INTERCEPTIVE ACTIONS IN ADULTS AND CHILDREN

Oscar Terence Giles

Submitted in accordance with the requirements for the degree of Doctor of Philosophy

The University of Leeds

School of Psychology

February 2017

Intellectual Property and Publications

The candidate confirms that the work submitted is his own and that appropriate credit has been given where reference has been made to the work of others.

Acknowledgements

I would like to express my utmost gratitude to my primary PhD supervisors, Richard Wilkie and Mark Mon-Williams. Their continuous encouragement, support and guidance made this thesis possible. I would also like to thank my supervisors in the School of Mechanical Engineering, Raymond Holt and Peter Culmer for their advice and support in developing the experimental apparatus used throughout this thesis.

Special thanks must go to all the postgraduate students, RAs and members of PACLab, who have made the last three years both enriching and enjoyable. I would particularly like to thank Callum, Radka, Ed, Earle, Katy, Faisal and George for the many insightful and interesting conversations that helped me in my research.

I would like to thank my parents and brother for their endless love and support. Special thanks to my partner Abigail, for patiently supporting me and getting me through the hard parts. Finally, I would like to dedicate this thesis to my loving and amazing grandmother, Irene Giles, who I will always miss dearly.

Abstract

Interceptive actions refer to goal directed movements in which people attempt to make a controlled collision with a moving object (e.g. catching a ball). Because interceptive actions often take place under severe temporal constraints, movements need to be executed with exquisite temporal accuracy and precision. To achieve this the sensorimotor system needs to: (i) accurately predict the motion of the target object, and (ii) move the intercepting effector (e.g. hand, bat) to a location through which the object will pass at just the right time. This presents the sensorimotor system with numerous computational challenges. Examining interceptive timing in adults provides insights into how these challenges are overcome in the developed sensorimotor system, while studying children can reveal how these abilities are acquired and how they are related to the development of other sensorimotor and cognitive processes. The first part of this thesis investigates the control of interceptive timing behaviours in adults. Chapter 3 provides evidence that online sensory information is combined with a-priori knowledge, using Bayesian integration, to optimise movement timing. Chapter 4 demonstrates that adults optimally time their movements to exploit a physical relationship between the speed and temporal precision of their movements. The second part of this thesis then examines interceptive timing abilities in children. Chapter 5 documents the developmental trajectory of interceptive timing abilities over childhood, revealing that performance is still far from adult levels by the time children finish primary school (age 11 years). Chapter 6 tests a common taxonomy of motor skills, revealing that interceptive timing tasks measure a somewhat distinct ‘motor construct’ from that measured by ‘fine’ and ‘gross’ motor tasks. Finally, chapter 7 reveals a relationship between interceptive timing abilities and academic attainment in mathematics, even after controlling for motor skills in other taxonomic domains. Together these experiments shed light on how humans are able to exquisitely time interceptive actions, and provide key insights into the ontogeny of this fundamental motor ability.

Contents

List of Figures

List of Tables

Abbreviations

BDA	Bayesian data analysis
CKAT	Clinical kinematic assessment tool
DCD	Developmental coordination disorder
ELPD	Expected log predictive density
HDI	Highest density interval
IT	Initiation time
MCMC	Markov Chain Monte Carlo
MT	Movement time
NHST	Null hypothesis significance testing
NI	National Instruments
NUTS	No-U-Turn algorithm
OC	Optimal control
PCT	Pre-programmed control of timing
SEN	Special educational needs
TTA	Time to arrival

Chapter 1

General Introduction

1.1 Overview

The evolutionary success of any animal species is ultimately dictated by their ability to interact with the environment. The control of complex movements endows animals with particularly advantageous evolutionary adaptations. Sophisticated control of movement engenders complex behaviours, from predator avoidance to feeding and hunting, as well as reproduction. Amongst the animal kingdom humans show a remarkable capacity to learn new sensorimotor actions, and so can perform an astonishing range of tasks, from throwing a spear to returning a tennis serve. Compared to the capabilities of some prey species (e.g. young deer can walk within a few hours of birth) humans are born with fairly limited hard-wired action capabilities and must spend many months learning to crawl, stand, walk and run. This developmental trajectory presumably has great evolutionary advantages since children quickly amass a huge skill repertoire that greatly surpasses what is currently possible at the cutting edge of intelligent robotics. At the heart of this sensorimotor prowess is the ability to skilfully operate in highly dynamic and unstructured environments (Lumelsky, 2006). Dynamic because the state of the environment changes over very short time scales (i.e. objects move) and unstructured because there is high uncertainty in how the environment will be arranged at any given point in time.

Dynamic environments present a major challenge to sensorimotor control. Objects in the world often move relative to the human actor, due to self-motion and/or the motion of objects themselves. The human will want to avoid colliding with some objects, while they may wish to bring about a controlled collision with others in order to achieve a desired outcome (e.g. hitting a falling ball). The latter interaction is referred to as an *interceptive action* and encompasses a broad range of goal directed behaviours. Interceptive actions all share a common prerequisite for achieving a successful collision: the intercepting effector (e.g. hand, bat) must coincide with the spatial location of the target object at the same point in time.

Successfully executing interceptive actions is difficult for a number of reasons. Unlike interactions with static objects, manual interceptive actions often take place in the context of severe temporal constraints. Consider the task of returning a serve from an elite tennis player. It is not uncommon for the ball to travel in excess of $50\mathrm{m\,s^{-1}}$, taking the ball less than half a second to travel the length of the court. This leaves little time for gathering sensory information, predicting the trajectory of the ball and deciding how to innervate the muscular-skeletal system in order to achieve an interception. Furthermore, the period of time in which a successful action can take place (referred henceforth as the *time window*) may be in the order of milliseconds, as a moving object will quickly move out of reach of the actor.

Skilled athletes with many years of training are able to successfully strike moving objects when the time window is shorter than 4ms, however even without this specific training healthy adults are able to time interceptive movements with exquisite precision, keeping timing errors within a standard deviation of ≈ 6ms (Brenner, van Dam, Berkhout, & Smeets, 2012). It is often noted in the literature that this level of performance is remarkable when compared to typical performance on other temporally demanding tasks (Brenner & Smeets, 2015b). For example, people show much worse temporal precision when attempting to tap two hands synchronously (Brenner et al., 2012), or when judging which of two events occurred first (Brenner & Smeets, 2010). Unsurprisingly researchers have taken great interest in how humans are able to control the timing of interceptive actions to achieve these levels of performance. Yet many questions remain regarding how adults perform interceptive tasks and very little is known about how these skills develop over childhood.

Studying interceptive actions in adults can reveal how interceptive timing behaviours are supported by the developed sensorimotor system. Yet examining the development of these abilities in children can illuminate not only how these skills are acquired, but also how they relate to development in other sensorimotor and cognitive domains. This chapter begins by reviewing the current literature on the timing of interceptive actions in adults and children, identifying some of the many remaining questions regarding how manual interceptive actions are performed and how the ontogeny of these abilities relates to more general motor and cognitive development.

1.2 Interceptive actions in adults

In interceptive actions the target object and intercepting effector can approach each other in many different ways. The limit case is moving the effector to intercept a non-moving object, but there are no real temporal constraints when performing this task (apart from a general requirement for actions to be achieved in a timely fashion). The reciprocal case is intercepting a moving object by placing the effector into its path. Here the timing constraints are weak and simply ensuring that the correct spatial position is adopted before the object arrives at that point will result in a successful interception. Thus getting into the right location as soon as possible will suffice as a control strategy. This might effectively achieve a blocking action, but tasks such as catching or hitting place much tighter spatiotemporal constraints over the interception with only a small time window in which the effector can arrive at the correct spatial location (as the object will quickly move out of range). These will be the sorts of skilled interceptive actions examined in the present thesis.

Successful skilled interceptive timing requires complex movements to be coordinated so that the effector (e.g. hand or bat) reaches a point in space where an interception can take place at precisely the right time. The challenges that must be overcome to achieve this are numerous (Franklin & Wolpert, 2011). Firstly, there are significant time delays present at every level of the nervous system, rendering visual signals about the objects location out of date with respect to the required action (by a non-trivial magnitude). Secondly, there is considerable uncertainty in the movement of the object, as visual information is insufficient to exactly specify the state of the object (i.e. its position, velocity, acceleration etc.). Thirdly, motor outputs are corrupted by noise, such that movements cannot be executed exactly as planned. In addition, coordinating interceptive actions requires the control of over 600 muscles, while non-stationarity in the sensorimotor system (e.g. neuromuscular fatigue) requires neural commands to be continuously adjusted in order to prevent systematic performance errors.

1.2.1 Time delays

Time delays in the sensorimotor system present a significant challenge to intercepting moving objects. These delays exist at every level of the nervous system (Nijhawan, 2008) and even the early stages of cortical visual processing are subject to significant delays. For example, average lags of $\approx$ 72ms have been recorded between the retina being stimulated and neurons in V1 responding (Lamme & Roelfsema, 2000). When reaching to static targets this does not necessarily present a major issue as the state of the object is unlikely to change during the lag. However, when objects are moving at high speeds the perceived state of the object may lag behind its true state by a substantial magnitude (see figure 1.1). To overcome these delays the sensorimotor system must make predictions about the object's likely trajectory (Nijhawan, 2008), allowing interceptive movements to be aimed to the position through which the object is likely to pass in the near future. The ability to predict the motion of objects over space and time allows adults to intercept moving objects, even when the time window is shorter than the lags typically observed in the conduction time of photoreceptors (Nijhawan, 2008).

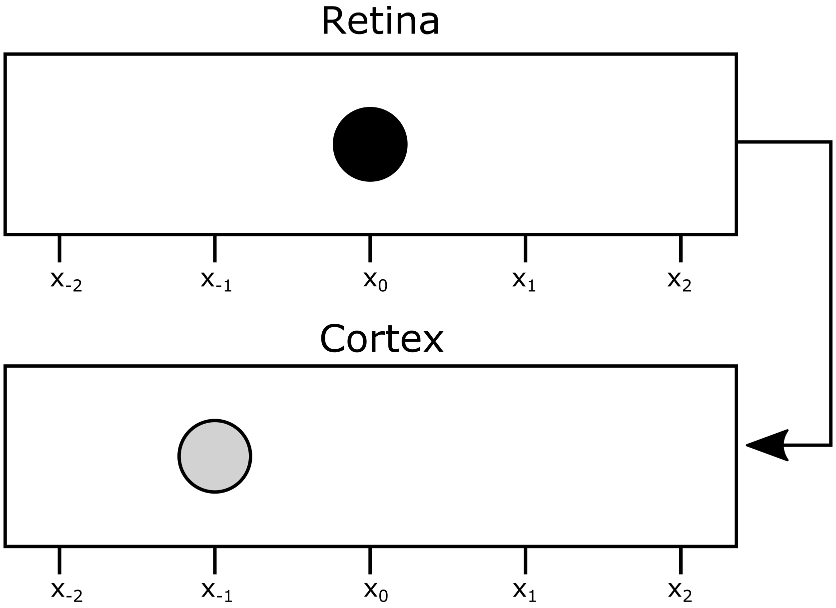

Figure 1.1: A one-dimensional representation of the retina mapped topographically onto the cortex. A ball moves from left to right at a constant velocity. A single instance in time is depicted in which the position of the ball is registered at x_{-1} on the cortex, while its true position is at x_0 on the retina. This lag is due to neural delays early in the sensorimotor system. Adapted from Nijhawan (2008).

1.2.2 Predicting target motion and control strategies

Predictions about the target object's motion are used to guide interceptive actions. Most viable control strategies suggest that people control the timing of their movements by predicting the trajectory of the object and then estimating the amount of time remaining until the target object reaches the spatial location of the planned interception, known as the time to arrival

(TTA) or time to contact. Predictive control models posit that movements are programmed on the basis of TTA estimates in advance of the movement onset, and once initiated are executed without the use of online sensory information. In contrast, *feedback control* models suggest that interceptive movements are controlled by continuously correcting for the error between the TTA of the target and the TTA of the intercepting effector. Both classes of models involve prediction of the target's motion in order to represent "where" the object can be intercepted and "when" the interception can take place. Models that involve separate representations of the spatial and temporal information are said to be *separable*.

Tresilian (2005) described a predictive control model, referred to here as the *pre-programmed control of timing* (PCT) model (see figure 1.2). The PCT model has two stages. First the desired movement time (MT) of the interceptive action is planned. Once a movement time has been programmed (MT_{prog}) the estimated TTA of the target ($TT\hat{A}_{tgt}$) is continuously monitored. The interceptive movement is then initiated at a point in time at which $TT\hat{A}_{tgt} = MT_{prog} + PT + TT$, where PT and TT are time delays associated with perception and the transmission of the descending motor commands. This simple model suggests that successful timing of the movement depends on the ability to accurately predict the target's TTA, compensate for delays in the sensorimotor system (PT and TT), initiate the movement at the right time and reliably produce a movement of the planned time duration (MT_{prog}).

Feedback control models have also been proposed which provide a possible strategy for the control of movement timing in interceptive actions. A very simple feedback control strategy specifies the force (F) that should be applied

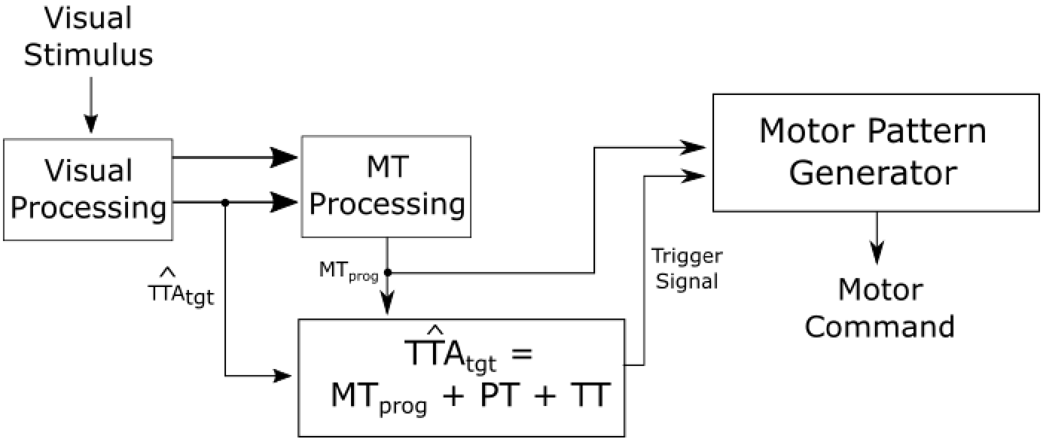

Figure 1.2: Block diagram of Tresilian's pre-programmed operational timing model. Adapted from Tresilian (2005).

to drive the effector to the interceptive location as follows,

$$F = \alpha(TT\hat{A}_{tgt} - TT\hat{A}_{eff}) \tag{1.1}$$

where α is the rate at which the error between the estimated TTA of the target and effector is corrected for (Tresilian, 2005). Most feedback models follow this general structure, although it's likely that first order approximations of TTA information are used (Lee, Young, Reddish, Lough, & Clayton, 1983; Senot, Prévost, & McIntyre, 2003). It has also been suggested that it is the error between the required velocity of the effector and its current velocity that drives the movement (Peper, Bootsma, Mestre, & Bakker, 1994; Tresilian, 1994). Of course hybrid schemes are also possible, where MT is planned in advance and the movement initiated in a similar fashion to the PCT model, but with the movements continuously corrected for on the basis of online visual information. Unfortunately it is often difficult to distinguish between these control schemes on the basis of behavioural data (Brouwer, Brenner, &

Smeets, 2003; Tresilian, 2005; Zago, McIntyre, Senot, & Lacquaniti, 2009), as behavioural 'markers' of feedback control, such as movement reversals (Montagne, Laurent, Durey, & Bootsma, 1999) can also be accounted for by single discrete corrections, rather than true continuous feedback control.

While there remains considerable debate over the control strategies used to intercept moving targets (Zago et al., 2009), the use of feedback control strategies seems implausible when intercepting very fast moving objects (Tresilian, 2005). This is because the time delays in the sensorimotor system mean that online corrections are of limited use when movements are very brief, because the movement will have ended before a correction can be implemented (Tresilian, 2012). In contrast the PCT model provides a simple solution for dealing with delays, in which they are corrected over repeated trials. Empirical support for this was provided by de la Malla, Lopez-Moliner, and Brenner (2012) who found that people were able to account for an experimentally imposed temporal delay, and found that visual information about the movement error at the end of the movement (indicating how much the target was missed by) was particularly important in correcting for delays. In addition it has been consistently found that adults make faster, briefer movements when the time window is smaller. This finding is difficult for feedback control models to explain as they would predict longer MTs as the temporal difficulty of the task increased (a Fitt's law type relationship between speed and accuracy). The PCT model can readily explain this finding as briefer MTs are associated with more precise movements in the absence of online corrections (Tresilian & Houseman, 2005).

1.2.3 Visual information and prior knowledge

Both the PCT and feedback control models suggest that the ability to estimate the TTA of the target object, and thus make predictions about the object's motion, is critical for successful interceptive actions. A large body of research has focused on the visual cues that may provide information regarding an object's TTA (Gray, 2009; Bootsma, Ledouit, Casanova, & Zaal, 2015; Brenner & Smeets, 2015a; Caljouw, van der Kamp, & Savelsbergh, 2004). Gibson suggested that all the information needed to support successful action is present in the visual stimulus (Gibson, 1961), without the need to invoke higher-level computations or representations. However, this position seems unlikely for rapid interceptive movements given the limitations of the human visual system. Sensory information is corrupted by noise (Faisal, Selen, & Wolpert, 2008; van Beers, Baraduc, & Wolpert, 2002) and can be ambiguous. This makes it particularly challenging to infer what environmental state resulted in the pattern of sensory information received (known as the *inverse problem of vision*). To overcome these limitations the nervous system may use prior knowledge of the statistical properties of the object's motion. Combining prior knowledge with online sensory information can help disambiguate sensory cues and reduce the impact of noise on perceptual judgements (Franklin & Wolpert, 2011).

Bayesian integration provides the statistically optimal method for combining prior knowledge with online sensory information, and there is growing evidence that adults' performance on some sensorimotor tasks can be explained by the brain acting as if it were performing Bayesian integration

(Körding & Wolpert, 2004; Vilares, Howard, Fernandes, Gottfried, & Kording, 2012; ?, ?; Verstynen & Sabes, 2011; Jazayeri & Shadlen, 2010; Franklin & Wolpert, 2011; Sato & Kording, 2014). Körding and Wolpert (2004) provide a simple hypothetical example of the use of Bayesian integration to estimate an object's velocity. When attempting to return a tennis serve the velocity of the ball cannot be known with certainty because sensory information is corrupted by noise. Over the course of the match the player could learn that certain velocities are more likely than others (the *prior*). The optimal estimate of the target's velocity is given by combining online visual information with this prior distribution.

1.2.4 Movement planning

Sensorimotor noise presents not only a problem for perceiving an object's state but also for the execution of planned movements. The PCT model of interceptive timing suggests that the ability to reliably execute movements of a pre-programmed duration is critical for success in interceptive actions. Even if the stimulus properties were estimated perfectly (e.g. TTA, position, velocity) the movements would not be timed perfectly over repeated trials. This is because the temporal dynamics of the descending motor commands may fluctuate over time, neural signals are corrupted by noise and the muscular-skeletal system does not respond in an identical way even when the same neural commands are received (Faisal et al., 2008). Feedback control mechanisms attempt to account for this variability by correcting for errors in the movement as they arise. However, as previously noted, online corrections may not be viable when intercepting fast moving targets. The PCT model provides a mechanism for minimising the effects of noise through the planning of the movement time (see figure 1.2). This is because making faster, briefer movements is known to improve the temporal precision (Brenner et al., 2012). This improvement occurs because errors, arising when moving at the wrong speed or misjudging the required amplitude of the movement, have less effect on the intercepting effector's movement time to the desired point of interception when moving at high speed (Brenner et al., 2012). It appears that people take this into account when planning their movements, making faster movements when the interceptive task requires greater temporal precision (Tresilian & Plooy, 2006; Tresilian, Plooy, & Carroll, 2004; Tresilian & Houseman, 2005).

The finding that people vary their movement times as a function of the temporal requirements of the task raises the question of why people do not simply move as fast as possible when intercepting all targets, as this would result in the highest level of temporal precision. A possible explanation is provided by theories of *optimal control*, which suggest that people seek to simultaneously minimise multiple costs when performing movements. Models of optimal control have been very successful in explaining behaviour on a large number of tasks (Todorov, 2004). In the case of interceptive timing the time window specifies the amount of time in which a successful interception can occur, and thus small timing errors may be acceptable so long as they remain within the bounds of the time window. Optimal controllers only correct for errors which interfere with the task goals, as error correction normally comes with some cost attached (Tresilian, 2012) (i.e. faster movements require more energy; less chance of online error corrections). Thus it may be that slower movement speeds are preferable when the reduced precision associated with them does not interfere with whether the target is hit or not.

1.3 Interceptive actions in children

While a large body of literature has sought to elucidate how adults achieve such amazing performance on interceptive timing tasks, few studies have examined how these timing abilities develop during childhood. Yet the possible benefits of understanding the ontogeny of these skills are numerous. Firstly, it can provide a window into how the motor system learns to overcome the complex computational challenges inherent in sensorimotor control. Secondly, the ability to perform interceptive actions is considered a fundamental movement skill (Foulkes et al., 2015) and poor sensorimotor abilities can provide a sensitive marker of both atypical sensorimotor and cognitive development (Purcell, Wann, Wilmut, & Poulter, 2012; Ament et al., 2014; Lefebvre & Reid, 1998; Whyatt & Craig, 2012). Thus understanding what information performance on interceptive timing tasks can provide regarding a child's sensorimotor status, over and above that assessed by other motor tasks (e.g. 'fine motor' skills), may provide insights into how to best identify and intervene in specific sensorimotor problems.

In addition, 'fine motor' skills are known to be predictive of children's academic attainment, particularly in mathematics (Son & Meisels, 2006). Thus if interceptive timing tasks capture a unique aspect of a child's sensorimotor status it may be that these measures can provide insights into the interrelated nature of these seemingly disparate domains of human performance. Thus understanding how interceptive timing abilities relate to other motor skills (i.e. whether they tap different ‘sensorimotor constructs’) and higher level cognitive faculties is a primary concern of developmental psychologists. Here we provide a brief overview of the existing data concerning the development of interceptive timing skills in children, before reviewing current literature on the relationship between interceptive timing and other motor and cognitive abilities.

1.3.1 Development of interceptive timing abilities

The ability to perform complex interceptive actions (i.e. catching a ball) appears later in a child’s developmental trajectory than manual skills with static targets (e.g. reaching to grasp a toy), yet the foundations of interceptive timing behaviours appear early in infancy. Infants begin to display smooth pursuit eye movements when tracking moving objects during their first few months of life (Shea & Aslin, 1990), a skill that appears to underpin the ability to predict target motion (Spering, Schütz, Braun, & Gegenfurtner, 2011). By 5 months of age infants begin to make catching movements towards the future locations of moving objects (Robin, Berthier, & Clifton, 1996; von Hofsten, Vishton, Spelke, Feng, & Rosander, 1998; von Hofsten, 1980), and appear to take account of their own movement ability when deciding whether an object is catchable (van Hof, van der Kamp, & Savelsbergh, 2008).

How the ability to time interceptive actions develops after infancy has not been well documented. While catching skills have been measured in children, one problem with these tasks is that they do not isolate 'timing' abilities well as errors may result from poor spatial control of the intercepting effector, rather than poor timing per se. In the adult literature the timing of interceptive movements has been studied by restricting arm movements to a single axis through the use of a 1-degree-of-freedom (DoF) manipulandum (see chapter 2). In these tasks only the amplitude of the movement needs to be controlled, which makes it easier to dissociate temporal errors from spatial ones (Tresilian, 2012). As well as allowing the timing aspects of interceptive movement to be examined, these tasks have a number of other benefits over the tasks typically studied in children (i.e. catching balls). First, they are capable of providing objective and detailed measures of performance, unlike the catching tasks typically used in standardised assessment tests which tend to use subjective measures of performance (Culmer, Levesley, Mon-Williams, & Williams, 2009). Second, measures like the MABC-2's catching task have to be adapted for different age groups because performance reaches ceiling and floor for the different age groups on the different tasks. Measuring performance with a single task is crucial for providing a parametric measure of interceptive timing performance.

Children's ability to time motor responses to external targets has been studied using a related class of tasks known as coincidence anticipation tasks. These typically involve viewing a runway of sequentially illuminating LEDs, or viewing a moving target. Participants then press a button to coincide with the last LED illuminating or the target reaching a specific location. Performance on these tasks may reach adult levels by age 11 (Haywood, 1980). However, the timing requirements of coincidence anticipation tasks are less demanding than interceptive timing tasks. Firstly, they do not require people to take account of the time it will take them to make a movement, beyond accounting for neural delays in pressing a button. Secondly, interceptive timing movements must be initiated much further in advance of the target reaching the desired point of interception, which requires accurate predictions of the object's TTA to be made over much longer time durations. Thus it is not clear when interceptive timing abilities will reach adult levels.

The dearth of research into the timing of interceptive actions in children is perhaps unsurprising given the current lack of a suitable research tool to measure interceptive timing abilities. A major barrier to the use of the objective kinematic measures used throughout the adult literature is that they often require expensive motion capture systems and are not particularly portable (Culmer et al., 2009). Portability provides a major advantage to collecting large samples of data as the measures can be deployed outside of the lab (i.e. in schools). Thus the development of a portable and objective measurement tool is vital in order to start answering critical questions regarding the development of interceptive timing abilities.

1.3.2 Interceptive timing, neuro-pathology and other motor abilities

Understanding the ontogeny of interceptive timing abilities in children is of particular interest given that performance on tasks which require excellent interceptive abilities (e.g. catching) can be indicative of neuro-developmental pathology. The MABC-2 (Barnett, Henderson, & Sugden, 2007), a popular standardised measure of motor skill, includes an ‘aiming and catching’ sub-scale, with a catching task that changes with the age group being tested. The inclusion of a catching task into the MABC-2 was initially based on subjective clinical insight (Schulz, Henderson, Sugden, & Barnett, 2011). However, catching tasks appear to provide a sensitive marker of several developmental disorders. For example, poor performance on interceptive timing and catching tasks is commonly observed in children with developmental coordination disorder (DCD) (P. H. Wilson, Ruddock, Smits-Engelsman, Polatajko, & Blank, 2013; Caçola, Ibana, Ricard, & Gabbard, 2016). In fact it is thought that children with DCD may have particular difficulties with motion prediction (Lefebvre & Reid, 1998) and with smooth pursuit eye movements (Langaas, Mon-Williams, Wann, Pascal, & Thompson, 1998) which are known to underpin successful interceptive timing. Children with other developmental disorders, such as autism spectrum disorder also seem to have particular problems with catching tasks (Whyatt & Craig, 2013).

Thus it seems likely that interceptive timing tasks may provide a useful measure of a child’s developmental status. Yet it is not known how interceptive timing tasks relate to other motor abilities. Interceptive tasks can share many similarities with other motor tasks, for example catching a ball requires excellent control of dynamic posture, while manually tracking a target requires dexterous movements of the arm and hand. However, interceptive tasks may provide a unique measure of children’s ability to predict the motion of objects in the world and may therefore tap into a distinct domain in which children may experience specific problems. Unfortunately, there is currently little experimental evidence regarding how different motor tasks meaningfully measure different classes of motor ability.

1.3.3 Interceptive timing and academic attainment

Piaget (1955) highlighted the importance of sensorimotor abilities in learning and developing abstract knowledge. The idea that sensorimotor development may have far reaching implications for development of higher order cognitive facets has recently gained popularity, with theories of embodied cognition suggesting that cognition is grounded in sensorimotor processes (M. Wilson, 2002; Gottwald, Achermann, Marciszko, Lindskog, & Gredebäck, 2016). A growing body of evidence suggests that sensorimotor processes play an important role in the development of numerical cognition (Crollen, Dormal, Seron, Lepore, & Collignon, 2013). This link may provide an explanation for why measures of fine motor ability can predict later academic attainment in mathematics (Grissmer, Grimm, Aiyer, Murrah, & Steele, 2010; Son & Meisels, 2006; Luo, Jose, Huntsinger, & Pigott, 2007).

Interestingly, Rigoli, Piek, Kane, and Oosterlaan (2012) reported that the 'aiming and catching' sub-scale of the MABC-2 showed a particularly strong relationship with mathematical attainment. A possible explanation for this is that representations of number, time and space are linked, possibly through a common representation of magnitude (Walsh, 2003). Thus it may be that representations of number are engendered by neural systems which underpin the ability to predict how objects move through time and space. Given that these abilities subserve interceptive timing, it is possible that performance on interceptive timing tasks may provide a useful measure for improving the predictive accuracy of statistical models which attempt to predict a child's academic performance. Circumstantial evidence for this hypothesis comes from clinical populations as children who typically perform poorly on interceptive timing tasks also tend to show poor mathematical ability (Pieters, Desoete, Van Waelvelde, Vanderswalmen, & Roeyers, 2012; Gomez et al., 2015; Simms et al., 2013; Hurks & Loosbroek, 2012), which is disproportionate to their problems in other domains (Simms et al., 2014). However, the relationship between interceptive timing and higher order cognitive skills (i.e mathematics abilities) has yet to be rigorously established even in non-clinical populations.

1.4 Thesis overview

This thesis examines how interceptive actions are performed by both adults and children. First, a brief review of the experimental methods used to study interceptive timing is provided in chapter 2. Given the current lack of an appropriate research tool for studying interceptive timing outside of the laboratory, the development of a bespoke interceptive timing research tool is documented. This chapter also provides an overview of the statistical methodologies utilised throughout this thesis, specifically Bayesian estimation. As many readers will be unfamiliar with Bayesian approaches to data analysis, a brief overview and primer are provided to aid readers in interpreting the analyses performed in the following experimental chapters.

The experimental work in this thesis is split into two sections. The first examines the mechanisms by which adults are able to achieve the high levels of temporal precision that have been well documented in the literature (Brenner & Smeets, 2015b). Chapter 3 explores how adults learn about the statistical properties of moving objects in order to improve interceptive timing performance and reduce the uncertainty in perceptual estimates. This builds on research showing that visual information may not be sufficient for the successful control of interceptive timing and that prior knowledge plays an important role (Zago et al., 2009). Chapter 4 then explores how adults try to minimise the effects of motor noise on movement timing by exploiting a physical relationship between movement speed and temporal precision. This builds on the work of Tresilian (Tresilian & Plooy, 2006; Tresilian et al., 2004; Tresilian & Houseman, 2005) who showed that adults appear to vary their movement time in response to changes in the time window.

The second experimental section explores the development of interceptive timing behaviours in children. While the abilities have been carefully examined in adults (Brenner & Smeets, 2015b), they have not been well documented in children. Chapter 5 develops a novel Bayesian model[1] to document how interceptive timing skills develop over childhood, revealing that the ability to time interceptive movements is far from adult levels by age 11, contrary to previous suggestions (Haywood, 1980). It also demonstrates that this timing deficit in young children is related to systematic errors in timing as well as changes in temporal precision. Next we examine how the development of interceptive timing abilities relate to other motor and cognitive facets. First chapter 6 explores the taxonomic relationship of interceptive timing to both 'fine' and 'gross' motor abilities. This was done to explore whether performance on interceptive timing tasks can provide information about a child's sensorimotor status over and above that provided by other common measures of sensorimotor skill. Chapter 7 then explores whether performance on an interceptive timing task is predictive of academic attainment, given the purported relationship between performance on the MABC-2's catching task and mathematical attainment (Rigoli et al., 2012). Finally, chapter 8 reviews the experimental findings and discusses areas of interest for future research.

---

[1] The model is a regression model for data analysis - as opposed to Bayesian models of behaviour which are explored in chapter 3.

Chapter 2

General methods and statistical analysis

2.1 Measuring interceptive timing

2.1.1 Overview

A diverse number of tasks and measures have been used throughout the literature to study interceptive timing behaviours. At one end of the literature, standardised measures of motor ability such as the MABC-2 (Barnett et al., 2007) simply require children to catch balls and the number of balls caught provides a measure of interceptive timing ability. This assessment has been useful in identifying children with poor sensorimotor abilities and is commonly used as a tool for the diagnosis of developmental coordination disorder (DCD). The MABC-2's catching task benefits from the fact that it can be easily conducted inside clinical and school environments. However, the measure's usefulness for studying interceptive timing behaviours has been questioned as researchers are not able to precisely control the trajectory of the ball and the measure appears to conflate throwing and catching abilities (Dirksen, De Lussanet, Zentgraf, Slupinski, & Wagner, 2016).

Tennis ball machines provide one method of obtaining greater control over the ball's trajectory (Van Waelvelde, De Weerdt, De Cock, & Engelsman, 2003), while more sophisticated measures of performance have used optical motion capture systems to obtain precise measures of interceptive timing abilities (see López-Moliner, Brenner, Louw, and Smeets (2010) for an example). One particularly useful interceptive task for examining the timing of interceptive movements has been developed by Tresilian et al. over a series of studies (Tresilian, 2012; Tresilian et al., 2004; Tresilian & Lonergan, 2002). In this task participants use a 1-DoF manipulandum to strike targets which are moving along a linear track orthogonal to the target's trajectory (see figure 2.1). Because the movement is restricted to a single axis the time window can be precisely controlled by changing the width and speed of the target and the width of the bat. As long as the interceptive movement achieves sufficient amplitude the bat is guaranteed to reach the interceptive point. Thus errors occur when the bat arrives too early or late to hit the target, allowing temporal errors to be investigated. In interceptive tasks with more degrees of freedom (2D or 3D movements), an error may occur because of errors in the spatial trajectory of the interceptive effector as well as timing errors. Again, this set-up uses optical tracking systems to record the position of the target and bat.

While optical tracking allows for sophisticated and detailed measurements

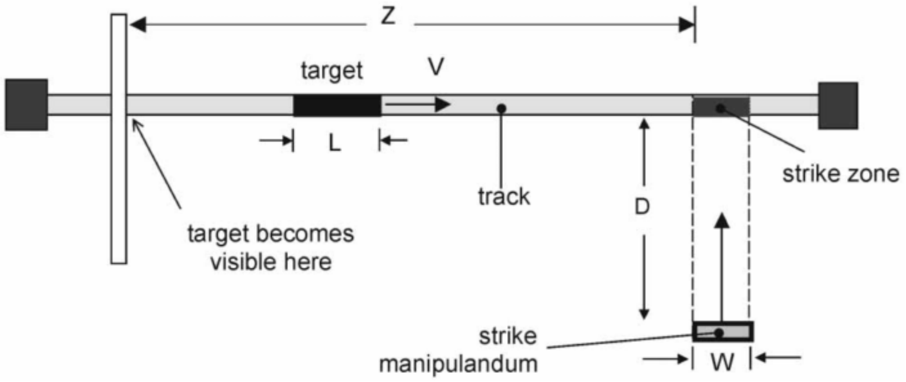

Figure 2.1: A simple interceptive timing task employed by Tresilian et al. A target moves along a linear track and the participant attempts to hit it using a manipulandum. The manipulandum moves along a linear track positioned perpendicular to the target's path. The target can only be struck when it is within the *strike zone*. The time the target is in the strike zone is known as the *time window*. In this task the time window is given by $\frac{L+W}{V}$, where V is the target's velocity.

to be made within the laboratory setting, they are generally not suitable for testing in clinical and school settings because of their large set up costs and lack of portability. Virtual interceptive timing tasks provide a potential method for studying interceptive timing outside of the laboratory. These tasks use computer displays to render virtual moving targets which people can intercept by controlling an on-screen bat (or cursor) via an interface (i.e. a mouse or stylus). These tasks can provide detailed kinematic measures of performance (see Brenner and Smeets (2015b) for an excellent example) as well as having the benefit of being much more portable than track based interceptive tasks.

2.1.2 Interceptive timing apparatus

In order to study interceptive timing abilities in adults and children a "virtual" version of the track based intercepting task used by Tresilian et al., (see figure 2.1) was developed. A large high speed gaming monitor (BenQ XL2720Z LCD display; Resolution: 1920 × 1080, size: 548 × 642mm, brightness: $300\text{cd}\,\text{m}^{-2}$, refresh rate: 144Hz) was used to present stimuli which were generated using a Python interface to the OpenGL graphics library. In most cases the stimuli consisted of a moving target and "bat" which could be used to intercept the target. The high frame rate of the display allowed for minimal lag between the position of the input device and the position of the on-screen bat.

A bespoke 3D printed 1-DoF manipulandum was developed to control the on-screen bat. The 3D manipulandum[1] was mounted on two linear tracks, allowing it to move smoothly along a single axis (see figure 2.2). Two iterations of the input device were developed. In the first iteration a high speed gaming mouse (Logitech G402 Hyperion Fury) was used to record the displacement of the manipulandum. This allowed the position of the bat to be estimated at 144Hz. This set up was used in experimental chapters 5, 6 and 7. The downside of this set-up was the position of the bat had to be calibrated at the start of each trial to ensure that the estimated position did not drift over repeated trials.

A second iteration of the manipulandum replaced the high speed mouse with a linear potentiometer and was used in experimental chapters 3 and 4. The potentiometer had a cable which was attached to the manipulandum.

---

[1] Thanks to Resolve Engineering (http://www.resolve-re.co.uk/) for help with the 3D modelling and printing.

Figure 2.2: The bespoke 3D printed manipulandum used throughout this thesis. Its displacement along the linear tracks could be measured via a linear potentiometer (depicted) or a high speed gaming mouse mounted under the handle.

As the cable extended it varied the voltage across an electrical circuit which was monitored by a National Instruments (NI) DAQ device. This allowed the position of the bat to be measured at high speeds (> 500Hz). A custom Python script allowed the position of the bat to be asynchronously polled and passed to the program controlling the on-screen stimulus and trial logic. A graphical representation of this system is shown in figure 2.3.

