π^* comprises actions $a = \arg\max_a Q(b, a)$ that maximize the expected reward. In tasks comprising sequences of actions, the optimal strategy, the ideal planner, incorporates rewards associated with future actions $V^*(b')$ into action selection. As a result, the sequence of actions that leads to the maximum total reward is chosen:

$$\pi^*_{\text{ideal planner}} = \underset{a_0, a_1, \ldots, a_n}{\operatorname{argmax}} \mathbb{E}[r_0 + \gamma r_1 + \cdots + \gamma^n r_n], \tag{8.2}$$

where γ is the discount factor, which controls how much future rewards influence the current action selection.

Ideal observer as a special case of the ideal planner. If we are only interested in the optimal next action ($\gamma = 0$) or if there is only a single action to perform Equation (8.1) simplifies to:

$$Q(b, a) = \int_{b'} P(b' \mid b, a) R(b', a) db' \tag{8.3}$$

where $P(b' \mid b, a)$ is the posterior over relevant quantities in the task and $R(s, a)$ is the cost or reward function. Therefore, if reduced to the next action alone, the ideal planner reduces to the ideal observer with an action selected to maximize task success after the next action. For sequences of actions, the sequential application of the ideal-observer paradigm leads to the action sequence:

$$\pi_{\text{ideal observer}} = \left(\underset{a_0}{\operatorname{argmax}} \mathbb{E}[r_0], \underset{a_1}{\operatorname{argmax}} \mathbb{E}[r_1], \ldots, \underset{a_n}{\operatorname{argmax}} \mathbb{E}[r_n] \right), \tag{8.4}$$

where $a_0, \ldots, a_n$ is the sequence of actions that yields the maximum expected return r_t for each time step t. Whether $\pi_{\text{ideal observer}}$ and $\pi^*_{\text{ideal planner}}$ lead to the same action sequence depends on the specific nature of the task. However, in general:

$$\pi^*_{\text{ideal planner}} \neq \pi_{\text{ideal observer}} \tag{8.5}$$

as can be seen in Figure 8.2. Ideal-observer approaches only lead to optimal actions if future rewards do not play a role, for example, if only a single action is concerned.

8.3.2 Ideal observer and ideal planner formalizations

Here we derive expressions that implement the general mechanisms of Equation (8.2) and (8.4) for our visual search task. According to our experimental design participants directed their gaze to suitable locations within a shape in order to decide if a target was present. Depending on the condition, the action sequence in our task comprised one (short condition) or two (long condition) fixation locations. Formally, the greedy policy of the ideal observer (Equation (8.4)) leads to the sequence of fixation locations $(x_0, y_0), (x_1, y_1), \ldots, (x_n, y_n)$ that maximizes the quality of the decision after each step. In the case of two fixations this leads to:

$$\pi_{\text{ideal observer}} := \left(\underset{(x_0, y_0)}{\operatorname{argmax}} p(\text{correct} \mid x_0, y_0), \underset{(x_1, y_1)}{\operatorname{argmax}} p(\text{correct} \mid x_0, y_0, x_1, y_1) \right) \tag{8.6}$$

where x_n, y_n are the coordinates of nth fixation location and $p(\text{correct}|x_n, y_n)$ denotes the probability of deciding correctly whether a target is present after the nth fixation.

The non greedy policy of the ideal planner can be derived from Equation (8.2) in a similar fashion. Again, we consider the case of two fixations (LI). Here, the next fixation location is determined by maximizing the reward simultaneously using the next two fixation locations:

$$\pi_{\text{ideal planner}} := \underset{(x_0,y_0),(x_1,y_1)}{\text{argmax}} \; p(\text{correct} \mid x_0, y_0, x_1, y_1) \qquad (8.7)$$

Thereby, (x_0, y_0) is the next location and (x_1, y_1) is the location thereafter. By jointly optimizing the entire sequence of fixation locations the ideal planner is always equal or better compared to the ideal observer. Intuitively, $\pi_{\text{ideal observer}}$ and $\pi_{\text{ideal planner}}$ yield the same action sequence if the sequence only contains a single action, i.e., a single fixation. Also, the first fixation location of the ideal observer is the same for both conditions. Crucially, this is not the case for $\pi_{\text{ideal planner}}$. By jointly maximization the reward over the whole action sequence, even the first fixation location can differ between the conditions.

Next, we derive the probability of a correct decision given a sequence of fixation locations since both proposed policies depend on the performance in the task, i.e., the detection probability. The probability of correctly judging the presence of a target is proportional to the area covered by the search. This can be computed as:

$$p(\text{correct} \mid x_n, y_n) \propto \sum_x \sum_y P_T(x, y) P_O(x, y | x_n, y_n) \qquad (8.8)$$

where $P_T(x, y)$ is the probability that the target is located at (x, y) and $P_O(x, y | x_n, y_n)$ is the probability that the location (x, y) is covered by the search given that the saccade was targeted at (x_n, y_n). The former is $1/N$ if (x, y) lies within the shape and zero otherwise, where N is the number of possible target locations. The latter depends on the distance between the saccadic target (x_n, y_n) and the target location (x, y). Therefore:

$$P_O(x, y | x_n, y_n) = \begin{cases} 1 & \text{if } || \, [x_n - x, y_n - y]^T \, || < \text{threshold} \\ 0 & else \end{cases} \qquad (8.9)$$

where the threshold is equal to the radius of the search area (6.5°).

8.3.3   Model extensions

The application of the two models to the shapes used in the experiment yields the gaze strategies shown in Figure 8.2. To take into account known cognitive and biological constraints we need to incorporate several well-known characteristics of the human visual system. We introduced costs on the saccade amplitude thus favoring smaller eye movements. As was shown by prior research, greater amplitudes lead to higher endpoint variability (van Beers, 2007) and longer saccade duration (Baloh, Sills, Kumley, & Honrubia, 1975). It has further been demonstrated that humans attempt to minimize endpoint variability when execution eye movements (Harris & Wolpert, 1998). Therefore, we hypothesized that subjects show a preference for smaller saccade amplitudes. Computationally, we obtain the total reward as a combination of performance and saccade amplitude

$$r_n(\alpha) = p(\text{correct}|x_n, y_n) - \alpha c(x_n, y_n) \qquad (8.10)$$
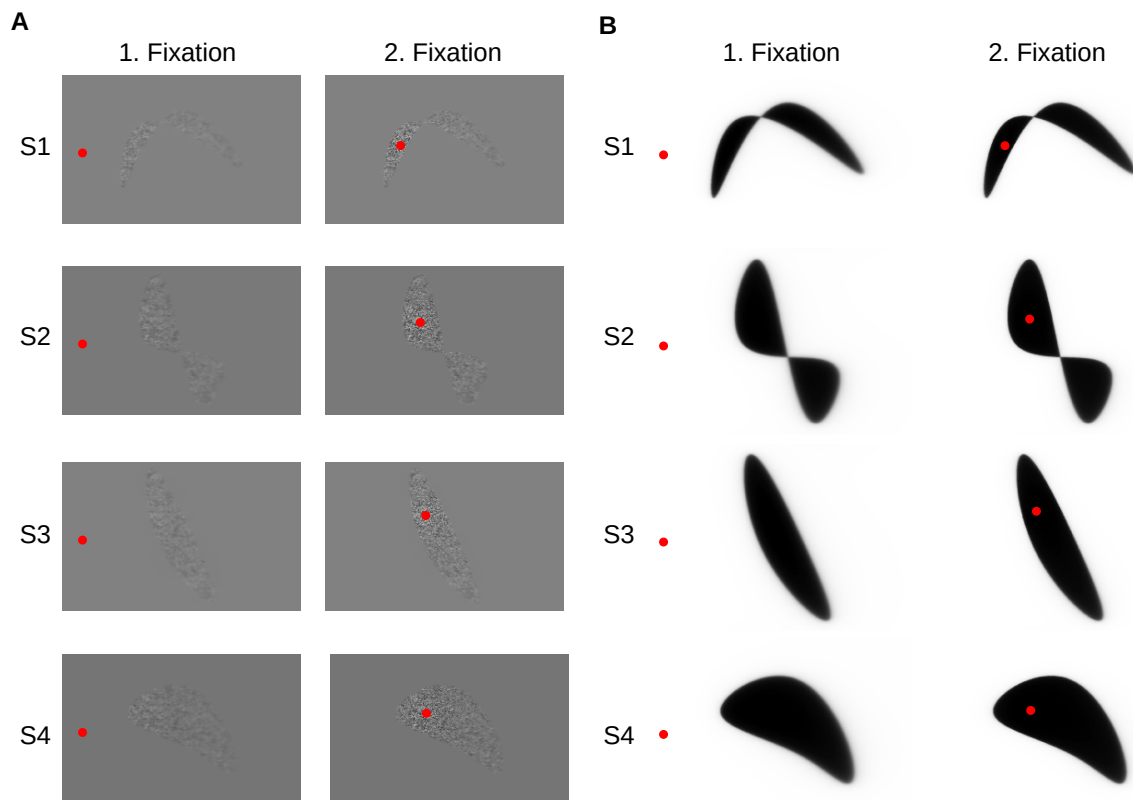
Figure 8.3: Foveated versions of the experimental stimuli. (A) Red circles correspond to the current fixation. All shapes depict the location of the planned policy as second fixation location. Shapes are filled with white noise to illustrate the decline of details with greater distance to the fixation location. (B) Same as (A). Instead of texture the probability of a location being part of a shape is shown. The black area inside the shape corresponds to a high probability of being part of the shape.

where c is a linear cost function returning the amplitude of the saccade. The parameter α determines how much detection probability a subject is willing to give up to decrease saccade amplitude (Hoppe & Rothkopf, 2016). This was estimated from the mean fixation locations of our participants using least squares.

Next, the human visual system does not have access to visual content at all locations in the field of view with unlimited precision. We accounted for the decline of visual acuity at peripheral locations (Figure 8.3). Therefore, foveated versions of the shapes were generated using the known human contrast sensitivity function (see Geisler and Perry (1998), Najemnik and Geisler (2005), Peterson and Eckstein (2012), for example). For the first fixation foveation was computed using the initial fixation location of the trial. As it was not computationally tractable to compute foveated images corresponding to the exact location of the first landing position, we approximated it by using the mean fixation location of our subjects instead. The same contrast sensitivity function was used for all participants, hence we did not account for interindividual differences in peripheral vision.

