THE ROLE OF CEREBELLUM IN ACTION ACQUISITION AND MOTOR LEARNING

by

MARIA DAGIOGLOU

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

DOCTOR OF PHILOSOPHY

School of Psychology

College of Life and Environmental Sciences

University of Birmingham

February 2014

ABSTRACT

The aim of the present thesis was to investigate the role of cerebellum in motor learning and action acquisition. This question was pursued by means of behavioural studies on healthy population. In a first study, the role of cerebellum in motor skill learning was explored by perturbing cerebellar activation with transcranial direct current stimulation. The involvement of cerebellum in action acquisition was studied in a paradigm that combined a visuomotor tracking task and an exploration task. The results of this study lead to chapter 4, where we investigated the impact of the tracking task in proprioceptive uncertainty. In a final study, the role of cerebellum, motor cortex and dorsolateral prefrontal cortex in action acquisition were investigated by modulating these brain areas using transcranial direct current stimulation. The results suggested that the cerebellum could be contributing in motor learning, not just by providing a state estimation but also by providing the uncertainty related to the estimates. However, based on the results of the final experimental chapter, we can conclude that, at least in the framework of the exploration task, motor cortex is more heavily involved than the cerebellum, perhaps via the cortico–basal–ganglia pathway, in reinforcement learning.

To Antonia, Dimitris and Stergios

AKNOWLEDGMENTS

I would like to express my gratitude to my supervisor Prof. Chris Miall for giving me the opportunity to pursue a PhD and also to thank him for mentoring me, teaching me research integrity and encouraging intellectual adventurousness.

The work in this thesis has been supported by the European Community 7^{th} Framework Programme. My PhD was part of the Initial Training Network: Cerebellar–Cortical Control: Cells, Circuits, Computations and Clinic (C^7). It has been a great and stimulating experience.

I would like to acknowledge Samuel Westwood and Alexandra Doherty for the data collection of the study presented in Chapter 5 and to thank them for our collaboration.

I would like to offer my special thanks to Prof. Jeremy Wyatt for all the stimulating discussions we had.

I would like to thank my colleagues in PRISM lab for the collaboration, the support and the discussions we had all these years. Special thanks to Jonathan Winter for all his help and support in technical issues.

I would also like to thank all my teachers and professors who have inspired me over the years. Special thanks to Prof. Maarten Frens for introducing me to the world of cerebellum.

I am very grateful to my friends in Birmingham who supported me all these years in various ways. Special thanks also to my friends abroad who have been there in every possible way.

Last but not least, I want to thank my family for supporting me, encouraging me and believing in me. Without them this PhD would never have happened!

NOTE

The entire Section 1.2 except from Subsection 1.2.2 was copied from the following publication:

Hardwick RM, Dagioglou M, Miall R (2013), State estimation and the cerebellum In Manto M, Schmahmann J, Rossi F, Gruol D, Koibuchi N, editors, Handbook of the Cerebellum and Cerebellar Disorders, pp. 1297–1313. Springer Netherlands.

The text copied corresponds to paragraphs: Introduction and Internal Models in Hardwick et al. (2013) and was my personal contribution to the publication.

Contents

List of Figures

List of Tables

Chapter 1
MOTOR FUNCTIONS: FROM CONTROL TO LEARNING

1.1 Introduction

The purpose of the present thesis was to investigate the involvement of cerebellum in motor learning processes that have not been traditionally related to the structure. The aim of this first chapter is to set the theoretical background that motivated the experiments in Chapters 2–5. Starting from the involvement of cerebellum to motor control, I move on by examining the role of cerebellum in motor learning

Before talking about motor learning, in the first section I begin by presenting several aspects one must consider when thinking about how the brain controls movements. The ability of the brain to control fast and accurate movements lies to a great extent in the existence of *internal models* that are able to predict future states of the body. Internal forward models are the topic of Section 1.2.1. Finally, in Section 1.2.2 I present experimental evidence that support the hypothesis that the cerebellum is the site of the brain where forward models are implemented.

In Section 1.3 the various aspects of learning are discussed. I look at hypotheses related to learning and how these are investigated through behavioural paradigms. I first briefly talk about information gathering and memory components related to sensorimotor learning. In Section 1.3.3, error–based, reinforcement and use–dependent learning processes are introduced. Error–based learning is the most well–studied process of learning and several issues related to it are presented via three categories of paradigms; force–field adaptation, visuomotor adaptation and adaptation to temporal delays. Following, I discuss the computational components of reinforcement learning and recent studies that have dealt with this kind of learning in motor tasks. Finally, in the last paragraph of Section 1.3 I briefly review the literature related to

use–dependent learning.

The cerebellum has often been related to error–based learning and the basal ganglia to reinforcement learning. In Section 1.4, I first review evidence of the involvement of the cerebellum in error–based learning. I then briefly discuss experimental evidence that support the contribution of the basal ganglia in reinforcement learning generally without confining to literature that involves motor tasks. In Section 1.4.4, studies that focused in differentiating the roles of the two structures in motor control and learning are discussed. Finally, anatomical evidence that connect the two areas and the functional implications of this connectivity are discussed.

1.2 Motor Control

Every day we perform a variety of movements, from very simple ones like pointing to a blackboard to very accurate and elegant ones like playing a musical instrument. The high performance of our motor system is viable because the central nervous system (*CNS*) controls the *biomechanical plant*[1] that is composed of the skeleton and the muscles attached to it. The cascade of events that underlies every single movement we perform is triggered when the CNS sends a signal (*a motor command*) to the muscles. The motor command causes a series of electrochemical reactions that result in the contraction of muscle fibres. The change in muscle length imposes forces on the skeleton that cause a new position and velocity of the limbs. The CNS is informed about the new position of the limbs by the sensory systems. Muscle length and velocity, as well as tension in the muscles are recorded by proprioceptive sensors (muscle spindles and Golgi tendon organs, respectively). Because movements of the limbs are accompanied by stretching or bending of the skin and contact with

[1] In terms of control theory, the term *plant* is used to describe a process and the controlled object (actuator). For example, the eye and its supporting tissues are often called *oculomotor plant* (Frens and Donchin, 2009)

external surfaces (particularly true for the fingers), tactile sensors add knowledge to the CNS about the status of the periphery. The CNS can be further updated about changing limb position by the visual system and, in some circumstances, by the auditory system (Boyer et al., 2013). The information these sensory systems send back to the CNS reflects not only the results of the motor commands but also any external (mechanical or visual) perturbations imposed on the limbs.

The ability of the CNS to control the biomechanical plant is based on the knowledge of its state at any moment. The *state* of the motor system can be defined as a set of variables (e.g., the position and velocity of different limb segments), the knowledge of which at any moment will allow the CNS to completely determine the future behaviour of the motor system, given that inputs to the system (all motor commands, any internal dependencies and any external perturbations) are known as well. So, given accurate knowledge of the current state of efferent motor commands and full knowledge of the behaviour of the biomechanical plant, the sequence of events described in the previous paragraph would allow the CNS to estimate the state of the motor system in the near future and hence control it, assuming there are no external perturbations. However, the reliability of this estimation is severely challenged by several factors. To begin with, it takes tens of milliseconds for afferent signals to be transmitted from the periphery to the CNS. Additional central processing delays add to these transport delays, leading to outdated central knowledge about the peripheral system. Moreover, in all physiological systems, signals are corrupted by neural noise both in the afferent and efferent paths. Delays and noise jeopardise the ability of the CNS to control movement and leave it exposed to inaccuracies and instability. Finally, the plant may change its behaviour in unpredictable ways, for example through fatigue, or hysteresis, or similar effects.

1.2.1 Internal Forward Models

It has been suggested and widely supported by experimental evidence that the CNS counteracts the consequences of delays and noise by being able to calculate a state prediction, which is an estimate of the state of the motor plant available before any sensory information is fed back by the periphery. State prediction is a product of forward internal models (Wolpert and Miall, 1996). Internal models are, in general, central neural representations of the motor system (Wolpert et al., 1998). Internal forward models capture the relationship between the input and the output of the biomechanical plant (the forward or causal relationship between motor commands and the new motor states), while internal inverse models capture the inverse relationship (the relationship between a desired new state and the motor commands necessary to achieve it). While the CNS can use inverse models as controllers to issue motor commands that drive a desired change in state, forward models are suitable for state prediction. What a state prediction embodies depends on the aspect of the biomechanical plant the forward model captures. A forward dynamic model is a representation of the mechanical properties of the biomechanical plant computing, for example, joint angles and velocities given the forces and torques applied to the limbs. A forward sensory output model gives as its output predictions of the different sensory modalities (tactile signals, proprioception, vision, etc.) capturing important properties of the biomechanical plant transformed by the receptive properties of the sensors. Hence, for example, joint angles are better encoded by the muscle spindles, albeit in a complex form, than by joint receptors, and so a sensory output model might predict spindle firing patterns rather than joint angles. This is often called sensory prediction. It is still not apparent how the brain deals with the different modalities of sensory representations (Haggard and Wolpert, 2005).

The CNS can benefit from state prediction in various ways. In the case of very short movements, such as saccades, sensory feedback cannot play an influential role due to intrinsic transmission delays. In this case control becomes the sole responsibility of the feedforward pathway, based on state prediction, confining the role of sensory feedback to keep the forward model updated about any time or history dependent changes in the plant (i.e., fatigue, ageing, etc.). On the other hand, limb movements are slow enough to allow sensory information to influence control. Were limb control to depend on the delayed sensory feedback alone, movements would have to be performed slowly or else instabilities would arise. In this context, the benefits of state prediction are exposed in rapid and skilled motor behaviour; such fast and accurate performance could not be achieved using feedback control. Another important use of state prediction is to distinguish between external and self–induced components of sensory inputs. When an afferent signal is transmitted back to the CNS, it contains two kinds of information. Part of the afferent signal (sometimes called *exafference*) is the result of any external perturbations applied to the limbs. The other part (which is called *reafference*) depicts the sensory outcome of our own actions, the result of the motor commands sent to the muscles (Wolpert et al., 1998). Predicting the sensory consequences of a motor command enables distinction of reafference. Any sensory discrepancy between a sensory prediction and an afferent signal should mirror an external perturbation (the exafference) and could potentially signal an error in performance that should be corrected. Alternatively, this sensory predictive mechanism can be used to suppress self-induced reafferent signals. Reafference is intrinsically uninformative, as it reflects the predictable consequences of motor commands, and so should not draw attention from or occupy processing resources. Sensory suppression is therefore valuable to avoid unnecessary central processing of reafferent signals (Blakemore et al., 1998). Thus, state predictions based on forward

models contribute to sensorimotor systems at several levels, from sensory processing to control.

Several control models that include forward models have been proposed to explain human motor performance. Broadly speaking, two different general frameworks can be distinguished. In the first, a forward model is used to calculate the current state, and errors between the state and a target desired trajectory are corrected by feedback control. The difference between the desired trajectory and the state prediction is used to drive the controller. When sensory feedback is available (after a feedback delay) the sensory prediction is compared with actual sensory feedback and any discrepancy (error signal) in this comparison drives further the output of a controller. An example of such control schemes is the *Smith Predictor* (Miall et al., 1993). Under this scheme, forward models can contribute to effectively compensate for feedback delays. However, the Smith Predictor control scheme is challenged if the dynamics of the biomechanical plant change. The forward model, unless updated to reflect the new dynamics, fails to accurately predict the future states, and so the feedback error signal that drives the controller becomes inaccurate. Bhushan and Shadmehr (1999) therefore proposed a model where both forward and inverse models are used to control movement. In this case, the error signal drives the output of an inverse model, which acts as an adaptive controller. They argue that this provides a better match to human movement data, especially when a robotic arm was used to impose forces upon the biomechanical plant, changing the dynamics of this combined system while the subject was moving toward the target.

The alternative general framework used to explain motor control is the optimal control theory (Todorov, 2004). Here, the aim is to minimise a cost function, and the task for the controller is to optimise motor commands to achieve this (Scott, 2004; Diedrichsen et al., 2010b). Under this framework, it is a state estimate that is

the input to the controller (rather than an error signal). The forward model provides the state prediction that is combined with sensory feedback to produce an optimal state estimation. A Kalman filter (Wolpert et al., 1995) is proposed to be used to calculate the state estimate by assigning weights to each source of information (state prediction and sensory feedback) based on the reliability of each. If sensory feedback is unreliable (e.g., visual feedback of the arm is obviously unreliable in the dark), then the state estimation can be based mainly on the state prediction. On the other hand, if the forward model's output is judged to be inappropriate or unreliable, then a higher weight can be attributed to the sensory feedback. It should be noted here that while both state prediction and state estimation are estimates of the state of the motor apparatus, they are the products of different procedures. State prediction is the outcome of the forward model, while state estimation is the combination between state prediction and sensory afference.

This description of the optimal weighting of prediction and sensory inputs gives us an appropriate point to introduce Bayesian statistics. Neuronal noise in sensory signals, the noise affecting motor commands resulting in state changes that deviate from the intended ones, and the constantly changing world produce uncertainty in the observations made. State predictions and state estimates in the frameworks described above are usually modelled as if taking discrete values, as they are usually the outcome of solving a system of differential equations (state space equations). The uncertainty due to noise translates the problem of calculating state estimates and predictions into a statistical inference problem (Bays and Wolpert, 2007). Bayesian statistics are used to assign a probability to each estimate and each prediction. All these probabilities are described by Gaussian distributions; by using Bayes rule we can calculate the probability of being in a particular state (a state estimation an *a posteriori* probability) given a particular observation, and given both the probability

of the observation acquired via sensory feedback and the probability of being in this state given the previous state and the action taken. Suppose we reach out to grasp and lift a carton of juice from a tabletop. At this time, we have to make certain judgments based on our prior experiences. In order to lift the carton, we must produce enough lifting force to overcome the mass of the carton. In addition, we need to create grip forces against the carton to avoid it slipping between our fingers; the grip required depends on the weight of the carton, on the friction of its surface (and our fingers), and the amount of lifting force we exert on it. To perform this action, we use prior experience of cartons to make judgments about their probable properties. However, you have probably experienced instances in which you expect the carton to be full, but it is actually empty, and you lift it far faster than intended. This mistaken action provides an example of a situation in which you use a false prior belief to plan your action. Moreover, in this example, the likelihood of a new observation of the mass of the carton results from two distinct modalities; proprioception and vision. These two cues can also be optimally combined using Bayesian inference, according to which each modality is weighted based on reliability. For example, if we lift the carton while wearing gloves, our proprioception is constrained and thus less reliable, leading to a reduction of its weighting, compared to vision. Conversely, if we lift the carton in the dark our vision becomes less trustworthy, and its weighting is reduced.

Returning to the two general control schemes outlined above, the robustness of control depends highly on the adaptive tuning of all components (the forward model, the state estimation, and the controller) to any changes in the dynamics of the biomechanical plant which can be caused by either intrinsic factors (i.e., fatigue, ageing, etc.) or extrinsic factors (e.g., the mass of a carried tool, or external applied loads, etc.). It has been suggested that the adaptation of forward models is achieved using sensory prediction errors – the discrepancies between sensory feedback and

sensory predictions (Shadmehr et al., 2010). However, the source of errors should be taken into consideration in appropriately adapting the forward model to achieve accurate state prediction (Berniker and Kording, 2008). External events that lead to unpredictable errors should not lead to changes in the forward model, whereas events that lead to predictable errors imply the model is not accurate and should be adjusted. This identification of the source of errors seems to become highly important for the generalisation of learning following adaptation, and bridges the gaps between different theories about how inter-limb and intra-limb generalisations occur and whether the internal models give an output analogous to intrinsic or extrinsic coordinate frames.

1.2.2 The role of Cerebellum in motor control

The cerebellum has been identified as the brain area that retains forward models. In this paragraph I present studies of cellular recordings, fMRI and Transcranial Magnetic Stimulation (TMS) that provide evidence of the cerebellar involvement in motor control via forward models. There is also plenty of evidence from patient and fMRI studies that show the contribution of cerebellar forward models in motor adaptation. This line of evidence is going to be separately discussed in the section about the role of cerebellum in motor learning.

Purkinje–cell simple spike activity in the cerebellar cortex has been found to be related to state prediction. Pasalar et al. (2006) recorded simple–spike firing of task–related Purkinje cells in two highly trained monkeys. The task involved the execution of circular manual tracking with a robotic manipulandum while force fields of various loads were applied. The type of the force field could be either viscous or elastic. Muscle activity was also recorded during the task. The results revealed that the hand forces and the EMG recordings reflected the magnitude and direction

of the force loads. On the other hand, simple–spike discharge of Purkinje cells in cerebellar lobules IV–VI was not significantly modulated by the force field type and the magnitude of load. These results were interpreted as evidence supporting the idea that Purkinje cells represent kinematic characteristics of the arm movements rather than an inverse dynamics model. However, the authors suggested that their results should not be generalised to areas other than the recorded cerebellar area. Liu et al. (2003) reported similar data for a visuomotor perturbation.

The findings of Pasalar et al. (2006) and Liu et al. (2003) support the idea of existence of forward models in the cerebellum that represent own body movements. A question that arises often is whether there are also forward representations of external objects of the environment. Cerminara et al. (2009) trained cats to perform a predictable visually guided reaching task. After 6-8 weeks of training period, Purkinje cells in Crus I were recorded during the performance of the task and their activity showed tonic simple spike activity that was related to the movement of the target and not to limb movement or eye movements. The tonic activity was sustained even when the target was temporary occluded which means that the cells recorded could predict the movements of the target even in the absence of the stimulus, providing evidence of the existence of forward models of external objects in the cerebellum.

Whether the cerebellum retains forward dynamic (arm movement properties) or a forward sensory output (sensory receptors properties) models or both is not clear yet (Ebner and Pasalar, 2008). The studies of Pasalar et al. (2006) and Roitman et al. (2005) do provide some evidence that simple spike activity predicted upcoming motor state rather than sensory reafferent signals.

Behavioural evidence in a Transcranial Magnetic Stimulation (TMS) study provided further evidence of the contribution of cerebellum in state estimation. In Miall

et al. (2007), TMS was used to disrupt ipsilateral cerebellar function during a fast reaching movement to a target. A short TMS pulse was delivered upon initiation of the reaching movement and caused directional deviation of the movement and increased end point error. Based on the average directional errors it was inferred that the reaching movements were planned based on around 138 ms out–of–date information of the state of the arm. So, TMS disruption of the cerebellum interfered with accurate state estimation.

Prediction signals are essential for visuomotor coordination. Miall et al. (2000) showed that that peak performance in a visuomotor task was achieved, if the eyes were following the same target trajectory with the hand with 75–100ms lead. These results were interpreted as evidence of predictive information being used; information coming form the ocular system could be used to improve hand tracking. In a subsequent imaging study, Miall et al. (2000) showed that cerebellum was active during visuomotor co–ordination and that it was also activated during large tracking errors.

1.3 Sensorimotor Learning

Very generally one could say that sensorimotor learning involves all the processes that need to occur in order to achieve sensorimotor control. When studying the literature about sensorimotor learning, it soon becomes clear that sensorimotor learning is a blanket term that could be used to describe quite different and complicated processes. We could distinguish three aspects/phases of learning. The first one, involves all the processes that are related to making task–relevant observations of the environment – the sensory information gathered before selecting an action. The second aspect of learning entails the processes that use the sensory information to select an action – the "neural algorithms" implemented by the brain to achieve the

goal. Finally, a third phase of sensorimotor learning includes storing the new–gained knowledge – the changes in motor and sensory plasticity as a result of learning that lead to improvement in future actions.

In the current literature survey I aim to only briefly discuss the first and third phases of sensorimotor learning (as described above) and give the main focus to the "neural algorithms" the brain might be using to induce changes in performance. In each case I consider both the computational frameworks (see for example Haith and Krakauer, 2013; Wolpert et al., 2011; Krakauer and Mazzoni, 2011) used to explain the behavioural evidence, and the anatomical areas and connections involved (see for example Bostan et al., 2013; Shmuelof and Krakauer, 2011; Doya, 2000).

1.3.1 Sensory Information and Knowledge of Facts

Sensorimotor learning begins at the moment we start exploring the environment in order to choose an action for achieving a goal. Given the incomplete knowledge of the environment the observations we make are uncertain. So the first step towards learning is to learn how to observe the environment in order to reduce the uncertainty (Rao, 2010). In other words, we need to learn how to use our sensors to maximise the quality of information gathered and extract only the appropriate task–specific features (Wolpert and Flanagan, 2010).

Decreasing the uncertainty over the states of the body and of the world is important for *credit assignment*. Credit assignment refers to identifying and estimating correctly the right source of motor error (Berniker and Kording, 2008). Consider the lifting the carton box example that we used before. Suppose that you attempt to lift the box and it slips from your hands. Several explanations are possible: either the mass of the box was more than expected or the grip force applied to the box was less than expected because of altered intrinsic properties of the body (e.g. fatigue)

or the friction was different (lower) than expected. Assigning appropriate credit to the potential sources of error is necessary for learning and generalising knowledge (Berniker and Kording, 2008, 2011).

While the aforementioned processes can occur implicitly, action selection can also depend on applying an explicit/cognitive strategy (Wolpert and Flanagan, 2010). Stanley and Krakauer (2013) make a distinction between "motor acuity", the refinement of movements over practice, and acquisition of a skill which involves knowledge of facts. Even when experts have achieved an automatic behaviour:

> "...a continuing symbiosis or bootstrapping between knowledge and non–knowledge is what we propose leads to greater skill overall." (Stanley and Krakauer, 2013)

So, sensory information and knowledge of facts will enable different action–related processes from action acquisition to action selection and movement refinement. These processes are proposed to be driven by different computational mechanisms and are related to distinct brain areas (Doya, 2000; Shmuelof and Krakauer, 2011; Redgrave et al., 2008).

Perhaps what one ultimately wants to be able to answer is what are the learning processes deprived in disease. Being able to better understand the learning mechanisms and how they interact will lead to important improvements of the existing rehabilitation methods (Bastian, 2008).

Before discussing in detail about the learning processes, we first review some aspects of the consolidation of motor learning.

1.3.2 Consolidation of Motor Learning

The acquisition of motor learning is a distinct process from the processes of retention and consolidation of learning. This is supported by evidence suggesting that

separate neural substrates are involved in each case. For example, it was shown in a transcranial Direct Current Stimulation (tDCS) study that during visuomotor adaptation the motor cortex retains what the cerebellum learns (Galea et al., 2011).

Consolidation of motor learning is used to describe both off–line improvement of the motor skill practiced but also the stabilisation of memories (Robertson et al., 2004). Cohen et al. (2005) suggested that off–line learning is supported by different mechanisms (at least in the framework of serial reaction time task): goal–based overnight improvement and movement–based over day improvement and that the engagement of this mechanisms depends on when consolidation takes place. Memory stabilisation refers to the resistance of motor skill acquisition to the interference of practising another new skill (Robertson et al., 2004).