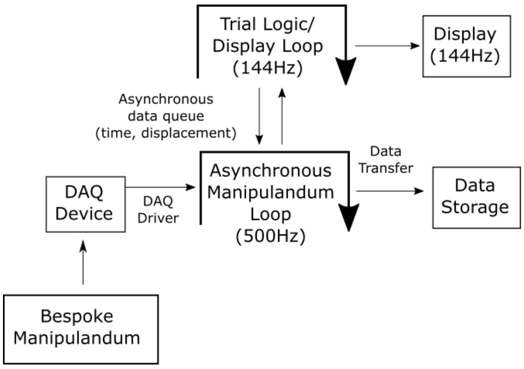

Figure 2.3: Structure of the interceptive timing task's apparatus. A bespoke 3D printed manipulandum with 1-DoF moved along a linear track. The displacement of the manipulandum was measured via a linear potentiometer. The potentiometer created a change in voltage across an electric circuit which was measured using a NI-DAQ device. A custom python script communicated with the DAQ device, carried out simple signal processing and converted the voltage measurement into a measure of displacement. The device achieved sub mm accuracy. The displacement data was recorded to disk at 500Hz while the display was updated at 144Hz. Stimuli were rendered to screen using a hardware accelerated OpenGL graphics layer.

2.2 Statistical methodology: Bayesian estimation

This thesis uses a statistical methodology known as Bayesian estimation. As this methodology is not currently widely used in the psychological sciences a brief justification for its use is provided. A number of texts exist which provide an excellent primer to the methods used throughout this thesis (Gelman, 2014; Kruschke, 2015, 2013). However, as readers may be unfamiliar with Bayesian estimation a quick overview of the general analysis approach is warranted.

2.2.1 Why Bayesian?

For the last century null hypothesis significance testing (NHST) has dominated statistical inference in the psychological sciences and many other disciplines. However, the application of NHST has been criticised for decades and has recently come under intense scrutiny (Dienes, 2011; Brooks, 2003; Kruschke & Liddell, 2017a; Goodman, 2008; Cumming, 2014; Cohen & Hubbard, 1995; Lambdin, 2012). Flexibility in the application of NHST can allow authors to present almost any finding as significant (Simmons, Nelson, & Simonsohn, 2011), a possible contributing factor to the crisis of replication in psychology (Open Science Collaboration, 2015). In addition, misunderstandings about p-values are pervasive among researchers (Goodman, 2008), while there is growing recognition of the need to move away from black and white thinking regarding the presence or absence of an effect, and move towards a cumulative science which incrementally improves estimates of effect magnitudes and uncertainty (Kruschke & Liddell, 2017b). In light of growing recognition of the problems associated with NHST, researchers, editors and journals are slowly moving away from NHST.

Several alternative approaches to NHST are commonly proposed, including frequentist estimation (Cumming, 2014), Bayes' factor (Dienes, 2014) and Bayesian estimation (Gelman, 2014). Kruschke and Liddell (2017b) provide an excellent discussion of the key differences between these approaches. In short, Bayes' factor is a Bayesian approach which compares the relative strength of the evidence for two competing models, and provides a method for hypothesis testing. A weakness of this approach is that the priors assigned to competing models can dramatically alter the conclusions drawn from the Bayes' factor. In addition, this approach normally proceeds by testing a null model (although it is not obligatory to do so), in which the model's prior is an infinitely dense spike at a point in the parameter space (e.g. precisely zero difference between two group means). Like NHST, this approach has been criticised as many null hypotheses are extremely implausible, in that we very rarely expect there to be exactly zero difference between groups.

Frequentist and Bayesian estimation take a very different approach to statistical analysis, attempting to quantify estimates of effect sizes and uncertainty. Estimation approaches make use of mathematical models which provide a description of the data generation process (i.e. how samples of data are generated from a target population). A statistical model can also be fit to a real dataset by finding values for the model's parameters which allow it to create simulated data which is as similar as possible to the real data. In frequentist estimation this is commonly achieved using a technique known as maximum likelihood estimation (MLE). Consider the following simple generative model $y \sim \text{Normal}(\mu = 10, \sigma = 5)$, which states that the data set y was sampled from a population of normally distributed values with a mean of 10 and standard deviation of 5. Now imagine that we had collected our data set y but didn't know the population mean (for simplicity we'll assume that we know the population's standard deviation). We could use MLE to find the value of μ which makes the data most likely. To do this we need to define a likelihood function as follows,

$$P(y|\mu, \sigma = 5) = \prod_{i}^{N} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(y_i - \mu)^2}{2\sigma^2}}. \tag{2.1}$$

This function simply says that the likelihood of dataset given a certain value of μ and σ is found by multiplying the probability of each data point under a normal distribution. In our example we assume that we already know σ so we simply plug in every possible value of μ from negative to positive infinity. The value of μ which results in the maximum value from the likelihood function is our MLE. This is illustrated in figure 2.4, which plots the likelihood function. The maximum likelihood function is shown as a red point.

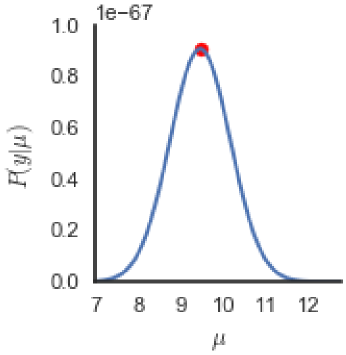

Figure 2.4: The likelihood function $P(y|\mu, \sigma = 5)$ for our sample from the target population.

In our very simple example the MLE is actually just the mean of the data sample y. However, for more complex models in which we may wish to simultaneously find MLEs for many model parameters we must employ optimisation algorithms which iteratively attempt to find the mode of the likelihood function. With more complex models we can use this approach to calculate effect sizes and estimate the magnitude of effects. For example, a more complex model may examine two groups each with their own means. We could use MLE to estimate the means of both groups and then examine the magnitude of the difference between them.

A key problem that arises from this approach is that the data set may still be very plausible under a whole range of other model parameter values. Cumming (2014) suggests that researchers use confidence intervals to quantify uncertainty in the model's parameter values. Unfortunately, despite their widespread use, confidence intervals are generally not able to quantify uncertainty in this way (Morey, Hoekstra, Rouder, Lee, & Wagenmakers, 2015). 95% confidence intervals are often incorrectly interpreted as a range of parameter values in which we believe there is a 95% probability that the true parameter value falls. What they actually provide is an interval which can be calculated for a given sample from a population. When we repeatedly sample from the population and calculate the 95% confidence interval for each sample we would expect 95% of the confidence intervals to contain the true parameter value.

Bayesian estimation provides the solution for quantifying uncertainty in parameter estimates. Rather than providing a single point estimate for the values of a model's parameters we would like to know what values of the model's parameters we consider plausible given our data set, and how plausible those values are relative to each other.

2.2.2 Calculating the posterior

Bayesian estimation uses Bayes' rule to reallocate "credibility across a space of candidate possibilities" (Kruschke, 2013), where the "possibilities" are possible values of parameters in a statistical model. Formally, Bayesian estimation calculates a joint posterior distribution over the parameters in a statistical model using Bayes' Rule,

$$P(\theta|y) = \frac{P(y|\theta)P(\theta)}{P(y)} \tag{2.2}$$

where θ is a vector of all the parameters in a statistical model and y is the data. The posterior distribution $P(\theta|y)$ is proportional to the likelihood distribution multiplied by the prior $P(\theta)$. Notice that the likelihood distribution is exactly the same as that used in MLE as defined in equation 2.1[2]. The prior distribution specifies how credible we believe different parameter values to be before we collect data.

For some simple statistical models we can calculate the posterior distribution analytically. However, for most of the models we are interested in it becomes impossible find analytical solutions. This is because the normalising term, $P(y)$, which ensures that the posterior distribution integrates to 1, can be difficult to calculate. It is primarily for this reason that Bayesian estimation has only recently started to gain wide spread popularity. The advent

---

[2] The likelihood function is also referred to as the likelihood distribution. They are equivalent

of Markov Chain Monte Carlo (MCMC) methods in the 1950's provided a method for approximating complex multidimensional integrals, such as $P(y)$, and advances in modern computational power allowed these methods to be used on your average desktop computer.

MCMC methods for Bayesian estimation work by constructing a Markov chain that has an equilibrium distribution which matches the joint posterior distribution (Gelman, 2014; Betancourt, 2017). In other words, MCMC algorithms randomly sample values for parameter values in statistical models, and they sample these values in proportion to their posterior probability. This is illustrated in figure 2.5. The left panel shows the first few steps of a Markov chain generated by an MCMC algorithm, in which it randomly samples values for a parameter (μ) in a statistical model. The right plot then shows a normalised histogram of the sampled values, with the actual posterior distribution superimposed over the top (red curve). We can see that the MCMC algorithm sample parameter values in proportion to their posterior probability, allowing us to approximate the posterior distribution.

This thesis makes use of a particular type of MCMC algorithm known as Hamiltonian Monte Carlo (HMC). Unlike many early MCMC algorithms, HMC scales well with model complexity, allowing complex hierarchical models to be fit within a reasonable time frame. Specifically this thesis makes use of a variant of HMC called the 'No-U-Turn' algorithm (NUTS) (Hoffman & Gelman, 2011), as implemented in the probabilist programming language Stan (Carpenter et al., 2016).

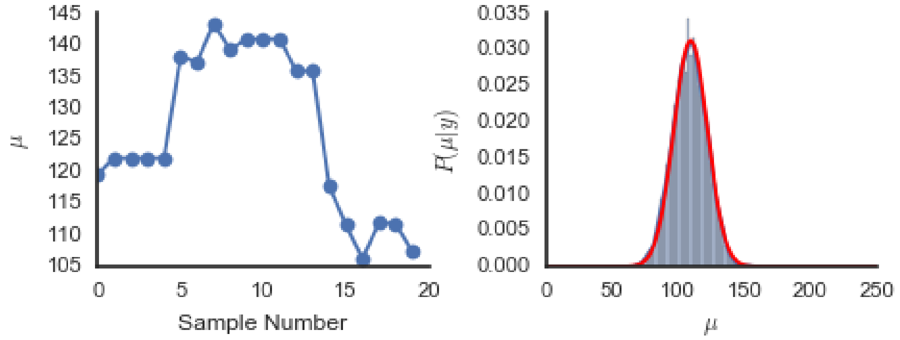

Figure 2.5: An example of an Markov Chain Monte Carlo algorithm. The left panel shows an MCMC algorithm randomly sampling values for a parameter in a statistical model. The right panel shows a normalised histogram of these values with the true posterior distribution superimposed (red curve). We can see that the MCMC algorithm is able to approximate the posterior distribution.

2.2.3 BDA in action

Bayesian estimation follows a simple work flow. First an appropriate statistical model is specified[3]. The model is then fit using the experimental data, which involves approximating the joint posterior distribution over the model's parameters. Once the model has been fit the posterior can be examined, summarised and used to guide decisions about the process under investigation.

For a simple example of Bayesian estimation in action, consider an experiment in which some quantity of interest was measured in two groups of participants. A researcher may then want to perform a statistical analysis to examine whether the means of the populations from which the participants were sampled were different. In frequentist statistics an independent t-test would commonly be used for this task.

---

[3] The models used throughout this thesis are all regression models

In Bayesian estimation we first need to define a model of how the sample data was generated from the populations. A simple model of the data might look like the following,

$$y \sim \text{Normal}(x\beta, \sigma) \tag{2.3}$$

where y is the sample data set, x is a matrix where each row indicates the group membership of a participant, β is a vector of regression coefficients corresponding to the mean of each group and σ is the variance parameter. The model simply states that the data set was sampled from a normal distribution with a mean that depended on the participant's group. This simple model has three unknown parameters which we would like to estimate, the two parameters of β (β_1 and β_2) and σ. We need to provide priors over these parameters. These priors should express all knowledge about the model at hand. Complete ignorance can easily be modelled as a uniform distribution over the parameter's support. In this case the modes of the marginal posterior estimates will be equal to the point estimates obtained through frequentist MLE. However, it is very uncommon to have a situation in which no prior information is known. Weakly informative priors are informed only by the scale of the data. For instance, when measuring someone's height in meters it is known a-priori that measurements above 2 metres are unlikely, and measurements above 5 meters are essentially impossible. For our example we could model the priors over the β parameters as follows,

$$\beta \sim \text{Normal}(0, 1000) \tag{2.4}$$

The choice of prior in this example is arbitrary, but in real applications the prior should reflect all knowledge about the problem at hand. In fact, the prior can also be improper (i.e. doesn't need to integrate to one) so long as the resulting posterior is proper. We also need to choose a prior for the model's standard deviation parameter σ. Because the standard deviation must be positive our prior must only have support over positive values. In this example we use a truncated Cauchy distribution,

$$\sigma \sim \text{Cauchy}^{+}(0, 50) \tag{2.5}$$

The Cauchy distribution is a continuous probability distribution commonly employed for scale parameters. It has much fatter tails (i.e. greater density in the tails of the distribution) than the normal distribution, allowing for priors to be specified with most of the density near zero, but with a non negligible amount of density over larger parameter values. Notice the + symbol in equation 2.5 which indicates that the distribution only has support over positive parameter values.

We now have a full generative model which we can fit using the data set y. For the purposes of this primer this model was fit using a simulated data set with known parameters. The group means of the population were set at $\beta_1 = 30$ and $\beta_2 = 34$, while the standard deviation was set at $\sigma = 5$. Of course in practise we would not already know these values, else we would not need to do any statistical analysis. The aim of Bayesian estimation is to quantify what parameter values we should consider plausible. For our simulated data set we should find that the posterior distribution contains the true parameter value.

Once the model has been fit using an appropriate MCMC algorithm (in this case HMC) we can example the approximate posterior distribution. Figure 2.6 shows the marginal posterior distribution over the fitted model's parameters. These provide complete distributional information over how credible we believe possible values of the model's parameters to be. These distributions can be summarised using a number of statistics, for example by their mean and standard deviation. Another useful summary statistic is the highest density interval (HDI). The $x\%$ HDI shows the interval in which there is a $x\%$ chance of the true parameter falling.

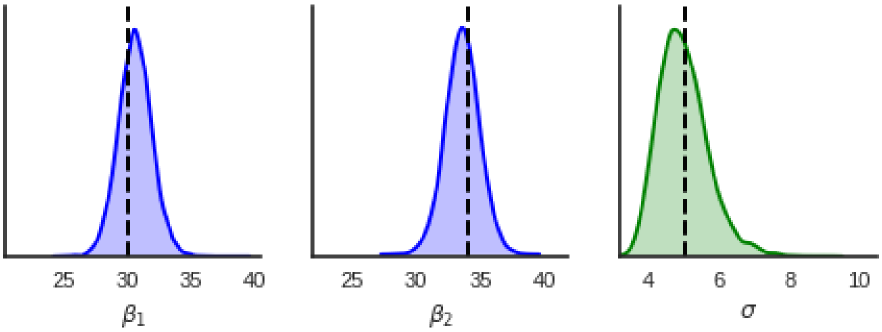

Figure 2.6: Marginal posterior distributions over the model's parameters. These provide complete distributional information regarding the credibility of the possible values of the parameter. The true values are shown as dotted black lines. It is clear that the posterior distribution contains the true model parameter.

Once the joint posterior distribution has been approximated we can make inferences from the data. The simplest question we can ask is, is there a difference between the two β parameters? In other words, are the population means of the two groups different? We can do so by comparing the two marginal distributions over β_1 and β_2 (i.e. by performing a contrast). In practise this is done by subtracting the samples generated by the MCMC algorithm for β_2 from those generated for β_1. Figure 2.7 shows the contrast between the two means ($\beta_1 - \beta_2$). This provides complete distributional information over the credible differences between the means. Again this distribution can be summarised as before using statistical functions (i.e. mean, SD, HDI).

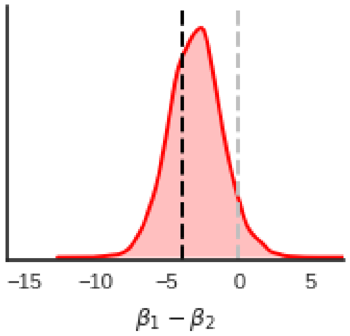

Figure 2.7: The magnitude of the difference between the two groups' mean parameters ($\beta_1 - \beta_2$). These provide complete distributional information regarding the credibility of the possible differences between groups' means. The black dotted line shows the true difference between the means. The lighter dotted line indicates the zero point.

2.2.4 Making inferences

A key question we often want to answer is whether a statistical effect exists, and whether it is of a substantive magnitude? In other words we want to use the posterior distribution to make a decision regarding whether to declare the existence of an effect. In the above example we might wish to ask whether there is a difference between the means of the populations we collected our data sample from. Answering this questions is surprisingly difficult and requires us to invoke decision theory (Kruschke, 2015) in which the costs of being incorrect are considered. However, this is often impractical owing to the difficulty in specifying appropriate cost functions, thus researchers often seek to make a decision from the posterior alone using decision heuristics.

One possible heuristic is to ask whether the HDI of the contrast contains zero. If it does not then we might conclude a non-zero difference. Of course, we could use the 95% HDI, or the 80% HDI or any other arbitrary interval. The 95% HDI is commonly used simply to match the conventions of confidence intervals. A related statistic is the proportion of the posterior distribution which falls on either side of zero (denoted η throughout this thesis). This allows us to make a statement about the probability that the effect is in the hypothesised direction. The 95% HDI will exclude zero when $\eta < 0.025$ or $\eta > 0.975$.

When the vast majority of the posterior density falls on one side of zero we can safely assume a non-zero difference in parameter values. However, in the case of our toy example the 95% HDI = [-7.31, 0.16]. Thus if we were following the 95% HDI heuristic we would conclude that we are not sufficiently certain of an effect. Examining the η statistic reveals there is a 97% chance of a negative difference between the two group means, thus if we were 0.5% more certain the 95% HDI would exclude zero. Depending on the specific statistical question and application we may conclude that we are happy that an effect exists. Unfortunately this raises the problem that we may allow ourselves the flexibility to conclude that any effect is in the direction we hypothesised, by slightly changing the value of η we will accept as ’significant’. Thus throughout this thesis an arbitrary threshold is used to decide whether a non-zero effect is ‘highly probable’. When the hypothesis is two tailed then we state that 97.5% of the contrast’s distribution must be on one side of zero ($\eta < 0.025$ or $\eta > 0.975$, or equivalently the 95% HDI excludes zero). When the hypothesis is one tailed then 95% of the contrast’s distribution must be on one side of zero ($\eta < 0.05$ or $\eta > 0.95$). When 90% of the posterior distribution falls on one side of zero we view this as evidence for ‘probable’ effect which strongly warrants further examination but should be treated with caution.

2.2.5 Conclusion

In summary Bayesian estimation allows us to estimate parameter values in statistical models. Unlike hypothesis testing (i.e. NHST and Bayes’ factor) this allows us to examine effect sizes and quantify the magnitude of statistical effects. Bayesian estimation differs from frequentist information in its use of probability to quantify uncertainty in parameter values.

Part I

Adults

Chapter 3

Turning down the noise in interceptive timing

Chapter Abstract

Humans demonstrate an exquisite ability to intercept moving targets. Interceptive timing behaviours require visual information, but there is growing evidence that vision is not informationally sufficient to allow accurate and precise behaviour in some interceptive timing tasks. Humans may use prior knowledge about the probability of a target's trajectory to optimise performance. Indeed, there is evidence that the target's speed on previous trials influences human interceptive timing behaviour. More specifically, when the prior distribution of target speeds (their trial-by-trial variability) is manipulated in a coincidence timing task, people show behaviour consistent with the brain acting as a Bayesian operator. Here we investigated whether an interceptive timing task would also be affected by the previous trial distribution structure. We explored whether the relationship between temporal errors and

target speed varied depending on the standard deviation of the target speed distribution, as predicted by Bayesian theories of sensorimotor control. This hypothesis suggests that people will make larger temporal errors when target speeds are less likely a-priori. Participants were exposed to either a narrow or wide target speed distribution (order counterbalanced). The distribution affected the relationship between temporal errors and target speed, although the effect was larger when participants were first exposed to a narrow speed distribution. The results are qualitatively consistent with the brain acting as a Bayesian operator, although a simple strategy of reinforcing successful movements on previous trials may provide a good approximation of Bayesian optimal behaviour.

3.1 Introduction

The interception of moving targets requires individuals to aim their movements at a point in space through which a target will pass in the future (the interception point). Most plausible control strategies involve estimating the future trajectory of the target, determining where it can be intercepted (the interception point), and gauging the time remaining before the object reaches the interception point (Tresilian, 2005). This task is made more difficult because the individual must not only determine when the object will arrive at the interception point, but how long it will take them to move their effector to that location. This is a non-trivial problem owing to the substantial time delays associated with neural conduction, muscle activation, receptor transduction and information processing (van Beers et al., 2002). Time delays are inherent at every level of sensorimotor processing. For example, the processing of visual information is subject to sizeable delays, with lags of 200ms recorded at various levels of the visual system (Nijhawan, 2008).

Gibson (1961) suggested that the environment contains all of the information required for successful interaction with objects in the world. From this viewpoint, the use of information from memory or the deployment of constructive processes are not necessary for action, as the stimulus contains ‘invariants’ that the perceptual system can detect directly. It is now widely considered that this Gibsonian viewpoint is not wholly correct and the information necessary for sensorimotor control is neither solely contained in invariants nor is it directly detected (Tresilian, 1999). There is widespread support, however, for Gibson’s suggestion that the stimulus is informationally sufficient in ecological conditions (Tresilian, 1999). Nevertheless, there are some tasks (including interceptive timing tasks) where visual information alone is not sufficient. In these situations, prior knowledge has the potential to improve performance.

Prior knowledge of the physical laws that govern the universe provides individuals with the ability to predict how objects move through space and time. The use of such information can explain the high levels of performance observed when humans intercept objects falling under gravity, even when vision is degraded and the object is small (Brenner, Driesen, & Smeets, 2014). This can be explained by the use of prior information about statistical regularities in the target object’s trajectory. In the case of falling balls, it is proposed that the nervous system stores information about gravitational acceleration. There is empirical evidence in support of such proposals as humans make predictable errors when trying to catch falling balls in different gravitational fields (Zago, McIntyre, Senot, & Lacquaniti, 2008). It has been suggested that prior knowledge is stored and utilised through the existence of 'internal models' that can learn how objects move through time and space to make predictions about an object's likely future state (Zago et al., 2008).

The ability to predict how objects move through space and time can help humans to optimise their behaviour in the face of substantial sensorimotor delays. In addition it may help overcome challenges posed by the fact that sensory signals are corrupted by noise (Faisal et al., 2008; van Beers et al., 2002) which places a limiting factor on perception. When sensory signals are corrupted by noise, statistically optimal inferences can be made by incorporating prior knowledge into perceptual estimates through the use of Bayesian integration. Thus, using knowledge about a target's likely states (inferred from many trials) may improve performance on sensorimotor tasks. Evidence in support of the hypothesis that humans use such prior knowledge was provided by Körding and Wolpert (2004) in an aiming task and force amplitude task. These results suggested that using prior knowledge about a target's likely states might be used across a variety of sensorimotor tasks, including timing tasks.

Miyazaki, Nozaki, and Nakajima (2005) tested the idea that people may show behaviour that is qualitatively consistent with that of a Bayesian operator by manipulating the structure of trial-to-trial variability across their experiment. Miyazaki et al.'s coincidence timing task required participants to press a button at the same time that the final LED from a line of three LEDs was illuminated. The duration between the illumination of each LED

(Δ_{LED}) was drawn from a normal distribution (referred to as the stimuli distribution). Participants experienced two conditions, one in which the standard deviation (SD) of this distribution was narrow and one in which it was wide. Following several hundred trials, it was found that the slope and intercept of the temporal errors plotted against Δ_{LED} depended on the stimuli distribution condition, which we refer to as the *stimulus distribution effect*. In the narrow distribution condition, the slope was greater than the wide distribution condition (and the intercept lower).

This result was predicted by Miyazaki et al. (2005) using a simple Bayesian model. The explanation was that the perceived value of Δ_{LED} on any single trial depends on both the current visual information and how credible the participant believed possible values of Δ_{LED} to be a-priori (their prior over the stimulus distribution). When the stimulus distribution prior is Gaussian it will draw perceptual estimates of Δ_{LED} towards the prior's mean by an amount that depends, in part, on the variance of the prior. Thus when a participant experiences a trial in which Δ_{LED} is far from the mean they will make a larger error when the value of Δ_{LED} is less likely a-priori (i.e. during the narrow distribution). Thus the finding of Miyazaki et al. (2005) is consistent with Bayesian integration of visual information and knowledge acquired over previous trials, and suggests that individuals are sensitive to the mean and variance of the Δ_{LED} distribution. Moreover, Miyazaki et al. found that the stimulus distribution effect was observable after a few hundred trials compared to the thousand-plus trials needed within the force amplitude task used by Körding and Wolpert (2004). This suggests that there may be large temporal differences in the rate at which priors are updated across sensorimotor tasks.

The Miyazaki et al. (2005) finding demonstrates that people are able to take account of the variability of the stimulus distribution. However, we wanted to explore whether the result was a function of the sparse visual environment associated with the task. It is possible that participants' estimates were dominated by the learnt priors due to information insufficiency in the 3 LED stimuli, which might not be present in a more ecological timing task with rich visual information. We also wished to determine whether the result would generalise to a more complex interceptive timing task. Interceptive timing tasks have greater degrees of complexity than coincidence timing tasks, as the individual must not only estimate the time that a moving object will arrive at a future destination, but also determine the time it will take them to move the effector to this location. The previous differences in the number of trials taken before the stimulus distribution effect could be detected across tasks (coincidence timing versus force-amplitude matching), meant that it was interesting to see whether different timing tasks would also show different rates of learning. We predicted that the priors may be learnt faster in an interceptive timing task due to the rich information contained in the stimuli.

We were particularly interested in exploring the asymmetric transfers reported previously when moving from a narrow to wide stimulus distribution and vice versa. Both Körding and Wolpert (2004) and Miyazaki et al. (2005) found that learning of the prior was much slower when moving from a wide to narrow stimulus distribution than when moving from a narrow to wide. Miyazaki et al. (2005) suggested that these asymmetries were not predicted by their Bayesian integration model per se, but suggested the asymmetries might be a “universal property of Bayesian integration in sensorimotor learning”.

Miyazaki’s et al. used a very simple Bayesian model which did not provide a mechanism for learning the prior over the target distribution. Instead they simply assumed that participants had learnt a prior that corresponded to the stimulus distribution. Bayesian state-space models provide a mechanism for learning the stimulus distribution, which does predict slower learning when switching from a wide to narrow distribution. The learning model is as follows,

$$\sigma_{\text{dist}_1} \sim \text{Cauchy}^{+}(0, \sigma_{\text{init}}) \tag{3.1}$$

$$\sigma_{\text{dist}_{n+1}} \sim \text{Normal}(\sigma_{\text{dist}_n}, \sigma_{\text{learn}}) \tag{3.2}$$

$$n \in 1 : N - 1 \tag{3.3}$$

$$\mu_i \sim \text{Normal}(M, \sigma_{\text{dist}_i}) \tag{3.4}$$

$$\hat{y_i} \sim \text{Normal}(\mu_i, \sigma) \tag{3.5}$$

$$i \in 1 : N \tag{3.6}$$

The model estimates the standard deviation of the stimulus distribution (σ_{dist}) as taking a random walk over the course of the trials, with each estimate of $\sigma_{\text{dist}_{n+1}}$ being drawn from a normal distribution with a mean centered on σ_{dist_n}. The perceptual estimate of the stimulus ($\hat{y_n}$) are then modelled as a noisy measurement of the true stimulus. This model[1] allows the estimate of σ_{dist} to vary on each trial, but regulates the learning rate using the parameter σ_{learn}.

---

[1] For simplicity M is fixed at the target distribution mean in equation 3.4.

The model captures the slower rate of learning when switching from the wide to the narrow target distribution. The reason for this is that when the stimulus distribution switched from narrow to wide the participant quickly experiences stimuli that are very unlikely under the current estimate of σ_{dist}, which results in the prior quickly updating. When moving from a wide to narrow stimulus distribution the participant experiences stimuli that are still very likely under the current estimate of σ_{dist}, resulting in slower learning. We were therefore interested in determining whether we would find asymmetries similar to those reported by Körding and Wolpert (2004) and Miyazaki et al. (2005) within our interceptive task.

3.2 Methods

3.2.1 Participants

Eleven right handed participants were recruited from the University of Leeds (2 male, 9 female, mean age = 25.46 years, SD age = 4.82). This sample size was based on the numbers recruited in previous studies (Körding & Wolpert, 2004; Miyazaki et al., 2005). Participants did not report any neurological or movement issues and had corrected-to-normal vision. All participants provided informed consent and the study was approved by the School of Psychology Ethics Committee, University of Leeds, UK (Ethics number: 15-0185, date: 17/07/2015).

3.2.2 Apparatus

The experimental hardware consisted of a bespoke 1-DoF manipulandum whose position was captured using a linear potentiometer, allowing for reliable measurements of the manipulandum position when it was moving at high speed. Participants sat at a desk and viewed a BenQ XL2720Z LCD gaming display (Resolution: 1920 × 1080, size: 548 × 642mm, brightness: 300cd m^{-2}, refresh rate: 144Hz), positioned 50 cm in front of them at eye level. The monitor was positioned vertically, and the manipulandum was placed on the desk 30cm to the right of the monitor (see figure 3.1A). The manipulandum controlled an onscreen bat (dimensions: 10 × 15mm; see figure 3.1B). All stimuli were generated in Python 2.7.9 using open source libraries. All coordinates were given in mm with the origin at the bottom left of the screen.

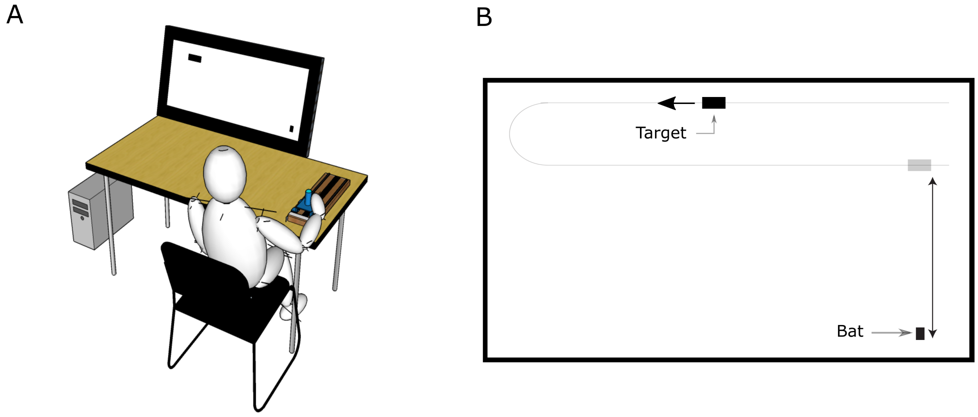

Figure 3.1: A) The experimental setup. The display was positioned in a vertical orientation 50cm in front of the participant's eye line. The manipulandum's position was measured using a linear potentiometer. B) The on screen display. A target moved around an invisible linear track (transparent line). Participants attempted to strike the target's underside using the bat.

3.2.3 Procedure

At the start of every trial, a start box appeared on screen and the participant was instructed to place the bat within it (dimensions: 15×20mm, coordinates: [500, 20mm]) by fully retracting the manipulandum. As soon as the bat was placed within the start box a target appeared (dimensions: 50 × 15mm) at a position that varied over trials. The target's trajectory was chosen so that it would always cross in front of the bat at the same location (coordinates: [500, 150mm]). The targets always moved along the same invisible track (see figure 3.1B) and its starting position on the track was chosen by first selecting the position that resulted in the target travelling for one second before reaching the interception point, which depended on the target's velocity (see below). This was to ensure that the target was visible for the same length of time on average. To try and prevent the target's start position providing a reliable cue of the target's speed, we displaced the target along the track from this position by a distance drawn from a normal distribution, $\text{Normal}(0, 80\text{mm})$, with the constraint that the target must always appear on screen. After the target had been visible for a duration drawn from the uniform distribution, $\text{Uniform}(0.5, 3\text{s})$, the target began to move. Participants were instructed to hit as many targets as possible. They were also told to make the striking action as one continuous movement, not as several smaller movements.

The target was successfully hit if the upper edge of the bat collided with the lower edge of the target. The target then stopped moving, turned red and span before disappearing, thereby providing motivating animated feedback of a successful strike. If the bat passed in front of the target's horizontal path the target immediately stopped moving and then remained on screen for 1 second, indicating a miss. Thus, participants could not simply move the bat in front of the target's path and wait for the target. If the bat crossed the target's path after the target had moved too far to be struck then the target stopped and remained visible for 1 second. The position of the bat and target was timestamped and saved to disk at 500Hz. The bat's positional data was filtered using a low pass second order zero-lag Butterworth filter with a cut off frequency of 10Hz. Spline interpolation was used to estimate the time at which the bat reached the interception point.

Targets moved at a speed that was drawn from a normal distribution, $\mathrm{Normal}(800, \sigma)$, where σ depended on the target speed condition. Participants first tried to hit 100 targets in which σ was 0 (speed was always $800\mathrm{mm\,s^{-1}}$) which acted as a baseline and practice condition. Next the participants were split into two groups. The N-first group ($n = 6$) experienced 200 trials in which the target speeds were first drawn from a narrow distribution in which $\sigma = 80$, followed by 200 trials in which the target speeds were drawn from a wide distribution where $\sigma = 200$. The W-First group ($n = 5$) experienced the wide distribution trials first and the narrow last.

3.2.4 Data analysis

The primary measure of interest was the temporal error (TE), defined as the distance between the centre of the target and centre of the bat, divided by the target speed, at the time that the upper edge of the bat reached the same y-coordinate as the lower edge of the target. TE was calculated for all trials, regardless of whether the target was hit or not. The TE data was modelled using a multilevel Bayesian model. Each participant had an intercept and slope for each target distribution condition (narrow and wide), where target speed was the predictor. The intercept and slope in each target distribution condition were themselves determined by a linear regression with trial number as a predictor. Thus, the intercept and slope were allowed to vary with trial number. Each participant's regression coefficients were distributed by a multivariate normal distribution with a mean vector that depended on the participant's group (N-First vs W-First).

More formally, the ith data point was distributed by a Student's distribution,

$$y_i \sim \text{Student}(\nu, \alpha_{0_{tv[i],p[i]}} + \alpha_{1_{tv[i],p[i]}} x_i, \sigma) \quad i \in 1:N \tag{3.7}$$

where $tv[i]$ and $p[i]$ index the target velocity distribution condition (wide or narrow SD) and participant respectively for the ith data point and x_i is the target speed. Both the participant level slope and intercept coefficients were modelled as depending of the trial number t, where t_1 is the first trial in a given condition,

$$\alpha_{0_{tv[i],p[i]}} = \beta_{1_{tv[i],p[i]}} + \beta_{2_{tv[i],p[i]}} t_i \tag{3.8}$$

$$\alpha_{1_{tv[i],p[i]}} = \beta_{3_{tv[i],p[i]}} + \beta_{4_{tv[i],p[i]}} t_i \tag{3.9}$$

Thus, each participant had eight regression coefficients, four for each target speed condition. The β_1 and β_2 parameters describe how the participant's intercept varies with trial number (t) and the β_3 and β_4 parameters describe how the participant's slope varies with trial number. This model is mathematically equivalent to a model with target speed and trial number as main effects and an interaction term between them.

Each participant's regression coefficients were modelled as a vector (β) of eight parameters (four for each target speed condition) distributed by a multivariate normal distribution,

$$\beta \sim \text{MultiNormal}(U_{p[i]}\gamma, \Sigma) \tag{3.10}$$

where $U_{p[i]}$ is a row vector specifying the group of participant p (N-First or W-First) and γ is a matrix in which each row contains the mean regression coefficients for each group. As recommended in the Stan user manual (Manual, 2013), Σ was calculated as,

$$\Sigma = \text{diagMatrix}(\tau)\Omega\text{diagMatrix}(\tau) \tag{3.11}$$

where τ is a vector of coefficient scales, Ω is the correlation matrix. $diagMatrix(\tau)$ is a square matrix in which all the elements are zero except for the diagonal which is filled with the values of τ. This approach allows us to place priors on the covariance and correlations independently.

3.2.5 Priors and implementation

The priors were set as weakly informative, being informed only by the scale of the data. Weakly informative priors help keep estimates away from clearly nonsensical estimates. Note that the Stan implementation was reparametrized to make it more amenable to the NUTS algorithm (see below for details), but is algebraically equivalent to the model specified here.

3.2.6 Sampling from the posterior

Bayes' rule was used to estimate the credible values of the model parameters (θ) given the data. The joint posterior distribution is given by

$$P(\theta|y) \propto P(y|\theta)P(\theta) \tag{3.12}$$

A representative sample was drawn from the posterior using NUTS (Hoffman & Gelman, 2011) implemented in PyStan 2.14 (Stan Development Team, 2016). Four chains of 5000 samples (warmup $N = 2500$) were started at random values in the joint posterior distribution. Convergence was assessed by visually examining the chains and computing $\hat{R}$ and effective sample size for each parameter.

3.3 Results

The model estimated the group level regression parameters as well as the parameters for individuals. The γ matrix in the model tells us the mean regression parameters where each row corresponds to the parameters for a group (N-First and W-First). For simplicity we can say that each group has a vector of parameters which we now refer to as $\gamma_1 \ldots \gamma_4$. For each group we can calculate the group level mean intercept $E(\alpha_0)$ and slope $E(\alpha_1)$ on trial

t as follows,

$$E(\alpha_{0_t}) = \gamma_1 + \gamma_2 t \tag{3.13}$$

$$E(\alpha_{1_t}) = \gamma_3 + \gamma_4 t \tag{3.14}$$

We can collapse across groups by taking the mean of these parameters over the two groups.

The main hypothesis was that the regression intercept $E(\alpha_0)$ would be lower in the narrow target speed distribution condition compared to the wide condition, and the slope $E(\alpha_1)$ would be greater in the narrow condition than in the wide (the *stimulus distribution effect*). These effects should emerge as the participants learn about the target distribution, thus there should be no difference between target speed conditions on the first trial, but there should be by the last trials.