Finally, prior studies have shown that saccades undershot target locations (Gillen, Weiler, & Heath, 2013). Initial landing positions are closer to the start location of a saccade. The final target is reached using subsequent corrective saccades. However, in our experiment there is no visible fixation target, therefore corrective saccades might not be present. To account for that,
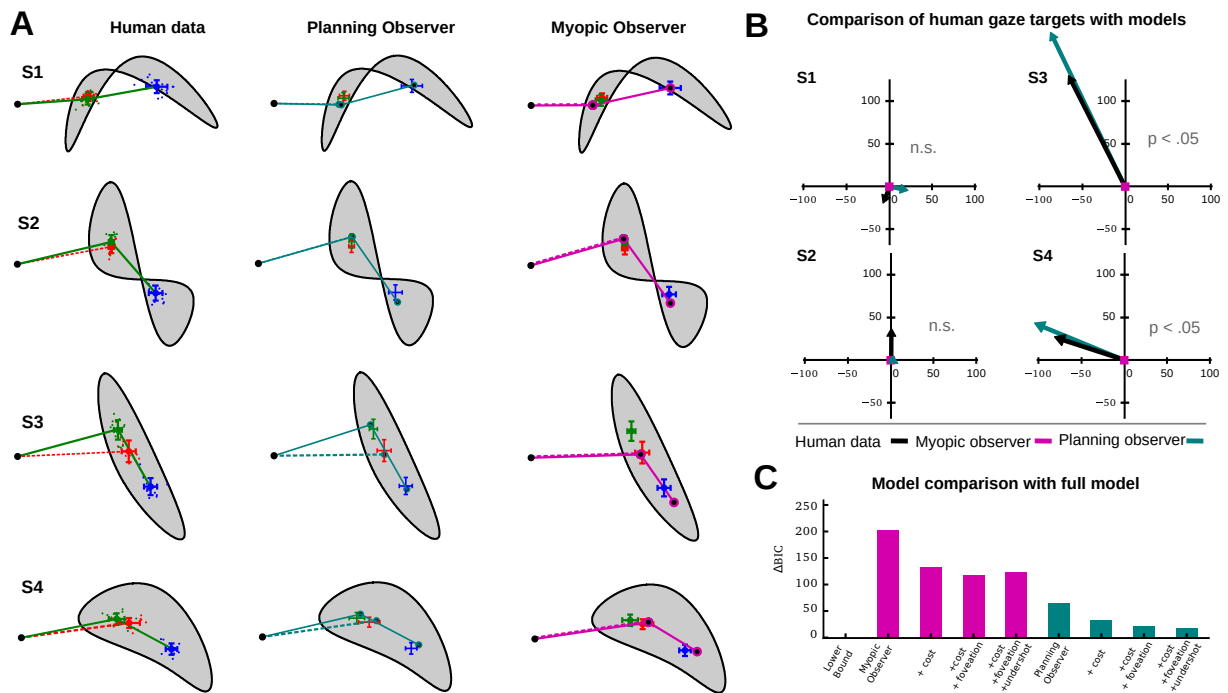
Figure 8.4: Behavioral und model results. (A) Mean human scanpaths for both conditions (solid lines correspond to long search intervals, dashed lines correspond to short search intervals) are shown in the left column. Colors refer to the condition and the position within the scanpath (red: short search interval, green: first fixation in the long search interval, and blue: second fixation in long search interval). Dots depict mean fixation locations aggregated for each subject individually, error bars show the standard deviation for the fixation location aggregated over all data. The scanpaths suggested by the best fitting models for the ideal planner and the ideal observer are shown in the center and the right column, respectively. Again, solid lines depict the strategy for the long search interval, dashed lines for the short search interval. Global means of the human data are also shown for reference (red, green, and blue). (B) Actual and predicted spatial relation of first saccades for all four shapes. Graphs are centered at the fixation location in the short search interval condition. Arrows depict the displacement of the first fixation location in the long search interval relative to the short interval. Arrow color corresponds to the data source. For the ideal observer, the first fixation location is the same for both conditions (indicated by the square centered at (0,0)). (C) Difference in BIC between all tested models. The lower bound corresponds to a model directly estimating the mean fixation locations for each shape and condition from the data (3×4 means).

we estimated the undershot from our data.

8.4 Results

8.4.1 Participants

Overall, 16 subjects (6 female) participated in the experiment. The subjects' age ranged from 18 to 30 years ($M = 21.8$, $SD = 3.1$). Participants either received monetary compensation or course credit for participation. All subjects had normal or corrected to normal vision (four wore contact lenses). One subject stated to have dyschromatopsia, which did not influence the experiment. Sufficient eye tracking quality was ensured for all data entering the analysis. In each trial, a single fixation location (short search interval) or a sequence of two fixation locations (long search interval) entered the analysis. Further, informed consent was obtained from all participants and

all experimental procedures were approved by the ethics committee of the Darmstadt University of Technology.

8.4.2 Preprocessing

First, fixations were extracted from the raw gaze signal using the software of the eye tracking device. Overall, 6400 trials (16 participants × 4 blocks × 100 trials per block) entered the preprocessing. 15 trials (0.23 %) were dismissed because the subjects failed to target gaze towards the shape. In these trials, subjects triggered the beginning of the trial by crossing the boundary, however did not engage in visual search. While search time was adjusted to enable subjects to perform a single saccade in the short condition and two saccades in the long condition, respectively, in 17 % of the trials subjects failed to do so. Since we are only interested in comparing the difference between strategies consisting of one or two targeted locations we only used the remaining 5288 trials. Next, we excluded trials where the target was present, regardless of whether it was found, leaving 2589 trials. Clearly, behavior after successfully finding the target is confounded and does no longer provide valid information about the search strategy. Also, trials in which a target was shown but not found are biased as they are more likely to occur in the context of inferior eye movement strategies.

Our analysis and our estimated model parameters rely on mean landing positions aggregated within subjects. Therefore, we need to make sure that the variation in landing positions arises due to saccadic endpoint variability or uncertainties the subject might have about the shape, but not from qualitatively different strategies. Shapes S1 and S2 consist of two separate parts, as a consequence the reward distribution is no longer unimodal across potential gaze targets (see Figure 8.5A). Indeed, qualitatively different strategies in the short condition were found for these stimuli (see Figure 8.5B). Using mean gaze locations therefore would have lead to misleading results as it implicitly implied unimodal variability in landing positions while the real data showed clear multi-modality. To further analyze the gaze targets of our participants, we first identified the strategy for each trial using a Gaussian mixture model. We only considered the most frequent strategy (see Figure 8.5C) for both shapes and discarded trials (10.6 %) deviating from the chosen strategy. However, our findings do not depend on the particular choice of strategy as shapes that revealed differences between the myopic observer and the planning observer (S3 and S4) did not elicit different strategies. The remaining 2313 trials were used for our analysis.

8.4.3 Behavioral and model results

For all shapes subjects showed higher detection performance in the condition with the long search interval (Figure 8.6A). The spatial distribution of detected and missed targets is shown in Figure 8.6B. The mean fixation location for each participant separately for all shapes and conditions is shown in Figure 8.4A. Also, fixation sequences for the best fit of the ideal observer (right column) and the ideal planner (center column) are depicted. Visual inspection suggests that the behavioral data is closer resembled by the results of the ideal planner. To test whether eye movements were planned, we compared the first fixation location in the short condition to the first fixation location in the long condition for all shapes. If subjects were capable of performing planning, we expected a difference in the first fixation location for Shape S3 and S4. We used Hotelling's T-test to compare the bivariate landing positions of the first saccade between the two search intervals (Table 8.1). Indeed, mean target locations for the first saccade were different in Shape S3 and S4. No significant differences, however, were found in shapes S1

Table 8.1: Descriptive statistics of landing positions (in screen pixel) for all shapes as well as inferential statistics (two-sided) for the comparison across conditions (long and short).

	N	Centroid		Covariance			Hotelling's T	p
		$\bar{x}$	$\bar{y}$	σ^2_{xx}	σ^2_{yy}	σ^2_{xy}		
S1								
Short	16	540	568	862	686	191	0.744	.485
Long		535	554	1 018	1040	513		
S2								
Short	16	669	622	267	1270	271	2.66	.086
Long		669	650	223	1574	210		
S3								
Short	16	766	554	1676	3991	-2136	17.43	1.07×10^{-5}*
Long		705	675	550	3207	-492		
S4								
Short	16	748	608	3903	815	-106	7.89	.002*
Long		677	631	2094	1066	-223		

and S2. Unlike the ideal observer, our ideal planner well predicted this behavior. In addition, the direction of the spatial difference of the first fixation location between the search interval conditions followed the course suggested by our ideal planner (Figure 8.4B).

8.4.4 Bounded actor extensions

We extended both the ideal observer as well as the ideal planner to yield a more realistic model for human visual search behavior, i.e., a bounded actor. We added additive costs for longer saccade amplitude, as they lead to longer scanpath duration (Baloh et al., 1975) and higher endpoint variability (van Beers, 2007), which humans have been shown to minimize (Harris & Wolpert, 1998), used foveated versions of the shapes to account for the decline of visual acuity in peripheral vision (Geisler & Perry, 1998), and accounted for the often reported fact, that human saccades undershoot their target (Harris, 1995; Gillen et al., 2013). To obtain a quantitative evaluation of the computational models, we employed model selection using the Bayesian information criterion (BIC). The two free parameters in the models, i.e. the magnitude of additive costs for saccade length and the magnitude of the undershot, were estimated using Maximum Likelihood with bivariate Gaussian error terms on subjects' empirical data. We also estimated the covariance matrices for the models' predictions and the behavioral gaze data to compute the BIC for each model.

Figure 8.4C shows the difference in BIC of all models compared to the best model. The lower bound was derived by computing the mean fixation locations directly from the data for each of the four shapes as well as for each of the three fixation locations. The difference in BIC values between two models is an approximation for the log-Bayes factor and a difference $\Delta BIC > 4.6$ is considered to be decisive Kass and Raftery, 1995. Results clearly favor the planning observer over the myopic observer ($\Delta BIC = 139$). Crucially, the planning observer without any parameter fitting still provided a better description of our human data than the myopic observer with all extensions ($\Delta BIC = 59$). Further, costs for saccade amplitudes and foveation did not only improve our model fit for the planning observer but were also favored by model selection, suggesting that they are needed for better describing the eye movement data

Figure 8.5: (A) Theoretical detection performance for the shapes used in the experiment. Contour lines depict the percentage of detected events dependant on the fixation location. Two shapes show unimodal detection performance maps, two have multiple modes. (B) Raw data for the two multimodal shapes. Gaussian mixture models were fitted to our subjects raw data. Model comparison revealed qualitatively different strategies. First fixations were distributed over the different modes for multimodal shapes. (C) Proportion of trials corresponding to the different strategies. Panels are aligned to (B). Color represent the choice of strategy according to the mixture model shown in (B).

in our experiment. For the saccadic undershot model comparison was less decisive but still in favor of the full model ($\Delta BIC = 3$ between planning observer with all extensions and planning observer without undershot).