The retention of learning is exposed behaviourally by being able to learn faster or starting from a better performance level when being exposed to a previous practised task, a phenomenon which is also called *savings*. Another phenomenon related to memory of learning is *spontaneous rebound*; if an extended period of adaptation is followed by a brief period of reverse–adaptation, in the subsequent error–clamped trials there is a recovery of the behaviour of the initial form of adaptation (Smith et al., 2006; Ethier et al., 2008).

Smith et al. (2006) have tried to capture learning after–effects, savings, interference, spontaneous rebound during motor adaptation paradigms in a computational model that entails two memory components. According to this model short–term motor adaptation is driven by a fast component which is sensitive to performance error and has poor retention, and a second slow component which is not sensitive to error but has good retention.

The various memory components are related to different brain processes and areas. Keisler and Shadmehr (2010) showed that the fast process in adaptation to

force perturbations was disrupted by a cognitive task that engaged the declarative system. However, the slow process remained intact. Moreover, impairing the fast component of learning, by executing the cognitive task, lead to to an enhanced slow (non declarative) process component after a six hours consolidation period. This result suggests that fast and slow memory components are processed differentially and points towards the direction of existence of an inhibitory relationship between the two of them.

1.3.3 Learning Processes

Terms from machine–learning have been used to describe the different mechanisms the brain might be using for sensorimotor learning. Depending on the learning signal we can distinguish the following processes: error–dependent supervised learning, reward–dependent reinforcement learning and repetition–based unsupervised learning (see for example Doya, 2000; Wolpert et al., 2011, 2001).

Error–Based Learning

In supervised or error–based learning an error between an observation and a predicted sensory signal drives learning (Shadmehr et al., 2010). A crucial characteristic of supervised learning is that the observed discrepancy does not only signal an error but also the direction of the corrective movement. The existence of forward models, that can predict the sensory consequences of the movements, are crucial for the computation of this error (Wolpert and Miall, 1996). In presence of a sensory prediction error or a motor error, the movements are adapted to minimise any discrepancies. In addition to motor errors, discrepancies across sensory modalities can drive sensory recalibration (Henriques and Cressman, 2012). The term sensorimotor adaptation has been equated to supervised–learning (Bastian, 2008). As discussed in Section 1.2.2, there is evidence that the cerebellum retains internal forward models,

and so it plays a crucial role in error–based learning.

To study sensorimotor adaptation in the lab, visual, proprioceptive or temporal perturbations that give rise to sensorimotor errors are introduced. Adaptation of movements to counteract the effects of the errors in these paradigms occurs fast, within minutes or hours. The degree of adaptation is usually evaluated by sudden removal of the perturbation which causes errors to the opposite direction, known as aftereffects. Finally, generalisation of learning to new locations of the workspace and to new tasks, as well as transfer to the opposite (untrained) limb are often tested to expose the nature of what has been learned. Examples of sensorimotor adaptation paradigms include saccadic adaptation (Jenkinson and Miall, 2010), vestibulo–ocular and optokinetic reflex adaptation (Frens and Donchin, 2009), visuomotor adaptation in the context of prism adaptation (Hatada et al., 2006) and of virtual environments (Mazzoni and Krakauer, 2006), adaptation to altered dynamic environments (Lackner and DiZio, 2005) and adaptation to delays (Miall and Jackson, 2006).

In the following paragraphs, I discuss several aspects of error–based learning through behavioural studies. I will focus on paradigms of adaptation to force field perturbations, to visuomotor rotations and to delays.

Adaptation to Dynamic Perturbations Adaptation to dynamic disturbances has been studied by applying motion–dependent force field perturbations via a manipulandum (Shadmehr and Mussa-Ivaldi, 1994). Initial application of forces initially affects reaching movements but over time participants learn to compensate the effect of forces and achieve performance similar to free–space hand trajectories. Sudden removal of the force field reveals *aftereffects* that highlight how the nervous system has been building a model to predict the force disturbances. Moreover, such model is not associated solely with the reaching movements to the targets that were vis-

ited during training but generalises across other regions of the workspace as well. The adaptation of a forward model has been suggested to play a dominant role in the error–based motor learning that underlies this process (Bhushan and Shadmehr, 1999).

A critical aspect of adaptation (relevant to all kinds of adaptation) is the predictability of perturbations. When there is high uncertainty about the nature of the perturbations then the central nervous system uses non–specific motor learning (Wei et al., 2011). By applying perturbations different in nature (both different force fields and visuomotor rotation trials) in random order, Wei et al. (2011) found that corrective movements in the following trial did not depend on the nature of the perturbation. The authors suggested that perturbation of varying nature that lack specificity or predictability in trial-by-trial adaptation could suggest that in order to form an accurate internal model, a consistent exposure to the same type of perturbation might be necessary. However, in reality the properties of the body do not change randomly – they can drift and be subject to random error in their measurement but the properties have predictable features. This observation implies that in order to build an accurate forward model for state prediction in a task there must be a level of certainty about the source of error and thus the state of the body, as opposed to that of the world.

When adapting to force field perturbation using a robotic arm, there is a question whether a new tool–specific model is built or an existing body–specific forward model is adapted. According to Kluzik et al. (2008) this depends on whether the error is assigned to one's own body or to the environment. In their study, participants performed reaching movements and adapted to force perturbations. The post–adaptation effect (generalisation) was tested under two different conditions; while holding the robot handle rigidly attached to the robotic arm (robot–null) and

while holding a robot handle which was detached from the robotic arm (free–space). In an altered adaptation condition the force perturbations were introduced gradually (as opposed to an abrupt presentation of the perturbation). The results revealed that adaptation to a novel tool generalises to free space and that this generalisation is more pronounced if the force field is introduced gradually. A possible explanation of this result is that in the gradual case the central nervous system could have identified the internal representation of the limb dynamics as a primary source of errors. On the other hand, for the abrupt case, most of the errors could have been attributed to an inadequate representation of the novel tool or the external environment.

Adaptation to force field perturbations is manifested through the adjustment of movements to counteract the applied disturbance. This change in movements provides evidence of motor learning, i.e. learning what movement is appropriate to use in order to achieve a goal (e.g. move in a straight line). In addition to motor learning, there is evidence that during adaptation to force fields, changes in sensory perception also take place. Haith et al. (2008) have suggested that an optimal Bayesian model that takes into account both motor and sensory contributions to perceived errors is essential to explain motor performance after effects and perceptual shifts during visuomotor and force–field adaptations. Surprisingly what is predicted by such a model, and what was further confirmed experimentally by Haith et al. (2008), is that perceptual shifts occur even in the absence of discrepancies between sensory modalities (vision and proprioception). Ostry et al. (2010) also found a change in sensed limb position following adaptation to force fields. This change in proprioception was only observed in participants that actively adapted in the task and not to participants that were exposed to passive movements, so it was due to sensorimotor learning.

In this section, I presented several examples of force field adaptation paradigms.

The predictability of the perturbation (Wei et al., 2011) and the attribution of the source of error (Kluzik et al., 2008) which is manipulated experimentally by abrupt or gradual presentation of the perturbation are some of the aspects of the error that influence the rate, the degree and generalisation of the adaptation process. Moreover, adaptation entails both a motor component and a shift in perception (Haith et al., 2008). Following, I discuss similar findings coming from paradigms of visuomotor rotation.

Adaptation to Visuomotor Perturbations Visuomotor perturbations have been traditionally studied in the context of prism adaptation (Martin et al., 1996). Here I focus on paradigms where visuomotor adaptation is studied in a virtual environment. In these experiments only the visual feedback of the arm position is perturbed as opposed to a whole world visual shift in case of prism goggles. As the name suggests in the visuomotor adaptation paradigms a discrepancy is induced between the felt and seen arm position. So, for example, after reaching straight ahead to a target, visual feedback is presented shifted to the right or left. This results in an error between the target and the seen end–point of the movement. Over the trials participants alter their movements in order to reduce the visual error. Although force field adaptation and visuomotor rotation paradigms are similar in terms of inducing adaptation to an error, there is a prominent difference between the two. In the case of the visuomotor adaptation paradigms, the introduction of errors leads to discrepancies between vision and proprioception.

Adaptation to visuomotor rotations could rely on either a sensory prediction error, i.e. a discrepancy between the expected and observed visual feedback, or on a motor correction. Tseng et al. (2007) compared the performance in adapting to a visuomotor perturbation, in “shooting movements” and “reaching movements”. In

the former, participants shot through the target position but did not correct their movement, in the latter condition, participants were given the chance to correct their movements on–line. On–line corrections did not give an advantage in the adaptation of subjects suggesting that visuomotor adaptation relies on sensory prediction errors.

Although on–line movement corrections do not facilitate or inhibit visuomotor adaptation (Tseng et al., 2007; Shabbott and Sainburg, 2010), simultaneous visual and proprioceptive feedback during movement was found to be crucial for adaptation (Shabbott and Sainburg, 2010). When visual feedback is presented at the end of the movement adaptation only occurs on a trial–by–trial basis, whereas on–line feedback allows for within trial error corrections. Presence or absence of visual feedback during the movement seems to have no influence to the adaptation process, although mean direction error is smaller when visual feedback is provided. However, as shown in Shabbott and Sainburg (2010) the absence of feedback and just the presentation of error at the end of each trial lead to an increased inter–trial variability during rotation exposure and also deprived generalisation.

Adaptation in the presence of sensory prediction errors occurs even if the errors are not task–specific. In Schaefer et al. (2012), participants were trained to perform reaching movements to an arc or a ray while gain or rotation visual perturbations were applied. Thus the perturbations could be task–goal relevant or irrelevant. For example, when participants were targeting to an arc a visuomotor rotation would be task–goal irrelevant as the goal was to generate the right extent of the movement, while the direction was less important. The significant aftereffects in all conditions revealed that motor adaptation occurred even if the perturbations were irrelevant to the task goal, which supports the idea that sensory predictions errors drive motor learning. However, catch trial aftereffects were more prominent in the task-relevant conditions, which could imply that task–relevant and sensory prediction errors might

be necessary for full adaptation.

Sudden introduction of visuomotor rotations increases awareness about the error. Could there be a cognitive component affecting the adaptation process? Mazzoni and Krakauer (2006) tested this hypothesis and showed that a cognitive strategy interfered with the implicit adaptation process. Participants that had explicit knowledge of the visuomotor rotation and the strategy needed to counteract it (they were instructed to point to a neighbouring target), initially, were successful in cancelling out the effects of the visuomotor rotation but over time they started making increasingly large directional errors. Moreover, the rate of the implicit adaptation to the neighbouring target was found to be comparable with the rate of adaptation of the participants that performed the task without explicit knowledge of the rotation. Mazzoni and colleagues concluded that adaptation occurs solely in an implicit manner.

However, if the adaptation phase is extended, in presence of large directional errors, participants manage to change strategy and restore the functional benefit of the adaptive drift. Taylor and Ivry (2011) showed, by means of a computational model and behavioural data, that implicit adaptation is driven by the error between actual and predicted movement (sensory prediction error), whereas the difference between the aiming and target location (aiming error) affects explicit learning. Moreover, in the absence of a target (towards which the explicit strategy is applied), the strategy became more effective as no directional drift was observed. Taylor and Ivry (2011) concluded that when the uncertainty about the aiming error was high the adaptation based on this signal was diminished.

It has been discussed before, in the section presenting the force field perturbation paradigms, that what follows adaptation is both motor learning and changes in perception. If sensory prediction errors drive motor adaptation in visuomotor

paradigms, discrepancies between the visual and proprioceptive estimates have been suggested to produce sensory recalibration. In a recent review that summarises the results of their lab, Henriques and Cressman (2012) report that based on their studies the change in felt position is similar for both kinds of distortions (rotation and lateral displacement) and that the shift is proportional to the 20% of the size of hand–cursor discrepancy. The sensory recalibration between seen and felt hand position has been proposed to depend on two different components describing transformations between visual proprioceptive and body–centred spatial representations (Simani et al., 2007).

Even when correcting for distorted visual feedback is not the goal of the task, sensory recalibration and motor learning do occur. In Synofzik et al. (2006) participants made reaching movements to self–selected positions and during an exposure phase a rotation was imposed on the visual feedback. So, in the context of the study, participants were not presented with a clear error between a target and a distorted visual feedback and thus were not expected to adapt their motor commands to the perturbation. The results showed perceptual changes of the finger position towards the rotated visual feedback. These changes in perception were not dependent on the on–line visual feedback indicating that subjects used an updated internal information about the outcome of their movements. Although not expected, participants also showed motor adaptation in the opposite direction of the imposed visual rotation when pointing to specific target after adaptation (after the exposure phase). This suggested that perceptual adaptation might have driven motor learning. The fact that the perceptual adaptation was only partial perhaps uncovers certain aspects of how the credit assignment problem might be solved; the proportion of adaptation could be the trade–off between an uncertain environment and change in own body intrinsic properties.

Sensory recalibration during visuomotor perturbation of movements was also shown to be dependent not on the absolute size of the sensory (visual) prediction error rather than on the proportion of the error that is attributed to internal causes (Wilke et al., 2013). Thus smaller errors will cause higher sensory recalibration compared to large errors which are discounted. Moreover, significant recalibration of sensory predictions occurs when sensory prediction errors are above a inherent noise level. Wilke et al. (2013) also showed that motor adaptation is driven by the recalibrated sensory prediction errors rather than the actual prediction errors, thus casual attribution of error information is vital for sensorimotor learning. It is perhaps worth noticing that the visuomotor perturbations in Wilke et al. (2013) were presented in a unsystematic way. However, systematic errors can further set a context within which causality is inferred.

Adaptation occurs in presence of uncertainty of the world and uncertainty that is inherent in the sensorimotor system in terms of motor noise, noisy receptors and delays. Depending on the inherent uncertainty, errors are treated differently in the course of adaptation. For example, when people perform short reaching movements for which they are more certain of their own state (as opposed to longer movements) they are more sensitive in attributing visuomotor errors to an external source rather than to an internal (Wei and Koerding, 2009). Thus adaptation (of any kind) does depend on uncertainty due to motor noise.

Bayesian statistics have been used to describe the process of visuomotor adaptation in face of uncertainty. Koerding and Wolpert (2004), exposed participants to a visuomotor rotation where the shift imposed was drawn from a gaussian distribution (prior). Moreover, under different conditions the uncertainty of the visual feedback ranged from zero (clear visual feedback without any blur) to infinite (no visual feedback). The motor behaviour of the participants indicated that over the

trials participants must have learned to internally represent both the statistical distribution of the task and sensory uncertainty and to then combine them according to a performance–optimising Bayesian process.

To sum up, in this last section, several aspects related to adaptation to visuomotor perturbations were presented. Sensory prediction errors rather than on on–line movement corrections drive this kind of adaptation (Tseng et al., 2007). Moreover, presence of visual feedback during the movement decreases inter–trial variability. Adaptation occurs even if the context of the task does not demand it (Schaefer et al., 2012). In addition, an explicit cognitive strategy interferes with implicit adaptation process in the initial stages of adaptation (Mazzoni and Krakauer, 2006) but not later on (Taylor and Ivry, 2011). Similar to the adaptation to dynamic perturbations, visuomotor adaptation is accompanied by changes in perception (Henriques and Cressman, 2012) which again can occur in a context–independent way (Synofzik et al., 2006). Actually, adaptation to visuomotor perturbations does not depend on the actual size of prediction error *per se* but to the proportion of the error attributed to internal causes (Wilke et al., 2013). In turn, inherent motor uncertainty also affects the degree of adaptation (Wei and Koerding, 2009). Nonetheless, people are able to adapt in face of world and motor apparatus uncertainties, perhaps by using a Bayesian–like process (Koerding and Wolpert, 2004).

Adaptation to Delays In the previous section I presented several aspects of sensorimotor learning in motor tasks where spatial and dynamic perturbations were applied. Another aspect important to motor learning and control is the timing of events. Delays can be deleterious to motor control, thus learning to perform in presence of asynchronous sensory information is crucial. Adaptation to externally imposed delays haven been studied in the context of manual tracking (Foulkes and

Miall, 2000; Miall and Jackson, 2006), driving in a high–fidelity simulator (Cunningham et al., 2001a), prism adaptation (Kitazawa et al., 1995; Kitazawa and Yin, 2002; Tanaka et al., 2011), visuomotor adaptation (Honda et al., 2012a,b) and force field adaptation (Levy et al., 2010).

Foulkes and Miall (2000) and Miall and Jackson (2006) have studied the effect of delays in target pursuit tasks. Participants pursued via a joystick a visual target that was moving in unpredictable trajectories and with a variable speed within each trial. Delays (0 ms, 200 ms or 300 ms) were imposed in the visual feedback of the joystick position. The results in Foulkes and Miall (2000) showed that subjects were able to adapt to the delays as revealed by a significant decrease in the tracking error. Moreover, the mean power spectra also indicated improvement in the performance in terms of smoother movement execution. Adapting to delays in visuomotor performance shows that appropriate predictive feedforward actions are modified to cope with the target movement, but it does not reflect a tracking strategy based on the concept of Smith Prediction (Miall and Jackson, 2006). Improvement in tracking performance continues with training over days (Foulkes and Miall, 2000; Miall and Jackson, 2006). This means that adaptation to delays in visuomotor tracking is more similar to acquiring a new motor skill rather than adapting to a perturbation (motor adaptation).

Cunningham et al. (2001a) studied the effect of temporal misalignments on a driving task. In this task participants manoeuvred a virtual car in a virtual street environment. In a first experiment it was tested if participants can adapt to driving delays (130ms, 230ms and 430ms). The aftereffects upon removal of the delays in a post–test showed that participants adapted to the delay. In a second experiment, participants were trained to drive one route with the delays. The post–test included driving in novel streets. The results revealed that adaptation to the delays was

generalised providing evidence of a temporal visuomotor adaptation rather than a single visuomotor transformation. It should be highlighted that although the speed was kept constant within a trial, the velocity was frequently changing (there were no straight sections of the "streets").

The effect of delays has also been studied by integrating delays in prism and visuomotor rotation adaptation paradigms. For example, Kitazawa et al. (1995) used a typical baseline – adaptation – de–adaptation paradigm using prism goggles. During the adaptation phase participants reached to a screen (they had no visual feedback of their movement) and after they touched it they were given visual feedback of their error with a delay between 0–10 second. Delays above 50msec had a significant effect in the rate and amount of adaptation. However, even with delays as long as 5000ms participants showed a $40-50\%$ of adaptation compared to that of 0 delays. Similarly, in a monkey study, Kitazawa and Yin (2002) showed that a delay above 50ms significantly decreased the rate of prism adaptation in the monkey and delays above 500ms completely impaired the adaptation.

Adaptation to delayed prism rotation was shown to be driven by the actual (physical delay) rather than the subjective delay (Tanaka et al., 2011). The authors first showed that 100ms visual delay imposed between reaching to a target and the target presentation (without any spatial distortion of the target position) caused a 40ms shift in timing perception. It was then tested if participants that were adapted to a 100ms delay prior to the execution of the visuomotor adaptation task would show learning rates similar to adapting to 100ms delay imposed on top of the visual distortion or to a 60ms delay (equivalent to 100ms of physical delay minus the 40ms of time perception shift). Comparison to control groups showed that the subjective shift did not facilitate a faster visuomotor adaptation rate, so the delayed visuomotor adaptation was driven by the actual (100ms) delay rather than the subjective delay

(60ms).

Contrary to the findings of Tanaka et al. (2011), Honda et al. (2012b) showed that prior adaptation to the delay (only) did actually result in higher adaptation rates during later visuomotor rotation when compared to a sudden introduction of the delay along with the visuomotor distortion. The authors attributed the differences in their results to those of Tanaka et al. (2011) in that in the case of the Tanaka and colleagues' paradigm the visuomotor rotation was introduced suddenly and the visual feedback was given only at the end. In contrast, Honda et al. (2012b) allowed on–line visual feedback (participants had feedback of their movements during execution) and imposed the visuomotor rotation gradually during the first 20 trials. However, something that is not discussed by the authors is the prominent difference in the amount of delays the participants were exposed to in the two studies (250ms in case of Honda et al. (2012b) versus 100ms in Tanaka et al. (2011)). It could be speculated that the shift in time perception in the Tanaka et al. (2011) study (40ms) lies in a time window within which delays are attributed to internal noise. Thus, an effect similar to the one observed in Honda et al. (2012b) could have been masked.

In a subsequent study (Honda et al., 2012a), Honda and colleagues studied the effect of prior adaptation (habituation) to a delay on a visuomotor adaptation paradigm where the rotated targets were presented with 0, 100, 200 and 300 ms of delay. The rotated trials were only presented in brief blocks and the after–affects were exposed in subsequent brief trials. When participants were not habituated in visual feedback delays and were exposed to the delays during the visuomotor rotation trials they showed decreased aftereffects with increased delay. This outcome was expected as the directional error (the error orthogonal to the direction of movement) decreased with increased delay. A similar response would be expected in the group that adapted to 200ms delays prior to exposure to visuomotor perturbations if the

participants formed an internal model of the delay (Miall et al., 1993). On the other hand, if participants became sensitive to the habituated delay (200ms) the should exhibit a peak aftereffect in the visuomotor rotation trials when presented the same amount of delays. A combination of the two hypothesis can then explain the results in Honda et al. (2012a).

Delays have also been incorporated in a deterministic force perturbation paradigm (Levy et al., 2010). Participants were able to adapt to a force that was dependent to the participant's velocity 50ms earlier. In catch trials, participants that were exposed to the delayed force perturbation deviated from straight line significantly later compared to the control group, showing that they adapted to compensate a perturbation that was depended on the current velocity but arrived with a delay. The 50ms delay falls within the temporal window that allows adaptation to delays. It would be interesting in the future to investigate the boundaries of such time window in a force adaptation task.

Summary on error–based learning In the previous sections I have reviewed three main categories of error-based learning paradigms: dynamic forces, visuomotor perturbations, and delays. The common characteristic that drives learning during all these paradigms is an error that arises either by sudden application of force, or by imposed visual shifts, or by the visual discrepancy due to the delay. Identifying the source of error plays a crucial role in the amount of adaptation and it depends on the uncertainty related to each source. Moreover, motor adaptation is accompanied by perceptual shifts.

Reinforcement Learning

If error–based learning can be considered as a processes of movement correction, reinforcement learning is the process of action/movement selection. In machine learning,

Sutton and Barto give the following definition of reinforcement learning (Sutton and Barto (1998), p.3-4):

> "Reinforcement learning is learning what to do – how to map situations to actions – so as to maximise a numerical reward. The learner is not told which actions to take,.., but instead must discover which actions yield the most reward by trying them."