The expected intercepts and slopes were first examined averaged over the groups. Figure 3.2 shows credible regression lines (calculated from every 10th sample from the posterior) plotted from these intercepts and slopes for the first trial ($t = 1$) and the last trial ($t = 200$) in the narrow and wide target speed distribution conditions. On the first trial no difference was found between the narrow and wide target speed distributions in terms of the intercept (contrast mean = $-7.42\mathrm{mm\,s^{-1}}$, SD = 26.23, 95% HDI = [-58.48, 46.40], $\eta = 0.39$)[2] or slope (contrast mean = $0.01\mathrm{mm\,s^{-1}}$, SD = 0.03, 95% HDI = [-0.06, 0.07], $\eta = 0.59$). Contrasts suggested that differences in

---

[2] see chapter 2 for an explanation of the contrasts

both the intercepts (contrast mean = $-35.63\text{mm}\,\text{s}^{-1}$, SD = 25.91, 95% HDI = [-88.14, 13.71], $\eta = 0.08$) and slopes (contrast mean = $0.044\text{mm}\,\text{s}^{-1}$, SD = 0.03, 95% HDI = [-0.02, 0.11], $\eta = 0.91$) were much more likely in the last trial. However, the η values did not reach the threshold for declaring a 'highly probable' difference (see section 2.2.4), with only 92% and 91% of the respective contrast distributions falling on the predicted side of zero.

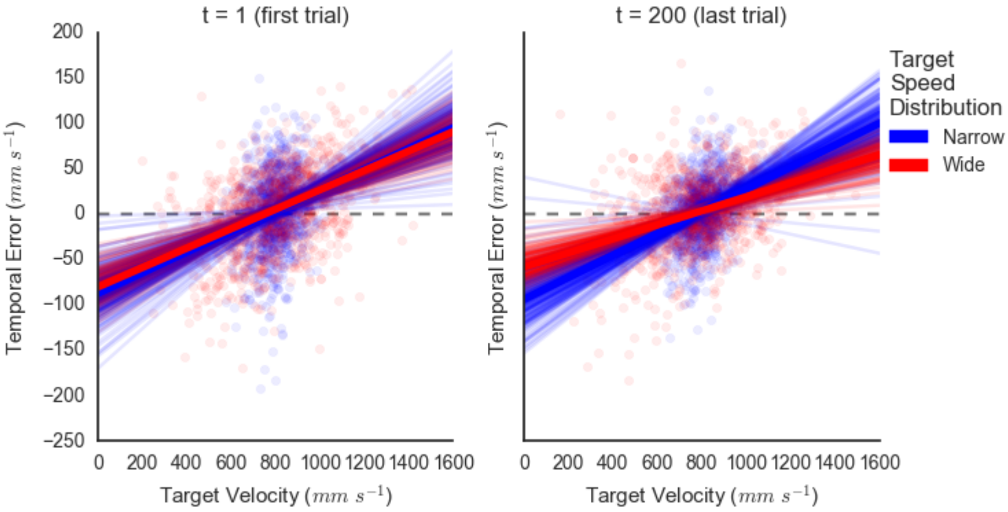

Figure 3.2: Posterior estimates of credible group level regression lines for the first trial ($t = 1$) and the last trial ($t = 200$). On the first trial regression lines were similar between the two target speed distribution conditions. By the last trial the intercept appeared lower and the slope greater in the narrow distribution condition, although the 95% HDI still included zero. The x-intercept is approximately the target velocity mean. Regression lines were plotted using every 50th sample from the posterior distribution. The raw data for the first 100 trials is plotted in the left panel. The raw data for the last 100 trials is plotted in the right panel.

Examining the credible regression lines in Figure 3.2 it could appear that there were greater temporal errors in the narrow condition, because of the steeper slope. However, it is important to remember that target speeds occurred less frequently at values far from the mean in the narrow target distribution condition, by design of the experiment. To examine whether errors were in fact smaller in the narrow target distribution condition, the root mean squared error ($RMSE$) of the temporal errors in the data was calculated. This conformed that temporal errors were smaller in the narrow target distribution condition than the wide (see figure 3.3).

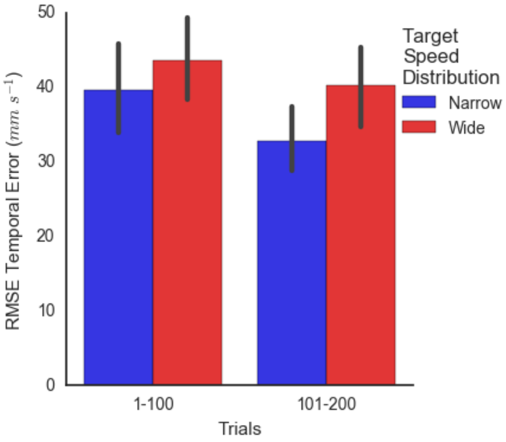

Figure 3.3: The RMSE of temporal errors as a function of trials block (trials 1-100 vs 101-200) and target speed distribution. This plot shows the data, not the posterior distribution. Error bars show the 95% confidence interval.

To explore the stimulus distribution effect further we examined the posterior estimates of the credible regression lines as a function of group (N-First vs W-First), plotting credible regression lines as a function of group, target speed distribution and trial number t (see figure 3.4). On the first trial there was no difference between intercepts in the narrow and wide target speed conditions in either the N-First group (contrast mean = 4.34, SD = 36.66, 95% HDI = [-67.137, 77.70], $\eta = 0.55$) or the W-First group (contrast mean = -18.95, SD = 38.27, 95% HDI = [-95.0, 55.821], $\eta = 0.303$). Likewise, no differences were found between the slopes in the N-First group (contrast mean = 0.008, SD =

0.05, 95% HDI = [-0.086, 0.09], $\eta = 0.515$) or the W-First group (contrast mean = 0.01, SD = 0.05, 95% HDI = [-0.082, 0.103], $\eta = 0.618$).

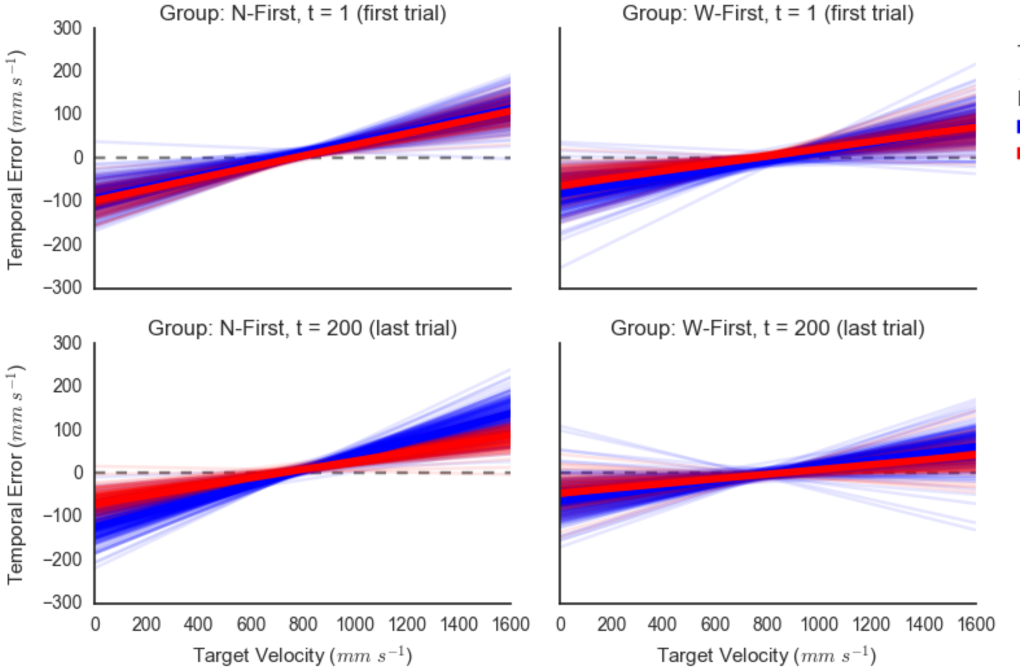

Figure 3.4: Posterior estimates of the expected regression lines for each group (N-First vs W-First) when $t = 1$ and $t = 200$. The left column shows the N-First group and the second shows the W-First group. The x-intercept is approximately at the target velocity mean. When $t = 1$ the plausible regression lines in the Narrow target speed distribution (blue lines) and Wide distribution (red lines) overlap. When $t = 200$ the blue lines show a lower intercept and greater slope than the red lines. The W-First group show more overlap than the N-First group.

By the last trial the N-First group probably had lower intercepts in the narrow target speed condition when compared to the wide (contrast mean = -53.48, SD = 35.6, 95% HDI = [-124.342, 16.652], $\eta = 0.062$) with 93.8% of the posterior mass being below zero $(1 - \eta)$. Again this fell just short of our criteria for a highly probable difference by 1.2%. The W-First group did not show a reliable difference in intercepts between narrow and wide target distribution conditions (contrast mean = -17.87, SD = 38.25, 95% HDI = [-92.942, 56.789], $\eta = 0.315$). The N-First group probably had greater slopes in the narrow condition when compared to the wide (contrast mean = 0.06, SD = 0.045, 95% HDI = [-0.027, 0.0150], $\eta = 0.932$), although η was just short of the 'highly probable' difference criteria (by 1.8%). No reliable difference was found between slopes for the W-First group (contrast mean = 0.023, SD = 0.048, 95% HDI = [-0.074, 0.116], $\eta = 0.693$). The results therefore tentatively suggest a possible stimuli distribution effect when moving from a narrow target speed distribution to a wide one, but not when moving from a wide to narrow distribution.

To examine the rate at which the intercepts and slopes were changing with trial we examined each of the group level parameters $\gamma_1 \ldots \gamma_4$. These group level parameters are the expected value of the regression coefficients (see equations 3.13 and 3.14). The first two parameters (γ_1, γ_2) tell us about how the intercept changes with trial number. These higher level parameters correspond to the intercept and the slope of the lower level intercept parameter (in multilevel models the parameters are themselves modelled as regressions; the lowest level of the model is given by equation 3.12 and the higher level by equations 3.13 and 3.14). The last two (γ_3, γ_4) tell us how the slope changes with trial number. These higher level parameters correspond to the intercept and slope of the lower levels slope parameter.

There was no difference in γ_1 between the narrow and wide conditions for either the N-First (contrast mean = 4.63, SD = 36.77, 95% HDI = [-97.18, 78.5], $\eta = .55$) or W-First group (contrast mean = -18.96, SD = 38.39,

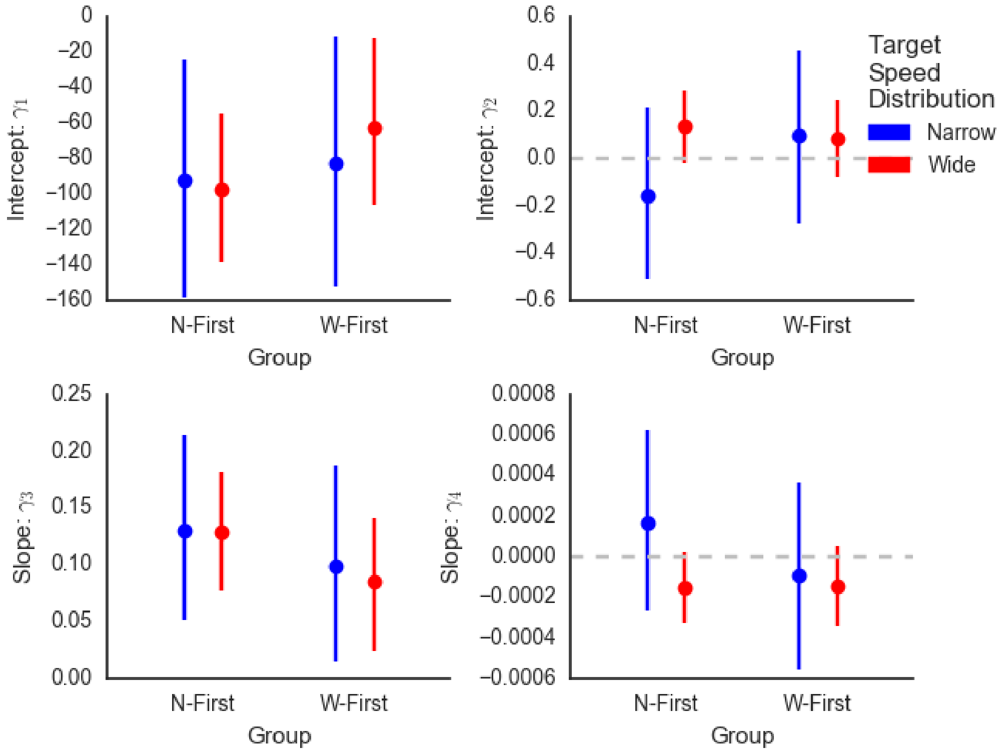

Figure 3.5: Posterior estimates of each γ parameter which correspond to the parameters in equations 3.13 and 3.14. Points show posterior means. Error bars show the posterior 95% HDI.

95% HDI = [-94.30, 56.37], $\eta = .30$). The γ_2 parameter revealed a possible difference between the narrow and wide target distribution conditions in the N-First group (contrast mean = -0.29, SD = 1.92, 95% HDI = [-0.674, 0.11], $\eta = 0.07$), but η just failed to reach the criteria for a 'highly probable' difference (by 2%). The γ_2 plot in figure 3.5 reveals that the intercept may have been decreasing over time in the narrow condition, although a substantial proportion of the marginal distribution spanned zero. However, the intercept was almost certainly increasing in the wide condition. No difference could be reliably detected in the W-First group (contrast mean = 0.01, SD = 0.21, 95% HDI = [-0.41, 0.41], $\eta = .51$).

The γ_3 parameter did not reveal a difference between narrow and wide conditions for either the N-First (contrast mean = 0.0, SD = 0.05, 95% HDI = [-0.09, 0.09], $\eta = .51$) or W-First group (contrast mean = 0.01, SD = 0.05, 95% HDI = [-0.08, 0.13], $\eta = .61$). The γ_4 parameter revealed a small possible difference between narrow and wide conditions of 3.2×10^{-4} (SD = 2.4×10^{-4}, 95% HDI = $[-1.94 \times 10^{-4}, 7.72 \times 10^{-5}]$, $\eta = 0.9$) in the N-First group, but again this did not reach the η criteria. The γ_4 plot in figure 3.5 reveals that the slope may have been increasing in the narrow condition, but decreasing in the wide condition. No differences in the γ_4 parameter could be found between the narrow and wide condition for the W-First group (contrast mean = 0.0, SD = $2.0times10^{-4}$, 95% HDI = $[-4.0 \times 10^{-3}, 6.0 \times 10^{-4}]$, $\eta = .57$).

3.3.1 Power analysis

While key contrasts suggested an asymmetrical stimulus distribution effect existed, whereby the differences in slope and intercepts between the target speed distribution conditions were probably larger in the group who experienced the narrow distribution condition first, their η values failed to reach our criteria for a 'highly probable' difference. Specifically, for the narrow first group, the η value fell short by 1.2% for the intercepts and 1.8% for the slopes. We attempted to estimate how many participants would be required to ensure that the hypothesised effect would be detected with high probability (power analysis).

In general Bayesian power analysis proceeds by simulating random samples of data from a generative statistical model, using hypothesised parameter values (in this case the marginal posterior means for each parameter was taken as a best guess of the model parameters; which suggested a asymmetrical stimulus distribution effect). The model with hypothesised parameter values can then be used to generate datasets of a fixed sample size (N). For each sample of size N the model can be refit and the key contrasts of interest repeated to assess whether the criteria for a highly probable difference was met (i.e. $\eta < 0.05$ or $\eta > 0.95$). This process can be repeated many times and the proportion of contrasts that meet the criteria for a high probable difference provides an estimate of power (Kruschke, 2015). By doing this for many different sample sizes it is theoretically possible to determine how many participants are required to obtain the desired statistical power (given the effect size and measurement noise).

Unfortunately this approach to power analysis is frustrated by the time required to repeatedly refit Bayesian models using MCMC. Hierarchical models can be particularly slow to fit, and in the current case refitting the model takes approximately 12 hours. Thus fitting the model for multiple sample sizes a sufficient number of times to provide a stable power estimate is not practically viable. Thus a different approach was taken in which five data samples, with $N = 10$ for each group (total N = 20) were generated. The model was then refit using these samples and the key contrasts were repeated (i.e. the four contrasts comparing the intercepts and slopes in both groups; see figure 3.4). The slope and intercept contrasts for the narrow first group reached the $\eta < 0.05$ (or $\eta > 0.95$) criteria in four out of the five simulations. Figure 3.6 shows the probability of obtaining $\eta < 0.05$ for four out of five samples as a function of statistical power. It is clear that high power is most probable, with the maximum likelihood estimate of 79% power. Thus it seems reasonable to suggest that an N of 10 would be highly likely to detect the hypothesised effect.

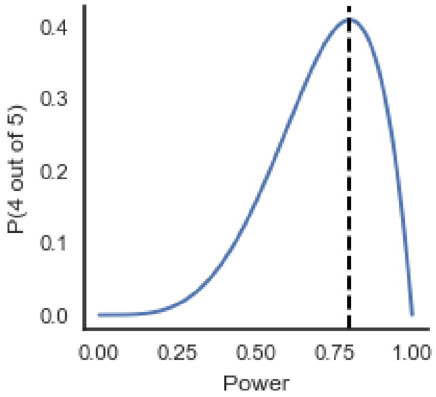

Figure 3.6: Probability of obtaining $\eta < 0.05$ for four out of the five simulated data sets when $N = 10$ for each of the two groups. The MLE of statistical power is 79% as indicated by the dotted vertical line. Probabilities were calculated using the binomial distribution.

3.4 Discussion

We explored whether an interceptive timing task would show evidence of behaviour being affected by the structure of the trial distribution (as previously found within a coincidence timing task). Participants were exposed to a narrow and wide distribution of target speeds and the resulting relationship between target speed and temporal error was examined. The prediction (from Körding and Wolpert (2004) and Miyazaki et al. (2005)) was that the slope would be steeper in the narrow speed distribution condition (and the intercept lower), as inferences are drawn more towards the mean of the prior when the standard deviation of the prior is smaller (the stimulus distribution effect). Unfortunately the η values for key comparisons did not reach the criteria for declaring an effect with high certainty, generally missing this criteria by a few percentage points. Thus the results suggest that a stimulus distribution effect was very likely present, but further experimental evidence may be required to achieve a high degree of certainty in the result. Power analysis suggested that 10 participants in each group would provide sufficient power.

Overall, the pattern of results tentatively suggests that individuals combine online visual information with knowledge acquired over previous trials when performing an interceptive task. This extends the findings of Miyazaki et al. (2005) and shows that these effects are most likely present in interceptive as well as coincidence timing tasks. The findings of both experiments are consistent with the interpretation that the brain acts as if it were a Bayesian operator providing optimal estimates of object arrival. The notion that the brain acts as a Bayesian operator can explain work by Brouwer, Brenner, and Smeets (2002) who asked participants to hit targets that moved at three speeds, while occluding the targets for varying lengths of time. When occlusions were long in duration, participants tended to hit ahead of the slow targets and behind the fast targets. However, this effect greatly reduced for short occlusions, with a decrease in hitting errors. This is an expected property of an optimal Bayesian operator. When occlusion times are long (i.e. view times are shorter), visual information will contain more uncertainty (i.e. the likelihood will be less precise) and thus the prior will dominate the posterior estimates of the target's velocity, making it appear as though velocity information is ignored. When the occlusion times are shorter, the likelihood will be more precise and will therefore dominate the posterior. Thus, the findings of Brouwer et al. (2002) compliment the current study, where, in Bayesian terms, we manipulated the prior whereas they manipulated the likelihood.

We also explored whether the order in which participants experienced the two target speed distributions influenced the strength of the stimulus distribution effect. We found that the differences in slope (and intercept) between target speed distribution conditions was probably stronger for the group who experienced the narrow distribution condition first, although again this effect just fell short of our η criteria required to declare a ‘highly probable’ effect (93.8% of the contrasts distribution was on the predicted size of zero). Despite this short coming, the most probable interpretation of the results is consistent with those of Miyazaki et al. (2005), who reported asymmetries when participants moved from narrow-to-wide versus wide-to-narrow distributions. In the Miyazaki et al. (2005) study, statistically significant differences were only observed when calculating the slope using a moving window that included all trials between the 241st and 480th trial. This may explain the weak effect in wide-to-narrow group from the current study where only 200 trials were used, as it may take longer for people to update their prior when moving from a wide distribution to a narrow distribution. Thus the 200 trials may have been insufficient for participants to update their prior over the target speed distribution when moving from the narrow to wide distribution.

The likely asymmetry in the learning of the prior over the stimulus distribution is consistent with a Bayesian state space model which updates estimates of the standard deviation of the stimulus distribution on a trial by trial basis. As outlined in the introduction learning is expected to be faster when moving from a narrow to wide stimulus distribution, as the new distribution suddenly produces stimuli that are very unexpected under the current prior. When moving from the wide to the narrow distribution, participants experience stimuli that are still consistent under the current prior, thus more trials are needed to learn that the stimulus distribution has changed.

Critical questions still remain about the nature of the mechanisms that underpin Bayesian inference. Our data suggests that individuals are probably able to learn new prior distributions in interceptive timing tasks within relatively few trials (≈ 200), a finding that has been found across a range of perceptual and motor tasks (Vilares & Kording, 2011). In contrast, several psychophysical phenomena related to speed perception, including the finding that low contrast stimuli appear to move slower than high contrast stimuli (Thompson, 1982), can be explained by assuming that people use a single prior that favours slower speeds (Freeman, Champion, & Warren, 2010; Stocker & Simoncelli, 2006). This appears to be incompatible with the finding that people rapidly update their priors over the course of a short experiment. However, more complex hierarchical Bayesian models (HBM) can readily account for both of these findings. HBMs are models which posit probability spaces over probability spaces, with priors upon priors. Each level of the model places a probability distribution on the level below (Tenenbaum, Kemp, Griffiths, & Goodman, 2011). For example, a HBM may have a prior that objects move slowly in the world, which constrains lower level priors regarding how individual objects move. A HBM may model the relationship between the speeds of different objects, for example, by learning that different objects have different speeds but their speeds tend to vary within a certain range of parameters. This confers two major benefits. Firstly, it allows for optimal transfer of learning from one object to another, and secondly it allows for abstract knowledge to emerge from lower levels of the hierarchy i.e. that most objects move slowly.

It should be highlighted that while the data are consistent with the brain acting as a Bayesian operator they do not necessarily mean that the system is actually functioning as a Bayesian operator (i.e. the system may not actually calculate or represent probabilities). One feature of any successful learning machine is that previously successful outputs influence subsequent outputs. It is possible that the learning mechanisms of the brain result in processes of discriminative classification. For example, it has been shown that under the right conditions, simple Hebbian learning in neural networks can approximate normative Bayesian models (Verstynen & Sabes, 2011). Thus behaviour that is near Bayesian optimal may be possible without the need to calculate or represent probabilities.

A clear weakness of the study is that critical contrasts did not reach our criteria for declaring an effect ‘highly probable’, generally falling short by a few percentage points. Thus while the results tentatively support our hypothesis, more data or replication of the study may be required to satisfy a strict threshold for declaring an effect. A major advantage of Bayesian estimation over frequentist approaches is that the posterior distribution from a study can be used as the prior for a replication, or the results of multiple studies can be combined using a multilevel model (Marsman et al., 2017). Thus Bayesian estimation provides a natural approach to cumulatively improving estimates of effects.

In conclusion, we have extended previous reports of individuals combining visual information with knowledge acquired over previous trials in a coincidence timing task, to show that the same effects can probably be seen in an interceptive task. The study also suggests that the previous observations of asymmetric transfer, where the effects are larger when participants move from a narrow-to-wide distribution rather than vice versa, are also present in interceptive timing tasks. The general pattern of behaviour is consistent with the brain acting as if it were a Bayesian operator. However, it may be that fundamental motor learning processes create Bayesian optimal behaviour but without the need to explicitly calculate or represent probabilities. Having explored how adults deal with uncertainty and noise in perceptual quantities, in the next chapter we move on to examine how adults make optimal choices in motor planning to reduce the effects of motor noise on task outcomes.

Chapter 4

The selection of ‘go faster strikes’ as a function of noise in interceptive timing tasks

Chapter Abstract

When intercepting moving targets people make temporal errors, sometimes arriving too early or late at the spatial location at which they intended to make the interception. Over repeated trials this pattern of errors can be described in terms of accuracy (the mean error) and precision (the reciprocal of the standard deviation of errors). Many interceptive tasks require excellent temporal precision because the window of time in which the target can be struck is in the order of milliseconds. Briefer, faster movements are associated with better temporal precision but these movements have associated costs, including greater energy expenditure. Theories of optimal control (OC) suggest that an optimal controller that minimises both temporal errors and

energy costs will select movements that have the longest possible duration whilst achieving sufficient temporal precision. An experimentally imposed perturbation that decreases temporal precision (but not accuracy) would push the system to increase movement speed if it is: (i) sensitive to the distribution of temporal errors over previous trials and (ii) using an OC strategy that attempts to minimise temporal errors and energy expenditure. To test whether this is the case, three groups of adults used a one degree of freedom manipulandum to hit virtual 'pucks' at computer generated moving targets. The puck moved at its strike speed plus noise drawn from a Gaussian distribution. The first group experienced no added noise and the results showed that their strike velocities decreased over trials. The second group experienced low noise and maintained reasonably constant strike velocity across trials. The third group were exposed to high noise and increased strike velocity over time. These data indicate that adults monitor the prior distribution of their temporal errors in an interceptive task and use this information for OC.

4.1 Introduction

Interceptive timing actions (e.g. striking a baseball) require the intercepting effector (e.g. hand, bat) to be in the right place (the same spatial location as the target object) at the right time. Spatial errors occur when the intercepting effector is moved to the wrong physical location whilst temporal errors arise when the effector arrives too early or too late at the correct spatial position. Both spatial and temporal errors occur for a number of reasons, including perceptual misjudgements of the target's current position, misestimates of the target's speed, and inaccuracies in predicting the target's future location. Even when perceptual estimates are veridical, errors may still occur in the execution of the movement due to noise in the motor system. It is often difficult to decouple spatial and temporal errors as there is no simple way to determine whether someone missed a target because their movement did not reach the planned interception location, or because they reached the planned location at the wrong time. However, by constraining movements to one degree of freedom, temporal errors can largely be isolated from spatial errors, since as long as the movement is of sufficient amplitude they will reach the spatial location at which an interception can take place. This allows for temporal accuracy and precision to be investigated in humans.

Skilled interceptive behaviours require individuals to learn how to minimise temporal errors. One possible strategy to minimise temporal errors is to use online feedback control where errors are detected and corrected within the course of the movement. There are a number of models of interceptive timing control that continuously correct for errors between the estimated time to arrival (TTA) of the target and estimated TTA of the effector at the planned point of interception (see Tresilian (2005) for a review). These models are feasible for intercepting slow moving targets but are not viable for fast interceptive actions, owing to the time delays that are inherent at every level of the sensorimotor system (Nijhawan, 2008). Thus fast interceptive actions may require ballistic movements without the aid of online corrections during a single movement.

Temporal errors can be minimised by using signed error signals to correct for the errors made on previous trials. Standard models of sensorimotor learning suggest that on trial n the expected sensory consequences ($\hat{y}_n$) of an action are predicted and compared to the true sensory consequences (y_n). An error is then generated by comparing the predicted and actual sensory consequences, where the error (e) is given by $e_n = y_n - \hat{y}_n$. The next action is then updated based on the prediction error and the system's learning rate. The ability to learn from signed error signals in this manner can explain the acquisition of a wide range of behaviours, including adaptation to sensory and motor perturbations (van Beers, 2009). This mechanism can explain how humans reduce their mean temporal error (improve temporal accuracy) but cannot explain how improvements are made in temporal precision. The limitation of such models is that they do not specify how the system selects the optimal movement when there is a choice of trajectories that will allow the participant to hit the target. It is often possible to hit moving targets using a wide range of combinations of movement (MT) and initiation times (IT). MT refers to the duration of the movement from its onset until the effector reaches the planned interception point, whilst IT is the time at which the movement starts relative to the time at which the target will reach the planned point of interception. In a 1-DoF interceptive timing task a perfect interception (zero temporal error) will then occur when $MT = IT$. This specifies the solution manifold, as any value of MT is viable so long as MT is shorter than the target's movement time (see figure 4.1).

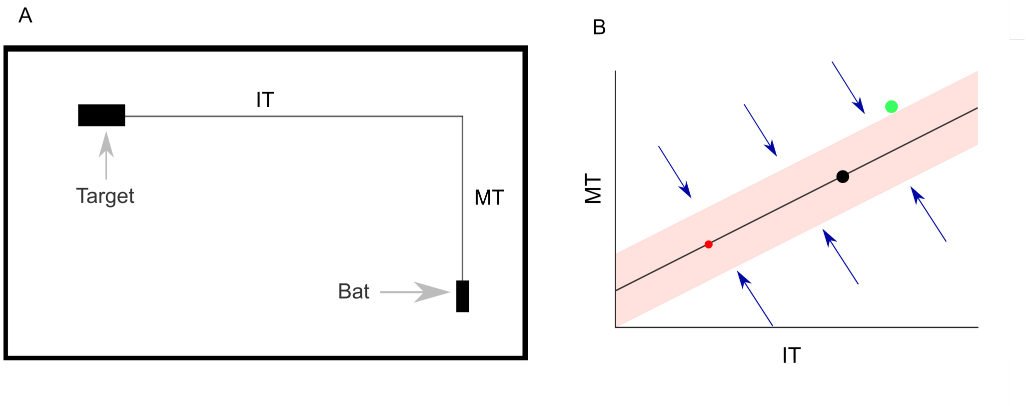

Figure 4.1: A) A 1-DoF interceptive timing task. The target moves from left to right across a screen while a participant attempts to hit the lower edge of the target with a bat that can only move orthogonal to the target's trajectory. As soon as the bat begins to move the IT can be defined as the target's time to arrival at the interception point (red dot). If $MT = IT$ then the center of the bat will strike the center of the target (zero temporal error). B) The solution manifold is specified by $MT = IT$ for any value of MT. The shaded region defines the deviation from zero temporal error that will still allow the target to be hit. Even when an appropriate combination of MT and IT are planned, errors in the movement's execution result in variability in the temporal errors. Smaller values of MT are associated with better temporal precision, but increased energy costs. The green dot shows a planned combination of MT and IT that result in the target being missed on average when errors are Gaussian and centered around the planned values of MT and IT. The black dot shows a solution that results in zero temporal error on average. The optimal solution (red dot) results in zero temporal error on average and falls along the solution manifold in a position that maximises temporal precision whilst minimising effort.

OC models suggest that movements are selected from viable alternatives to minimise a mixed cost function - a function that specifies costs associated with both errors and energy (Todorov & Jordan, 2002; Todorov, 2004). This raises the issue of the rewards and penalties associated with the choice of the trajectory selected from the solution manifold in interceptive tasks. It appears that faster movements have improved precision (Newell, Carlton, & Kim, 1994; Tresilian, 2012), suggesting there are benefits to selecting shorter duration movements. There are several explanations for the improved precision with decreased duration. First, shorter MTs allows the target to be viewed for longer before the movement is initiated, which may result in better estimates of the target object's TTA at the planned interception point. Secondly, when movements are very fast even large deviations in the planned movement speed result in small changes in MT. Thus errors in the movement execution have less of an impact on temporal errors. However, moving faster has associated penalties. First, faster movements provide less possibility for any modification of the movement on the basis of online error detection. This appears to be a consideration in timing behaviour as increasing spatial accuracy constraints causes a decrease in temporal precision within a two degree of freedom timing task (Tresilian, Plooy, & Marinovic, 2009). Second, faster movements tend to have higher energy demands (Todorov, 2004). The rewards and penalties associated with shorter duration movements suggest that optimal performance will involve selection of faster movements when the temporal precision requirements of a task are high, but slower movements when the temporal constraints are reduced.

In a series of studies, Tresilian and colleagues examined participants' MTs when using a one degree of freedom manipulandum to intercept targets of different widths (Tresilian et al., 2004) and targets moving at different speeds (Tresilian & Lonergan, 2002). It was found that MT was shorter when the target speed was faster, even after controlling for the target's view time (Tresilian & Lonergan, 2002). It was also found that MT was shorter when the targets were narrower (Tresilian et al., 2004). These findings are consistent with results from experiments using a two degree of freedom interceptive task (Brenner & Smeets, 2015b). It was later shown that the changes in MT were directly related to the time window (the period in which the target can be struck), which is determined in part by the target's speed and size (Tresilian & Houseman, 2005). Critically, the temporal constraints of the task are relaxed when the time window is larger, so larger temporal errors can still result in the target being hit. These results therefore suggest that skilled adult performers are attempting to minimise multiple costs (e.g. error, effort and energy) as predicted by OC theory. In addition, adults appear to only attempt to reduce their errors when they interfere with the task goal of hitting the target, making slower, more variable movements when the time window is larger.

Thus, the existing empirical evidence indicates that skilled performers use OC when intercepting moving targets. It is not clear, however, whether adults have learnt this optimal behaviour over a large time period (i.e. whether participants have learned to adjust their movement times in response to the stimulus over many months and years of experience) or whether adult humans have the flexibility to implement optimal strategies within the time course of a discrete task (e.g. during the course of a tennis match). The ability to adopt such a strategy within a task would require participants to monitor their distribution of errors over the duration of the task. There is some evidence that participants are sensitive to the distribution of target speeds within an interceptive task (see chapter 3), so it is plausible that participants might be likewise sensitive to the distribution of their temporal errors over a series of trials within a relatively short time period.

In this study we wished to test whether participants were sensitive to the distribution of their temporal errors over previous trials when planning their movements. OC predicts that adults should adjust the timing of their movements in response to an increase in temporal variability (a decrease in precision). We therefore examined whether people would adjust the timing of their interceptive movements in response to externally imposed execution noise (noise which increased the variability but not the mean accuracy of the temporal errors). We reasoned that the increased noise would worsen performance but participants would be unable to improve their performance by simple error correction mechanisms (because the average accuracy was unaffected by the addition of the noise). It would be possible, however, for participants to improve performance by increasing the speed of their movements (Tresilian, 2012), because faster movements are associated with greater temporal precision. This allowed us to test the hypothesis that participants use knowledge of their previous temporal errors within an interceptive task to select the optimal movement duration (where shorter durations improve precision but at the cost of higher energy expenditure).

We used a simple interceptive timing task in which participants struck virtual pucks with a bat, launching them towards moving computer generated targets travelling orthogonal to the puck's path. Three groups of participants hit pucks towards targets that were identical in size and speed. On average, the puck moved at the strike velocity, but we added Gaussian noise to the puck's velocity in two of the groups: a *high noise* and a *low noise* group. The remaining group experience no added noise (*no noise* group). If participants used OC then we would expect the *no noise* group to decrease their strike velocities as the experiment progressed (as they became more skilful through practising the task) whereas the *high noise* group should increase their strike velocities (to counteract the decrements in performance created through the addition of the noise). The *low noise* group's behaviour would be expected to fall between the other two groups.

4.2 Methods

4.2.1 Participants

Twenty six participants were recruited from the University of Leeds (4 male, 22 female, mean age = 27.21 years, SD age = 4.98). Participants did not report any neurological or movement problems and had normal or corrected-to-normal vision. All participants provided informed consent and the study was approved by the School of Psychology Ethics Committee, University of Leeds, UK (Ethics number: 16-0119, date: 19/04/2016).

4.2.2 Apparatus

All dimensions and coordinates are given in millimeters. A custom built 1-DoF manipulandum was used to control an onscreen bat. A linear potentiometer and a NI-DAQ device was used to capture the manipulandum's displacement along a linear track at 500Hz. Participants sat at a desk and viewed a BenQ XL2720Z LCD gaming display (Resolution: 1920 × 1080, size: 548 × 642mm, brightness: 300cd m^{-2}, refresh rate: 144Hz), positioned 50cm in front of them at eye level. The monitor was positioned vertically, and the manipulandum was placed on the desk to the right of the monitor (see figure 4.2A). The position of the manipulandum controlled an onscreen bat (dimensions: 10 × 15mm). All stimuli were generated using Python 2.7.9 using open source libraries. All coordinates were given in mm with the origin at the bottom left of the screen.

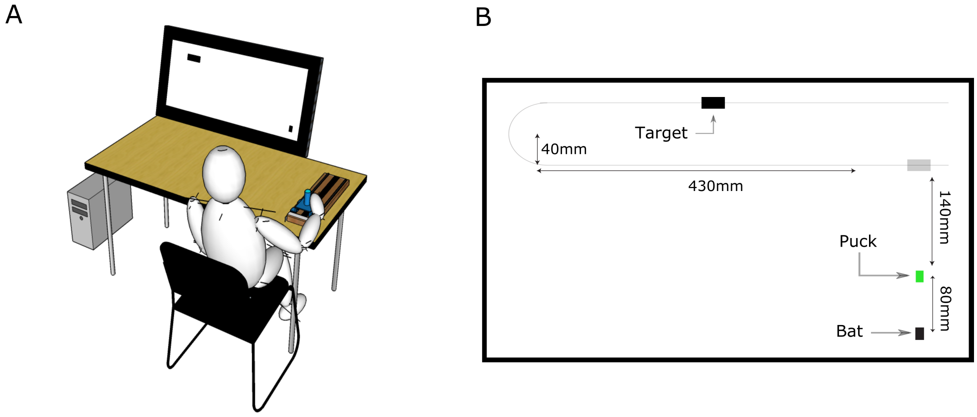

Figure 4.2: The experimental setup. A) Participants used an onscreen bat to hit pucks towards moving targets. The display was positioned in a vertical orientation and the manipulandum's position was measured using a linear potentiometer. B) The stimulus. A target moved around an invisible track to increase the viewing time.