Parameter estimates for the saccadic undershot were similar for the myopic observer (2.9 %) and the planning observer (3.2 %). The influence of the costs for longer saccades was higher for the myopic observer (0.69 DP / Deg) compared to the planning observer (0.34 DP/Deg). The unit of the costs is detection performance (DP in %) per degree (Deg) and states, how much performance subjects were willing to give up to shorten saccade amplitudes by one visual degree. It is important to note, that both factors, costs and saccadic undershot, represent distinct computational concepts. The influence of the costs does not depend on the amplitude of the saccade directly, but on the reward structure of different potential landing locations. Hence, for two different shapes the same costs can have very different effect on where to target gaze. On the other hand, the undershot is relative to and only depends on the amplitude of the saccade and does not depend on the reward structure and therefore the shape. We also estimated the

A

B

Detected targets   Undetected targets

Figure 8.6: (A) Human detection performance for all shapes and conditions separately. (B) Error distribution for all four shapes and both conditions. Green crosses depict locations where a target was successfully detected. Red cross depict locations of targets that were missed. The upper panel shows the error distribution for the short condition. The errors for the long condition are shown in the lower panel.

radius of the circular gaze contingent search shape centered at the current fixation. Parameter estimation yielded values very close to the true radius and did not improve model quality for neither the planning observer nor the myopic observer.

8.5 Discussion

It has been unclear whether sequences of human eye movements are planned ahead of time. Prior studies indicate that multiple saccadic targets are jointly prepared as a scanpath and that cueing new targets during execution of eye movements results in longer execution times (Zingale & Kowler, 1987; Baldauf & Deubel, 2008a; De Vries et al., 2014). However, to our knowledge, there has been no experimental evidence that eye movements are chosen by considering more than one step ahead into the future. Instead, the ideal-observer paradigm, that models human eye movements as sequential Bayesian decisions has been the predominant approach.

In our study, we tested whether the implicit assumptions that accompany the ideal observer are justified. Therefore, we contrasted the ideal observer with the more general ideal planner that was formalized as a Markov Decision Process (Sutton & Barto, 1998) with partially observable states (Kaelbling et al., 1998). We formalized policies for the ideal observer, only considering the immediate reward for action selection, and for the ideal planner, which also considers future

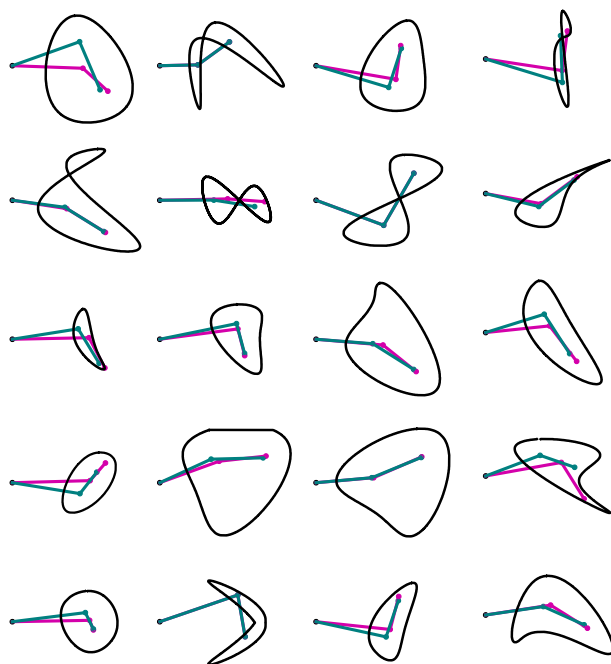Figure 8.7: Shapes generated by the algorithm with random initialization. Green and pink lines correspond to the optimal two fixation search strategy of the ideal planner and the ideal observer, respectively.

rewards. Next, we derived the specific circumstances under which the models produce different policies. Ultimately, we used these insights to manufacture stimuli that maximized the behavioral differences elicited by the different cognitive strategies and also obtained stimuli that show very similar strategies. Thus, the experiment was highly suitable for examining which cognitive strategy was adopted by our subjects.

We developed a visual search task where we expected different behavioral sequences depending on the cognitive strategy of our subjects. In particular, we investigated whether subjects adjust their scanpath during visual search depending on the duration of the search interval. Therefore, we controlled the length of the saccadic sequence. The short search interval allowed subjects to execute a single saccade, while in the long search interval subjects were able to fixate two locations.

Our results suggest that eye movements are indeed planned. Subjects' scanpath was very well predicted by the ideal planner while showing severe deviations from the scanpath proposed by the ideal observer. Crucially, this was the case even if the sequence required planning. We found fixation locations to be different depending on the duration of the search interval. This difference is only expected under the ideal planner and cannot be explained by the ideal observer. Finally, model comparison favored the ideal planner and its extensions over the ideal observer by a large margin. Furthermore, extending our ideal planner model to a bounded planner, we found evidence that subjects traded off task performance and saccade amplitude. Including additive costs for saccades with great amplitude into the ideal planner and accounting for saccadic undershot was best capable of explaining our data further.

Finding and executing near optimal gaze sequences is crucial for many extended sequential every-day tasks (Hayhoe & Ballard, 2005; Land & Hayhoe, 2001). The capability of humans to plan behavioral sequences gives further insights into why we can solve so many tasks with ease, which are extremely difficult from a computational perspective. In many visuomotor tasks

coordinated action sequences are needed rather than single isolated actions (Hayhoe, 2017). This leads to delayed rewards and thus a complex policy is required rather than an action that directly maximizes the performance after the next single gaze switch. Additionally, our findings have implications for future models of human eye movements. While numerous influential past models have considered planning (Najemnik & Geisler, 2005; Peterson & Eckstein, 2012; Hoppe & Rothkopf, 2016; Navalpakkam et al., 2010), our results indicate that in the case of visual search humans are capable of including future states into the selection of a suitable scan path.

The broader significance of the present results beyond the understanding of eye movements lies in the fact that human behavior in our experiment was best described by a computational model of a bounded probabilistic planner including perceptual uncertainty. In this framework, sensory measurements and goal-directed actions are inseparably intertwined (Gottlieb, 2012; Yang, Wolpert, & Lengyel, 2016). So far, the predominant approach to probabilistic models in perception has been the ideal observer (Geisler, 2003, 2011), which can be formalized in the Bayesian framework (Knill & Richards, 1996; Kersten et al., 2004) as inferring latent causes in the environment giving rise to sensory observations. Models of eye movements selection have so far used ideal observers (Najemnik & Geisler, 2005; Peterson & Eckstein, 2012; Hoppe & Rothkopf, 2016) without planning.

A limitation of the current study is that it does not disambiguate between open-loop and closed-loop planning (Russell, Norvig, & Davis, 2010). The distinction between these two types of planning lies in the way future observations are utilized within the planning process. While open-loop algorithms plan a sequence of actions but disregard the outcome of future observations, closed-loop algorithms are much more sophisticated by taking all possible future observations after each action in the entire sequence into account within the planning process. As such, closed-loop planning is even more demanding computationally than open-loop planning. The current experiments cannot disambiguate whether human behavior is better explained by either of these two planning algorithms, because for the second fixation in the long search interval condition the belief after the first fixation only depends on whether the target was found. Thus, subjects terminating the search in the long search interval condition after finding the target after the first fixation may be the only support for closed-loop control in our experiments. Future work will need to address, which of these two types of planning better describes human gaze selection.

Probabilistic, Bayesian formulations of optimality in perceptual tasks (Ernst & Banks, 2002; Körding & Wolpert, 2004), cognitive tasks (Oaksford & Chater, 2007; Gershman, Horvitz, & Tenenbaum, 2015), reasoning (Tenenbaum, Griffiths, & Kemp, 2006), motor control (Todorov & Jordan, 2002), learning (Daw, Niv, & Dayan, 2005), and planning (Huys et al., 2015) have led to a better understanding of human behavior and the quest to unravel, how the brain could implement these computations (Ma et al., 2006; Fiser et al., 2010; Sanborn & Chater, 2017), which are known in general to be intractable (Kwisthout & Van Rooij, 2013). Our results extend the current understanding by demonstrating that planning under perceptual uncertainty is also part of the repertoire of human visual behavior and open up the possibility to understand recent neurophysiological results (Foley, Kelly, Mhatre, Lopes, & Gottlieb, 2017) within the planning under uncertainty framework.

General discussion

9.1 Summary of the findings

The present thesis investigated how eye movements are scheduled in dynamic environments. To this end, we developed a family of tasks that facilitated us to probe our subjects' visual system while having full knowledge of the temporal statistics and the reward structure of our experiments. Crucially, these two components are usually unknown when considering natural behavior. This enabled us to develop normative computational models for temporally complex visual behavior. Using these models, we were able to gain new insights into several aspects of the human visual information processing.

   We found new insights into the mechanisms underlying the control of blinks. Using an event detection task together with known temporal statistics we were able to explain the spatial as well as the temporal properties of human blinking behavior. Using the information about the task properties, we predicted when blinks should occur and also the distribution of times between two consecutive blinks, the interblink intervals. Importantly, this yielded a computational explanation for the characteristic interblink interval distributions proposed by Ponder and Kennedy (1927) and repeatedly found by various investigations. Surprisingly, our results suggest, that costs for blink suppression do not increase with the duration since the last blink. Instead, they can be sufficiently explained by a memoryless stochastic process.

   Further, our findings suggest that temporal eye movements are guided by internal representations of the stochastic quantities in our environment. Crucially, we were able to understand the processes involved in learning these quantities using computational approaches. Further, our results suggest that eye movement strategies take into account the temporal variability in fixation durations when deciding for a particular scheduling of eye movements. It has been shown for pointing (Trommershäuser et al., 2006) that humans perform nearly optimal despite the variability of actions. Our data and computational model results suggest that this is also the case for fixation durations and saccades. We showed that the scalar law of biological timing, a well-documented phenomenon in psychophysical tasks, also plays a part in action selection for eye movements. Also, we presented work on how physiological limits bound the visual system, e.g., the processing time, that is needed to perform the computational steps in order to decide for a subsequent action.

   Having continuous access to the environmental statistics during a dynamic task provides an opportunity to study the variability of eye movements. We created stimuli that appeared different but shared the same statistical structure in dimensions relevant to the task. Using this approach, we were able to probe subjects' with same stimulus sequences multiple times while avoiding learning and training effects. We found that variability of eye movements is related to the expected loss in the task. This means, that if the stakes are high, subjects show similar

eye movement behavior while variability increases if the expected loss is smaller. Again, we substantiated our results using a computational model that closely reflected the eye movements of our subjects.

