

Investigating Interlimb Generalisation of Bayesian Sensorimotor Learning

Chris Hewitson

Principal Supervisor: Dr. David Kaplan

Associate Supervisor: Associate Professor Paul Sowman

A thesis submitted in partial fulfilment of the requirements for the degree of
Master of Research

Department of Cognitive Science
Faculty of Human Sciences
Macquarie University

Submitted on
October 10th 2016

Contents	
Abstract	4
Declaration	5
Acknowledgements	Error! Bookmark not defined.
1. General Introduction	7
1.1 Types of sensorimotor learning	9
1.1.1 Error-based versus reinforcement-based sensorimotor learning	10
1.1.2 Kinematic versus dynamic sensorimotor learning	10
1.1.3 Intrinsic versus extrinsic reference frames for sensorimotor learning ...	13
1.1.4 Intralimb versus interlimb generalisation of sensorimotor learning	14
1.2 Key experimental results	15
1.2.1 Intralimb generalisation of sensorimotor learning	15
1.2.2 Interlimb generalisation of sensorimotor learning	17
1.3 Sensorimotor learning under uncertainty: Does the brain perform Bayesian integration?	20
1.3.1 Bayesian sensorimotor learning	21
1.3.2 Recent experimental work and next steps	22
2. Investigating Interlimb Generalisation of Bayesian Sensorimotor Learning	24
2.1 Introduction	24
2.2 General Material and Methods	27
2.2.1 Ethics	27
2.2.2 Participants	27
2.2.4 General Experimental procedure	28
2.2.5 Analysis	31
2.3 Experiment 1: Bayesian visuomotor adaptation (VA)	34
2.3.1 Models of visuomotor adaptation	34
2.3.1.1 Full compensation model	35
2.3.1.2 Minimal mapping model	35
2.3.1.3 Bayesian estimation model	36
2.3.2 Predictions	38
2.3.3 Results	39
2.3.4 Discussion	40
2.4 Experiment 2: Extrinsic vs. intrinsic transfer	43
2.4.1 Models and predictions of generalisation	45
2.4.1.1 Full-Bayesian model	45

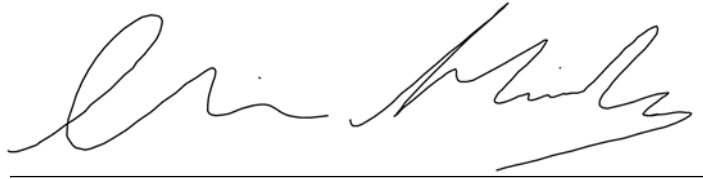
2.4.1.2 <i>Partial-Bayesian model</i>	46
2.4.1.3 <i>Quarantine model</i>	46
2.4.2 <i>Results</i>	48
2.4.3 <i>Discussion</i>	55
2.4.3.1 <i>Generalisation of the learned prior</i>	56
2.4.3.2 <i>Integrating visual uncertainty</i>	64
2.5 <i>Experiment 3: Temporal constraints on Bayesian sensorimotor learning</i>	66
2.5.1 <i>Predictions</i>	67
2.5.2 <i>Results</i>	67
2.5.3 <i>Discussion</i>	68
3. General Discussion	70
3.1 <i>The Probabilistic Population Code (PPC) model</i>	71
3.2 <i>The Neural Weighting (NW) model</i>	73
3.3 <i>A Fundamental Tension Between the Two Models</i>	74
3.4 <i>What kind of constraints can behavioural evidence provide?</i>	75
3.5 <i>Future Directions</i>	78
References	80
 Appendix	93
A.1. <i>Consent form</i>	93
A.2 <i>The KINARM in detail</i>	96
A.3 <i>Protocol scripts</i>	101

Abstract

An emerging paradigm shift is currently underway in neuroscience involving the modelling of neural systems using the mathematical framework of Bayesian decision theory, and more significantly, treatment of the brain itself as a Bayesian machine. Recent work suggests that the brain represents probability distributions and performs Bayesian integration during sensorimotor learning, but the evidence remains inconclusive. In this study, we seek to accrue additional behavioural evidence for Bayesian sensorimotor learning. Using a novel variation of an interlimb generalisation paradigm involving a stochastic visuomotor perturbation, we tested the hypothesis that Bayesian sensorimotor learning transfers to the other limb, and relatedly, that the representation of this learned visuomotor perturbation is encoded in an extrinsic reference frame. Consistent with our hypothesis, we found that interlimb transfer of learning is nearly complete when the visuomotor perturbation is congruent in extrinsic coordinates and relatively incomplete when congruent in intrinsic coordinates. Interestingly, we also found that although the learned prior distribution transfers to the untrained limb, the likelihood is not optimally integrated by the untrained limb, indicating that the prior and likelihood are represented independently of one another. This study provides valuable information about the nature of the representations underlying Bayesian integration in sensorimotor learning and opens up a number of intriguing paths for future investigation.

Declaration

I declare that this work has not been submitted elsewhere as either part of a degree or for publication, in any university or institution. The sources of information for this work was wholly based on my own research, except where referenced. The writing of this thesis was independent and is free from plagiarism. This research was approved by the Macquarie University Ethics Committee (protocol number: 5201600282)

A handwritten signature in black ink, appearing to read 'Chris Hewitson', is written above a horizontal line.

Chris Hewitson

Acknowledgements

First and foremost, I would like to acknowledge my principal supervisor, Dr. David Kaplan. David has a wealth of knowledge and experience in both the Philosophical and Empirical domains of Cognitive Science, and it was this breadth and depth of experience that motivated me to study under him. Over the course of this incredibly short and intense project, David has encouraged and supported me in my pursuit of a project that encompasses core elements from both fields, which, although immensely challenging at times, has enabled me to develop a diverse range of skills and ideas ripe for future investigation. I have been incredibly fortunate for the time David has contributed – above and beyond – to furthering my learning and development, and for this, I am eternally grateful.

Secondly, I would like to thank Associate Professor Paul Sowman, who, although being officially announced as my associate supervisor only late in the project, provided invaluable advice and timely programming support at crucial points throughout the entire project. For this, I thank him very much.

Thirdly, I would like to thank Jordan Wehrman, a PhD student in Cognitive Science who also provided invaluable programming support, in addition to a patient and receptive ear toward which I could air my (often strange) ideas in relative safety.

Further, acknowledgement is deserved of my fellow Master of Research students, with whom I could share my woes (and successes) in an unconditionally empathetic (and sympathetic) manner. I congratulate them on the submissions of their own work and wish them all the very best for their future endeavours.

Furthermore, I would like to thank the members of the Macquarie University Cognitive Science Department, especially the administrative staff who work tirelessly behind the scenes to ensure that students like me have the luxury to attend to their projects. Neither the department, nor my completed thesis, would exist without their efforts.

Finally, I would like to thank Sara, whose patience and support knows no bounds.

1. General Introduction

There is a major paradigm shift underway in contemporary cognitive science and neuroscience towards thinking about the brain as a Bayesian machine that encodes probability distributions and performs probabilistic inference (Rao et al. 2007; Clark 2013, 2015; Pouget et al. 2013). This idea has been termed the “Bayesian coding hypothesis” (Knill and Pouget 2004). Consistent with this trend, Bayesian modelling approaches are becoming increasingly prevalent (Knill and Richards 1996; Knill and Pouget 2004; Doya et al. 2006; Körding 2007, 2014; Vilares and Körding 2011; Ma 2012, Ma and Rahmati 2013; Lee and Wagenmakers 2014). These approaches combine the powerful mathematical frameworks of Bayesian statistics (Bayes and Price 1763) and Bayesian decision theory (Berger 2013). Bayesian statistics formalises how new information should be combined with prior beliefs and decision theory specifies how those updated beliefs lead to optimal decisions. These approaches have been employed to model aspects of vision (Knill and Richards 1996; Kersten et al. 2004; Stocker and Simoncelli 2008), multisensory integration (van Beers 1996, 1999; Ernst and Banks 2002; Fetch et al. 2011, 2013), sensorimotor control (Körding and Wolpert 2006; Bays and Wolpert 2007; Orbán and Wolpert 2011; Franklin and Wolpert 2011; Wolpert 2007), and even neural coding and computation (Eliasmith and Anderson 2003; Ma et al. 2013; Pouget, Beck, Ma and Latham 2013; Ma and Jazayeri 2014).

Despite their popularity, Bayesian approaches have attracted their share of criticism (for a review, see Hahn 2014).¹ One persistent challenge centres on the claim that the neuroscientific evidence supporting Bayesian estimation in the brain is at best inconclusive and at worst extremely weak (Jones and Love 2011; Bowers and Davis 2012a, b). For instance, Bowers and Davis (2012a) disparage Bayesian models as providing highly unconstrained “just-so stories”

¹ It is important to note up front that the thesis does not address a very different yet similar sounding debate in contemporary cognitive science and neuroscience concerning the appropriateness of Bayesian statistical analysis methods (Bayes factors) over null hypothesis significance testing (see, e.g., Kruschke 2010; Cumming 2013; Körding 2015; Puga et al 2015). Instead, the central focus here is whether Bayesian ideas provide a fruitful framework for understanding how the brain represents or encodes information, namely, in the form of probability distributions (i.e., the Bayesian coding hypothesis).

that are only weakly tied to neural data. Another challenge emphasises how the extreme flexibility or arbitrariness in the choice of model parameters including priors, likelihoods, and utility functions can make Bayesian models virtually unfalsifiable (Bowers and Davis 2012a, b). Other criticisms focus on the predictive and/or explanatory limitations of Bayesian models. Bowers and Davis (2012a) maintain that Bayesian models are, in many circumstances, predictively equivalent to non-Bayesian models. Colombo and Seriès (2012) and Colombo and Hartmann (2015) argue that Bayesian models are non-mechanistic and lack explanatory power when they fail to describe underlying neural mechanisms.

The aims of the project are to shed further light on the nature of Bayesian computation and representation in the brain, and to deepen our understanding of Bayesian modelling approaches in neuroscience. The primary aim of the project is to test a Bayesian hypothesis concerning human sensorimotor learning. Recent work suggests that the brain both represents probability distributions and performs Bayesian integration during bouts of sensorimotor learning involving visuomotor perturbations (e.g., Körding and Wolpert 2004; Pouget et al. 2013). However, evidence for these claims remains inconclusive.

In this study, we seek to provide additional behavioural evidence to further constrain Bayesian accounts of sensorimotor learning. Specifically, we will test whether Bayesian integration during sensorimotor learning is limb-specific by asking whether subjects who are exposed to and learn the statistical distribution of a visuomotor perturbation task using one arm will show similar compensation when tested with the other arm.

Although interlimb transfer is a well-established paradigm for investigating visuomotor learning (e.g., Sainburg and Wang 2002; Wang and Sainburg 2003, 2004, 2005; Taylor et al. 2011; Carroll et al. 2014, 2016), it remains unknown whether the internal statistical representation of the task domain that subjects learn using one limb transfers across limbs, and if so, how rapidly and completely this transfer occurs. Accordingly, this study will provide critical information about aspects of the underlying processes and representations through which sensorimotor learning is achieved. Specifically, the study will help to determine whether

probability distributions are represented in an effector-specific or effector-general manner in the motor system.

A secondary aim of the project is to reflect more broadly on the role that Bayesian approaches play in contemporary neuroscience (Jones and Love 2011). Some have claimed that Bayesian models are not falsifiable (Bowers and Davis 2012a). Others have claimed that the approach is explanatorily bankrupt without supporting neural evidence. Based on the current study and related studies, it is argued that Bayesian models of behavioural phenomena can and do play important roles in cognitive and brain science. Despite their limitations, well designed behavioural studies can provide useful constraints on the discovery of underlying neural mechanisms and representations supporting Bayesian integration.

1.1 Types of sensorimotor learning

Before discussing Bayesian approaches to sensorimotor learning, it is important to understand in general terms what sensorimotor learning is and how it has been experimentally interrogated in recent decades. Sensorimotor learning can be informally defined as the learning of new “mappings” or “transformations” between sensory and motor variables (Krakauer and Mazzoni 2011; Wolpert et al. 2011). A tennis player learning to return a serve – which involves learning to transform visual information about the motion and speed of the ball into appropriate swings of the racket – is a prototypical example of sensorimotor learning. A novice computer user learning to control a mouse cursor is another, more mundane, example. Sensorimotor learning has been extensively studied (Wolpert et al. 2011).

1.1.1 Error-based versus reinforcement-based sensorimotor learning

Another important distinction for understanding the nature of sensorimotor learning is that between error-based and reinforcement-based learning (Wolpert et al. 2011). All forms of supervised learning (Jordan and Rumelhart 1992) depend on the availability of an error signal that can be used to improve behaviour over time. Error-based learning is a form of supervised learning in which the error signal is “signed”, and therefore provides information about both the direction and magnitude of the error between the actual and desired behavioural outcome (Jordan and Rumelhart 1992; Kuperstein 1988; Kawato et al. 1987). Error-based learning is thought to underlie improvements in behavioural performance in many well-established motor adaptation paradigms including prism adaptation (Martin et al. 1996), visuomotor adaptation (Krakauer et al. 2000), and force field perturbations (Shadmehrand Mussa-Ivaldi 1994; Flanagan and Wing 1997). Learning in these paradigms is often extremely rapid, occurring over the course of individual trials and single experimental sessions (Lackner and Dizio 1994; Shadmehr and Mussa-Ivaldi 1994). While error-based learning is of primary relevance to the current project, there are other learning processes that have been investigated in the sensorimotor domain. For instance, reinforcement learning is another (less relevant) form of supervised learning in which the error signal is inherently “unsigned”, providing information about success or failure without supplying information about the direction or magnitude of the error to update behaviour on subsequent trials (Wolpert et al. 2001; Taylor et al. 2014; Wu et al. 2014).

1.1.2 Kinematic versus dynamic sensorimotor learning

Another important distinction is between learning kinematic transformations versus learning dynamic transformations (or simply, kinematics learning versus dynamics learning) (Atkeson 1989; Shadmehr and Wise 2005). Kinematic learning refers to learning the transformations or mappings between positional variables such as joint angles and hand position (Atkeson 1989; Shadmehr and Wise 2005). For example, learning to move the hand to a target involves, at a

minimum, learning what joint angles the muscles need to achieve in order to move the limb segments to that target position. These learned sensorimotor mappings are sometimes distinguished in terms of the involvement of different kinds of internal representations or internal models. Because they capture the inverse of the normal causal relationship between motor variables (actions) and their consequences, they are often referred to as inverse internal models (Kawato 1999; Wolpert and Ghahramani 2000). Investigations of kinematics learning have relied predominantly on prism adaptation (Martin et al 1996) and visuomotor adaptation paradigms (Krakauer et al. 2000).

In the standard visuomotor adaptation paradigm, a visual perturbation is introduced that distorts the normal mapping between movements and their visual consequences, while proprioceptive feedback remains unchanged (for review, see Krakauer 2009). For example, visual feedback of the subject's hand position might be rotated clockwise by 90 degrees so that straight ahead reaches result in feedback consistent with what is normally experienced during rightward movements. Anyone who has inadvertently had their wireless mouse rotated 180 degrees on their desk, has informally experienced a visuomotor perturbation. In experimental contexts, visuomotor adaptation is typically achieved either by having subjects wear prism goggles (Held and Freedman 1963; Martin et al. 1996a,b) or by withholding visual feedback of the real hand and having subjects move a cursor to targets on the screen in which the relationship between cursor position and hand position is manipulated (Krakauer et al. 1999, 2000). This manipulation alters the normal kinematic mapping between motor output and sensory feedback. For example, having subjects wear prism goggles that shift the visual field by x degrees laterally will produce initial reaches that are off by x degrees. With practice, subjects will gradually minimize their motor errors and eventually learn to reach targets accurately. This recovery of normal motor performance provides evidence that they have successfully learned the new kinematic mapping. An additional source of evidence that sensorimotor learning has occurred comes from the observation that when goggles are removed, subjects make inaccurate reaches

(missing by $-x$ degrees). These training-induced errors, known as *aftereffects*, are signatures of the sensorimotor learning process.

Kinematic transformations are not, however, the only type of transformation that must be learned to generate accurate and appropriate motor behaviour. Consider the difference between learning to lift a full cup of coffee compared to learning to lift an empty cup. Although the hand trajectories (kinematics) to be learned in the two cases may be the same, the appropriate torques on the elbow that need to be learned will differ. Dynamics learning refers to the process by which the motor system learns precisely how much force/torque to produce on each joint to generate the desired movement (Atkeson 1989; Shadmehr and Wise 2005). This fundamentally concerns learning the appropriate mapping that relates input forces and/or torques on the joints to output motions of the limb.

Dynamics learning has been primarily studied using another adaptation paradigm, called force-field adaptation, in which the dynamics of a task can be systematically altered while kinematics remain unchanged (Shadmehr and Mussa-Ivaldi 1994; Shadmehr et al., 2010). Force-field paradigms involve subjects reaching to visual targets in the presence of an externally-imposed force field, which alters the normal mapping between input forces on the limb joints (motor commands) and desired movement outputs (as reflected in visual and proprioceptive feedback). As indicated above, this mapping or transformation is termed dynamics. Force-field adaptation experiments are typically performed by having subjects grasp the handle of a robotic manipulandum that produces forces on the hand that change with hand motion (Shadmehr and Mussa-Ivaldi 1994; Shadmehr et al. 2013). Because force fields change the dynamics of the task, subjects will initially make inaccurate reaches (compared to those made in the absence of externally-imposed forces, termed the so-called null field). Subjects will eventually learn the appropriate dynamical mapping and produce accurate reaches in the force field. After training, subjects will also show pronounced aftereffects in response to the sudden removal of the force field, providing additional evidence for motor adaptation or learning.

Although the theoretical distinction between kinematics and dynamics learning is clear enough, it remains controversial whether this tidy division is respected in the brain. Some investigators maintain that kinematic and dynamic transformations are learned independently, and are represented and stored in distinct areas of working memory (Krakauer et al. 1999). Others claim that kinematics learning can be the source of interference for dynamics learning and vice versa, which provides evidence that they are not represented independently (Tong et al. 2002). Consequently, it remains unclear the degree to which kinematics and dynamics learning interact, and whether the underlying representational substrates for these types of learning are distinct or shared.

1.1.3 Intrinsic versus extrinsic reference frames for sensorimotor learning

The brain must also learn motor tasks and plan movements in the appropriate reference frame (or reference frames) (Soetiching and Flanders 1992, 1995). For example, if the task is to reach for the coffee cup on the table in front of you, information about the spatial location of the cup might be represented in an extrinsic reference frame relative to some external feature of the workspace such as the centre of the table surface. Critically, spatial information represented in this way would be insensitive to changes in arm configuration (for instance, the cup remains at the same table-centred coordinates even if you were to move your arm). The same movement can also be represented in an intrinsic reference frame such as one centred on the shoulder or elbow joint of the limb being moved. In the toy example (Fig. 1), the location of the cup defined in intrinsic joint coordinates changes as a function of limb posture. The concepts of intrinsic and extrinsic reference frames are particularly important for studies of sensorimotor learning involving generalisation across limbs and will therefore be addressed in more detail below.



Fig. 1. Intrinsic and extrinsic reference frames.

1.1.4 Intralimb versus interlimb generalisation of sensorimotor learning

Generalisation patterns are frequently used to test and further constrain hypotheses about the processes and representations underlying sensorimotor learning (Thoroughman and Shadmehr 2000; Shadmehr 2004; Paz and Vaadia 2009). Generalisation may be defined as the process by which a skill (sensorimotor or otherwise) gained through training or experience in one context changes performance in another (Poggio and Bizzi 2004; Criscimagna-Hemminger et al. 2003). When generalisation is beneficial, it is termed transfer. When it is detrimental, it is termed interference. A motor skill may generalise across similar tasks within the same limb, referred to as intralimb generalisation (Wang and Sainburg 2004). For example, a subject may learn a new visuomotor mapping for a restricted set of movement directions in the workspace and then generalize this learned mapping to new directions (Bedford 1989; Ghilardi et al. 1995; Ghahramani et al. 1996; Wu and Smith 2013). A learned motor skill may also generalise to the other limb, referred to as interlimb generalisation (Criscimagna-Hemminger et al. 2003; Wang and Sainburg 2003; Taylor et al. 2011; Joiner et al. 2013; Stockinger et al. 2015; Shadmehr et al. 2010; Parmer et al. 2015).

1.2 Key experimental results

1.2.1 Intralimb generalisation of sensorimotor learning

Using versions of the standard visuomotor adaptation paradigm, a number of studies of reaching movements provide converging evidence for the hypothesis that hand kinematics are learned from errors in extent and direction in an extrinsically-defined reference frame (Flanagan and Rao 1995; Wolpert et al. 1995; Krakauer 2009). Results such as these are typically established by training subjects to learn a distorted visuomotor mapping in one region of the workspace (e.g., in a region to right of the body midline) and then immediately test how this learning generalizes to other regions of the workspace (e.g., in a region to the left of the body midline). The error patterns generated during generalisation provide useful (albeit limited) information about how sensorimotor learning is represented in the nervous system (Thoroughman and Shadmehr 2000; Shadmehr 2004; Paz and Vaadia 2009). One specific piece of information that can be gleaned from these experiments is whether the learned information is represented in an intrinsic or extrinsic reference frame.