4.2.3 Procedure

At the beginning of every trial a start box (dimensions: 15×20, coordinates: $[500, 20]$) appeared onscreen and the participants were instructed to place the bat within it by fully retracting the manipulandum. The x-coordinate of the bat was always 500, and its y-coordinate depended on the position of the manipulandum. The displacement of the manipulandum resulted in an identical displacement of the bat along the y-axis. Once the bat was placed within the start box a rectangle representing a puck (dimensions: 10, coordinates: $[500, 100]$) then appeared directly above the bat. Participants could hit the puck with the bat, which upon being struck would launch up the screen, at a velocity that depended on the strike velocity and the condition (see below). The participants launched the puck to hit the target (dimensions: 50×15, coordinates: $[264.34, 320]$) that was displayed on the screen. There was a delay before the target began to move along an invisible track at 500mm/s. The delay was drawn from a uniform distribution, Uniform$(0.5, 3.0\text{s})$. The centre of the target passed in front of the bat after moving 750mm, at the coordinates $[500, 240]$. The track (see figure 4.2) enabled the target to be viewed for 1.5 seconds.

Participants were instructed to strike the puck in order to launch it at the target, and were instructed to hit the target with the puck as many times as possible. The puck successfully hit the target if its upper edge collided with the lower edge of the target. The target then stopped moving, turned red and span before disappearing, thereby providing motivating animated feedback. If the bat passed in front of the target's horizontal path the target immediately stopped moving and then remained on screen for 1 second. Thus, participants could not simply move the bat in front of the target's path and wait for the target. The position of the bat was timestamped and saved to disk at 500Hz. The bat's positional data were filtered using a low pass second order zero-lag Butterworth filter with a cut off frequency of 10Hz. Spline interpolation was used to estimate the time at which the bat struck the puck.

Participants were split into three groups, a *no noise* ($n = 8$), *low noise* ($n = 9$), and *high noise* ($n = 9$) group. All participants completed 100 pre-noise trials in which they attempted to hit the target. For these trials the puck always moved at the same speed that it was struck at (the strike velocity). After the pre-noise trials, all participants completed 200 noise trials. For the no noise group these trials were identical to the baseline. For the low noise and high noise groups the puck moved at a speed drawn from the normal distribution, Normal(μ_{strike}, σ), where μ_{strike} was the strike velocity and σ was 100 for the *low noise* group and σ was 200 for the *high noise* group. Following these 200 trials, all participants completed a further 100 post-noise trials that were identical to those in the pre-noise condition.

4.2.4 Data analysis

Two performance measures were examined. The first was the strike velocity (SV) - the velocity at which the bat was travelling when it struck the puck. The second was whether the target was successfully hit or not, which was used to estimate the probability of hitting the targets (P-Hit). The SV was modelled using a multilevel Bayesian model, adapted from the Stan user manual (Manual, 2013). Each participant was modelled as having a mean SV for the pre-noise and post-noise condition. The noise trials were split into blocks of 50, and each participant was modelled as having four mean SVs, one for each block. The means for each participant were distributed through a multivariate Gaussian, with a mean vector that depended on the participant's group (*no noise*, *low noise* or *high noise*).

Strike velocity model

More formally we modelled each data point y_i as belonging to participant p. A row vector x_i specified the trial block of the ith data point. The likelihood function for datum y_i is then given by,

$$y_i \sim \text{Normal}(x_i \beta_{p[i]}, \sigma) \quad i \in 1 : N \tag{4.1}$$

where $\beta_{p[i]}$ is a vector of regression coefficients for the participant p who generated the ith data point and N is the total number of data points. Each participant's vector of regression coefficients was distributed by a multivariate normal distribution,

$$\beta_p \sim \text{MultiNormal}(M_{g[p]}, \Sigma), \quad p \in 1 : P \tag{4.2}$$

where $M_{g[p]}$ is mean vector of coefficients for the group g to which participant p belongs. The covariance matrix Σ is specified as follows,

$$\tau_k \sim \text{Cauchy}^+(0, 500), \quad k \in 1 : K \tag{4.3}$$

$$\Omega \sim \text{LKJCorr}(2) \tag{4.4}$$

$$\Sigma = \text{diag\_matrix}(\tau)\Omega\text{diag\_matrix}(\tau) \tag{4.5}$$

where K is the number of regression coefficients. This formulation allows the prior over Σ to be specified as a correlation matrix (Ω) and vector of scale parameters (τ). The prior distribution over the correlation matrix ($LKJCorr$) is defined in the Stan user manual (Manual, 2013).

This model was re-parametrised to make it easier for the Hamiltonian Monte Carlo algorithm used by Stan to explore, but is equivalent to the model here.

Probability of target hit model

We were also interested in the probability of hitting the target (P-Hit), which provided a simple measure of performance. This was modelled in the same way as SV except the likelihood function (see equation 4.2) was replaced with a Bernoulli distribution with a logistic link function, to account for the fact that the outcome variable was discrete (either a hit or a miss).

$$\mu = x_n \beta_{p[i]} \tag{4.6}$$

$$\text{logit}^{-1}(\mu) = \frac{1}{1 + \exp(-\mu)} \tag{4.7}$$

$$y \sim \text{Bernoulli}(\text{logit}^{-1}(\mu)) \tag{4.8}$$

The priors were chosen to be weakly informative, being informed only by the scale of the data. Bayes' rule was used to estimate the credible values of the model parameters (θ) given the data (y). The joint posterior distribution is given by $P(\theta|y) \propto P(y|\theta)P(\theta)$.

A representative sample was drawn from the posterior using the "No-U-Turn sampler" (Hoffman & Gelman, 2011) implemented in PyStan 2.14 (Stan Development Team, 2016). Four chains of 5000 samples (warmup $N = 2500$) were started at random values in the joint posterior distribution. Convergence was assessed by visually examining the chains and computing $\hat{R}$ and effective sample size for each parameter.

4.3 Results

Strike Velocity

Figure 4.3 shows the posterior estimates of the group means as a function of the trial block and noise group, where the grey shaded area shows the added noise blocks. The *no noise* group possibly showed a slower mean strike velocity at baseline than the *high noise* (contrast mean = $-166.79\,\text{mm}\,\text{s}^{-1}$, SD

= 114.65, 95% HDI = [-391.90, 54.93], 50% HDI: [-242.73, -91.61], $\eta = 0.07$) and *low noise* (contrast mean = $-122.87\mathrm{mm\,s^{-1}}$, SD = 116.30, 95% HDI = [-344.57, 113.36], 50% HDI: [-197.07, -46.28], $\eta = .14$) group, although the η values did not reach the criteria for a 'highly probable' difference. Despite this the posterior means suggested that the differences between the groups were potentially quite large. The difference between the low and high noise groups at baseline was most likely quite small in comparison (contrast mean = $-43.92\mathrm{mm\,s^{-1}}$, SD = 112.90, 95% HDI = [-273.36, 175.11], 50% HDI: [-119.72, 26.98], $\eta = .35$), although the large posterior SD means that there was high uncertainty in this estimate (i.e. large differences between these groups were not improbable).

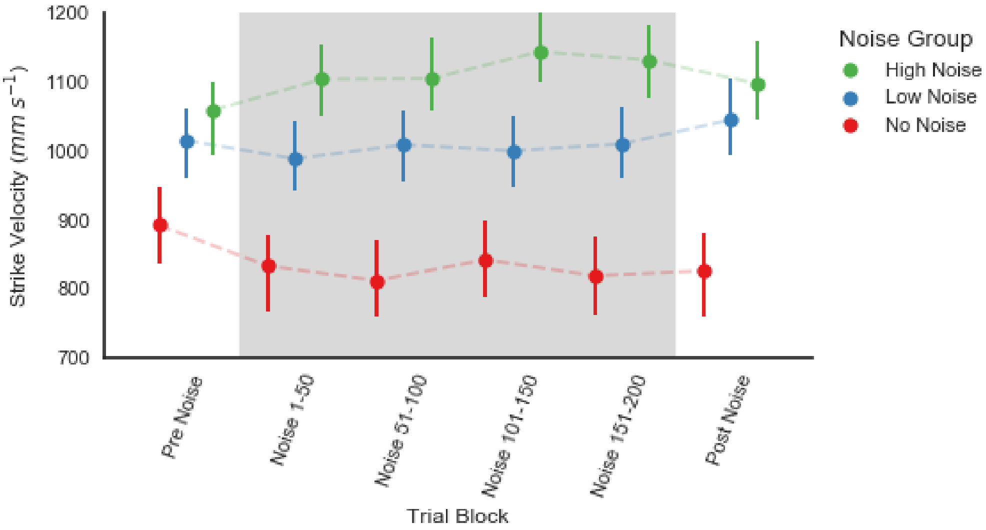

Figure 4.3: Posterior estimates of the group means as a function of noise group and trial block. Points show posterior means and error bars show the 50% HDI.

We predicted that the no noise group would decrease their strike velocity

Table 4.1: *No Noise* group. Contrasts between added noise trial blocks and baseline. Note that the noise trials had zero added noise for the no noise group. The operator :: indicates a contrast.

Contrast: SV	Mean	SD	95% HDI	50% HDI	η
Noise 1-50 :: BL	-59.00	51.29	-157.85, 42.12	-89.02, -21.15	.13
Noise 51-100 :: BL	-80.26	53.00	-187.44, 22.11	-159.35, -45.99	.07
Noise 101-150 :: BL	-50.10	53.87	-159.37, 50.50	-82.08, -10.50	.18
Noise 151-200 :: BL	-76.27	62.46	-200.90, 45.61	-110.01, -26.69	.12
Post-Noise :: BL	–66.26	72.84	-208.40, 79.44	-112.85, -19.03	.18

over time, so we compared the baseline conditions to each block of 50 trials post baseline. Contrasts failed to reach the η criterion for a 'highly probable' effect (see table 4.1). However the 51-100 trial block most likely had a mean strike velocity lower than the baseline, with the η value falling short of our strict criteria by 2% (i.e. there was a 93% chance of an effect in the hypothesised direction). The other comparisons were less convincing, with a large amount of the contrasts distribution spanning zero.

For the *low noise* group, there was no clear change in strike velocity between the baseline block and any of the other blocks (see table 4.2). The 95% and 50% HDI always spanned zero, although the contrast mean and SD indicated that differences of a reasonable magnitude were possible, but less likely than differences of a low magnitude. This was because the means of the contrast's distributions were centered around low values, but their standard deviations (and 95% HDIs) were reasonably large.

In the *high noise* group, the contrasts suggested that the average strike velocity increased compared to baseline. For the first and second noise block the η value failed to reach the threshold, but a 'highly' probable difference

Table 4.2: *Low Noise* group. Contrasts between added noise trial blocks and baseline.

Contrast: SV	Mean	SD	95% HDI	50% HDI	η
Noise 1-50 :: BL	-26.52	49.75	-123.98, 71.14	-56.43, 9.54	.30
Noise 51-100 :: BL	-6.27	52.70	-113.78, 97.36	-39.62, 29.79	.45
Noise 101-150 :: BL	-14.85	52.23	-113.32, 91.99	-50.34, 18.00	.38
Noise 151-200 :: BL	-5.12	61.33	-130.94, 111.19	-48.141, 32.02	.46
Post-Noise :: BL	29.80	71.08	-117.62, 163.92	-23.48, 70.52	.66

Table 4.3: *High Noise* group. Contrasts between added noise trial blocks and baseline.

Contrast: SV	Mean	SD	95% HDI	50% HDI	η
Noise 1-50 :: BL	44.78	48.76	-49.82, 143.00	9.94, 74.82	.82
Noise 51-100 :: BL	45.71	51.18	-55.47, 147.36	8.65, 77.60	.82
Noise 101-150 :: BL	85.01	50.70	-18.97, 183.07	49.27, 115.33	.95
Noise 151-200 :: BL	72.61	60.16	-46.68, 192.62	33.31, 112.53	.89
Post-Noise :: BL	38.20	70.50	-104.98, 173.25	-9.19, 83.70	.71

was observed between the baseline and the 101-150 noise block (see table 4.3), suggesting that strike velocity increased over the trials.

4.3.1 Probability of hitting the target

Figure 4.4 shows the mean probability of hitting the targets as a function of group and trial block. At baseline, the mean probability of hitting the targets may have been lower in the no noise group than in the low noise group (contrast mean = -0.06, SD = 0.06, 95% HDI = [-0.19, 0.07], 50% HDI: [-0.11, -0.02], η = .18) and the high noise group (contrast mean = -0.05, SD = 0.066, 95% HDI = [-0.18, 0.07], 50% HDI: [-0.09, -0.00], η = .24), although there was high uncertainty in the contrasts and the η value did not reach the criteria for a 'highly probably' effect. There was no evidence of a difference between the low noise and high noise groups at baseline (contrast mean = -0.01, SD = 0.06, 95% HDI = [-0.11, 0.13], 50% HDI: [-0.03, -0.04], η = .58).

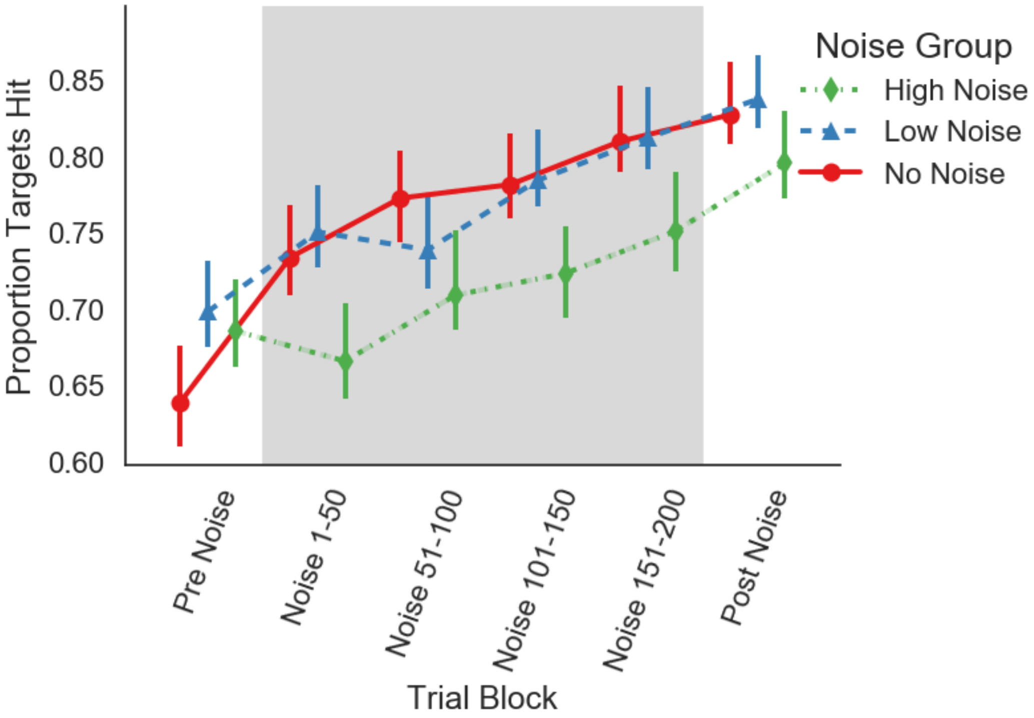

Figure 4.4: Posterior estimates of the group mean probability of hitting the targets as a function of noise group and trial block. Points show posterior means and errors bars show the 50% HDI.

We also examined the difference between the groups, averaged over all the noise blocks, and found no evidence of a difference between the no noise and low noise groups (contrast mean = 0.00, SD = 0.05, 95% HDI = [-0.1, 0.10], 50% HDI: [-0.03, 0.4], η = .53). There was, however, a possible difference between the no noise and high noise groups (contrast mean = 0.06, SD = 0.05, 95% HDI = [-0.04, 0.17], 50% HDI = [.02, 0.9], η = .88), as well as between

Table 4.4: *No Noise P(Hit)* group. Contrasts between added noise trial blocks and baseline.

Contrast: SV	Mean	SD	95% HDI	50% HDI	η
Noise 1-50 :: BL	0.09	0.04	0.02, 0.16	0.07, 0.11	$> .99$
Noise 51-100 :: BL	0.13	0.04	0.05, 0.22	0.10, 0.16	$> .99$
Noise 101-150 :: BL	0.14	0.04	0.07, 0.21	0.12, 0.16	$> .99$
Noise 151-200 :: BL	0.17	0.05	0.08, 0.26	0.14, 0.20	$> .99$
Post-Noise :: BL	0.18	0.05	0.11, 0.27	0.16, 0.21	$> .99$

the low noise and high noise groups (contrast mean = 0.05, SD = 0.05, 95% HDI = [-0.04, 0.16], 50% HDI: [-0.02, 0.9], η = .88). The differences in both of these contrasts were centred around 0.05, although 12% of the contrast distribution lay on the other side of zero, thus the η value did not reach the level of a 'highly probable' difference.

We also examined how the probability of hitting the targets improved over time. For the no noise group, the probability of hitting the targets clearly improved compared to baseline over all the trial blocks with an increase of around 1% over each block (see table 4.4 for contrasts). For the low noise group, the difference between baseline and noise block was reliably detectable by the 101-150 block (see table 4.5 for contrasts). For the high noise group, a clear difference was observed by the final post-noise block (see table 4.6 for contrasts).

4.3.2 Power analysis

The key finding was that strike velocity appeared to increase over time for the high added noise group. The results also suggested that the low noise group may have reduced their strike velocity over time, particularly between

Table 4.5: *Low Noise P(Hit)* group. Contrasts between added noise trial blocks and baseline.

Contrast: SV	Mean	SD	95% HDI	50% HDI	η
Noise 1-50 :: BL	0.05	0.03	-0.01, 0.11	0.03, 0.08	.95
Noise 51-100 :: BL	0.04	0.04	-0.04, 0.12	0.02, 0.07	.85
Noise 101-150 :: BL	0.09	0.03	0.02, 0.15	0.07, 0.11	$>$ .99
Noise 151-200 :: BL	0.11	0.04	0.03, 0.20	0.09, 0.14	$>$ .99
Post-Noise :: BL	0.14	0.04	0.07, 0.21	0.12, 0.02	$>$ .99

Table 4.6: *High Noise P(Hit)* group. Contrasts between added noise trial blocks and baseline.

Contrast: SV	Mean	SD	95% HDI	50% HDI	η
Noise 1-50 :: BL	-0.02	0.03	-0.09, 0.04	-0.04, 0.00	.28
Noise 51-100 :: BL	0.02	0.04	-0.06, 0.10	-0.00, 0.05	.72
Noise 101-150 :: BL	0.04	0.03	-0.03, 0.10	0.01, 0.06	.87
Noise 151-200 :: BL	0.07	0.05	-0.03, 0.16	0.04, 0.10	.91
Post-Noise :: BL	0.11	0.04	0.03, 0.19	0.08, 0.01	$>$ .99

the baseline and 51-100 trial block, but the η value for this contrast did not reach the threshold for a 'highly probable' difference.

As in section 4.3.2 a Bayesian power analysis was conducted to test how many participants would be needed to provide sufficient power to detect the hypothesised difference between baseline and the 51-100 trial block for the no-noise group. As in section 4.3.2 the mean of the marginal posterior distributions over the model parameters was taken as our estimate of the true effect. We then generated 5 data sets with $N = 15$ for each group ($N = 30$ in total) and refit the model with each data set. A sample size of 15 was chosen as this was a small increase from the sample size used in this study. As in section 4.3.2, 4 out of 5 of the model fits returned $\eta < 0.05$. The maximum likelihood estimate suggested that statistical power with $N = 15$ would be approximately 79% for this contrast.

4.4 Discussion

This experiment was designed to explore whether there was evidence for adults using an OC strategy within the course of a short interceptive task. In order to address this issue, we investigated whether external Gaussian noise would cause predictable changes in the strike velocities selected by participants when performing a one degree of freedom interception task. Participants completed 100 baseline trials with no added noise followed by experimental trials in which different amounts of noise were added to the velocity of a puck that the participants launched at a moving target. When no noise was added (*no noise* group), participants may have reduced their speed between the pre-noise and first block of added noise trials, but the contrasts did not reach the threshold for a 'highly probable' difference. The probability of this group hitting the target increased over the course of the experiment. When low noise was added (*low noise* group) participants maintained a fairly constant strike velocity, whilst also increasing their probability of hitting the targets across the blocks. When high noise was added (*high noise* group), participants increased their strike velocity, reaching a peak strike velocity by the third block of trials. This group also increased their probability of hitting the targets over the block. Contrasts suggested that the *high noise* group may have had a lower probability of hitting the targets than the other groups (the contrast mean suggested a difference of approximately 5%), but the η values did not reach the criteria for a 'highly probable' difference.

Overall these findings support the hypothesis that adults vary their movement speeds in order to accommodate the effects of noise on their temporal precision whilst maintaining energy efficiency, as predicted by OC. This complements the finding that adults make faster movements when attempting to hit faster or narrower targets (Tresilian & Houseman, 2005; Tresilian & Plooy, 2006). The present work extends previous observations by showing that the system can use OC over a relatively short time period (i.e. the duration of the experiment: $\sim$ 40min). This ability can help to explain the exquisite levels of performance reached by skilled performers in tasks such as tennis.

It appears that more participants may have been required to confirm that people reduce their strike velocity in the absence of added noise. This would complement the findings and suggest that OC tailors behaviour over a relatively short learning period. A Bayesian power analysis suggested that 15 participants would provide approximately 79% power to detect the hypothesised effect. It seems plausible that as participants became better at the task with practise their temporal precision would have improved, allowing them to reduce their strike velocities until precision was at the lowest acceptable level. This strategy would allow energy costs to be minimised. While the results lend support to this hypothesis a replication with a larger sample size is required to confirm the effect.

The fact that participants showed behaviour consistent with an OC strategy suggests that they were sensitive to the distribution of their temporal errors, or a proxy of this variable (i.e. the proportion of targets they were hitting). The mechanism that provided sensitivity to the prior performance distribution is not addressed within the current study. One possibility is that the distribution is explicitly represented within the system, and this representation is used to influence behaviour when precision decreases. An alternative possibility is that the system was simply biased by the preceding trial when a successful hit was made. The fact that faster movements increase precision means that trials where the movement was faster than average would have a greater probability of hitting the target. This bias would result in faster movements being selected on average as there would be more preceding trials associated with a fast movement than a slow one. The net result would be an increase in fast movements over the course of the trials, but without the system ever possessing an explicit representation of the distribution.

While there is a clear signal (missing the target) that can explain why the system would select faster movements, it is less clear how the system would sense that the faster movements had higher energy costs than necessitated by the task constraints. It is possible that the pressure for selecting slower movements comes from the increased possibility of implementing online feedback corrections afforded by a decrease in speed. The potential for feedback correction is limited in fast interceptive timing actions but is nevertheless present. The ability to slightly increase or decrease dampening through co-activation of the arm's muscles provides a means by which small adjustments in the time of arrival of the arm might be made online. These online adjustments might be of limited effectiveness but even small improvements could affect performance and thus drive change in a system striving to find benefits at the edge of performance capability. If this conjecture has any merit then it suggests that the improvements in energy efficiency caused by the selection of slower movements may be driven by the small enhancements in accuracy that can be obtained by decreasing movement speed when task constraints allow. This would mean that a single variable (hitting performance) could provide the necessary signal to drive the increases and decreases in movement duration that underpin optimal control in interceptive timing behaviours.

It is also possible that the system is using simple heuristics in order to optimise movement speed. For example, the adult system might have learned over the preceding decades that optimality results from decreasing speed when hitting targets and increasing speed when missing. The implementation of such a heuristic could offer some advantages. One advantage of increasing speed when missing targets is that the increased temporal precision might make it easier to detect an error signal indicating systematic inaccuracy (van Beers, 2012).

One strong test of the idea that adults are demonstrating OC within the current experimental task would involve adding signal dependant noise to the movements. This manipulation would allow the creation of an unnatural situation where slow movements were associated with less variability than fast movements in an interceptive task (to fast moving targets). Unfortunately implementing such as task is difficult as the reduction in temporal variability with increased movement speed is much greater than effect of adding signal dependent noise, even when the noise increases exponentially. This is because when moving at high speeds, even large errors in movement speed result in very small deviations in movement time[1]. One way around this would be to use an interceptive timing task in which the feedback provided to participants

---

[1] This is because the error in movement time (MT_e) is given by $MT_e = \frac{D}{S_e}$, where D is the distance to move and S_e is the error in movement speed.

is manipulated. This could be done by making the target disappear shortly before it reaches the interception point (the spatial location at which the bat could intercept the target) and then manipulating the position it reappears at once the bat (or puck) also reaches the interception point. This would allow for slow movements to be associated with less temporal variability. If participants are as effective in using OC as the present experiments suggest, then the results should show participants quickly learning to select slower movements.

The time course of such action selection would be of great interest. On the one hand, there is the possibility that the system has an inbuilt bias to increase speed when a task's temporal constraints become stricter or temporal errors increase. This would suggest that it will take longer for participants to reduce their movement speeds when better performance results from the selection of slower movements. On the other hand, the rewards associated with the selection of slower movements would include energy efficiency as well as improved precision, and this dual benefit might drive the system to show even faster adaptation than observed in the current experiment. It is always possible, however, that these factors (an inbuilt bias and a shift in the relative rewards) would counteract one another. These are empirical questions that can nonetheless be addressed readily through a relatively simple adjustment to the current experimental design.

In conclusion, we examined whether participants would show behaviour consistent with OC during the performance of a simple interceptive timing task. The task design allowed external noise to be added to the speed of a puck. We found that participants increased their strike velocity when high levels of noise were added to the puck's speed. These findings are consistent with adults adopting an optimal control policy, so that participants maximised the number of targets they hit whilst simultaneously minimising effort in response to the exact requirements of the task.

This chapter and chapter 3 examined the mechanisms by which adults achieve their amazing levels of temporal precision, a skill level that has been well documented in the scientific literature (Brenner & Smeets, 2015b). However, these abilities have not been well documented in children. The next part of this thesis examines interceptive timing in children, beginning with an examination of how the ability to precisely time interceptive actions develops over childhood.

Part II

Children

Chapter 5

The development of interceptive timing abilities

Chapter Abstract

Intercepting moving targets underpins a myriad of human activities, with adults evidencing exquisite interceptive abilities. While interceptive timing performance has been carefully documented in adults, the ontogeny of these abilities in children remains unclear. Understanding the typical developmental trends is useful as interceptive timing abilities may provide a marker of neurodevelopmental pathology in addition to providing insights into sensorimotor development. We used a cross-sectional design to examine the development of interceptive timing abilities in primary school children (aged 5-11 years; n = 309) and adults (n = 22), using a striking task involving 54 moving virtual targets. We manipulated task difficulty by altering target speed and width and used a one degree-of-freedom manipulandum so we could isolate temporal error from spatial inaccuracy. The results showed clear developmental trends with

the probability of hitting the targets improving with age. This improvement in performance was due to refinements in both accuracy (reduction in temporal bias) and precision (less variability in the errors), but even the oldest children were far from adult levels. Interestingly, accuracy was similar between the adults and all other age groups when intercepting fast targets. However, for the slow and medium speed targets there was a systematic bias to hit too early, with the bias increasing in the younger age groups. The systematic bias can be explained by the use of a ‘dumb but smart’ heuristic where perceptual uncertainty is associated with fast approaching targets.

5.1 Introduction

The sensorimotor skill repertoire of humans is remarkable, both in terms of the number of tasks that people are able to engage in and the incredible levels of performance they demonstrate. The ability to hit fast moving targets provides an exemplar of human sensorimotor abilities. Successfully striking moving objects requires movements to be timed so that the intercepting effector arrives at the same spatial location as the target at just the right time. When objects are moving at high speed, the time window in which the target object can be struck is often in the order of milliseconds. Thus temporal errors in the range of tens of milliseconds can result in failure to hit the target. Temporal errors over multiple trials can be described in terms of temporal bias (accuracy) and variable errors (precision). Temporal bias refers to systematic errors in the timing of movements (i.e. hitting too early or late on average), while variable error refers to the precision in the timing errors over repeated trials (typically the standard deviation of the temporal errors). To repeatedly hit moving objects people must therefore achieve both low temporal bias and variable error (accuracy and precision).

The fact that human adults readily achieve this is remarkable given the challenges associated with neural delays (Nijhawan, 2008) and noise in the sensorimotor system (Faisal et al., 2008). Typical adults show incredible interceptive timing abilities, hitting falling balls with a temporal standard deviation (variable error) of $\approx$ 6ms (Brenner & Smeets, 2015b; Brenner et al., 2012). This is particularly impressive given that adults show much larger temporal standard deviation when performing other tasks (e.g. achieving a standard deviation of only $\approx$ 20ms when gauging which of two visual targets appears first (Brenner & Smeets, 2010)).

To study interceptive timing behaviours in adults, researchers have used tasks in which the possible trajectory of the interception is restricted to a single axis (1-DoF interceptive tasks). This allows temporal errors to be isolated from errors in the spatial trajectory, since the presence of spatial error makes it difficult to study timing behaviour in unconstrained interceptive tasks. Thus, tasks that allow exploration of interceptive timing behaviour use apparatus where targets move along a linear track and participants hit the moving target with a 'bat' which can only move orthogonally to the target's track (Tresilian & Lonergan, 2002). The temporal difficulty of such tasks can be controlled by manipulating the time window (the period in which an interception can be made). The time window is a function of the width of the target and the target's velocity (if the size of the effector remains constant). This form of task is required for investigating how timing abilities develop throughout childhood because the difficulty of the task can be parametrically varied to make it suitable for different age groups. Unfortunately, there is a dearth of research utilising such methodology with regard to interceptive timing in children.

A number of studies have investigated how children's catching abilities change with age, and several standardised motor tests include such measures as part of the assessment battery (e.g. Movement Assessment Battery for Children, Henderson and Sugden (1992)). These tasks reveal clear improvements in performance as children get older. However, the complexity of the catching task (involving a number of skilled components including spatial coordination) does not allow one to draw inferences about the ability of children to intercept moving targets as a function of age. The only relevant timing data that has been collected with children comprise of coincidence-anticipation tasks (Haywood, 1980). Coincidence-anticipation tasks involve participants viewing a linear track of equally spaced LEDs which illuminate sequentially. Participants attempt to press a button to coincide with the last illuminated LED, and the time between the button press and the LED illumination provides a measure of timing ability. In one such study it was reported that coincidence timing performance reached adult-like levels by approximately 11-13 years (Haywood, 1980). However, this task does not require participants to take account of the time it takes them to make a movement, and requires estimates of the target's time to arrival (TTA) to be made over a much shorter time duration than in a true interceptive task, where movements must be initiated while the target is further away from the planned interception point. In another study, Kim, Nauhaus, Glazek, Young, and Lin (2013) used a coincidence timing task to explore aiming movements. Kim et al. found that variable errors were higher in 14-16 year olds than 17-18 year olds. Unfortunately, it is impossible to determine whether the poorer performance in the younger age group was due to differences in the timing judgements or the control of the aiming movements. Thus, there currently appears to be no reliable data showing how interceptive timing abilities change over childhood.

In order to shed light on the ontogeny of interceptive timing, we developed a 1-DoF interceptive timing task in which participants hit virtual moving targets by controlling an on-screen bat via a custom-built manipulandum. The task involved hitting a series of targets of varying difficulty (three speeds and three target widths) so the same task could be used to measure interceptive timing abilities in young children (ages 5-11 years) and an adult comparison group. A Bayesian model was developed which related temporal bias and variable errors to the probability of hitting the target (P-Hit), which provided a simple measure of performance. We were first interested in documenting how P-Hit changed as a function of age and whether adult levels of performance were reached by 11 years. We also examined whether changes in P-Hit were related to refinements in the precision of the movements (low variable error), in accuracy (low temporal bias), or both. Adults show smaller temporal errors (better precision) when the time window is shorter (when hitting small or fast targets) (Tresilian & Houseman, 2005; Tresilian & Plooy, 2006) which they achieve by making faster, briefer movements. We therefore examined movement time (MT) to examine whether children also exploit this relationship to reduce variable errors when the time window is shorter.

5.2 Methods

5.2.1 Participants

Participants were recruited from a state primary school in Bradford, West Yorkshire, UK. There were 368 children in UK school years 1 to 6 (aged 5-11 years) at the time of testing. All children were invited to take part in the study. The children completed two test sessions in which they completed a range of motor and cognitive tasks. All sensorimotor tasks took place in the first session. Ethical approval was obtained from the University of Leeds Ethics and Research committee.

From the 368 children at the school, 309 full data sets were included in the data analysis. Eleven children were removed from the 368 because they were classed as having special education needs (SEN) by the school. Twenty-nine were excluded because the experimenter recorded that they did not complete one or more tasks. Fourteen were excluded because they did not provide data on the interception task and five did not provide data on postural control.

An adult comparison group (n = 22; 15 Female; Mean Age = 24.76 years, SD Age = 4.70) was also recruited from the University of Leeds.

5.2.2 Task

Children completed a computer based interception task in which they hit moving targets by controlling a custom-made 1-DoF joystick (see figure 5.1). The joystick was placed next to a horizontally positioned BenQ XL2720Z LCD gaming display (Resolution: 1920×1080, size: 548×642mm, brightness:

300cd m^{-2}, refresh rate: 144Hz). The position of the joystick was represented on screen by a black rectangular 'bat' (dimensions: 10 × 15mm) that was always in line with the joystick. All stimuli were generated using Python 2.7.9 using open source libraries.

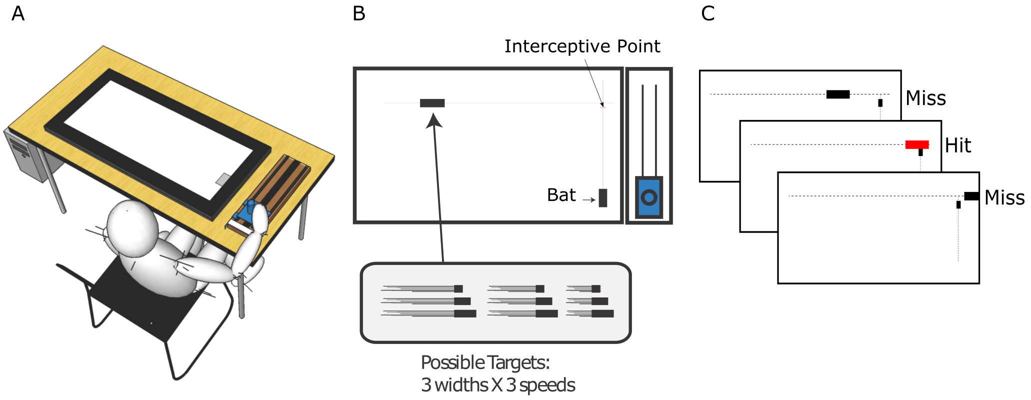

Figure 5.1: A) The experimental setup: children viewed a horizontally oriented monitor while controlling an onscreen 'bat' via a 1-DoF manipulandum (placed on the left of the display for left handed participants with stimuli reversed). B) A schematic of the experimental display. Targets moved from left to right across the screen. Participants were instructed to hit the target from beneath. C) Possible outcomes. In the upper pane the participant has arrived too early. In the middle the participant successfully hit the target on its underside. In the lower pane the participant was too late.

All coordinates are given in millimetres, with the origin of coordinate system at the bottom left of the screen. A 'start box' appeared on screen at the start of every trial and the participant was instructed to place the bat within it (coordinates [570, 20mm]; coordinate origin at top left of screen). A black target (height: 15mm) then appeared at the left hand side of the screen (coordinates [0,150mm] (for left handed participants the apparatus and stimuli were reversed, with the manipulandum placed on the left side of the screen). After a delay drawn from a uniform distribution U(0.25, 3.0s) the target moved from left to right at a constant speed. The centre of the target passed in front of the centre of the bat after moving 570mm. The children were instructed to hit the target with the bat. The target was successfully hit if the upper edge of the bat collided with the lower edge of the target (see figure 5.1). The target then stopped moving, turned red and span before disappearing, thereby providing motivating animated feedback for the children (the children enjoyed ‘playing the game’). If the bat passed in front of the target’s horizontal path the target immediately stopped moving and then remained on screen for 1 second. Thus, participants could not simply move the bat in front of the target’s path and wait for the target. If the bat crossed the target’s path after the target had moved too far to be struck then the target stopped and remained visible for 1 second. The position of the bat and target was timestamped and saved to disk at 144Hz. The bat’s positional data were filtered using a low pass second order zero-lag Butterworth filter with a cut off frequency of 10Hz. Spline interpolation was used to estimate the time at which the bat reached the interception point. The total number of targets hit by each participant provided our measure of interceptive timing ability which is referred to as IntT.

Children performed 54 trials in which the target speed (250mm s^{-1}, 400mm s^{-1}, 550mm s^{-1}) and target width (30mm, 40mm, 50mm) varied (9 trial types x 6). Each target type was presented in a block of 3 trials, with 2 blocks for each trial type. The blocks were pseudorandomly ordered with the constraint that two blocks of the same kind could not occur sequentially. All participants experienced an identical sequence of blocks. The order of the targets (by speed and width) is given in table 5.1. This was done to make it

Table 5.1: Sequence of target type blocks experienced by all participants. Participants attempted to hit three targets of the same type within each block.

Block	Target Speed ($\mathrm{mm\,s^{-1}}$)	Target width (mm)
1	250	40
2	250	30
3	250	40
4	400	30
5	400	40
6	400	50
7	400	30
8	550	30
9	400	40
10	250	50
11	550	50
12	400	50
13	550	40
14	250	30
15	550	50
16	550	30
17	550	40
18	250	50

easier to compare children's performance in later chapters, but means that order effects may not be accounted for in the data analysis.

5.2.3 Data analysis and measures

We were interested in the probability of hitting the target (P-Hit), the temporal bias (accuracy) and the variable errors (precision) of the movements. P-Hit, temporal bias and variable errors can be estimated from the participant's temporal errors (TE). TEs describe how early or late a movement was (negative values mean the bat arrived at the interceptive point too early), and is calculated as

$$TE = \frac{\text{Target}_{\text{center}} - \text{Bat}_{\text{center}}}{\text{Target}_{\text{speed}}} \tag{5.1}$$

where the numerator provides the spatial distance (along the x-axis) between the target and the bat, at the point in time at which the upper edge of the bat shares the same y coordinate as the lower edge of the target (the point at which an interception could have taken place). We were also interested in the movement time (MT), defined as the time at which the bat's velocity exceeded $40\text{mm}\,\text{s}^{-1}$ until the point at which the bat reached the interceptive point. Finally we also examined the movement initiation time (IT), defined as the time delay between the target starting to move and the participant initiating their movement, in seconds.