Finally, normative computational models understand action selection in terms of reward maximization. In a simple two-alternative forced-choice paradigm, there is only a single action and therefore the action with the maximum reward should be taken. This is not the case when considering sequences of eye movements, as the accumulated reward and therefore the optimal action sequence depends on the length of the sequence. However, state-of-the-art models for eye movements do not take this into account (Najemnik & Geisler, 2005; Navalpakkam et al., 2010; Schütz et al., 2012; Peterson & Eckstein, 2012; Hoppe & Rothkopf, 2016; Yang, Lengyel, & Wolpert, 2016). We contrasted two different models, an ideal observer taking the action with the maximal immediate reward, and an ideal planner taking the action that maximizes the accumulated total reward. Taking future actions into account for action selection leads to higher rewards but also to higher computational costs. While humans have shown to be able to trade off how far ahead they plan (Krueger, Lieder, & Griffiths, 2017), and hand and eye movements can be affected by future rewards in foraging tasks (Diamond, Wolpert, & Flanagan, 2017), there has been no normative evidence that eye movements are scheduled using planning. We derived a visual search task with search shapes designed to maximally discriminate between the two models. Our results showed that our subjects' scanpath was closely reflected by the ideal planner while deviating severely from the ideal observer. This suggests that the visual system is capable of planning ahead at least two eye movements. Our study also shows, that whether or not future rewards need to be accounted for highly depends on the specific choice of stimuli.

9.2 Influence of costs on action selection

All the empirical data in the present thesis showed that visual behavior can be explained as a trade-off between multiple sources of costs and rewards. We used inverse reinforcement learning, i.e., we inferred the goal structure from action sequences, to show the mechanics of this trade-off. In our experiments, the task-related reward that could be earned through suitable actions was controlled and known. We then explored the physiological and motivational costs that influence action selection.

Outside of the lab, we constantly have to trade-off multiple different costs and benefits. Various studies have investigated the trade off between eye movements and memory. The results show, that whether subjects acquire information by targeting their sensory system towards the region of interest compared to retrieving the information from memory depends on the cost for storing the information in memory (Droll & Hayhoe, 2007) as well as the cost for making the explorative eye movement (Ballard et al., 1995). In a block sorting task, subjects more frequently relied on information stored in memory if the distance of the eye movement needed to gather information was increased (Ballard et al., 1995). On the other hand, Droll and Hayhoe (2007) increased the cost of solving a sorting task using memory by increasing the number of features that needed to be remembered. As a result, they observed an increase in the use of strategies relying on online information. Both studies show that by manipulating either of the dimensions of the reward function (memory and eye movements) behavioral changes can be observed that reflect an internal trade-off. While there exists evidence pointing to the influence of costs on behavior, computational accounts that quantitatively show how costs affect action selection are scarce.

Here, we provided quantitative evidence for the trade-off between task-related costs, benefits

and physiological costs. We reported costs for the frequency of saccades (Chapter 6), the suppression of blinks (Chapter 5), the amplitude of a saccade (Chapter 8), and simply how closely a subject pays attention to the stimulus (Chapter 7). We are still a long way from deciphering the complex interplay of the different sources of rewards and costs that produce complex everyday behavior in natural contexts. Future work should focus on quantifying the behavioral costs and trade-offs for different actions. One possibility could be to measure quantities related to energy consumption for a movement.

For example, the physiology of the eye affects blinking behavior and various measures for aberrations have been proposed, e.g., the tear breakup time (Al-Abdulmunem, 1999). These physiological measures give insight into the cost structure for different actions. In their seminal work on cue combination, Ernst and Banks (2002) showed, that humans optimally combine different sensory measurements and weigh them according to their variability to yield a single estimator. Crucially, they measured the uncertainty in both modalities - audio and visual - separately. In the same fashion, future work should attempt to unfold the cost structure using separate diagnostics and plug these measurements into our computational models.

9.3 Temporal dynamics of human visual behavior

While there exists a rich literature on human time perception (James, 1890; Allan, 1979; Grondin, 2010; Mauk & Buonomano, 2004; Brown & Merchant, 2007) little research has applied this principle to behavioral data that go beyond time estimation tasks. There are studies from the field of engineering that deal with action selection in dynamic scenarios (Levinson et al., 2011), however, technical systems do not usually share the same characteristics as the human visual system. In a series of experiments, we showed that humans are capable of performing nearly optimal despite the temporal uncertainty that is omnipresent in natural environments. In particular, Weber's law for time estimation plays a key role in scheduling actions with critical timing. Humans even accounted for it during learning the temporal statistics from experience. We extended past research showing that humans are near optimal in spatial tasks (Najemnik & Geisler, 2005; Peterson & Eckstein, 2012; Navalpakkam et al., 2010; Chukoskie et al., 2013) to the temporal domain. In our experiments, subjects performed close to computational models based on optimal decision making, if uncertainties and costs were considered. Future work should attempt to unite the advances of models in both domains, spatial and temporal, to make predictions for eye movements in more complex situations.

9.4 Alternative models for eye movements

In the present thesis, we proposed models based on optimal decision making using Bayesian decision theory and generalized this to partially observable Markov decision processes. We found that the developed models were able to relate the behavioral data to biologically plausible concepts on a computational level. However, various other methods have been proposed and applied to predicting eye movements in the past (see Chapter 4). In their influential work on Bayesian surprise, Itti and Baldi (2006) suggest that eye movements are targeted at locations with maximum gain in information. However, models using information theoretic criteria instead of maximizing a reward only lead to optimal behavior if every bit of information that a subject can harvest at any time is equally valuable in terms of task performance. This property is violated, for example, in the experiment presented in Chapter 7. The dots at all three of the presented regions moved according to a random walk with the same step size. Hence, when not

being fixated, the uncertainty about the current location of the dot grows with the same rate for all three locations. As a consequence, the amount of information associated with pointing gaze to a location does not depend on the last seen location of the dot but depends solely on the time since the last fixation. However, in terms of reward, the information of a dot last seen closer to the boundary is much more valuable than information of a dot closer to the center.

Saliency models comprise a prominent family of computational models relating low-level visual features to the probability of attracting gaze. However, they do not incorporate temporal characteristics of behavior, hence they can not explain any of the experiments described in the present thesis. Saliency is a spatial concept and, crucially, does not change over time. This is clearly not the case in dynamic environments, where a single location can attract gaze in one moment and stop to do so in the next.

9.5 Connection to natural stimuli

We used stimuli designed explicitly for studying temporal environments without the spatial complexity present in real-world tasks. This approach was motivated in Chapter 4 as it facilitates developing computational models for temporal eye movements. How do our experimental designs relate to real-world tasks and how can the developed models be applied to everyday problems? First, while the spatial properties of the stimuli were kept as simple as possible, the temporal statistics of the stimuli followed complex dynamics. The different statistics used in the experiments are connected to real-world dynamics and cover a wide range of everyday activities. The events' durations used in our experiments ranged from 50 ms (blinking) to 1.5 s (TED). Many everyday tasks have events with similar durations between a couple milliseconds and a few seconds including facial expressions during conversation (Yan et al., 2013), traffic light changes during locomotion, or speed changes of the car ahead during driving. Second, we covered two different monitoring tasks with clear connections to everyday activities. The TED task (Chapter 6) where observing that the event has not yet occurred does not give any information about the likelihood of the event in the future. This, for example, is similar to waiting for a traffic light turning or a bus coming. Second, in the monitoring task in Chapter 7 observing the location of a dot, even when the event has not occurred yet, changes the probability of the event at that location, i.e., the probability becomes very low if the dot is near the center. The same is true for monitoring the vehicle in front of us in driving.

All tasks required minimal practice (5 minutes for the TED task, 1 minute for the planning task, no training for the remaining two experiments) despite the reduced spatial complexity. This indicates that participants resorted to eye movement strategies they also use in daily tasks with similar temporal dynamics instead of learning eye movement patterns specifically adapted to the experimental settings. Further evidence comes from the fact that the visual behavior during the experiments shared common statistical properties when compared to natural behavior. The temporal statistics of the behavioral data were similar to previous findings, i.e., the blinking rates (similar to Garcia et al., 2011), the distribution of interblink intervals (similar to Ponder & Kennedy, 1927), even if they contained more natural stimuli, i.e., the distribution of fixation durations (similar to Dorr et al., 2010).

9.6 Application in technical systems

Results about human eye movement patterns are essential to advanced systems that infer the attentional state of drivers (see Dong, Hu, Uchimura, & Murayama, 2011, for an overview).

According to the World Health Organization (WHO) some 1.25 million people die in traffic each year (World Health Organization, 2015). Hence, improvement of the systems preventing road accidents by monitoring the state of the driver has great benefits. Many of the indicators of inattention and fatigue reported in the literature use static thresholds to determine the state of the driver (see the references in Dong et al., 2011). Besides static thresholds, machine learning algorithms (Liu, Yang, Huang, Yeo, & Lin, 2016) as well as deep neural networks (Streiffer, Raghavendra, Benson, & Srivatsa, 2017) have been proposed to infer whether a driver is distracted. However, all reviewed systems focus solely on data about the driver while neglecting both the environmental dynamics as well as the current task.

The percentage of eyelid closure (PERCLOS) is a measure for drowsiness comprising three metrics: the proportion of time the eyes are closed at least 70 and 80 percent, respectively, and the mean square percentage (Dinges & Grace, 1998). It is frequently applied to the detection of fatigue as higher PERCLOS scores have found to be an indicator for fatigue (Sommer & Golz, 2010). However, studies have suggested that driver fatigue can be detected better when using EEG-based systems and not reducing eyelid movements to the PERCLOS measures, our results reveal two further weaknesses of this commonly used metric. First, our results show that blinking rates change according to the current demands. Second, interindividual differences in the physiology of the eyes lead to great blinking variability across subjects. Both aspects have shown to heavily influence blinking behavior and therefore directly affect the percentage of eyelid closure.

Our results emphasize, that task-related properties greatly influence where people target their gaze. This was the case both for the event detection task (Chapter 6) as well as for the temporal monitoring task (Chapter 7). Both psychophysical approaches are related to visual tasks during driving: keeping an eye on the car in front as well as checking the mirror for passing vehicles. Models that include task-relevant effects have been proposed in the past (Wickens, Helleberg, Goh, Xu, & J. Horrey, 2001; Feuerstack & Wortelen, 2017). According to the Wickens et al. (2001) the probability that a region is attended through eye movements depends on the respective saliency, the effort of making a saccade to the location, the expectancy for receiving new information at the location, and the value.

The Human Efficiency Evaluator (Feuerstack & Wortelen, 2017) is based on the Saliency Effort Expectancy Value (SEEV) model. SEEV models have been used to model in various scenarios related but not limited to driving (see Feuerstack & Wortelen, 2017, and the references therein). The parameters for the models are estimated using lowest ordinal heuristic. However, they require human judgments in order to estimate the free parameters.

While our results link human eye movement behavior to task-related rewards, our data show great interindividual variability. In our experiments, both for blinking behavior as well as for monitoring, we have shown that differences in physiological costs across participants lead to severe behavioral differences. This indicates, that interindividual differences need to be taken into account to correctly infer the current state from visual behavior. In addition, state of the art computational models could benefit from the methodology described in the present thesis as all reported characteristics of dynamic gaze behavior ground on biologically plausible mechanisms, i.e., the scalar law of biological timing and physiological cost for blink suppression, which have been reported repeatedly in the literature.