As indicated above, one can define a movement in extrinsic coordinates by specifying its starting point and a movement vector. For a straight point-to-point movement, the movement vector is just the directed line segment connecting its start and end points. One can also define a movement in intrinsic coordinates by specifying the joint angle changes required to bring about the movement. Consider a simplified example of how the reference frame of sensorimotor learning can be tested. Suppose that during the training period, subjects are required to learn a new visuomotor mapping by performing reaches in a region of the workspace to the right of the body midline (right workspace). During the test period, subjects are required to make movements in a region of the workspace to the left of the body midline (left workspace). In one test condition, the visuomotor mapping experienced over a number of trials in the left workspace is the same as the previously learned mapping in the right workspace only when specified in extrinsic coordinates. In another test condition, the visuomotor mapping is the same as the previously learned mapping only in intrinsic coordinates. If prior learning generalizes

(transfers) more readily or more completely in one condition over the other, this is taken as evidence that the observed sensorimotor learning is internally represented in that reference frame.

Along these lines, Krakauer et al. (1999) demonstrated that intralimb generalisation (and more specifically, transfer) of kinematic learning predominantly occurs in an extrinsic reference frame. Krakauer (2000) demonstrating that adaptation to a 30° rotation resulted in 100% generalisation to novel targets that deviated as much as 22.5° from the trained directions. Wang and Sainburg (2005) extended these findings by demonstrating generalization of 85% when reaching to targets that deviated by 45° from the trained directions, and 60% generalization to a direction that deviated nearly 90° from its closest training direction.

The claim that kinematic adaptation is encoded in an extrinsic reference frame was strengthened by Krakauer et al. (2000) who demonstrated that accuracy in reaching movements after adaptation to visuomotor rotations is achieved by using errors in extent and direction to adaptively update a vectorial representation of intended movement in extrinsic coordinates. Wang and Sainburg (2005) went on to confirm this vectorial hypothesis of movement planning by demonstrating that the brain integrates information regarding the location of the target and the initial position of the hand to form a hand-centred plan of the intended movement trajectory as a direction and extent in extrinsic space.

In another seminal motor learning study, Shadmehr and Mussa-Ivaldi (1994) developed the force-field adaptation paradigm to test intralimb generalisation of dynamics learning. They showed that subjects generalized what they had learned about the force field environment in an intrinsic (joint-based) reference frame. In other words, subjects generalized learning when the force field encountered during the test period was the same as the field encountered during training when specified in intrinsic, joint-based coordinates. No transfer of sensorimotor learning was observed when the training and test phases were only invariant in extrinsic coordinates and did not match in terms of arm joint configurations). The representation of dynamics learning in intrinsic, joint-based coordinates has subsequently been confirmed in a

number of studies (Brashers-Krug et al.1996; Gandolfo et al. 1996; Conditt et al. 1997; Flanagan et al. 1999; Shadmehr and Mussa-Ivaldi 2000;).

These studies all reveal a general pattern that holds up with a high degree of consistency across studies: Kinematic learning is predominantly encoded in extrinsic coordinates, while dynamic adaptation is encoded in intrinsic coordinates.

More recent studies suggest a more complex picture. For example, Brayanov et al. (2012) provided behavioural evidence that visuomotor learning relies on a gain-field combination of representations in both intrinsic and extrinsic coordinates. Berniker et al. (2014) similarly demonstrated that no single coordinate frame accounted for the generalization pattern following force-field adaptation. Instead, they argued that the representation of learning involved a mixture of reference frames.

1.2.2 Interlimb generalisation of sensorimotor learning

The above studies specifically test generalisation of sensorimotor learning within a single limb. Another important set of findings about sensorimotor learning relate to generalisation of learning from one limb to the other, known as interlimb generalisation. Subjects who adapt and learn new kinematic or dynamic mappings using one hand may show no adaptation, full adaptation, or some degree of partial adaptation in between these two extremes if asked to switch to performing the task with the opposite hand. The degree of adaptation (or transfer or interference) provides valuable information about the characteristics of the internal representation acquired by the subject during the previous episode of sensorimotor learning.

In contrast to studies of intralimb generalisation, there is considerably more divergence in the findings concerning interlimb generalisation of prior sensorimotor learning.

The transfer of visuomotor adaptation, whereby subjects maintained a fraction of their adapted performance when completing the same task with their opposite limb, has been observed predominantly across extrinsic coordinates (Imamizu and Shimojo 1995; Sainburg and Wang 2002; Wang and Sainburg 2004, 2005; Taylor et al. 2011).

In contrast, when performing reaches with their opposite limb to opposite rotations (incongruent across extrinsic coordinates) transfer was absent altogether (Wang and Sainburg, 2004; Wang 2008). Further, while numerous studies have indicated that kinematic transfer is asymmetrical; from non-dominant to dominant limb but not vice versa (Anguera et al. 2007; Sainburg and Wang 2002; Wang and Sainburg 2004) some evidence exists to the contrary (Balitsky and Henriques 2010; Isaias et al. 2011; Taylor et al. 2011; Poh et al. 2016).

In order to test interlimb generalisation of dynamics, Criscimagna-Hemminger et al (2003) adapted the standard force-field paradigm used to investigate intralimb generalisation. They report that interlimb transfer only occurs in extrinsic coordinates, a finding reiterated by Stockinger et al. (2015). These findings are striking because it is seemingly at odds with the conventional wisdom in the field that dynamics are represented and learned (at least as far as intralimb generalisation is concerned) in intrinsic coordinates (Shadmehr and Mussa-Ivaldi 1994, 2000). Interestingly, dynamics learning transfers also occurred asymmetrically from the dominant to non-dominant limb, but not vice versa; and transfer even occurred in one subject for whom the corpus callosum had been completely severed (Criscimagna-Hemminger et al. 2003).

The jury is, however, still out on the nature of interlimb generalisation of sensorimotor learning. Several more recent studies have reported contradictory findings. For example, Wang and Sainburg (2004) report significant transfer in joint-based coordinates, and Burgess et al. (2007) describe weak but significant non-dominant to dominant transfer. Recently, Carroll et al. (2014, 2016) has responded by insisting that these studies are inconclusive because they both share a key methodological limitation. Specifically, they argue that these studies uniformly rely on a corrupted measure of transfer – size of the movement kinematic (movement) error. This measure, they maintain, is problematic because a number of other processes could contribute to straighter trajectories (smaller kinematic errors), which have little if nothing to do with transfer. These include changes in limb impedance and/or changes in feedback gains. Hence,

inferring transfer on the basis of the size of the kinematic error alone may not provide definitive information about the degree to which a novel sensorimotor map of dynamics – learned by one limb – is available to the other limb.

Joiner et al. (2013) recently re-examined dynamical transfer using an “error-clamp” paradigm specifically designed to avoid this limitation. They restricted all movement kinematics to a straight path to the target, and measured the lateral forces made by subjects against the walls of a virtual force channel. Accordingly, it was hypothesized that adaptive impedance compensation would no longer occur and hence the lateral forces measured should more directly reflect the transfer of learned dynamics to the opposite limb. Using this approach, they report a small but significant transfer (approx. 12%) in an extrinsic reference frame.

Given that dynamics are conventionally thought to be represented in joint-based coordinates, the findings of Joiner et al. (2013) are also surprising. One possible explanation for these results involves appeal to the idea that sensorimotor learning or transformations do not generalize globally according to any single reference frame, but rather appear to involve flexible or context-dependent representation in some combination of multiple reference frames (Sober and Sabes 2003, 2005; Berniker et al. 2014; Parmar et al. 2015). This raises the question of how sensorimotor learning (and especially interlimb generalisation) is affected when the reference frames in which the tasks are learned are aligned or misaligned across the limbs.

To investigate this issue, Carroll et al. (2016) test interlimb transfer of learning of a velocity-dependent force field oriented in the sagittal plane (aligned in both intrinsic and extrinsic reference frames across the limbs) and compared it to interlimb transfer of learning of a velocity-dependent force field oriented in the transverse plane (misaligned in intrinsic but not extrinsic reference frames across the limbs). They report that the interlimb transfer of learning in the sagittal force field was significantly greater than the transfer of learning for the transverse field. Consequently, Carroll et al. (2016) conclude that dynamic transfer is contingent upon an intrinsic reference frame after all, since it is reduced when the learned dynamics are misaligned in intrinsic coordinates across the two limbs. Their results are consistent with the recent

hypothesis that learned dynamics are internally represented according to both the intrinsic and extrinsic coordinates of the sensorimotor context experienced during adaptation.

One potential limitation of all visuomotor and force field adaptation studies discussed above is that subjects are required to learn and internally represent what we might call stationary (non-changing) perturbations or task contexts. For instance, in one experiment the learned visuomotor perturbation might be a fixed rotation of visual feedback information by 45 degrees or lateral displacement of 2 cm. Similarly, the velocity-dependent force field subjects are required to learn may be a constant 10 nM force in one direction (e.g., perpendicular to the direction of movement). Accordingly, most of the generalisation results that have been reported in the literature involve generalisation over deterministic tasks. But what about learning and generalisation in probabilistic tasks? It seems one may reasonably contend that although perturbations can be stationary in a controlled experimental environment, this is an exception and not the rule when it comes to sensorimotor learning in the real world rife with variability and uncertainty. If true, this raises questions about the external validity of the results canvassed above. As a minimum, it opens up new questions about what sensorimotor learning might look like in task contexts involving uncertainty.

1.3 Sensorimotor learning under uncertainty: Does the brain perform Bayesian integration?

Uncertainty about the state of the environment and our bodies reflects both intrinsic sensory and motor variability or noise, and places a fundamental constraint on all information processing in the brain (Faisal et al. 2008). For example, a descending fog may increase the degree of uncertainty about whether the silhouette is an oncoming vehicle or a tree. Moreover, even if noise could somehow be eliminated, other physical limitations in our sensory systems entail that we only receive partial and ambiguous information about environmental properties (Yuille and Kersten, 2006; Purves and Lotto 2011).

The stochastic physical and chemical nature of our sensory receptors and the intrinsic noise associated with neural signalling and transformations means that neural signals plagued

with uncertainty and noise are the rule and not the exception in the brain (Faisal et al. 2008). Noise is present at all stages of sensorimotor control, from sensory processing, through planning, to the outputs of the motor system (Faisal et al. 2008). Sensory noise contributes to variability in estimating both internal states of the body (e.g., position of our hand in space) and external states of the world (the location of a cup on a table). Noise also contaminates the planning process leading to variability in movement endpoints (Gordon et al. 1994; Vindras and Viviani, 1998; Harris and Wolpert 1998) and is reflected in neuronal variability of cortical neurons that can predict variability in reaching (Churchland et al. 2006). In addition, variability in action can arise through noise in motor commands (van Beers et al. 2004). There is evidence that the major reason for the signal-dependent nature of this variability may come from the size principle of motor unit requirement (Jones et al. 2002).

Given these basic facts, it seems undeniable that the brain must routinely contend with uncertainty. One well established way to make inferences (or more generally, process information) under uncertainty is to operate according to Bayesian principles.

1.3.1 Bayesian sensorimotor learning

Bayesian statistics formalizes the optimal inferential strategy under circumstances where uncertainty or noise are present. Consider estimating the trajectory of an incoming tennis ball based on noisy or incomplete visual information (such as when playing at dusk). One way to supplement this impoverished visual input is to augment it with information based on your prior history. This includes your prior interactions with tennis balls and their behaviour (e.g., that they eventually fall to the ground, that they move along continuous rather than discontinuous trajectories), your prior history playing the game of tennis (e.g., that it is against the rules to win a serve by hitting over the line), and even your prior history playing against this specific opponent (that they favour their backhand). The distribution characterizing the probability assigned to each possible ball trajectory at a given time constitutes the posterior probability, $P(A|B)$. Calculating the posterior requires two pieces of information. First, on the basis of prior

experience, you need to know (or have an estimate of) the probability of each potential ball trajectory occurring, termed the prior distribution $P(A)$. Second, for any given sensory input, you need an estimate of how probable that particular input is for different possible states of the world that might have generated it. This is termed the likelihood, $P(B|A)$. For example, sensing a curved trajectory is more likely if the ball is spinning than if it is not. From the prior and likelihood, the posterior can then be estimated simply by multiplying the prior by the likelihood (and normalizing the result to sum to a probability of one) according to Bayes' rule:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)} \quad (1)$$

1.3.2 Recent experimental work and next steps

Despite their growing popularity in cognitive science and neuroscience, the Bayesian coding hypothesis has been challenged along multiple fronts including, most critically, that it is poorly supported by neural evidence (Jones and Love 2011; Bowers and Davis 2012a). Although this remains an open challenge, there is a growing and increasingly diverse body of computational modelling and experimental research on how Bayesian inference might actually be implemented in neural systems (e.g., Zemel et al. 1998; Ma et al. 2006, 2008; Rao et al. 2002; Rao 2004; Fetsch et al. 2009, 2011, 2013; Vilares et al. 2012; Pouget et al. 2013; Ma and Jazayeri 2014; Rich et al. 2015; Ting et al. 2015; van Bergen et al. 2015; Dekleva et al. 2016; Tan et al. 2016; Palmer et al. 2016).

In addition to this important work, a broad range of findings from sensory and motor psychophysical experiments provide another complementary – albeit less direct – line of empirical support for the Bayesian coding hypothesis (e.g., Knill et al. 1996; Ernst and Banks 2002; Schrater and Kersten 2002; Adams et al. 2004; Körding and Wolpert 2004; Stocker and Simoncelli 2006; Tassarini et al. 2006; Maloney and Mamassian 2009; Brayanov and Smith 2010; Verstynen and Sabes 2011; Fernandes et al. 2012, 2014; Sato and Körding 2014; Acuna

2015). Although the inferences that may be drawn on the basis of behavioural evidence alone are limited, it is our contention that well-designed psychophysical experiments can nonetheless provide highly useful constraints on the search for neural mechanisms (Colombo and Hartmann 2015), and in doing so can help supply additional support for the Bayesian coding hypothesis.

It is for this reason that we sought to use behavioural experiments to test and further constrain Bayesian accounts of sensorimotor learning (Körding and Wolpert 2004; Fernandes 2012, 2014). In this study, we take the next step forward in understanding Bayesian sensorimotor learning by asking about the reference frame in which this learning occurs and whether the learned representations are limb- or effector-specific. We do this by probing whether subjects who learn the statistical distribution of a visuomotor perturbation task using one arm will show similar compensation when tested with the other arm. Although interlimb transfer is a well-established paradigm for investigating visuomotor learning and the nature of the underlying representations involved (e.g., Sainburg and Wang 2002; Wang and Sainburg 2003, 2004, 2005; Taylor et al. 2011; Carroll et al. 2014, 2016), it remains unknown whether the internal statistical representation of the task domain that subjects learn using one limb transfers across limbs, and if so, how readily this transfer occurs. The current study provides valuable information about whether Bayesian integration in sensorimotor learning is represented in an effector-specific manner in the human motor system, and relatedly, whether these representations are encoded in an extrinsic or intrinsic reference frame.

2. Investigating Interlimb Generalisation of Bayesian Sensorimotor Learning

2.1 Introduction

Recent work suggests that the brain both represents probability distributions and performs Bayesian integration during episodes of sensorimotor learning involving visuomotor perturbations (e.g., Körding and Wolpert 2004; Tassinari et al. 2006; Fernandes et al. 2012, 2014). Understanding precisely how these statistical distributions are represented and learned remains a key open challenge for the field. Although direct neural evidence remains limited and difficult to interpret (Vilares et al. 2012), other more indirect routes for investigating these underlying representations are available.

Generalisation – the process by which training or experience in one context changes performance in another – is widely thought to provide a useful window into the representational changes underlying other forms of sensorimotor learning (Thoroughman and Shadmehr 2000; Poggio and Bizzi 2004; Shadmehr 2004; Paz and Vaadia 2009). It is well-established that sensorimotor learning generalizes across similar tasks using the same limb (intralimb generalisation), such as when subjects learn a novel visuomotor mapping (kinematic transformations from desired limb position into joint angle changes) for a restricted set of movement directions in the workspace and then generalize this learning to untrained directions (Bedford 1989; Ghilardi et al. 1995; Ghahramani et al. 1996; Malfait et al. 2002; Wu and Smith 2013). Intralimb generalisation of prior learning of limb dynamics (dynamic transformation from desired limb trajectories into joint torques and forces) resulting from exposure to novel force field environments has also been reported (Shadmehr and Mussa-Ivaldi 1994, 2000; Thoroughman and Shadmehr 2000). Sensorimotor learning of new kinematic and dynamic transformations also generalizes from the trained to the untrained limb (Criscimagna-Hemminger et al. 2003; Wang and Sainburg 2003, 2004; Taylor et al. 2011; Joiner et al. 2013; Stockinger et al. 2015; Shadmehr et al. 2010; Parmer et al. 2015 Carroll et al. 2014, 2016).

Generalisation studies can provide valuable information about the reference frame in which the learning occurs as well as the specificity of the underlying representation of the learned task. Specificity is here defined as the degree to which training effects are restricted to the specific conditions experienced, which includes the local region of the workspace or the limb used during training. Investigating the reference frame in which a motor task is learned is important as the brain appears to represent a given task in terms of one or more coordinate systems, for instance, extrinsic (i.e., screen-centred), or intrinsic (i.e., joint-centred) (Sober and Sabes 2003, 2005; Carroll et al. 2014).

Error patterns generated during generalisation from training are typically used to determine how sensorimotor learning is represented in the nervous system including whether the learned information is encoded in an intrinsic or extrinsic reference frame. Consider a simplified example that illustrates how this is done. Suppose that during the training period, subjects are required to learn a new visuomotor mapping (e.g., a counterclockwise rotation of visual feedback relative to the real hand position defined in extrinsic, screen-based coordinates). In one test condition, the visuomotor mapping experienced over a number of novel trials to new untrained directions is the same as the previously learned mapping only when specified in extrinsic (e.g., screen-based) coordinates. In another test condition, the visuomotor mapping is only the same as the previously learned mapping when specified in intrinsic (e.g., joint-based) coordinates (i.e., relative to the joint configuration of the arm). If prior learning generalizes (transfers) more readily or more completely in one condition over the other, this provides evidence that the observed sensorimotor learning is internally represented in that reference frame. If prior learning fails to generalize, this suggests that the reference frame in which learning occurred is different and incompatible from that in which testing occurred, and perhaps that the internal representation may be limb- or effector-specific.

A number of visuomotor adaptation studies involving a single limb provide converging evidence for the hypothesis that hand kinematics are learned from errors in extent and direction in

an extrinsically-defined reference frame such as one anchored to the centre of the screen or physical workspace (Flanagan and Rao 1995; Wolpert et al. 1995; Krakauer 1999, 2000; Vetter et al. 1999).

Clearly, a considerable amount is now known about how learning of fixed or deterministic visuomotor perturbations (involving a single fixed rotation or 1-dimensional shift in the visuomotor mapping) generalise (or fail to generalise) to new contexts with the same limb or the other limb. By comparison, relatively little is known about intralimb generalisation of prior learning of indeterministic or stochastic visuomotor perturbations (Körding and Wolpert 2004; Tassarini 2006; Fernandes et al. 2012, 2014), involving a probability distribution of rotations or shifts in the visuomotor mapping.

While no previous study has investigated interlimb transfer of a statistical adaptation, there is some behavioural evidence for distinct representations of the prior and likelihood (Beierholm 2009; Shams 2012) as well as limited neural evidence (Vilares et al. 2011). Fernandez et al. (2012) investigated how prior uncertainty generalized during reaching by having different groups of subjects groups adapt to a visual rotation with the same mean (30°) but different levels of feedback uncertainty (standard deviations of 0° , 4° or 12° for the different experimental groups). They found that subjects rapidly adapted to the mean, however, as the uncertainty of the perturbation increased, the rate of adaptation slowed. Despite this slower rate, the breadth of generalisation was the same across all groups, limited to a width of $\pm 30^\circ$. Fernandez et al. concluded that generalisation of a stochastic visuomotor perturbation is independent of feedback uncertainty, which supports the hypothesis that the prior and likelihood are independently represented. While these findings pertain to intralimb generalization, to the best of our knowledge, no one has investigated interlimb transfer of a statistical prior or integration of the likelihood distribution by the untrained limb.

2.2 General Material and Methods

2.2.1 Ethics

All participants gave informed consent before the experiment.² All experimental protocols were approved by the Macquarie University Human Research Ethics Committee (protocol number: 5201600282).

2.2.2 Participants

35 right-handed subjects (22 males, 13 females, age 17-49 years) with normal or corrected to normal vision and no history of motor impairments participated in the experimental study. Subjects were either paid and recruited from the Macquarie University Cognitive Science Participant Register (<https://mq-cogsci.sona-systems.com>) or were Macquarie University students participating for course credit. All participants were randomly assigned to 1 of 5 experimental groups across 3 experimental paradigms. Seven subjects participated in Experiment 1, which consisted of a stochastic visuomotor adaptation task. Fourteen subjects participated in Experiment 2, which sought to test interlimb transfer of visuomotor learning using the same basic task from Experiment 1. Seven subjects participated in Experiment 3, which consisted of a variation of the basic visuomotor adaptation task used in Experiment 1.

2.2.3 Equipment

A unimanual KINARM endpoint robot³ (BKIN Technologies, Kingston, Ontario, Canada) was utilized in all experiments (Fig. 2).⁴ The KINARM has a single graspable manipulandum that permits unrestricted 2D arm movement in a horizontal 2D plane (the movement plane). A

² See Appendix 1 for ethic approval and a copy of the consent form used.

³ See Appendix 2 for additional detail pertaining to the KINARM setup.

⁴ Because the KINARM was brand new to our lab, all programming of the KINARM experimental paradigm was done from scratch by C.H. Programming was done using MATLAB Simulink.

projection-mirror system facilitates presentation of visual stimuli that appear in the movement plane. Subjects received visual feedback about their hand position via a cursor (white sphere, 1 cm diameter) that was controlled in real-time by moving the manipulandum. Mirror placement and an opaque apron attached to the participant ensured that visual feedback from the real hand was not available for the duration of the experiment.