5.2.4 Temporal error model (bias and variable error)

A Bayesian multilevel model was used to model the TE data. The model provided estimates of P-Hit, temporal bias and variable error. Each temporal error (y_i) was modelled as belonging to participant p. Each participant had a vector of regression coefficients β_p which determined their temporal bias for each of the target widths and speed. The likelihood function for datum y_i was given by:

$$\mu_i = x_i \beta_{p[i]} \tag{5.2}$$

$$y_i = \text{Student's t}(\nu_{ag[i],w[i],s[i]}, \mu_i, \sigma_{ag[i],w[i],s[i]}) \tag{5.3}$$

where x_i is a row vector specifying the target's width and speed on the ith trial. Thus μ_i gives the temporal bias of participant p for the target width and speed that was encountered on the ith trial. Each participant's vector of regression coefficients (β_p) was distributed by a multivariate normal distribution,

$$\beta_p \sim \text{MultiNormal}(M_{g[p]}, \Sigma) \tag{5.4}$$

where $M_{g[p]}$ is the mean vector for the group g to which participant p belongs. The covariance matrix Σ was specified as

$$\tau_k \sim \text{Cauchy}^{+}(0, 2.5), \quad k \in 1 : K \tag{5.5}$$

$$\Omega \sim \text{LKJCorr}(2) \tag{5.6}$$

$$\Sigma = \text{diag\_matrix}(\tau)\Omega\text{diag\_matrix}(\tau) \tag{5.7}$$

where K is the number of regression coefficients. This formulation allows the prior over Σ to be specified as a correlation matrix (Ω) and vector of scale parameters (τ). The prior distribution over the correlation matrix ($LKJCorr$) is defined in the Stan user manual (Manual, 2013).

The degrees of freedom and scale parameters for the ith data point (equation 5.3) depended on the age group, target width and target speed of the ith trial ($ag[i]$, $w[i]$, $s[i]$ respectively). The scale parameter provided our measure of the variable error.

The bias and variable error parameters (μ and σ) in the model can be combined to provide an estimate of the probability of hitting the target (P-Hit) as a function of age group, target width and target speed. The time window (TW) defines the duration of time in which the target can be struck. Multiplying this quantity by 0.5 yields the upper and lower TEs that will still result in the target being struck ($TW_{0.5}$), and is given by,

$$TW_{0.5} = \pm 0.5 \times \frac{\text{Target}_{\text{width}} - \text{Bat}_{\text{width}}}{\text{Target}_{\text{speed}}} \tag{5.8}$$

On any given trial, the TE and TW determined whether a target was hit or not. The probability of hitting the ith target given its width and speed is provided by the following equation:

$$\begin{aligned} \text{P-Hit}_i = {} & P(y_i \leq -TW_{0.5} | \mu_i, \sigma_{ag[i],w[i],s[i]}, \nu_{ag[i],w[i],s[i]}) \\ & -P(y_i \leq +TW_{0.5} | \mu_i, \sigma_{ag[i],w[i],s[i]}, \nu_{ag[i],w[i],s[i]}) \end{aligned} \tag{5.9}$$

The major benefit of estimating P-Hit this way over other modelling approaches (e.g. logistic regression) is that changes in P-Hit between age groups, target widths and target speeds can be attributed to changes in temporal bias and variable error.

Initiation time model

Exactly the same model was used to model the IT data as the TE data.

5.2.5 Movement time model

The same model was used for MT except that a lognormal distribution replaced the student's t distribution because MT was always positive and could have positive tails (the parameter ν was dropped from the model). The mean and standard deviation were then calculated as,

$$MT_{\text{mean}} = e^{\mu + \frac{\sigma^2}{2}} \tag{5.10}$$

$$MT_{SD} = \sqrt{(e^{\sigma^2} - 1)e^{2\mu + \sigma^2}} \tag{5.11}$$

Bayes' rule was used to estimate the credible values of the model parameters (θ) given the data. A representative sample was drawn from the posterior using the 'No-U-Turn sampler' (Hoffman & Gelman, 2011) implemented in PyStan 2.12. Four chains were started at random values of θ, taking 2500 warmup iterations followed by 2500 samples each. Convergence was assessed by visually examining the chains. All $\hat{R}$ values were below 1.1.

5.3 Results

We report the model's marginal posterior distributions over group level parameters unless otherwise stated. Marginal posterior distributions are plotted with points showing the mean of the distribution while error bars denote the 95% highest probable density (HDI). For a unimodal distribution of mass M, the HDI is the narrowest possible interval of M (Kruschke, 2015).

The 95% HDI is the interval in which there is 95% probability that the true parameter value falls.

5.3.1 Probability of hitting the targets as a function of age

P-Hit (the estimated probability of hitting the target) provides a measure of interceptive timing ability. Figure 5.2A shows the posterior mean and 95% HDI for the probability of hitting the target as a function of age group, averaged over all target widths and speeds. P-Hit was lowest in the youngest age group (< 6 years) with a posterior mean of 38% (SD = 0.012, 95% HDI = [.366, .390]) and reached 60% (SD = 0.014, 95% HDI = [.580, .628]) in the oldest children (10+ years old). In contrast, the adults performed much better, with a posterior mean P-Hit of 87% percent (posterior SD = 0.011, 95% HDI = [.854. .897]), indicating that the oldest children tested were far from adult performance levels, hitting 27% percent (posterior SD = 0.017, 95% HDI = [.237, .307]) fewer targets.

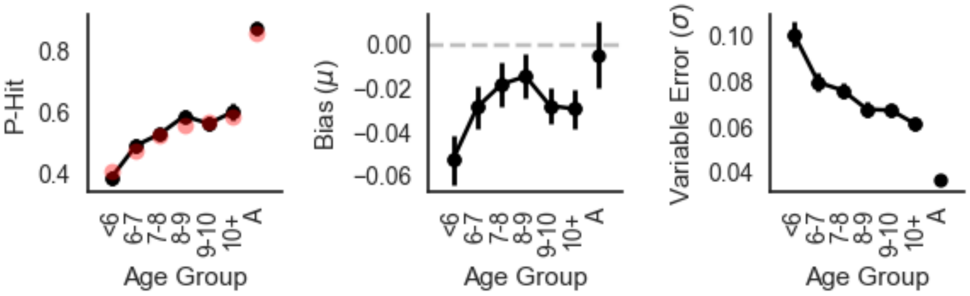

Figure 5.2: A)P-Hit (Probability of hitting the target), B) Temporal Bias and C) Temporal Variable Errors as a function of age group, averaged over all target widths and speeds. Black points show posterior mean. Error bars show the 95% HDI. Red points show the data means for the P-Hit metric. These are not shown for the bias or variable error metric.

5.3.2 Temporal bias and variable error as a function of age

Figure 5.2B and 5.2C reveal that changes in P-Hit were due to both a reduction in temporal bias and variable errors with increasing age groups. All the children tended to hit early with this bias decreasing until the 8-9 year old age group. The adult group were the only group who showed near zero bias (posterior mean = -0.004, SD = 0.006, 95% HDI = [-0.002, 0.01]). Variable errors also appeared to decrease across the age groups, with the largest decrease between consecutive year groups (in the children) seen between the < 6 and 6-7 years age groups (contrast mean = -0.02, contrast SD = 0.003, 95% HDI = [-0.023, -0.019]) and smaller differences seen thereafter (mean difference ≈ -0.003). Again the adults showed much lower variable errors than even the oldest age group of children (contrast mean = 0.024, contrast

SD = 0.002, 95% HDI = [-0.025, -0.023]).

5.3.3 P-Hit as a function of age group, target width and speed

To test whether changes in P-Hit with age group depended on the target's width and speed (i.e. whether certain targets are better able to discriminate between certain age groups) P-Hit was examined across the age groups as a function of target width and target speed (see figure 5.3). The fastest targets (right panels in figure 5.3) revealed approximately linear increases in P-Hit across consecutive age groups for the children (increasing by approximately 3% with each age group) which was consistent across the three target widths. In contrast, the slowest targets revealed a different pattern, with much larger differences in P-Hit between the youngest age groups. For example, in the 250mm s^{-1}, 30mm target, P-Hit increased by approximately 10% for each consecutive age group, until age group 8-9 years, beyond which no improvements were seen in the children. However, the adult group still hit 32% (posterior SD = 3.68, 95% HDI = [30.4, 35.9]) more targets than even the oldest children. This general trend was also seen for the other target widths, with the largest increase in P-Hit found between the youngest age groups and smaller differences found thereafter (for the child age groups).

The effect of target width on P-Hit was then examined as a function of age group and target speed. To keep the number of contrasts performed to a manageable number P-Hit was compared between the widest (50mm) and narrowest (30mm) targets, for every age group and target speed. As expected

P-Hit was considerably higher when hitting the wider targets irrespective of age group and target speed (see figure 5.4; values are all positive indicating higher P-Hit on widest target). Thus performance was better on the widest targets than the narrowest, regardless of age group, with the adults showing increases in P-Hit across the two widths of approximately the same magnitude as the children.

The effect of target speed on P-Hit was also examined as a function of age group and target width. Contrasts between the slowest (250mm s^{-1}) and fastest (550mm s^{-1}) targets (figure 5.5) revealed that P-Hit was higher for the slowest targets in the adults and for most of the children (as indicated by positive values), with a few notable exceptions. First, the youngest group (< 6 years old) showed no advantage in P-Hit for the slowest target when they were 30mm and 50mm wide (95% HDI spanned zero), and actually showed worse performance on the slowest target when the target was 40mm wide (indicated by a negative value). Three other age groups (6-7, 9-10, 10+) showed no difference in P-hit with target speed when striking the 40mm target. Notably, while P-Hit was higher for the slowest target for the remaining group, the magnitude of the difference in P-Hit was smaller for the 40mm targets for all the children age groups. This appeared to be because P-Hit was lower for the 250mm s^{-1} 40mm target than for the 550mm s^{-1} 40mm target. While this result was not expected it was likely because of a practice effect, as the former target appeared at the start of the sequence of targets (see table 5.1).

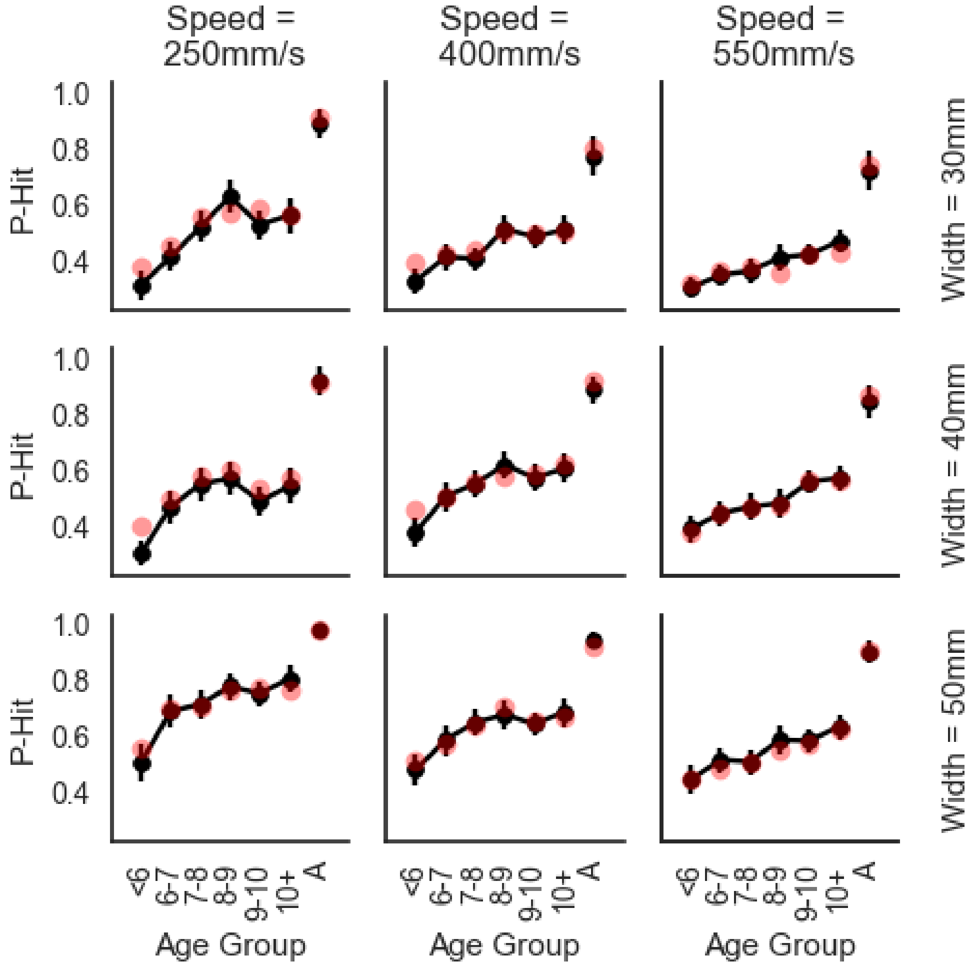

Figure 5.3: P-Hit as a function of age group, target width (rows) and speed (columns). Black points show posterior means. Errors bars show 95% HDI. Red points show the data means for the P-Hit metric. These are not shown for the bias or variable error metric.

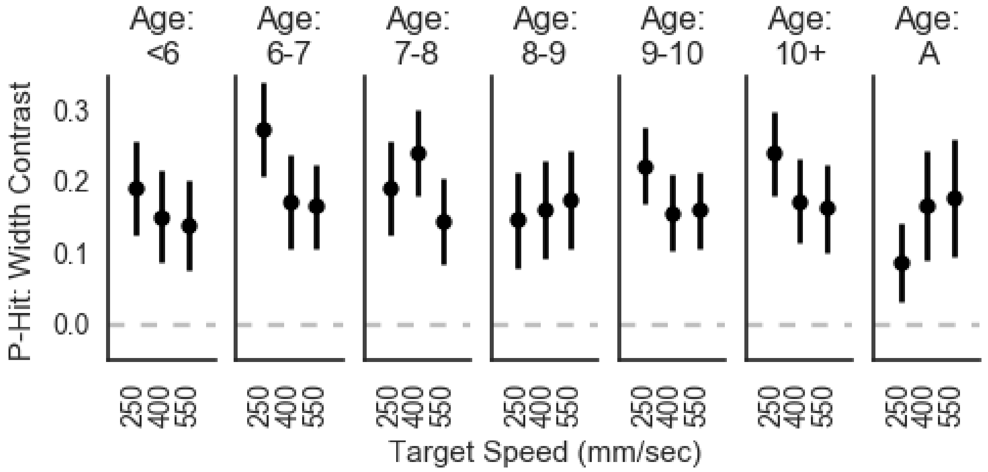

Figure 5.4: P-Hit contrasts: Width. Magnitude of the difference in P-Hit between the widest (50mm) and narrowest (30mm) targets for every age group (columns) and target speed (x-axis). Positive values indicate that P-Hit was higher on the widest target. Points show the posterior mean. Error bars show the 95% HDI.

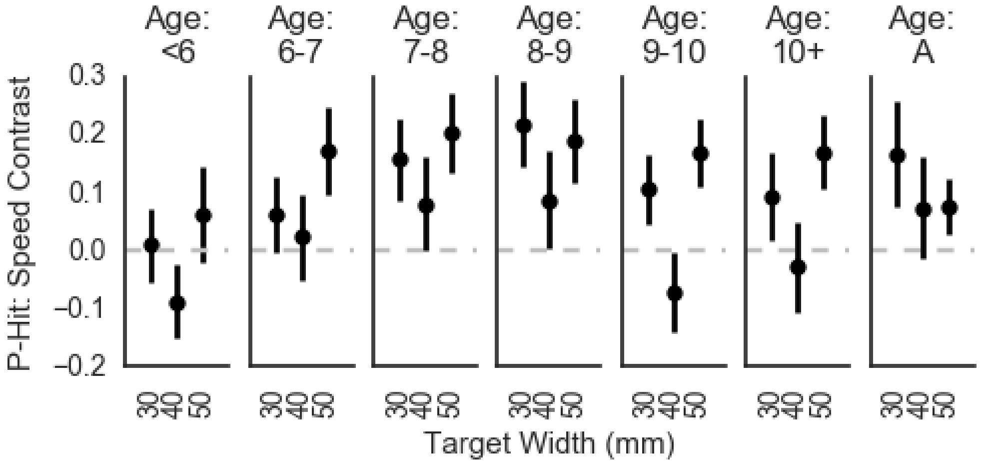

Figure 5.5: P-Hit Contrasts: Speed. Magnitude of the difference in P-Hit between the fastest (550mm/s) and slowest (250mm/s) target speed for every age group (columns) and target speed (x-axis). Positive values indicate that P-Hit was higher for the slowest target. Points show posterior mean. Error bars show the 95% HDI.

5.3.4 Temporal bias as a function of age group, target width and speed

Temporal bias was examined to see whether age related changes in temporal bias could account for changes observed in P-Hit. Temporal bias was examined as a function of age group, target width and target speed (see figure 5.6). Contrasts between the widest and narrowest targets for all age groups and target speeds (see figure 5.7) did not reveal a clear effect of target width on temporal bias, regardless of age group or target speed.

For the younger age groups there was a tendency to strike ahead of the slowest target (indicated by the large negative values in figure 5.6). This was confirmed by contrasts between the fastest and slowest targets (see figure 5.8). This tendency reduced across the subsequent age groups. However, only the adults showed no difference in bias between the fast and slow targets (they showed zero bias for all target types). Again figure 5.8 revealed a tendency for the magnitude of the difference between the fastest and slowest target to be greater (more negative) for the 40mm targets than the other target widths. This was because participants tended to hit too early on the 250mm s^{-1} 40mm target (they also hit this target less; see above). Similar to the pattern observed in P-Hit this was likely due to this target type appearing early in the sequence of targets.

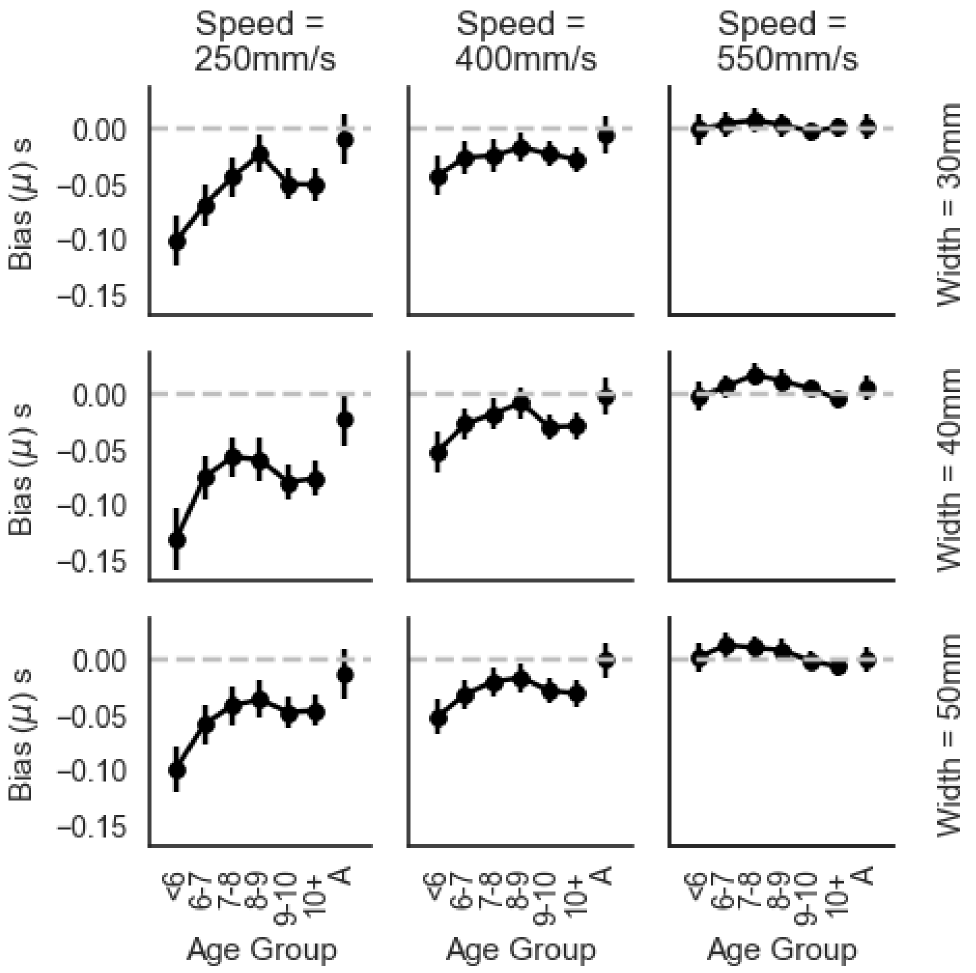

Figure 5.6: Temporal Bias (mean temporal error) as a function of age group, target speed (columns) and target width (rows). Dashed lines indicate perfect accuracy. Points show the posterior mean. Error bars show the 95% HDI (unless smaller than the symbol size).

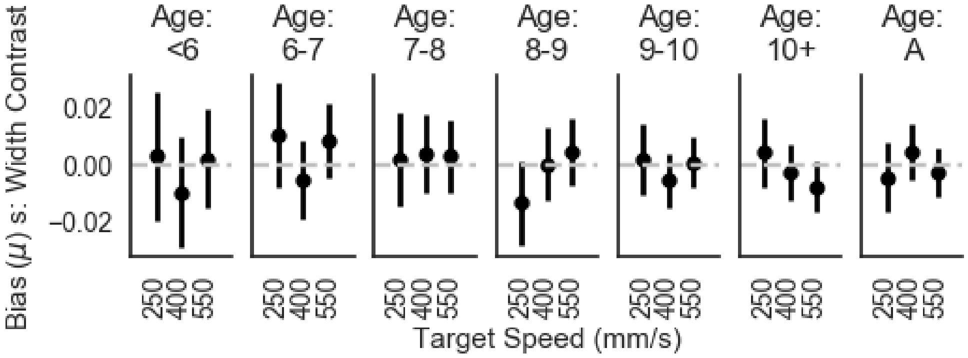

Figure 5.7: Bias (mean temporal error) contrasts: Width. Magnitude of the difference in temporal bias between the widest (50mm) and narrowest (30mm) targets for every age group (columns) and target speed (x-axis). Positive values indicate that accuracy was higher for the widest target. Points show the posterior mean. Error bars show the 95% HDI.

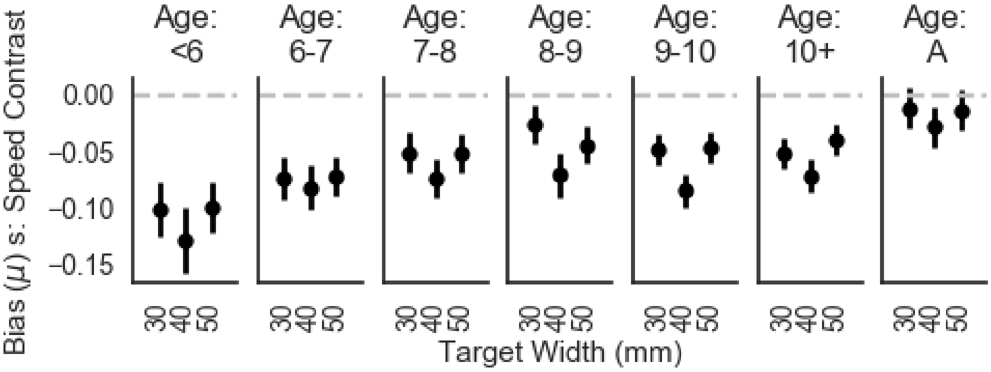

Figure 5.8: Bias (mean temporal error) contrasts: Speed. Magnitude of the difference in temporal between the fastest (550mm/s) and slowest (250mm/s) target speed for every age group (columns) and target speed (x-axis). Negative values indicate that participants were striking early for the slowest target relative to the fastest. Points show posterior mean. Error bars show the 95% HDI.

5.3.5 Variable errors as a function of age group, target width and speed

Variable error was examined to see whether precision could also account for changes in P-Hit (see figure 5.9). We were also interested in whether variable error would change as a function of the temporal window (lower variable error on narrower and faster targets). It was clear that variable errors decreased with age group for all the target widths and speeds, although the adult group showed lower variable errors than the child age groups. Contrasts between the widest and narrowest targets did not reveal an effect of target width on variable errors regardless of age group or target speed (see figure 5.10). However, target speed did appear to effect variable errors as revealed by contrasts between the fastest and slowest targets (see figure 5.11). Variable errors were smaller when intercepting the fastest target, and the magnitude of the difference between the fastest and slowest target was similar across all age groups, including the adult group. There was a trend for the contrast between the fastest and slowest target speed to reveal a difference of a larger magnitude for the 40mm target than the other target widths. As with the P-Hit and temporal bias, this was likely due to a practice effect, with variable errors being particularly large for the 40mm, 250mm s^{-1} target which appeared early in the target presentation sequence (see table 5.1).

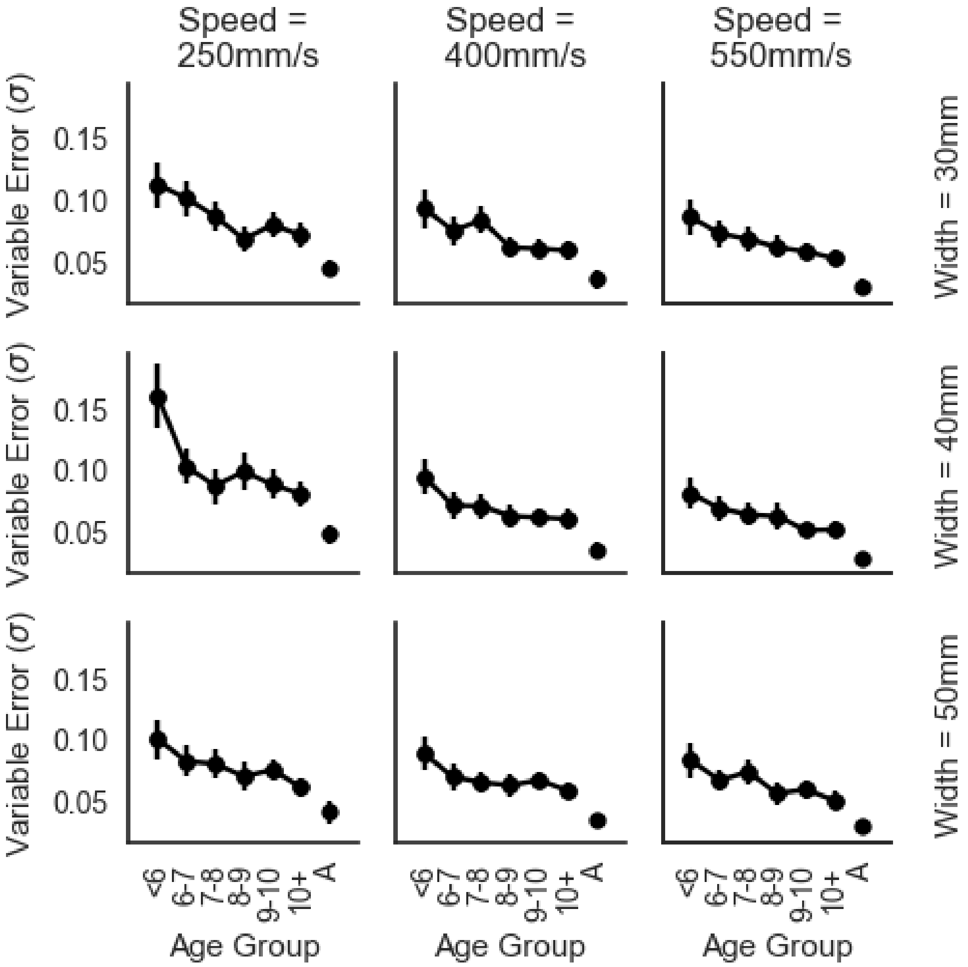

Figure 5.9: Variable Error as a function of age group, target speed (columns) and target width (rows). Points show the posterior mean. Error bars show the 95% HDI (unless smaller than the symbol size).

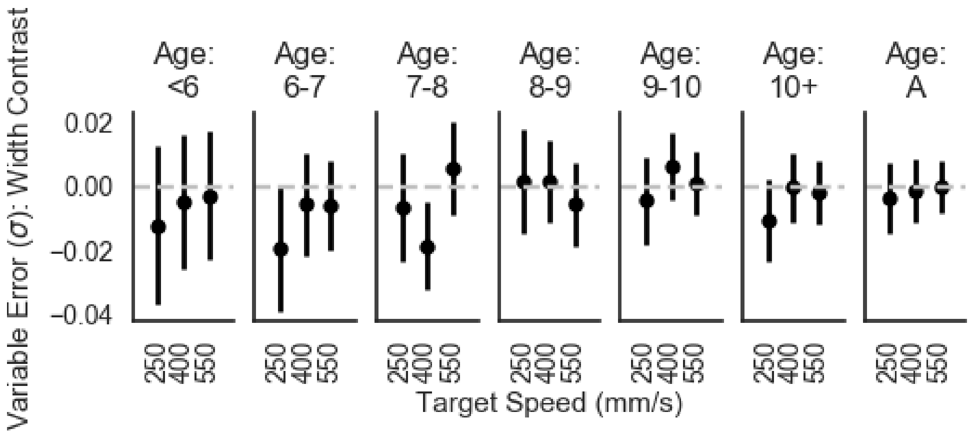

Figure 5.10: Variable Error contrasts: Width. Magnitude of the difference in variable error between the widest (50mm) and narrowest (30mm) targets for every age group (columns) and target speed (x-axis). Positive values indicate that variable errors were larger for the widest target. Points show the posterior mean. Error bars show the 95% HDI.

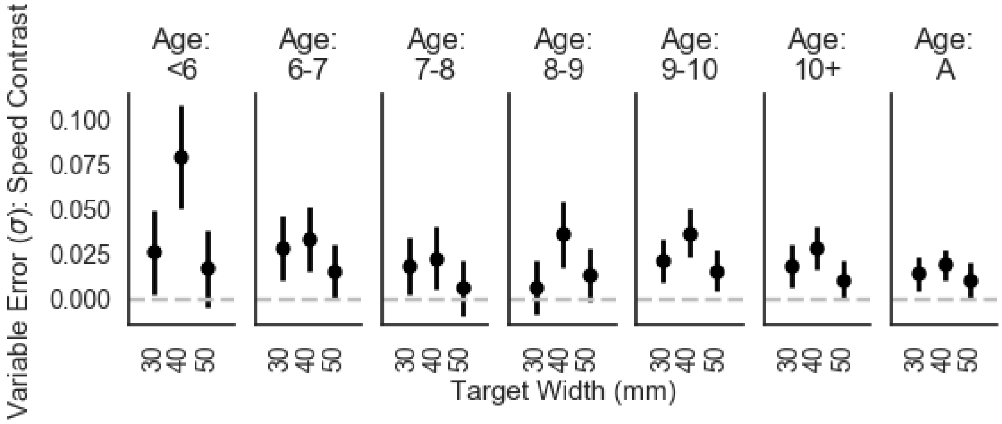

Figure 5.11: Variable error contrasts: Speed. Magnitude of the difference in variable error between the fastest (550mm/s) and slowest (250mm/s) target speed for every age group (columns) and target speed (x-axis). Positive values indicate that variable error was greater for the slowest target. Points show posterior mean. Error bars show the 95% HDI.

5.3.6 Movement time (MT)

Given that variable errors appeared to reduce with target speed but not target width we explored whether these changes could be attributed to differences in MT, as briefer movements are associated with lower variable error (Tresilian, 2012). MT generally remained constant across target widths and this was confirmed by examining the difference between the widest and narrowest target at each target speed (see figure 5.12). However there were a few exceptions, with certain contrasts excluding zero for the < 6, 6-7 and 9-10 years age groups. There was not a clear pattern to the results and the values suggested that MT was actually longer for the narrowest targets in these cases.

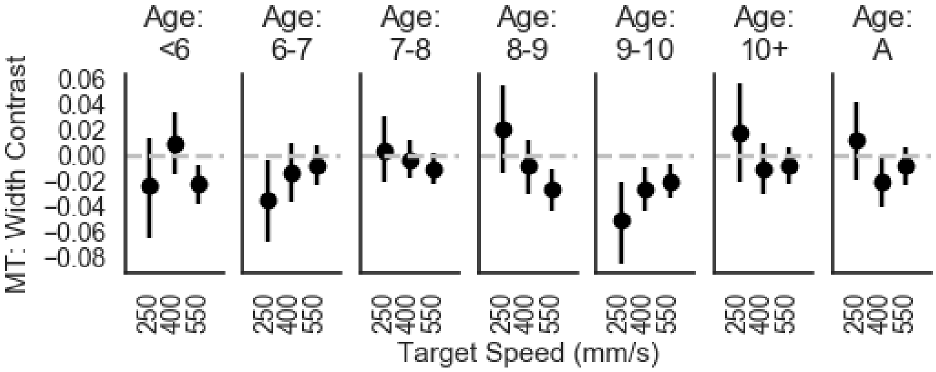

Figure 5.12: MT contrasts: Width. Magnitude of the difference in MT between the widest (50mm) and narrowest (30mm) targets for every age group (columns) and target speed (x-axis). Positive values indicate that MT was longer for the widest target. Points show the posterior mean. Error bars show the 95% HDI.

In line with Tresilian and Houseman (2005) MT decreased with target speed as shown in figure 5.13. The magnitude of the difference between the fastest and slowest target was roughly the same across target widths but showed some variation between the age groups.

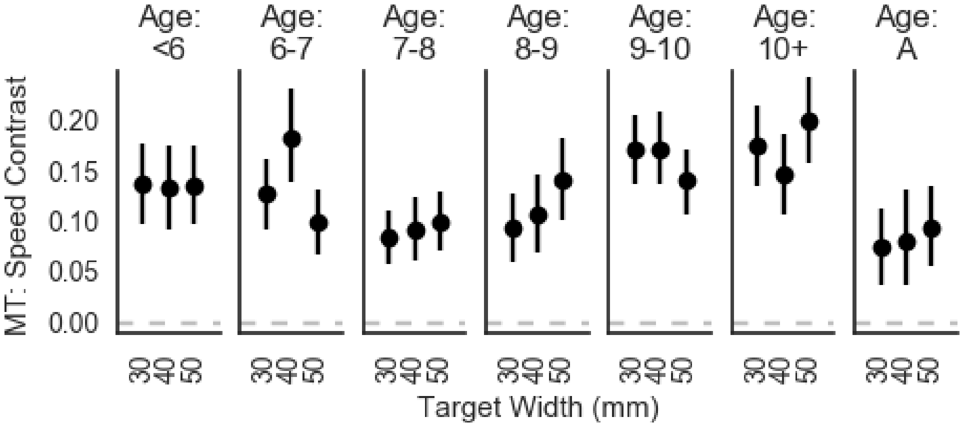

Figure 5.13: MT contrasts: Speed. Magnitude of the difference in MT between the fastest (550mm/s) and slowest (250mm/s) target speed for every age group (columns) and target speed (x-axis). Positive values indicate that MT was longer for the slowest target. Points show posterior mean. Error bars show the 95% HDI.

5.3.7 Initiation time as a function of age group, target width and speed

Initiation time was examined to see if systematic differences in the time of movement initiation would exist between age groups. It may be that the younger children showed a larger temporal bias (see figure 5.6) because they simply triggered their movements as soon as they perceived the target moving. Figure 5.14 suggests that this wasn't the case, with no clear differences between age groups in terms of initiation time. It is clear from figure 5.16 that initiation times were earlier in the target's motion when the target was faster, as would be expected as the target actually moved for less time (all targets moved a fixed distance). This was confirmed by contrasts in figure 5.16. Initiation time appeared to vary across target widths for the youngest two age groups (¡6 years and 6-7 years), as well as the 9-10 years age group.

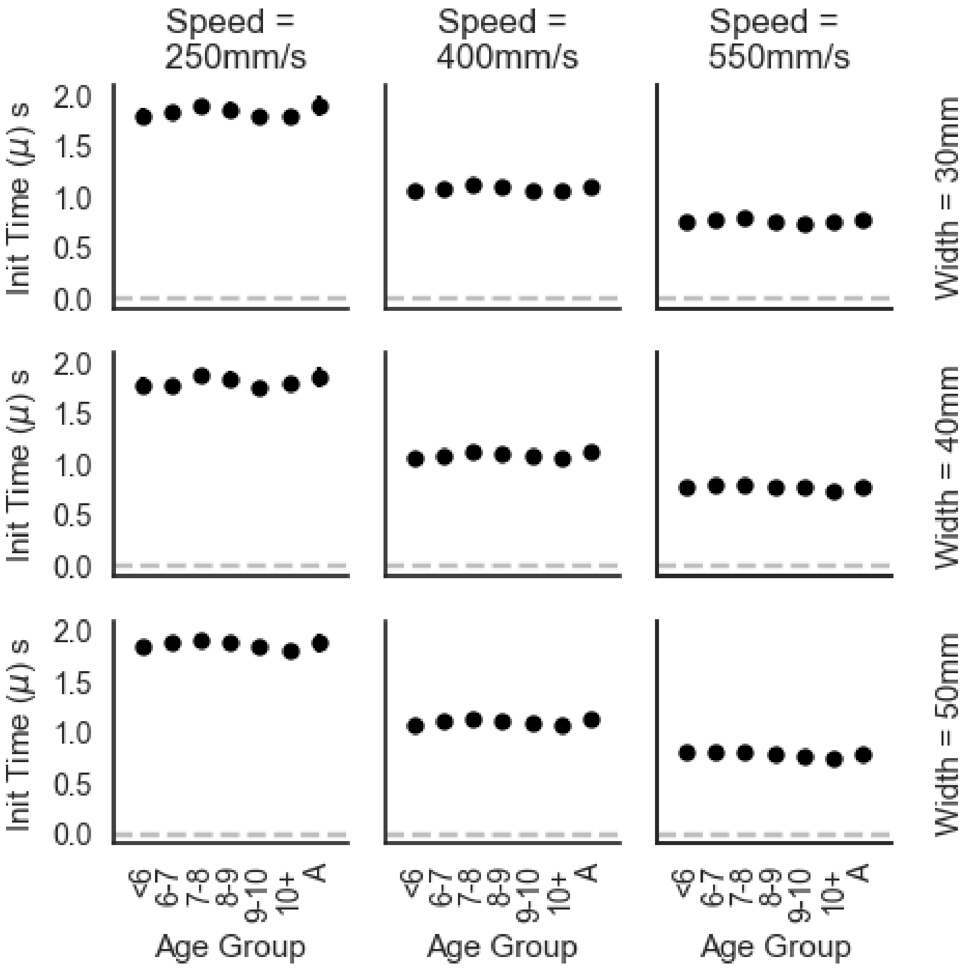

Figure 5.14: Initiation time as a function of age group, target speed (columns) and target width (rows). Points show the posterior mean. Error bars show the 95% HDI (unless smaller than the symbol size).