CHAPTER 10

Conclusion

The study of human behavior lags behind when it comes to general explanations for how humans behave in their natural environment compared to other disciplines like physics or chemistry, where detailed mathematical descriptions for complex phenomena exist. In order to understand human behavior in natural settings, we must understand the statistics of the respective environment. Imagine a ball rolling through a hilly landscape. The ball's trajectory can be described by application of Newton's mechanics which is essentially that $F = ma$. However, without knowing the topology of the landscape, the observed trajectory is of great complexity. The present thesis has shown, that visual behavior highly depends on the environmental regularities and can only be explained if those regularities are taken into account. In a series of experiments we found a clear connection between behavioral action selection and environmental statistics. This observation is crucial as it makes finding general laws for the generation of behavior impossible, unless they comprise the complexity of natural statistics. However, it does not contradict the existence of psychological laws. Instead, we might not be in possession of the right quantities for them so far. For example, in our experiment investigating eye movement variability we found a simple linear relationship between variability and task properties. However, without access to the environmental statistics and the computational theory this relationship would have been untraceable or would seem of great complexity. Therefore, it might be that complex phenomena retain a law-like structure once we have more knowledge of the complex environmental statistics and task structures that humans are designed to solve. Recent advances in intelligent algorithms as well as easy access to large quantities of environmental data open a chance to discover more about these problems.

References

Al-Abdulmunem, M. (1999). Relation between tear breakup time and spontaneous blink rate. *International Contact Lens Clinic*, *26*(5), 117–120.

Allan, L. G. (1979). The perception of time. *Perception & Psychophysics*, *26*(5), 340–354.

Archibald, S. (2008). Ways of seeing. *Cabinet*, *9*(2), 550–570.

Baldauf, D. & Deubel, H. (2008a). Properties of attentional selection during the preparation of sequential saccades. *Experimental Brain Research*, *184*(3), 411–425.

Baldauf, D. & Deubel, H. (2008b). Visual attention during the preparation of bimanual movements. *Vision Research*, *48*(4), 549–563.

Ballard, D. H., Hayhoe, M. M., & Pelz, J. (1995). Memory representations in natural tasks. *Journal of Cognitive Neuroscience*, *7*(1), 68–82.

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

Battaglia, P. W., Kersten, D., & Schrater, P. R. (2011). How haptic size sensations improve distance perception. *PLoS computational biology*, *7*(6), e1002080.

Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E., & Pouget, A. (2012). Not noisy, just wrong: The role of suboptimal inference in behavioral variability. *Neuron*, *74*(1), 30–39.

Bellman, R. (1957). A markovian decision process. *Journal of Mathematics and Mechanics*, 679–684.

Berman, B. D., Horovitz, S. G., Morel, B., & Hallett, M. (2012). Neural correlates of blink suppression and the buildup of a natural bodily urge. *NeuroImage*, *59*(2), 1441–1450.

Berret, B., Chiovetto, E., Nori, F., & Pozzo, T. (2011). Evidence for composite cost functions in arm movement planning: An inverse optimal control approach. *PLoS computational biology*, *7*(10), e1002183.

Biro, P. A. & Stamps, J. A. (2010). Do consistent individual differences in metabolic rate promote consistent individual differences in behavior? *Trends in Ecology & Evolution*, *25*(11), 653–659.

Borji, A. (2012). Boosting bottom-up and top-down visual features for saliency estimation. In *Computer vision and pattern recognition (cvpr), 2012 ieee conference on* (pp. 438–445). IEEE.

Borji, A., Sihite, D. N., & Itti, L. (2013). What stands out in a scene? a study of human explicit saliency judgment. *Vision Research*, *91*(0), 62–77.

Brainard, D. H. & Freeman, W. T. (1997). Bayesian color constancy. *JOSA A*, *14*(7), 1393–1411.

Bressman, S. B., Cassetta, E., & Bentivoglio, A. R. (1997). Analysis of blink rate patterns in normal subjects. *Movement*, *12*(6), 1–7.

Bristow, D., Haynes, J. D., Sylvester, R., Frith, C. D., & Rees, G. (2005). Blinking suppresses the neural response to unchanging retinal stimulation. *Current Biology*, *15*(14), 1296–1300.

Brown, S. W. & Merchant, S. M. (2007). Processing resources in timing and sequencing tasks. *Perception & psychophysics*, *69*(3), 439–449.

Bruce, N. D. & Tsotsos, J. K. (2009). Saliency, attention, and visual search: An information theoretic approach. *Journal of vision*, *9*(3), 5–5.

Burnham, K. P. & Anderson, D. R. (2004). Multimodel inference. *Sociological Methods & Research*, *33*(2), 261–304.

Castelhano, M. S., Mack, M. L., & Henderson, J. M. (2009). Viewing task influences eye movement control during active scene perception. *Journal of vision*, *9*(3), 6–6.

Chang, C.-J. & Jazayeri, M. (2018). Integration of speed and time for estimating time to contact. *Proceedings of the National Academy of Sciences*, 201713316.

Chukoskie, L., Snider, J., Mozer, M. C., Krauzlis, R. J., & Sejnowski, T. J. (2013). Learning where to look for a hidden target. *Proceedings of the National Academy of Sciences USA*. *110*(Supplement 2), 10438–10445.

Clarke, A. D. & Hunt, A. R. (2015). Failure of intuition when choosing whether to invest in a single goal or split resources between two goals. *Psychological science.*

Cruz, A. A., Garcia, D. M., Pinto, C. T., & Cechetti, S. P. (2011). Spontaneous eyeblink activity. *The ocular surface*, *9*(1), 29–41.

Darwin, C. (1859). *On the origin of species by means of natural selection.* London: Murray.

Daw, N. D. & Doya, K. (2006). The computational neurobiology of learning and reward. *Current opinion in neurobiology*, *16*(2), 199–204.

Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature neuroscience*, *8*(12), 1704–1711.

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

De Vries, J. P., Hooge, I. T., & Verstraten, F. A. (2014). Saccades toward the target are planned as sequences rather than as single steps. *Psychological science*, *25*(1), 215–223.

DeGroot, M. H. (2005). *Optimal statistical decisions.* John Wiley & Sons.

Deubel, H., Shimojo, S., & Paprotta, I. (1997). The preparation of goal-directed movements requires selective visual attention: Evidence from the line-motion illusion. *Perception*, *26*(1), 124–124.

Diamond, J. S., Wolpert, D. M., & Flanagan, J. R. (2017). Rapid target foraging with reach or gaze: The hand looks further ahead than the eye. *PLOS Computational Biology*, *13*(7), e1005504.

Diaz, G., Cooper, J., Rothkopf, C., & Hayhoe, M. (2013). Saccades to future ball location reveal memory-based prediction in a virtual-reality interception task. *Journal of Vision*, *13*(1), 20.

Dimitrakakis, C. & Rothkopf, C. A. (2011). Bayesian multitask inverse reinforcement learning. In *European workshop on reinforcement learning* (pp. 273–284). Springer.

Dinges, D. F. & Grace, R. (1998). Perclos: A valid psychophysiological measure of alertness as assessed by psychomotor vigilance. *US Department of Transportation, Federal Highway Administration, Publication Number FHWA-MCRT-98-006.*

Dinstein, I., Heeger, D. J., & Behrmann, M. (2015). Neural variability: Friend or foe? *Trends in cognitive sciences*, *19*(6), 322–328.

Dong, Y., Hu, Z., Uchimura, K., & Murayama, N. (2011). Driver Inattention Monitoring System for Intelligent Vehicles: A Review. *IEEE Transactions on Intelligent Transportation Systems*, *12*(2), 596–614.

Dorr, M., Martinetz, T., Gegenfurtner, K. R., & Barth, E. (2010). Variability of eye movements when viewing dynamic natural scenes. *Journal of vision*, *10*(10), 28.

Dorr, M., Vig, E., & Barth, E. (2012). Eye movement prediction and variability on natural video data sets. *Visual cognition*, *20*(4-5), 495–514.

Doughty, M. J. (2002). Further assessment of gender- and blink pattern-related differences in the spontaneous eyeblink activity in primary gaze in young adult humans. *Optometry and vision science : official publication of the American Academy of Optometry, 79*(7), 439–447.

Doughty, M. J. (2001). Consideration of three types of spontaneous eyeblink activity in normal humans: during reading and video display terminal use, in primary gaze, and while in conversation. *Optometry and vision science : official publication of the American Academy of Optometry, 78*(October), 712–725.

Droll, J. A. & Hayhoe, M. M. (2007). Trade-offs between gaze and working memory use. *Journal of Experimental Psychology: Human Perception and Performance, 33*(6), 1352.

Eagleman, D. M. (2008). Human time perception and its illusions. *Current Opinion in Neurobiology, 18*(2), 131–136.

Eisler, H., Eisler, A. D., & Hellström, Å. (2008). Psychophysical issues in the study of time perception. *Psychology of time*, 75–109.

Ernst, M. O. & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415*(6870), 429–433.

Evinger, C., Manning, K. a., Pellegrini, J. J., Basso, M. a., Powers, a. S., & Sibony, P. a. (1994). Not looking while leaping: the linkage of blinking and saccadic gaze shifts. *Experimental brain research. Experimentelle Hirnforschung. Expérimentation cérébrale, 100*(2), 337–44.

Fabre-Thorpe, M., Delorme, A., Marlot, C., & Thorpe, S. (2001). A limit to the speed of processing in ultra-rapid visual categorization of novel natural scenes. *Journal of Cognitive Neuroscience, 13*(2), 171–180.

Faisal, A. A., Selen, L. P., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature reviews neuroscience, 9*(4), 292.

Farrell, S., Ludwig, C. J., Ellis, L. A., & Gilchrist, I. D. (2010). Influence of environmental statistics on inhibition of saccadic return. *Proceedings of the National Academy of Sciences, 107*(2), 929–934.

Feuerstack, S. & Wortelen, B. (2017). A model-driven tool for getting insights into car drivers' monitoring behavior. In *Intelligent Vehicles Symposium (IV), 2017 IEEE* (pp. 861–868). IEEE.

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

Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: From behavior to neural representations. *Trends in cognitive sciences, 14*(3), 119–130.

Flanagan, J. R. & Johansson, R. S. (2003). Action plans used in action observation. *Nature, 424*(6950), 769–771.

Fogarty, C. & Stern, J. A. (1989). Eye movements and blinks: Their relationship to higher cognitive processes. *International Journal of Psychophysiology, 8*(1), 35–42.

Foley, N. C., Kelly, S. P., Mhatre, H., Lopes, M., & Gottlieb, J. (2017). Parietal neurons encode expected gains in instrumental information. *Proceedings of the National Academy of Sciences, 114*(16), E3315–E3323.