This definition points out that the differences between supervised and reinforcement learning are both the signal that drives the learning (sensory prediction error or reward) and the information given by the "teacher" signal in each kind of learning. The error in the one case also indicates the direction of the corrective movement, in reinforcement learning no such information is provided. The basic elements of reinforcement learning are: a *policy* that will define the behaviour and must be learned, a *reward function* that attributes rewards to states, a *value function* that specifies the long–term reward accumulation and a *model of the world* (where by world is meant a model of the environment and the body). Having a model of the world is not necessary for a reinforcement learning system and so we could distinguish between reinforcement learning algorithms that are model–based (indirect) and model–free (or direct). In the first case, experience improves the model and the output of the model (via planning processes) the value function and the policy. In the latter case the value function and the policy are changed directly via experience (Sutton and Barto, 1998).

Recently, the terms "model–based" and "non–model–based" learning (inspired from the aforementioned framework) have been used to differentiate among different kinds of motor learning (Haith and Krakauer, 2013). In their review, Haith and colleagues equated the term model–based learning with what I described before as error–based learning and the term non–model–based with model–free reinforcement

learning and the use–dependent learning (I will refer to this kind of learning later in this section). However, it is not very clear if, when using the term model–based learning, Haith and Krakauer (2013) mean generally the involvement of a forward model in learning (so that error–based learning is model–based) or they imply a model–based reinforcement learning process. Converting a supervised learning problem into a reinforcement learning one is possible (but not the other way around) but perhaps would add unnecessary complexity (Barto and Dietterich, 2004). In this review, the terms model–based and model–free are going to be use only to describe the two different reinforcement learning cases.

As described above, there are critical steps that most of the reinforcement algorithms consider: evaluation of actions, selection of actions and learning from experience (Daw and Doya, 2006). In the first step, a value is attributed to each action which defines the expected utility if this action is taken. An action is then selected based on the action values. The action selection in reinforcement learning entails a dilemma; while the objective is to maximise future reward, always selecting (exploiting) the actions with the higher value, or actions that in the past have been shown to be effective, would prevent the agent from discovering new actions that could obtain more reward in the future. So, an agent must consider a trade-off between exploring novel actions and exploiting the established ones (Sutton and Barto, 1998; Stafford et al., 2012). Equally, in an uncertain environment an agent needs to make a decision whether more sampling (information gathering) is needed before acting. The trade–off between sampling and acting is learnt and altered during learning (Rao, 2010).

Upon action completion the received reward is used to update the values of the actions and thus experience is used to learn. Learning is driven by trying to minimise the difference between the expected and received reward. This reward

prediction error is central in Temporal Difference (TD) algorithms of reinforcement learning. The expected reward can be computed based on the expected rewards of several future states (forward view of TD) which constitute an eligibility trace. There is also another view of the eligibility traces, that is called backward view and is essential for temporal credit assignment. That is, in the presence of a reinforcing event the eligibility trace is going to define if a past state has contributed to receiving a reward and the degree of its contribution (Sutton and Barto, 1998).

The temporal assignment problem is a critical issue in reinforcement learning problems and it is of vital importance when we consider action acquisition. At this point we should differentiate between action selection and action acquisition. In operant (or instrumental) conditioning paradigms the actions that are going to be reinforced and selected are in a sense predetermined. However, a question rises about how these actions were acquired at first place (Stafford et al., 2012) or else how we discover novel actions, i.e. "a movement that has a predicted outcome" (Redgrave et al., 2008, pp.331).

Reinforcement learning has been mainly studied in the framework of operant conditioning as introduced by Skinner and in study of decision making (e.g. Glaescher et al., 2010). However, it is only recently that researchers have started exploring the characteristics of reinforcement learning in human motor paradigms.

Recent evidence showed that saccadic movement characteristics (duration and velocity) are affected by temporal discounting of future reward (Shadmehr, 2010). In another study, Dam et al. (2013) used a paradigm where hand trajectories were perturbed in two ways (features): direction–wise and curvature–wise. Participants performed reaching movements, where no targets were provided, and received monetary reward feedback for either correcting for curvature or for direction. The results showed that the subjects were able to extract the rewarding features and learn the

implicit value function.

In the framework of visuomotor rotations, Izawa and Shadmehr (2011) studied the consequences of sensory prediction errors and reward prediction errors. Participants were given either high quality visual feedback of their movement (a cursor throughout the movement) or low quality visual feedback (information about their position at the end of the movement) or no visual feedback. In all cases they received a binary reward for successful (or not) reaching–to–target movement. The reward and sensory prediction errors were found to have differential effects on adaptation. After exposure to high quality sensory feedback, adaptation driven by the sensory prediction error lead to a remapping of motor commands to sensory consequences and to broad generalisation. The same was not true when participants received low quality visual feedback or when the relied only on the reward prediction error to adapt. Another interesting finding of this study was that variability during adaptation, which was larger for the reward–only group, was related with exploration of motor commands in order to maximise rewards.

Stafford et al. (2012) have recently developed a paradigm for studying acquisition of novel actions in humans. Participants need to explore the workspace and discover the movement that elicits reward. This paradigm sets an appropriate framework to study several aspects of reinforcement learning, such as the exploration–exploitation trade–off. It was also shown by the same group, how delays and sensory modality play an important role in acquiring novel actions and solving the temporal credit assignment problem (Walton et al., 2013). Moreover, in this context movements appear to be learnt as trajectories rather than in terms of a reaching target (Thirkettle et al., 2013).

Use–dependent Learning

Pure repetition of movements, in absence of errors or rewards, can lead to changes in performance. This kind of learning has been referred to as use–dependent learning (or sometimes unsupervised learning) (Wolpert et al., 2001, 2011; Haith and Krakauer, 2013). Use–dependent learning can be related to achieving skilled performance. For example, even when a pianist has put together the finger sequences to be executed in a piece of music, speed can only arise via exhausting repetition. To become skilled is a long–term procedure and and as many as 10000 hours of practice might be necessary to conquer perfection (Miall, 2013).

Recent experience can change the statistical properties of movement even in the absence of errors (Verstynen and Sabes, 2011). This experience–dependent learning produces a variance–bias trade off, which is compatible with Bayesian estimation processes. Verstynen and Sabes (2011), in a visually guided reaching movements task, showed that repeated practise of reaching movements to targets near a particular target location decreased variance in subsequent execution of movements to that particular target location but increased the bias (decreased accuracy) for reaches to other targets. They further demonstrated, by means of a computational model, that these results can be explained by Hebbian (associative) learning.

Error–based learning and use–dependent learning can occur at the same time; while the first acts to cancel the effects of a disturbance the latter associates the current executed movement to the last one. In Diedrichsen et al. (2010a), participants performed reaching movements in a redundant–task design (reaching to a bar rather than a target). At first place, use–dependent learning was induced by passively guiding movements in a hand path directed 8^o laterally. Subsequent free movements were biased towards the direction of the manipulation. However, when participants actively executed movements being physically constrained in the same

hand path, it was shown that after removal of the constraint there was a brief, error–based aftereffect against the direction of the force channel and a longer–lasting use–dependent after effect in the direction of the imposed path. Finally, when comparing the behaviour in a typical force field adaptation paradigm under a redundant–task design and a standard goal–directed design, Diedrichsen and colleagues found that while participants equally adapted to the perturbation in both tasks, so error–based learning was not reduced in the redundant task, they only showed use–dependent learning in the redundant task. According to the authors that indicated that that use–dependent learning is task sensitive.

Huang et al. (2011) tried to differentiate the effects of the different kinds of learning in a study where they used variants of a visuomotor rotation, in which movement repetition was either eliminated or exaggerated. Based on the result of their study, the authors proposed that different learning process underlie visuomotor adaptation and subsequent savings. In particular, model–based (error–based) learning will attempt to reduce errors in visual space, use–dependent learning (model–free) will induce directional biases toward the repeated movement but savings will only occur via operant reinforcement (model–free).

Specific areas in the cortex have been related to motor skill acquisition. For example, participants that were trained over days to a new motor skill while receiving M1 anodal transcranial direct current stimulation (tDCS) achieved greater performance compared to the participants that received sham stimulation (Reis et al., 2009). In another study, Transcranial Magnetic Stimulation (TMS) was applied on primary motor cortex in a group of skilled musicians and non–musicians (Gentner et al., 2010). Both these results implied that experience–dependent motor skills are represented in the primary motor cortex area. In a recent imaging study, Wiestler and Diedrichsen (2013) showed that learning of sequential movements leads to the

development of specialised neuronal circuits in primary secondary motor areas and mainly in the supplementary motor area. Long–term training (over years) in monkeys was found to have lead in reduction in the synaptic activity needed to perform sequences of movements (Picard et al., 2013) in the primary motor cortex, suggesting an increased efficiency of activation of the neurones in this area.

1.4 The Cerebellum and the Basal Ganglia in Sensorimotor Learning

1.4.1 Cerebellum

I mentioned previously (Section 1.2.2) that the cerebellum retains forward models. Thus, as it would be expected, cerebellar damage causes deficiencies in adaptation related to problems in calculating a sensory prediction error. It was discussed before how Tseng et al. (2007) showed that on–line corrections did not add any advantage in the adaptation of control subjects and so the use of sensory prediction errors is essential for adaptation. In the same study, a group of cerebellar patients showed deficits in adapting, in both shooting and reaching movements, confirming the crucial role of the cerebellum in visuomotor adaptation via calculating sensory prediction errors.

It was suggested earlier that gradual introduction of a distortion leads to different patterns of generalisation possibly through a mechanism that attributes small and large errors to different sources (Kluzik et al., 2008). So a related question is whether the cerebellum plays a different role in presence of large or small sensory prediction errors. Criscimagna-Hemminger et al. (2010) tested the effects of abrupt versus gradual adaptation in force–field adaptation in a group of severely affected cerebellar patients. The patients showed an improved performance only when perturbations were introduced gradually. Moreover, the motor memory after

gradual training persisted for longer compared to when the perturbations were introduced abruptly. These results were interpreted as an indication of distinct neural mechanisms in motor learning which are differentially driven by large and small errors. However, in a more recent study Gibo et al. (2013) showed that environment dynamics were more important than error size. Specifically, cerebellar patients performed much better in a clock–wise (CW) curl force field (compared to counter clock–wise (CCW) one) that was assisting their movement, both in abrupt and gradual introduction of the force field. The authors attributed the differences between Criscimagna-Hemminger et al. (2010) and Gibo et al. (2013) in that in the former study the effect of error size observed could have been a result of the force fields applied as the gradual adaptation was performed with a CW force field and the abrupt condition with a CCW direction. Alternatively, differences in the protocols may account for the differences observed. So even though the error determines the kind of forward model to be learnt (of a tool or of body), cerebellar patients seem to be equally impaired in both cases.

When talking about cerebellar contributions to adaptation it is also important to consider not only deficiencies caused by damage in the cerebellar cortex but also in deep cerebellar nuclei. For example, inactivation of dentate nucleus impaired differentially the adaptation to visuomotor rotations when the perturbations were presented gradually (significant dysfunction) and abruptly (Robertson and Miall, 1999).

The studies that have examined the effect of size of errors in visuomotor rotations have given controversial results. Schlerf et al. (2013) observed that cerebellar patients showed deficits in adapting to a visuomotor rotation both when the perturbations were presented abruptly or gradually. On the other hand, Izawa et al. (2012) found that patients and controls did not differ significantly in terms of reach

direction (they adapted in a similar degree). Moreover, Izawa and colleagues measured generalisation patterns (generalisation of adaptation to neighbouring targets) after adaptation, as a mean of assessing changes in motor commands depending on the sensory goals (inverse model adaptation), and they found that patients and controls exhibited similar results. However, in an inter–limb localisation task, only the control group showed altered sensory consequences of their motor commands. According to the authors this indicated that cerebellar damage impaired the ability to learn the visual consequences of motor commands. The difference in the results between the two studies could be searched in differences of the gradual adaptation protocol but also in the differences among individual patients. A prominent difference though between the two studies is that Schlerf et al. (2013) were only giving visual feedback at the end of the trial (during the adaptation phase) whereas in the study of Izawa et al. (2012) visual feedback of the cursor was provided during the reaching movement as well. As a result it could be that it was easier for the patients in Izawa et al. (2012) to infer and use a strategy. It was previously shown in a study by Taylor et al. (2010) that cerebellar patients were able to adapt to a visuomotor rotation using a explicit cognitive strategy, in contrast to healthy control subjects who were not dependent on a strategic solution, confirming the role of cerebellum in implicit adaptive processes.

Although the cerebellum is related to adaptation in both force–field perturbations and visuomotor rotations during reaching movements, different areas in the cerebellum are involved in each kind of adaptation, and the performance in one task is independent from the performance in the other (Rabe et al., 2009). Specifically, as shown by Rabe and colleagues, degeneration in the intermediate cerebellar zone of the posterior lobe is related to deficits in visuomotor adaptation whereas degeneration in the intermediate and lateral zones of the anterior cortex leads to deficits

in force–field adaptation.

Cerebellar activation has contributions related not only to error performance but also to changes in performance. In an imaging study, Miall and Jenkinson (2005) measured the change in brain activation levels before and after training to a novel eye–hand tracking task. There were two groups of subjects that were trained with or without temporal delays between the hand and eye movements. The changes of activation in the cerebellum as a result of training support the hypothesis of the acquisition of an internal model that allows predictive control of eye-hand movements.

Evidence of the involvement of the cerebellum in updating internal predictions of the sensory consequences of actions are also supported by patient studies. In Synofzik et al. (2008) cerebellar patients appeared to exhibit perceptual adaptation in presence of visual feedback but they seemed unable to use this information in order to update an internal model. Thus, in the absence of feedback they showed a significantly smaller perceptual adaptation compared to the healthy controls. In addition, they exhibited no motor adaptation.

On the other hand, there are findings indicating that the cerebellum is not necessarilly involved in sensory recalibration. In a recent study by Block and Bastian (2012), cerebellar patients showed similar amount of sensory realignment compared to healthy controls in a "sensory task" where no feedback about the reaching hand position was ever presented. Sensory recalibration was driven by an error between the proprioceptive and visual information of the target hand.

The cerebellum has been also shown to retain forward models of external objects of the environment (Cerminara et al., 2009). Similarly, cerebellar damage causes deficiency in recalibrating sensory predictions about external sensory events (Roth et al., 2013). Participants in Roth et al. (2013), had to predict the time of reappearance of a moving target that temporally disappeared behind an occluder. Both healthy

controls and cerebellar patients could predict this time at baseline. However, when a delay was introduced during the occlusion phase cerebellar patients were unable to recalibrate their predictions which indicated that they were unable to model the delay and keep their spatiotemporal predictions accurate. So, the cerebellum appears to play an important role during learning not only in the motor but also in the sensory domain.

To sum up, the cerebellum is essential for calculating sensory prediction errors via internal forward models. How the cerebellum is involved in adaptation depending on the size of errors is still an open question and could be related to the question of the involvement of cerebellum in the credit assignment problem. The internal models in the cerebellum must be adapted and new models could be formed in the course of sensorimotor adaptation to capture changes in the environment and the body. Although cerebellum is not necessary for sensory recalibration, changes in sensory perception must be integrated in internal predictions *of* the cerebellum in order to keep accurate the motor–to–sensory mappings.

1.4.2 Basal Ganglia

The basal ganglia are the brain structures mostly related to reinforcement learning. The various nuclei within the basal ganglia are associated with different components of reinforcement learning (Samejima and Doya, 2007). For example, action re–selection based on previous experience could be implemented by strengthening the pre– or post–synaptic inputs to the striatum (input nuclei) (Redgrave et al., 2008). Phasic dopamine neurones in substantia nigra (output nuclei) are believed to provide the reward prediction error (Schultz, 2013). An alternative hypothesis is that phasic dopamine activity, given its temporal characteristics, might be a signal that provides a "time–stamp" determining which actions were relevant to an elicited

reward (Redgrave and Gurney, 2006). Time–stamping an action is important for determining the agent as the cause of an event and attributing an outcome to a novel action (Redgrave et al., 2008).

Giving a detailed review of the aforementioned issues goes beyond the scope of the current chapter. Here, I am mostly interested in reviewing paradigms of motor learning that entail reinforcement learning characteristics. However, to my knowledge there is no imaging or patient study on this topic. Of course there are studies that have looked into the role of basal ganglia in motor control and adaptation paradigms but I will discuss such examples in a later paragraph. However, I will mention few imaging studies that have looked into how reinforcement learning might be implemented in the basal ganglia in non–motor tasks.

In a fMRI study by Seymour et al. (2004), participants were presented with two sequential visual cues. Based on the second cue a predictable pain stimulus was delivered. However, the first cue only allowed a guess about the stimulus to be made. This manipulation allowed the authors to study both the encoding of expectations but also of prediction errors. The results revealed that prediction error was highly correlated with activity in putamen, caudate and the substantia nigra but also in the cerebellum. Temporal difference value was found to be represented in the anterior insula cortex.

Glaescher et al. (2010) studied the existence of model–based and model–free reinforcement learning signals in human brain during a sequential two–choice decision task. States were presented as fractal images and were associated with values. Participants first learned the model of state transitions and the reward contingencies. In a second session participants were allowed to make free choices (actions). The results showed that activation in the lateral prefrontal cortex was related to state prediction error, whereas ventral striatum was related to reward prediction error.

1.4.3 Cerebellum and Basal Ganglia

The classical view of the connections between the cerebellum and the basal ganglia held that the cerebro–cerebellar and cerebro–basal ganglia loops are anatomically and functionally distinct and that any interactions between the two loops occur primary at cortical level (Doya, 2000). As a result many of the studies that looked at both areas simultaneously have focused on differentiating the roles of the two brain areas. In the first part of this subsection I review some studies that have focused on differentiating the roles of the two brain areas in motor control and learning.

However, recent anatomical evidence (Bostan et al., 2013) showed that the cerebellum and the basal ganglia are actually connected in a disynaptic way. These evidence are discussed in the second part of this subsection. So, although the two areas contribute differently in motor control and motor learning the existence of direct connections between them poses the question of how these areas collaborate. For example, activation of both the cerebellum and the basal ganglia has been reported during the early phase of visuomotor adaptation (Seidler et al., 2006). What is the relative contribution of each area to the adapation process? In this last subsection I briefly mention how the existence of direct connections between the two areas offers a new background to explain clinical symptoms of diseases and the implications of these findings in alternative treatments.

1.4.4 Different functional roles

Jueptner and colleagues in a series of PET studies of CBF (Jueptner et al., 1996, 1997a,b,c) studied the role of cerebellum and basal ganglia in the control and learning of movements. These studies provided evidence that the basal ganglia (but not the cerebellum) are involved in movement selection whereas the cerebellum (but not the basal ganglia) is involved in sensory information processing (Jueptner and Weiller,

1998).

In a set of motor learning studies (Jueptner et al., 1997b,c), participants had to learn a sequence of finger movements under different conditions. In a *New sequence learning* (new) condition, the subjects learnt a new sequence of finger movements by trial and error. This characteristic of the task made learning explicit (in contrast to implicit motor learning of finger movements sequences where participants are not aware of the sequence). In another condition participants were required to perform in the scanner a *pre–learned sequence* (pre). During the *free selection* condition (free) participants were pressing keys randomly, while in the *repetitive* condition (rep) they were required to press a single key repetitively with their middle finger. The results showed that both the basal ganglia and the cerebellum were active in the learning of motor sequences (new minus pre) and in the improvement of motor performance (new minus free), although the cerebellar areas were more activated in the latter contrast. However, the contrast between the free selection and repetition conditions (free minus rep) showed that only the basal ganglia were involved in the decision making and movement selection procedures. Based on these results Jueptner et al. (1997c) suggested that the basal ganglia may be involved in specifying the movements on the basis of a mapping of the movements and their outcomes, whereas the cerebellum may deal more directly with movement execution changes.

In a visuomotor co–ordination study, Jueptner et al. (1996) studied the roles of the cerebellum and basal ganglia to the sensory guidance of continuous movements. Participants performed four different conditions of the a visuomotor task in the scanner. During the first condition (draw) they were asked to draw a series of lines on a computer screen using a computer mouse. After that, participants had to track the same lines with a mouse pointer on the screen – that was the copy condition (copy). In a third condition (eyes), participants were asked to follow with their eyes

the movements of a pointer that was presenting the same lines as in the previous conditions. Finally, the subjects completed a fixation condition where they had to fixate their eyes at a central point and ignore the sequence of presented lines. The results showed that during the tracking condition there was a massive activation of the cerebellum compared to the freely drawing condition (copy minus draw) but not activation of the basal ganglia. On the other hand the basal ganglia were activated equally in both conditions. Based on these results, the authors suggested that the cerebellum is differentiated into correcting movements using sensory information while the basal ganglia are concerned with self–generated movements and movements driven by external cues.

Finally, Jueptner et al. (1997a) looked at the relevance of sensory input for the cerebellar and basal ganglia control of movements. In this study there were three conditions. In the first condition (active) participants performed flexion movements, with their arms fixed to a hinge. During the passive movement condition (passive) a motor was inducing flexion and extension movements of the elbow. The third condition was a rest condition (rest) where the participants kept their eyes closed and executed no movements. The results showed that only the cerebellum was activated in the passive condition (passive vs rest contrast). Both the cerebellum and the basal ganglia showed activations during active movements (active vs rest) and in the direct comparison between active and passive movements (active minus passive). Thus, it was concluded that only the cerebellum and not the basal ganglia are concerned with monitoring the outcome of the movements (afferent sensory component).

The distinct roles of the cerebellum and basal ganglia has also been studied via recording and inactivation studies in monkeys (van Donkelaar et al., 1999, 2000). The animals were trained to perform visually triggered and internally generated limb movements. The results showed that the part of motor thalamus that receives

input from the cerebellum was related to the visually triggered movements whereas the part of thalamus receiving input from the basal ganglia was related to internally generated movements. So, cerebello– and basal ganglio–thalamo–cortical loops are differentially involved in movements depending on the presence on external or internal cues.

1.4.5 Anatomical connections and functional consequences

Recent anatomical data (Hoshi et al., 2005; Bostan et al., 2010) which have provided evidence for direct connections between the two areas. Specific strains of viruses that move from one neurone to another exclusively at synapses allowed for the examination of circuits that are multisynaptic (Bostan and Strick, 2010). These neurotropic viruses move in a time–dependent way and by adjusting the survival time after the injection it is possible to track neural connections composed of more than two synaptically connected neurones. Hoshi et al. (2005) and Bostan et al. (2010) used rabies virus that is transported solely in a retrograde direction in a time–dependent way and showed that the cerebellum and the basal ganglia are linked directly with disynaptic pathways.

Hoshi et al. (2005) injected N2C strain of rabies virus into the putamen and external part of globus pallidus (GPe) of macaque monkeys to explore links between the dentate nucleus – output of cerebellar processing. They found the existence of a disynaptic projection to the striatum originating from both the motor and non–motor domains of the dentate nucleus and also some less substantial inputs from fastigial and interpositus. Moreover, there was a tri–synaptic link between the dentate nucleus and the GPe.