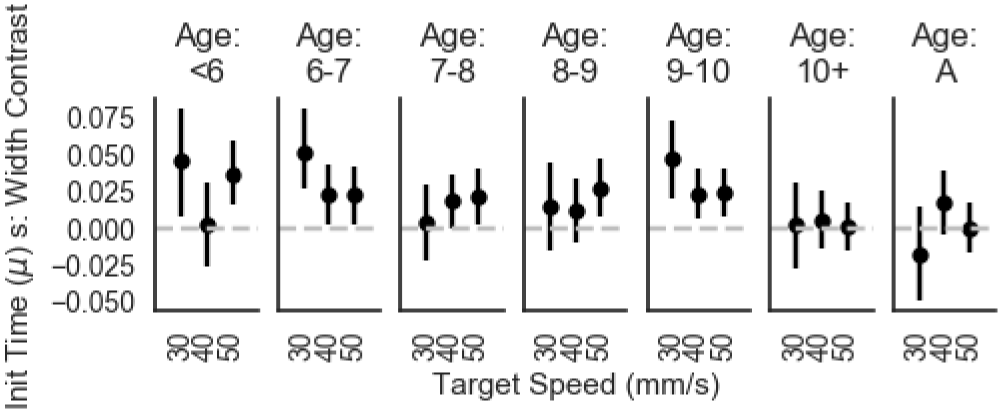

Figure 5.15: Initiation time contrasts: Width. Magnitude of the difference in initiation time between the widest (50mm) and narrowest (30mm) targets for every age group (columns) and target speed (x-axis). Points show the posterior mean. Error bars show the 95% HDI.

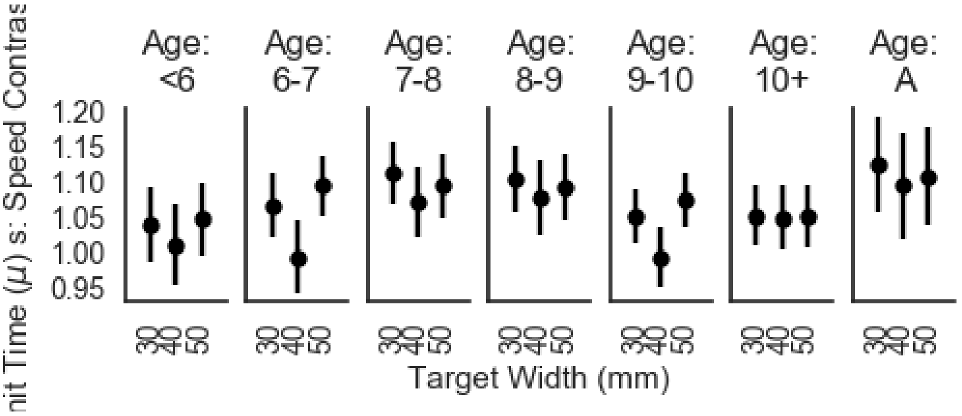

Figure 5.16: Initiation time contrasts: Speed. Magnitude of the difference in initiation time between the fastest (550mm/s) and slowest (250mm/s) target speed for every age group (columns) and target speed (x-axis).Points show posterior mean. Error bars show the 95% HDI..

5.3.8 Controlling for trial order

It is clear from figure 5.3 that performance was generally poorer on the $250\mathrm{mm\,s^{-1}}$ - $40mm$ target than for the faster targets of the same width. This pattern of results is evident in the P-Hit speed contrasts (see figure 5.5), resulting in lower values for the 40mm contrasts than for the 30mm and 50mm contrast. This was almost certainly due to the order of the target blocks with all the $250\mathrm{mm\,s^{-1}}$ - $40mm$ targets appearing within the first three blocks. To attempt to control for this the model was refit with block number as a predictor. The P-Hit speed contrasts when controlling for trial block are shown in figure 5.17. It is clear that the statistical model was unable to account for the bias caused by the order effects. This may be because both blocks of the $250\mathrm{mm\,s^{-1}}$ - $40mm$ targets appeared early in the trial sequence. Performance improved rapidly over these trials so it may be difficult from a statistical perspective to discriminate between learning to perform the task per se, or just performing poorly on these early target types.

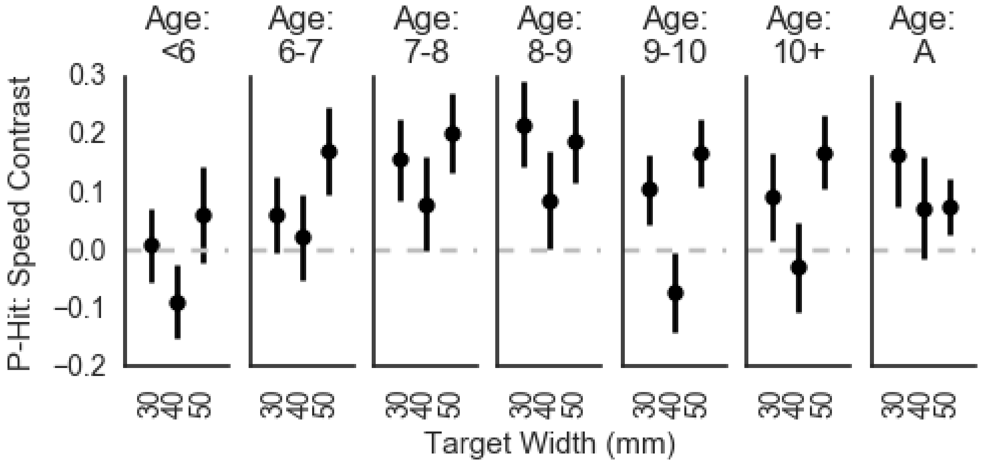

Figure 5.17: P-Hit Contrasts: Speed - controlling for trial block. Magnitude of the difference in P-Hit between the fastest (550mm/s) and slowest (250mm/s) target speed for every age group (columns) and target speed (x-axis). Positive values indicate that P-Hit was higher for the slowest target. Points show posterior mean. Error bars show the 95% HDI.

5.4 Discussion

In this experiment we explored how performance on an interceptive timing task varied over a cross-section of primary school children (ages 5-11 years) and an adult comparison group. We also examined whether changes in the probability of hitting the targets (P-Hit) were due to changes in temporal bias (accuracy), variable error (precision), or both. The data showed clear developmental trends with P-Hit improving gradually with age group. Nevertheless, the adults showed superior performance to even the oldest children. This contrasts with studies employing coincidence-anticipation timing tasks, which have reported that coincidence timing reaches adult levels by 11 years (Haywood, 1980). Thus, the present data provides the first definitive demonstration that interceptive timing abilities are not fully developed by 11 years of age. The data revealed that both bias and variable error reduced over childhood. This suggests that as children become older they become better able to time their actions in response to perceptual input, and develop an improved ability to generate the necessary motor actions in a reliable manner.

Many of the tasks used to measure interceptive timing abilities in children (i.e. the MABC-2) use different tests for different age groups to avoid ceiling and floor effects, making it difficult to compare across age groups. However, the interceptive timing task created for our experiment worked well. By carefully selecting the range of target widths and speeds, the task was designed to allow the same measure to be used with both young children and adults. The task successfully differentiated adults from the youngest children, but provided a challenge for the adults (who only hit 88% of the targets), whilst pitched at a level that allowed the youngest children to succeed (hitting 38% of the targets on average).

Unsurprisingly, the probability of hitting the targets (P-Hit) was greater for wider targets, and this pattern of results was seen for all age groups including the adults. The probability of hitting the target showed a more complex relationship with target speed, which varied as a function of age. For the adults and the older children (with a few exceptions), the probability of hitting the target was higher for slower targets. Surprisingly, the youngest children (ages < 6 years) did not have a higher probability of hitting the slower targets. The reason for this appears to be that the younger children showed a large bias to hit ahead of the slowest target, and a similar but smaller bias was found for the medium speed targets. This bias to hit slower targets earlier than the fastest target was seen in all the age groups except the adults, but the magnitude of the effect decreased with age (i.e. the oldest children showed a very small bias). This finding is consistent with previous reports of children responding early to slow targets in coincidence-anticipation tasks (Benguigui, Broderick, Baurès, & Amorim, 2008), and adults striking targets early even when there is only one target speed (Brenner, Cañal-Bruland, & van Beers, 2013).

The results raise the issue of why there was a bias to hit slower moving targets early relative to faster ones. Optimal performance on our task involves participants detecting the perceptual information specifying the TTA of the approaching target at the interception point, and using this information to time their movements. From this perspective, sensorimotor learning describes the process in which humans become 'tuned' to the sensory invariants that link stimuli to optimal action selection. It seems reasonable to suggest that this ability will improve over childhood - and indeed the present data indicate that this is the case. It is probable, however, that the noise inherent in the stimulus will lead to uncertainty about the TTA on a number of trials. In this situation (i.e. when perceptual uncertainty exists), participants might use a simple heuristic to time their response. One heuristic would be to assume that the object is approaching quickly. Indeed, Rushton and Wann (1999) have shown previously that adults have a bias towards the earliest estimate of TTA when TTA information is available from multiple sensory cues. This heuristic is sensible because the perceptual information associated with faster moving targets will typically be associated with more uncertainty, and in many interceptive timing tasks it is better to be early than late. This heuristic would serve the participants well when faced with fast moving targets, but would lead to systematic biases with slower targets. Furthermore, children's interceptive behaviours may be more influenced by the use of such a heuristic than adults, assuming they are less able to accurately predict the TTA of the target. Indeed, the bias towards hitting the slow targets early (relative to the fast targets) was most apparent in the youngest children and then steadily decreased until it was only just evident within the adult group. Most notably, the youngest group of children showed the greatest effect even when the largest target moved at the slowest speed. These results suggest that the youngest children have difficulty in using sensory information to provide reliable estimates of the TTA of an approaching target, and have an intrinsic bias to assume that targets are approaching quickly. This can explain the otherwise counter intuitive finding that the youngest group of children showed the same levels of performance for the fast targets as for the slow ones.

The data also showed that temporal precision (lower variable error) was better in the older age groups, and for faster targets. The increased precision with target speed can be explained by the decreased movement duration observed with faster targets. It has been shown that there is lower temporal variability associated with faster movements (Brouwer et al., 2002; Tresilian & Plooy, 2006). The data also showed that accuracy and precision did not vary with target width for any of the age groups, including the adults. This contrasts with previous studies where it was found that movement time decreased (and thus precision increased) with smaller targets (Tresilian & Houseman, 2005). The difference between studies might be due to smaller changes in width within the present study or the smaller number of trials (Tresilian & Houseman, 2005). Thus it may be that more trials are required for participants to optimally tailor their MT to smaller changes in target width.

One limitation of the current study is that performance was worse for the middle sized, slowest speed target. As this target type appeared at the start of the sequence of targets this pattern of results was probably due to learning over the course of the experiment. The same fixed order of target presentation was used across all of the participants. This design ensured that we could compare between individuals (see chapters 6 and 7) and age groups. It did mean, however, that the early trials were subject to initial improvements in performance as the participants became increasingly familiar with the task.

These data establish for the first time that children continue to develop their interceptive timing capabilities beyond 11 years of age, given that performance was far from adult levels by this age. The improvements in performance on our interceptive timing task entailed reductions in both bias and improvements in precision (smaller variable errors). Thus, childhood is marked by an improved capacity to tailor actions to perceptual information and a greater ability to reliably execute the motor responses. The present study indicates that young children may rely on a simple heuristic whereby the presence of perceptual uncertainty regarding a target's TTA triggers an earlier movement initiation. This heuristic created systematic errors with slower targets but would have increased the probability of hitting the faster targets (despite the reduced opportunity to use perceptual information because of the shorter viewing window). The presence of such response biases suggests a mechanism through which participants might increasingly tailor their movements to the the task statistics. In the next chapter we move on to examine how the interceptive timing abilities documented in this chapter relate to performance on other sensorimotor tasks.

Chapter 6

Testing a common taxonomy of motor skill in children

Chapter Abstract

Human development is characterised by the acquisition of a large repertoire of sensorimotor abilities. Children progressively learn to crawl, stand and walk over their first few years of life, while also developing the ability to grasp and dexterously manipulate objects. This skill repertoire develops over childhood, both in terms of the number of sensorimotor tasks that children can engage in and also the proficiency with which children perform these tasks. These skills are often subjected to a binary taxonomic scheme in which they are classified as being either ‘fine’ (e.g. writing, grasping) or ‘gross’ (e.g. standing, walking etc). One difficulty with classifying actions in this manner is that many behaviours do not fall neatly within one or the other domain. For example, catching a ball can require both gross movements of the torso and fine coordination of the hands. Thus it may be useful to classify tasks into a larger number of more specific categories. It may be that interceptive timing tasks (e.g. hitting, catching) tap into children's ability to predict target motion, providing additional information regarding a child's 'sensorimotor status' beyond that provided by typical 'fine' and 'gross' motor tasks (i.e. task not involving moving objects). We took three canonical measures of skill within these domains and obtained precise and accurate objective measures of children completing these actions (n = 309). We explored the child's ability to: (i) manipulate a hand held stylus when interacting with visual stimuli; (ii) maintain posture with eyes open and closed; (iii) intercept moving targets. Bayesian analysis indicated that it is meaningful to consider these actions as falling within different categories, though the number of useful domains and taxonomic ranks for classifying motor skills remains to be determined. The results suggest that measures of interceptive timing ability provide a useful measure of 'sensorimotor status' when used alongside 'fine' and 'gross' measures.

6.1 Introduction

There is a bewildering array of skilled behaviours evidenced by adult humans. Humans are able to perform tasks as diverse as controlling vehicles, preparing food, and performing laparoscopic surgery. Even young children show sensorimotor behaviours several orders of magnitude more sophisticated than those currently achieved at the frontiers of robotics, both in terms of the number of tasks they can perform and the proficiency with which they do so. The sheer number of motor tasks in which humans engage makes it difficult to catalogue the specific skills found across and within cultures throughout the world. Yet developing a taxonomy of sensorimotor abilities is of substantial importance to behavioural scientists and clinicians alike. Researchers often address hypotheses pertaining to broad sensorimotor constructs using only a single task, while clinicians attempt to obtain measures of an individual's sensorimotor 'status' using a handful of tasks. Thus understanding how sensorimotor skills relate to one another is of huge practical importance.

As the idea of a 'general intelligence' flourished at the start of the 20th century, researchers began to posit the existence of an analogous 'general motor' ability which underpins performance across a wide range of tasks (Brace, 1927; McCloy, 1934). This hypothesis was later challenged (Bachman, 1961), yet the literature still often groups performance on different motor tasks under broad categories of abilities. One commonly deployed scheme categorises motor skills as being either 'fine' or 'gross' in nature. Skills classified as being 'fine motor' generally involve dexterous manual manipulation whilst 'gross motor' is typically used to describe actions involving movement of the torso and locomotion (e.g. running, walking, standing).

There are some difficulties with such a taxonomic scheme. Flatters, Mushtaq, Hill, Holt, et al. (2014) suggested that categorising motor tasks as either 'fine' or 'gross' may not reflect the complexity of many tasks, as many skilled behaviours require both fine and gross actions. For example, many manual dexterity tasks require postural adjustments to be made, owing to the consequent shifts in the centre of gravity following arm movements (Huang & Brown, 2013; Thelen & Spencer, 1998). Indeed, postural stability is likely to be a pre-requisite for obtaining reliable sensory information (holding the head steady so that clear and stable vision is available), information which is vital for the performance of many manual dexterity tasks. In support of this notion, it has been shown that individuals reduce their postural sway when performing a precise manual task (Balasubramaniam, Riley, & Turvey, 2000), and the amount of observed sway varies with the spatial precision requirements of the supra-postural task in both adults (Haddad, Ryu, Seaman, & Ponto, 2010) and children (Flatters, Mushtaq, Hill, Rossiter, et al., 2014). The interactions that must exist between postural and manual control suggest that the development of these different skills will have a degree of synergy, and thus skills in one domain will not be independent of skills in the other. Nevertheless, there appears to be face validity in distinguishing between 'fine' and 'gross' skills as there are a number of tasks that involve control of the torso but little involvement of the hands, and vice versa.

In order to test the empirical support for the binary classification scheme, Flatters, Mushtaq, Hill, Holt, et al. (2014) tested a large sample of children on a test of manual dexterity (Culmer et al., 2009) and a test of postural stability (Flatters, Culmer, Holt, Wilkie, & Mon-Williams, 2014). Flatters, Mushtaq, Hill, Holt, et al. (2014) found there were weak to moderate correlations between the tests, supporting the idea that gross and fine motor abilities have a degree of interdependence, but indicating that it is meaningful to consider actions as falling within one or other category. Other lines of evidence suggest that some children display poor abilities in either fine or gross motor tasks (Zwicker, Missiuna, Harris, & Boyd, 2012), whilst a study employing confirmatory factor analysis gave support to fine and gross motor abilities reflecting at least partially separable constructs (Schulz et al., 2011). In summary, classifying tasks as being either gross or fine appears to have some empirical support, although the two are at least partially related (Flatters, Mushtaq, Hill, Holt, et al., 2014).

The usefulness of the binary classification scheme relates to the identification of children with motor difficulties where the problems differentially affect gross or fine motor skills. It is therefore unsurprising that the common binary motor skill classification scheme is reflected in the design of standardised tests of motor ability (Barnett et al., 2007). One of the most popular tools for assessing motor skill in children is the Movement Assessment Battery for Children -2 (MABC-2). First published in 1992, the original MABC was developed from a Test of Motor Impairment. The test was designed to assess the movement skills of children ages 4-12 years, with separate test items for different age groups. While the items vary by age group they involve similar skills. The MABC originally had 4 age groups, but this was reduced to three in its successor, the MABC-2 (Barnett et al., 2007). The MABC / MABC-2 has generally been found to have good test-retest reliability, with Croce, Horvat, and McCarthy (2001) reporting intra-class correlation coefficients of 0.92-0.98 depending on the age group being examined. Similarly high values have been reported when examining children in Hong Kong (Chow, 2003). In addition, it has been found that physiotherapists show high inter-rater reliability, with kappa coefficients in the region of 0.99 to 1. In addition the MABC-2 has been found to be responsive to physical therapy interventions in children with developmental coordination disorder (Wuang, Su, & Su, 2012).

The MABC-2 groups actions into ‘manual dexterity’, ‘static and dynamic balance’ and ‘ball skills’. The division of tasks into these categories was justified initially on the basis of subjective “common sense and clinical experience” (Henderson & Barnett, 1998), and there was only limited empirical evidence to support such a classification (Schulz et al., 2011). The work of Flatters, Mushtaq, Hill, Holt, et al. (2014), however, provides some support for this division, with the ‘manual dexterity’ section of the MABC-2 capturing ‘fine’ motor abilities and the ‘static and dynamic balance’ tasks reflecting ‘gross’ motor performance.

There is less empirical support for a separate category of ‘ball skills’ within standardised assessment batteries. For example, the MABC-2 has a ball skill section which includes tasks involving bouncing and catching a ball, or throwing a ball against a wall and catching it (the task varying as a function of age group). The difficulty is that such tasks have concurrent demands on other postulated categories of movement (i.e. ‘gross’ and ‘fine’). Many catching tasks involve complex coordination of the head, neck, trunk, arms and legs, and may require locomotion to a position from which the target can be caught (McLeod & Dienes, 1996). In addition, rapid arm movements require prospective and reactive adjustments in posture (Van Der Fits, Klip, Van Eykern, & Hadders-Algra, 1998), and thus poor posture may limit catching performance. In support of this notion, it has been found that children who lag behind their peers in terms of catching performance can show significant improvements in catching ability when provided with external postural support (Savelsbergh, Bennett, Angelakopoulos, & Davids, 2005). In addition, it has been suggested that the MABC-2’s gross motor measures correlate moderately with the measures of ball skill in young children, although no inferential statistics or measures of uncertainty were provided for the correlations (Schulz et al., 2011). Likewise, skilled catching and throwing behaviours require precise control over the fingers of the hand. These considerations raise questions about the validity of introducing a third division into the classification of motor skill.

There is, however, a critical difference between ball skills and many other 'gross motor' tasks which can provide a theoretical justification for creating a different taxonomic class. Catching tasks can be considered to fall within a category of 'interceptive timing' behaviour (Tresilian, 2005). Interceptive timing skills can be classified as those actions which involve interacting with objects that are in motion relative to the actor, including catching or hitting moving objects as well as object avoidance (e.g. avoiding cars while crossing a road). Interceptive timing abilities are fundamental to performing many activities that are ubiquitous in daily lives. The successful interception of a moving object requires the interceptive effector (e.g. hand, bat) to arrive at the same spatial location at the same time as the object. This requires movements to be aimed towards a position through which the target will pass at some point in the near future. Thus, the ability to predict the future motion of moving objects underpins performance on many interceptive timing tasks (Zago et al., 2009). In addition, most models of interceptive timing suggest that individuals must estimate the time remaining until the target reaches a specific point along its trajectory where it can be intercepted (Tresilian, 2005). These specific task attributes suggest that 'ball skill' might be a useful taxonomic category as these tasks have a fundamental component (interceptive timing) that is not captured in many canonical examples of fine and motor ability.

The usefulness of the tripartite scheme of classification ultimately hinges on an empirical demonstration that these three domains can be empirically distinguished. We therefore tested a large sample of children (ages 4-11 years) on an interceptive timing task, a manual dexterity battery (fine motor) (Culmer et al., 2009) and a test of postural stability (gross motor) which required participants to stand still with both eyes open and eyes closed (Flatters, Culmer, et al., 2014). The manual dexterity battery included three uni-manual fine motor tasks. One of these tasks involved tracking a moving target with a stylus. Like interceptive timing tasks, performance on tracking tasks is thought to depend on the ability to predict the motion of the tracked target (Barnes & Asselman, 1991; Dallos & Jones, 1963), and thus these tasks may tap overlapping sensorimotor constructs. Thus we predicted that while the largest correlations between tasks would be found within the manual dexterity battery, a larger correlation should be found between the interceptive timing task and the tracking task than between interceptive timing and the other fine motor measures which did not involve moving targets. We expected only small correlations between tasks in the different domains. If the domains turn out to be indistinct, as demonstrated by large correlations between all tasks, then it would suggest that less extensive testing is required to identify children with motor problems.

6.2 Methods

6.2.1 Participants

Participants were recruited from a state primary school in Bradford, West Yorkshire, UK. From the 368 children at the school, 309 full data sets were included in the data analysis. More information on the participants can be found in chapter 5.

6.2.2 Apparatus

Interceptive timing task (IntT)

Details of the interceptive timing task can be found in chapter 5. Participants performed 54 trials in which the target speed (250mm s^{-1}, 400mm s^{-1}, 500mm s^{-1}) and target width (30mm, 40m, 50mm) varied (9 trial types x 6). Each target type was presented in a block of 3 trials, with 2 blocks for each trial type. The blocks were pseudorandomly ordered with the constraint that two blocks of the same kind could not occur sequentially. All participants experienced an identical sequence of blocks. The number of targets hit (IntT) provided a simple measure of interceptive timing performance

Manual dexterity

Manual dexteriy was measured using the Clinical Kinematic Assessment Tool (CKAT) (Culmer et al., 2009). CKAT consists of three sensorimotor tasks that are presented on a tablet computer screen (Toshiba Portege M700-13p tablet, screen: 260 × 163mm, 1200 × 800 pixels, 60Hz refresh rate) and completed using a hand-held stylus. The planar position of the stylus was recorded at 120Hz and smoothed using a 10Hz dual-pass Butterworth filter at the end of each testing session.

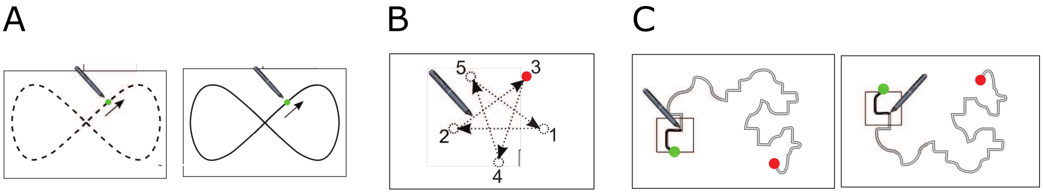

Figure 6.1: CKAT Tasks A) Tracking task. Participants followed a green dot with the stylus. In the first trial the dot followed the dotted (invisible) path. In the second trial the guide track was visible. In each trial the dot made three revolutions of the figure of eight pattern at each speed: fast, medium and slow. B) The Aiming Task. Participants made movements to sequentially appearing targets. C) Tracing task: Participants traced a path using the stylus, while staying within a moving box.

CKAT - tracking task (with and without spatial guide)

Participants completed two trials in the tracking task (figure 6.1A). In the first trial, they placed the stylus on a static dot (10mm diameter) displayed on the centre of the screen. After one second the dot began to move across the screen in a figure-of-8 pattern. Participants were instructed to keep the tip of the stylus as close as possible to the dot's centre for the duration of the trial. The dot completed nine revolutions of the figure-of-8 pattern. The dot moved at a 'slow' pace during the first three revolutions. In the next three revolutions the dot moved at a 'medium' pace and in the last three the dot moved at a 'fast' pace (see Flatters, Mushtaq, Hill, Rossiter, et al. (2014) for details). Participants then completed a second trial which was identical to the first except that a black 3mm wide ‘guide’ line was displayed on the screen, indicating the path which the dot would follow.

The root mean square error (RMSE) provided a measure of the participant’s spatio-temporal accuracy, where the error was the straight line distance in millimetres between the centre of the target dot and the stylus. A separate RMSE score was calculated for each target speed within each trial. The median value of these was taken to provide an overall measure of performance on the tracking task.

CKAT - aiming task

The aiming task (figure 6.1B) required participants to make 50 aiming movements to sequentially appearing circular targets (5mm diameter). Once the participant successfully moved the stylus to the target dot then that target disappeared and the next target appeared (see Flatters, Mushtaq, Hill, Rossiter, et al. (2014) for details). Movement time (MT) was the measure of interest and was defined as the time between arriving at one target location and arriving at the next, defined as when the stylus entered the target. The mean MT over the first 50 trials provided our measure of ‘aiming’ performance.

CKAT - tracing task

The tracing task required participants to trace a path displayed on the tablet (figure 6.1C). A box moved along the path every 5 seconds. Participants were told to trace the path as accurately as possible while ensuring they stayed within the moving box at all times. At each time point (120Hz) the minimum 2D distance between an idealised reference path and the stylus was calculated.

The arithmetic mean was calculated for these values across each trial, giving a measure of path accuracy (PA). The ideal trial time if the participant remained within the moving box was 36 seconds. To normalise path accuracy for task time, PA was inflated by the percentage that participant's actual MT deviated from the ideal 36 seconds value (adjusted path accuracy). Adjusted path accuracy was then used as the measure of performance on the tracing task.

Postural control task

Postural movements were measured using a custom built motion-capture rig (Flatters, Mushtaq, Hill, Rossiter, et al., 2014; Flatters, Mushtaq, Hill, Holt, et al., 2014). Children stood with their feet shoulder width apart on a Nintendo Wii Fit board, which recorded the participant's centre of pressure (COP) at 60Hz. The data was filtered using a wavelet filter as described in Flatters, Culmer, et al. (2014). The 2D path length subtended by the COP (in mm) provided a measure of postural stability. Two measurements were taken. In the first the children stood fixating a target on a wall directly in front of them (posture eyes open). In the second task the children stood with their eyes closed (posture eyes closed).

6.2.3 Statistical models

Correlation model

Several Bayesian correlation models were specified to explore the relationships between performance on the motor tasks. The first two models were used to examine the correlations between performance on all the motor tasks without controlling for age. The first of these modelled the data as being distributed by a multivariate normal distribution (normal model), while the second modelled the data as coming from a multivariate student's t-distribution (robust model). We assumed that this model would be more robust to outliers in the data. The normal model with priors was defined as follows:

$$\mu \sim \text{Normal}(0, 10) \tag{6.1}$$

$$\tau \sim \text{Cauchy}^{+}(0, 10) \tag{6.2}$$

$$\Omega \sim \text{LKJCorr}(2) \tag{6.3}$$

$$\Sigma = \text{diagMatrix}(\tau)\Omega\text{diagMatrix}(\tau) \tag{6.4}$$

$$y \sim \text{multiNormal}(\mu, \Sigma) \tag{6.5}$$

where Ω is a correlation matrix, τ is a vector of coefficient scales and μ is a vector of variable means. Thus the model naturally decomposes the covariance matrix into a correlation matrix (Ω) and vector of scale parameters (τ) (Manual, 2013). Priors were chosen to be weakly informative, based on the scale of the data. A half cauchy prior was placed over τ, and an LKJCorr prior was placed over Ω, as recommended by Gelman (Manual, 2013).

The robust model was identical except the multivariate normal likelihood (equation 6.5) was replaced with a multivariate student's t distribution,

$$y \sim \text{multiStudent}(\nu, \mu, \Sigma) \tag{6.6}$$

and the degrees of freedom prior was given a prior distribution,

$$\nu \sim \text{exponential}(\frac{1}{30}) \tag{6.7}$$

Partial correlation model

A final model was specified that allowed for correlations between variables to be examined after controlling for age[1] (partial age model). This model was identical to the robust-model except that the mean vector μ (equation 6.1) was replaced with a mean for each participant that depended on a matrix of regression coefficients (β) and the participant's age as follows,

$$\beta \sim \text{Normal}(0, 10) \tag{6.8}$$

$$\mu_i = \beta x_i \tag{6.9}$$

where x was a design matrix with age as the only predictor. This allowed us to partial out age, leaving the correlation matrix Ω as the correlations between all the motor task variables after controlling for age. Several versions of the correlation models were fit using different data transformations and compared against one another, as described in the results section.

The priors were informed only by the scale of the data. Bayes rule was used to estimate the credible values of each of the model's parameters (θ) given the data. The joint posterior distribution is given by

---

[1] The *robust model* was also use to explore the relationships between variables after controlling for age (see the z-model in the results section)

$$P(\theta|y) \propto P(y|\theta)P(\theta) \tag{6.10}$$

A representative sample was drawn from the posterior using the ‘No-U-Turn sampler’ (Hoffman & Gelman, 2011) implemented in PyStan 2.14 (Stan Development Team, 2016). Four chains of 5000 samples (warmup N = 2500) were started at random values in the joint posterior distribution for each model. Convergence was assessed by visually examining the chains and computing $\hat{R}$ and effective sample size for each parameter.

6.3 Results

6.3.1 Motor task correlations

The normal and robust models were fit several times with different data transformations applied. As the data were not normally distributed this was done to explore how to best model the data. The first model applied no transformation and used the normal-model specification, and is referred to as the identity model (ID normal). The second two models transformed the data using the natural logarithm and used the normal and robust specifications (log-normal, log-robust). A final model was fit to the transformed data using the common logarithm (log10-robust) and only used the robust specification.

Approximate leave-one-out cross validation was used to compare the expected out-of-sample predictive accuracy of the models (Vehtari, Gelman, & Gabry, 2015), as defined by the expected log pointwise predictive density (*elpd*). This provides a simple method of model comparison. First the *elpd* was compared for the ID normal and log-normal models, revealing a strong preference for the latter ($elpd_{\text{diff}} = 3568.37$, $SE = 100.477$). This was unsurprising given that the most of the outcome measures were skewed positively. The log-normal model was then compared to the log-robust model, revealing a preference for the latter ($elpd_{\text{diff}} = 152.08$, $SE = 49.56$), which suggests that this model was better accounting for outliers in the data set. Finally, we compared the log-robust model to the log10-robust model, revealing a preference for the later ($elpd_{\text{diff}} = 1546.66$, $SE = 0.17$), suggesting that the log10-robust transformation was the best. Thus, the results of the log10-robust model were used to make inferences from the data.

Figure 6.2 shows a matrix of plots revealing the marginal posterior distributions over the correlation coefficients in the log10-robust model. There was only a very small correlation between IntT and both posture measures. The 95% HDIs spanned zero, although it was a negative relationship for the IntT and posture eyes open measure. Larger correlations were seen between IntT and the three manual dexterity measures. The posture measures showed small to medium correlations with the manual dexterity measures. These correlations are unsurprising given that all of these measures are expected to improve with age. Therefore, the more interesting question is whether these correlations hold once we control for age.

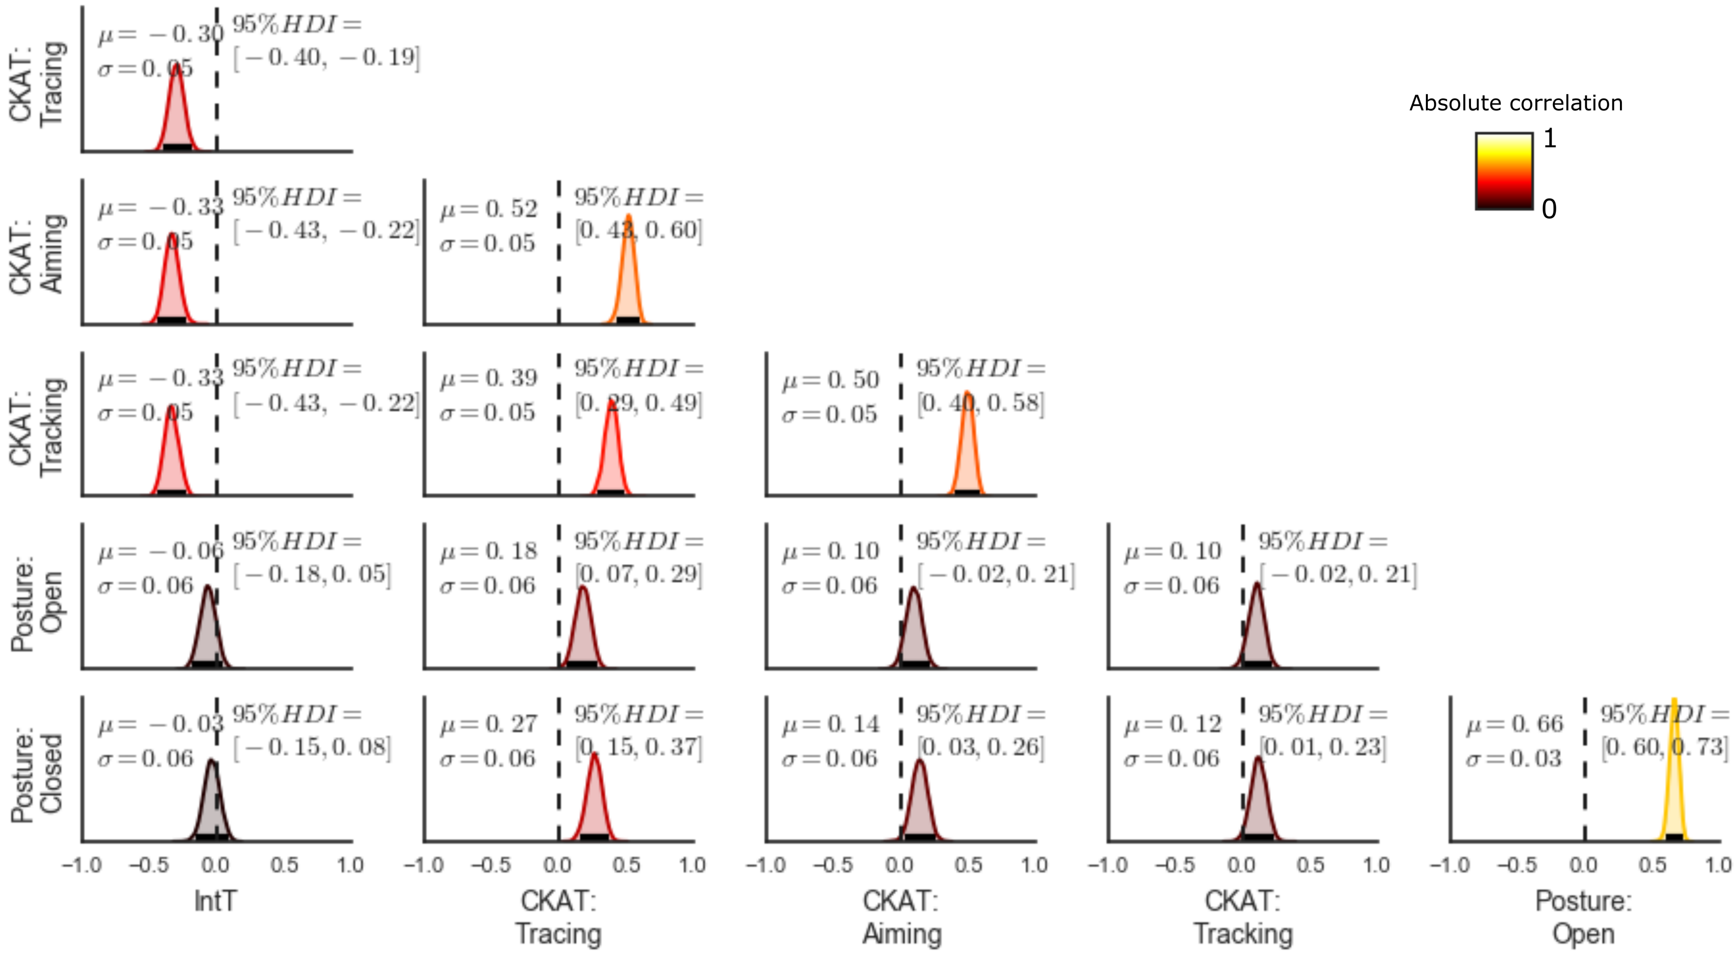

Figure 6.2: Bayesian correlation matrix. Marginal posterior estimates over the correlation coefficients for each motor task pair. The x-axis value denotes the correlation coefficient. Posterior means (μ), standard deviations (σ) and 95% HDI are shown in each panel. Dotted lines show the zero point. Colours indicate the strength of relationship according to the posterior mean, with brighter colours indicating a stronger relationship (see colour bar). Black horizontal bars indicate the 95% HDI.