Foulsham, T. (2015). Eye movements and their functions in everyday tasks. *Eye, 29*(2), 196–199.

Franchak, J. M., Heeger, D. J., Hasson, U., & Adolph, K. E. (2016). Free viewing gaze behavior in infants and adults. *Infancy, 21*(3), 262–287.

Friston, K., Adams, R., Perrinet, L., & Breakspear, M. (2012). Perceptions as hypotheses: Saccades as experiments. *Frontiers in psychology, 3*, 151.

Gandhi, N. J. (2012). Interactions between gaze-evoked blinks and gaze shifts in monkeys. *Experimental Brain Research*, *216*(3), 321–339.

Garcia, D., Pinto, C. T., Barbosa, J. C., & Cruz, A. A. V. (2011). Spontaneous interblink time distributions in patients with Graves' orbitopathy and normal subjects. *Investigative Ophthalmology and Visual Science*, *52*(6), 3419–3424.

Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychological Review*, *96*(2), 267.

Geisler, W. S. (2003). Ideal observer analysis. *The visual neurosciences*, *10*(7), 12–12.

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

Geisler, W. S. & Perry, J. S. (1998). Real-time foveated multiresolution system for low-bandwidth video communication. In *Human vision and electronic imaging* (Vol. 3299, pp. 294–305).

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian data analysis*. CRC press Boca Raton, FL.

Gershman, S. J. (2016). Empirical priors for reinforcement learning models. *Journal of Mathematical Psychology*, *71*, 1–6.

Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, *349*(6245), 273–278.

Gibbon, J. (1977). Scalar expectancy theory and weber's law in animal timing. *Psychological Review*, *84*(3), 279.

Gilbert, C. D. & Li, W. (2013). Top-down influences on visual processing. *Nature Reviews Neuroscience*, *14*(5), 350.

Gillen, C., Weiler, J., & Heath, M. (2013). Stimulus-driven saccades are characterized by an invariant undershooting bias: No evidence for a range effect. *Experimental Brain Research*, *230*(2), 165–174.

Godijn, R. & Theeuwes, J. (2003). Parallel allocation of attention prior to the execution of saccade sequences. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(5), 882–896.

Goldstein, R. B., Woods, R. L., & Peli, E. (2007). Where people look when watching movies: Do all viewers look at the same place? *Computers in biology and medicine*, *37*(7), 957–964.

Gottlieb, J. (2012). Attention, learning, and the value of information. *Neuron*, *76*(2), 281–295.

Gregory, R. L. (1952). Variations in blink rate during non-visual tasks. *Quarterly Journal of Experimental Psychology*, *4*(4), 165–169.

Grondin, S. (2010). Timing and time perception: A review of recent behavioral and neuroscience findings and theoretical directions. *Attention, Perception, & Psychophysics*, *72*(3), 561–582.

Gupta, S. & Kim, H.-W. (2010). Value-driven internet shopping: The mental accounting theory perspective. *Psychology & Marketing*, *27*(1), 13–35.

Hacisalihzade, S. S., Stark, L. W., & Allen, J. S. (1992). Visual perception and sequences of eye movement fixations: A stochastic modeling approach. *IEEE Transactions on systems, man, and cybernetics*, *22*(3), 474–481.

Hall, C., Figueroa, A., Fernhall, B., & Kanaley, J. A. (2004). Energy expenditure of walking and running: Comparison with prediction equations. *Medicine and science in sports and exercise*, *36*, 2128–2134.

Halvor Teigen, K. (2002). One hundred years of laws in psychology. *The American journal of psychology*.

Harris, C. M. (1995). Does saccadic undershoot minimize saccadic flight-time? a monte-carlo study. *Vision research*, *35*(5), 691–701.

Harris, C. M. & Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, *394*(6695), 780.

Hausknecht, M. & Stone, P. (2015). Deep recurrent q-learning for partially observable mdps. *CoRR, abs/1507.06527*, *7*(1).

Hayes, D. J. & Greenshaw, A. J. (2011). Neuroscience and Biobehavioral Reviews. *Biobehavioral Reviews*, *35*, 1419–1449.

Hayhoe, M. M. (2017). Vision and action. *Annual Review of Vision Science*, *3*(1), 389–413. PMID: 28715958.

Hayhoe, M. M., McKinney, T., Chajka, K., & Pelz, J. B. (2012). Predictive eye movements in natural vision. *Experimental Brain Research*, *217*(1), 125–136.

Hayhoe, M. & Ballard, D. (2005). Eye movements in natural behavior. *Trends in Cognitive Sciences*, *9*(4), 188–194.

Hayhoe, M. & Ballard, D. (2014). Modeling task control of eye movements. *Curr. Biol. 24*(13), 622–628.

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

Henderson, J. M., Shinkareva, S. V., Wang, J., Luke, S. G., & Olejarczyk, J. (2013). Predicting cognitive state from eye movements. *PloS one*, *8*(5), e64937.

Hergenhahn, B. R. & Henley, T. (2013). *An introduction to the history of psychology*. Cengage Learning.

Hoffman, J. E. & Subramaniam, B. (1995). The role of visual attention in saccadic eye movements. *Attention, Perception, & Psychophysics*, *57*(6), 787–795.

Hommel, B., Pratt, J., Colzato, L., & Godijn, R. (2001). Symbolic control of visual attention. *Psychological Science*, *12*(5), 360–365.

Hoppe, D., Helfmann, S., & Rothkopf, C. A. (2018). Humans quickly learn to blink strategically in response to environmental task demands. *Proceedings of the National Academy of Sciences*, 201714220.

Hoppe, D. & Rothkopf, C. A. (2019). Multi-step planning of eye movements in visual search. *Scientific reports*, *9*(1), 144.

Hoppe, D. & Rothkopf, C. A. (2016). Learning rational temporal eye movement strategies. *Proceedings of the National Academy of Sciences*, *113*(29), 8332–8337.

Hudson, T., Maloney, L., Landy, M., & Friston, K. J. (2008). Optimal compensation for temporal uncertainty in movement planning. *PLOS Computational Biology*, *4*(7), e1000130.

Huys, Q. J. M., Lally, N., Faulkner, P., Eshel, N., Seifritz, E., Gershman, S. J., . . . Roiser, J. P. (2015). Interplay of approximate planning strategies. *Proceedings of the National Academy of Sciences*, *112*(10), 3098–3103.

Irwin, D. E. (2011). Where does attention go when you blink? *Attention, Perception, & Psychophysics*, *73*(5), 1374–1384.

Itti, L. & Baldi, P. F. (2006). Bayesian surprise attracts human attention. In *Advances in neural information processing systems* (pp. 547–554).

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

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

Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on pattern analysis and machine intelligence*, *20*(11), 1254–1259.

Ivry, R. B. (1996). The representation of temporal information in perception and motor control. *Current opinion in neurobiology*, *6*(6), 851–857.

Ivry, R. B. & Schlerf, J. E. (2008). Dedicated and intrinsic models of time perception. *Trends in Cognitive Sciences*, *12*(7), 273–280.

James, W. (1890). The principles of psychology new york. *Holt and company*.

Jazayeri, M. & Shadlen, M. N. (2010). Temporal context calibrates interval timing. *Nature Neuroscience*, *13*(8), 1020–1026.

Johansson, R. S., Westling, G., Bäckström, A., & Flanagan, J. R. (2001). Eye–hand coordination in object manipulation. *Journal of Neuroscience*, *21*(17), 6917–6932.

Jongkees, B. J. & Colzato, L. S. (2016). Spontaneous eye blink rate as predictor of dopamine-related cognitive function—a review. *Neuroscience & Biobehavioral Reviews*, *71*, 58–82.

Jordan, M. I. (2003). An introduction to probabilistic graphical models. preparation.

Judd, T., Ehinger, K., Durand, F., & Torralba, A. (2009). Learning to predict where humans look. In *Computer vision, 2009 ieee 12th international conference on* (pp. 2106–2113). IEEE.

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

Kaminer, J., Powers, A. S., Horn, K. G., Hui, C., & Evinger, C. (2011). Characterizing the Spontaneous Blink Generator: An Animal Model. *31*(31), 11256–11267.

Kass, R. E. & Raftery, A. E. (1995). Bayes factors. *Journal of the american statistical association*, *90*(430), 773–795.

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annual Review of Psychology*, *55*, 271–304.

Khoramshahi, M., Shukla, A., Raffard, S., Bardy, B. G., & Billard, A. (2016). Role of gaze cues in interpersonal motor coordination: Towards higher affiliation in human-robot interaction. *PloS one*, *11*(6), e0156874.

Kleinke, C. L. (1986). Gaze and eye contact: A research review. *Psychological bulletin*, *100*(1), 78.

Knill, D. C. & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*(12), 712–719.

Knill, D. C. & Richards, W. (1996). *Perception as Bayesian inference*. Cambridge University Press.

Koh, S., Maeda, N., Hirohara, Y., Mihashi, T., Bessho, K., Hori, Y., . . . Tano, Y. (2008). Serial measurements of higher-order aberrations after blinking in patients with dry eye. *Investigative Ophthalmology and Visual Science*, *49*(1), 133–138.

Koh, S., Maeda, N., Hirohara, Y., Mihashi, T., & Ninomiya, S. (2006). Serial Measurements of Higher-Order Aberrations after Blinking in Normal Subjects. *47*(8), 3318–3324.

Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision making and the avoidance of cognitive demand. *Journal of Experimental Psychology: General*, *139*(4), 665.

Körding, K. P. & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*(6971), 244–247.

Körding, K. P. & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in cognitive sciences*, *10*(7), 319–326.

Krolak, A. & Strumillo, P. (2008). Vision-based eye blink monitoring system for human-computer interfacing. *2008 Conference on Human System Interaction, HSI 2008*, 994–998.

Krueger, P. M., Lieder, F., & Griffiths, T. L. (2017). Enhancing metacognitive reinforcement learning using reward structures and feedback. In *Proceedings of the the 39th annual meeting of the cognitive science society. austin, tx: Cognitive science society.*

Kuhn, G., Tatler, B. W., & Cole, G. G. (2009). You look where i look! Effect of gaze cues on overt and covert attention in misdirection. *Visual Cognition*, *17*(6-7), 925–944.

Kwisthout, J. & Van Rooij, I. (2013). Bridging the gap between theory and practice of approximate Bayesian inference. *Cognitive Systems Research*, *24*, 2–8.

Land, M. F. & McLeod, P. (2000). From eye movements to actions: How batsmen hit the ball. *Nature Neuroscience*, *3*, 1340–1345.

Land, M. F. & Nilsson, D. (2002). *Animal eyes.* Oxford University Press.

Land, M. & Hayhoe, M. (2001). In what ways do eye movements contribute to everyday activities? *Vision Research*, *41*, 3559–3566.

Land, M. & Tatler, B. (2009). *Looking and acting: Vision and eye movements in natural behaviour.* Oxford University Press.