In another study of the same group, Bostan et al. (2010) injected rabies virus in Crus IIp and the hemispheric expansion of lobule VIIB (HVIIB) of the cerebellar

cortex of cebus monkeys. They found second–order neurones labeled with the virus in the subthalamic nucleus (STN). Moreover, this connection between the cerebellar cortex and the STN is topographically organised, with most of the STN neurones that project to Crus IIp to be located in its associative territory and those projecting to HVIIB to be located in the sensorimotor territory of STN.

The existence of direct connections between the basal ganglia and the cerebellum raises further questions about how these two areas might collaborate. In a recent review Wu and Hallett (2013) discuss about the role of cerebellum in Parkinson's disease which is related to basal ganglia (dopaminergic) degeneration. Dopaminergic degeneration, abnormal drives from the basal ganglia and dopaminergic treatment might all cause pathological changes in the cerebellum which can explain some of the clinical symptoms in Parkinsons disease. On the other hand, compensatory cerebellar mechanisms can assist maintaining motor and non–motor functions.

Dystonia is another condition that has been related to basal ganglia. However, recent evidence suggests that both the cerebellum and the basal ganglia play a role in the generation of dystonia and that aberrant cerebellar activity (perhaps compensatory) is an important parameter with the pathophysiology of dystonia (Sadnicka et al., 2012). For example, abnormal cerebellar activity altered basal ganglia activity in a mouse model. However, when the connection between the cerebellum and the basal ganglia were severed dystonia was alleviated (Calderon et al., 2011). All these findings are of great significance in terms of potential therapeutic treatments.

1.5 Conclusions

As discussed in this literature review the role of cerebellum in error–based learning is well–appreciated. Of course, there are still a lot of questions related to adaptation that need to be answered. However, in this thesis I focused on researching about the

role of cerebellum in other kinds of learning.

The first question posed was what is the contribution of cerebellum in a novel motor skill paradigm. In Chapter 2, we present a where study transcranial Direct Current Stimulation (tDCS) was used to modulate cerebellar activity during a finger tapping task.

We then shifted our interest to exploring the role of cerebellum in reinforcement learning. As presented earlier, reinforcement learning is gaining a lot of attention recently and how this mechanism might be implemented by the brain in motor tasks is unknown. In light of anatomical connections between the cerebellum and the basal ganglia, an interesting question is what is the relative contribution of each area and how they might collaborate.

In order to pursuit these questions, we first designed an experiment that combined the exploration task developed by Stafford et al. (2012) and a visuomotor tracking task (Foulkes and Miall, 2000). The motivation was to use a basal ganglia and a cerebellum dependent task and explore the interactions between the two. The aim was to establish a behavioural paradigm that could be transferred to an imaging study. In Chapter 3 we present the results of this combined Exploration–Tracking study. In Chapter 5, we present another study across the same lines, where we investigated the contribution of cerebellum but also of dorsolateral prefrontal cortex and motor cortex in the exploration task.

The results of Chapter 3 indicated that the effect of the tracking task to the exploration task was not the expected. The result lead us to the conclusion that during motor adaptation in the tracking task there might be an adaptive process occurring that was not considered in the first place. Thus, in Chapter 4 we investigated how the tracking task might be changing proprioceptive uncertainty.

Chapter 6 presents an overview of the work completed, and discusses the poten-

tial role of cerebellum in action acquisition and motor learning. Moreover, I discuss some considerations to be taken into account when designing a tDCS study and some possible limitations of the analyses used.

Chapter 2
THE EFFECTS OF CEREBELLAR TDCS IN A KEYPRESS MOTOR TASK

2.1 Introduction

Mastering fast and accurate motor behaviours is achieved by learning to control our body and limbs. Motor learning is the process of establishing appropriate control mechanisms that enable skill proficiency. Learning is accomplished by fine tuning of different components, such as information extraction, decision making, strategy selection and control implementation (Wolpert et al., 2011). It is apparent that these components need to be explored and exploited differently depending on the nature of the skill. For example, learning to dance and playing the piano can be quite different to learning to play the violin or to playing tennis, in that in the first case one needs to learn to control only parts of her body, whereas in the latter case the dynamics of a hand-held object need to be integrated in the execution of a movement. Another important aspect of learning is the process or processes that are recruited during skill acquisition. Learning a new motor behaviour is based on different mechanisms compared to modifying a movement based on trial by trial errors (motor adaptation) (Bastian, 2008).

The cerebellum has long been identified as a brain structure that plays a critical role in motor control, adaptation and learning. It has been proposed that the cerebellum retains internal representations of our motor system. Internal models allow for predictions of the motor states that are essential for motor control, in order to counteract sensory transport delays and biological noise (Wolpert and Miall, 1996). State predictions are also essential for the computation of motor prediction errors that drive motor adaptation (Shadmehr et al., 2010). The role of the cerebellum in other processes of learning is also appreciated because of its involvement in motor

functions. However, most of the human behavioural studies to date have focused in understanding the role of cerebellum in motor control and adaptation rather than in other types of learning, such the acquisition of a motor skill.

One approach to investigate the role of cerebellum in motor learning is through brain stimulation. *Transcranial direct current stimulation (tDCS)* is a non invasive method that is used to manipulate neural activity. tDCS over motor cortex drives a polarity specific effect; cortical anodal stimulation enhances excitability whereas cathodal tDCS has the opposite effect (Jacobson et al., 2012). Galea et al. (2009) showed that tDCS modulates in a polarity specific way the excitability of the cerebellum, as well. Purkinje cells (PC - output neurons of the cerebellum) activity exerts an inhibitory tone over M1, through the PCs inhibitory connections to dentate cerebellar nucleus (Middleton and Strick, 2000). This is referred to as *cerebello–brain inhibition (CBI)*. The results of Galea et al. (2009) showed that anodal tDCS enhanced CBI which resulted in a stronger PC inhibitory tone exerted to M1 (thus M1 was stronger inhibited), whereas cathodal tDCS resulted in a decrease of CBI (more excitation observed at M1). In another study, Galea et al. (2011) used tDCS to study the role of the cerebellum in visuomotor adaptation. Their results showed that anodal tDCS caused faster adaptation to the visuomotor transformation, depicted in the reduction of movement errors. Similarly, anodal cerebellar tDCS applied during locomotor adaptation expedited the adaptive process while cathodal tDCS slowed it down (Jayaram et al., 2012). Finally, Ferrucci et al. (2013) showed that anodal cerebellar tDCS influenced implicit learning processes in a serial reaction time task (SRTT).

In the presented study, we investigated the role of the cerebellum in learning a novel motor skill by using both anodal and cathodal tDCS. The task involved learning a chord–like finger tapping sequence that demanded the coordination of fingers.

The participants had no or minimal experience in key instruments, so they were naive in this type of motor movement. We hypothesised that the participants who received anodal stimulation would achieve better performance compared to sham, and those that received cathodal would either show greater difficulty in performing the task or exhibit the same performance as the sham group.

2.2 Materials and Methods

2.2.1 Participants

70 right handed participants (age range: 18-36, mean=21, 53 female) participated in the study. They were informed about all the aspects of the experiment and gave informed written consent. A screening questionnaire was filled in to assure safe application of tDCS. The experimental protocol was approved by the Science, Technology, Engineering and Mathematics Ethical Review Committee of the University of Birmingham Ethics Committee.

2.2.2 Task

Participants performed a chord-like keypress task. All of them had no or minimal experience in playing key instruments. They were asked to use their right hand and simultaneously press with their thumb and middle finger 'space' and 'B' keys of a keyboard and then change to 'V' and 'N' with their index and ring finger.

The total time of tDCS stimulation was 20 minutes. The task started 10 minutes after the initiation of stimulation and ended a minute after the stimulation had finished. There were 7 Blocks. During a first sub-block participants performed as many correct trials as possible for 15 seconds (*performance phase*), followed by a small break of 5 seconds. In a second sub-block, they were instructed to change pair of keys every one second following an auditory stimulus of 1Hz, for a minute

(*synchronization phase*). A larger break of 15 seconds followed every *synchronization phase* sub-block. This sequence of events was repeated 7 times during the study (see Figure 2.1). Participants received no feedback of their performance during the study.

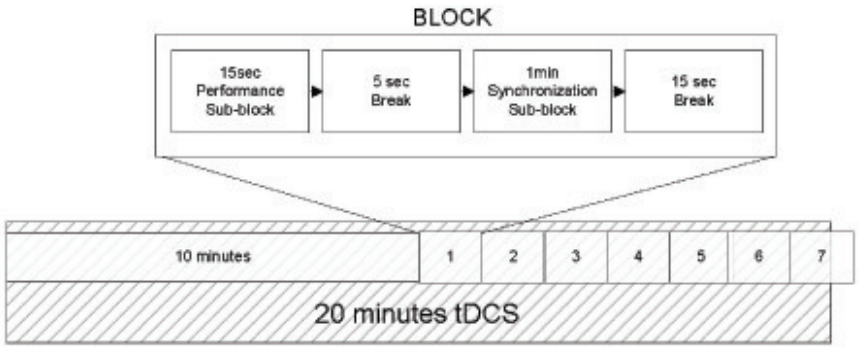

Figure 2.1: Participants received stimulation for 20 minutes. They performed the task during the last 10 minutes of the stimulation. There were two main phases of task execution: a *synchronization* phase and a *performance* phase.

2.2.3 tDCS

Direct current was delivered via a pair of conductive rubber electrodes (thickness d$\simeq 1mm$ surface A$\simeq 3.7 \times 3.7\ cm^2$) that were inserted in NaCl-moistened sponge pockets ($A_{sponge} \sim 4.7 \times 5.72\ cm^2$). For stimulation of the lateral cerebellum (right hemisphere), one electrode was centred $3cm$ lateral and $1cm$ below the inion (Miall et al., 2007). The other electrode was positioned on the ipsilateral buscinator muscle (Galea et al., 2009). A Magstim Eldith DC-stimulator was used to deliver a current of $2mA$ (Galea et al., 2009; Ferrucci et al., 2008). The 0.075 mA/cm^2 current density (total charge 890 C/m^2) is within the safety limits (Liebetanz et al., 2009). Current intensity was raised in a ramp-like way over the first 10 seconds until the desired level and faded out the same way at the end of stimulation. The stimulation lasted 20 minutes, in anodal and cathodal conditions. During sham condition, stimulation was applied for 15 seconds at the beginning and the end of the 20 minutes. Depend-

ing on the type of stimulation received (anodal, cathodal or sham), participants were assigned into three different groups.

2.2.4 Data Analysis

The time of keyboard keypresses was recorded using PsychotoolBox in Matlab vR2007b. Suppose the times $t_{1_{space}}$, t_{1_B} of a first simultaneous pair of keypresses (thumb and middle finger pressing space bar and 'B' key) followed by a second pair (index and ring finger pressing 'V' and 'N' keys) with times t_{2_V} and t_{2_N} (index and ring finger – 'V' and 'N' keys). The *Intertap Interval* (IeTI) was determined as the time difference Δt=t_{2_V}-t_{1_B} (i.e. the time difference between the first keypress of the second pair and the second keypress of the first pair). In order to account for errors, especially in the *performance phase* sub–blocks, the IeTI was calculated based on successful consecutive trials only. The *Intratap Interval* (IaTI) was defined as the time between the two keypresses of one pair (e.g. t_{1_B}- $t_{1_{space}}$). IeTIs and IaTIs with a value larger than the mean plus/minus 2 standard deviations (per participant and per block) were omitted from the analysis. For each sub–block of the *performance phase* the mean, standard deviation (SD) and Coefficient of Variation (CoV) of IeTI were calculated. To deal with the variability in individual data, the CoV was used as an indicator of the relative precision of the temporal measurements and was calculated as the standard deviation divided by the mean estimate. For each sub–block of the *synchronization phase* the mean, standard deviation (SD) and Coefficient of Variation (CoV) of both the IeTI and the IaTI were calculated.

In addition to the aforementioned analysis, Matlab's Curve Fitting Toolobox was used to fit the mean IeTI data into learning curves. Both individual and group data were used. The power function used to fit the data was of the type $ax^b + c$, where b is the learning rate of the curve. The fitting method was that of non linear Least

Squares and the algorithm used was that of Trust Region (default Matlab setting). The values of the parameters a and c were bound to be positive, whereas the value of parameter b was bound to be negative. The learning rates (parameter b) of the fitted curves and the fitted IeTI values were used for further data exploration.

From the 70 participants that participated in the study, only 36 (12 in each group) were used for the analysis. 11 participants were initially excluded because the number of keypresses in at least one *synchronization phase* sub–block was outside the range of $55 - 65$ pair of keypresses per minute. Since during these sub–blocks participants had to follow a 1Hz auditory stimulus and thus the task was easy and clear, failing to perform in these blocks could indicate that either they had not understood the task or that they were indifferent in performing. Moreover, we can not exclude that in some instances the keyboard–software interface could have failed. The remaining 23 participants were excluded because they showed no learning or they were performing worse over time in the *performance phase* sub–blocks. For this reason, we calculated the normalised (to the mean of the first block) means of IeTI during those sub–blocks and we excluded the participants that in blocks 2–7 had mean normalised IeTI above one in at least two blocks. As the number of participants that showed this kind of behaviour was almost equal in all groups (6 Sham, 9 Anodal, 8 Cathodal), we rule out the possibility that the 'non-learning' behaviour was an effect of tDCS application. Of course the 'non-learning' behaviour has to do with the time participants were exposed to the task, but the design flaws will be discussed later on.

2.2.5 Statistical Analysis

SPSS 16.0 was used for the statistical analysis. The means, SDs and CoVs for both IeTIs and IaTIs were analysed using 3×7 mixed ANOVAs, with a between par-

ticipant factor of *Stimulation* (with three conditions of anodal, cathodal and sham tDCS) and a within participant factor of learning over *Time* (with 7 conditions each of the seven blocks of the study). In case of a significant or close to significant interaction between the two factors further analysis of the simple main effects was conducted. In case of a significant or close to significant main effect of the between participants factor of tDCS, post–hoc analysis using Tukey HSD tests was conducted. The mean IeTI fitted data, in case of the *performance phase* sub–blocks were also analysed using a 3×7 mixed ANOVA. The learning rate (coefficient b of the fitted curves) was analysed separately using a one-way ANOVA among the three tDCS groups. In all the statistical tests we conducted the Mauchly's test of sphericity was violated (p–values were always $p < .001$), therefore the degrees of freedom were corrected using Greenhous–Geisser estimates of sphericity.

2.3 Results

2.3.1 Performance Phase

Figure 2.2 shows the mean IeTI values over the course of the seven blocks in each stimulation group. It is obvious that participants in all three groups got faster in performing the task but there is no apparent difference in the learning rates of the three groups. It is also worth mentioning the differences in average group behaviour at Block 1.

The ANOVA showed that there was a significant effect of learning over time ($F(2.24, 74.1) = 81.6$, $p < 0.001$, $\eta^2 = .712$). However, there was no significant interaction of Time and Stimulation ($F(4.49, 74.1) = 1.3$, $p = .263$, $\eta^2 = .075$). Moreover, no significant effect of Stimulation in the IeTI ($F(2, 33) = .399$, $p = .647$, $\eta^2 = .024$) was found.

The SDs across the groups are shown in Figure 2.3. There is a noticeable differ-

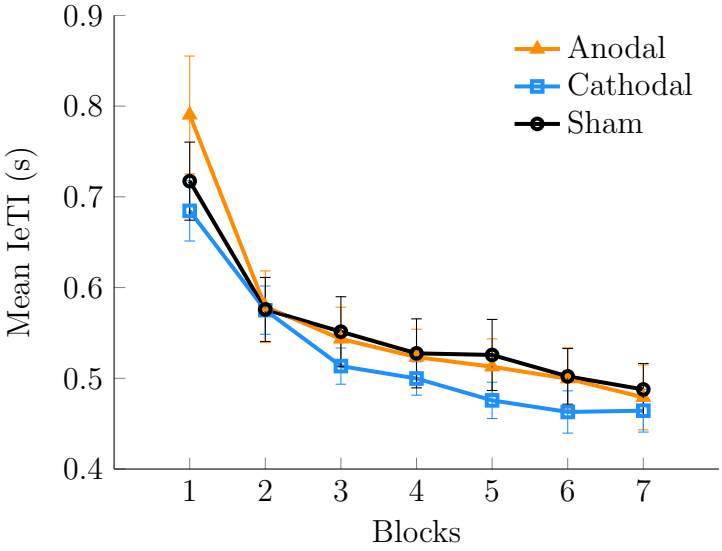

Figure 2.2: Mean IeTI (in seconds) in *Performance Phase* sub–blocks under the three different stimulation conditions. The mean IeTI drops over the the seven blocks but there is no evident difference among the three groups. The error bars show the standard error of the mean.

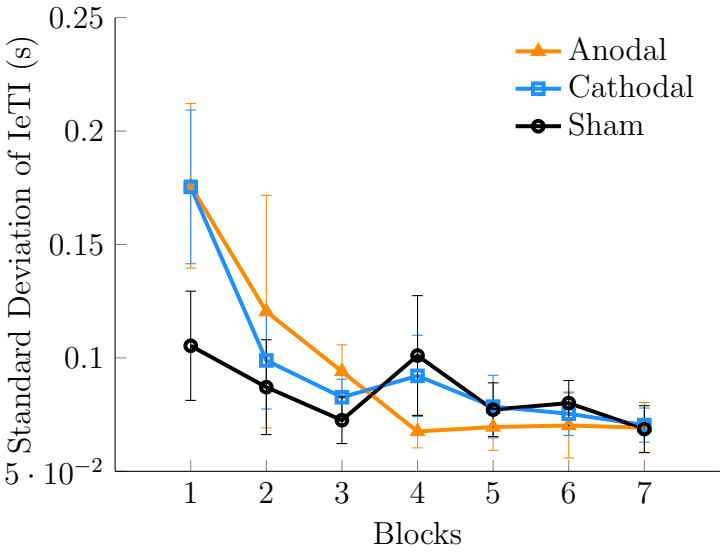

Figure 2.3: SD IeTI (in seconds) in *Performance Phase* sub–blocks under the three different stimulation conditions. Participants became more accurate over the course of the study. However, there is no significant difference among the groups. The error bars show the standard error of the mean.

ence of the variance in Block 1 between the stimulation groups and the sham group. Although, participants in the stimulation groups started on average with higher variance they managed to reach the same variance levels as the participants in the sham group. The ANOVA showed that the effect of Time was significant ($F(2.19, 72.3) = 14.54$, $p < 0.001$, $\eta^2 = .306$). However, there was no significant interaction of Time and Stimulation ($F(4.38, 72.3) = 1.9$, $p = .102$, $\eta^2 = .107$). Neither was there a significant effect of stimulation on the SDs ($F(2, 33) = .08$, $p = .923$, $\eta^2 = .005$).

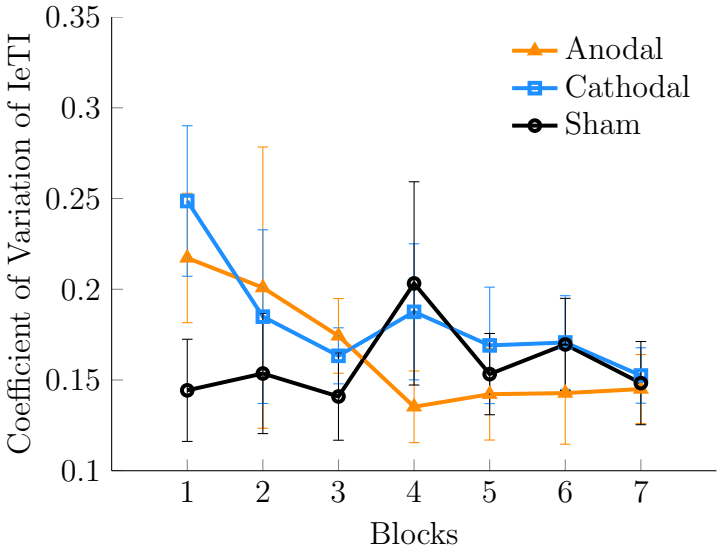

Figure 2.4: CoV IeTI in *Performance Phase* sub–blocks under the three different stimulation conditions. Participants became more precise over the course of the study. However, there is no significant difference among the groups. The error bars show the standard error of the mean.

Finally, Figure 2.4 presents the CoV across time in the three groups. This figure is quite similar to Figure 2.3. The statistical analysis showed that the effect of Time on CoV was significant ($F(3.57, 117.8) = 2.9$, $p = .029$, $\eta^2 = .081$). The interaction of tDCS over Time was not significant ($F(7.14, 117.8) = 1.8$, $p = .09$, $\eta^2 = .099$). However, we considered that this p–value ($p = .09$) was close to significance and we explored the simple main effects. We found that this close to significant value was

mainly driven because of a near significant ($p = .054$, *Sidak corrected*) difference of CoV between Block1 and Block 2 (all groups averaged). TDCS had no overall effect ($F(2, 33) = .225$, $p = .80$, $\eta^2 = .013$).

2.3.2 Fitted *Performance Phase* Data

Individual Data

The mean IeTI values of each participant across the blocks were used to fit learning curves. Figures 2.5, 2.6, and 2.7 show the results of the fitted learning curves for each participant of each group and Table 2.1 lists the parameters of each curve and the goodness of the fits.

The main observation of this analysis was that the behaviour of the participants in each group was quite variable. Moreover, the data points over which each curve is fitted were very few leading occasionally to non robust results. In some of the cases R^2 was low (Participant 2 in the Anodal Group, Participants 4 &10 in the Cathodal Group). The 3×7 ANOVA of the fitted on the learning curves data was similar to that of the mean recorded data. There was a significant effect of time ($F(1.32, 43.7) = 116.352$, $p < .001$, $\eta^2 = .903$) but no significant interaction of Time and Stimulation ($F(2.65, 43.7) = 1.3$, $p = .266$, $\eta^2 = .077$) and no significant effect of tDCS ($F(2, 33) = .398$, $p = .675$, $\eta^2 = .024$).

We also analysed the parameters of the fitted curves. For the learning rates (parameter b) of the fitted curves, Levene's Test for Homogeneity of Variances was significant ($\mathrm{F}(2, 33) = 9.409$, $p = .001$). Thus, we ran Kruskal – Wallis non parametric statistics test. There was no significant effect of the tDCS on the learning rates ($\chi^2 = 3.389$ with an associated probability of p=.184). For the a and c parameters, Levene's Test for Homogeneity of Variances was not significant ($\mathrm{F}(2, 33) = .425$, $p = .657$ and $\mathrm{F}(2, 33) = 1.518$, $p = .234$, accordingly). There was no signif-

icant effect of the tDCS on parameters a and c ($F(2, 33) = 2.51$, $p = .780$ and $F(2, 33) = 2.007$, $p = .150$, accordingly).