Fig 2. Experimental setup and paradigms.

2.2.4 General Experimental procedure

Subjects were instructed to perform fast and accurate reaching movements with the dominant (right) arm using cursor feedback, whenever it was available. Reaches were from a start target located at the centre of the workspace to a single end target located 20 cm away (Fig. 3a). When subjects moved the cursor within the boundaries of the start target, its colour changed from blue to red indicating the start of a trial. Subjects were free to reach at any time after the target colour changed. Once the cursor exited the start target, cursor feedback was extinguished and laterally

shifted to the right of the true hand position (positive in the x -plane) by an amount drawn at random on each trial from a Gaussian distribution with mean of 1 cm and standard deviation of 0.5 cm (the *true prior*). At the midpoint of the movement, displaced cursor feedback was provided for 100 ms (*midpoint feedback*).

To test whether Bayesian integration occurs during sensorimotor learning, following K rding and Wolpert (2004), the reliability of the sensory feedback information provided about the true cursor position at the reach midpoint was varied by introducing different amounts of visual noise or blur (thus changing the degree of sensory uncertainty) on each trial. This allowed us to assess the subjects' reliance on their previously experienced distribution of shifts. One of four visual uncertainty conditions (σ_0 , σ_M , σ_L , σ_∞) (Fig. 3b) was selected at random on each trial with the frequencies of the 3, 1, 1, 1, respectively⁵. In the zero uncertainty condition (σ_0), midpoint feedback was a single white sphere (1 cm diameter), identical to the initial cursor. In the moderate uncertainty condition (σ_M), midpoint feedback was one of ten randomly generated point clouds⁶ comprised of 50 small translucent spheres (0.2 cm diameter) distributed as a two-dimensional Gaussian with a standard deviation of 1 cm and a mean centred over the true (displaced) cursor position on the current trial. In the large uncertainty condition (σ_L), everything was analogous to the moderate uncertainty condition (σ_M) except that the point clouds had a standard deviation of 2 cm. In the unlimited uncertainty condition (σ_∞), no midpoint feedback was provided. Cursor feedback was again extinguished for the remainder of the reach to the end target. Cursor feedback at the endpoint of the reach (endpoint feedback) was provided only in the zero uncertainty (σ_0) condition for a duration of 100 ms. After movement offset, there was a delay of 150 ms before the

⁵ Specifically, the clear condition σ_0 was presented 3 times more than the σ_M , σ_L , and σ_∞ conditions respectively. For instance, in a 1000 trial block, approximately 498 σ_0 trials and 166 σ_M , σ_L , and σ_∞ trials would be presented. As 6 does not divide into 1000 neatly, we chose to run 1080 trials per block (540 σ_0 /180 σ_M /180 σ_L /180 σ_∞).

⁶ See Appendix 3.1 for *point cloud* generation script.

start target reinitialised the next trial by changing colour from red back to blue. The maximum allowable time to complete a reach was 4000 ms⁷. Irrespective of the cursor's position along the x -axis, if subjects did not cross the lower bound of the end target along the y -axis (dashed line, Fig. 3a) the trial would time out. Timeouts were signalled by the disappearance of the end target and the start target changing back to blue.




Fig 3. Experimental paradigm. (a) KINARM workspace with example hand and cursor paths shown when a 2 cm visual shift is applied. Dashed white lines indicate feedback windows. (b) Midpoint feedback conditions with different amounts of visual uncertainty.

⁷ Despite explicitly instructing the subjects to complete the reaches as rapidly and accurately as possible, the maximum allowable reach time of 4000ms is conservative in comparison to a timeout period of 500ms-1s as previously reported during other visuomotor pointing tasks (Goodbody and Wolpert 1999; Baraduc and Wolpert 2002; Knill and Kersten 2004). In contrast to other pointing tasks that require subjects to reach unencumbered (whereby the motion of a reach trajectory is remotely tracked), our paradigm required subjects to grip accelerate and decelerate the kinarm manipulandum, hence additional reach time was deemed appropriate. Importantly, post-hoc safeguards were applied in order to discard reach outliers (see section 2.2.5 Analysis).

2.2.5 Analysis

Kinematic data including hand position and velocity was recorded for all trials using BKIN's Dexterit-E experimental control and data acquisition software (BKIN Technologies, Kingston, Ontario, Canada). Hand position data was recorded at 200Hz and logged in Dexterit-E. Custom scripts for data processing and analysis were written in MATLAB.⁸ Hand position, velocity, and cursor shift values were extracted from the c3d files in MATLAB as integers. A combined spatial- and velocity-based criterion was used to determine movement offset and corresponding reach endpoints (Georgopoulos et al. 1982; Moran and Schwartz 1999; Scott et al. 2001). Specifically, movement offset was defined as the first point in time t at which the movement dropped below a minimum velocity threshold ($<5\%$ of peak velocity), after a minimum reach of 19 cm from the start target in the y -plane. Reach endpoints were defined as the x - and y -values at time t . The additional spatial criterion ensured that data from the start of the trial (also $<5\%$ of peak velocity) was not included in subsequent analysis.

Since the visual shift was systematically applied along the x -axis, the primary measure of the subject's estimate of the visuomotor perturbation (the *estimated prior*) was their mean hand position (x -coordinate only) at the end of the reach (henceforth *endpoint*) for all reaches completed during the unlimited uncertainty (σ_∞) condition. If a subject estimates the lateral shift to be $+n$ degrees, then successfully acquiring the target requires the final hand position to be at $-n$ degrees. The unlimited uncertainty (σ_∞) condition was used because it provides the most uncontaminated measure of the subject's prior, as endpoints on those trials are not influenced by midpoint feedback.

Endpoint indicates the extent to which the prior distribution of shifts was integrated and provides an estimate of lateral shift. In Experiment 1, mean endpoint was computed across the entire testing phase (trial 1080-2160; testing phase) (Fig. 4a). In Experiment 2, the degree of

⁸ See Appendix 3 for representative Matlab scripts.

generalisation was assessed by comparing the mean endpoint at the end of the training phase (trial 980-1080; late training phase) with the mean endpoint at the start of the testing phase (trial 1080-1180; early testing phase) for the respective groups.




Fig 4. Experimental design.

The second measure of statistical learning was deviation from target (*cursor error*⁹) as a function of the applied shift. Specifically, we compared the slopes of the linear fits for these plots, sorted by visual uncertainty condition, in order to determine the degree to which subjects compensated for visual uncertainty by changing their reliance on their stored prior. In Experiment 1, cursor error as a function of shift (slope) was determined by averaging across the entire testing phase (trials 1080-2160). In Experiment 2, slope was assessed at the end of the training phase (trials 980-1080; late training phase) at the start of the testing phase (trials 1080-1180; early testing phase) for the respective groups (Fig. 4b,c). If subjects compensate fully for the visual feedback, then the average deviation from target for all visual conditions should be zero. If, however, subjects integrate the statistically learned prior and current visual evidence, then endpoint should move towards the mean of the prior and depend on sensory uncertainty (Körding and Wolpert 2004). Hence, for a Gaussian distribution of visual uncertainty, if subjects perform Bayesian estimation, a linear relationship between cursor error and the imposed shift is expected. Specifically, the linear fit will intercept the abscissa at the mean of the prior (1 cm) and have a slope that increases as a function of visual uncertainty.

A repeated measures ANOVA with planned pairwise comparisons was used to analyse mean endpoints across all subjects within experimental groups, and the slopes for all experiments. The Mauchly test was used to assess the sphericity of repeated measures effects of visual condition as it constitutes a four-level factor. If sphericity was violated, Greenhouse-Geisser degree of freedom corrections were applied. The significance level for all non-corrected contrasts was $\alpha < .05$. Statistical analysis was performed using SPSS v22.0 for Windows.

⁹ Hand position and cursor position are related as a function of the imposed shift. Specifically, cursor position is calculating as hand position + shift. *Cursor error* is the difference between the cursor position and the end target location at movement offset.

2.3 Experiment 1: Bayesian visuomotor adaptation (VA)

The primary aim of Experiment 1 was to test whether subjects learn to compensate for the imposed stochastic visuomotor perturbation (lateral shifts drawn from a distribution with fixed mean and SD). A secondary aim was to provide an approximate replication of the findings reported by Körding and Wolpert (2004). Before the experiment started, each subject performed 10 familiarization trials in which cursor feedback was always provided and no lateral shift was imposed. Further, for 2 of the 7 subjects tested, an additional Baseline task was run to measure each subject's baseline motor variability and directional biases when reaching with each hand. The Baseline task used the same basic paradigm as the other experiments and consisted of the following sequence: 10 feedback trials with the right hand (cursor feedback always provided; no lateral shift imposed), 10 no-feedback trials with the right hand (no cursor feedback provided; no lateral shift imposed), 10 feedback trials with the left hand (cursor feedback always provided; no lateral shift imposed), and 10 no-feedback trials with the left hand (no cursor feedback provided; no lateral shift imposed). After completing the Familiarization and Baseline tasks, subjects completed 2160 trials of the Visuomotor Adaptation (VA) task with their right hand (Fig 2a,b). For the purposes of comparison with Körding and Wolpert (2004), and comparison with data from the transfer experiment (Experiments 2) we nominally defined the training phase as the first 1080 trials in each session and the testing phase as second 1080 trials in each session (Fig. 4a). In this experiment, there were no objective differences between these phases.

2.3.1 Models of visuomotor adaptation

Following Körding and Wolpert (2004), we consider three possible computational models of sensorimotor integration that subjects could use to reach accurately to the target on the basis of the visual feedback provided.

2.3.1.1 Full compensation model

One possibility is that subjects fully compensate for the sensed lateral shift (Fig. 5a). According to this model, increasing the uncertainty of the feedback for an imposed shift would increase endpoint variability (variance) without changing the mean. Importantly, this model does not require subjects to estimate either visual uncertainty or the prior distribution of shifts applied.

2.3.1.2 Minimal mapping model

The minimal mapping model involves an iterative mapping from visual feedback about cursor error to an estimate of the imposed shift. This crucial error signal can be reduced over repeated trials, and an accurate estimate of the shift can be attained. While this model predicts a mean endpoint of 1cm to the left of the target (for a 1 cm rightward shift), indicating that the mean of the prior had been learned, it does not require an explicit representation of either the prior distribution or visual uncertainty (Körding and Wolpert 2004). All that is required in order to learn this mapping is an indication of cursor error at the end of the movement. However, in our paradigm, cursor error is only provided for the clear feedback condition (σ_0). Therefore, a mapping may only be learned based on this condition and then applied to all other conditions ($\sigma_M, \sigma_L, \sigma_\infty$) (hence the term minimal, for minimal condition mapping). Importantly, the minimal mapping model predicts a compensation pattern that is the same for all trials, regardless of visual uncertainty (Fig. 5b).

2.3.1.3 Bayesian estimation model

The final model considered is the Bayesian estimation model¹⁰, according to which subjects use information about the prior distribution and the uncertainty of the visual feedback in order to estimate the imposed shift. The *posterior probability distribution* can be obtained by applying Bayes' rule as follows:

$$P(x_{true}|x_{sensed}) = \frac{P(x_{sensed}|x_{true})P(x_{true})}{P(x_{sensed})} \quad (2)$$

Where x_{true} is the imposed shift, x_{sensed} is the sensed shift (the visual evidence) and $P(x_{true})$ the prior distribution of shifts. Assuming that the noise of each measurement is independently Gaussian (Fig. 6) then the optimal estimate of the imposed shift is a sum of the mean of the prior and the sensed feedback position ($\mu_{estimate}$) weighted by their relative variances [(σ_p^2) and (σ_s^2) respectively]:

$$\mu_{estimate} = \frac{\sigma_s^2}{\sigma_s^2 + \sigma_p^2} [1cm] + \frac{\sigma_p^2}{\sigma_s^2 + \sigma_p^2} x_{sensed} \quad (3)$$

Where $(\frac{\sigma_s^2}{\sigma_s^2 + \sigma_p^2})$ and $(\frac{\sigma_p^2}{\sigma_s^2 + \sigma_p^2})$ is the 'weighting' (degree of influence) attributed to the prior and visual information relative to their respective variance. Accordingly, the joint variance (σ_{sp}^2) of the posterior is given by:

¹⁰ It is important to note that while we are not employing Bayesian statistics to analyse our psychophysical data (we employ, instead, a frequentist approach the choice of statistical methods by which to analyse our data, in no way influences the fundamental hypothesis that we are testing; namely that the brain performs Bayesian sensorimotor integration. The credibility of this hypothesis is contingent upon whether or not subjects (a) learn the prior from a statistical distribution of priors and (b) integrate the degree of sensory uncertainty with the learned prior in order to estimate the imposed shift.

$$\sigma_{sp}^2 = \frac{\sigma_s^2 \sigma_p^2}{\sigma_s^2 + \sigma_p^2} \quad (4)$$

The Bayesian estimation model predicts that as visual uncertainty increases, the subject's estimate of the imposed shift moves away from the sensed shift and tends towards the mean of the learned prior distribution (Fig. 5c). For example, consider an imposed shift of 2cm. Given sensory uncertainty there are multiple shifts that can produce a sensed shift of approximately 2 cm (i.e., within the range of 1.8-2.2 cm). However, if visual uncertainty is a function of Gaussian noise on the visual feedback, then, according to the Bayesian model, the most probable shift is less than 2 cm, due to the influence of the learned prior. Hence, the estimated shift will tend toward the prior by an amount that depends on both the prior distribution and the degree of uncertainty in the visual feedback (Figs. 5c & 6). Furthermore, without visual feedback (σ_∞) the estimate should approximate the mean of the learned prior (because the likelihood distribution is flat).




Fig 5. Computational models. The average lateral cursor deviation from the target (cursor error) as a function of the imposed shift for the models. (a) Full compensation model, (b) minimal mapping model, and (c) Bayesian estimation model. (Transparent bands indicate the relative degree of variability in estimation). The colours of the linear fits correspond to the visual condition (matching Fig. 3b), as do the bands of variability in (c).




Fig 6. The experimentally imposed prior distribution of shifts is Gaussian with a mean of 1 cm (in black). The probability distribution of possible visually experienced shifts under the clear and the moderate and large uncertainty conditions are represented with solid lines (colours as in Fig. 3b) for a trial in which the imposed shift is 2 cm. The Bayes-optimal estimate of the shift that combines the prior with the evidence is represented by dashed lines (colours also as in Fig 3b).

2.3.2 Predictions

Based on the previous results of Körding and Wolpert (2004), we predicted that subjects would not only learn the prior distribution of imposed shifts, but would apply it in a fashion consistent with the Bayesian estimation model. Accordingly, we predicted that the (sign-inverted) mean endpoint across the entire testing block (trials 1080-2160) would closely approximate the mean of the learned prior of 1 cm, and that subjects would integrate the degree of visual uncertainty when estimating the imposed shift. It was also expected that cursor error would increase as a function of visual uncertainty as depicted in Fig. 5c (where increasing error is indicated by a larger slope). That is, subjects will estimate the imposed shift with a greater degree of accuracy during trials in which visual feedback is more reliable, and with accuracy decreasing across less reliable visual feedback conditions (accuracy during $\sigma_0 > \sigma_m > \sigma_L > \sigma_\infty$).

2.3.3 Results

The mean endpoint (σ_{∞}) across the experimental group was $-1.51 \pm 0.15\text{cm}$ (mean \pm SD) to the left of the target indicating that subjects had learned the average shift of 1 cm experienced over the ensemble of trials (Fig. 7).




Fig 7. Mean endpoints over the entire testing phase of Experiment 1 (VA). A mean end point of $-1.51 \pm 0.15\text{cm}$ (SD) indicates that the prior distribution had been learned. Note that endpoint error is plotted in screen coordinates, where 0 along the x-axis corresponds to the body midline.

Cursor error as a function of shift was averaged across 11 bins of applied shift values and plotted for all visual feedback conditions (representative subject in Fig. 8a). The slope of the linear fit was analysed in order to investigate the relationship between cursor error and the imposed shift. Mauchly's test for sphericity was violated ($p = .01$), requiring Greenhouse-Giesser correction. According to the corrected repeated measures ANOVA, slope increased significantly ($F_{3,81} = 14.1$ $p = .002$) with increasing uncertainty in the visual feedback (Fig. 8b).




Fig 8. Results for a Gaussian distribution. (A) Cursor error at the end of the trial as a function of the imposed shift for a representative subject. Colours as in Fig. 3b. Values represent Cartesian screen coordinates. Error bars denote standard deviation. The horizontal dotted line indicates the full compensation prediction, while the dashed line is the fit for the minimal mapping model that ignores sensory feedback on the current trial, correcting for the mean over all trials only. The solid line provides the Bayesian estimation, fitted to the data, relative to the sensory uncertainty. **(B)** The slopes of the linear fits for all subjects in Experiment 1 (the first bar in each grouping corresponds to the subject represented in panel a). A Greenhouse-Geisser corrected repeated-measures ANOVA of the slope, with visual uncertainty as the main factor, indicates a significant difference across visual conditions ($F_{3,81} = 14.1$ $p=.002$). A planned comparison of the slopes between visual conditions indicated significant differences for all visual conditions except the σ_0 conditions which was similar ($p = .46$) to condition σ_M .¹¹

2.3.4 Discussion

On average, across the entire 1080 trial testing block, subjects consistently reached to the left of the target with a mean endpoint of -1.51 ± 0.15 cm (mean \pm SD across subjects), indicating that they had learned the average visual shift of +1 cm (mean of the imposed prior) over the course of the testing phase. The mean overshoot of -0.51 cm may be explained by the size of the cursor relative to that of the end target. The use of a 2 cm diameter target, and a 1 cm diameter cursor may

¹¹ Due to trial scheduling statistics, the applied shift values differ slightly across each subject. Hence a lack of cursor error values for small shifts in the σ_L and σ_∞ conditions is evident for the individual subject depicted in Fig 8a. Shift values in this range were represented across other subjects, and critically, every subject experienced the same overall statistical distribution of shifts during the training phase.

have induced a spectrum of subjective accuracy. Although subjects were instructed to reach as accurately – to the origin of the target – as possible, no special feedback was provided on correct trials. Consequently, bringing the cursor anywhere within the circumference of the target may have been perceived as an ‘accurate’ reach by the subject (leading to no further corrections or compensation). Any pre-existing bias in any direction up to $\pm 1\text{cm}$ (the radius of the target) might therefore remain uncorrected through the experiment. To determine if this was the case, we collected baseline reach data for the last 2 of the 7 subjects in Experiment 1.¹² Baseline-adjusted mean endpoint averaged over those 2 subjects indicate a compensation of $-1.16 \pm 0.12\text{cm}$ (mean \pm SD) (Fig. 9), providing a closer correspondence to the mean of the true prior (and closer to the results reported by Körding and Wolpert 2004).




Fig 9. Baseline-adjusted mean endpoints across the entire testing phase (1080-2160) for subjects 6 and 7 from each group only. Mean baseline endpoints were $-0.35 \pm 0.03\text{cm}$ for right hand reaches, and $0.126 \pm 0.02\text{cm}$ for left hand reaches.

We then examined the relationship between imposed shift and cursor error. On trials in which feedback was provided, there was compensation during the second half of the movement (Fig. 3a, example hand and cursor paths for a trial with an imposed shift of 2 cm shown in blue).

¹² In retrospect, it would have been useful to collect baseline data for all subjects. We initially expected that any directional biases would be eliminated due to the sheer number of exposure trials in the experimental session. After data from 5 subjects across all the experiments was collected and analysed, we came to appreciate the potential target-cursor size tolerance issue and the role for baseline corrections.

The visual feedback midway through the movement provides information about the imposed shift. Cursor error as a function of shift is explicitly shown for a representative subject in Fig. 8a. It is also reflected in the slope plots for all experimental subjects in Fig. 8b. These plots show slopes that increase with increasing uncertainty, a pattern that is inconsistent with both the *Full Compensation* and *Minimal Mapping* models. In contrast, the *Bayesian Estimation Model* predicts slopes that increase with increasing uncertainty. Specifically, as visual uncertainty increases the influence of visual feedback on cursor endpoint is expected to decrease and reliance on the learned prior increases (Fig. 8a,b).

The influence of visual uncertainty on cursor error is evident across the population of subjects tested. For all 7 subjects, the slope increases significantly with visual uncertainty across three conditions ($\sigma_M, \sigma_L, \sigma_\infty$) but not for the clear condition (σ_0) (Fig. 8b). One possible explanation for this is that the clear and moderate uncertainty conditions provide similar information about the imposed shift (Fig. 3b). Although the stimuli used for the moderate uncertainty condition were randomly generated Gaussian point cloud distributions of 25 translucent spheres with a standard deviation of 1 cm, the origin of the moderate uncertainty feedback is still relatively easy to discern. Hence, from the subject's perspective, there might have been no effective difference between the clear (solid white sphere of 1cm diameter) and moderate uncertainty feedback conditions. This could have produced the similar slopes observed across these two conditions. Nevertheless, the influence of visual uncertainty on cursor error remains significant for all other comparisons. This indicates that subjects integrate their degree of uncertainty in the current visual evidence with the learned prior distribution in order to generate motor behaviour that compensates for the imposed shift, as predicted by the Bayesian estimation model.

2.4 Experiment 2: Extrinsic vs. intrinsic transfer

In comparison to intralimb generalisation studies, studies investigating the generalisation of visuomotor learning across limbs have been met with mixed results. For example, partial transfer of visuomotor adaptation has been reported when the perturbation imposed across limbs is congruent in an extrinsic reference frame (Imamizu and Shimojo 1995; Sainburg and Wang 2002; Wang and Sainburg 2004, 2005; Taylor et al. 2011;), while others note its absence (Wang and Sainburg, 2004; Wang 2008). Furthermore, while most studies report that transfer is asymmetrical, there is disagreement about the direction of transfer (dominant hand to non-dominant hand or vice versa) [Anguera et al. 2007; Sainburg and Wang 2002; Wang and Sainburg 2004]). Finally, Carroll et al. (2014) observe interlimb transfer exclusively when the perturbation is congruent across both extrinsic and intrinsic reference frames during an isometric force-aiming task.