Laubrock, J., Cajar, A., & Engbert, R. (2013). Control of fixation duration during scene viewing by interaction of foveal and peripheral processing. *Journal of Vision*, *13*(12), 11.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436–444.

Lee, M. D. & Wagenmakers, E.-J. (2014). *Bayesian cognitive modeling: A practical course.* Cambridge university press.

Lennie, P. (2003). The cost of cortical computation. *Current biology*, *13*(6), 493–497.

Levinson, J., Askeland, J., Becker, J., Dolson, J., Held, D., Kammel, S., ... Pratt, V., et al. (2011). Towards fully autonomous driving: Systems and algorithms. In *Intelligent vehicles symposium (iv), 2011 ieee* (pp. 163–168). IEEE.

Liu, T., Yang, Y., Huang, G.-B., Yeo, Y. K., & Lin, Z. (2016). Driver distraction detection using semi-supervised machine learning. *IEEE transactions on intelligent transportation systems*, *17*(4), 1108–1120.

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

Luke, S. G., Nuthmann, A., & Henderson, J. M. (2013). Eye movement control in scene viewing and reading: Evidence from the stimulus onset delay paradigm. *Journal of Experimental Psychology Human Perception & Performance*, *39*(1), 10.

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature neuroscience*, *9*(11), 1432–1438.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information.* New York, NY, USA: Henry Holt and Co., Inc.

Mauk, M. D. & Buonomano, D. V. (2004). The neural basis of temporal processing. *Annual Review of Neuroscience*, *27*, 307–340.

Merchant, H., Harrington, D. L., & Meck, W. H. (2013). Neural basis of the perception and estimation of time. *Annual Review of Neuroscience*, *36*, 313–336.

Miall, R. C., Christensen, L. O., Cain, O., & Stanley, J. (2007). Disruption of state estimation in the human lateral cerebellum. *PLoS biology*, *5*(11), e316.

Mital, P. K., Smith, T. J., Hill, R. L., & Henderson, J. M. (2011). Clustering of gaze during dynamic scene viewing is predicted by motion. *Cognitive Computation*, *3*(1), 5–24.

Montés-Micó, R. (2007). Role of the tear film in the optical quality of the human eye. *Journal of Cataract and Refractive Surgery*, *33*(9), 1631–1635.

Moraitis, T. & Ghosh, A. (2014). Withdrawal of voluntary inhibition unravels the off state of the spontaneous blink generator. *Neuropsychologia, 65,* 279–286.

Morvan, C. & Maloney, L. (2009). Suboptimal selection of initial saccade in a visual search task. *Journal of Vision, 9*(8), 444–444.

Morvan, C. & Maloney, L. T. (2012). Human visual search does not maximize the post-saccadic probability of identifying targets. *PLoS Computational Biology, 8*(2), e1002342.

Motala, A., Heron, J., McGraw, P. V., Roach, N. W., & Whitaker, D. (2018). Rate after-effects fail to transfer cross-modally: Evidence for distributed sensory timing mechanisms. *Scientific Reports, 8*(1).

Murphy, K. P. (2000). *A survey of pomdp solution techniques.* UC Berkeley.

Naase, T., Doughty, M. J., & Button, N. F. (2005). An assessment of the pattern of spontaneous eyeblink activity under the influence of topical ocular anaesthesia. *Graefe's Archive for Clinical and Experimental Ophthalmology, 243*(4), 306–312.

Najemnik, J. & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature, 434*(7031), 387.

Nakano, T., Kato, M., Morito, Y., Itoi, S., & Kitazawa, S. (2013). Blink-related momentary activation of the default mode network while viewing videos. *Proceedings of the National Academy of Sciences USA. 110*(2), 702–706.

Nakano, T. & Kitazawa, S. (2010). Eyeblink entrainment at breakpoints of speech. *Experimental Brain Research, 205*(4), 577–581.

Nakashima, R. & Kumada, T. (2017). The whereabouts of visual attention: Involuntary attentional bias toward the default gaze direction. *Attention, Perception, & Psychophysics, 79*(6), 1666–1673.

Navalpakkam, V., Koch, C., Rangel, A., & Perona, P. (2010). Optimal reward harvesting in complex perceptual environments. *Proceedings of the National Academy of Sciences, 107*(11), 5232–5237.

Neiman, T. & Loewenstein, Y. (2011). Reinforcement learning in professional basketball players. *Nature communications, 2,* 569.

Ng, A. Y., Russell, S. J. et al. (2000). Algorithms for inverse reinforcement learning. In *Icml* (pp. 663–670).

Nienborg, H. & Cumming, B. G. (2009). Decision-related activity in sensory neurons reflects more than a neurons causal effect. *Nature, 459*(7243), 89–92.

Oaksford, M. & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning.* Oxford University Press.

Oh, J., Han, M., Peterson, B. S., & Jeong, J. (2012). Spontaneous eyeblinks are correlated with responses during the stroop task. *PLoS ONE, 7*(4).

Oliva, A., Torralba, A., Castelhano, M. S., & Henderson, J. M. (2003). Top-down control of visual attention in object detection. In *Image processing, 2003. icip 2003. proceedings. 2003 international conference on* (Vol. 1, pp. I–253). IEEE.

Oliveri, M., Vicario, C. M., Salerno, S., Koch, G., Turriziani, P., Mangano, R., ... Caltagirone, C. (2008). Perceiving numbers alters time perception. *Neuroscience Letters, 438*(3), 308–311.

Paeye, C. & Madelain, L. (2014). Reinforcing saccadic amplitude variability in a visual search task. *J Vis, 14*(13), 20.

Palmer, E. M., Horowitz, T. S., Torralba, A., & Wolfe, J. M. (2011). What are the shapes of response time distributions in visual search? *Journal of Experimental Psychology Human Perception & Performance, 37*(1), 58.

Pariyadath, V. & Eagleman, D. (2007). The effect of predictability on subjective duration. *PloS one*, *2*(11), e1264.

Pekny, S. E., Izawa, J., & Shadmehr, R. (2015). Reward-dependent modulation of movement variability. *Journal of Neuroscience*, *35*(9), 4015–4024.

Perreault, L., Bernier, J., Bobée, B., & Parent, E. (2000). Bayesian change-point analysis in hydrometeorological time series. part 1. the normal model revisited. *Journal of Hydrology*, *235*(3), 221–241.

Peterson, M. F. & Eckstein, M. P. (2012). Looking just below the eyes is optimal across face recognition tasks. *Proceedings of the National Academy of Sciences USA. 109*(48), 3314–3323.

Peterson, M. S., Kramer, A. F., & Irwin, D. E. (2004). Covert shifts of attention precede involuntary eye movements. *Attention, Perception, & Psychophysics*, *66*(3), 398–405.

Ponder, E. & Kennedy, W. (1927). On the act of blinking. *Experimental Physiology*, *18*(2), 89–110.

Posner, M. I. (1980). Orienting of attention. *Quarterly journal of experimental psychology*, *32*(1), 3–25.

Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: Knowns and unknowns. *Nature Neuroscience*, *16*(9), 1170–1178.

Ratcliff, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychol. Bull. 86*(3), 446.

Rayner, K. (2009). Eye movements and attention in reading, scene perception, and visual search. *Quarterly journal of experimental psychology*, *62*(8), 1457–1506.

Remington W., R., Wu, S.-C., & Pashler, H. (2011). What determines saccade timing in sequences of coordinated eye and hand movements? *Psychonomic Bulletin & Review*, *18*(3), 538–543.

Renart, A. & Machens, C. K. (2014). Variability in neural activity and behavior. *Current opinion in neurobiology*, *25*, 211–220.

Renninger, L. W., Coughlan, J. M., Verghese, P., & Malik, J. (2005). An information maximization model of eye movements. In *Advances in neural information processing systems* (pp. 1121–1128).

Renninger, L. W., Verghese, P., & Coughlan, J. (2007). Where to look next? Eye movements reduce local uncertainty. *Journal of Vision*, *7*(3), 6.

Rothkopf, C. A. & Ballard, D. H. (2013). Modular inverse reinforcement learning for visuomotor behavior. *Biological cybernetics*, *107*(4), 477–490.

Rothkopf, C. A., Ballard, D. H., & Hayhoe, M. M. (2007). Task and context determine where you look. *Journal of vision*, *7*(14), 16.

Rothkopf, C. & Dimitrakakis, C. (2011). Preference elicitation and inverse reinforcement learning. *arXiv:1104.5687 [cs, stat]*. arXiv: 1104.5687.

Russell, S. J., Norvig, P., & Davis, E. (2010). *Artificial intelligence: A modern approach* (3rd ed). Prentice Hall series in artificial intelligence. Upper Saddle River: Prentice Hall.

Sanborn, A. N. & Chater, N. (2016). Bayesian brains without probabilities. *Trends in cognitive sciences*, *20*(12), 883–893.

Sanborn, A. N. & Chater, N. (2017). The sampling brain. *Trends in Cognitive Sciences*, *21*(7), 492–493.

Schieppati, M. & Schmid, M. (2011). Human Movement Science. *Human Movement Science*, *30*(2), 151–152.

Schiffman, H. R. (1990). *Sensation and perception: An integrated approach*. John Wiley & Sons.

Schmitt, F., Bieg, H.-J., Herman, M., & Rothkopf, C. A. (2017). I see what you see: Inferring sensor and policy models of human real-world motor behavior. In *Aaai* (pp. 3797–3803).

Schütz, A. C., Trommershauser, J., & Gegenfurtner, K. R. (2012). Dynamic integration of information about salience and value for saccadic eye movements. *Proceedings of the National Academy of Sciences*, *109*(19), 7547–7552.

Schütz, A. C., Lossin, F., & Kerzel, D. (2013). Temporal stimulus properties that attract gaze to the periphery and repel gaze from fixation. *Journal of Vision*, *13*(5), 6.

Schwartenbeck, P., FitzGerald, T. H., Mathys, C., Dolan, R., Kronbichler, M., & Friston, K. (2015). Evidence for surprise minimization over value maximization in choice behavior. *Scientific reports*, *5*, 16575.

Senju, A. & Csibra, G. (2008). Gaze following in human infants depends on communicative signals. *Current Biology*, *18*(9), 668–671.

Shepherd, S. V., Steckenfinger, S. A., Hasson, U., & Ghazanfar, A. A. (2010). Human-monkey gaze correlations reveal convergent and divergent patterns of movie viewing. *Current Biology*, *20*(7), 649–656.

Shi, Z., Church, R. M., & Meck, W. H. (2013). Bayesian optimization of time perception. *Trends in Cognitive Sciences*, *17*(11), 556–564.

Shin, Y. S., Chang, W. D., Park, J., Im, C. H., Lee, S. I., Kim, I. Y., & Jang, D. P. (2015). Correlation between inter-blink interval and episodic encoding during movie watching. *PLoS ONE*, *10*(11), 1–10.