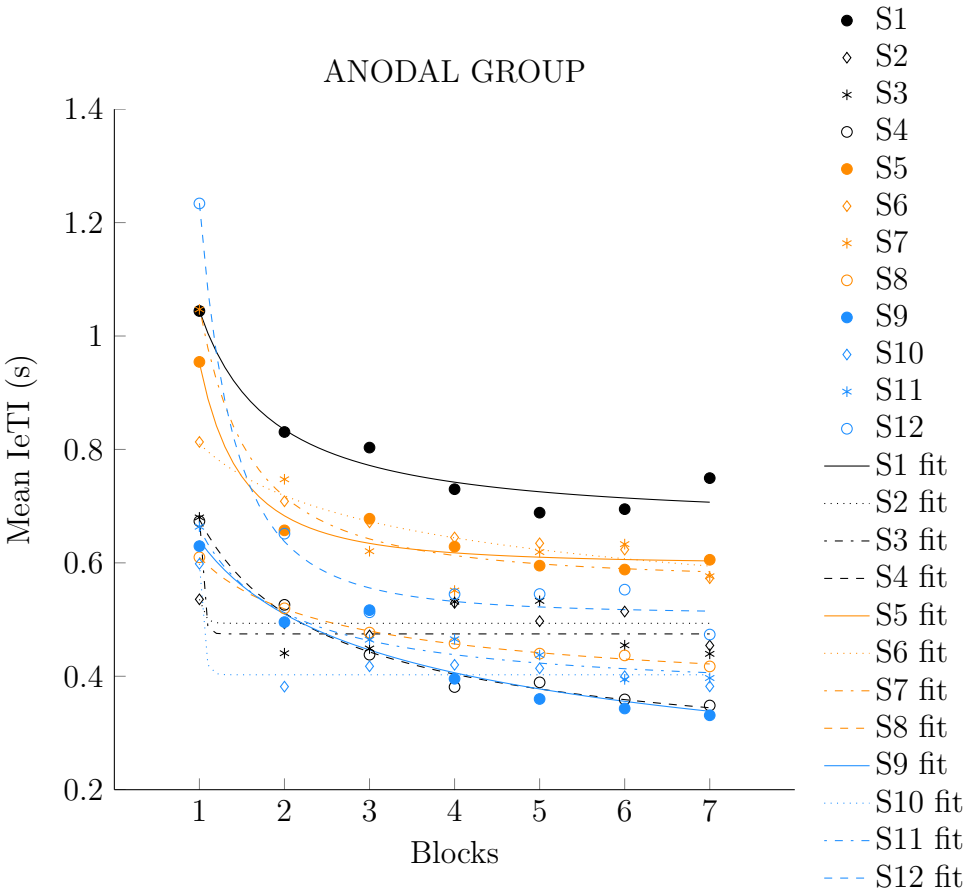

Figure 2.5: Fitted individual curves and recorded data of Mean IeTI in *Performance Phase* subblocks -Anodal Group. The lines are the fitted curves and the marker points are the calculated mean IeTI for each participant in each block.

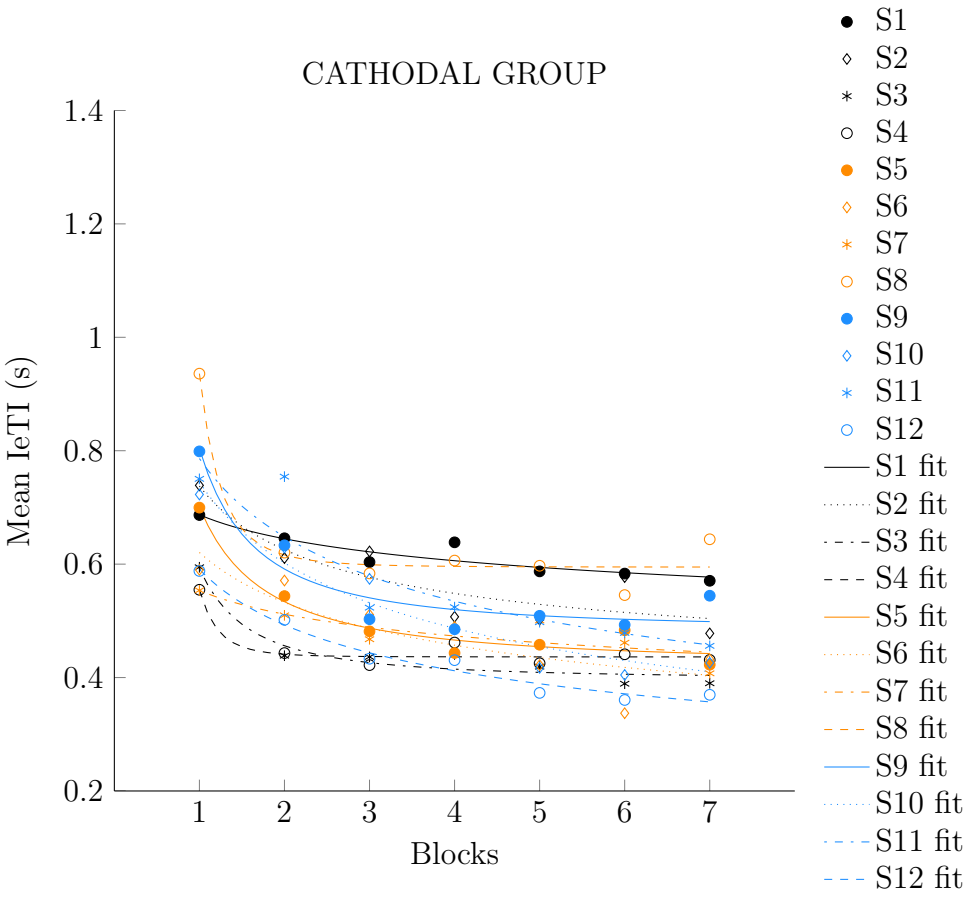

Figure 2.6: Fitted individual curves and recorded data of Mean IeTI in *Performance Phase* subblocks -Cathodal Group. The lines are the fitted curves and the marker points are the calculated mean IeTI for each participant in each block.

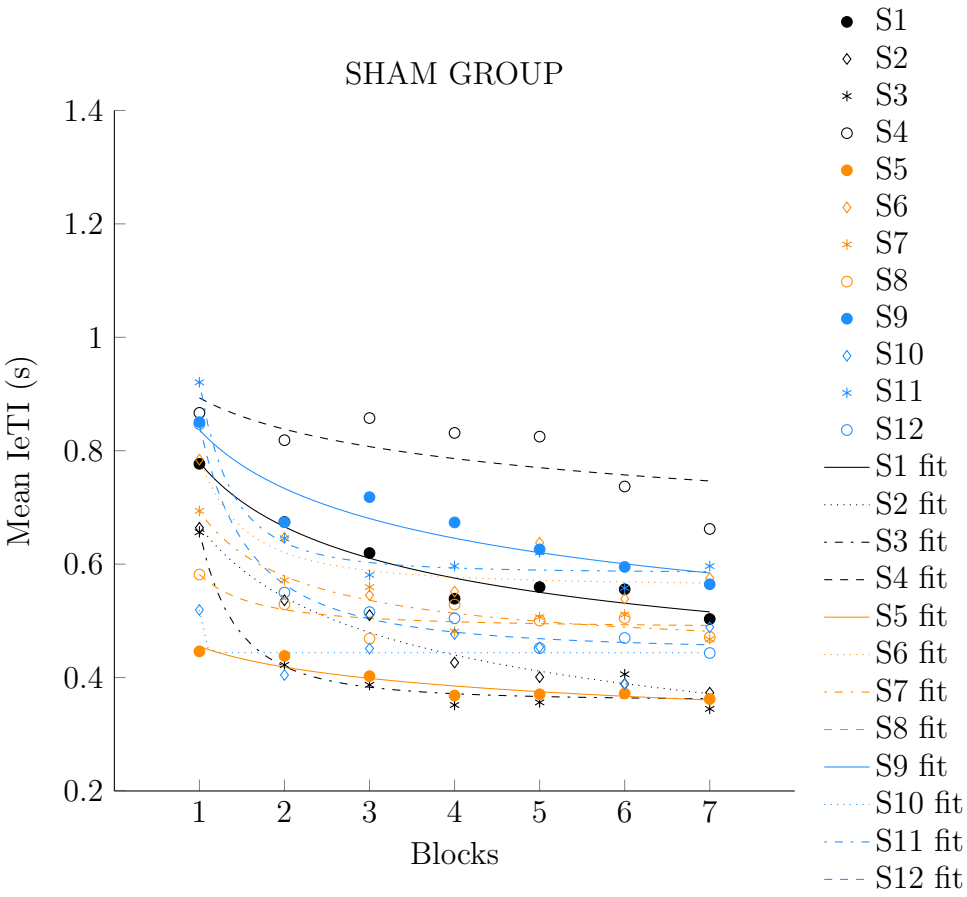

Figure 2.7: Fitted individual curves and recorded data of Mean IeTI in *Performance Phase* subblocks -Sham Group. The lines are the fitted curves and the marker points are the calculated mean IeTI for each participant in each block.

Table 2.1: Coefficients and Goodness of Fitted Curves of Individual participants in the Anodal (A) Group, Cathodal (C) Group and Sham (S) Group.

		Coefficients			Goodness of Fit	
Group	Participant	a	b	c	RMSE	R^2
A	1	0.376	-1.182	0.669	0.034	0.948
A	2	0.042	-20.529	0.493	0.031	0.292
A	3	0.205	-29.679	0.475	0.050	0.782
A	4	0.476	-0.613	0.200	0.015	0.989
A	5	0.356	-2.046	0.597	0.028	0.968
A	6	0.490	-0.298	0.320	0.015	0.973
A	7	0.484	-1.684	0.566	0.044	0.957
A	8	0.285	-0.555	0.325	0.004	0.997
A	9	0.639	-0.327	0.000	0.034	0.921
A	10	0.196	-25.348	0.403	0.020	0.956
A	11	0.296	-0.998	0.364	0.021	0.965
A	12	0.725	-2.500	0.509	0.038	0.987
C	1	0.500	-0.128	0.187	0.024	0.957
C	2	0.359	-0.540	0.378	0.018	0.981
C	3	0.195	-1.724	0.397	0.025	0.964
C	4	0.118	-4.645	0.436	0.064	0.504
C	5	0.281	-1.305	0.420	0.014	0.892
C	6	0.621	-0.221	0.000	0.046	0.816
C	7	0.502	-0.126	0.052	0.024	0.934
C	8	0.341	-3.989	0.595	0.027	0.685
C	9	0.320	-1.571	0.484	0.040	0.880
C	10	0.745	-0.306	0.000	0.044	0.383
C	11	0.786	-0.278	0.000	0.025	0.973
C	12	0.563	-0.276	0.028	0.019	0.987
S	1	0.562	-0.326	0.218	0.024	0.957
S	2	0.667	-0.301	0.000	0.018	0.981
S	3	0.297	-2.303	0.359	0.025	0.964
S	4	0.893	-0.092	0.000	0.064	0.504
S	5	0.425	-0.129	0.030	0.014	0.892
S	6	0.225	-1.925	0.561	0.046	0.816
S	7	0.267	-0.808	0.426	0.024	0.934
S	8	0.095	-1.543	0.487	0.027	0.685
S	9	0.692	-0.233	0.144	0.040	0.880
S	10	0.075	-22.887	0.444	0.044	0.383
S	11	0.336	-2.663	0.585	0.025	0.973
S	12	0.401	-1.740	0.444	0.019	0.987

Group Data

A final step to explore the data was to fit curves to the group data. Figure 2.8 shows the fitted curves of the group average mean IeTI. The fitting coefficients and the goodness of the fits are presented in Table 2.2. The Anodal group appears to have acquired the higher learning rate ($b = -1.475, R^2 = 0.9933$) while the Cathodal group the lower ($b = -0.6983, R^2 = 0.9949$). Although the R^2 values are high for all three groups, the 95% confidence bounds on the parameters (see Table 2.2) are quite overlapping (especially for the Cathodal and Sham group).

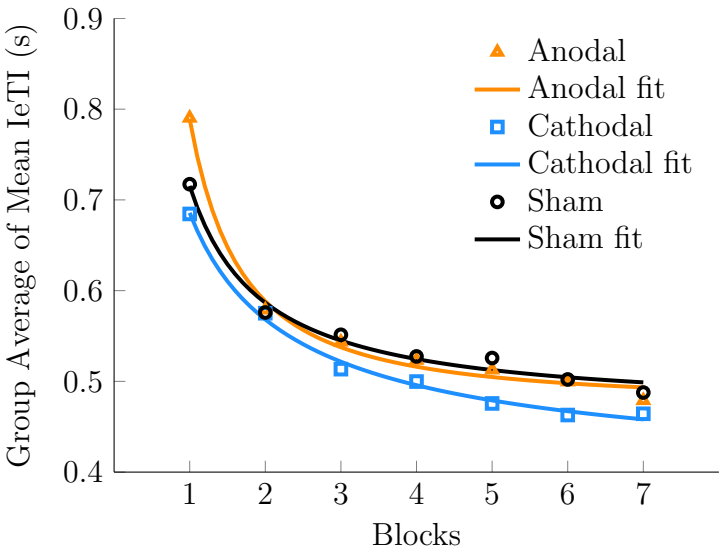

Figure 2.8: Fitted curves and recorded data of group average Mean IeTI in *Performance Phase* sub–blocks.

Table 2.2: *Performance Phase*: Coefficients of group fitted curves (with 95% confidence bounds) and goodness of the fits.

Group	**Coefficients** (95% **confidence bounds**)			**Goodness of Fit**	
	a	b	c	RMSE	R^2
Anodal	0.3133 (0.2693, 0.3574)	-1.475 (-2.056, -0.8935)	0.4756 (0.4386, 0.5126)	0.01061	0.9933
Cathodal	0.3063 (0.2214, 0.3912)	-0.6983 (-1.076, -0.3206)	0.3795 (0.2907, 0.4682)	0.007047	0.9949
Sham	0.2485 (0.1827, 0.3142)	-1.054 (-1.738, -0.3701)	0.4669 (0.4017, 0.5322)	0.01088	0.9867

2.3.3 Synchronisation Phase

The IeTI interval in *Synchronisation phase* sub–blocks, if participants were performing accurately, would be one second. The majority of the participants (after exclusions) managed to follow the imposed synchronisation frequency quite faithfully as shown in Figure 2.9. As expected there was no change of the mean IeTI over time ($F(1.72, 56.8) = .985$, $p = .275$, $\eta^2 = .038$). Moreover, the interactions showed that tDCS did not have any effect over Time ($F(3.44, 56.8) = .985$, $p = .415$, $\eta^2 = .056$). Finally, there was no significant difference among the groups ($F(2, 33) = 1.288$, $p = .289$, $\eta^2 = .072$).

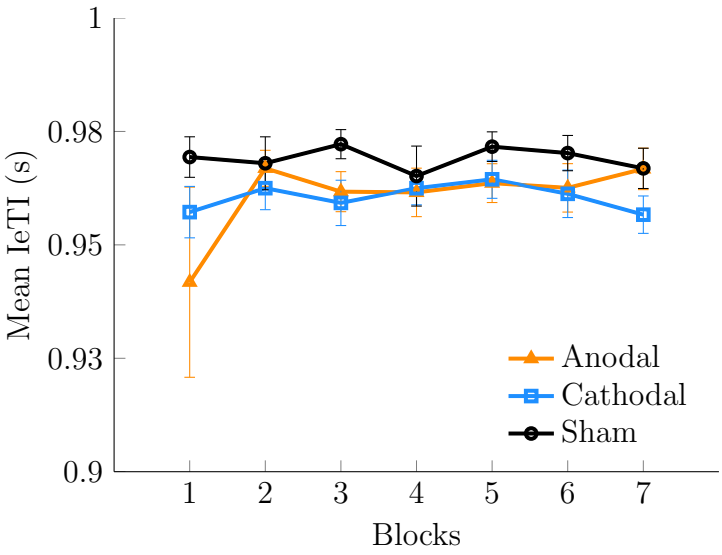

Figure 2.9: Mean IeTI (in seconds) in *Synchronization Phase* sub–blocks under the three different stimulation conditions. Participants had to synchronise to a 1Hz stimulus. The majority of them managed to follow the tempo in all three groups. The error bars show the standard error of the mean.

The group SDs in the IeTI during *Synchronisation phase* sub–blocks are presented in Figure 2.10. The variance over Time did not change ($F(4.06, 134.1) = 1.284$, $p = .279$, $\eta^2 = .037$)). The statistical analysis showed a significant interaction of tDCS and Time ($F(8.12, 134.1) = 2.036$, $p = .046$, $\eta^2 = .110$).

The simple main effects analysis revealed that the interaction arose because of a significant difference in Block 7 between the Anodal and the Cathodal group ($p = .018$, *Sidak corrected*) and between the Cathodal and Sham group ($p = .049$, *Sidak corrected*). In the same block the difference between the Anodal and Sham group was not significant ($p = .756$, *Sidak corrected*). No significant differences were found in Block1. TDCS did not affect the variance across the groups ($F(2, 33) = .615$, $p = .547$, $\eta^2 = .036$).

Analogous observations can be conducted by inspecting the CoV (at group level) of IeTI in Figure 2.11. Time did not play a role in the participants' performance ($F(3.03, 100.0) = 1.074$, $p = .346$, $\eta^2 = .032$). The interaction of tDCS and time was close to be significant ($F(6.06, 100.0) = 1.979$, $p = .075$, $\eta^2 = .107$). The analysis of the simple main effects showed that the effect was driven by a significant difference ($p = .018$, *Sidak corrected*) between Anodal and the Cathodal group in Block 7.

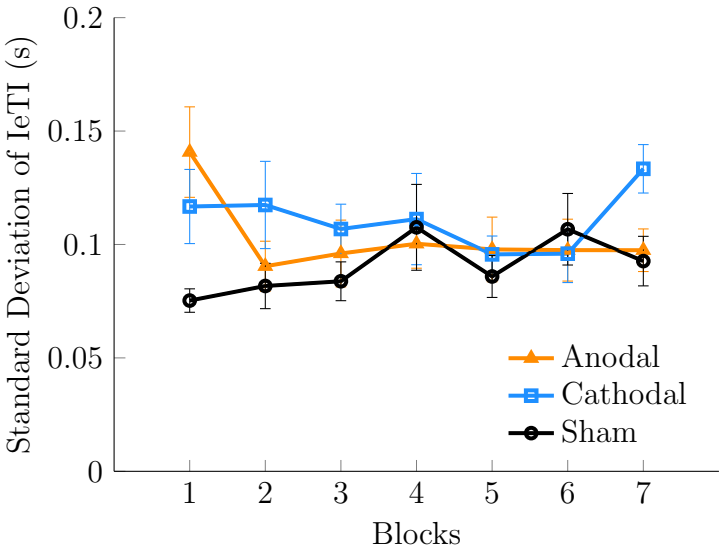

Figure 2.10: SD IeTI (in seconds) in *Synchronization Phase* sub–blocks under the three different stimulation conditions. The variance of participants appears to be different among the different groups. The error bars show the standard error of the mean.

TDCS as a between participants factor had no effect on the CoV ($F(2, 33) = .59, p = .560. \eta^2 = .035$).

In the *Synchronization Phase* sub–blocks we were also interested in the Intra Tap Intervals (IaTI). Figure 2.12 presents the mean IaTI. There was no improvement of the performance over time ($F(2.35, 77.7) = 1.402$, $p = .252$, $\eta^2 = .041$) and tDCS did not interact with Time ($F(4.71, 77.7) = 1.085$, $p = .374$, $\eta^2 = .062$). The effect of tDCS across the groups was found marginally significant ($F(2, 33) = 3.198$, $p = .054$, $\eta^2 = .162$). Tukey's HSD post hoc analysis revealed no significant interactions. The marginally significant effect seems to have arisen because of a difference between the cathodal and sham groups ($p = .090$) and a difference between sham and anodal groups ($p = .086$). As depicted in Figure 2.12 these differences are towards the same direction and this is why the difference between anodal and cathodal groups was

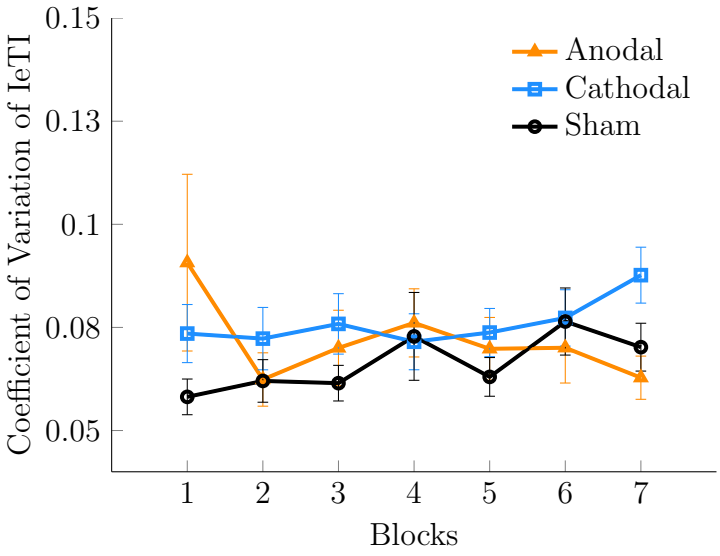

Figure 2.11: Coefficient of Variation of IeTI in *Synchronisation phase* sub–blocks. Participants in the sham group seem to have been more consistent across all blocks, both in terms of mean CoV across time but also in terms of variance. This was not the case for the cathodal and especially the anodal group. The error bars show the standard error of the mean.

.

not significant ($p = 1.00$).

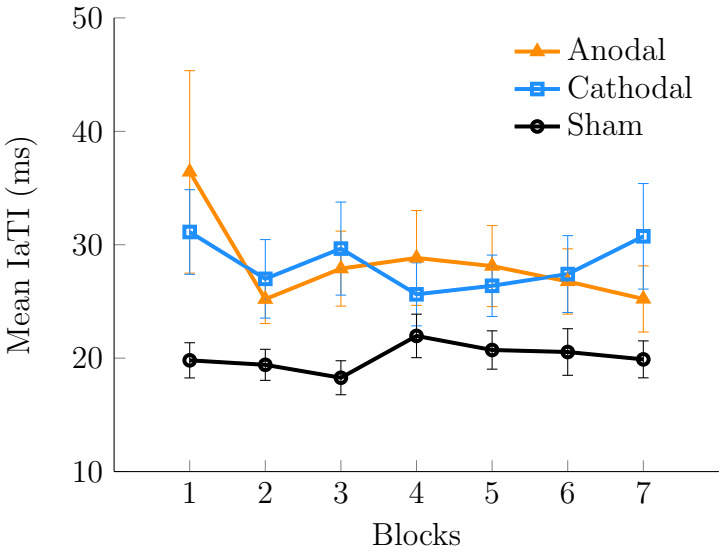

Figure 2.12: Mean IaTI (in milliseconds) in *Synchronisation phase* sub–blocks. Participants in the sham group appear to have been able to coordinate better the pair of fingers used. This was not the case for the cathodal and anodal groups. The error bars show the standard error of the mean.