Despite this broad range of research on interlimb generalisation, to our knowledge, virtually nothing is known about whether and to what extent learning generalises across limbs in statistical contexts. The aim of Experiment 2 was therefore to build on the results of Experiment 1, and test whether Bayesian sensorimotor learning generalises from one limb to the other. Relatedly, we were interested in testing whether the initial visuomotor learning occurring during training is represented in extrinsic or intrinsic coordinates.

Like Experiment 1, Experiment 2 started with 10 trials of a familiarization task in which cursor feedback was always provided and no lateral shift was imposed. After completing a training phase with their right hand (1080 trials), subjects completed a testing phase (1080 trials) using their left hand in which they experienced cursor feedback sampled from the same Gaussian distribution as experienced previously with the right hand (mean of 1cm, SD of 0.5cm) (Fig. 4 b,c). To assess the reference frame in which transfer occurs, 7 subjects experienced a congruent-extrinsic (CE) condition in which the cursor was shifted in the same visual direction across both the training phase with the right arm (Fig. 2c) and the testing phase with the left arm (Fig. 2d). By design, the imposed

visuomotor perturbation was congruent in extrinsic (screen-based) coordinates (rightward lateral shift), yet incongruent in intrinsic coordinates (requiring an elbow joint flexion in the right arm and an elbow joint extension in the left to compensate for the shift). Another 7 subjects experienced a congruent-intrinsic (CI) condition in which the cursor was shifted in opposite visual directions for each arm (rightward shift for the right arm during the training phase [Fig. 2e], and leftward shift for the left arm during the testing phase [Fig. 2f]). This time, the visuomotor perturbation imposed across both the training and testing phases was congruent in intrinsic coordinates (requiring joint flexion in both right and left arms), yet incongruent in extrinsic coordinates.

If the learning that occurs during the training phase is represented in an extrinsic reference frame, this predicts that transfer will be relatively strong and complete in the CE condition and relatively weak and incomplete in the CI condition. If the learning that occurs during the training phase is represented in an intrinsic reference frame, this predicts that transfer will be relatively strong and complete in the CI condition and relatively weak and incomplete in the CE condition. Here, we define *strong* transfer as indicated by early LH endpoints that are greater than or approximate to 50% of late RH endpoints, and *weak* transfer as endpoints that are less than 50%, according to the following generalisation equation¹³:

$$\% \text{ Generalisation} = \frac{\text{mean early LH endpoints}}{\text{mean late RH endpoints}} \times 100 \quad (5)$$

Finally, it is important to note that since we were primarily interested in assessing the degree of generalization of the learned prior, only endpoints from the (σ_∞) trials were used in the analysis

¹³ Assessing generalisation via the comparison of early LH reaches and late RH reaches is in line with the methods used by Wang and Sainburg (2005) and Brayanov et al. (2015). Another way of testing generalisation is via the comparison of early RH and LH reaches, as employed by Carroll et al. (2014), which provides a comparison of learning rate between training and testing reaches. Given that we are not explicitly investigating the rate of learning, comparing late RH with early LH reaches is preferable as it indicates the percentage generalisation after complete RH adaptation.

(as in Experiment 1). In these trials, the influence on the prior should be least contaminated by current sensory evidence (see General Methods [Section 2.2.5]).

2.4.1 Models and predictions of generalisation

Having established that subjects behave in a fashion consistent with the Bayesian estimation model in Experiment 1, we considered three variants of the Bayesian estimation model as applied to interlimb transfer.

2.4.1.1 Full-Bayesian model

According to the Full-Bayesian model, the internal representation of statistical sensorimotor learning is limb- or effector-independent. Consequently, this model predicts that the mean of the prior will transfer to the other limb, and the likelihood will be optimally integrated. For Experiment 2, the specific predictions are that left arm (LH) trials early in the testing phase should resemble right arm (RH) trials late in the training phase (once learning has stabilised) both with respect to endpoints (indicating transfer of the previously learned prior) and sensitivity to degree of visual feedback uncertainty (indicating integration of the likelihood). The Full-Bayesian model therefore predicts that learning will transfer and endpoints will accurately estimate the imposed prior, and cursor error as a function of shift (slopes) will increase proportionally to increasing visual uncertainty.

Importantly, there are two sub-varieties of this model (and also the Partial-Bayesian model) that relate to the different reference frames in which the imposed prior might be encoded: an *Extrinsic reference frame model* and an *Intrinsic reference frame model*. According to the Extrinsic model, the prior is initially learned and therefore generalised to the opposite limb in extrinsic coordinates. Because the imposed prior learned in RH training has a positive value (shifts are rightward along the x -axis), this model predicts a negatively-signed mean endpoint in early LH trials for both CE and CI conditions (adapted endpoints should be leftward along the x -axis).

According to the Intrinsic model, the prior is initially learned and therefore generalised to the opposite limb in intrinsic (joint-based) coordinates. Because the imposed prior learned in RH training has a positive value, requiring a compensatory flexion in the right limb, this model predicts an analogous flexion in the left limb resulting in a positively-signed mean endpoint in early LH trials for both CE and CI conditions (adapted endpoints should be leftward along the x -axis).

2.4.1.2 Partial-Bayesian model

According to the Partial-Bayesian model, some but not all of the internal representations of the component distributions involved in Bayesian integration are limb- or effector-specific. This model is based on recent evidence that the mean and variance of probability distributions involved in Bayesian integration may exhibit different patterns of generalization (Fernandes et al. 2014), and that prior and likelihood distributions might have quite distinct representational substrates in the human brain (Vilares et al. 2012). The Partial-Bayesian model predicts that either the mean of the prior, or its likelihood distribution will transfer, but not both. For Experiment 2, this model predicts that if the mean transfers, then early LH trials should be highly similar to late RH trials with respect to endpoints but visual uncertainty will not immediately be integrated. In contrast, if the likelihood transfers, then subjects will rapidly integrate visual uncertainty in early LH trials, but will fail to generate reach endpoints that compensate for the imposed shift.

2.4.1.3 Quarantine model

The quarantine model predicts that all internal representations learned under high levels of uncertainty will be “quarantined” in the sense that generalisation to new, untrained contexts will not occur (Fernandes et al. 2012). Although the prior distribution of experienced shifts and the likelihood distribution of visual uncertainty may be learned during the RH training phase, these representations would be restricted in their application to the specific context in which they were


learned. Consequently, they could not be recruited for the purposes of generalisation including for opposite limb reaches. According to the quarantine model, subjects will not compensate for the imposed shift with any degree of efficiency, nor will they integrate visual uncertainty. For Experiment 2, this model specifically predicts that early LH trials during the testing phase for both CE and CI groups will be highly similar to naïve performance exhibited very early in the initial RH training phase.

2.4.2 Results

In this experiment, a training phase of 1080 RH trials was followed by 1080 LH trials for both the CE and CI conditions (Fig. 4b,c). The percentage of generalisation was determined by comparing the mean endpoint during late training (trials 980-1080) against the mean endpoint from the early LH trials (1080-1180), as per equation (5). To rule out the possibility that training differences between subjects participating in Experiments 1 and 2 could influence our results, mean endpoints during late RH training (980-1080) were compared across VA, CE, and CI groups. Endpoints were highly similar: $-1.52 \pm 0.2\text{cm}$, $-1.23 \pm 0.32\text{cm}$ and $-1.34 \pm 0.22\text{cm}$ (mean \pm SD in all cases) for VA, CE and CI groups, respectively (Fig. 10a), and the observed differences were not significant (VA vs CE, $p=.069$; CE vs CI, $p=.31$; VA vs CI, $p=.15$). This result indicates that learning of the prior distribution for CE and CI subjects was comparable to learning for VA subjects.¹⁴

During early LH trials (1080-1180), a mean endpoint (σ_∞) of $-1.22 \pm 0.1\text{cm}$ (mean \pm SD) was observed for subjects in the CE group, which is highly similar to VA endpoints during the same period ($p=.094$) (Fig. 10b). This indicates strong (98%)¹⁵ transfer of the learned prior when the visual perturbation is congruent in an extrinsic reference frame.

¹⁴ In addition, the mean endpoint (σ_∞) during late training (980-1080) for VA subjects was similar to both late training (1980-2160) (see footnot figure), and the entire block of 1080 trials from the testing phase, thus reinforcing the use of late testing reaches (980-1080) as a suitable indicator of prior learning.



¹⁵ Calculated according to equation (5) as follows:

In contrast, an endpoint of $0.31 \pm 0.26\text{cm}$ (mean \pm SD) during early LH reaches was observed for CI subjects, which is significantly different to both CE and VA endpoints during the same period ($p=.0001$ in both cases) (Fig. 10b). This indicates that the learned prior incompletely generalised (23%) when the perturbation is only congruent in an intrinsic reference frame. Transfer in the CI group was relatively weak in comparison to the nearly complete transfer observed in the CE group, suggesting that the learned prior is encoded in extrinsic coordinates.




Fig 10. A comparison of endpoints for all experimental groups. (A) Mean endpoint during late RH training. (B) Mean endpoints during early LH testing.

Although the mean endpoint averaged across the first 100 trials of the testing phase may provide some indication of transfer, it may also reflect new learning with the opposite limb. Hence, an analysis of early LH reaches with higher temporal resolution was required.¹⁶

Subsequently, a moving average analysis was performed on early LH reaches (Fig. 11). For the CE group, the moving average reveals a significant degree of rapid adaptation from a mean endpoint of $-0.56 \pm 0.12\text{cm}$ (45%) in the first 5 trials and $-0.86\text{cm} \pm 0.18\text{cm}$ (68%) between 5 and

$$\frac{\text{mean LH CE reaches (1080 - 1180)}}{\text{mean RH CE reaches (980 - 180)}} = \frac{-1.21}{-1.23} \times 100 = 45\%$$

The same calculation is performed to calculate the percentage for early CI reaches as well (not shown here).

¹⁶ Ideally, interlimb generalisation would be measured by running a block of probe trials with the left hand, immediately following right hand training, in which no visual feedback is provided (σ_{∞} condition). Due to differences in our paradigm, especially those stemming from our aim of replicating the findings of Körding and Wolpert (2004), such probe trials were not included. Probe trials will be included in future studies.

10 trials (providing an average of 0.7cm [56%] over the first 10 trials), before stabilising from 25 trials onward toward a mean of $-1.27 \pm 0.13\text{cm}$ (103%) after 100 trials (Fig. 11a).

Mean endpoints in CI significantly and rapidly adapt to a mean imposed shift of -1cm (left, in extrinsic, screen-based coordinates); from $-0.6 \pm 0.13\text{cm}$ (45%) in the first 5 trials and $-0.75 \pm 0.16\text{cm}$ (55%) between 5 and 10 trials (providing an average of 0.67cm [50%] over the first 10 trials). However, Despite this rapid adaptation, endpoints plateau toward a mean of $0.41 \pm 0.06\text{cm}$ (30%) after 100 trials (Fig. 11b). Interestingly, there was no significant difference between mean endpoints for both CE and CI groups after 5 trials ($p=.65$) or after 10 trials ($p=.45$). A significant difference between CE and CI groups only began to emerge after 15 trials ($p=.002$), a difference which was maintained over the remaining trials.




Fig 11. Moving average plot for early LH trials. (a) A moving average of endpoints across the first 100 trials in the LH testing block for CE (b) and for CI (b). Each bar represents the average across 7 subjects using a window

size of 5 trials.¹⁷ The dashed line represents the mean across the 100 reaches. Note that all values represent endpoints in horizontal screen-based coordinates.

Following an investigation of early LH reaches, remaining reaches over the course of the LH testing phase were assessed via a moving analysis of 50 trial windows (Fig. 12). For later trials (1180-2160), mean endpoints for the CE group remain constant, settling at $-1.28 \pm 0.05\text{cm}$ in the last 80 trials of the testing phase, representing a mean of -1.22cm over the entire 1080 trial testing block (Fig. 12 in green). In contrast, mean endpoints for the CI group did not recover toward a mean of 1cm , but rather, settled on a diminished value of $0.42 \pm 0.04\text{cm}$ (31% compared to late RH CI) in the last 80 trials of the testing phase, representing a mean of 0.51cm over the entire 1080 trial testing block (Fig. 12 in orange).¹⁸

¹⁷ Although each block of 5 trials constitutes a mean over a theoretical maximum of 35 endpoints (5 trials x 7 subjects), each block actually represents a mean of less trials due to the frequency of σ_∞ trials represented (average probability of $\frac{1}{6}$ from four possible trial types with a frequency of 3,1,1,1). From analysis, the range of trials was determined to be 6-8 per block across the data set.

¹⁸ It remains possible that the observed pattern of over-correction in the CE group and incomplete correction for the CI group could reflect a generalised leftward bias in aiming that applies across both arms. Such a bias may be the result of an implicit consequence of dominant versus non-dominant limb dynamics, or as a result of recruiting an explicit strategy learned during the training phase. To reduce the influence of such effects, in future experiments we can introduce more task complexity including counter-balancing the direction of the cursor shifts (within each limb), varying the target direction, and/or using more than one reach target.




Fig 12. Moving average plot for the testing phase (trials 1080-2160). Same conventions as in Fig. 11, except that window size is 50 trials. Mean values over the entire 1080 testing block are represented by dashed horizontal lines for the respective groups. Also, note that the last bar constitutes the mean over 80 rather than 50 trials).

It is important to note that a repeated measures analysis of variance across the averaged windows was not possible for the data displayed in Figures 11 and 12. As both the number of trials and the subjects represented differ across windows, neither trial-matching, nor repeated measure matching was possible.¹⁹ Accordingly, the significance values reported in Figures 11 and 12 correspond to a pairwise analysis of the group mean and standard deviation between consecutive windows. Although not a comprehensive statistical analysis of the data, such pairwise analysis confirms the apparent difference between the CE and CI data displayed in the figures, which, given the consistency evident in the data, does not detract from the overarching trend that is evident.

¹⁹Although each window constitutes theoretical maximum of 350 endpoints (50 trials x 7 subjects), the average number of trials per window was determined to be 56 ± 9 across the data set. This variation was a result of the way in which the visual feedback condition was randomly allocated for a given trial. Subsequently, the number of trials per subject varied across windows, and hence, a pairwise comparison of means across windows was not possible.

Nonetheless, future studies will include trial-matched conditions across subjects in order to provide a more robust statistical analysis of the data.

Next we assessed the slopes of cursor error as a function of the imposed shift across the different visual uncertainty conditions. Slopes increased significantly as a function of visual uncertainty for all experimental groups during late RH training trials (980-1080), indicating that CE and CI subjects were comparable to VA subjects with respect to integrating visual uncertainty at the end of the RH training phase (Fig. 13).




Fig. 13. Similar integration of visual uncertainty across subjects in Experiments 1 (VA) and 2 (CE and CI).

Figure 14 indicates cursor error as a function of the imposed shift for both experimental groups, averaged over the entire LH testing phase (panels a,g), and averaged across 100 trials at different junctures throughout the testing phase (panels b-l). There is no readily discernible trend across CE reaches (Fig 14a, $p = .06$), however, there is some evidence of a trend of increasing slope as a function of visual uncertainty on cursor error across the CI testing phase (Fig. 14g, $p = .01$). Although there are significant differences across visual conditions for early reach trials during the 1080-1180 block in the CE group (Fig. 14b, $p = .01$), no consistent pattern is evident. The same is also true for CE subjects in the 1280-1380 block (Fig. 14c, $p = .01$). In contrast, a significant trend of increasing slopes emerges in the CI group during the 1280-1380 block (Fig. 14i; $p = .001$). The

same trend of increasing slopes as a function of visual uncertainty continues for CI subjects during the late 1780-1880 block (Fig.14k, $p=.025$), but then largely disappears during the subsequent 1960-2160 block (Fig, 14l; $p=0.2$).




Fig 14. Mean slopes (cursor error as a function of shift) for CE and CI groups. Plots for the entire testing phase (top row) and four 100-trial blocks extracted from different parts of the LH testing phase.

2.4.3 Discussion

2.4.3.1 Generalisation of the learned prior

In this experiment, we investigated the interlimb generalisation of Bayesian sensorimotor learning by having subjects learn to perform reaches with one limb in a task involving a probability distribution of visuomotor shifts (with a fixed mean and SD) before testing the degree to which this learning generalises to the opposite limb. This paradigm allows us to assess how well subjects adapt their movements to the average experienced visuomotor perturbation (mean of the prior), the uncertainty in the visual feedback (visual uncertainty or likelihood), and most importantly, the reference frame in which learning generalises to the opposing limb.

The moving average analysis of early LH reach endpoints in early CE trials revealed a significant degree of rapid adaptation from a partially generalised mean endpoint of -0.7cm [56%] over the first 10 trials to -1.29 ± 0.14 cm after 25 trials before stabilising around a mean of -1.4 ± 0.09 cm after 100 trials (Fig. 11a). Similarly, LH reach endpoints in early CI trials significantly and rapidly adapt from a mean endpoint of -0.67cm [50%] over the first 10 trials to 0.6 ± 0.07 cm after 15 trials, before reaching a plateau toward a mean of 0.41 ± 0.06 cm after 100 trials (Fig. 11b). The negatively-signed mean endpoints in early LH trials across both CE and CI conditions match the predictions of the *Extrinsic reference frame* model. The observed generalisation pattern across both early CE and CI trials therefore indicates that the prior learned during RH training is represented in a limb-general format available to both limbs, and relatedly, is encoded in an extrinsic reference frame. At a minimum, this disqualifies the *quarantine* model.

We then considered our results in Bayesian terms. Firstly, as a simplifying assumption, we considered the case of 100% generalisation of the prior learned during RH adaptation.²⁰ Specifically, we modelled the result of Bayesian estimation during early LH reaches for both the

²⁰ This is not to be confused with complete generalisation of the *true* prior of +1cm. Instead, according to equation (5), 100% generalisation implies that the prior learned during RH adaptation generalises completely to early LH reaches.

CE and CI conditions assuming a mean of the learned prior of 1.23cm^{21} (SD = 0.5cm) for CE (Fig 10a), and 1.34cm (SD = 0.5cm) for CI (Fig10a), respectively; values which were obtained directly from our data. These priors were integrated with a likelihood reflecting an imposed mean visual shift of $+1\text{cm}$ and a standard deviation of 0.7cm , the latter of which represents the average variance across all experienced visual uncertainty conditions²² ($\sigma_0, \sigma_M, \sigma_M, \sigma_\infty$) (Fig. 15).




Fig 15. Model of Bayesian estimation for early CE and CI reaches assuming 100% generalisation of an extrinsically encoded prior.

A posterior estimate²³ of 1.15cm was calculated for the simulated early CE reach trials, and 0.55cm for the early CI trials, representing endpoints of -1.15cm and -0.55cm , respectively. While the

²¹ Priors and likelihood distributions are described in terms of extrinsic, workspace-based coordinates, where positive values indicate right of the workspace midline and negative values indicate left of midline.

²² In Matlab, the likelihood was generated as follows (shown for CI):

```
LIKELIHOOD = makedist('Normal','mu',-1,'sigma',0.7);
LIKELIHOOD = pdf(LIKELIHOOD,X);
```

²³ In Matlab, the Posterior estimate was generated as follows:

```
POST = (LIKELIHOOD.*PRIOR)/max(LIKELIHOOD);
```

Where $\max(\text{LIKELIHOOD})$ is the probability of the imposed shift $P(D)$. And the prior, generated for the CE example is as follows:

```
PRIOR = makedist('Normal','mu',1.23,'sigma',0.5);
PRIOR = pdf(LIKE,X);
```

model approximates the average endpoint of -0.67cm observed over the first 10 CI reaches²⁴, it does not accurately predict the average endpoint of -0.70cm observed over the first 10 CE reaches. The magnitude and sign of the discrepancy between the modelled endpoint and that observed during early CE reaches suggests that the learned prior is not generalised completely from RH adaptation to LH reaches in the CE condition, thus strengthening the hypothesis of partial generalisation across extrinsic coordinates as observed previously (Imamizu and Shimojo 1995; Sainburg and Wang 2002; Wang and Sainburg 2004, 2005; Taylor et al. 2011).

In order to test this hypothesis more rigorously, we ran the same Bayesian model, this time assuming partial generalisation of the prior (as observed experimentally). Specifically, we assumed a prior of 0.7cm (SD = 0.5cm) was integrated in CE reaches and a prior of 0.67cm (SD = 0.5cm) was integrated in CI reaches²⁵, where the same likelihood reflecting an imposed mean visual shift of +1cm and a standard deviation of 0.7cm was applied (Fig. 16).

²⁴ Note that the posterior estimate has a lower probability than either the prior or the likelihood of the visual feedback, which, although a product of Bayesian estimation, is far from Bayes-optimal in this case. Yet such a sub-optimal estimate is exactly what is expected if a statistical prior, encoded in extrinsic coordinates is erroneously applied to an incongruent shift defined in extrinsic coordinates (as in the CI condition). As an interesting consequence, the detection of sub-optimal Bayesian estimation may provide a means by which to categorise the degree of congruence across task domains that are otherwise indistinguishable, whereby a posterior estimate that is less probable than the prior may be used to identify the contextual bounds of a given task domain.

²⁵ A prior -0.70cm and -0.67cm represents a partial generalisation of 57% and 50% for CE and CI respectively, obtained directly from our data (based on late RH [980-1080] endpoints).