Shultz, S., Klin, A., & Jones, W. (2011). Inhibition of eye blinking reveals subjective perceptions of stimulus salience. *Proceedings of the National Academy of Sciences USA. 108*(52), 21270–21275.

Simoncelli, E. P. & Olshausen, B. A. (2001). Natural image statistics and neural representation. *Annual review of neuroscience*, *24*(1), 1193–1216.

Sims, C. R., Jacobs, R. A., & Knill, D. C. (2011). Adaptive allocation of vision under competing task demands. *Journal of Neuroscience*, *31*(3), 928–943.

Slagter, H. A., Georgopoulou, K., & Frank, M. J. (2015). Spontaneous eye blink rate predicts learning from negative, but not positive, outcomes. *Neuropsychologia*, *71*, 126–132.

Smilek, D., Carriere, J. S. A., & Cheyne, J. A. (2010). Out of Mind, Out of Sight: Eye Blinking as Indicator and Embodiment of Mind Wandering. *Psychological Science*, *21*(6), 786–789.

Smith, D. T. & Schenk, T. (2012). The Premotor theory of attention: Time to move on? *Neuropsychologia*, *50*(6), 1104–1114.

Smith, P. L. & Ratcliff, R. (2009). An integrated theory of attention and decision making in visual signal detection. *Psychological Review*, *116*(2), 283–317.

Sommer, D. & Golz, M. (2010). Evaluation of perclos based current fatigue monitoring technologies. In *Engineering in medicine and biology society (embc), 2010 annual international conference of the ieee* (pp. 4456–4459). IEEE.

Stanford, T. R., Shankar, S., Massoglia, D. P., Costello, M. G., & Salinas, E. (2010). Perceptual decision making in less than 30 milliseconds. *Nature Neuroscience*, *13*(3), 379–385.

Staub, A. (2011). The effect of lexical predictability on distributions of eye fixation durations. *Psychonomic Bulletin & Review*, *18*(2), 371–376.

Stern, J. A., Boyer, D., & Schroeder, D. (1994). Blink rate: A possible measure of fatigue. *Human factors*, *36*(2), 285–297.

Sternad, D. (2018). It's not (only) the mean that matters: Variability, noise and exploration in skill learning. *Current opinion in behavioral sciences*, *20*, 183–195.

Streiffer, C., Raghavendra, R., Benson, T., & Srivatsa, M. (2017). Darnet: A deep learning solution for distracted driving detection. In *Proceedings of the 18th acm/ifip/usenix middleware conference: Industrial track* (pp. 22–28). ACM.

Sutton, R. S. & Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT press Cambridge.

Sweeney, D. F., Millar, T. J., & Raju, S. R. (2013). Tear film stability: A review. *Experimental eye research*, *117*, 28–38.

Tanner, J. & Itti, L. (2017). Goal relevance as a quantitative model of human task relevance. *Psychological Review*, *124*(2), 168–178.

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

Tatler, B. W., Baddeley, R. J., & Vincent, B. T. (2006). The long and the short of it: Spatial statistics at fixation vary with saccade amplitude and task. *Vision research*, *46*(12), 1857–1862.

Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based bayesian models of inductive learning and reasoning. *Trends in cognitive sciences*, *10*(7), 309–318.

Terhune, D. B., Sullivan, J. G., & Simola, J. M. (2016). Time dilates after spontaneous blinking. *Current Biology*, *26*(11), R459–R460.

Thorpe, S., Fize, D., Marlot, C., et al. (1996). Speed of processing in the human visual system. *Nature*, *381*(6582), 520–522.

Thrun, S., Burgard, W., & Fox, D. (2005). *Probabilistic robotics (intelligent robotics and autonomous agents)*. The MIT Press.

Todorov, E. & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature neuroscience*, *5*(11), 1226–1235.

Todorov, E. & Jordan, M. I. (2003). A minimal intervention principle for coordinated movement. In *Advances in neural information processing systems* (pp. 27–34).

Torralba, A., Oliva, A., Castelhano, M. S., & Henderson, J. M. (2006). Contextual guidance of eye movements and attention in real-world scenes: The role of global features in object search. *Psychological review*, *113*(4), 766.

Trommershäuser, J., Landy, M. S., & Maloney, L. T. (2006). Humans rapidly estimate expected gain in movement planning. *Psychological Science*, *17*(11), 981–988.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in cognitive sciences*, *12*(8), 291–297.

Trukenbrod, H. & Engbert, R. (2014). ICAt: A computational model for the adaptive control of fixation durations. *Psychonomic Bulletin & Review*, *21*(4), 907–934.

Tutt, R., Bradley, A., Begley, C., & Thibos, L. N. (2000). Optical and visual impact of tear break-up in human eyes. *Investigative Ophthalmology and Visual Science*, *41*(13), 4117–4123.

Valero-Cuevas, F. J., Venkadesan, M., & Todorov, E. (2009). Structured variability of muscle activations supports the minimal intervention principle of motor control. *Journal of neurophysiology*, *102*(1), 59–68.

van Beers, R. J. (2007). The Sources of Variability in Saccadic Eye Movements. *Journal of Neuroscience*, *27*(33), 8757–8770.

Volkmann, F. C., Riggs, L. a., & Moore, R. K. (1980). Eyeblinks and visual suppression. *Science*, *207*(4433), 900–902.

Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? optimal decisions from very few samples. *Cognitive science*, *38*(4), 599–637.

Wagenmakers, E.-J. & Farrell, S. (2004). Aic model selection using akaike weights. *Psychonomic Bulletin & Review*, *11*(1), 192–196.

Wang, H. X., Freeman, J., Merriam, E. P., Hasson, U., & Heeger, D. J. (2012). Temporal eye movement strategies during naturalistic viewing. *Journal of vision*, *12*(1), 16–16.

Wang, Y., Toor, S. S., Gautam, R., & Henson, D. B. (2011). Blink frequency and duration during perimetry and their relationship to test–retest threshold variability. *Investigative ophthalmology & visual science*, *52*(7), 4546–4550.

Warren Jr, W. H., Kay, B. A., Zosh, W. D., Duchon, A. P., & Sahuc, S. (2001). Optic flow is used to control human walking. *Nature neuroscience*, *4*(2), 213.

Wass, S. V. & Smith, T. J. (2014). Individual differences in infant oculomotor behavior during the viewing of complex naturalistic scenes. *Infancy*, *19*(4), 352–384.

Wickens, C., Helleberg, J., Goh, J., Xu, X., & J. Horrey, W. (2001). Pilot task management: Testing an attentional expected value model of visual scanning.

Williamson, S. S. (2004). Modulation of Gaze-Evoked Blinks Depends Primarily on Extraretinal Factors. *Journal of Neurophysiology*, *93*(1), 627–632.

Wilming, N., Harst, S., Schmidt, N., & König, P. (2013). Saccadic momentum and facilitation of return saccades contribute to an optimal foraging strategy. *PLoS computational biology*, *9*(1), e1002871.

Wittmann, M. (2013). The inner sense of time: How the brain creates a representation of duration. *Nature Reviews Neuroscience*, *14*(3), 217–223.

Wolfe, J. M. (1998). Visual search. *Attention*, *1*, 13–73.

World Health Organization. (2015). *Global status report on road safety 2015*. Geneva: World Health Organization.

Yan, W. J., Wu, Q., Liang, J., Chen, Y. H., & Fu, X. (2013). How Fast are the Leaked Facial Expressions: The Duration of Micro-Expressions. *Journal of Nonverbal Behavior*, *37*(4), 217–230.

Yang, S. C.-H., Lengyel, M., & Wolpert, D. M. (2016). Active sensing in the categorization of visual patterns. *Elife*, *5*.

Yang, S. C.-H., Wolpert, D. M., & Lengyel, M. (2016). Theoretical perspectives on active sensing. *Current Opinion in Behavioral Sciences*, *11*, 100–108.

Yarbus, A. (1967). *Eye movements and vision*. New York: Plenum Press.

Yorzinski, J. L. (2016). Eye blinking in an avian species is associated with gaze shifts. *Scientific Reports*, *6*, 32471.

Yu, C., Ballard, D. H., & Aslin, R. N. (2005). The role of embodied intention in early lexical acquisition. *Cognitive science*, *29*(6), 961–1005.

Yuille, A. & Kersten, D. (2006). Vision as bayesian inference: Analysis by synthesis? *Trends in cognitive sciences*, *10*(7), 301–308.

Zago, M., McIntyre, J., Senot, P., & Lacquaniti, F. (2008). Internal models and prediction of visual gravitational motion. *Vision Research*, *48*(14), 1532–1538.

Zaman, M. L. & Doughty, M. J. (1997). Some methodological issues in the assessment of the spontaneous eyeblink frequency in man. *Ophthalmic and Physiological Optics*, *17*(5), 421–432.

Zednik, C. & Jäkel, F. (2016). Bayesian reverse-engineering considered as a research strategy for cognitive science. *Synthese*, *193*(12), 3951–3985.

Zhao, Q. & Koch, C. (2013). Learning saliency-based visual attention: A review. *Signal Processing*, *93*(6), 1401–1407.

Zhao, R., Ouyang, W., Li, H., & Wang, X. (2015). Saliency detection by multi-context deep learning. In *Proceedings of the ieee conference on computer vision and pattern recognition* (pp. 1265–1274).

Zingale, C. M. & Kowler, E. (1987). Planning sequences of saccades. *Vision research, 27*(8), 1327–1341.

Curriculum Vitae

David Hoppe
Alexanderstraße 10
64283 Darmstadt, Germany
hoppe@psychologie.tu-darmstadt.de
+49 6151 16-23366

GENERAL

NAME	David Hoppe
DATE OF BIRTH	September 12th 1986 in Munich
NATIONALITY	German

EDUCATION

06/14 –	**PhD Student** *Technical University Darmstadt*
10/12 – 10/14	**Computer Science** MASTER OF SCIENCE *Technical University Darmstadt*
10/12 – 04/14	**Psychology** MASTER OF SCIENCE *Technical University Darmstadt*
10/08 – 08/12	**Computer Science** BACHELOR OF SCIENCE *University of Innsbruck, Austria*
10/08 – 08/12	**Psychology** BACHELOR OF SCIENCE *University of Innsbruck, Austria*
10/05 – 04/08	**Electrical engineering** *Technical University Munich*
– 06/05	**Gymnasium** ABITUR *Carl-Orff Gymnasium Unterschleißheim*

Obligatory Declaration

I declare that I have developed and written the enclosed doctoral thesis entiteled "Eye movements in dynamic environments" completly by myself, and have not used sources or means without declaration in the text. Any thoughts from others or literal quotations are clearly marked. This thesis was not used in the same or in a similar version to achieve an academic grading or is being published elsewhere.

Darmstadt, May 25, 2019

....
David Hoppe