6.3.2 Age controlled correlations

Two models were developed to account for age. The first was exactly the same as the log10-robust model except that the data were split into groups by age and then transformed to z-scores, using the means and standard deviations of each age group (log10-Robust-Age). This model is analogous to the analysis methods employed by (Flatters, Mushtaq, Hill, Holt, et al., 2014). For the next analysis, age was used as a predictor in a linear regression model with all the motor tasks as the outcome measures, allowing examination of the relationships between these variables after accounting for age (partial correlation coefficients). Again we compared these models using approximate leave-one-out cross validation and found that the partial correlation model was favoured ($elpd_{\text{diff}} = 2993.97$, SE=19.11). We therefore made inferences from this model.

For clarity we split the correlation matrix plots for the partial correlation model into separate figures. Figure 6.3 shows the marginal posterior distributions over the correlations between the three CKAT measures and figure 6.4 shows the correlation between the two posture measures (eyes open and eyes closed). Weak to moderate relationships were found between performance on the CKAT tasks and a large correlation was found between the two posture measure (eyes open and eyes closed).

Figure 6.5 shows the relationship between the interceptive timing measure (IntT) and all other motor measures. After controlling for age small correlations between IntT and the CKAT aiming and tracing tasks were likely. As predicted, the correlation between IntT and the CKAT tracking

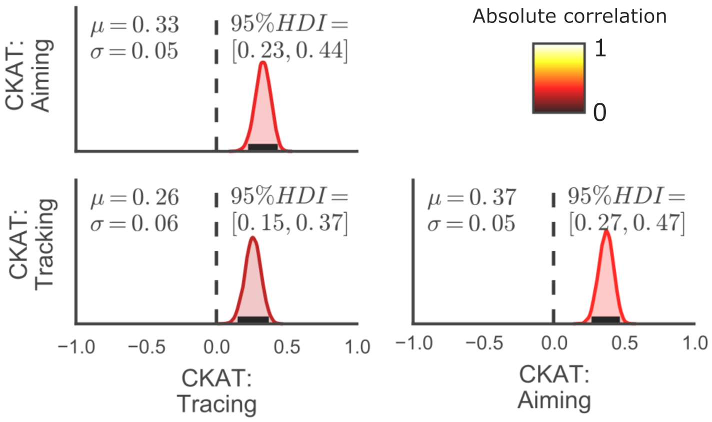

Figure 6.3: Bayesian partial correlation plots. Marginal posterior estimates over the correlation coefficients between the CKAT tasks after controlling for age. Posterior means (μ), standard deviations (σ) and 95% HDI are shown in each panel. Dotted lines show the zero point. Colours indicate the strength of relationship according to the posterior mean, with brighter colours indicating a stronger relationship. Black horizontal bars indicate the 95% HDI.

task appeared to be larger (posterior mean -0.22) than the relationship better IntT and aiming and tracing, with the 95% HDI spanning a higher range of values. Contrasts between the IntT-aiming and IntT-tracking partial correlation coefficients confirmed that the IntT-tracking correlation was larger (contrast mean = -.07, SD = 0.07, 95% HDI = [-.20, .06], 50% HDI = [-.12, -.27], η = .16), and the same was found when comparing the IntT-tracing and IntT-tracking correlation coefficients (contrast mean = -.09, sd = 0.07, 95% HDI = [-.24, .05], 50% HDI = [-.14, -.43], η = .11). Figure 6.5 suggested that only very small correlations were likely between IntT and the posture measures. The posterior means were close to zero and the 95% HDI spanned zero.

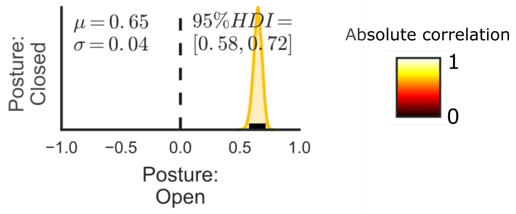

Figure 6.4: Bayesian partial correlation plots. Marginal posterior estimates over the correlation coefficients between the posture tasks (eyes open; eyes closed) after controlling for age. Posterior means (μ), standard deviations (σ) and 95% HDI are shown in each panel. Dotted lines show the zero point. Colours indicate the strength of relationship according to the posterior mean, with brighter colours indicating a stronger relationship. Black horizontal bars indicate the 95% HDI.

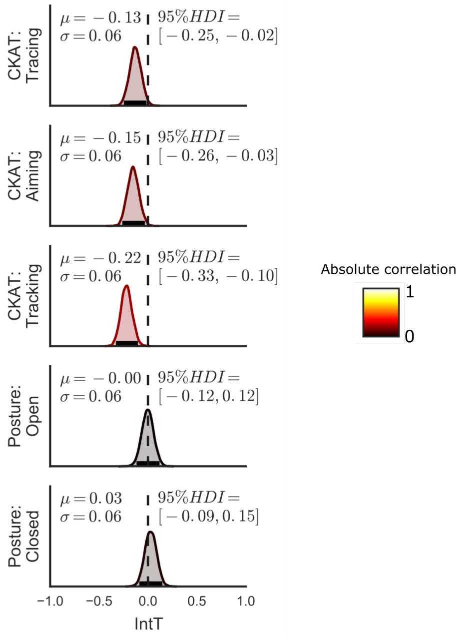

Figure 6.5: Bayesian partial correlation plots. Marginal posterior estimates over the correlation coefficients between the IntT task and all other motor measures (CKAT and posture) after controlling for age. Posterior means (μ), standard deviations (σ) and 95% HDI are shown in each panel. Dotted lines show the zero point. Colours indicate the strength of relationship according to the posterior mean, with brighter colours indicating a stronger relationship. Black horizontal bars indicate the 95% HDI.

Finally the correlations between the CKAT measures and both posture measures are shown in figure 6.6. The marginal posterior distributions revealed that the correlations were likely positive but very small, although a larger correlation was found between the CKAT tracing task and the eyes-closed posture task.

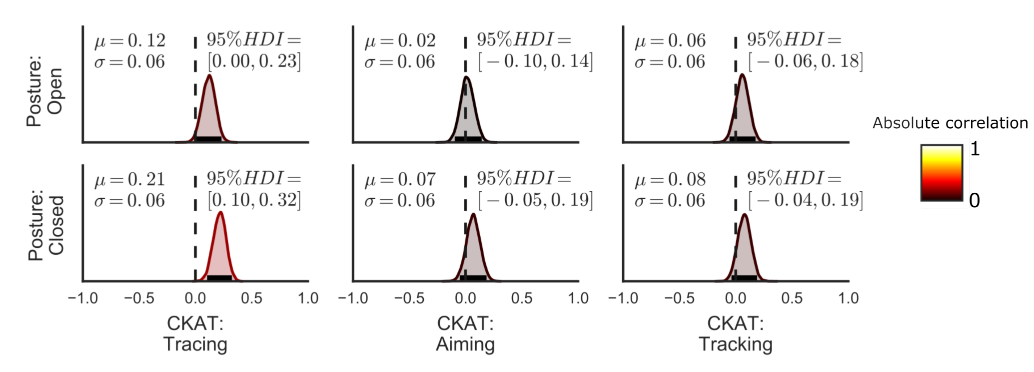

Figure 6.6: Bayesian Partial Correlation Plots. Marginal posterior estimates over the correlation coefficients between the CKAT and posture tasks after controlling for age. Posterior means (μ), standard deviations (σ) and 95% HDI are shown in each panel. Dotted lines show the zero point. Colours indicate the strength of relationship according to the posterior mean, with brighter colours indicating a stronger relationship. Black horizontal bars indicate the 95% HDI.

To examine how the results would differ if the z-scored model was used instead of the partial correlation model, the results of both are plotted in figure 6.7. The filled coloured curves show the correlations in the partial correlation model, while the black curves show the correlations from the z-scored model. It is clear that there was very close agreement between both models.

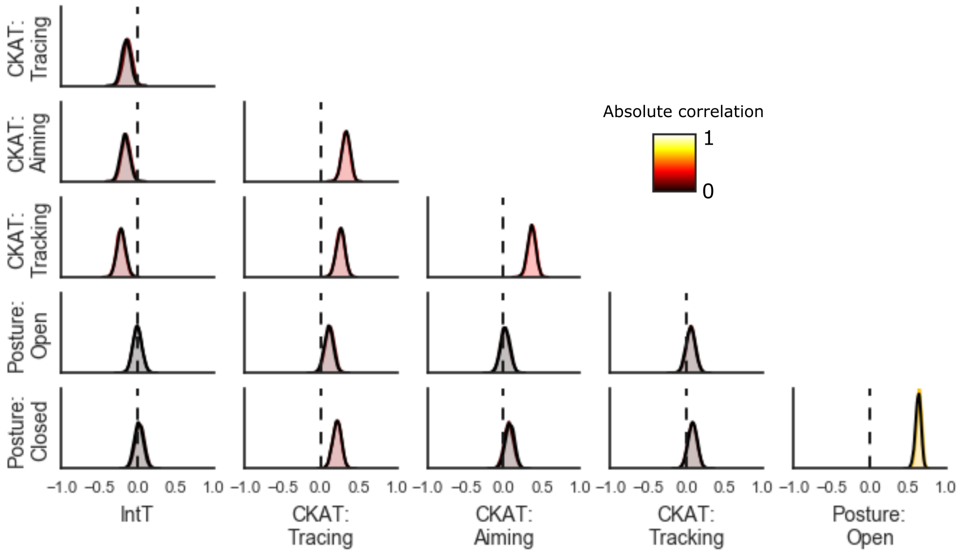

Figure 6.7: Bayesian Partial Correlation Matrix. The filled coloured plots show the marginal posterior estimates over the correlation parameters in the partial correlation model. The solid black lines show the same for the z-scored model. Both posteriors shows very close agreement.

6.4 Discussion

We tested a large sample of primary school children on an interceptive timing task whilst also taking measures of manual dexterity and postural control. As expected the largest correlations were found between the subtasks of the CKAT battery and between the two posture measures. The results suggested that performance on the interceptive timing task did not allow inference about performance on the posture tasks. In other words, there was support for the hypothesis that these tasks fall within different domains (operationalised within the MABC-2). A number of ball skills require the individual to control posture whilst engaged in an interceptive task, and this could conceivably give rise to a strong relationship between interceptive timing and postural control ability, but this was not found to be the case. Notably, we minimised the postural demands of the interceptive task by ensuring that the task was performed whilst the participant was seated. This arrangement decreased the need for participants to make postural adjustments relative to standing, where adjustments are required to compensate for the forces generated when making rapid arm movements (Massion, 1992). In contrast to postural control, we did find a relationship between interceptive timing performance and manual dexterity. This relationship was seen primarily with the tracking task, which may be due to this task having an element of temporal prediction (as we discuss later). However, the size of the relationships between interceptive timing performance and manual dexterity was small, which provides support for the idea that these tasks fall within different categories. Finally, there were small to moderate correlations found between posture and manual dexterity, suggesting a degree of dependency, but again supporting the idea that these skills can be categorised separately.

The moderate correlations found between posture and manual dexterity provide support for the conclusions drawn by Flatters, Mushtaq, Hill, Holt, et al. (2014). Flatters et al., reported statistically significant relationships between centre of pressure measures (eyes open and eyes closed) and the tracking task, but not the aiming and tracing task. The correlation coefficients reported by Flatters et al., are broadly consistent with the posterior distribution of credible values reported in the current study. Our study suggests that very small positive correlations are most likely between the postural measures and the aiming and tracing task. For the tracing task, Flatters et al. reported substantially larger correlations when posture was measured using a head tracker, possibly because the centre of pressure measurements act as a proxy for changes in the centre of gravity, but do not map directly to postural maintenance (Flatters, Hill, Williams, Barber, & Mon-Williams, 2014). Thus, it is possible that a stronger relationship exists between manual dexterity and postural stability than we were able to detect with the measurements used in the present study.

The results of the present study provide support for the notion that motor skills can be usefully classified as falling within three domains of 'gross', 'fine' and 'interceptive timing'. The findings therefore give empirical justification for the use of the categories of 'manual dexterity', 'static and dynamic balance' and 'ball skills' within the Movement ABC-2. This does not, however, allow the conclusion that this is an optimal taxonomic classification scheme. It might be argued, for instance, that a greater number of categories is required. In this context, it is worth exploring the extent to which tasks within the categories correlate with each other. It has been shown previously that catching and throwing tasks are moderately correlated (Van Waelvelde, De Weerdt, De Cock, & Smits-Engelsman, 2004a). Likewise, performance on a catching task has substantial correlations with other interceptive timing tasks (Van Waelvelde, De Weerdt, De Cock, & Smits-Engelsman, 2004b). These results suggest that there is some merit to the use of a general category of 'ball skill'. The present study showed large correlations between the eyes open and closed postural measures, while moderate correlations were found between the manual dexterity measures, suggesting these tasks can be placed within a 'gross' and 'fine' motor category respectively. Thus, there is empirical support for the broad domains encompassing a number of related tasks.

The findings from the present study suggest, however, that the tri-domain classification scheme may have limitations, with some tasks showing a mixed pattern of cross- and within-domain correlations. Our results showed that performance on the interceptive timing task had a stronger correlation with the tracking task than with the aiming and tracing tasks. This makes sense as both the tracking and interceptive timing task involved interacting with a moving target. Thus, both the interceptive timing and tracking task may benefit from an underlying ability to predict how objects move through space and time, extrapolating from the object's current position to its likely future location at specific points of time. This suggests that a better taxonomic classification scheme might be based around the sensorimotor processes (e.g. tracking, steering, intercepting approaching targets, using visual information to control posture etc,) that underpin motor tasks. It remains to be seen whether actions might be better classified according to their underlying sensorimotor processes, at least within one rank of a taxonomic scheme. The advent of recent technology (such as that used within the present study) makes precise measures of these processes a realistic proposition and thus it is possible to address these issues in future research.

In summary, the present results support the broad classification of motor skills within the three domains of 'gross', 'fine' and 'ball skills', as operationalised in tests such as the Movement Assessment Battery for Children. Nevertheless, it remains to be seen whether adoption of these domains is sufficient for pinpointing childhood motor deficits. For example, there is growing interest in the concept of 'fundamental movement skills' where specific actions are identified as key developmental building blocks for the majority of human movement skills (Robinson et al., 2015). From this viewpoint, it could be argued that ‘gross’ motor skills might be better assessed through investigation of a larger number of fundamental actions (such as walking, running, jumping, hopping, kicking and climbing) rather than sampling a sub-set of these behaviours. In a similar vein, ‘fine’ motor skills might be better assessed through a detailed exploration of fundamental actions such as reach-to-grasp behaviour and the control of fingertip forces. The data collection and Bayesian analysis techniques reported in the current study will allow future exploration of such issues.

In this chapter it was found that the interceptive timing task provides a measure of motor ability, which is somewhat distinct from that of ‘fine’ and ‘gross’ motor measures. In the next chapter we explore how this distinct measure relates to academic attainment, specifically in mathematics.

Chapter 7

Interceptive timing and mathematics

Chapter Abstract

Interceptive timing is a fundamental ability underpinning numerous actions (e.g. ball catching), but its development and relationship with other cognitive functions remains poorly understood. Piaget (1955) suggested that children need to learn the physical rules that govern their environment before they can represent abstract concepts such as number and time. This leads to the hypothesis that the neural processes involved in learning how objects move in space and time could underpin the development of abstract representations related to these concepts (i.e. mathematics). To test this hypothesis, we captured objective measures of interceptive timing abilities in 309 primary school children (5-11 years), alongside national standardised academic attainment scores and general motor skill. Bayesian estimation showed that performance on an interceptive timing task predicted mathematical ability

even after controlling for motor skill and age. These findings suggest that attempts to educate children should not neglect sensorimotor development within the physical body that houses the cognitive phenotype.

7.1 Introduction

Interceptive timing is a fundamental human sensorimotor ability that underpins actions where the goal is to make contact with a moving target (e.g hitting a baseball). These tasks require both spatial and temporal accuracy, and proficiency in these tasks appears later in a child's developmental history than skills with minimal temporal constraints (e.g. reaching to static objects). Neurologically intact adult humans show exquisite precision in interceptive timing, with elite baseball batters able to swing their bat to a spatial accuracy of ±1.5cm and a temporal accuracy of ±10ms (Tresilian, 1999). The interceptive timing skills of humans are a testimony to the incredible learning capacity of the sensorimotor system and its ability to overcome the challenges involved in controlling over 600 muscles with the inherent difficulties of nonlinearity, nonstationarity, information delays, and noise whilst operating within an uncertain world (Franklin & Wolpert, 2011). The temporal delays involved in processing perceptual information and transmitting motor commands are particularly problematic in interceptive timing tasks and require the individual to make predictions about where the object and the limb will be at the time of desired contact (Tresilian, 2012). These predictions require precise estimates of how the object will move over time, together with state estimates of the neuromuscular system.

It is widely believed that sensorimotor prediction relies on internal models within the sensorimotor system. Internal models allow for prediction of object motion through space and time (Merfeld, Zupan, & Peterka, 1999) with forward models used to estimate the sensory consequences of motor commands (Flanagan & Wing, 1997; Wolpert, Miall, & Kawato, 1998). Thus the development of these models is central to the ontogenetic acquisition of interceptive timing skills. The deleterious impact of developmental delays in motor prediction can be readily imagined with regard to a child's ability to engage in physical activity. But it is possible that sensorimotor impairments have consequences for a child's cognitive capabilities in a manner that is not so readily appreciated by educational authorities. Such proposals are consistent with the view that the phylogenetic emergence of higher-order cognitive abilities was built upon the evolutionary platform provided by the motor system, particularly with respect to estimating the future state of the environment and physical body (Desmurget & Grafton, 2000).

The idea that higher-order cognitive processes emerged from sensorimotor abilities is attractive (M. Wilson, 2002). It has been suggested that the fundamental importance of sensorimotor substrates to cognition extends both to the individual as well as the species, with Piaget (1955) suggesting that ontogeny recapitulates phylogeny in this regard. Thus, Piaget proposed that sensorimotor interactions with the environment underpin the development of cognitive representations, including our understanding of number. This idea has received a surge of support over the last decade, with evidence that abstract representations of number are grounded in early interactions with objects and an understanding of physical space (de Hevia & Spelke, 2010; Nieder & Dehaene, 2009). It appears that number representations are spatially orientated (Fias, van Dijck, & Gevers, 2011) with representations of number and space sharing overlapping neural circuitry (Hubbard, Piazza, Pinel, & Dehaene, 2005) and being closely related to representations of time (Bueti & Walsh, 2009; Burr, Ross, Binda, & Morrone, 2011; Chang, Tzeng, Hung, & Wu, 2011; Lourenco & Longo, 2010; Srinivasan & Carey, 2010).

The putative relationship between sensorimotor ability and cognitive capacity leads to the hypothesis that a child's interceptive timing skills will be related to the development of their ability to represent space, time and number (i.e. their mathematical ability). A robust test of this hypothesis is to measure interceptive timing skill and relate it to standardized school mathematical measures. A failure to find a relationship would allow us to reject the hypothesis, whilst a more general relationship between interceptive timing skill and cognitive ability (in reading and writing) would suggest that there is no specific functional relationship between mathematics and interceptive timing ability over and above general academic achievement.

Thus, we developed an interceptive timing task with 54 moving targets to test 368 primary school children (aged 5-11 years). Three target speeds and three target widths were presented (9 trial types) with a sufficient range to challenge older children whilst allowing younger children to also succeed. The number of targets hit (IntT score) was the primary measure of interest. We also measured the manual dexterity and postural control abilities of the children to distinguish between general motor skill and interceptive timing abilities. Mathematics ability was obtained from the children's nationally standardised mathematics attainment scores (1-14 scale). These, along with reading and writing scores, were provided by the school.

7.2 Methods

7.2.1 Participants

Participants were recruited from a state primary school in Bradford, West Yorkshire, UK. From the 368 children at the school, 309 full data sets were included in the data analysis. More information on the participants can be found in chapter 5.

7.2.2 Apparatus

Interceptive timing task

Details of the interceptive timing task can be found in chapter 5. The metric used to index interceptive timing ability was the number of targets hit (IntT) out of a total of 54.

Manual dexterity

Manual dexterity was measured using the clinical kinematic assessment tool (CKAT). All details can be found in chapter 6.

Postural control task

Postural movements were measured using a custom built motion-capture rig (Flatters, Mushtaq, Hill, Rossiter, et al., 2014; Flatters, Mushtaq, Hill, Holt, et al., 2014). All details are provided in chapter 6.

Academic attainment

Nationally standardised academic attainment scores for mathematics, reading and writing were provided by the school (https://www.gov.uk/national-curriculum/ overview). Children were graded on a scale from 1 to 14 which map to UK standardised scores.

7.2.3 Data analysis

Bayesian estimation techniques were employed to make inferences from the data and conducted using PyStan 2.8.0. Bayesian estimation uses Bayes' rule to yield complete distributional information about the relative credibility of all possible parameter values in a statistical model. Formally, Bayes' rule provides the posterior distribution $P(\theta|y, X)$, where θ is a vector of model parameters, y is the data, and X is a matrix of predictors. The marginal posterior distribution can be summarized by the highest density interval (HDI). For example, for a given parameter the 95% HDI gives the upper and lower bounds of the interval which has 95% probability of containing the true parameter value.

We employed an ordered-probit regression to model the data. Ordered-probit regression is appropriate when the dependent variable is ordinal, as in the case of the academic attainment metric. The model was fit separately for each attainment outcome. The model linearly combines predictor variables (IntT, manual dexterity, posture and age) to give an expected latent academic attainment score (μ). Thus, the model assumes that academic attainment is on a continuous and linear scale. The model then maps the expected latent attainment score to the ordinal observed attainment metric via a thresholded cumulative-normal inverse link function (see figure 7.1). The model is analogous to frequentist ordered-probit regression which can be conducted with several R packages (e.g. polr; ordered).

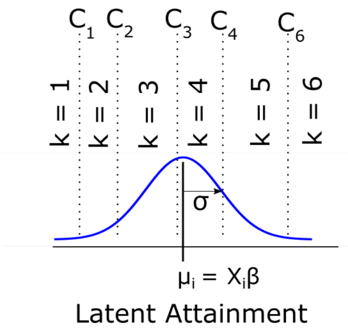

Figure 7.1: Illustration of an ordered probit model. The ith participant's predictor scores (the ith row of X) are multiplied by a vector of regression coefficients (β) to give an expected latent attainment score (μ_i). The latent attainment score is dispersed by a normal distribution centered at μ_i with standard deviation σ. Thresholds ($C_1 \ldots C_{K-1}$) slice through the distribution and the area under the curve between consecutive thresholds provides the probability of obtaining each observed academic attainment scores ($k \in 1 : K$). The distance between thresholds is not necessarily equal. Note that the number of possible attainment outcomes in this illustration is lower than in the actual model.

IntT, age, tracking, aiming, tracing and postural scores (eyes open and eyes closed) were entered as predictors. The model was based on Kruschke (2015), and its full specification with priors is given below.

$$\beta \sim \text{Normal}(0, K) \tag{7.1}$$

$$\mu = X\beta \tag{7.2}$$

$$C_1 \equiv 1.5 \tag{7.3}$$

$$C_t \sim \text{Normal}(t + 0.5, K) \quad t \in 2 : K - 1 \tag{7.4}$$

$$\sigma \sim \text{Cauchy}^{+}(0, 100) \tag{7.5}$$

$$\theta_{i,k} = \begin{cases} 1 - \phi(\frac{\mu_i - C_1}{\sigma}), & \text{if } k = 1 \\ \phi(\frac{\mu_i - C_{k-1}}{\sigma}) - \phi(\frac{\mu_i - C_k}{\sigma}), & \text{if } 1 < k < K \\ \phi(\frac{\mu_i - C_{k-1}}{\sigma}), & \text{if } k = K \end{cases} \tag{7.6}$$

$$y_i \sim \text{Categorical}(\theta_i) \quad i \in 1 : N \tag{7.7}$$

where N is the number of participants, X is an $N \times 7$ matrix of predictor variables where the first column is equal to 1. θ is an $N \times K$ matrix where the ith row specifies the probabilities of obtaining each academic score for the ith participant. ϕ is the cumulative normal function and μ represents a continuous latent attainment outcome.

The first and last threshold value C_1 and C_{K-1} were fixed in order to identify the model. All priors were chosen to be weakly informative on the scale of the data. For each model a representative sample was taken from the posterior distribution using the NUTS algorithm (Hoffman & Gelman, 2011) implemented in Stan. Four chains of 10000 samples were started at random locations of the joint posterior parameter space. Each chain first took 5000 warm up samples which were then discarded. Convergence was assessed by visually inspecting the chains and examining the Gelman-Rubin statistic ($\hat{R}$) (Gelman & Basbøll, 2014) and effective sample size of all parameters. All $\hat{R}$ values were close to 1 and the effective sample size was > 6000 for all parameters.

7.3 Results

Figure 7.2A-D shows the marginal posterior distributions over the key model parameters for the mathematics attainment model. The credible values of the IntT score slope (β_2) are displayed in figure 7.2C. The slope was non-zero, as indicated by the 95% highest density interval (HDI; horizontal black line in figure 7.2C) with a mean estimate of 0.03 (95% HDI = [0.01, 0.5]). This means that for every five extra targets hit, the model estimates an average increase of 0.15 on latent mathematics score for that individual. The age slope (β_1) was non-zero with a marginal posterior at 1.1 (95% HDI = [0.9, 1.22]) (figure 7.2B), suggesting that the latent mathematics score increased by more than one point for every year at school. For the reading and writing models the 95% HDI over the β_2 (IntT Score) parameters contained zero (see figure 7.2G and 7.2K) indicating that IntT score is unlikely to have predictive value for reading and writing scores.

The mathematics model β_2 parameter shows that mathematics attainment increases with IntT score. However, in order to estimate the probability of obtaining a given observed mathematics score ($k \in 1 : 14$) given a set of predictor values, the model also takes account of the SD (σ) and threshold (C) parameters. While the model suggests that IntT score influences mathematics

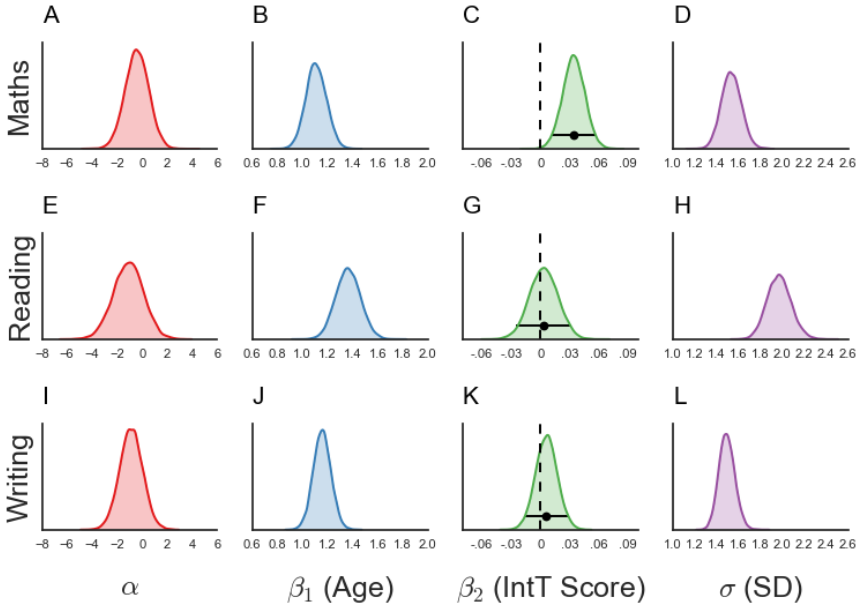

Figure 7.2: Kernel density estimates of marginal posterior distributions over key attainment model parameters: A, E, I) The regression intercepts (α) for each attainment model; B, F, J) The age slopes (β_1); C,G,K) The IntT score slopes (β_2) with an error bar showing 95% HDI and a dot indicating the mean. The vertical black dashed line indicates the zero point. HDI is clearly non-zero in the mathematics model; D, H, L) SD parameters (σ).

ability, its influence on the observed attainment measure depends on the value of all other predictors. For example, it can be seen in Figure 7.1 that if the other predictor variables result in the expected latent attainment score being high enough to place the normal distribution's mass above the last threshold, IntT score will have no discernible effect on our outcome measure (because the outcome measure is at ceiling). It is also possible that the change in mathematics ability required to move up a grade may not be equal at all levels of the attainment outcome. The model readily accounts for these possibilities, but it means that we need to hold other predictor variables constant in order to explore the model's predictions as IntT score increases.

To explore the predictions made by the model we plotted the posterior mean estimate of the probability of having an observed mathematics attainment score greater than or equal to different values, as a function of age and IntT score, where each value is shown as a separate surface (see figure 7.3A). All other predictor variables were fixed at the school median. It is clear that the probability of obtaining an observed mathematics score equal to or greater than k increased with both age and interceptive score. As expected, the effect of interceptive score depended on both age and the value of k. For example, the probability of having a mathematics score greater than or equal to 6 increased with IntT score in 9 year olds much more than in 6 year olds (for 6 year olds the probability of a score above 4 increased rapidly with IntT score).

This is further illustrated by figure 7.3B which shows the flattened surface $k \geq 6$, with the addition of the 95% HDI around the posterior mean (with all other predictors fixed at the school median). For children aged nine years, the model predicted that the probability of an observed mathematics score greater than or equal to 6 (the school median) increased by $\approx 4\%$ for every five hit targets on the interception task. For children aged six years, the IntT score had little influence on the probability of the mathematics score being ≥ 6 when the other predictors were fixed. In reality the predictors were correlated to various degrees such that good IntT scores were generally associated with good scores on the other motor tasks.

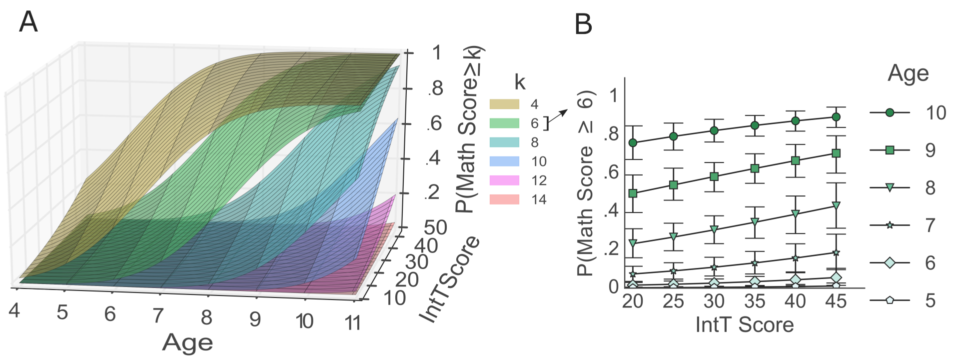

Figure 7.3: A) Mean posterior estimate of the probability of an observed mathematics attainment score equal to or greater than k as a function of age and IntT score, with all other predictors fixed at the school median. In practice these predictors covary with age and each other. B) The surface $k \geq 6$ with error bars showing the 95% HDI.

7.4 Discussion

This study demonstrates that interceptive timing ability can predict mathematical performance in primary school children. This finding is consistent with human sensorimotor systems and cognitive abilities being intrinsically linked. Correlational studies always raise questions about the direction of causality. One possibility is that learning mathematics actually improves interceptive timing ability. However, longitudinal studies have found a predictive relationship between ‘fine motor’ skills before starting school and later mathematics attainment (Grissmer et al., 2010). Thus it seems more plausible that interceptive abilities influence mathematics ability. This present finding could be extended either by replicating the effect in a longitudinal study or by training interceptive timing skills and examining the impact on mathematics (or vice versa).

Another possibility is that sensorimotor performance is a proxy measure of psychopathology, as populations with clinical motor control deficits sometimes exhibit poor mathematics ability (Tinelli et al., 2015; Van Rooijen, Verhoeven, & Steenbergen, 2011). Indeed, 'fine-motor skills' can predict measures of mathematics ability in healthy children (Carlson, Rowe, & Curby, 2013; Grissmer et al., 2010; Luo et al., 2007; Pagani, Fitzpatrick, Archambault, & Janosz, 2010; Son & Meisels, 2006). However, we controlled for general motor skills and still found IntT score was predictive of mathematics attainment but not reading or writing attainment. These controls also rule out simplistic explanations based on interceptive timing skills acting as a proxy measure for parental involvement, access to technology, or social economic status (Ritchie & Bates, 2013).

The findings are consistent with the idea that number representations are linked with concepts of time and space, possibly through a common representation of magnitude (Walsh, 2003). It is possible that children must first learn the physical rules that govern how objects move before they can form related abstract representations (Piaget, 1955). The ability to learn the physical rules is likely to vary between individuals, and our findings may reflect variance in the development of the neural structures that underpin predictive learning regarding how objects move in space and time. In this regard, our results are consistent with recent findings showing that basic spatial processing abilities in infants relate to later mathematical ability (Lauer & Lourenco, 2016).

The relationship between interceptive ability and mathematics is likely to be complex, since it is likely that not all elite sports people are excellent mathematicians, whilst people with physical disability may excel in mathematics. However, the results of the present study suggest that we should not neglect the importance of sensorimotor development in young children (given that the environment - broadly construed - is known to exert a large influence on sensorimotor ability). Indeed, the present work complements reports that physical activity can exert positive benefits on cognitive processing, even if the mechanisms remain opaque (Hill, Williams, Aucott, Thomson, & Mon-Williams, 2011). Thus, our study agrees with a growing body of evidence that suggests the quality of early sensorimotor interactions with the environment may have a direct impact on children's cognitive and education outcomes.

Chapter 8

Discussion and conclusion

8.1 Introduction

The ability to successfully intercept rapid moving objects provides a quintessential example of the capacity of humans to perform complex sensorimotor actions in dynamic and unstructured environments. Unlike interactions with static objects, interceptive actions require exquisite control over the timing of movements. That humans are able to time interceptive actions with millisecond precision, despite the numerous challenges presented by information delays, noise, uncertainty, redundancy and controlling over 600 muscles (Franklin & Wolpert, 2011) is a testament to humans' sensorimotor prowess. These skills may have played a vital role in our species' evolutionary success, allow us to engage in survival-critical behaviours and underpin many everyday tasks and sports. Interceptive timing has fascinated researchers for decades yet many questions remain regarding how interceptive actions are controlled and how these skills are learned and developed. The preceding chapters examined critical questions regarding how adults are able to achieve such high levels of performance in interceptive timing tasks, as well as exploring the ontogeny of these skills, their relationship to other motor skills, and their impact on other life outcomes (specifically academic attainment).

8.2 Summary of experimental findings

Examining interceptive timing abilities in adults and large numbers of children required a portable experimental tool which could provide detailed measures of performance. To this end a 'virtual' interceptive timing task, based on the track style experimental apparatus used in a series of studies by Tresilian et al., (e.g. Tresilian and Lonergan (2002)), was developed as documented in chapter 2. The system used a bespoke 3D printed, 1-DoF manipulandum to control an on screen bat to hit virtual moving targets. This allowed for careful control over the experimental stimuli, provided detailed and objective measures of interceptive timing performance, and yet was portable enough to be used outside of the laboratory in school environments.

The thesis then explored two interrelated themes. The first part examined the control of interceptive actions in adults, specifically how adults are able to achieve the temporal precision required to successfully strike rapid moving targets, while the second examined the ontogeny of interceptive timing abilities in children. Chapter 3 explored how adults use prior information about the motion of targets to reduce the impact of sensorimotor noise on movement timing. It has been previously found that adults make systematic timing errors in coincidence anticipations tasks, which is consistent with the brain performing Bayesian inference (Miyazaki et al., 2005). This was replicated in the interceptive timing task, with adults learning the prior distribution over target speeds within several hundred trials.

Chapter 4 then explored how people minimise errors in the execution of their movements. Online corrections are unlikely to be helpful in rapid interceptive actions, but errors can be reduced by carefully planning the pre-programmed duration of the movement (Tresilian, 2012), as faster and briefer movements are associated with better temporal precision (Newell et al., 1994). In the experiment participants struck 'pucks' at moving targets. This allowed Gaussian noise to be added to the speed of the puck. The results revealed that people compensated for the added noise by increasing the speed of their movements. However, in the absence of added noise participants reduced their movement speeds over the course of the experiment, while the probability of them hitting the targets increased. This suggests that people will choose to make slower movements when they are able to do so without compromising performance. This pattern of results is consistent with models of optimal control (Todorov, 2004), with the sensorimotor system attempting to simultaneously minimise both temporal errors and energy costs (Tresilian, 2012).

While the timing abilities of adults have been well documented (Brenner & Smeets, 2015b; Brenner et al., 2014; Tresilian et al., 2009; Tresilian, 2012) very little is know regarding how interceptive timing abilities develop over childhood. Chapter 5 examined the developmental trajectory of interceptive timing in a large cross-section of children (aged 5-11) and an adult comparison group. The study revealed that, contrary to the findings of previous studies which used coincidence timing tasks (Haywood, 1980), children were still far from adult levels of performance by age 11. The study also revealed that young children not only show poor temporal precision, but also make systematic timing errors, striking too far ahead of moving targets on average.