Fig 16 Models of Bayesian estimation for early CE and CI reaches assuming partial generalisation of an extrinsically encoded prior.

This time, assuming partial generalisation, a posterior estimate of 0.75cm was calculated for the simulated CE reaches, and 0.12cm for early CI reaches, representing endpoints of -0.75cm and -0.12cm respectively. In contrast to the 100% generalisation model, this model of partial generalisation provides a closer approximation of the observed mean endpoint across the first 10 CE trials = -0.70cm, but the same is not true for CI endpoints (mean endpoint across the first 10 CI trials = -0.67cm).

Consequently, while the model of 100% generalisation accurately predicts early CI endpoints, it does not predict early CE endpoints, thus indicating partial generalisation of the learned prior. However, when assuming partial generalisation, the model accurately predicts early CE but not CI endpoints. These opposed findings are puzzling, indicating the possibility of distinct processes of adaptation across the CE and CI conditions, the interaction of other unmodelled factors, or both. A more refined experimental paradigm, including a richer set of target directions, might help to resolve this puzzle and shall be considered in future work.

This issue notwithstanding, both early CE and CI endpoints indicate that the prior learned during RH adaptation is encoded in an extrinsic reference frame in a limb-general manner, and is

therefore at least partially available to influence learning with the opposite limb. These reference frame findings are striking and closely resemble a number of earlier results indicating intralimb learning and generalisation of visuomotor rotations in extrinsic coordinates (Flanagan and Rao 1995; Wolpert 1995; Vetter et al. 1999; Krakauer et al. 2000) and interlimb generalisation in extrinsic coordinates (Imamizu and Shimojo 1995; Sainburg and Wang 2002; Wang and Sainburg 2004, 2005; Taylor et al. 2011). Nevertheless, they differ from more recent findings. For instance, Carroll et al. (2014) reported strong and immediate interlimb transfer only when a static visual perturbation was congruent across both intrinsic and extrinsic reference frames. Transfer was highly limited when the visuomotor perturbation was only congruent in single reference frame (19% during an extrinsic-congruent condition and 8% during an intrinsic-congruent condition). Clearly, it is their findings about lack of transfer in an extrinsic reference frame congruent condition that is most relevant for present purposes.

One important difference between our paradigm and the one employed by Carroll et al. that could potentially explain the divergent results is that theirs was an isometric visuomotor rotation task in which forces or torques produced by a static hand are mapped into cursor movement. Isometric movements involve muscle contraction without corresponding changes in joint angle and muscle length, and are known to differ in terms of the muscle activity required from that used in natural multi-joint movements (Sergio et al. 2005). It is therefore possible that learning a visuomotor rotation in an isometric task is more closely linked to learning arm dynamics than learning arm kinematics, and involves at least partial coding in intrinsic coordinates (Shadmehr and Mussa-Ivaldi 1994).

Using a similar isometric visuomotor task, Rotella et al. (2015) provide evidence that the learned visuomotor rotation transferred predominantly in intrinsic (joint-based) coordinates. Given this, the observation by Carroll et al. (2014) of minimal transfer when the visuomotor rotation was only congruent across an extrinsic reference frame may be less surprising. Accordingly, transfer in

an isometric visuomotor task would be expected only if there was congruence in intrinsic coordinates. However, since the transfer reported by Rotella et al. (2015) was incomplete, they could not rule out the possibility that the visuomotor learning was encoded by a mixture of intrinsic and extrinsic reference frames (Brayanov et al. 2012).

More recently, Poh et al. (2016) reinforce the possibility of a mixed reference frame representation underlying visuomotor learning using a more standard (non-isometric) reaching task. By implementing an innovative variation on the standard visuomotor rotation paradigm, Poh et al. (2016) were able to concurrently investigate both the evolving degree of adaptation in the trained limb and transfer to the untrained limb. They report stronger transfer when the visuomotor rotation is congruent in both extrinsic and intrinsic coordinates as compared to when the perturbation is aligned in only a single reference frame. They claim that these results implicate a role for mixed reference frames in the interlimb transfer of visuomotor learning. Although these results appear at odds with our own, their focus seems to be on the degree of transfer relatively late in the testing phase with the untrained limb (specifically the last two blocks of “probe” trials; see their Figure 4e). By contrast, our primary focus was on early transfer. Although paradigm differences make a direct comparison difficult, the pattern of transfer Poh et al. (2016) observed in early blocks of probe trials appears less consistent with their mixed reference frame conclusion than the transfer pattern observed in late blocks. Further investigation is therefore required. Specifically, ruling out the possibility of a combined intrinsic-extrinsic representation (Brayanov et al. 2012) underlying learning and transfer in our paradigm will require a further experimental condition in which the mapping learned across the limbs is aligned in both reference frames (see Section 3.5).

In addition to the relatively strong generalisation of an extrinsically encoded prior, another phenomenon of interest in our results is evident during early LH reaches pertaining to the rate of subsequent adaptation. While early CE reaches adapt as expected toward a mean endpoint of -1cm,

the pattern of adaptation during early CI reaches is unexpected. Contrary to a gradual shift from erroneous endpoints toward a mean of +1cm (as predicted by an automatically applied prior), an abrupt error-corrective switch is observed after 15 reaches (Fig. 11b). One plausible explanation for this rapid shift is that some explicit motor learning strategy is invoked (Taylor and Ivry, 2013). According to Taylor and Ivry (2013), generalisation between effectors and across workspaces may involve both *implicit* and *explicit* learning processes. In fact, Poh et al. (2016) investigated the contribution of explicit processes to transfer of visuomotor learning and found that explicit learning is typically encoded in extrinsic coordinates and is fully available to the opposite limb. With these distinct learning processes in mind, one explanation may be that after implicitly generalising the extrinsically coded prior (reflected in endpoints for initial CI trials), subjects soon explicitly recognize a change in task context and thus rapidly adopt a new explicit strategy (e.g., reach to the right of the target). Although an interesting source of speculation, our paradigm was not designed to address this issue.

Another especially interesting finding is that while CE reaches stabilise toward a mean endpoint of -1.22cm over the course of LH reaches, CI reaches plateau on a diminished mean value of 0.51cm over the course of the LH reaches. This finding indicates a clear distinction between the CE and CI conditions in terms of the extent of subsequent adaptation over the course the LH testing block. One possible explanation for this CI-specific effect might be anterograde interference²⁶ in the extent of subsequent adaptation. For example, interference has been demonstrated if, at some time after initial visuomotor learning, a counter-rotation equal in magnitude but opposite in direction is learned (Wigmore et al. 2002; Krakauer et al. 2005). Given our results indicating that the statistical visuomotor perturbation is learned and generalised to the other limb in extrinsic-

²⁶ In the context of sensorimotor adaptation, anterograde interference is defined as a negative change in performance of task B in virtue of adaptation during task A, which may constitute an increase in early endpoint error, an increase in the rate, and/or a decrease in the overall extent of subsequent adaptation (Shadmehr and Brashers-Krug 1997; Miall et al. 2004; Krakauer et al. 2009).

coordinates, it is possible that the transferred prior, erroneously applied to visual shifts that are opposite in extrinsic coordinates during the CI condition gives rise to interference.

According to Krakauer (2009) motor memory, undergoes a process of consolidation, whereby a newly acquired internal model becomes increasingly resistant to modification by a competing model. Not only does this consolidation commence rapidly (Krakauer et al. [2005] report significant consolidation after 5 minutes), but it appears to strengthen as a function of time and is strongly correlated with number of adaptation trials performed (Krakauer et al. 2005). When investigating visuomotor adaptation in non-human primates, Yin and Kitazawa (2001) demonstrated that although a similar extent of adaptation occurred between testing blocks of 250 versus 500 trials, the duration of aftereffects resulting from the initial adaptation varied. Monkeys trained on 250 trials showed evidence of aftereffects up to 24 hours later, whereas animals trained on 500 trials showed aftereffects up to 3 days later. Interestingly, nothing further was learned on these additional exposure trials. According to Yin and Kitazawa (2001), “the error during the exposure period had already decreased to nil asymptote by the 250th trial” (2001, p. 253). Nevertheless, it seems that the “additional repetition of reaching with approximately zero errors that was crucial for triggering the [more robust] consolidation” (ibid, p. 253). Given that our subjects adapted to the extrinsically encoded perturbation over a large number of trials (1080), the process of consolidation may render the prior robust to change which thus slowly – and incompletely – updates toward the mean of -1cm. Future experiments involving the systematic manipulation of the number of exposure trials are needed to test the consolidation theory.

Another possible explanation for the observed plateau (lower asymptote) for CI reaches is that the recruitment of an explicit strategy may override implicit learning (Taylor and Ivry, 2013; Taylor et al. 2014) and lead subjects to persevere on left-shifted reaches with the left limb. For example, if subjects come to adopt a simple high-level cognitive strategy such as “land to the left of the target” during the training phase with the right limb, this explicit strategy might carry over

to the testing phase and result in incomplete correction over the course of CI reaches with the left limb. Additionally, the large cursor-target size tolerance (footnote 12) might also allow a leftward (or rightward) bias to persist for longer than if the tolerance were smaller and more frequent error feedback was given.

In summary, experiment 2 provides compelling evidence of three distinct phenomena including; the significant generalisation of the learned prior across extrinsic reference frames, rapid, contextual generalisation of the prior across intrinsic reference frames, and the presence of a consolidated prior. Furthermore, while both rapid contextualisation and consolidation strengthens the hypothesis of generalisation of the prior across extrinsic reference frames, they are mutually exclusive as explanations of their target phenomena, thus requiring additional investigation in order to disentangle what is actually occurring over the course of CI reaches.

However, despite such puzzling findings, what is clearly evident according to early CE and CI endpoints is that the prior learned during RH adaptation is encoded in extrinsic coordinates and generalized across limbs. This finding rules out the proposed *quarantine* model. However, these results alone do not allow us to arbitrate between the *partial-Bayesian model* and *full-Bayesian model*. Next, we discuss the degree to which information about the experienced distribution of visual uncertainty (the likelihood) is generalised during early LH reaches.

2.4.3.2 Integrating visual uncertainty

Interestingly, our results indicate that visual uncertainty is not being integrated by the CE group at any time throughout the LH testing phase (Fig 14a-f). While the average slopes across the entire testing phase are large for all visual conditions (Fig. 14a), there is no discernible pattern over the course of reaches, culminating in a mixture of small and negative slopes during the last 100 reaches. Similarly, early CI reaches indicate that visual uncertainty is not being integrated. Taken together with the previously discussed results, this supports the *partial-Bayesian model* of transfer. These

findings are broadly consistent with recent work suggesting separate representations of the prior and likelihood in the human brain (Vilares et al. 2012), as well as evidence that the “components” (mean and variance) of a prior distribution involved in Bayesian computation might be represented and generalized independently (Fernandez et al. 2012).

Although our results indicate that the likelihood of visual uncertainty is not optimally integrated with the untrained limb it is expected that over time CE subjects recommence integrating uncertainty. This is not the case, however. Instead, CE subjects appear to cease integrating visual uncertainty over the entire LH testing phase (Fig. 14a-f). Interestingly, CI subjects do recommence the integration of visual uncertainty over the course of the LH testing phase (Fig. 14g-l). One possible explanation is that differing attentional demands across the CE and CI conditions may be partially responsible. The incongruence of the imposed visual perturbation across extrinsic reference frames encountered in the CI condition may provide enough of a discrepancy between task contexts to require sustained attention toward the degree of visual uncertainty present in the midpoint feedback during individual trials. Further experiments are needed to test this hypothesis.

Another possibility is that a limb-dominance effect contributes to the lack of likelihood integration during reaches with the non-dominant limb, which in turn gives rise to differences in the results observed between the CE and CI conditions. For instance, according to Berniker and Körding (2008), motor errors experienced during training of the nondominant limb are more frequently attributed to a misestimate of limb dynamics as opposed to uncertainty in the applied perturbation. At this stage, we do not know whether optimal integration of likelihood information occurs when subjects are exposed to a stochastic visuomotor rotation for the first time with the non-dominant limb (i.e., without prior training experience with the dominant limb). Future experiments testing non-dominant to dominant transfer are therefore important.

2.5 Experiment 3: Temporal constraints on Bayesian sensorimotor learning

The aim of Experiment 3 was to investigate the influence of timing on the stability of the learning effects observed in Experiment 1 (VA). Specifically, we aimed to test whether an inter-trial delay was required for subjects to show evidence of Bayesian integration in sensorimotor learning. Based on previous studies of motor learning (e.g., Krakauer 2009), it seemed plausible to assume that there might be some temporal constraints on motor memory consolidation while learning a new task, and consequently, that certain timing changes to the task (or interval between task trials) might therefore influence their ability to learn. Accordingly, a new cohort of subjects ($n = 7$) completed an identical task to that of the Experiment 1, but with zero (or minimal) delay in between subsequent trials (the *truncated zero-delay* condition).²⁷ Instead of providing a delay interval between trials matched to the main experiment (150 ms), visual feedback was maintained for the same 100 ms after movement offset in the zero uncertainty (σ_0) condition as in the main experiment, but, this time, with no delay before the start target reappeared, and only a minimal delay (corresponding to the amount of time required for the hand to return to the start target), before a new trial commenced (Fig. 17).

We also wanted to determine whether it was the specific presence of an inter-trial delay or simply longer feedback processing time at the end of the trial that was required for the observed learning. To address this, we tested another cohort of subject ($n = 7$) using a different timing condition. As in the truncated zero-delay condition, these subjects were given no time in between subsequent trials (Fig. 17). However, instead of simply truncating all trials so that they ended with the cessation of visual feedback 100 ms after movement offset, visual feedback was provided for a full 250 ms at the end of zero uncertainty (σ_0) reaches followed by no delay before the start target

²⁷ Importantly, the name of this condition reflects the absence of a delay interval between trial events – the offset of visual feedback at the end of the previous trial and the appearance of the start target indicating the start of the next trial. Unavoidably, there was a minimal delay corresponding to the time required for subjects to bring their hands back to the start target.

reappeared and a new trial commenced (*long-feedback zero-delay* condition). For comparison purposes, the 250 ms feedback duration matched the total duration of the 100 ms of visual feedback + 150 ms inter-trial delay that subjects experienced in Experiment

	<i>Feedback Duration</i>	<i>Inter-trial Delay</i>	<i>TOTAL</i>
VA	100ms	150ms	250ms
Truncated	100ms	0ms	100ms
250ms	250ms	0ms	250ms

Fig. 17 The differences in visual feedback and trial delay periods across the experimental conditions. For comparison purposes, visual feedback duration and inter-trial interval are provided for Experiment 1 (VA).

2.5.1 Predictions

If an inter-trial delay is required for Bayesian sensorimotor learning, subjects experiencing either the *truncated zero-delay* condition or the *long-feedback zero-delay* condition will show no evidence of Bayesian integration in sensorimotor learning (when compared against subject data collected in Experiment 1). For example, the lack of delay between trials – no matter how long the feedback was delivered – could interfere with learning and consolidation of the imposed prior. If Bayesian sensorimotor learning simply requires more processing time than 100 ms after movement offset (even in the absence of an inter-trial interval), adaptation comparable to that exhibited by subjects in Experiment 1 should occur in the *long-feedback zero-delay* condition but not in the *truncated zero-delay* condition.

2.5.2 Results

Across the testing phase (1080-2160), a mean endpoint of $-1.57 \pm 0.38\text{cm}$ for the truncated group is significantly similar ($p=.28$) to a mean of $-1.37 \pm 0.28\text{cm}$ for the 250ms group, while both truncated and 250ms endpoints are significantly similar to their VA counterparts (Fig. 18a). These results indicate that the mean of the imposed prior distribution of shifts was equivalently learned

across both groups, suggesting that, at least within the tested ranges, inter-trial intervals and specific feedback durations are not required for Bayesian sensorimotor learning.




Fig. 18 (A). A comparison of endpoints across the delay conditions in experiment 3, indicating that the mean of the prior is learned by all groups (all endpoint are significantly similar). (B & C) Cursor error as a function of shift during early test reaches. While there is a significant trend toward increasing slopes for the 250ms group (C), no such trend exists for the Truncated group (B).

Although significantly different across visual conditions, cursor error as a function of visual uncertainty for the truncated condition does not show a consistent increase as is expected for subjects performing Bayesian integration ($F_{3,81} = 3.26, p=.046$; Fig. 18b). In contrast, there is a significant trend of increasing slopes across visual uncertainty conditions for the 250ms group ($F_{3,81} = 24.6, p=.001$; Fig. 18c).

2.5.3 Discussion

Interestingly, while both the truncated and long-feedback groups learn the mean of the prior, only the long-feedback group showed evidence of integrating visual uncertainty (comparable to subject data in Experiment 1). The lack of uncertainty integration in the truncated group is puzzling and difficult to explain. Initially, we expected that the absence of a delay period between the end of the current trial and the start of the next would hinder the subject's ability to integrate feedback about cursor error at movement offset and consequently learn to adapt their movements to the visual shift.

In both zero-delay conditions, endpoint feedback (on unlimited uncertainty trials) and the onset of the start target to initiate the next trial occur very close together in time. Consequently, this could have resulted in a kind of *dual-task interference* (Pashler, 1994; Frensch 1998; Schmidt 1988; Taylor and Thoroughman, 2007), which is known to occur when subjects are required to perform two concurrent attention-demanding tasks. Moreover, Taylor and Thoroughman (2007) specifically argue that such interference can sometimes occur in motor learning, resulting in the disruption of “proper encoding and transformation of previously experienced errors into changes in predictive control” (ibid p. 325), which thereby limits adaptation in motor learning. Strikingly, our results show that while adaptation to the mean of the prior is not affected in either of the zero-delay conditions, the capacity to integrate visual uncertainty about midpoint feedback appears to be disrupted. This finding is particularly puzzling, given the fact that the mid-point feedback conditions were analogous between the truncated and long-feedback groups. One plausible explanation is that, due to a lack of delay-period, subjects performing the truncated task became overly fatigued as a result of repetition, in the absence of the minimal delay period afforded to the long-feedback group, who, once integrated endpoint feedback availed themselves of the additional 150ms to rest between trials. If this interpretation is on track, it is plausible to propose that a visual feedback period of $\geq 100\text{ms}$ is required in order to adapt to endpoint error. Another plausible explanation for the observed differences between delay conditions is that two different processes (operating over two different timescales) are in play (e.g., Smith et al. 2006). For instance, the brain may only update its estimate of visual uncertainty given an extended period of visual feedback, while the estimate of the mean shift (the peak of the prior distribution) may be updated relatively rapidly via a different process such as a high-level explicit strategy (Taylor et al. 2014). In this case, both estimates would be updated during the long- feedback trials (hence both uncertainty and the mean shift would be integrated), but only the mean shift would be updated during truncated

trials. Clearly, more experimental work is needed to disentangle the time-dependence of these observed effects.

Reassuringly, we observed highly similar results between subject performance in the long-feedback condition compared to Experiment 1. This implies that additional visual feedback (>150ms) neither inhibits nor aids adaptation, and the capacity to integrate visual uncertainty is not disrupted by the extended window of visual feedback. More experimental studies are needed to investigate the precise temporal switch point that determines whether Bayesian estimation is strengthened or inhibited.

3. General Discussion

In stark contrast to the moderate claims of Bayesian modelling approaches, which employ the mathematical framework of Bayesian statistics to model psychological and neural phenomena, the Bayesian coding hypothesis (*henceforth BCH*) makes stronger claims about the brain. Specifically, the BCH proposes that cognitive, perceptual, and/or motor processes are Bayes-optimal, and that the brain encodes probability distributions and performs probabilistic inference (Knill and Pouget 2004; Körding 2014). According to the proponents of the BCH there exists “strong behavioural and physiological evidence that the brain both represents probability distributions and performs probabilistic inference” (Pouget et al. 2013, p. 1).

On the back of their behavioural findings, which have inspired much of the experimental work reported in this thesis, Körding and Wolpert (2004) claim to have shown that “subjects internally represent both the statistical distribution of the task and their sensory uncertainty, combining them in a manner consistent with a performance-optimizing Bayesian process.” Thus concluding that “[t]he central nervous system therefore employs probabilistic models during sensorimotor learning.” (2004, 1). Clearly, drawing such a strong conclusion about Bayesian inference being implemented in the brain on the basis of behavioural evidence alone is at very least

premature. While such behavioural evidence does serve to indicate that the brain integrates sensory information in a Bayes-optimal way, it does not, on its own, provide conclusive evidence for BCH. Substantiating BCH clearly requires information about how Bayesian computations are neurally implemented.

In an attempt to provide direct neural evidence in support of BCH, several neuroimaging studies have identified brain regions specifically associated with the integration of prior learning and sensory likelihood during sensorimotor (Vilares et al. 2012; d’Acremont et al. 2013). For instance, Vilares et al. (2012) employed a task in which they could independently vary prior and likelihood and found that were differentially represented in the human brain. They report selective activation of early areas in the visuomotor processing pathway when likelihood uncertainty was manipulated, and activation of a diverse range of areas outside the visuomotor pathway when the standard deviation of learned priors was manipulated including amygdala, insula, and orbitofrontal cortex. Although neuroimaging localization studies like these begin to tell us about where in the brain Bayesian integration is performed, they provide scant information about how neural circuits actually implement the required computations. Given the well-known spatial and temporal resolution limitations inherent to BOLD fMRI (Logothetis 2008), neuroimaging studies provide inconclusive evidence about how Bayesian computations are performed by the brain. In light of the lack of direct neural evidence, models of computational models have also been relied upon to provide useful constraints. Two models have been particularly influential for theorising about the neural basis of Bayesian integration.