.

Figure 2.13 shows the SD of IaTI in the three groups. There was no improvement of the variance over time ($F(1.98, 65.3) = .732$, $p = .483$, $\eta^2 = .022$), neither an interaction between tDCS and Time ($F(3.96, 65.3) = .976$, $p = .426$, $\eta^2 = .056$). However, we can see a great variability in the Anodal and Cathodal Group, depicted in a marginally significant effect of tDCS ($F(2, 33) = 3.213$, $p = .053$, $\eta^2 = .163$). Tukey's HSD post hoc analysis showed no significant effects between the groups ($p = .984$ between Anodal – Cathodal, $p = .104$ between Cathodal – Sham and $p = .073$ between Sham and Anodal group).

Finally, the CoV of the IaTI (Figure 2.14) did not reveal any effect of time ($F(4.65, 153.5) = .630$, $p = .661$, $\eta^2 = .019$), nor an interaction of tDCS over Time ($F(9.30, 153.5) = .952$, $p = .952$, $\eta^2 = .022$). TDCS did not impose any differences across the groups in the CoV ($F(2, 33) = .765$, $p = .473$, $\eta^2 = .044$).

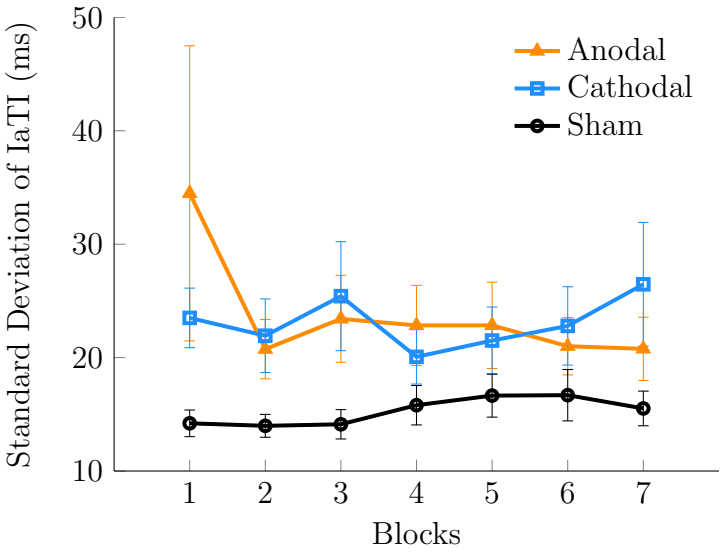

Figure 2.13: SD IaTI (in milliseconds) in *Synchronisation phase* sub–blocks. Participants in the sham group were more consistent. The performance of the participants in anodal and cathodal groups was much more variable. The error bars show the standard error of the mean.

.

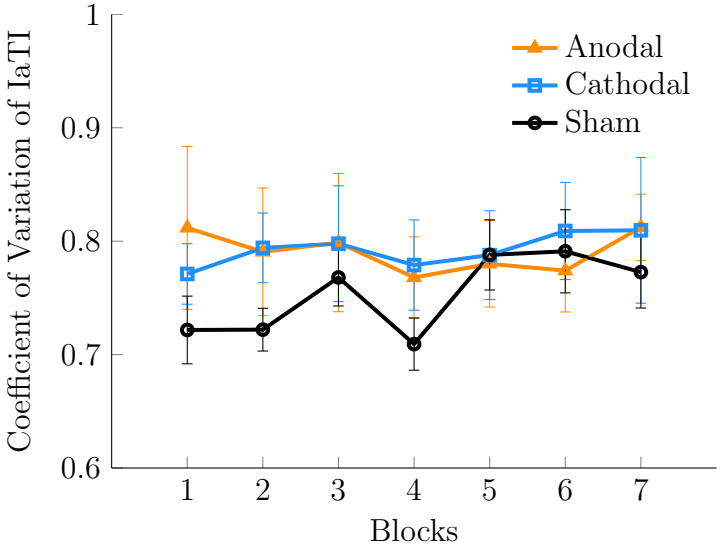

Figure 2.14: CoV IaTI in *Synchronisation phase* sub–blocks. The error bars show the standard error of the mean.

.

2.4 Discussion

The study was designed to test whether the application of tDCS over the cerebellum would induce the same effect in a novel motor task as in a visuomotor adaptation task (Galea et al., 2011). No significant effect of cerebellar tDCS modulation was revealed in the *Performance* phase sub–blocks. However, the results seem to imply a faster learning rate in the Anodal group. Moreover, tDCS seem to have influenced the performance of participants during the training (*Synchronization*) sub–blocks by increasing the variance of the IaTI. This result could have been partly expected. rTMS on the medial cerebellum (but not the lateral) was shown to increase variability in a paced single finger tapping task (Thoret et al., 2001). All in all, it would be unwise to interpret the results as an indication of the cerebellum being heavily engaged –or not– in the learning of this novel motor task. This is because of weaknesses in the design of the study but also because of the insufficient knowledge of tDCS mechanisms.

One of the drawbacks of the design was the lack of a baseline measure of the participants performance before the initiation of tDCS. This choice was made on purpose to ensure that participants were completely naive to the task and there was no learning before the application of tDCS. Thus, the differences observed in the first block of mean IeTI during the *Performance* phase sub–blocks could have been a result of the application of tDCS rather than a baseline difference in the groups' behaviour. On the other hand, the measures during the synchronisation condition could depict group differences rather than a tDCS effect. Another major weakness of the design was its duration. Unlike adaptation, motor learning is a slower procedure that demands more time and is perhaps difficult to capture over the course of some minutes. To our knowledge the only study that used tDCS over M1 to study motor

skill acquisition was conducted over the course of several days (Reis et al., 2009). The 20 minutes of having used tDCS here may have been insufficient to draw out differences in learning.

The marginally significant effect of the increased variance of IaTI in the anodal and cathodal groups during the training (*Synchronization*) sub–blocks could be an indication of different mechanisms by which the cerebellum contributes to motor learning and to motor control. Variance is not necessarily an indication of worse performance. It could as well be an indication of an exploration phase that learning undergoes before establishing the desired new motor patterns (see Chapter 3). Moreover, provided that the mean fitted curves of the mean IeTI in the *Performance* phase sub–blocks have captured some true tendencies (Figure 2.8), we could speculate that whereas anodal tDCS group showed a faster learning rate, there seems to be a thresholding over time. On the other hand, cathodal group had a slower learning rate but no thresholding is apparent. The behaviour of sham group could be a trade–off between learning rate and stabilisation of some sort of motor pattern. Finally, it is possible that our study is simply underpowered. A number of publications have used considerably large groups (Pope and Miall, 2012).

tDCS Effects

Based on the results of previous studies of cerebellar tDCS that involved a motor task (Jayaram et al., 2012; Galea et al., 2011; Ferrucci et al., 2013), even though the stimulation electrode was centred $3cm$ lateral and $1cm$ below the inion (Miall et al., 2007), it is to some extent likely that tDCS in our study did affect the cerebellum. In Jayaram et al. (2012) and Galea et al. (2011) the elcetrode was centred just cross $3cm$ lateral to the inion and in Ferrucci et al. (2013) 2 cm below the inion (centred on the median line) with its lateral borders about 1 cm medially to the mastoid

apophysis. So, why did we fail to see a pronounced effect of tDCS on the learning behaviour? As mentioned before, it might be that in order to see an effect of tDCS on learning we need to use a protocol over several days (Reis et al., 2009). Moreover, another critical component of tDCS application is whether it is applied before, or during the execution of the task (Stagg et al., 2011). In our study, we chose to apply tDCS for ten minutes before the initiation of the task and keep it on during the execution of the study. Interestingly, in the studies of Jayaram et al. (2012) and Galea et al. (2011) the tDCS stimulation and the adaptation tasks started at the same time. In both, studies, differences in the learning rates were observed in the initial stages of adaptation but not later on. This could imply that the initial interference of tDCS with the cerebellum is critical and we could have lost it by starting the task several minutes later. However, our task was different in nature. To completely understand and explain why tDCS did not drive significant changes in the performance we need to integrate the timing properties of tDCS with the timing and dynamics of the contribution of cerebellum in learning.

In order to understand how tDCS might affect cerebellum we need to take a look at the proposed mechanisms underlying the interaction of tDCS with the neuronal populations. One of the interpretations given, based on protocols that used Transcranial Magnetic Stimulation (TMS) to access the excitability of motor cortex after the application of tDCS (Nitsche and Paulus, 2000), indicates that it is possible that tDCS modifies neuronal excitability in a polarity specific way; anodal tDCS enhances the excitability whereas cathodal tDCS has opposite results. In a pharmacological approach, Liebetanz et al. (2002) suggested that polarity-driven alterations of resting membrane potentials represent the crucial mechanisms of the DC–induced after–effects, leading to both an alteration of spontaneous discharge rates and to a change in NMDA–receptor activation.

Another possible interpretation of tDCS effects is that tDCS interferes with mechanisms of homeostatic plasticity. The homeostatic regulation of plasticity prevents neurones becoming hyper– or hypo– polarised. The changes occur relatively slowly so as to not dampen the moment–to–moment fluctuations, but to keep up with the slow changes in drive produced by other plasticity mechanisms (Turrigiano and Nelson, 2004). In this context, the timing of tDCS and its repetition seem to allow different processes to be developed. Monte-Silva et al. (2010) applied cortical cathodal tDCS. They showed that if a second stimulation is performed during the after–effects of the first stimulation, it prolongs the effects of tDCS compared to doubling the duration of stimulation without interstimulation break. Moreover, when a second stimulation was performed several hours after the first, the effects of cathodal tDCS were attenuated. Monte–Silva and colleagues interpreted their results as indicative of stimulation timing–dependent plasticity regulation in motor cortex. In an other study, Fricke et al. (2011) used a similar protocol to that of Monte-Silva et al. (2010). Although the results of these two studies are different because of the duration of stimulations and the repetition intervals, the effects of tDCS seem to be governed by rules of homeostatic plasticity. Siebner et al. (2004) preconditioned the application of 1Hz rTMs with tDCS. They found that changing the initial state of the motor cortex by tDCS reversed the conditioning effects of rTMS, suggesting the existence of homeostatic plasticity mechanisms. L–type voltage–gated Ca^{2+} channels (L–VGCC) seem to play a key role in metaplasticity mechanisms, that among others seem to regulate homeostatic plasticity. Wankerl et al. (2010) showed that blocking L–VGCCs reversed the effect of continuous Theta Burst Stimulation (cTBS). In another study, Hasan et al. (2012) used a protocol that combined tDCS and cTBS. It was shown, among other effects, that cathodal tDCS suppressed and even reversed the effects of cTBS. Thus, it is possible that tDCS can also have an

effect on Ca^{2+} dynamics and as a result it can affect plasticity mechanisms involved in learning.

The only study that has explored explicitly the interaction of tDCS with the cerebellum is that of Galea et al. (2009). They assessed cerebellar–brain Inhibition (CBI) by paired pulses of TMS and found that anodal tDCS enhanced the CBI whereas cathodal tDCS decreased it. They interpreted their results based on neuronal excitability changes, in that anodal tDCS increased the excitability of the inhibitory Purkinje cells (PC) leading to increased CBI, whereas cathodal tDCS further inhibited the output of the cerebellum leading to a decreased CBI. However, there are no studies that have studied explicitly the effect of tDCS in cerebellum in protocols (similar to the ones mentioned above for the motor cortex) using different durations of stimulation and various repetition intervals that could relate homeostatic plasticity mechanisms involved in cerebellar tDCS. Nonetheless, PCs do have voltage–gated calcium channels (although of a different P–type) that are involved in the synaptic regulation of PCs activity and could be affected by tDCS. On the other hand, PCs do not have NMDA receptors, which makes it ambiguous to interpret the results of cerebellar tDCS in a similar way to those of cortical tDCS studies.

The aforementioned studies did not include any behavioural task. However, it is really important to think of the changes imposed to cortical excitability by the afferent signals during a motor task. Nitsche et al. (2007) induced facilitatory associative plasticity by Paired Associative Stimulation (PAS). They found out that changes in background activity induced by tDCS had an effect on the efficacy of PAS. More interestingly, if tDCS was applied before the application of PAS it resulted in an opposite effect compared to concurrent application of PAS and tDCS in accordance with homeostatic plasticity mechanisms. Analogous findings were observed by Gentner et al. (2008) in a study where they used continuous Theta Burst Stimulation

(cTBS). Prior voluntarily motor activation reversed the effects of cTBS.

To sum up, there is indication suggesting that tDCS affects both neuronal excitability and mechanisms of homeostatic plasticity. Moreover, the effect of tDCS changes depending on the afferent input to the stimulated area during a motor task. All these reasons make crucial the timing of applying tDCS and practising a task during a study. However, given the cytoarchitectural and functional differences among brain areas, tDCS can be inducing differential effects over different brain areas.

Generally speaking, the polarity specific effect of tDCS in motor studies seems to be quite consistent; cortical anodal stimulation (in motor areas) enhances excitability whereas cathodal tDCS has the opposite effect (see Jacobson et al. (2012) for a meta – analytical review). However, it is very difficult, if not impossible, to draw any further conclusions about the interaction of tDCS with behavioural tasks. Table 2.3 shows a series of studies that involved motor adaptation and learning. The important parameters of tDCS application are the current intensity (mA), the duration of stimulation (minutes), the brain area of interest, the area where the reference electrode is placed, the sequence of the application of tDCS and the task performance. It is apparent that the inhomogeneity in the protocols invalidates any comparison.

Conclusions

In the current study, we wanted to investigate the involvement of the cerebellum in learning a novel motor skill. Our results are inconclusive mainly because of the lack of understanding of tDCS mechanisms but also due to some weaknesses of the design. We propose that tCDS needs to be used in paradigms that are well studied and give clear and pronounced behavioural measurements. The uncertainty surrounding the

Table 2.3: Effect of tDCS in motor learning studies.

Authors	Task	tDCS					Results
		mA	minutes	Sequence	Area	Reference	
Nitsche et al. (2003)	SRTT Implicit Motor learning (IML)	1	15	During	MC, PMC, Lateral/Medial Prefrontal Cortex	Contralateral supraorbital area	AtDCS in MC improved performance in the acquisition and early consolidation phase of IML
Antal et al. (2004)	Visuomotor coordination task	1	10	During	M1,V5 Primary Visual cortexl	Cz	AtDCS improved in early phase of learning in V5 & M1
Hunter et al. (2009)	Force Field Adaptation	1	17	During	M1	Supraorbital Area	AtDCS induced overshoot errors in deadaptation phase
Reis et al. (2009)	$S_{equential}V_{isual}I_{sometric}$ $P_{inch}T_{ask}$	1	20	During	M1	Supraorbital area	AtDCS enhanced skill acquisition & total learning
Galea et al. (2011)	Visuomotor adaptation	2	15	During	M1 Cerebellum	Supraorbital area , Right Buccinator muscle	Crebellar AtDCS enhanced adaptation/ M1 AtDCS increased retention
Stagg et al. (2011)	SRTT Explicit Motor learning (EML)	1	10	Before & During	M1	Contralateral supraorbital area	tDCS during task modulated learning rates in a polarity specific way, tDCS decreased rates prior to learning
Schambra et al. (2011)	$S_{equential}V_{isual}I_{sometric}$ $P_{inch}T_{ask}$	1	20	During	M1	Ipsilateral Deltoid Muscle	Left M1 tDCS improved acquisition in both hands
Orban de Xivry et al. (2011)	Force Field Adaptation	1	20	During	M1 PPC	Contralateral Supraorbital area	A&C tDCS in M1 increased generalisation in intrinsic cordinates
Jayaram et al. (2012)	Locomotor Adaptation	2	15	During	Cerebellum	Ipsilateral Bucciinator Muscle	AtDCS increased CtDCS slowed down the adaptive process
Ferrucci et al. (2013)	SRTT Implicit Motor learning	2	20	Before & 35min Post	Cerebellum	Right arm	AtDCS improved implicit motor learning (IML)

mechanisms of tDCS discourages the application of tDCS in combination with novel motor paradigms as a way to explore basic motor functions. On the other hand, the use of tDCS in well studied paradigms can reveal interactions useful for the use of tDCS as a rehabilitation tool.

Chapter 3
THE EFFECT OF TRACKING IN A NOVEL EXPLORATION TASK

3.1 Introduction

An important question relative to reinforcement learning is how we discover novel actions (Redgrave et al., 2008). A critical component of action acquisition is to be able to associate which part of the movement was responsible for the rewarding outcome. It has been suggested (Redgrave et al., 2008) that short–latency phasic dopamine neurones in basal ganglia provide a "time–stamp" that could serve to associate the delivery of rewards with preceding actions.

Recently, Stafford et al. (2012) developed a paradigm that focused on how actions are acquired rather than how predetermined actions are modified depending on the frequency of reinforcement signals. In this task participants are required to explore a given area until they find an action that brings an expected outcome (reinforcement). This procedure requires several exploratory actions before choosing a final one as being responsible for the delivery of the reinforcement. Thus, it is a suitable tool for studying several behaviours related to reinforcement learning rules such as the exploration–exploitation trade–off (Sutton and Barto, 1998). Now, if a reward is delivered with a delay then the last part of the movement that is not contingent to the behaviour leading to the reward is accidentally reinforced, giving rise to so–called "superstitious behaviour" (Skinner, 1948). This non–contingent motor output has a negative impact during learning of new actions (Walton, 2011).

Delays severely affect behaviour in paradigms of motor adaptation and control, as well. For example, delayed visual feedback in manual tracking impairs the performance of the participants (Foulkes and Miall, 2000; Miall and Jackson, 2006). However, participants show improvements in performance both in terms of a reduc-

tion in the error between the hand and the delayed hand representation but also in terms of performing smoother movements. Adaptation to tracking delays is mediated by the cerebellum via the update of predictive forward models (Miall and Jenkinson, 2005).

It was recently shown that cerebellum and the basal ganglia are anatomically connected in a disynaptic way (Bostan et al., 2013). In the present study we aimed to investigated if we could observe any behavioural interactions between two tasks where the one is cerebellum related and the other basal ganglia related. Our motivation was to explore the nature of behavioural mechanisms that could emerge from the anatomical connections between the cerebellum and the basal ganglia.

To do so, we combined the exploration task designed by Stafford et al. (2012) as a task that is highly related to the functions of basal ganglia with a manual tracking task as described by (Foulkes and Miall, 2000) as a task related to the cerebellum. The main question was if the superstitious behaviour that arises from a delayed delivery of the reinforcement during the exploration task could be alleviated if participants adapted their motor behaviour to the same delays during the tracking task. In other words, after adaptation to the delays in a certain environment (manipulation of a robotic arm in our case) can relevant actions (and not non–contingent motor output) be better associated to an outcome?

We hypothesised that participants can perform better in an exploration task, where detection of a target is via a tool whose state is uncertain because of the delay, if they can first learn the relationship between tool movements and outcomes, during the tracking task.

3.2 Materials and Methods

3.2.1 Participants

30 right handed subjects (age range: 18–45, mean=22, female:22) participated in the study. All of them had normal or corrected to normal vision. They did not have any restricted mobility or suffer form any neurological condition. They were informed about all the aspects of the experiment and gave informed written consent. The experimental protocol was approved by the Science, Technology, Engineering and Mathematics Ethical Review Committee of the University of Birmingham Ethics Committee. Participants received either cash or credits upon the completion of the study.

3.2.2 Apparatus

A vBOT robotic manipulandum (Howard et al., 2009) was used as an input device. An 30-inch Apple Cinema HD Display monitor was used to display visual objects viewed via the projection mirror surface (see Figure 4.1 below).

Subjects were seated on a stool and held the robotic arm with their right arm. The height of the stool was adjusted so that the shoulders of the subject were at the same height as the horizontal projection surface they were looking at. This allowed an almost fixed body posture throughout the study. The vision of the arm was obscured during the study.

The background of the experimental area displayed on the monitor was black. A white circle of 13 cm radius marked the experimental area where the visual objects appeared. Participants were physically constrained in the circle by a high force field: $F[i] = 30 * (13 - |x|) * \frac{x[i]}{|x|} - 0.015 * v[i]$, where i is the x and y components, F is the Force, p the position and v the velocity. This gave the feeling of bumping into

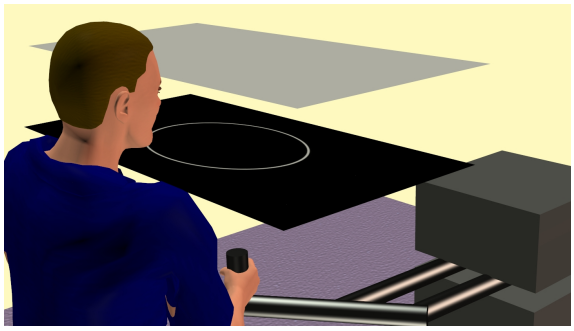

Figure 3.1: The Set–up. Subjects the robotic arm with their right arm. A monitor (grey surface) displayed visual objects viewed via the projection mirror surface (black surface).

a wall colocated with the circle. The velocity dependent term provided damping.

The experiment consisted of two different tasks; an exploration task and a tracking task.

3.2.3 Exploration Task

During the exploration task, participants had to explore the circular workspace and find a hidden target area (Stafford et al., 2012). At the beginning of each trial a white cross appeared at the centre of the workspace and participants had to bring the vBOT on the cross. Once within 0.5cm of the centre, a black dot appeared on the cross indicating that the hand reached the starting position. After that the cross and the dot disappeared and participants had no longer visual feedback of their forearm. They were then free to start exploring the workspace. They were instructed that somewhere in the workspace there was a circular target area. They had to explore within the experimental area and try to find the target area. When they placed the vBOT in the target area, they received non–spatial reinforcement for a brief flash of the whole workspace. However, contrary to what they were instructed, participants received the reinforcement flash **300ms after** they encountered the target area (see Appendix A.1 for the choice of the delay). They then returned to the central cross.

For the next nine trials the target area remained at the same location and they had to relocate the target area *as efficiently as possible and within as short a time as possible*. It is also worth mentioning that the participants were not given any specific instructions on the strategy they had to use to explore the workspace. Each trial lasted a maximum of 30 seconds. If they had not located the target within that period they returned to the central cross. After 10 trials the target area changed location and participants were instructed to re–explore the workspace to find the new location. From now on we will refer to each set of 10 trials at the same target area location as a Batch of the Exploration task. The size of the target area was 0.25% of the size of the total workspace (see Appendix A.1 for the choice of the size of the target area). The distance of the target area from the centre of the workspace varied between $5 - 11cm$.