After establishing how interceptive timing abilities develop over early childhood, chapters 6 and 7 examined the relationship between these abilities and other motor and cognitive facets. Chapter 6 explored the relationship between interceptive timing and measures of ‘fine’ and ‘gross’ motor skills. The results suggested that interceptive timing measures may tap into a ‘motor construct’ that is somewhat distinct from those measured by motor tasks in other domains, possibly because interceptive timing tasks capture people’s ability to predict the motion of moving objects, while tasks involving static objects do not. These findings justify the inclusion of interceptive tasks in standardised measures of motor ability (i.e. the MABC-2) as they appear to capture a unique aspect of a child’s sensorimotor ‘status’.

Finally, chapter 7 examined how children’s interceptive timing abilities relate to academic attainment. A growing body of evidence suggests that sensorimotor skills, particularly in ‘fine’ motor tasks, are predictive of academic attainment (Grissmer et al., 2010; Son & Meisels, 2006). However, it is not known whether interceptive timing may add some unique predictive value when attempting to forecast children’s attainment. It was predicted that interceptive timing abilities may be particularly related to mathematics attainment, given evidence that representations of number are grounded in sensorimotor processes (M. Wilson, 2002; Crollen et al., 2013; Hubbard et al., 2005). This was confirmed with interceptive ability predicting attainment in mathematics but not reading and writing, even after controlling for both 'fine' and 'gross' motor abilities. This finding is consistent with the idea that number representations are linked with concepts of time and space (Walsh, 2003), which may play an important role in motion prediction - a fundamental component of interceptive timing performance.

8.3 Future work

The first part of this thesis explored the mechanisms by which adults minimise temporal errors in interceptive timing, while the second part documented the development of these skills in children, as well as investigating their relationship to other motor abilities and academic attainment. The results raise a number of interesting questions and suggest areas in which progress could be made. These are discussed below.

Now that the developmental trends in interceptive timing have been documented (see chapter 5), a clear line of enquiry for future research is to investigate the mechanisms driving improvements in interceptive timing with age. Chapters 3 and 4 suggest possible mechanisms by which performance may improve over childhood. Chapter 3 demonstrated that adults integrate visual information with prior knowledge in a Bayesian manner. Yet it is not known at what age this behaviour may become observable. Is Bayesian integration a fundamental function of the nervous system which is present very early in life, or does it take years for the sensorimotor system to begin operating in this fashion? Gori, Del Viva, Sandini, and Burr (2008) found that children do not appear to integrate information across multiple sensory modalities until at least age 8. However, this does not preclude the use of Bayesian integration within modalities at an earlier age. The authors suggested that a possible reason for the lack of multisensory integration may be that the developing sensorimotor systems need to constantly recalibrate to account for physical growth (e.g. lengthening of limbs) and different senses may be used to recalibrate each other. If the need to calibrate senses to one another precludes multisensory integration, it may be that within modality priors relating to physical attributes of the world may still be used by young children, as physical changes to the body would not be expected to change perception of the physical attributes of the environment (e.g. object movement speeds).

Chapter 5 demonstrated that young children show a large bias in the timing of their movements, striking too far ahead of moving targets, which may suggest that the youngest children are not performing in a Bayes optimal way. However, this was impossible to confirm in the study. A simple follow up study could use the experimental design of chapter 3 to test this hypothesis in children directly. However, one obvious challenge in doing so is that the experimental procedure requires a large number of trials, making the tasks unsuitable for testing in schools.

If children do perform Bayesian inference in a similar fashion to the adults, a number of predictions can be made. For example, we may expect children's visual estimates of the target's speed to be noisier than adults (Deutsch & Newell, 2005). If this is the case then children should take longer to learn the distribution of target speeds, as priors update more slowly when visual information is more uncertain. We may also expect children to be more reliant on their priors than adults because of the uncertainty in their online visual information. This would be observable as a larger slope and lower intercept, when regressing temporal error on target speed, than was seen in the adults in chapter 3. In addition, if adults are able to better predict how objects will move over time and space than children, we may expect to see a greater decrement in children's performance when the quality of visual input is degraded (e.g. by refractive blur). In fact, it is known that adults with normal vision show interceptive skills that are fairly robust to simulated refractive blur (Mann, Abernethy, & Farrow, 2010; Mann, Ho, De Souza, Watson, & Taylor, 2007) possibly because of their ability to predict target motion with only sparse visual information.

Further studies are also needed to determine how adults and children program their movement times in order to reduce temporal errors. Chapter 4 provided a novel method for studying this, which was examined in adults only. The results of chapter 4 demonstrate that people are most likely minimising multiple costs (i.e. errors and effort). In adults this raises the question of whether performance might be improved by simply motivating the participant to choose faster strike speeds. For example, would participants reduce their movement times if they received a reward for doing so, as would be expected from studies that show adults make decisions in movement planning which maximise the expected gain of the movement (Trommershäuser, Maloney, & Landy, 2009). Extending this research to children would also be of particular value. Childhood is characterised by a rapid increase in physical strength, thus it may be that the costs associated with certain movements may change with age (i.e. faster movements may become less effortful with age). Thus the cost functions implicit in the decisions involved in planning movement times may fluctuate over the developmental trajectory.

As well as furthering our understanding of the mechanisms driving improvements in interceptive timing with age, chapters 6 and 7 suggest that research into the complex interplay between intercept timing, general sensorimotor abilities and child development is also of critical importance. It is clear that early sensorimotor skills are related to life outcomes, including academic attainment (chapter 7) and even higher level cognitive abilities (Gottwald et al., 2016). While existing studies have examined the relationship between motor skills and academic attainment, these have generally involved using a broad number of complex motor tasks with scores averaged across them. Together, chapters 6 and 7 suggest that different tasks not only tap different motor 'constructs', but that they also contribute uniquely to predictive models of academic attainment. In other words, there is not a 'general' motor skill which predicts attainment, but rather different motor tasks (including interceptive timing tasks) appear to measure different constructs which relate to academic attainment.

In understanding the link between interceptive timing and academic attainment in mathematics there are a number of challenges which need to be overcome. A key limitation of chapter 7 was the observational nature of the study, making it difficult to establish a causal relationship. In addition, it is not possible to completely rule out the prospect of a third variable accounting for the relationship (e.g. access to technology, parental involvement), although this seemed unlikely given the pattern of results. Yet the explanation that the sensorimotor processes which underpin interceptive timing tasks may provide a foundation upon which mathematical abilities arise is compelling (Hubbard et al., 2005; Walsh, 2003). One possible method of establishing a causal relationship would be to train interceptive timing skills and examine whether any improvements result in performance on approximate or symbolic number tasks. While this may initially seem implausible, it is known that even brief training on approximate number tasks improve symbolic mathematical ability (Hyde, Khanum, & Spelke, 2014), and neural representations of number, which are presumably recruited in approximate number tasks, may be related to representations of space and time (Burr et al., 2011).

Establishing whether shared neural substrates exist between the processes involved in mathematics and interceptive timing skills is a clear challenge for future research. Evidence suggests that the posterior parietal cortex may be a possible substrate for processes critical in mathematics and interceptive timing abilities. Walsh (2003) proposed that the representation of number, time and space may be linked through a shared representation of magnitude. In support of this finding, Burr et al. (2011) found that saccades compress perceptions of spatial, temporal and numerical magnitudes, an effect which has been linked to predictive remapping of retinotopic neurons during saccades (Duhamel, Colby, & Goldberg, 1992), particularly in the lateral intraparietal sulcus (LIP). In addition neurons in LIP and the ventral intraparietal sulcus (VIP) have been found to be sensitive to both numerosity (Roitman, Brannon, & Platt, 2007; Nieder & Miller, 2004) and time durations (Leon, Leon, Shadlen, & Shadlen, 2003; Hayashi et al., 2013), while damage to the posterior parietal cortex is often associated with deficits in numerical and temporal judgements (Bueti & Walsh, 2009). Given that estimation of time durations is critical in interceptive timing tasks it seems plausible that the same neuronal processes may be recruited when intercepting moving objects. In fact, neurons in LIP appear to play a role in visual motion extrapolation (Bosco et al., 2015), an ability which is critical in interceptive tasks and requires excellent estimation of time.

8.4 Concluding remarks

The results of this thesis provided insights into the mechanisms by which adults can achieve amazing levels of temporal precision, how these abilities develop in childhood, and how they relate to other sensorimotor and cognitive domains. It is clear that much more work needs to be done to understand how these skills develop and influence cognitive facets and academic attainment. Using sophisticated portable research tools (such as that described in chapter 2), which move beyond the crude and noisy measures found in standardised motor tests, should provide a useful first step to achieving this.

References

Ament, K., Mejia, A., Buhlman, R., Erklin, S., Caffo, B., Mostofsky, S., & Wodka, E. (2014). Evidence for specificity of motor impairments in catching and balance in children with autism. *Journal of Autism and Developmental Disorders*, *45*(3), 742–751.

Bachman, J. C. (1961). Specificity vs. generality in learning and performing two large muscle motor tasks. *Research Quarterly. American Association for Health, Physical Education and Recreation*, *32*(1), 3–11.

Balasubramaniam, R., Riley, M. A., & Turvey, M. T. (2000). Specificity of postural sway to the demands of a precision task. *Gait and Posture*, *11*(1), 12–24.

Barnes, G. R., & Asselman, P. T. (1991). The mechanism of prediction in human smooth pursuit eye movements. *The Journal of Physiology*, *439*, 439–461.

Barnett, A., Henderson, S. E., & Sugden, D. A. (2007). *Movement assessment battery for children - second edition (movement abc-2)*. Pearson.

Benguigui, N., Broderick, M. P., Baurès, R., & Amorim, M.-A. (2008). Motion prediction and the velocity effect in children. *British Journal of Developmental Psychology*, *26*(3), 389–407.

Betancourt, M. (2017). A conceptual introduction to hamiltonian monte carlo.

Bootsma, R. J., Ledouit, S., Casanova, R., & Zaal, F. T. J. M. (2015). Fractional-order information in the visual control of lateral locomotor interception. *Journal of Experimental Psychology: Human Perception and Performance*, *42*(4), 517–529.

Bosco, G., Monache, S. D., Gravano, S., Indovina, I., La Scaleia, B., Maffei, V., … Lacquaniti, F. (2015). Filling gaps in visual motion for target capture. *Frontiers in Integrative Neuroscience*, *9*(February), 1–17.

Brace, D. K. (1927). *Measuring motor ability: a sale of motor ability tests*. New York: A. S. Barnes.

Brenner, E., Cañal-Bruland, R., & van Beers, R. J. (2013). How the required precision influences the way we intercept a moving object. *Experimental Brain Research*, *230*(2), 207–18.

Brenner, E., Driesen, B., & Smeets, J. B. J. (2014). Precise timing when hitting falling balls. *Frontiers in Human Neuroscience*, *8*, 342.

Brenner, E., & Smeets, J. B. J. (2010). How well can people judge when something happened? *Vision Research*, *50*(12), 1101–8.

Brenner, E., & Smeets, J. B. J. (2015a). How moving backgrounds influence interception. *PloS One*, *10*(3), e0119903.

Brenner, E., & Smeets, J. B. J. (2015b). How people achieve their amazing temporal precision in interception. *Journal of Vision*, *15*(3), 1–21.

Brenner, E., van Dam, M., Berkhout, S., & Smeets, J. B. (2012). Timing the moment of impact in fast human movements. *Acta Psychologica*, *141*(1), 104–111.

Brooks, S. P. (2003). Bayesian computation: a statistical revolution. *Philosophical Transactions. Series A, Mathematical, Physical, and Engineering Sciences*, *361*(1813), 2681–97.

Brouwer, A.-M., Brenner, E., & Smeets, J. B. J. (2002). Hitting moving objects: is target speed used in guiding the hand? *Experimental Brain Research*, *143*(2), 198–211.

Brouwer, A.-M., Brenner, E., & Smeets, J. B. J. (2003). When is behavioral data evidence for a control theory? tau-coupling revisited. *Motor control*, *7*(2), 103–110.

Bueti, D., & Walsh, V. (2009). The parietal cortex and the representation of time, space, number and other magnitudes. *Philosophical Transactions of the Royal Society of London. Series B, Biological sciences*, *364*(1525), 1831–1840.

Burr, D. C., Ross, J., Binda, P., & Morrone, M. (2011). Saccades compress space, time, and number. In *Space, time and number in the brain* (Vol. 14, pp. 175–186). Elsevier.

Caçola, P., Ibana, M., Ricard, M., & Gabbard, C. (2016). Children with developmental coordination disorder demonstrate a spatial mismatch when estimating coincident-timing ability with tools. *Research in Developmental Disabilities*, *48*, 124–131.

Caljouw, S. R., van der Kamp, J., & Savelsbergh, G. J. P. (2004). Catching optical information for the regulation of timing. *Experimental Brain Research*, *155*(4), 427–38.

Carlson, A. G., Rowe, E., & Curby, T. W. (2013). Disentangling fine motor skills' relations to academic achievement: the relative contributions of visual-spatial integration and visual-motor coordination. *The Journal of Genetic Psychology*, *174*(5), 514–533.

Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., … Riddell, A. (2016). Stan: a probabilistic programming language. *Journal of Statistical Software*, *VV*(Ii).

Chang, a. Y.-C., Tzeng, O. J. L., Hung, D. L., & Wu, D. H. (2011). Big time is not always long: numerical magnitude automatically affects time reproduction. *Psychological Science*, *22*(12), 1567–1573.

Chow, S. M. K. (2003). Interrater and test-retest reliability of the movement assessment battery for chinese preschool children. *Physical Therapy*, *57*(5).

Cohen, J., & Hubbard, R. (1995). The earth is round. *American Psychologist*, *49*(12), 997–1003.

Croce, R. V., Horvat, M., & McCarthy, E. (2001). Reliability and concurrent validity of the movement assessment battery for children. *Perceptual and Motor Skills*, *93*(1), 275–280.

Crollen, V., Dormal, G., Seron, X., Lepore, F., & Collignon, O. (2013). Embodied numbers: the role of vision in the development of number-space interactions. *Cortex*, *49*(1), 276–283.

Culmer, P. R., Levesley, M. C., Mon-Williams, M., & Williams, J. H. G. (2009). A new tool for assessing human movement: the kinematic assessment tool. *Journal of Neuroscience Methods*, *184*(1), 184–192.

Cumming, G. (2014). The new statistics. *Psychological Science*, *25*(1), 7–29.

Dallos, P. J., & Jones, R. W. (1963). Learning behavior of the eye fixation control system. *IEEE Transactions on Automatic Control*, *8*(3), 218–227.

de Hevia, M. D., & Spelke, E. S. (2010). Number-space mapping in human infants. *Psychological Science*, *21*(5), 653–660.

de la Malla, C., Lopez-Moliner, J., & Brenner, E. (2012). Seeing the last part of a hitting movement is enough to adapt to a temporal delay. *Journal of Vision*, *12*(10), 4–4.

Desmurget, M., & Grafton, S. (2000). Forward modeling allows feedback control for fast reaching movements. *Trends in Cognitive Sciences*, *4*(11), 423–431.

Deutsch, K. M., & Newell, K. M. (2005). Noise, variability, and the development of children's perceptual-motor skills. *Developmental Review*, *25*(2), 155–180.

Dienes, Z. (2011). Bayesian versus orthodox statistics: which side are you on? *Perspectives on Psychological Science*, *6*(3), 274–290.

Dienes, Z. (2014). Using bayes to get the most out of non-significant results. *Frontiers in Psychology*, *5*(July), 1–17.

Dirksen, T., De Lussanet, M. H. E., Zentgraf, K., Slupinski, L., & Wagner, H. (2016). Increased throwing accuracy improves children's catching performance in a ball-catching task from the movement assessment battery (mabc-2). *Frontiers in Psychology*, *7*, 1–8.

Duhamel, J.-r., Colby, C. L., & Goldberg, M. E. (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science*, *255*(5040), 90–92. Retrieved from `http://www.jstor.org/stable/2876068`

Faisal, a. A., Selen, L. P. J., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature reviews. Neuroscience*, *9*(4), 292–303.

Fias, W., van Dijck, J. P., & Gevers, W. (2011). How is number associated with space? the role of working memory. In *Space, time and number in the brain* (pp. 133–148).

Flanagan, J. R., & Wing, A. M. (1997). The role of internal models in motion planning and control: evidence from grip force adjustments during movements of hand-held loads. *The Journal of Neuroscience*, *17*(4), 1519–1528.

Flatters, I., Culmer, P., Holt, R. J., Wilkie, R. M., & Mon-Williams, M. (2014). A new tool for assessing head movements and postural sway in children. *Behavior Research Methods*, *46*(4), 950–959.

Flatters, I., Hill, L. J. B., Williams, J. H. G., Barber, S. E., & Mon-Williams, M. (2014). Manual control age and sex differences in 4 to 11 year old children. *PloS One*, *9*(2), e88692.

Flatters, I., Mushtaq, F., Hill, L. J. B., Holt, R. J., Wilkie, R. M., & Mon-Williams, M. (2014). The relationship between a child's postural stability and manual dexterity. *Experimental Brain Research*, *232*(9), 2907–17.

Flatters, I., Mushtaq, F., Hill, L. J. B., Rossiter, A., Jarrett-Peet, K., Culmer, P., . . . Mon-Williams, M. (2014). Children's head movements and postural stability as a function of task. *Experimental Brain Research*, *232*(6), 1953–70.

Foulkes, J. D., Knowles, Z., Fairclough, S. J., Stratton, G., O'Dwyer, M., Ridgers, N. D., & Foweather, L. (2015). Fundamental movement skills of preschool children in northwest england. *Perceptual and Motor Skills*, *121*(1), 260–283.

Franklin, D. W., & Wolpert, D. M. (2011). Computational mechanisms of sensorimotor control. *Neuron*, *72*(3), 425–42.

Freeman, T. C., Champion, R. a., & Warren, P. a. (2010). A bayesian model of perceived head-centered velocity during smooth pursuit eye movement. *Current Biology*, *20*(8), 757–762.

Gelman, A. (2014). *Bayesian data analysis* (3rd ed.). Boca Raton: Taylor & Francis Group.

Gelman, A., & Basbøll, T. (2014). When do stories work? evidence and illustration in the social sciences. *Sociological Methods & Research*, 1–24.

Gibson, J. J. (1961). Ecological optics. *Vision Research*, *1*(3-4), 253–262.

Gomez, A., Piazza, M., Jobert, A., Dehaene-Lambertz, G., Dehaene, S., & Huron, C. (2015). Mathematical difficulties in developmental coordination disorder: symbolic and nonsymbolic number processing. *Research in Developmental Disabilities*, *43-44*, 167–178.

Goodman, S. (2008). A dirty dozen: twelve p-value misconceptions. *Seminars in hematology*, *45*(3), 135–40.

Gori, M., Del Viva, M., Sandini, G., & Burr, D. C. (2008). Young children do not integrate visual and haptic form information. *Current Biology*, *18*(9), 694–698.

Gottwald, J. M., Achermann, S., Marciszko, C., Lindskog, M., & Gredebäck, G. (2016). An embodied account of early executive-function development. *Psychological Science*, *27*(12), 1600–1610.

Gray, R. (2009). Intercepting moving objects: fundamental principles learned from baseball. *Reviews of Human Factors and Ergonomics*, *5*(1), 114–139.

Grissmer, D., Grimm, K. J., Aiyer, S. M., Murrah, W. M., & Steele, J. S. (2010). Fine motor skills and early comprehension of the world: two new school readiness indicators. *Developmental Psychology*, *46*(5), 1008–1017.

Haddad, J. M., Ryu, J. H., Seaman, J. M., & Ponto, K. C. (2010). Time-to-contact measures capture modulations in posture based on the precision demands of a manual task. *Gait and Posture*, *32*(4), 592–596.

Hayashi, M. J., Kanai, R., Tanabe, H. C., Yoshida, Y., Carlson, S., Walsh, V., & Sadato, N. (2013). Interaction of numerosity and time in prefrontal and parietal cortex. *The Journal of Neuroscience*, *33*(3), 883–93.

Haywood, K. M. (1980). Coincidence-anticipation accuracy across the life span. *Experimental Aging Research*, *6*(5), 451–462.

Henderson, S. E., & Barnett, A. L. (1998). The classification of specific motor coordination disorders in children: some problems to be solved. *Human Movement Science*, *17*(4-5), 449–469.

Henderson, S. E., & Sugden, D. A. (1992). *The movement assessment battery for children.* London: The Psychological Corporation.

Hill, L. J. B., Williams, J. H. G., Aucott, L., Thomson, J., & Mon-Williams, M. (2011). How does exercise benefit performance on cognitive tests in primary-school pupils? *Developmental Medicine and Child Neurology*, *53*(7), 630–635.

Hoffman, M., & Gelman, A. (2011). The no-u-turn sampler: adaptively setting path lengths in hamiltonian monte carlo. *Journal of Machine Learning Research*, *15*, 1351–1381.

Huang, M. H., & Brown, S. H. (2013). Age differences in the control of postural stability during reaching tasks. *Gait and Posture*, *38*(4), 837–842.

Hubbard, E. M., Piazza, M., Pinel, P., & Dehaene, S. (2005). Interactions between number and space in parietal cortex. *Nature reviews. Neuroscience*, *6*(6), 435–448.

Hurks, P. P. M., & Loosbroek, E. V. (2012). Time estimation deficits in childhood mathematics difficulties. *Journal of Learning Disabilities*, *47*, 450–461.

Hyde, D. C., Khanum, S., & Spelke, E. S. (2014). Brief non-symbolic, approximate number practice enhances subsequent exact symbolic arithmetic in children. *Cognition*, *131*(1), 92–107.

Jazayeri, M., & Shadlen, M. N. (2010). Temporal context calibrates interval timing. *Nature Neuroscience*, *13*(8), 1020–6.

Kim, R., Nauhaus, G., Glazek, K., Young, D., & Lin, S. (2013). Development of coincidence-anticipation timing in a catching task. *Perceptual and Motor Skills*, *117*(1), 319–338.

Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*(6971), 244–7.

Kruschke, J. K. (2013). Bayesian estimation supersedes the t test. *Journal of Experimental Psychology: General*, *142*(2), 573–603.

Kruschke, J. K. (2015). *Doing bayesian data analysis* (2nd ed.). London: Elsevier.

Kruschke, J. K., & Liddell, T. M. (2017a). The bayesian new statistics:

hypothesis testing, estimation, meta-analysis, and power analysis from a bayesian perspective. *Psychonomic Bulletin & Review*, 1–21.

Kruschke, J. K., & Liddell, T. M. (2017b). The bayesian new statistics: hypothesis testing, estimation, meta-analysis, and power analysis from a bayesian perspective. *Psychonomic Bulletin & Review*, 1–28.

Lambdin, C. (2012). Significance tests as sorcery: science is empirical–significance tests are not. *Theory & Psychology*, *22*(1), 67–90.

Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, *23*(11), 571–579.

Langaas, T., Mon-Williams, M., Wann, J., Pascal, E., & Thompson, C. (1998). Eye movements, prematurity and developmental co-ordination disorder. *Vision Research*, *38*(12), 1817–1826.

Lauer, J. E., & Lourenco, S. F. (2016). Spatial processing in infancy predicts both spatial and mathematical aptitude in childhood. *Psychological Science*, *27*(10), 1291–1298.

Lee, D. N., Young, D. S., Reddish, P. E., Lough, S., & Clayton, T. M. H. (1983). Visual timing in hitting an accelerating ball. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, *35A*(2), 333–346.

Lefebvre, C., & Reid, G. (1998). Prediction in ball catching by children with and without a developmental coordination disorder. *Adapted Physical Activity Quarterly*, *15*(4), 299–315.

Leon, M. I., Leon, M. I., Shadlen, M. N., & Shadlen, M. N. (2003). Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron*, *38*(2), 317–27. Retrieved from `http://www.ncbi.nlm.nih.gov/pubmed/12718864`

López-Moliner, J., Brenner, E., Louw, S., & Smeets, J. B. J. (2010). Catching a gently thrown ball. *Experimental Brain Research*, *206*(4), 409–17.

Lourenco, S. F., & Longo, M. R. (2010). General magnitude representation in human infants. *Psychological Science*, *21*(6), 873–81.

Lumelsky. (2006). *Sensing, intelligence, motion: how robots and humand move in an unstructured world.* Hoboken, New Jersey: John Wiley & Sons, Inc.

Luo, Z., Jose, P. E., Huntsinger, C. S., & Pigott, T. D. (2007). Fine motor skills and mathematics achievement in east asian american and european american kindergartners and first graders. *British Journal of Developmental Psychology*, *25*(4), 595–614.

Mann, D. L., Abernethy, B., & Farrow, D. (2010). The resilience of natural interceptive actions to refractive blur. *Human Movement Science*, *29*(3), 386–400.

Mann, D. L., Ho, N. Y., De Souza, N. J., Watson, D. R., & Taylor, S. J. (2007). Is optimal vision required for the successful execution of an interceptive task? *Human Movement Science*, *26*(3), 343–356.

Manual, R. (2013). Stan modeling language user ’ s guide and reference manual.

Marsman, M., Schönbrodt, F. D., Morey, R. D., Yao, Y., Gelman, A., & Wagenmakers, E.-J. (2017). A bayesian bird’s eye view of âĂŸreplications of important results in social psychology’. *Royal Society Open Science*, *4*(1), 160426.

Massion, J. (1992). Movement, posture and equilibrium: interaction and coordination. *Progress in Neurobiology*, *38*(1), 35–56.

McCloy, C. H. (1934). The measurement of general motor capacity and general motor ability. *Research Quarterly. American Physical Education Association*, *5*, 46–61.

McLeod, P., & Dienes, Z. (1996). Do fielders know where to go to catch the ball or only how to get there? *Journal of Experimental Psychology: Human Perception and Performance*, *22*(3), 531–543.

Merfeld, D. M., Zupan, L., & Peterka, R. J. (1999). Humans use internal models to estimate gravity and linear acceleration. *Nature*, *398*(6728), 615–618.

Miyazaki, M., Nozaki, D., & Nakajima, Y. (2005). Testing bayesian models of human coincidence timing. *Journal of Neurophysiology*, *94*(1), 395–9.

Montagne, G., Laurent, M., Durey, a., & Bootsma, R. (1999). Movement reversals in ball catching. *Experimental Brain Research*, *129*(1), 87–92.

Morey, R. D., Hoekstra, R., Rouder, J. N., Lee, M. D., & Wagenmakers, E.-J. (2015). The fallacy of placing confidence in confidence intervals. *Psychonomic Bulletin & Review*.

Newell, K., Carlton, L. G., & Kim, S. (1994). Time and space-time movement accuracy. *Human Performance*, *7*(1), 1–21.

Nieder, A., & Dehaene, S. (2009). Representation of number in the brain. *Annual review of Neuroscience*, *32*, 185–208.

Nieder, A., & Miller, E. K. (2004). A parieto-frontal network for visual numerical information in the monkey. *Proceedings of the National Academy of Sciences of the United States of America*, *101*(19), 7457–

7462.

Nijhawan, R. (2008). Visual prediction: psychophysics and neurophysiology of compensation for time delays. *Behavioral and Brain Sciences*, *31*(02), 179–198.

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, *349*(6251), aac4716–aac4716.

Pagani, L. S., Fitzpatrick, C., Archambault, I., & Janosz, M. (2010). School readiness and later achievement: a french canadian replication and extension. *Developmental Psychology*, *46*(5), 984–994.

Peper, L., Bootsma, R. J., Mestre, D. R., & Bakker, F. C. (1994). Catching balls: how to get the hand to the right place at the right time. *Journal of Experimental Psychology: Human Perception and Performance*, *20*(3), 591–612.

Piaget, J. (1955). *The construction of reality in the child.* Routledge & K. Paul.

Pieters, S., Desoete, A., Van Waelvelde, H., Vanderswalmen, R., & Roeyers, H. (2012). Mathematical problems in children with developmental coordination disorder. *Research in Developmental Disabilities*, *33*(4), 1128–1135.

Purcell, C., Wann, J. P., Wilmut, K., & Poulter, D. (2012). Reduced looming sensitivity in primary school children with developmental co-ordination disorder. *Developmental Science*, *15*(3), 299–306.

Rigoli, D., Piek, J. P., Kane, R., & Oosterlaan, J. (2012). Motor coordination, working memory, and academic achievement in a normative adolescent sample: testing a mediation model. *Archives of Clinical Neuropsychology*, *27*(7), 766–780.

Ritchie, S. J., & Bates, T. C. (2013). Enduring links from childhood mathematics and reading achievement to adult socioeconomic status. *Psychological Science*, *24*(7), 1301–8.

Robin, D. J., Berthier, N. E., & Clifton, R. K. (1996). Infants' predictive reaching for moving objects in the dark. *Developmental Psychology*, *32*(5), 824–835.

Robinson, L. E., Stodden, D. F., Barnett, L. M., Lopes, V. P., Logan, S. W., Rodrigues, L. P., & D'Hondt, E. (2015). Motor competence and its effect on positive developmental trajectories of health. *Sports Medicine*, *45*(9), 1273–1284.

Roitman, J. D., Brannon, E. M., & Platt, M. L. (2007). Monotonic coding of numerosity in macaque lateral intraparietal area. *PLoS Biology*, *5*(8), 1672–1682.

Rushton, S. K., & Wann, J. P. (1999). Weighted combination of size and disparity: a computational model for timing a ball catch. *Nature Neuroscience*, *2*(2), 186–190.

Sato, Y., & Kording, K. P. (2014). How much to trust the senses: likelihood learning. *Journal of Vision*, *14*(13), 13–13.

Savelsbergh, G. J. P., Bennett, S. J., Angelakopoulos, G. T., & Davids, K. (2005). Perceptual-motor organization of children's catching behaviour under different postural constraints. *Neuroscience letters*, *373*(2), 153–8.

Schulz, J., Henderson, S. E., Sugden, D. a., & Barnett, A. L. (2011). Structural validity of the movement abc-2 test: factor structure comparisons across three age groups. *Research in Developmental Disabilities*, *32*(4), 1361–1369.

Senot, P., Prévost, P., & McIntyre, J. (2003). Estimating time to contact and impact velocity when catching an accelerating object with the hand. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(1), 219–237.

Shea, S. L., & Aslin, R. N. (1990). Oculomotor responses to step-ramp targets by young human infants. *Vision Research*, *30*(7), 1077–1092.

Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, *22*(11), 1359–66.

Simms, V., Gilmore, C., Cragg, L., Clayton, S., Marlow, N., & Johnson, S. (2014). Nature and origins of mathematics difficulties in very preterm children: a different etiology than developmental dyscalculia. *Pediatric Research*, *77*(2), 389–395.

Simms, V., Gilmore, C., Cragg, L., Marlow, N., Wolke, D., & Johnson, S. (2013). Mathematics difficulties in extremely preterm children: evidence of a specific deficit in basic mathematics processing. *Pediatric Research*, *73*(2), 236–44.

Son, S.-H., & Meisels, S. J. (2006). The relationship of young children's motor skills to later school achievement. *Merrill-Palmer Quarterly*, *52*(4), 755–778.

Spering, M., Schütz, A. C., Braun, D. I., & Gegenfurtner, K. R. (2011). Keep your eyes on the ball: smooth pursuit eye movements enhance prediction of visual motion. *Journal of Neurophysiology*, *105*(4), 1756–1767.

Srinivasan, M., & Carey, S. (2010). The long and the short of it: on the nature and origin of functional overlap between representations of space and time. *Cognition*, *116*(2), 217–241.

Stan Development Team. (2016). Pystan: the python interface to stan, version 2.9.0.

Stocker, A. a., & Simoncelli, E. P. (2006). Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience*, *9*(4), 578–585.

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: statistics, structure, and abstraction. *Science*, *331*(6022), 1279–85.

Thelen, E., & Spencer, J. P. (1998). Postural control during reaching in young infants: a dynamic systems approach. *Neuroscience and Biobehavioral Reviews*, *22*(4), 507–514.

Thompson, P. (1982). Perceived rate of movement depends on contrast. *Vision Research*, *22*(3), 377–380.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2009). Nih public access. , *12*(8), 291–297.

Tinelli, F., Anobile, G., Gori, M., Aagten-Murphy, D., Bartoli, M., Burr, D. C., … Concetta Morrone, M. (2015). Time, number and attention in very low birth weight children. *Neuropsychologia*, *73*, 60–69.

Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience*, *7*(9), 907–15.

Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, *5*(11), 1226–35.

Tresilian, J. R. (1994). Perceptual and motor processes in interceptive timing. *Human Movement Science*, *13*(3-4), 335–373.

Tresilian, J. R. (1999). Visually timed action: time-out for 'tau'? *Trends in Cognitive Sciences*, *3*(8), 301–310.

Tresilian, J. R. (2005). Hitting a moving target: perception and action in the timing of rapid interceptions. *Perception & Psychophysics*, *67*(1), 129–49.

Tresilian, J. R. (2012). Sensorimotor control and learning. In *Sensorimotor control and learning.* London: Palgrave Macmillan.

Tresilian, J. R., & Houseman, J. H. (2005). Systematic variation in performance of an interceptive action with changes in the temporal constraints. *The Quarterly Journal of Experimental Psychology. A, Human Experimental Psychology*, *58*(3), 447–66.

Tresilian, J. R., & Lonergan, A. (2002). Intercepting a moving target: effects of temporal precision constraints and movement amplitude. *Experimental Brain Research*, *142*(2), 193–207.

Tresilian, J. R., & Plooy, A. (2006). Systematic changes in the duration and precision of interception in response to variation of amplitude and effector size. *Experimental Brain Research*, *171*(4), 421–35.

Tresilian, J. R., Plooy, a., & Carroll, T. J. (2004). Constraints on the spatiotemporal accuracy of interceptive action: effects of target size on hitting a moving target. *Experimental Brain Research*, *155*(4), 509–26.

Tresilian, J. R., Plooy, a. M., & Marinovic, W. (2009). Manual interception of moving targets in two dimensions: performance and space-time accuracy. *Brain Research*, *1250*, 202–17.

Van Der Fits, I. B. M., Klip, A. W. J., Van Eykern, L. A., & Hadders-Algra, M. (1998). Postural adjustments accompanying fast pointing movements in standing, sitting and lying adults. *Experimental Brain Research*, *120*(2), 202–216.

Van Rooijen, M., Verhoeven, L., & Steenbergen, B. (2011). Early numeracy in cerebral palsy: review and future research. *Developmental Medicine and Child Neurology*, *53*(3), 202–209.

Van Waelvelde, H., De Weerdt, W., De Cock, P., & Engelsman, B. C. M. S. (2003). Ball catching. can it be measured? *Physiotherapy Theory and Practice*, *19*(4), 259–267.

Van Waelvelde, H., De Weerdt, W., De Cock, P., & Smits-Engelsman, B. C. M. (2004a). Aspects of the validity of the movement assessment battery for children. *Human Movement Science*, *23*(1), 49–60.

Van Waelvelde, H., De Weerdt, W., De Cock, P., & Smits-Engelsman, B. C. M. (2004b). Association between visual perceptual deficits and motor deficits in children with developmental coordination disorder. *Developmental Medicine and Child Neurology*, *46*(10), 661–6.

van Beers, R. J. (2012). How does our motor system determine its learning rate? *PloS One*, *7*(11), e49373.

van Beers, R. J., Baraduc, P., & Wolpert, D. M. (2002). Role of uncertainty in sensorimotor control. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, *357*(1424), 1137–45.

van Hof, P., van der Kamp, J., & Savelsbergh, G. J. P. (2008). The relation between infants' perception of catchableness and the control of catching. *Developmental Psychology*, *44*(1), 182–194.

Vehtari, A., Gelman, A., & Gabry, J. (2015). Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Statistics and Computing*(July).

Verstynen, T., & Sabes, P. N. (2011). How each movement changes the next: an experimental and theoretical study of fast adaptive priors in reaching. *The Journal of Neuroscience*, *31*(27), 10050–9.

Vilares, I., Howard, J. D., Fernandes, H. L., Gottfried, J. A., & Kording, K. P. (2012). Differential representations of prior and likelihood uncertainty in the human brain. *Current Biology*, *22*(18), 1641–1648.

Vilares, I., & Kording, K. (2011). Bayesian models: the structure of the world, uncertainty, behavior, and the brain. *Annals of the New York Academy of Sciences*, *1224*, 22–39.

von Hofsten, C. (1980). Predictive reaching for moving objects by human infants. *Journal of Experimental Child Psychology*, *30*(3), 369–382.

von Hofsten, C., Vishton, P., Spelke, E. S., Feng, Q., & Rosander, K. (1998). Predictive action in infancy: tracking and reaching for moving objects. *Cognition*, *67*(3), 255–285.

Walsh, V. (2003). A theory of magnitude: common cortical metrics of time, space and quantity. *Trends in Cognitive Sciences*, *7*(11), 483–488.

Whyatt, C., & Craig, C. M. (2012). Motor skills in children aged 7-10 years, diagnosed with autism spectrum disorder. *Journal of autism and developmental disorders*, *42*(9), 1799–809.

Whyatt, C., & Craig, C. M. (2013). Interceptive skills in children aged 9âĂŞ11 years, diagnosed with autism spectrum disorder. *Research in Autism Spectrum Disorders*, *7*(5), 613–623.

Wilson, M. (2002). Six views of embodied cognition. *Psychonomic bulletin & review*, *9*(4), 625–636.

Wilson, P. H., Ruddock, S., Smits-Engelsman, B., Polatajko, H., & Blank, R. (2013). Understanding performance deficits in developmental coordination disorder: a meta-analysis of recent research. *Developmental Medicine and Child Neurology*, *55*(3), 217–28.

Wolpert, D. M., Miall, R. C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences*, *2*(9), 338–347.

Wuang, Y.-P., Su, J.-H., & Su, C.-Y. (2012). Reliability and responsiveness of the movement assessment battery for children-second edition test in children with developmental coordination disorder. *Developmental Medicine & Child Neurology*, *54*(2), 160–165.

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

Zago, M., McIntyre, J., Senot, P., & Lacquaniti, F. (2009). Visuo-motor coordination and internal models for object interception. *Experimental Brain Research*, *192*(4), 571–604.

Zwicker, J. G., Missiuna, C., Harris, S. R., & Boyd, L. a. (2012). Developmental coordination disorder: a review and update. *European Journal of Paediatric Neurology*, *16*(6), 573–581.