3.1 The Probabilistic Population Code (PPC) model

In a prominent modelling study, Ma et al., (2006) proposed their “probabilistic population code” (PPC) model (Fig 19). The PPC model involves two basic assumptions. First, Bayesian integration is performed by relatively small neural populations rather than in individual neurons or entire brain

regions. Second, the firing rates of individual neurons in the relevant population must be highly variable to the extent that they approximately obey Poisson statistics. In extreme cases where variability is essentially random and Poisson-like, the mean response of a neuron for a given condition might be equal to or even exceed its variance (Fano factor ≥ 1). Ma et al.'s critical insight is that this variability is not a nuisance or unwanted noise, but rather that neural populations automatically encode probability distributions in virtue of this Poisson-like variability. More specifically, because of the variability in individual neuron responses to some stimulus, s , the overall response of the population made up of these neurons, \mathbf{r} , to s is best described in terms of a probability distribution, $p(\mathbf{r}|s)$, rather than a deterministic mapping from s onto a single value of \mathbf{r} . Importantly, $p(\mathbf{r}|s)$ may be thought of as the likelihood distribution from Bayes' rule. With information about the prior distribution $p(s)$, Bayes' rule can be used to recover information about the probability of the stimulus given the population response, $p(s|\mathbf{r})$, the posterior distribution.

Ma et al. also assume that each distribution is represented by the activity of a distinct neural population (Fig. 19). According to their model, the mean and variance of each distribution are encoded by population activity in such a way that a population with a higher mean will naturally have lower variance (lower uncertainty) and vice versa. Accordingly, summing two population responses (representing an individual prior and individual likelihood) in a new population response (representing the posterior) will be skewed towards the population response with the larger mean (lower variance), just as one would expect for Bayesian integration. In the toy example depicted in Fig. 19, the population response representing the likelihood has lower variance (uncertainty), so would have a proportionally larger influence on the downstream population response representing the posterior.




Figure 19. Probabilistic population model of Bayesian integration. Adapted from Ma and Pouget (2008).

Despite its aspirations as a model of neural implementation, it remains to be seen whether real neural systems perform Bayesian computations in the same way. Currently, there is no way to map key aspects of the PPC model onto real neural structures and activities. Obviously, the ultimate arbiter of the success of the PPC model will be the extent to which it is consistent with neural data. Unfortunately, at the moment, there is a serious lack of data to test the model.

3.2 The Neural Weighting (NW) model

In a set of recent empirical studies, Fetsch et al. (2009, 2012) investigated the neural correlates of Bayesian multisensory integration.²⁸ They recorded single-unit activity while monkeys performed a heading discrimination task in which the reliability of the visual motion information was

²⁸ In this case, Bayesian inference occurs in the absence of a prior. Instead, it is the likelihood of multiple sensory cues that are summed together instead of summing a single cue with an existing prior (as per our experiments). Nonetheless, the process is analogous in both cases.

manipulated by varying the percentage of dots in the stimulus moving coherently in a single direction. During the experiment, monkeys were presented with either a single-cue (visual or vestibular) indicating motion, or a combined-cue (visual plus vestibular) where the two cues were in conflict with one another. Behavioural thresholds from the single-cue conditions were used to estimate cue reliability (the inverse of uncertainty or variance), and the weighting that an ideal observer should apply to each cue. Psychometric data collected during cue conflict trials were then compared to theoretical estimates predicted by Bayes' rule. Monkey's choices were significantly biased towards the more reliable cue, indicating that monkeys are integrating sensory information according to its variance, in a Bayesian manner.

To explore the neural basis of these behaviours, Fetsch et al. recorded single unit activity in dorsal medial superior temporal (MSTd) – an area involved in self-motion perception – while monkeys performed the heading task described above. When they modelled the firing rates, they found that many MSTd neurons encoded information about cue reliability (sensory uncertainty). Specifically, the modelled weights on individual MSTd neurons varied with cue reliability on a trial-by-trial basis in a manner consistent with optimal Bayesian integration.

3.3 A Fundamental Tension Between the Two Models

There is a fundamental tension between the two neural implementation models. According to the PPC, Bayesian integration is implemented at the population level, implying that cue reliability is reflected in the population activity. This means that each neuron within the population has the same fixed reliability-based weight = 1 (for instance, the visual weight of each neuron in the visual population is equal to 1 and the weight of each neuron in the population representing the prior is equal to 1).

This assumption is in tension with the NW model proposed by Fetsch et al. in which reliability-dependent reweighting occurs at the level of individual neurons and hence are not fixed

and uniform across a given population. Subsequently, the principles laid out by the PPC and the NW models imply vastly different – even incompatible – ways of implementing Bayesian integration. Needless to say, exactly how the brain might implement Bayesian computations remains an open question that will continue to occupy neuroscientists and computational modellers well into the future.

An important consequence of these different assumptions is that they both embody rather different views about what we should be looking for in the brain. The PPC model predicts that there are specialised brain regions that separately encode and process priors versus sensory uncertainty (likelihoods). By contrast, the NW model predicts that information about both prior and uncertainty is localised with individual neurons primarily involved in coding task-relevant sensory or motor variables.

3.4 What kind of constraints can behavioural evidence provide?

Given the paucity of solid neural evidence supporting BCH, there is good reason to look elsewhere to find potential support. Although less direct than neural evidence, carefully designed behavioural experiments can also provide informative constraints on the representations involved in Bayesian estimation (Fiser et al. 2010). In the context of sensorimotor learning, generalisation studies have proven to be a highly useful tool for revealing underlying representations (Thoroughman and Shadmehr 2000; Poggio and Bizzi 2004; Shadmehr 2004; Paz and Vaadia 2009). These studies have the capacity to provide useful constraints on the format in which learning represented (and specifically, whether the learned information is encoded in an intrinsic or extrinsic coordinate frame), and the degree to which prior learning generalises to novel contexts. In light of the difficulty of investigating the neural basis of Bayesian estimation directly, researchers have recently started to use behavioural experiments to explore the issue. For instance, Fernandez et al. (2012) investigated how prior uncertainty generalized during reaching by having different groups of

subjects groups adapt to a visual rotation with the same mean but different levels of feedback uncertainty. They conclude that the generalisation of a stochastic perturbation is independent of the uncertainty of visual feedback, indicating that the prior and sensory likelihood are independently represented. Despite the usefulness of this study, Fernandez et al.'s results do not provide information about the reference frame in which learning and generalization occurs. By limiting generalization to an intralimb setting, the reference frame in which generalization occurred did not change, and thus the format in which such adaptation is generalized could not be investigated.

To provide useful constraints on the aspects of the underlying processes and representations through which sensorimotor learning is achieved, especially the frame of reference in which such learning occurs and the degree to which the learned representation is limb-specific or effector-general, is readily investigated via interlimb generalisation studies like ours.

Accordingly, our results indicate that Bayesian integration in visuomotor learning exhibits some degree of effector-specificity in the human motor system, and relatedly, that the underlying representations are predominantly encoded in an extrinsic reference frame. Furthermore, the strong transference of the prior and simultaneous lack of integration of visual uncertainty during early CE reaches indicates that the prior and sensory likelihood may be represented independently. Furthermore, our results provide useful constraints on both the PPC and NW models.

With reference to the PPC model, our results make distinct neural predictions for both CE and CI reaches. Although neuroimaging or neural recording was not a part of the current study, it is interesting to speculate about the neural correlates of the sensorimotor learning we observed. In particular, activity within relevant visuomotor populations that significantly correlate with both RH adaptation and LH CE reaches would be plausible candidates for limb-general representations of the prior, but not the likelihood. In contrast, population activity that correlates significantly with RH adaptation reaches exclusively, but not LH CE reaches, would be natural candidates for the

limb-specific representation of sensory likelihood. Of course, these conjectures will ultimately have to be judged according to the weight of neural evidence, when it is available.

In terms of the NW model, our results predict that the activity of neurons in the same population should be significantly correlated during both RH and LH reaches (CE and CI). Further, if that single population was to be identified (in some region in the sensorimotor processing network), it is expected that single-unit recording should identify a distinction in neural weighting. Specifically, in light of the NW model, we might expect that individual neurons within this population flexibly update their weights on the prior and sensory likelihood based on the reliability of the midpoint feedback provided about the true cursor position – i.e., the four visual uncertainty conditions (σ_0 , σ_M , σ_L , σ_∞). As uncertainty increases (reliability decreases), we might expect the modelled weights on the firing rates of these individual neurons to shift more toward the prior. And as uncertainty decreases, we might expect weights to shift toward the likelihood.

One obvious area to begin to test these hypotheses is in the neural regions known to be engaged during sensorimotor learning and interlimb generalisation (inferred from PET and fMRI) (Imamizu et al., 2000; Inoue et al., 2000; Miall et al., 2001; Krakauer et al., 2004; Graydon et al., 2005; Anguera et al., 2007). Most recently, Anguera et al., (2007) use fMRI to investigate whether any regions engaged during early visuomotor learning were also involved during transfer. They showed that during early learning of a visuomotor adaptation task increased activation occurred in frontal and parietal regions, including bilateral dorsal premotor cortex. During transfer, activation was seen in the temporal cortex as well as the right medial frontal gyrus and the middle occipital gyrus. Although these findings are far from conclusive, they do suggest that frontal cortex might be a useful place to start looking for neural populations involved in the Bayesian integration during sensorimotor learning and during transfer.

3.5 Future Directions

Having demonstrated how well-designed behavioural studies have the capacity to levy fruitful constraints on neural models of Bayesian integration, it goes without saying that more progress remains to be made in this area. Although we investigated transfer across intrinsically and extrinsically congruent reference frames, we did not investigate the degree of transfer across mixed reference frames (i.e., when the perturbation is constant across both extrinsic and intrinsic reference frames simultaneously). The fact that a number of recent studies (Carroll et al. 2014, 2016; Poh et al. 2016) indicate that transfer may be more complete when the visuomotor perturbation is aligned across both intrinsic and extrinsic reference frames highlights the importance of conducting similar investigations in a Bayesian context. Furthermore, in light of the fact that the findings of Carroll et al. (2014 and 2016) pertain to isometric force aiming and the findings of Poh et al. (2016) pertain predominantly to late transfer, further investigations utilising more naturalistic reach settings that account for all stages of transfer are needed.

Accordingly, we have designed and will soon run a protocol that investigates the extent to which a stochastic visuomotor adaptation transfers across limbs in congruent reference frames. Analogous to the visuomotor rotation paradigm run in our Experiments 1-3, subjects will complete reaches to one of several targets under four different visual feedback conditions in which visual uncertainty is varied ($\sigma_0, \sigma_M, \sigma_L, \sigma_\infty$) at the midpoint. Like the current experiments, visual feedback will only be provided during the veridical condition (σ_0), and the shift applied on each trial will be randomly drawn from a normal distribution (mean of 1 cm and standard deviation of 0.5 cm). Unlike our previous experiments, however, subjects will perform reaches starting from either the left or the right towards a target aligned with the body midline while adapting to a visual shift imposed along the sagittal plane (the y-plane of the screen) (Fig 20). Accordingly, the applied visual shift will be congruent across both intrinsic and extrinsic reference frames between left -and right-hand reaches. Further, we shall alter the trial scheduling in order to deliver more precisely

matched visual conditions across subjects during both training and testing phases. Unfettered by the goal of replicating Körding and Wolpert (2004), we will also incorporate a block of no-feedback (σ_∞) “probe” trials at the beginning of the testing phase to obtain an improved measure of interlimb transfer over the current experiments. Finally, borrowing from the innovative methodology developed by Poh et al. (2016), we shall interleave no-feedback probe trials (σ_∞) involving the trained and untrained limbs as a means by which to assess concurrent adaptation and transfer rates. Finally, as an important control, we also plan to run a (non-Bayesian) version of this experiment in which the imposed visuomotor perturbation is fixed.




Fig. 20 A) Paradigm for investigating visuomotor transfer across jointly congruent reference frames. The visual shift is applied positively in the y-plane for both left and right reaches. Accordingly, compensation requires an elbow flexion for both left and right reaches, which are thus congruent in both extrinsic and intrinsic reference frames. B) Midpoint feedback conditions with different amounts of visual uncertainty.

In order to further investigate the nature and representation of statistical priors, and the degree to which these representations generalise across the limbs, we also intend on implementing analogous protocols in which the mean and variance of the imposed prior are independently varied during learning (Fernandez et al. 2014). The degree of adaptation and interlimb generalisation observed across these follow-up experiments will help shed further light on the nature of Bayesian integration in sensorimotor learning.

References

- Acuna, D. E., Berniker, M., Fernandes, H. L., & Körding, K. P. (2015). Using psychophysics to ask if the brain samples or maximizes. *Journal of vision*, 15(3), 7-7.
- Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the 'light-from-above' prior. *Nature neuroscience*, 7(10), 1057-1058.
- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98(3), 409.
- Anguera, J. A., Russell, C. A., Noll, D. C., & Seidler, R. D. (2007). Neural correlates associated with intermanual transfer of sensorimotor adaptation. *Brain research*, 1185, 136-151.
- Atkeson, C. G. (1989). Learning arm kinematics and dynamics. *Annual review of neuroscience*, 12(1), 157-183.
- Bayes, M., & Price, M. (1763). An essay towards solving a problem in the doctrine of chances. by the late rev. mr. bayes, frs communicated by mr. price, in a letter to john canton, amfrs. *Philosophical Transactions (1683-1775)*, 370-418.
- Bays, P. M., & Wolpert, D. M. (2007). Computational principles of sensorimotor control that minimize uncertainty and variability. *The Journal of physiology*, 578(2), 387-396.
- Bechtel, W. (2008). *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*: Taylor & Francis.
- Bedford, F. L. (1989). Constraints on learning new mappings between perceptual dimensions. *Journal of Experimental Psychology: Human Perception and Performance*, 15(2), 232.
- Beierholm, U., Shams, L., Ma, W. J., & Körding, K. (2007). *Comparing Bayesian models for multisensory cue combination without mandatory integration*. Paper presented at the Advances in neural information processing systems.
- Beierholm, U. R., Quartz, S. R., & Shams, L. (2009). Bayesian priors are encoded independently from likelihoods in human multisensory perception. *Journal of vision*, 9(5), 23-23.
- Berger, J. O. (2013). *Statistical decision theory and Bayesian analysis*: Springer Science & Business Media.
- Berniker, M., Franklin, D. W., Flanagan, J. R., Wolpert, D. M., & Körding, K. (2014). Motor learning of novel dynamics is not represented in a single global coordinate system: evaluation of mixed coordinate representations and local learning. *Journal of neurophysiology*, 111(6), 1165-1182.
- Berniker, M., & Körding, K. (2008). Estimating the sources of motor errors for adaptation and generalization. *Nature neuroscience*, 11(12), 1454-1461.
- Berniker, M., & Körding, K. (2011). Bayesian approaches to sensory integration for motor control. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(4), 419-428.

- Bickle, J. (2006). Reducing mind to molecular pathways: Explicating the reductionism implicit in current cellular and molecular neuroscience. *Synthese*, 151(3), 411-434.
- Bowers, J. S., & Davis, C. J. (2012a). Bayesian just-so stories in psychology and neuroscience. *Psychological bulletin*, 138(3), 389.
- Bowers, J. S., & Davis, C. J. (2012b). Is that what Bayesians believe? reply to Griffiths, Chater, Norris, and Pouget (2012).
- Brainard, D. H., & Freeman, W. T. (1997). Bayesian color constancy. *Journal of the Optical Society of America*, 14(7), 1393-1411.
- Brashers-Krug, T., Shadmehr, R., & Bizzi, E. (1996). Consolidation in human motor memory. *Nature*, 382(6588), 252-255.
- Brayanov, J. B., & Smith, M. A. (2010). Bayesian and “anti-Bayesian” biases in sensory integration for action and perception in the size–weight illusion. *Journal of neurophysiology*, 103(3), 1518-1531.
- Brayanov, J. B., Press, D. Z., & Smith, M. A. (2012). Motor memory is encoded as a gain-field combination of intrinsic and extrinsic action representations. *Journal of Neuroscience*, 32(43), 14951-14965.
- Burgess, N., Doeller, C., Barry, C., & Jeffery, K. (2007). S. 03.02 The role of the hippocampus in visuospatial and episodic memory: animal and human models. *European Neuropsychopharmacology*, 17, S178-S179.
- Carroll, T. J., De Rugy, A., Howard, I. S., Ingram, J. N., & Wolpert, D. M. (2016). Enhanced cross-limb transfer of force-field learning for dynamics that are identical in extrinsic and joint-based coordinates for both limbs. *Journal of neurophysiology*, 115(1), 445-456.
- Carroll, T. J., Poh, E., & de Rugy, A. (2014). New visuomotor maps are immediately available to the opposite limb. *Journal of neurophysiology*, 111(11), 2232-2243.
- Chater, N., & Oaksford, M. (2008). *The probabilistic mind: Prospects for Bayesian cognitive science*: Oxford University Press, USA.
- Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in cognitive sciences*, 10(7), 287-291.
- Churchland, M. M., Afshar, A., & Shenoy, K. V. (2006). A central source of movement variability. *Neuron*, 52(6), 1085-1096.
- Cisek, P., Crammond, D. J., & Kalaska, J. F. (2003). Neural activity in primary motor and dorsal premotor cortex in reaching tasks with the contralateral versus ipsilateral arm. *Journal of neurophysiology*, 89(2), 922-942.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(03), 181-204.

- Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*: Oxford University Press.
- Colombo, M., & Hartmann, S. (2015). Bayesian cognitive science, unification, and explanation. *The British Journal for the Philosophy of Science*, axv036.
- Colombo, M., & Seriès, P. (2012). Bayes in the brain—on Bayesian modelling in neuroscience. *The British Journal for the Philosophy of Science*, 63(3), 697-723.
- Cothros, N., Wong, J., & Gribble, P. (2006). Are there distinct neural representations of object and limb dynamics? *Experimental Brain Research*, 173(4), 689-697.
- Craver, C. F. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*: Oxford University Press.
- Criscimagna-Hemminger, S. E., Donchin, O., Gazzaniga, M. S., & Shadmehr, R. (2003). Learned dynamics of reaching movements generalize from dominant to nondominant arm. *Journal of neurophysiology*, 89(1), 168-176.
- Cumming, G. (2013). *Understanding the new statistics: Effect sizes, confidence intervals, and meta-analysis*: Routledge.
- Cunningham, H., & Welch, R. B. (1994). Multiple concurrent visual-motor mappings: implications for models of adaptation. *Journal of Experimental Psychology: Human Perception and Performance*, 20(5), 987.
- d'Acremont, M., Schultz, W., & Bossaerts, P. (2013). The human brain encodes event frequencies while forming subjective beliefs. *The Journal of neuroscience*, 33(26), 10887-10897.
- Danks, D. (2007). Theory unification and graphical models in human categorization. *Causal learning: Psychology, philosophy, and computation*, 173-189.
- Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience* (Vol. 10): Cambridge, MA: MIT Press.
- Dekleva, B. M., Ramkumar, P., Wanda, P. A., Körding, K. P., & Miller, L. E. (2016). Uncertainty leads to persistent effects on reach representations in dorsal premotor cortex. *eLife*, 5, e14316.
- Doya, K. (2000). Reinforcement learning in continuous time and space. *Neural computation*, 12(1), 219-245.
- Doya, K. (2007). *Bayesian brain: Probabilistic approaches to neural coding*: MIT press.
- Eliasmith, C., & Anderson, C. H. (2003). Neural engineering. *Massachusetts Institute of Technology*.
- Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870), 429-433.

- Faisal, A. A., Selen, L. P., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature reviews neuroscience*, 9(4), 292-303.
- Fernandes, H. L. (2013). Uncertainty, generalization, and neural representation of relevant variables for decision making.
- Fernandes, H. L., Stevenson, I. H., & Körding, K. P. (2012). Generalization of stochastic visuomotor rotations. *PloS one*, 7(8), e43016.
- Fernandes, H. L., Stevenson, I. H., Vilares, I., & Körding, K. P. (2014). The generalization of prior uncertainty during reaching. *The Journal of neuroscience*, 34(34), 11470-11484.
- Fetsch, C. R., DeAngelis, G. C., & Angelaki, D. E. (2013). Bridging the gap between theories of sensory cue integration and the physiology of multisensory neurons. *Nature reviews neuroscience*, 14(6), 429-442.
- Fetsch, C. R., Pouget, A., DeAngelis, G. C., & Angelaki, D. E. (2012). Neural correlates of reliability-based cue weighting during multisensory integration. *Nature neuroscience*, 15(1), 146-154.
- Fetsch, C. R., Turner, A. H., DeAngelis, G. C., & Angelaki, D. E. (2009). Dynamic reweighting of visual and vestibular cues during self-motion perception. *The Journal of neuroscience*, 29(49), 15601-15612.
- Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends in cognitive sciences*, 14(3), 119-130.
- Flanagan, J. R., Nakano, E., Imamizu, H., Osu, R., Yoshioka, T., & Kawato, M. (1999). Composition and decomposition of internal models in motor learning under altered kinematic and dynamic environments. *Journal of Neuroscience*, 19, RC34 (31-35).
- Flanagan, J. R., & Rao, A. K. (1995). Trajectory adaptation to a nonlinear visuomotor transformation: evidence of motion planning in visually perceived space. *Journal of neurophysiology*, 74(5), 2174-2178.
- 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.
- Fletcher, P. C., & Frith, C. D. (2009). Perceiving is believing: a Bayesian approach to explaining the positive symptoms of schizophrenia. *Nature reviews neuroscience*, 10(1), 48-58.
- Franklin, D. W., & Wolpert, D. M. (2011). Computational mechanisms of sensorimotor control. *Neuron*, 72(3), 425-442.
- Frensch, P. A., Lin, J., & Buchner, A. (1998). Learning versus behavioral expression of the learned: The effects of a secondary tone-counting task on implicit learning in the serial reaction task. *Psychological Research*, 61(2), 83-98.