3.2.4 Tracking Task

In the tracking task, participants had to pursue a circle target (4mm in diameter) that was smoothly moving in an unpredictable 2-D path. They were to track the target with a cursor (6 mm in diameter) that reflected the position of the vBot handle. An example of such a target trajectory is shown in Figure 4.2. Each tracking trial lasted for 45 seconds. At the end of each trial participants returned to the centre of the workspace and they were given three seconds of rest. Moreover, they could take a longer break between the blocks of tracking trials. In this task participants did have visual feedback of their hand and they were instructed to track the target as accurately as possible. The pseudorandom target trajectories were generated in each axis independently as the sum of four non–harmonic sinusoids (0.11, 0.23, 0.35 and 0.42Hz) whose relative phases were randomised for each trial (Foulkes and Miall, 2000). This means that the target was moving with *variable speeds*. The speeds at

which the target was moving were between 1-20 cm/sec, with an average speed at around 8.5cm/sec.

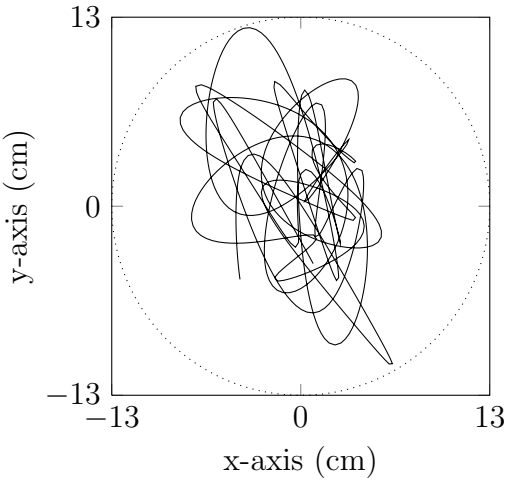

Figure 3.2: Example of a tracking target trajectory. The circle encloses the workspace.

Conditions. Participants were randomly assigned to two different Tracking Groups. Participants in Group A practised tracking without any visual feedback delays imposed on their hand cursor. However, participants in Group B were presented with a cursor that represented a 300ms delayed version of their hand position.

When there were visual delays, the forces developed on the circle at the edge of the workspace were those of a viscous force field: $F[i] = -0.015 * v[i]$, where i is the x and y components, F is the Force and v the velocity. We used a viscous force field, without a still position dependent wall, so that the delayed visual and proprioceptive sensation were approximately aligned. However, the participants were not supposed to experience these forces as they were expected to remain within the workspace, and in fact only rarely crossed over the circular boundary.

3.2.5 Design

Participants came to the lab on two subsequent days and repeated the same protocol on both days. They first completed 10 Batches of the Exploration task (from now on referred to as an Exploration Block) which lasted on average 12 minutes. They then practised the tracking task for an hour, completing 3 Blocks of 24 tracking trials each. Finally, they did one more Block of the Exploration task. The design is summarised in Table 3.1.

Table 3.1: Experimental Design: Participants completed the same experimental protocol on two subsequent days. Those who were assigned to Group A completed the tracking task without visual feedback delays, whereas those in Group B tracked with visual feedback delays imposed on their hand cursor. Both Groups received the reinforcement in the Exploration task with 300ms delay.

Task	Exploration	Tracking	Exploration
	10 Batches x 10 Trials	3 Blocks x 24 Trials	10 Batches x 10 Trials
Day 1 **Day 2**	**Block1** **Block3**		**Block2** **Block4**
Delays	Reinforcement	Visual Feedback	Reinforcement
Group A	300ms	0ms	300ms
Group B	300ms	300ms	300ms

3.2.6 Data Analysis

Matlab R2011b was used to process the data recorded from the vBots. Unless otherwise indicated, SPSS was used for running the statistical analysis.

Tracking task

Two different indices of performance were used to evaluate the improvement of the tracking behaviour. The first one was the RMS error between the target and the cursor (tracking error). Adaptation to tracking delays would be indicated by

reduction of the tracking error. For each tracking trial, the RMS error between the target and the cursor was calculated after excluding the first 2 seconds recorded.

The second index of performance was the Power Spectral Density (PSD) of speed. Participants were expected to perform smoother movements while adapting to the delays. This would show on the PSDs as a reduction of the spectral content. The de–meaned positional data were smoothed using a Hanning filter. A 4th order zero phase low pass Butterworth filter with a cut-off frequency at 10 Hz was used to filter high frequencies. We then calculated the velocity vector by taking the gradient of the position data independently in each axis. The PSD was taken using a Fast Fourier Transform. The mean PSD was then calculated in the frequencies between 0.1Hz and 3Hz.

For each participant we calculated the mean error and average mean power spectral density in 6 Blocks of 24 trials each (three Blocks were executed the first day and three on the second day of the study). The mean error and the average mean PSD were analysed separately for Group A (no visual feedback delays in tracking) and for Group B (visual feedback delays in tracking), using a repeated measures ANOVA with one within subject factor of *Block* (time) with 6 levels (Block 1 to Block 6).

Exploration task

Improvement or deterioration of performance in the exploration task could be described by several different indices.

- **Mean Irrelevant Distance.** The first index of performance we used was the Irrelevant Distance that the participants travelled until they encountered the Target Area. The Irrelevant Distance was defined as the Raw Distance travelled until reaching the target minus the Euclidean distance from the centre of the

workspace to the target (Walton, 2011). This metric takes into account the between–block different distances of each target area from the centre of the workspace. Generally, it is expected that participants, once they discover the target area, will achieve lower Irrelevant Distances during the revisits to the same target within a batch. The optimal behaviour in the task is zero Irrelevant Distance, meaning that participants travel in a straight line from the centre of the workspace to the target area. We hypothesised that if participants who tracked with visual feedback delays became aware of / or adapted to the delay, and were able to apply this information to the exploration task, they would be better able to locate the target area despite the reinforcement delay and so they would find a shorter path to it. That would be depicted in lower Irrelevant Distances for Group B compared to Group A (no tracking delays).

- **Post–Discovery Distance**. If participants in Group B inferred the exact location of the target area, as a result of the adaptation to the tracking delays, we would expect them to stop on the target area and wait for the delayed reinforcement without travelling any further. To test this hypothesis we calculated the Post–Discovery Distance, which is defined as the distance travelled from the moment the target area was encountered until the moment the reinforcement flash was received (i.e. the end of the trial).

  However, to accurately interpret the results of the analysis of this variable, we had to investigate if the average speed adopted by the participants across the Blocks of the Exploration task remained steady. Unless participants maintained an equal speed during all Blocks of the Exploration task, then any potential differences exposed in the Post–Discovery Distance analysis could be attributed to either their change in speed or, to their adaptation to the delay.

- **Overall Speed.** We calculated the Overall Speed of each trial by taking the ratio of the total distance travelled over the duration of the trial. One reason for examining this variable was to interpret the results of the Post–Discovery Distance analysis, as explained above. Moreover, the Overall Speed could expose any effects of the tracking task to the exploration task. An expected effect could be that the participants who tracked with delays would lower their speeds in the exploration task, to compensate for the consequences of the delay. That is, if they had adapted to tracking with delayed visual feedback and if they had attributed it as a property of the vBot generalised across tracking and exploration. Another possible effect of the tracking task, with or without delays, could be that participants shifted their speed during the exploration task towards the average target speeds of the tracking task.

- **Post–Discovery Speed.** We calculated the Post–Discovery Speed of each trial by taking the ratio of the total Post–Discovery Distance travelled over 300ms (reinforcement delay = post–discovery duration). The motor output produced in the post–discovery period (non–contingent output) has a great impact on learning a new action, as mentioned in the introduction. This also became obvious by repeating the exploration task with different delays (see Appendix A.1). Generally, the longer the delay the more non–contingent output is produced and so it becomes more difficult for the participant to discriminate which part of the movement performed brought the desired result (i.e. to place the vBot handle in the target area).

  Under the same delay, high speeds result in longer post–discovery movements. Thus the end–points of the movements are, most probably, further away from the target area (more non–contingent output). Suppose then that participants

on a particular trial try to reach the end–point of the previous trial. End–points further away from the target area, on a given trial, are expected to give rise to longer Irrelevant Distances in the subsequent trial.

In this analysis, using Post–Discovery Speed or Post–Discovery Distance would be the same. Remember that Post–Discovery Speed of each trial is calculated by taking the ratio of the total Post–Discovery Distance travelled over the reinforcement delay. We used the Post–Discovery speed because it is the independent factor in the post–discovery period (higher speeds are causing longer post–discovery distances).

Thus, it was of great interest to study if there was any correlation between the Post–Discovery Speed (or Distance) and the Irrelevant Distance. If such a correlation existed, it could be further predicted that participants that tracked with visual feedback delays, if adapted to them, would be less affected by the non–contingent output. This would be true either because these participants lowered their Overall Speed, and so their Post–Discovery Speeds, producing less non–contingent output, or because they could better infer the target area location after taking into account the delay.

To investigate the impact of the Post–Discovery Speed to the Irrelevant Distance we ran four 1 x 2 ANCOVAS, one for each of the four exploration Blocks (see Table 3.1), with the Irrelevant Distance in Trial n ($where\, n = 2, \ldots, 10$) as a dependent variable, the Post Discovery Speed in Trial $n-1$ ($where\, n = 2, \ldots, 10$) as a covariate and with one between–subject factor of Tracking Group (Group A–no visual feedback delays and Group B–visual feedback delays). Matlab aoctool was used for this statistical analysis.

- **Exploration–Exploitation trade–off.** According to learning theory (Sut-

ton and Barto, 1998), it is expected that greater exploration is related to better final performance. So, in the context of the Exploration Task, greater variability in the first trials is expected to lead to better average performance in the last trials (Stafford et al., 2012). We ran this analysis to study if the tracking experience influenced the Exploration–Exploitation trade off.

For each participant, we calculated the Standard Deviation of the Irrelevant Distance in Trials 2–5 and the Mean Irrelevant Distance in Trials 6–10. Each Exploration Block was analysed separately using a 1x2 ANCOVA with Mean Irrelevant Distance as a dependent variable, the Standard Deviation of the Irrelevant Distance as a covariate and one between subject factor of Tracking Group (Group A–no visual feedback delays and Group B–visual feedback delays). All ten Batches of each participant in each Block were submitted in the analysis. Before running the analysis, we visually inspected how each participant's data were distributed within the group distribution to ensure that the outcome of the analysis was not a product of local clustering of the participants across the regression line. Matlab aoctool was used for this analysis.

For all the indices above, unless otherwise described, we followed the same preprocessing and analysis steps. We first plotted the frequency distributions to check if the raw data were normal. Especially for the Irrelevant Distance, we expected that the distributions would be positively skewed (Walton, 2011). If the raw data were not normal they were log10 transformed. The values submitted to the statistical analysis were calculated as follows. As described earlier, participants were given ten trials in each target location. The first trial was excluded from the analysis as it mirrored naive performance. The mean of the index in Trials 2–10 of each Batch (target area) was taken. For each participant, we further calculated the average of

the variable across the ten Batches of each Block. After this preprocessing, the dependent variable was analysed with a 2x2x2 mixed ANOVA, with one within subject factor of Day (Day 1 and Day 2), one within subject factor of Block (pre and post tracking within each day) and one between subject factor of Tracking Group (Group A–no visual feedback delays and Group B–visual feedback delays).

3.3 Results

3.3.1 Tracking

General Observations

Typical tracking behaviour of individual trials in both groups is presented in the graphs of Figure 3.3.

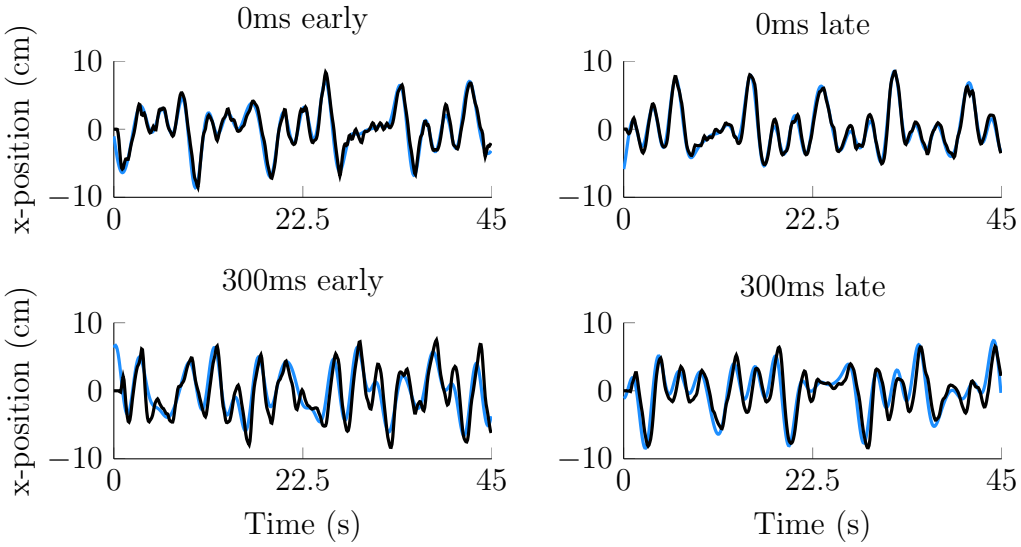

Figure 3.3: Error between the target (blue line) and the cursor (black line) in early and late trials, for individuals that tracked without (0ms) and with (300ms) visual feedback delays.

In early trials, as expected, the error when tracking with visual feedback delays is larger compared to an early tracking trial without any visual feedback delays. Over

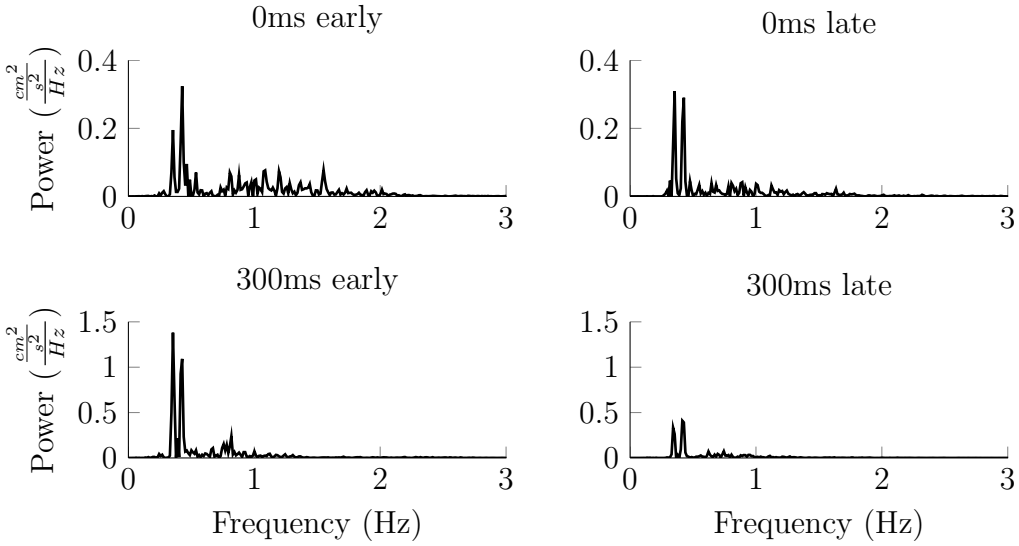

Figure 3.4: Speed Power Spectral Densities of individual early and late trials of participants that tracked without (0ms) and with (300ms) visual feedback delays. Note the change in vertical axis scale (top row versus bottom row).

practice, there seems to be no difference in the condition without visual feedback delays in terms of the error between the target and the cursor. On the other hand, in the condition with visual feedback delays there is a noticeable improvement which is mainly depicted in the turning points of the target.

Aside from the tracking error, the improvement in the tracking behaviour is mirrored in how smooth the movements become over practice. The power spectra in Figure 3.4 show that there is a clear improvement from early to late trials both without and with visual feedback delays. The spectra have less power in late trials as a result of the smoother movements achieved towards the end of training. In line with previous studies (Foulkes and Miall, 2000), during early trials there is a power band around the frequency of 1.2Hz for participants that tracked without visual feedback delays and around the frequency of 0.7Hz for participants that tracked with visual feedback delays. These components are less dominant in late trials.

Improvement in tracking performance

The aforementioned observations drawn from individual trials do generalise on a group level. There was no significant effect of practice on the RMS error between the target and the cursor in Group A (no visual feedback delays) as shown in Figure 3.5. This was further confirmed by the repeated measures ANOVA which showed no effect of Block ($F(5, 25) = 1.4$, $p = .264$, $\eta^2 = .091$) for Group A. On the other hand, there was a significant effect of Block in Group B that tracked with visual feedback delays ($F(5, 34) = 8.11$, $p = .001$, $\eta^2 = .091$), which is also seen in Figure 3.5.

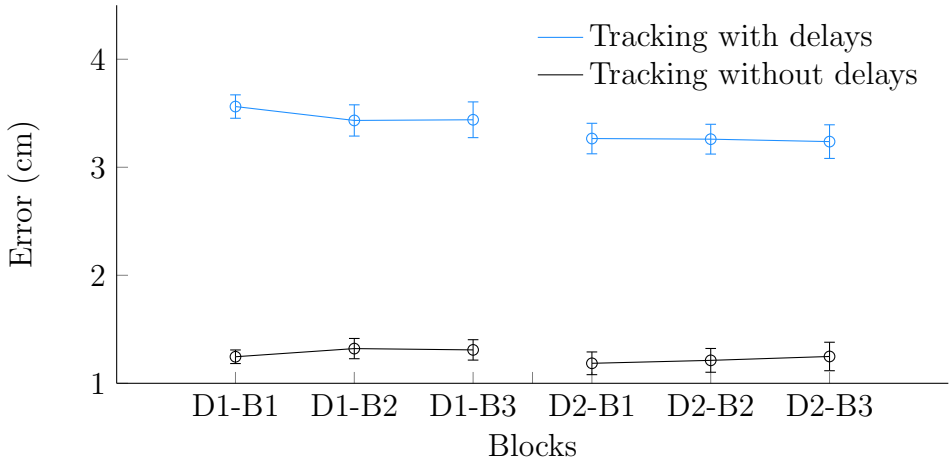

Figure 3.5: Mean RMS error between the target and the cursor in Group A that tracked without visual feedback delays (black line) and in Group B that tracked with visual feedback delays (blue line), in the 3 Blocks of Day 1 (D1 B1-3) and in Day 2 (D2 B1-3). Only Group B showed a significant improvement of tracking behaviour in terms of the error. The error bars show the standard error of the mean.

Participants in both groups executed smoother movements after practise. Figure 3.6 shows that participants in both groups achieved lower spectral densities, as expected (Foulkes and Miall, 2000). The repeated measures ANOVAs conducted separately for each group revealed that there was a significant effect of Block both in the group that tracked without visual feedback delays ($F(5, 40) = 18.628, p < 0.001, \eta^2 = .571$) and with visual feedback delays ($F(5, 36) = 12.567$, $p < 0.001$, $\eta^2 = .473$).

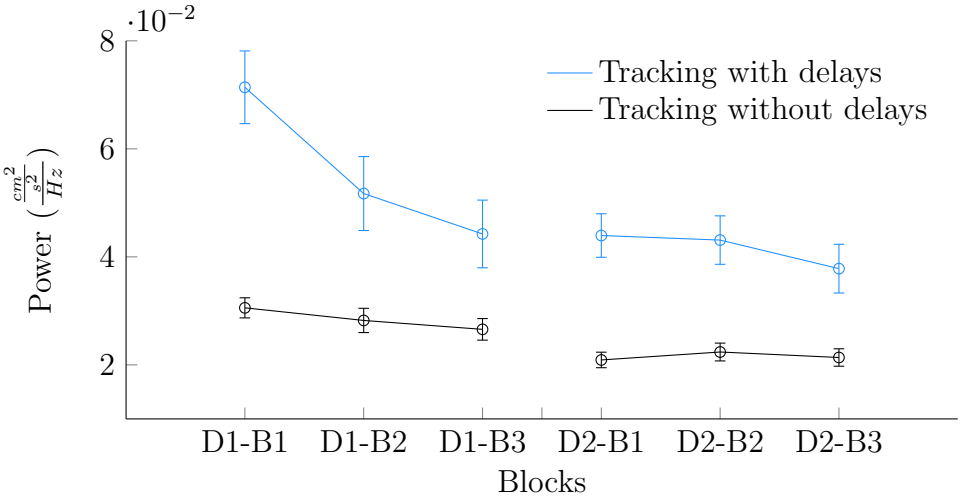

Figure 3.6: Average mean Power Spectral Density of Speed in Group A that tracked without visual feedback delays (black line) and in Group B that tracked with visual feedback delays (blue line), in the 3 Blocks of Day 1 (D1 -B1-3) and in Day 2 (D2 B1-3). Both Groups did improve over practice. The error bars show the standard error of the mean.

3.3.2 Exploration

General Observations

A rare ideal Exploration Behaviour of a participant in a given Batch is shown in Figure 3.7. However, not all the Exploration Batches looked the same. One frequently met behaviour is the superstitious behaviour (Skinner, 1948) and is shown in Figure 3.8. One of our expectations was that via the exposure to the delays in the tracking task (Group B) participants would not show this sort of behaviour. That would be depicted in the data as a drop in the Irrelevant Distance.

As explained earlier in the methods of the task, participants were not instructed to use a specific exploration strategy. This was a source of great variance among the participants. Some of the exploration methods used by the subjects are shown schematically in Figure 3.9. It is clear that some of the exploration techniques would bias early or late detection, depending on target location (e.g. Figure 3.9.e) whereas others could even possibly cancel out the effects of the delay (e.g. Figure 3.9.d) as

participants moved through the target in a straight line.

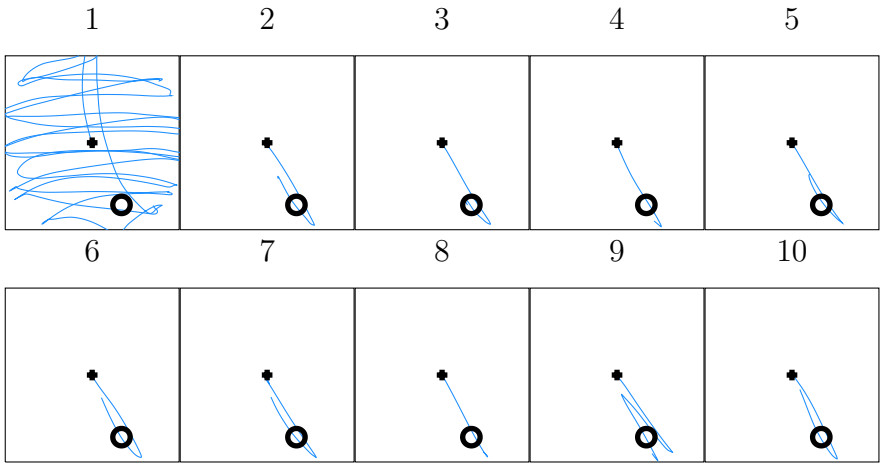

Figure 3.7: Ideal Exploration Task Behaviour. Participants started from the centre of the workspace and explored the workspace until they found the target area (black circle). For the next nine trials they had to relocate the same target area. That was defined as an Exploration Batch. Ideally, once they found the target area they would move to it in a straight line.