- Gandolfo, F., Mussa-Ivaldi, F., & Bizzi, E. (1996). Motor learning by field approximation. *Proceedings of the National Academy of Sciences*, 93(9), 3843-3846.
- Georgopoulos, A. P., Kalaska, J. F., Caminiti, R., & Massey, J. T. (1982). On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *The Journal of neuroscience*, 2(11), 1527-1537.
- Ghahramani, Z., Wolpert, D. M., & Jordan, M. I. (1996). Generalization to local remappings of the visuomotor coordinate transformation. *The Journal of neuroscience*, 16(21), 7085-7096.
- Ghilardi, M. F., Gordon, J., & Ghez, C. (1995). Learning a visuomotor transformation in a local area of work space produces directional biases in other areas. *Journal of neurophysiology*, 73(6), 2535-2539.
- Gordon, J., Ghilardi, M. F., Cooper, S. E., & Ghez, C. (1994). Accuracy of planar reaching movements. *Experimental Brain Research*, 99(1), 112-130.
- Graydon, F. X., Friston, K. J., Thomas, C. G., Brooks, V. B., & Menon, R. S. (2005). Learning-related fMRI activation associated with a rotational visuo-motor transformation. *Cognitive Brain Research*, 22(3), 373-383.
- Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in cognitive sciences*, 14(8), 357-364.
- Griffiths, T. L., & Kalish, M. L. (2007). Language evolution by iterated learning with Bayesian agents. *Cognitive science*, 31(3), 441-480.
- Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review*, 116(4), 661.
- Hahn, U. (2014). The Bayesian boom: good thing or bad? *Frontiers in psychology*, 5, 765.
- Harris, C. M., & Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, 394(6695), 780-784.
- Held, R., & Freedman, S. J. (1963). Plasticity in human sensorimotor control. *Science*, 142(3591), 455-462.
- Imamizu, H., Miyauchi, S., Tamada, T., Sasaki, Y., Takino, R., Pütz, B., & Kawato, M. (2000). Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature*, 403(6766), 192-195.
- Imamizu, H., & Shimojo, S. (1995). The locus of visual-motor learning at the task or manipulator level: implications from intermanual transfer. *Journal of Experimental Psychology: Human Perception and Performance*, 21(4), 719.
- Inoue, K., Kawashima, R., Satoh, K., Kinomura, S., Goto, R., Sugiura, M., & Fukuda, H. (1997). Activity in the parietal area during visuomotor learning with optical rotation. *Neuroreport*, 8(18), 3979-3983.

- Isaias, I. U., Moisello, C., Marotta, G., Schiavella, M., Canesi, M., Perfetti, B., & Ghilardi, M. F. (2011). Dopaminergic striatal innervation predicts interlimb transfer of a visuomotor skill. *The Journal of neuroscience*, 31(41), 14458-14462.
- Jacobs, R. A. (1999). Optimal integration of texture and motion cues to depth. *Vision research*, 39(21), 3621-3629.
- Joiner, W. M., Brayanov, J. B., & Smith, M. A. (2013). The training schedule affects the stability, not the magnitude, of the interlimb transfer of learned dynamics. *Journal of neurophysiology*, 110(4), 984-998.
- Jones, K. E., Hamilton, A. F. d. C., & Wolpert, D. M. (2002). Sources of signal-dependent noise during isometric force production. *Journal of neurophysiology*, 88(3), 1533-1544.
- Jones, M., & Love, B. C. (2011). Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and Brain Sciences*, 34(04), 169-188.
- Jordan, M. I., & Rumelhart, D. E. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive science*, 16(3), 307-354.
- Kaplan, D. M. (2011). Explanation and description in computational neuroscience. *Synthese*, 183(3), 339-373.
- Kaplan, D. M. (2015). Explanation and Levels in Cognitive Neuroscience *Handbook of Neuroethics* (pp. 9-29): Springer.
- Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current opinion in neurobiology*, 9(6), 718-727.
- Kawato, M., Furukawa, K., & Suzuki, R. (1987). A hierarchical neural-network model for control and learning of voluntary movement. *Biological cybernetics*, 57(3), 169-185.
- Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annual Review of Psychology*, 55, 271-304.
- Kluzik, J., Diedrichsen, J., Shadmehr, R., & Bastian, A. J. (2008). Reach adaptation: what determines whether we learn an internal model of the tool or adapt the model of our arm? *Journal of neurophysiology*, 100(3), 1455-1464.
- Knill, D. C., Kersten, D., & Yuille, A. (1996). Introduction: A Bayesian formulation of visual perception. *Perception as Bayesian inference*, 1-21.
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *TRENDS in Neurosciences*, 27(12), 712-719.
- Knill, D. C., & Richards, W. (1996). *Perception as Bayesian inference*: Cambridge University Press.
- Knill, D. C., & Saunders, J. A. (2003). Do humans optimally integrate stereo and texture information for judgments of surface slant? *Vision research*, 43(24), 2539-2558.

- Körding, K. (2007). Decision theory: what" should" the nervous system do? *Science*, 318(5850), 606-610.
- Körding, K. P. (2014). Bayesian statistics: relevant for the brain? *Current opinion in neurobiology*, 25, 130-133.
- Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244-247.
- Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in cognitive sciences*, 10(7), 319-326.
- Krakauer, J. W. (2009). Motor learning and consolidation: the case of visuomotor rotation. *Progress in motor control* (pp. 405-421): Springer.
- Krakauer, J. W., Ghez, C., & Ghilardi, M. F. (2005). Adaptation to visuomotor transformations: consolidation, interference, and forgetting. *The Journal of neuroscience*, 25(2), 473-478.
- Krakauer, J. W., Ghilardi, M.-F., & Ghez, C. (1999). Independent learning of internal models for kinematic and dynamic control of reaching. *Nature neuroscience*, 2(11), 1026-1031.
- Krakauer, J. W., Ghilardi, M.-F., Mentis, M., Barnes, A., Veytsman, M., Eidelberg, D., & Ghez, C. (2004). Differential cortical and subcortical activations in learning rotations and gains for reaching: a PET study. *Journal of neurophysiology*, 91(2), 924-933.
- Krakauer, J. W., & Mazzoni, P. (2011). Human sensorimotor learning: adaptation, skill, and beyond. *Current opinion in neurobiology*, 21(4), 636-644.
- Krakauer, J. W., Mazzoni, P., Ghazizadeh, A., Ravindran, R., & Shadmehr, R. (2006). Generalization of motor learning depends on the history of prior action. *PLOS Biol*, 4(10), e316.
- Krakauer, J. W., Pine, Z. M., Ghilardi, M.-F., & Ghez, C. (2000). Learning of visuomotor transformations for vectorial planning of reaching trajectories. *The Journal of neuroscience*, 20(23), 8916-8924.
- Kruschke, J. K. (2008). Models of categorization. *The Cambridge handbook of computational psychology*, 267-301.
- Kruschke, J. K. (2010). What to believe: Bayesian methods for data analysis. *Trends in cognitive sciences*, 14(7), 293-300.
- Kuperstein, M. (1988). Neural model of adaptive hand-eye coordination for single postures. *Science*, 239(4845), 1308-1311.
- Lackner, J. R., & Dizio, P. (1994). Rapid adaptation to Coriolis force perturbations of arm trajectory. *Journal of neurophysiology*, 72(1), 299-313.
- Landy, M. S., Maloney, L. T., Johnston, E. B., & Young, M. (1995). Measurement and modeling of depth cue combination: in defense of weak fusion. *Vision research*, 35(3), 389-412.

- Lee, M. D., & Wagenmakers, E.-J. (2014). *Bayesian cognitive modeling: A practical course*: Cambridge University Press.
- Logothetis, N. K. (2008). What we can do and what we cannot do with fMRI. *Nature*, 453(7197), 869-878.
- Ma, W. J. (2012). Organizing probabilistic models of perception. *Trends in cognitive sciences*, 16(10), 511-518.
- Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature neuroscience*, 9(11), 1432-1438.
- Ma, W. J., Beck, J. M., & Pouget, A. (2008). Spiking networks for Bayesian inference and choice. *Current opinion in neurobiology*, 18(2), 217-222.
- Ma, W. J., & Jazayeri, M. (2014). Neural coding of uncertainty and probability. *Annual review of neuroscience*, 37, 205-220.
- Ma, W. J., & Rahmati, M. (2013). Towards a neural implementation of causal inference in cue combination. *Multisens Res*, 26, 159-176.
- Malfait, N., Shiller, D. M., & Ostry, D. J. (2002). Transfer of motor learning across arm configurations. *The Journal of neuroscience*, 22(22), 9656-9660.
- Maloney, L. T. (2002). *Statistical decision theory and biological vision*. In (Eds.) Dieter Heyer and Rainer Mauself. Perception and the physical world: Psychological and philosophical issues in perception. Wiley & Sons, Ltd.
- Maloney, L. T., & Mamassian, P. (2009). Bayesian decision theory as a model of human visual perception: testing Bayesian transfer. *Visual neuroscience*, 26(01), 147-155.
- Martin, T., Keating, J., Goodkin, H., Bastian, A., & Thach, W. (1996a). Throwing while looking through prisms. *Brain*, 119(4), 1183-1198.
- Martin, T., Keating, J., Goodkin, H., Bastian, A., & Thach, W. (1996b). Throwing while looking through prisms. II. Specificity and storage of multiple gaze-throw calibrations. *Brain*, 119(4), 1199-1212.
- Mauk, M. D. (2000). The potential effectiveness of simulations versus phenomenological models. *Nature neuroscience*, 3(7), 649-651.
- Miall, R., Reckess, G., & Imamizu, H. (2001). The cerebellum coordinates eye and hand tracking movements. *Nature neuroscience*, 4(6), 638-644.
- Moran, D. W., & Schwartz, A. B. (1999). Motor cortical representation of speed and direction during reaching. *Journal of neurophysiology*, 82(5), 2676-2692.
- Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the brain: a diffusion model analysis of prior probability and potential payoff. *The Journal of neuroscience*, 32(7), 2335-2343.

- Orbán, G., & Wolpert, D. M. (2011). Representations of uncertainty in sensorimotor control. *Current opinion in neurobiology*, 21(4), 629-635.
- Palmer, C., Zapparoli, L., & Kilner, J. M. (2016). A New Framework to Explain Sensorimotor Beta Oscillations. *Trends in cognitive sciences*, 20(5), 321-323.
- Parmar, P. N., Huang, F. C., & Patton, J. L. (2015). Evidence of multiple coordinate representations during generalization of motor learning. *Experimental Brain Research*, 233(1), 1-13.
- Pashler, H. (1994). Dual-task interference in simple tasks: data and theory. *Psychological bulletin*, 116(2), 220.
- Poggio, T., & Bizzi, E. (2004). Generalization in vision and motor control. *Nature*, 431(7010), 768-774.
- Poh, E., Carroll, T. J., & Taylor, J. A. (2016). Effect of coordinate frame compatibility on the transfer of implicit and explicit learning across limbs. *Journal of Neurophysiology*, 116(3), 1239-1249.
- Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: knowns and unknowns. *Nature neuroscience*, 16(9), 1170-1178.
- Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature reviews neuroscience*, 1(2), 125-132.
- Puga, J. L., Krzywinski, M., & Altman, N. (2015). Points of Significance: Bayes' theorem. *Nature Methods*, 12(4), 277-278.
- Purves, D., Wojtach, W. T., & Lotto, R. B. (2011). Understanding vision in wholly empirical terms. *Proceedings of the National Academy of Sciences*, 108(Supplement 3), 15588-15595.
- Rahnev, D., Lau, H., & de Lange, F. P. (2011). Prior expectation modulates the interaction between sensory and prefrontal regions in the human brain. *The Journal of neuroscience*, 31(29), 10741-10748.
- Rao, R. P. (2004). Bayesian computation in recurrent neural circuits. *Neural computation*, 16(1), 1-38.
- Rao, R. P. (2007). Neural Models of Bayesian Belief Propagation. *Bayesian brain: Probabilistic approaches to neural coding*, 239.
- Raposo, D., Sheppard, J. P., Schrater, P. R., & Churchland, A. K. (2012). Multisensory decision-making in rats and humans. *The Journal of neuroscience*, 32(11), 3726-3735.
- Rich, D., Cazettes, F., Wang, Y., Peña, J. L., & Fischer, B. J. (2015). Neural representation of probabilities for Bayesian inference. *Journal of computational neuroscience*, 38(2), 315-323.

- Rotella, M. F., Nisky, I., Koehler, M., Rinderknecht, M. D., Bastian, A. J., & Okamura, A. M. (2015). Learning and generalization in an isometric visuomotor task. *Journal of neurophysiology*, 113(6), 1873-1884.
- Sabes, P. N. (2000). The planning and control of reaching movements. *Current opinion in neurobiology*, 10(6), 740-746.
- Sainburg, R. L., & Wang, J. (2002). Interlimb transfer of visuomotor rotations: independence of direction and final position information. *Experimental Brain Research*, 145(4), 437-447.
- Sato, Y., & Körding, K. P. (2014). How much to trust the senses: Likelihood learning. *Journal of vision*, 14(13), 13-13.
- Schmidt, R. A., & Lee, T. (1988). *Motor control and learning: Human kinetics*.
- Schrater, P., & Kersten, D. (2002). 2 Vision, Psychophysics and Bayes. *Probabilistic Models of the Brain*, 37.
- Scott, S. H., Gribble, P. L., Graham, K. M., & Cabel, D. W. (2001). Dissociation between hand motion and population vectors from neural activity in motor cortex. *Nature*, 413(6852), 161-165.
- Sergio, L. E., Hamel-Pâquet, C., & Kalaska, J. F. (2005). Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. *Journal of neurophysiology*, 94(4), 2353-2378.
- Shadmehr, R. (2004). Generalization as a behavioral window to the neural mechanisms of learning internal models. *Human movement science*, 23(5), 543-568.
- Shadmehr, R., & Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *The Journal of neuroscience*, 14(5), 3208-3224.
- Shadmehr, R., & Mussa-Ivaldi, S. (2012). *Biological learning and control: how the brain builds representations, predicts events, and makes decisions*: MIT Press.
- Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual review of neuroscience*, 33, 89-108.
- Shadmehr, R., & Wise, S. P. (2005). *The computational neurobiology of reaching and pointing: a foundation for motor learning*: MIT press.
- Shams, L. (2012). Early integration and Bayesian causal inference in multisensory perception.
- Smith, M. A., Ghazizadeh, A., & Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol*, 4(6), e179.
- Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11(01), 1-23.
- Sober, S. J., & Sabes, P. N. (2003). Multisensory integration during motor planning. *The Journal of neuroscience*, 23(18), 6982-6992.

- Sober, S. J., & Sabes, P. N. (2005). Flexible strategies for sensory integration during motor planning. *Nature neuroscience*, 8(4), 490-497.
- Soechting, J., & Flanders, M. (1992). Moving in three-dimensional space: frames of reference, vectors, and coordinate systems. *Annual review of neuroscience*, 15(1), 167-191.
- Soechting, J. F., Buneo, C. A., Herrmann, U., & Flanders, M. (1995). Moving effortlessly in three dimensions: does Donders' law apply to arm movement? *The Journal of neuroscience*, 15(9), 6271-6280.
- Stein, B. E., & Meredith, M. A. (1993). *The merging of the senses*: Cambridge, MA: MIT Press.
- Stocker, A. A., & Simoncelli, E. P. (2006). Noise characteristics and prior expectations in human visual speed perception. *Nature neuroscience*, 9(4), 578-585.
- Stocker, A. A., & Simoncelli, E. P. (2008). *A Bayesian model of conditioned perception*. In (Eds.) J. C. Platt, D. Koller, Y. Singer, & S. Roweis. *Advances in Neural Information Processing Systems*, 1409–1416. Cambridge, MA: MIT Press.
- Stockinger, C., Thüerer, B., Focke, A., & Stein, T. (2015). Intermanual transfer characteristics of dynamic learning: direction, coordinate frame, and consolidation of interlimb generalization. *Journal of neurophysiology*, 114(6), 3166-3176.
- Tan, H., Wade, C., & Brown, P. (2016). Post-movement beta activity in sensorimotor cortex indexes confidence in the estimations from internal models. *The Journal of neuroscience*, 36(5), 1516-1528.
- Tassinari, H., Hudson, T. E., & Landy, M. S. (2006). Combining priors and noisy visual cues in a rapid pointing task. *The Journal of neuroscience*, 26(40), 10154-10163.
- Taylor, J. A., & Ivry, R. B. (2013). Context-dependent generalization. *Frontiers in Human Neuroscience*, 7(171).
- Taylor, J. A., Krakauer, J. W., & Ivry, R. B. (2014). Explicit and implicit contributions to learning in a sensorimotor adaptation task. *The Journal of neuroscience*, 34(8), 3023-3032.
- Taylor, J. A., & Thoroughman, K. A. (2007). Divided attention impairs human motor adaptation but not feedback control. *Journal of neurophysiology*, 98(1), 317-326.
- Taylor, J. A., Wojaczynski, G. J., & Ivry, R. B. (2011). Trial-by-trial analysis of intermanual transfer during visuomotor adaptation. *Journal of neurophysiology*, 106(6), 3157-3172.
- Thompson, A. A., & Henriques, D. Y. (2010). Locations of serial reach targets are coded in multiple reference frames. *Vision research*, 50(24), 2651-2660.
- Thoroughman, K. A., & Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature*, 407(6805), 742-747.

- Ting, C.-C., Yu, C.-C., Maloney, L. T., & Wu, S.-W. (2015). Neural mechanisms for integrating prior knowledge and likelihood in value-based probabilistic inference. *The Journal of neuroscience*, 35(4), 1792-1805.
- Tong, C., Wolpert, D. M., & Flanagan, J. R. (2002). Kinematics and dynamics are not represented independently in motor working memory: evidence from an interference study. *The Journal of neuroscience*, 22(3), 1108-1113.
- van Beers, R. J., Sittig, A. C., & van der Gon Denier, J. J. (1996). How humans combine simultaneous proprioceptive and visual position information. *Experimental Brain Research*, 111(2), 253-261.
- van Beers, R. J., Sittig, A. C., & van Der Gon, J. J. D. (1999). Integration of proprioceptive and visual position-information: An experimentally supported model. *Journal of neurophysiology*, 81(3), 1355-1364.
- van Bergen, R. S., Ma, W. J., Pratte, M. S., & Jehee, J. F. (2015). Sensory uncertainty decoded from visual cortex predicts behavior. *Nature neuroscience*.
- 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), 10050-10059.
- Vetter, P., Goodbody, S. J., & Wolpert, D. M. (1999). Evidence for an eye-centered spherical representation of the visuomotor map. *Journal of neurophysiology*, 81(2), 935-939.
- Vilares, I., Howard, J. D., Fernandes, H. L., Gottfried, J. A., & Körding, K. P. (2012). Differential representations of prior and likelihood uncertainty in the human brain. *Current Biology*, 22(18), 1641-1648.
- Vilares, I., & Körding, K. (2011). Bayesian models: the structure of the world, uncertainty, behavior, and the brain. *Annals of the New York Academy of Sciences*, 1224(1), 22-39.
- Vindras, P., & Viviani, P. (1998). Frames of reference and control parameters in visuomanual pointing. *Journal of Experimental Psychology: Human Perception and Performance*, 24(2), 569.
- Wang, J. (2008). A dissociation between visual and motor workspace inhibits generalization of visuomotor adaptation across the limbs. *Experimental Brain Research*, 187(3), 483-490.
- Wang, J., & Sainburg, R. L. (2003). Mechanisms underlying interlimb transfer of visuomotor rotations. *Experimental Brain Research*, 149(4), 520-526.
- Wang, J., & Sainburg, R. L. (2004a). Interlimb transfer of novel inertial dynamics is asymmetrical. *Journal of neurophysiology*, 92(1), 349-360.
- Wang, J., & Sainburg, R. L. (2004b). Limitations in interlimb transfer of visuomotor rotations. *Experimental Brain Research*, 155(1), 1-8.
- Wang, J., & Sainburg, R. L. (2005). Adaptation to visuomotor rotations remaps movement vectors, not final positions. *The Journal of neuroscience*, 25(16), 4024-4030.

- Wang, J., & Sainburg, R. L. (2006). The symmetry of interlimb transfer depends on workspace locations. *Experimental Brain Research*, 170(4), 464-471.
- Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature neuroscience*, 5(6), 598-604.
- Wigmore, V., Tong, C., & Flanagan, J. R. (2002). Visuomotor rotations of varying size and direction compete for a single internal model in a motor working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 28(2), 447.
- Wolpert, D. M. (2007). Probabilistic models in human sensorimotor control. *Human movement science*, 26(4), 511-524.
- Wolpert, D. M., Diedrichsen, J., & Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nature reviews neuroscience*, 12(12), 739-751.
- Wolpert, D. M., & Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature neuroscience*, 3, 1212-1217.
- Wolpert, D. M., Ghahramani, Z., & Flanagan, J. R. (2001). Perspectives and problems in motor learning. *Trends in cognitive sciences*, 5(11), 487-494.
- Wolpert, D. M., Ghahramani, Z., & Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science*, 269(5232), 1880.
- Wu, H. G., & Smith, M. A. (2013). The generalization of visuomotor learning to untrained movements and movement sequences based on movement vector and goal location remapping. *The Journal of neuroscience*, 33(26), 10772-10789.
- Yang, J.-l., Chang, C.-w., Chen, S.-y., & Lin, J.-j. (2008). Shoulder kinematic features using arm elevation and rotation tests for classifying patients with frozen shoulder syndrome who respond to physical therapy. *Manual therapy*, 13(6), 544-551.
- Young, M. J., Landy, M. S., & Maloney, L. T. (1993). A perturbation analysis of depth perception from combinations of texture and motion cues. *Vision research*, 33(18), 2685-2696.
- Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? *Trends in cognitive sciences*, 10(7), 301-308.
- Zemel, R. S., Dayan, P., & Pouget, A. (1998). Probabilistic interpretation of population codes. *Neural computation*, 10(2), 403-430.

Appendix

A.1. Ethics approval and consent form

Office of the Deputy Vice-Chancellor
(Research)

Research Office
Research Hub, Building C5C East
Macquarie University
NSW 2109 Australia
T: +61 (2) 9850 4459
<http://www.research.mq.edu.au/>
ABN 90 952 801 237



MACQUARIE
University
SYDNEY · AUSTRALIA

12 May 2016

Dear Dr Kaplan

Reference No: 5201600282

Title: *Using interlimb transfer to test and constrain the Bayesian approach to sensorimotor learning*

Thank you for submitting the above application for ethical and scientific review. Your application was considered by the Macquarie University Human Research Ethics Committee (HREC (Medical Sciences)).

I am pleased to advise that ethical and scientific approval has been granted for this project to be conducted at:

- Macquarie University

This research meets the requirements set out in the *National Statement on Ethical Conduct in Human Research* (2007 – Updated May 2015) (the *National Statement*).

Standard Conditions of Approval:

1. Continuing compliance with the requirements of the *National Statement*, which is available at the following website:

<http://www.nhmrc.gov.au/book/national-statement-ethical-conduct-human-research>

2. This approval is valid for five (5) years, subject to the submission of annual reports. Please submit your reports on the anniversary of the approval for this protocol.

3. All adverse events, including events which might affect the continued ethical and scientific acceptability of the project, must be reported to the HREC within 72 hours.

4. Proposed changes to the protocol and associated documents must be submitted to the Committee for approval before implementation.

It is the responsibility of the Chief investigator to retain a copy of all documentation related to this project and to forward a copy of this approval letter to all personnel listed on the project.

Should you have any queries regarding your project, please contact the Ethics Secretariat on 9850 4194 or by email ethics.secretariat@mq.edu.au

The HREC (Medical Sciences) Terms of Reference and Standard Operating Procedures are available from the Research Office website at:

http://www.research.mq.edu.au/for/researchers/how_to_obtain_ethics_approval/human_research_ethics

The HREC (Medical Sciences) wishes you every success in your research.

Yours sincerely



Professor Tony Eyers

Chair, Macquarie University Human Research Ethics Committee (Medical Sciences)

This HREC is constituted and operates in accordance with the National Health and Medical Research Council's (NHMRC) *National Statement on Ethical Conduct in Human Research* (2007) and the *CPMP/ICH Note for Guidance on Good Clinical Practice*.

Details of this approval are as follows:

Approval Date: 10 May 2016

The following documentation has been reviewed and approved by the HREC (Medical Sciences):

Documents reviewed	Version no.	Date
Macquarie University Ethics Application Form		Received 11/4/2016
Correspondence responding to the issues raised by the HREC (Medical Sciences)		Received 3/5/2016
MQ Participant Information and Consent Form (PICF) entitled <i>Interlimb transfer and Bayesian sensorimotor learning</i>	1.0	3/5/2016
CSR Advertisement	1.0	3/5/2016

***If the document has no version date listed one will be created for you. Please ensure the footer of these documents are updated to include this version date to ensure ongoing version control.**



Information and Consent Form: *Interlimb transfer and Bayesian sensorimotor learning*

You are invited to participate in a study that investigates human sensorimotor learning and seeks to test the extent to which statistical information learned about a motor task using one limb is available to the other limb.

This research is being conducted by:

Chris Hewitson (Cognitive Science, christopher.hewitson1@students.mq.edu.au)

Dr David Kaplan (Cognitive Science, david.kaplan@mq.edu.au)

If you decide to participate, you will be asked to view and respond to various visual stimuli by making arm movements. The movements you make and the timing associated with these movements will be recorded. The experimental session may take up to 60 minutes to complete. No risks are expected to result from participation. You will receive \$15 per hour (or pro rata) for your participation. If you are participating for course credit, you will receive one credit for each half-hour of your participation.

Any information or personal details (e.g. age, gender) gathered in the course of the study are kept confidential. Access to the data is limited to persons directly involved in the research. You may request feedback after the study has been completed and the analysis of the results is available. No individual will be identified in any publication of the results.

Participation in this study is voluntary and you are free to withdraw from further participation in the research at any time without having to give a reason and without consequence. Macquarie University students who are participating as part of their course requirements will not forfeit their course credits if they choose to withdraw from the research.

I, _____ have read (*or, where appropriate, have had read to me*) and understand the information above and any questions I have asked have been answered to my satisfaction. I agree to participate in this research, knowing that I can withdraw from further participation in the research at any time without consequence. I have been given a copy of this form to keep.

Participant's Name: _____
(Block letters)

Participant's Signature: _____ Date: _____


Investigator's Name: _____
(Block letters)

Investigator's Signature: _____ Date: _____

The ethical aspects of this study have been approved by the Macquarie University Human Research Ethics Committee. If you have any complaints or reservations about any ethical aspect of your participation in this research, you may contact the Committee through the Director, Research Ethics (telephone (02) 9850 7854; email ethics@mq.edu.au). Any complaint you make will be treated in confidence and investigated, and you will be informed of the outcome.

A.2 The KINARM in detail

The unimanual version of the KINARM endpoint robot was utilized in all experiments (Appendix figure 1a). The KINARM includes a single handle manipulandum which provides 2D movement across a horizontal plane. The handle is located beneath a 2D virtual reality display, which allows the presentation of visual stimuli such that the stimuli (e.g., targets for reaching movements, and a cursor that represents handle location) appear at the same horizontal level as the hand (Appendix Figure 1b). Feedback of hand position is acquired through incremental encoders that log position incrementally at 200Hz.



Appendix Figure 1: A The KINARM endpoint Robot. B A Schematic of the KINARM in use.

A.2.1 Dexterit-E TM experimental control and data acquisition software

The Dexterit-E TM experimental control and data acquisition software runs on a multi-computer system. Dexterit-E itself runs on a Windows-based computer, in which it effectively acts as a user-interface for choosing task protocols, providing visual feedback to the operator, and saving data. The chosen task protocol is associated with a real-time computer, which is used to control the task. The real-time computer runs Xpc Target, a Simulink model (Mathworks Corporation), with which Dexterit-E interfaces. During the execution of a reaching task, communication between the real-time computer to the Windows-based computer allows the Windows-based computer to offer online feedback to the operator.

A.2.2 Creating a task program


Task programs are created to define and control the system behaviour that can occur during a single trial of a reaching task implemented by Dexterit-E. For a generic reaching task, the task protocol may be defined as follows:

1. A start target will turn on during a trial.
2. Once a subject reaches that target, the target will change state and a second (end) target will turn on.
3. The subject reaches to the second target, which will turn off when the subject reaches the target.
4. trial is over.
5. The start target reverts to its initial state.
6. Back to step 1.

The task program does not define the details of the task, such as the target location, color, inter-trial period, visual feedback conditions or the number of trials. These parameters are specified through the BKIN Dexterit-E's windows-based interface. Creating a task program involves modelling the task parameters in Simulink via the Stateflow toolboxes. Simulink is a block diagram programming environment for a model based design in which task programs are developed and represented as a graph of data flow in the task.

Stateflow is a graphical design tool for developing event-driven state machine that allows transitions between the states defined in the task. Task programs built in stateflow are compiled by Matlab and the xPC Target toolbox, using a generic C/C++ compiler.

The figure below (Appendix Figure 2) shows a schematic of the stateflow chart that controls task parameters during a single trial in our experiments.



Appendix figure 2: Simulink State flow of the steps implemented during experimental protocol.

In this flow chart, states are represented by ovals and transitions between states are represented with arrows. An event can only exist in one state at a time, and the event must be in one of the defined states. The event transfer from one state to another can only occur when there is a transition between the states, and the conditions for that transition are true. The eleven states in the flow state are as follows:

1. **Initialize** – When this state is entered, all target and cursor variables are initialised with empty states.
2. **FirstTargetOn** – When this state is entered, the first target will be turned on. In addition, three logical windows (invisible windows within which visual feedback is modified) will be turned on. The event will stay in this state until the cursor enters the first target.
3. **StayAtFirstTarget** – When this state is entered, data logging is initialized, indicating that the cursor has entered the first target. At this point, the first target changes state (color). Once the condition “wait for a specific time” is true, the event will switch to the “ShowSecondTarget” state.
4. **ApplyShift** – This state is entered when the cursor leaves the first target. At this point, visual feedback of the cursor is extinguished, and a generated shift (from veridical hand position) is applied to the cursor position.
5. **ShowSecondTarget** – When this state is entered, the second target will be turned on, indicating to the subject that they may commence their reaching movement. The condition “MaxReachTime” will turn on, initializing a count-down timer signaling the maximum trial duration. The event will stay in this state until the cursor exits the first target, or max trial duration has past
6. **ShowCursor** – This state is entered when the hand position enters the logical window at the mid-point of the reach, at which point the designated visual feedback condition is turned on at the shifted position.
7. **SwitchOfCursor2** – When the cursor leaves the mid-point logical window, visual feedback of the cursor position is turned off.
8. **StayAtSecondTarget** – When the hand position enters the second target, this state is initialised. While the target state (colour) remains the same, visual feedback of the cursor

will change according to the designated feedback type. of the cursor entered Once the condition “wait at second target for a specific time” is true, the event will switch to the “EndOfTrial” state.

9. **EndOfTrial** – When this state is entered, the states of all initialised variables within the previous trial return to zero, and the next trial commences. – When this state is entered, the target will be turned off. This state will switch to “Between Trials” state.
10. **InterTrialState** – When this state is entered, State flow sends a signal to the task program that trial is over and to provide a specific time delay, allowing the Task Program to update the Trial Protocol for the next trial. Exiting from this state back to the “Target On state” for the next trial occurs after a specified time delay.

A.3 Protocol scripts

A.3.1 Point cloud function

```
function makePointclouds(num_figs,number_of_points,meanc,stdc)

%use as makePointclouds(num_figs) where num_figs is the number of figures
%you want created, number of points is the # points in the cloud drawn from
distribution meanc, stdc

close all;
meanc=meanc*100;
stdc=stdc*100;

for i=1:num_figs
    n=number_of_points; % number of points
    radius = meanc + stdc.*randn(n,1); %draw n radii
    xc = randn(n,1); %
    yc = randn(n,1);

    xc=xc';
    yc=yc';

    theta = rand(1,n)*(2*pi);
    r = sqrt(rand(1,n)).*radius';
    x = xc + r.*cos(theta);
    y = yc + r.*sin(theta);

    hs=scatter(x,y,10);

    set(hs,'MarkerFaceColor','w');
    set(hs,'MarkerEdgeColor','w');

    set(hs,'MarkerFaceAlpha',0.5);
    set(hs,'MarkerEdgeAlpha',0.0);

    set(gcf, 'Position', [0 0 400 400]);

    axis equal;
    axis off;

    xlim([-400 400]);
    ylim([-400 400]);

    title([]);

    %hold on

    export_fig(['pointcloud_',num2str(stdc/100),'_',num2str(i)],...
        '-transparent','-png','-native','-nocrop');
end

end;
```

A.3.2 Read in .c3d, sort data and designate required bin size

```
%Read in c3d data
% Trial type 1 = sigma_1, 10-1 = sigma_m, 20-29 = sigma_L, 30 = sigma_0;
right_hand_veridical = c3d_load('*_01_*.c3d');
right_hand_low = c3d_load('*_10_*.c3d', '*_11_*.c3d', '*_12_*.c3d',...
'*_13_*.c3d', '*_14_*.c3d', '*_15_*.c3d', '*_16_*.c3d', '*_17_*.c3d',
'*_18_*.c3d', '*_19_*.c3d');
right_hand_high = c3d_load('*_20_*.c3d', '*_21_*.c3d', '*_22_*.c3d',
'*_23_*.c3d', '*_24_*.c3d', '*_25_*.c3d', '*_26_*.c3d', '*_27_*.c3d',...
'*_28_*.c3d', '*_29_*.c3d');
right_hand_unlimited = c3d_load('*_30_*.c3d');

%Sort data by ascending trial number
for idx=1:size(right_hand_veridical, 2)
    right_hand_veridical(1, idx).TRIAL_NUM =
right_hand_veridical(1,idx).TRIAL.TRIAL_NUM;
end
right_hand_veridical = nestedSortStruct(right_hand_veridical, 'TRIAL_NUM');

for idx=1:size(right_hand_low, 2)
    right_hand_low(1, idx).TRIAL_NUM = right_hand_low(1,idx).TRIAL.TRIAL_NUM;
end
right_hand_low = nestedSortStruct(right_hand_low, 'TRIAL_NUM');

for idx=1:size(right_hand_high, 2)
    right_hand_high(1, idx).TRIAL_NUM =
right_hand_high(1,idx).TRIAL.TRIAL_NUM;
end
right_hand_high = nestedSortStruct(right_hand_high, 'TRIAL_NUM');

for idx=1:size(right_hand_unlimited, 2)
    right_hand_unlimited(1, idx).TRIAL_NUM =
right_hand_unlimited(1,idx).TRIAL.TRIAL_NUM;
end
right_hand_unlimited = nestedSortStruct(right_hand_unlimited, 'TRIAL_NUM');

% Select data window (i.e Trial 0-100 shown here)

idx = [right_hand_veridical(:).TRIAL_NUM]>0 &
[right_hand_veridical(:).TRIAL_NUM]<101;
right_data_veridical= right_hand_veridical(idx);

idx = [right_hand_low(:).TRIAL_NUM]>0 & [right_hand_low(:).TRIAL_NUM]<101; %
Those in right_data_low= right_hand_low(idx);

idx = [right_hand_high(:).TRIAL_NUM]>0 & [right_hand_high(:).TRIAL_NUM]<101; %
Those in
right_data_high= right_hand_high(idx)

idx = [right_hand_unlimited(:).TRIAL_NUM]>0 &
[right_hand_unlimited(:).TRIAL_NUM]<101; %
right_data_unlimited= right_hand_unlimited(idx);
```

A.3.3 Unpack desired data field

```

function UndesignatedField = UnpackField(FieldNames,data_in)

%%Unpack required field "FieldNames" from c3d array "data_in",

fieldData = {};

for n=1:length(data_in)
    fieldData{n} = getfield(data_in(n), FieldNames);
end

minLen = length(fieldData{1});
for n=2:length(data_in)
    minLen = min(minLen, length(fieldData{n}));
end

arr = fieldData{1};
matrixVersion = arr(end-minLen+1:end);
for n=2:length(data_in)
    arr = fieldData{n};
    maxLen=length(arr);
    deltaLen=maxLen-(minLen);

    if deltaLen < 1
        arr=arr;
    else
        arr= arr(end-minLen+1:end);
    end

    matrixVersion = [matrixVersion, arr];
end

UndesignatedField = matrixVersion;

```

A.3.4 Unpack reach coordinates and velocity profiles

```
%% Calculate reach trajectories for each visual condition type

%sigma_1 data
RH_X_veridical = UnpackField('Right_HandX', right_data_veridical);
RH_Y_veridical = UnpackField('Right_HandY', right_data_veridical);

%sigma_m data
RH_X_low = UnpackField('Right_HandX', right_data_low);
RH_Y_low = UnpackField('Right_HandY', right_data_low);

%sigma_L data
RH_X_high = UnpackField('Right_HandX', right_data_high);
RH_Y_high = UnpackField('Right_HandY', right_data_high);

%Sigma_0 data
RH_X_unlimited = UnpackField('Right_HandX', right_data_unlimited);
RH_Y_unlimited = UnpackField('Right_HandY', right_data_unlimited);

%% Calculate Velocity
RH_vel_veridical = UnpackField('Right_L1Vel', right_data_veridical);
RH_vel_low = UnpackField('Right_L1Vel', right_data_low);
RH_vel_high = UnpackField('Right_L1Vel', right_data_high);
RH_vel_unlimited = UnpackField('Right_L1Vel', right_data_unlimited);
```


A.3.5 Apply position and velocity criteria to calculate endpoints

```
%% Endpoint Capture (minimum velocity criteria) [Example for sigma_1 only]

%Find the maximum velocity within the trial vector
maxima_veridical=squeeze(max(RH_vel_veridical,[],1));

%populate an array of possible endpoints @ 5% max velocity
threshold_veridical=0.05*maxima_veridical;

for i=1:length(maxima_veridical)
    maxima_veridical_IDX(i)=min(find(ismember(RH_vel_veridical(:,i),maxima_veridical(i))));
end

for i=1:length(threshold_veridical)
    if
        isempty(find(RH_vel_veridical(maxima_veridical_IDX(i):end,i)<threshold_veridical(i),1,'first'))
            fivepc_veridical_IDX(i)=0;
        else

            fivepc_veridical_IDX(i)=find(RH_vel_veridical(maxima_veridical_IDX(i):end,i)<threshold_veridical(i),1,'first')+maxima_veridical_IDX(i)-1;
        end
    end

%find min velocity endpoint after 20cm reach (y>34cm)
logical_veridical=RH_Y_veridical>0.34;

for i=1:size(logical_veridical,2)
    if isempty(find(logical_veridical(:,i),1,'first'))
        y_veridical_IDX(i)=999;
    else
        y_veridical_IDX(i)=find(logical_veridical(:,i),1,'first');
    end
end

index_veridical=fivepc_veridical_IDX(fivepc_veridical_IDX > y_veridical_IDX);
idx_veridical=(fivepc_veridical_IDX > y_veridical_IDX);
X_veridical_temp=RH_X_veridical(:,idx_veridical);
Y_veridical_temp=RH_Y_veridical(:,idx_veridical);

for i=1:size(Y_veridical_temp,2);
    Y_veridical(i)=Y_veridical_temp(index_veridical(i),i);
end
Y_veridical=Y_veridical';

for i=1:size(Y_veridical_temp,2);
    endpoints_veridical(i)=X_veridical_temp(index_veridical(i),i);
end
endpoints_veridical = endpoints_veridical';
```

A.3.8 Unpack the imposed shift, calculate endpoints and cursor position

```
%% Calculate X_Shifts
shift_veridical=UnpackField('XSHIFT', right_data_veridical);
shift_veridical=shift_veridical';
shift_veridical=shift_veridical(:,end);
shift_veridical = shift_veridical(idx_veridical);

shift_low=UnpackField('XSHIFT', right_data_low);
shift_low=shift_low';
shift_low=shift_low(:,end);
shift_low = shift_low(idx_low);

shift_high=UnpackField('XSHIFT', right_data_high);
shift_high=shift_high';
shift_high=shift_high(:,end);
shift_high = shift_high(idx_high);

shift_unlimited=UnpackField('XSHIFT', right_data_unlimited);
shift_unlimited=shift_unlimited';
shift_unlimited=shift_unlimited(:,end);
shift_unlimited = shift_unlimited(idx_unlimited);

%% Calculate reach endpoints
endpoints_veridical = [];
endpoints_low = [];
endpoints_high = [];
endpoints_unlimited = [];

endpoints_combined;

mean_endpoints_veridical=mean(endpoints_veridical);
mean_endpoints_low=mean(endpoints_low);
mean_endpoints_high=mean(endpoints_high);
mean_endpoints_unlimited=mean(endpoints_unlimited);

%% Calculate cursor position (endpoint + shift)
cursor_position_veridical = endpoints_veridical + shift_veridical;
cursor_position_low = endpoints_low + shift_low;
cursor_position_high = endpoints_high + shift_high;
cursor_position_unlimited = endpoints_unlimited + shift_unlimited;

mean_cursor_veridical = mean(cursor_position_veridical);
mean_cursor_low = mean(cursor_position_low);
mean_cursor_high = mean(cursor_position_high);
mean_cursor_unlimited = mean(cursor_position_unlimited);
```

A.3.6 calculate and plot cursor error as a function of the imposed shift

```
%Calculate and plot cursor error as a function of applied shift
%Exemplar - Sigma_1 (repeated for all visual conditions)

% bin data, calculate mean of bins [11 bins];
[N,edges,bins] = histcounts(shift_veridical(:,1),11);
X_vals_veridical = accumarray(bins(:),shift_veridical,[],@mean);
Y_vals_veridical = accumarray(bins(:),cursor_position_veridical,[],@mean);
e_veridical=accumarray(bins(:),cursor_position_veridical,[],@std);
P_veridical = polyfit(X_vals_veridical,Y_vals_veridical,1);
yfit_veridical = P_veridical(1)*X_vals_veridical+P_veridical(2);

trend_cursor_binned_veridical = fit(-X_vals_veridical,-
Y_vals_veridical,'poly1');

figure
scatter(-X_vals_veridical,-Y_vals_veridical);
hold on;
plot(-X_vals_veridical,-yfit_veridical, 'r-.');
ylabel('Deviation from target (cursor poition)');
xlabel('true veridical shift values');
title ('Deviation from target as a function of true lateral Shift (sigma_1)');
axis([-0.01 0.02 -0.02 0.02]);
errorbar(-X_vals_veridical,-Y_vals_veridical,-e_veridical);

%calculate slope of trendline
slope_cursor_veridical = trend_cursor_binned_veridical(1);
intercept_binned_cursor_veridical = trend_cursor_binned_veridical(2);
trendline_cursor_veridical=
trend_cursor_binned_veridical(1)*shift_veridical+trend_cursor_binned_veridical
(2);
disp(['veridical slope' ' ' (num2str(slope_cursor_veridical)) ' ' 'metres'])
```