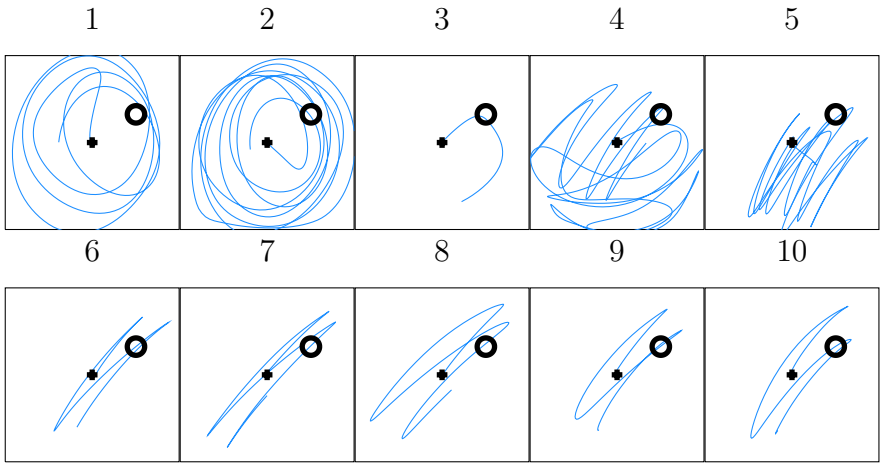

Figure 3.8: Superstitious Exploration Task Behaviour. Participants started from the centre of the workspace and explored until they found the target area (black circle). For the next nine trials they had to relocate the same target area. That was defined as an Exploration Batch. In this example, the participant found a movement that brought her to the target (Trial 6) and she repeated the same redundant movement for the rest of the trials without trying to find the optimal one.

For example, if the target was located in the middle of the workspace and participants were moving in a straight line through the target area, they would receive the reinforcement flash while still travelling on the same line. If in the next trial they aimed to go back to the final point of the previous movement, with the same straight line movement, they would inevitably pass through the target area. This created in some cases "false positives" of optimal Irrelevant Distance behaviour. That would also be the case if participants went through the target accidentally when starting to explore. However, this strategy could also make it easy to miss the target, if it fell between two "spokes". Nonetheless, we decided not to exclude any Batches of the analysis and to accept this variance as a characteristic of the Exploration task.

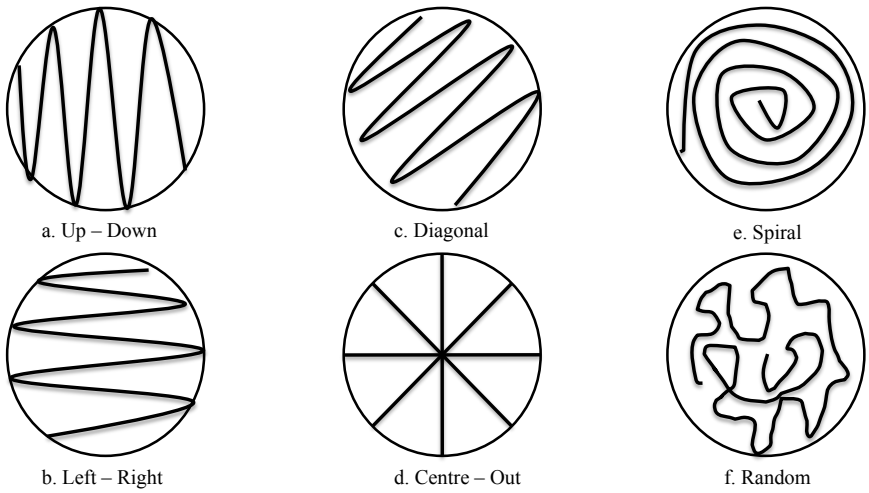

Figure 3.9: Exploration Task Techniques. Participants used different techniques to explore the workspace and find the target area. These are idealised versions of frequently observed patterns.

Mean Irrelevant Distance

Before calculating the mean Irrelevant Distance in Trials 2 –10, we explored the raw data set by plotting frequency histograms of the Irrelevant Distances. As expected,

the data were highly positively skewed. For this reason, the Irrelevant Distance of each Trial was log10–transformed. However, to our surprise, the resulting distribution was not a normal one but bimodal. Figure 3.10 shows the log10–transformed frequency histogram with all the Irrelevant Distances (Trials) of all 30 participants in Block 1 of the Exploration task. We ran Hartigan's dip test (Hartigan and Hartigan, 1985) which tests the null hypothesis that the data come from a unimodal distribution. The test came out significant ($p = 0.0226$) indicating that we should reject the null hypothesis of a unimodal distribution.

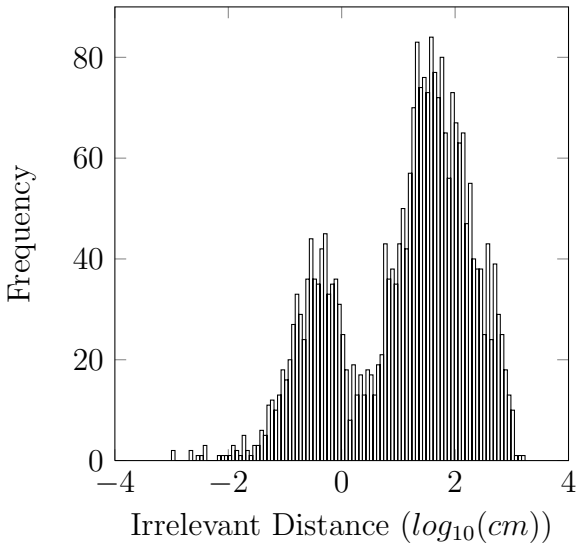

Figure 3.10: Irrelevant Distance Distribution ($log_{10}(cm)$). After log_{10} transforming the raw Irrelevant Distances the resulting distribution is a bimodal one. The histogram contains all the data from all 30 participants in Block 1 (10 Trials*10 Batches for each participant).

In the frequency histogram of all the raw (not transformed) data (12000 trials; including all participants (30), blocks (4), batches (10) and trials(10) – not presented here), there was a dominant first bin that concentrated 25% of the trials. These were trials that had an Irrelevant Distance below 1.1cm. These are the trials that give rise to the lefthand lobe distribution in Figure 3.10 (values below $0.04 = log_{10}(1.1)$).

We further ran Hartigan's dip test for each Block and each participant separately to check individual distributions and investigate if they are related to the exploration technique participants used. The test came out significant for 45% of the Block distributions. However, it was not the case that the data of any one participant across the Blocks would consistently be unimodal or bimodal. Moreover, we did not observe any consistent effect of the exploration technique to the distribution. That is partly because the same participant could switch between several techniques within the same block, or even within the same batch.

Although we tried to analyse the data taking into the account the bimodality of each participant and Block, we decided that the indices chosen were quite removed from the original raw data, and thus too difficult to be interpreted. This is why, keeping the diverse individual behaviours in mind, we chose to use the mean of the log10–transformed Irrelevant Distance in Trials 2-10 as a dependent variable.

Figure 3.11 shows the changes of the mean Irrelevant Distance in Trials 2–10 across the Exploration Blocks. Both Groups achieved much lower Irrelevant Distances in Day 2 compared to Day 1. However, contrary to the expectation that the group that tracked with delays would perform better, there was a tendency to the opposite. Group A (no tracking delays) seemed to have improved on Day 1 and remained the same on Day 2, whereas Group B (tracking delays) did not change their performance in Day 1 and got worse from Block 3 to Block 4, on Day 2.

The 2x2x2 ANOVA showed that only the factor of Day had a significant effect ($F(1, 28) = 12.014$, $p = .002$, $\eta^2 = .300$) on the Irrelevant Distance. There was not any other significant effect of the main factors to the Irrelevant Distance: Block ($F(1, 28) = .010$, $p = .923$, $\eta^2 < .001$) and Group ($F(1, 28) = .234$, $p = .632$, $\eta^2 = .008$) or of any of the interactions between the factors: Block*Group ($F(1, 28) = .821$, $p = .373$, $\eta^2 = .028$), Day*Group ($F(1, 28) = .012$, $p = .913$, $\eta^2 < .001$),

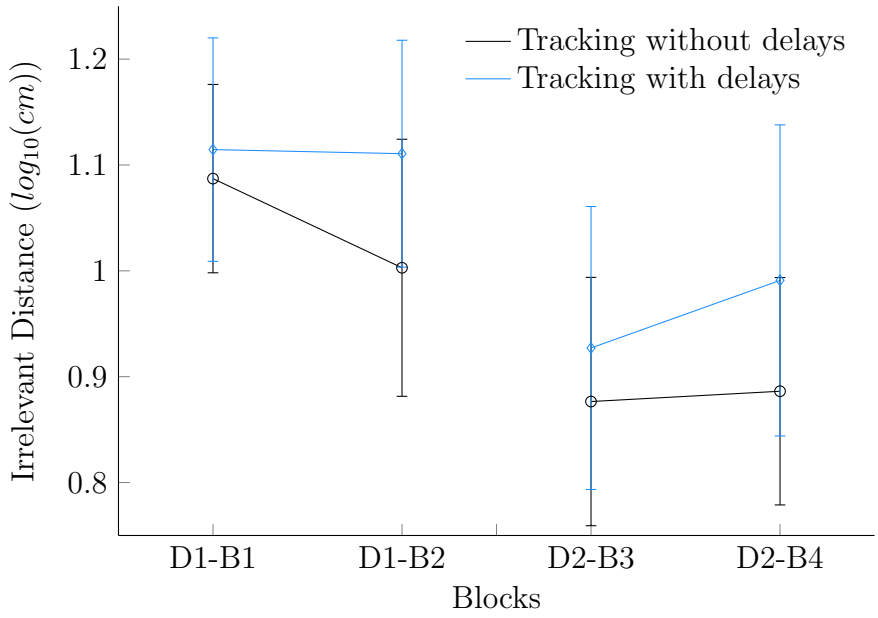

Figure 3.11: Mean Irrelevant Distance. Participants completed four Blocks of the Exploration task: in Day 1 pre (D1-B1) and post (D1-B2) tracking and in Day 2 pre (D1-B3) and post (D1-B4) tracking. The error bars show the average mean Irrelevant Distance (log10–transformed) in Trials 2-10 in the Group that tracked without visual feedback delays (black) and in the Group that tracked with visual feedback delays (blue). The error bars show the standard error of the mean.

Block*Day ($F(1, 28) = 1.085$, $p = .306$, $\eta^2 = .037$), Block*Day*Group ($F(1, 28) = .029$, $p = .867$, $\eta^2 = .001$).

We assumed that, given the variability induced by the different exploration techniques used by the participants and the number of participants we had in each group, we may not have had enough power to find any significant effect of tracking (Block) on the Irrelevant Distance. Instead we ran two paired sample t-tests, separately for each group, comparing Exploration Block 1 and Exploration Block 4, as these were the Blocks where the maximum tracking effect should be exposed. The results showed that Group A (no tracking delays) got significantly better in terms of Irrelevant distance ($t = 2.525$, $p = .024$) whereas Group B (tracking delays) did not ($t = 1.415$, $p = .179$). Although this last analysis is not statistically appropriate, it

suggests that our initial hypothesis that the group that tracked with delays would improve because of the tracking experience, is not confirmed. Instead, exploration seems not to improve, or even to become worse, with extended delayed tracking experience. On the other hand the group that tracked without visual feedback delays improved, also counter to our hypothesis.

Post–Discovery Distance

The Post–Discovery Distances distribution was positively skewed. After the transformation the data fell into a normal distribution. The mean Post–Discovery Distance across the four Exploration Blocks is shown in Figure 3.12. Participants tended to travel longer Post–Discovery Distances across the Exploration Blocks. Although Group A (no tracking delays) appeared to have a tendency towards longer distances than Group B, this was not confirmed by the statistical analysis.

The results of the 2x2x2 ANOVA revealed that there was a significant effect of Block ($F(1, 28) = 20.306$, $p < .001$, $\eta^2 = .420$) and of the interaction Block*Day ($F(1, 28) = 7.235$, $p = .012$, $\eta^2 = .205$) on the Post–Discovery Distance. This interaction was because of a significant increase in the Post–Discovery Distance in Day 1 ($p = .012$–adjusted for multiple comparisons Sidak). There was not any significant effect of the other main factors on the Post–Discovery Distance: Day ($F(1, 28) = 1.937$, $p = .175$, $\eta^2 = .065$) and Group ($F(1, 28) = .409$, $p = .528$, $\eta^2 = .065$) or of any of the interactions between the factors: Block*Group ($F(1, 28) = .920$, $p = .346$, $\eta^2 = .032$), Day*Group ($F(1, 28) = 1.189$, $p = .285\,\eta^2 = .041$) and Block*Day*Group ($F(1, 28) = 1.109$, $p = .301$, $\eta^2 = .038$).

Based just on these results, there is no sign of adaptation to delays in terms of the participants stopping on the target area and not travelling further. However, no firm conclusion can be derived without looking at the Overall Speed first.

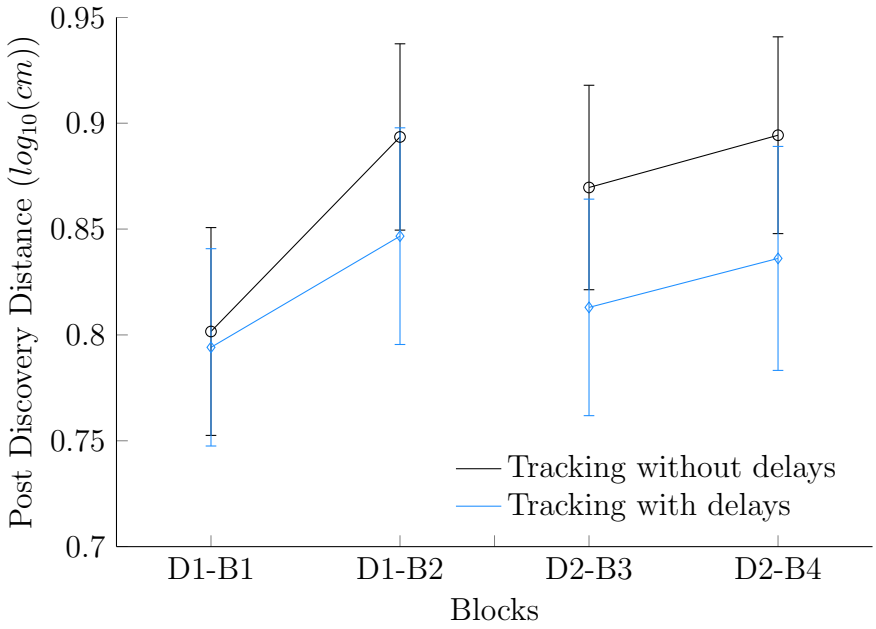

Figure 3.12: Post–Discovery Distance. Participants completed four Blocks of the Exploration task: in Day 1 pre (D1-B1) and post (D1-B2) tracking and in Day 2 pre (D1-B3) and post (D1-B4) tracking. The error bars show the average mean Post–Discovery Distance (log10–transformed) in Trials 2-10 in the Group that tracked without visual feedback delays (black) and in the Group that tracked with visual feedback delays (blue). The error bars show the standard error of the mean.

Overall Speed

The distribution of the raw Overall Speed data was also positively skewed. After the transformation the data fell into a normal distribution. Figure 3.13 shows the variation of Overall Speed across the Blocks.

There was an increase of movement speed across the Exploration Blocks. Participants in Group A (no tracking delays) tended to explore with higher speeds compared to the Group B (tracking delays). The group average speeds varied from $15.85cm/s$ in Block 1 to $20cm/s$ in Block 4 (the values in the graph are log10 values). These values are much higher compared to the average target speed in the tracking task ($8.5cm/s$). So, we can exclude the possibility that the changes in speed during

the exploration task were caused by adoption of the tracking target speeds, and instead may reflect growing familiarity with the task across the four blocks.

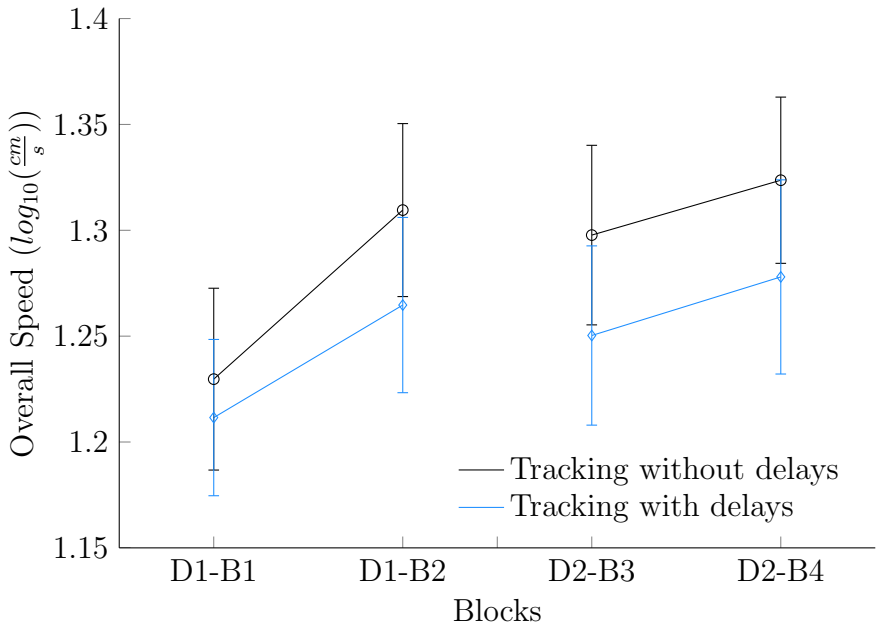

Figure 3.13: Overall Speed. Participants completed four Blocks of the Exploration task: in Day 1 pre (D1-B1) and post (D1-B2) tracking and in Day 2 pre (D1-B3) and post (D1-B4) tracking. The error bars show the average mean Overall Speed Distance (log10–transformed) in Trials 2-10 in the Group that tracked without visual feedback delays (black) and in the Group that tracked with visual feedback delays (blue). The error bars show the standard error of the mean.

The statistical analysis showed that there was a significant effect of Day ($F(1, 28) = 8.813$, $p = .006$, $\eta^2 = .239$), Block ($F(1, 28) = 20.358$, $p < .001$, $\eta^2 = .421$) and their interaction Block*Day ($F(1, 28) = 6.278$, $p = .018$, $\eta^2 = .183$) on the Overall Speed. Following this result we ran a simple main effects analysis. The interaction found was because of a significant increase in the Overall Speed in Day 1 ($p < .001$, Sidak–adjusted) but not in Day 2 ($p = .060$, Sidak–adjusted). Moreover, the increase of speed between Block 1 and Block 3 was also significant ($p = .001$, Sidak–adjusted).

There was not any significant effect of the other main factors to the Post–Discovery Distance: Group ($F(1, 28) = .482$, $p = .493$, $\eta^2 = .017$) or of any of

the following interactions between the factors: Block*Group ($F(1, 28) = .364$, $p = .551$, $\eta^2 = .013$), Day*Group ($F(1, 28) = .443$, $p = .511$, $\eta^2 = .016$) and Block*Day*Group ($F(1, 28) = .806$, $p = .377$, $\eta^2 = .028$).

Post–Discovery Speed

As explained in Methods paragraph 3.2.6, in order to study the impact of the Post–Discovery Speed on the Irrelevant Distance we used four 1x2 ANCOVAs, one for each of the four exploration Blocks, with the Irrelevant Distance in Trial n ($where\, n = 2, \ldots, 10$) as a dependent variable, the Post Discovery Speed in Trial $n-1 (where n = 2, \ldots, 10)$ as a covariate and with one between subject factor of Tracking Group (Group A – no visual feedback delays and Group B – visual feedback delays). The reason for using all the trials in the analysis (and not taking averages of Batches and Blocks) was to avoid obscuring any trial by trial effect. In the following paragraphs, the ANCOVA results are presented Block by Block.

Figure 3.14 shows the distribution of the two groups in **Block 1** and the fitted regression lines. The two groups were identical and a correlation between the Irrelevant Distance and the Post–Discovery Speed seemed to exist. This was confirmed by the ANCOVA which showed that the effect of the covariate was significant ($F(1, 2696) = 298.93$, $p < .001$) and that there was no Group effect ($F(1, 2696) = .98$, $p = .3229$) or any interaction of Group and Post–Discovery Speed ($F(1, 2696) = .01$, $p = .9157$).

The Matlab aoctool also returned the parameter estimates of the regression lines. The model fitted to the data, chosen as input to the aoctool command, was that of "Separate Lines". This model is defined as $y = (\alpha + \alpha_i)x + (\beta + \beta_i)$, where i is the index for each group. Table 3.2 presents the coefficients of these models, as given by the Matlab aoctool output. The results are interpreted as follows: the lines

relating the Irrelevant Distance to Post–Discovery Speed have an *Intercept* close to $\beta = -0.3538$ and a *Slope* close to $\alpha = 1.478$. Each group's coefficients are offset from these values somewhat by the rise of α_i or β_i. For example, the slope for Group A line was equal to $\alpha + \alpha_{GroupA} = 1.478 - 0.0091 = 1.4689$. The p–values show the probability of the true group specific coefficients to be zero. In this case, the probability of the α_{GroupA}, α_{GroupB}, β_{GroupA}, β_{GroupB} to be zero is very high. Thus, the slopes for the two groups were not significantly different from each other, which was to be expected by the lack of interaction in the ANCOVA output.

We also separately calculated the R^2 for the goodness of fit of the regression line of each group (given in Table 3.2 next to the slope coefficient of each group). Based on the very low R^2 values, we can conclude that the Post–Discovery Speed can only explain a very small percentage of the variance in the Irrelevant Distances.

The distributions of the groups in **Block 2** are shown in Figure 3.15. The 1x2 ANCOVA revealed that the effect of the Post –Discovery Speed was significant ($F(1, 2696) = 295.88$, $p < .001$). Moreover, there was a significant effect of Group ($F(1, 2696) = 17.61$, $p < .001$) and an interaction of Group*Post–Discovery Speed ($F(1, 2696) = 5.96, p = .0147$). Because of this interaction the coefficients Slopes and Intercepts of the regression lines for Group A and B were significantly different from zero. That meant that a different slope described better each group. By observing the slopes we could say that the group that tracked with delays had higher values of Irrelevant Distance for the same Post –Discovery Speed. The R^2 values were very low, like in the previous Block (see Table 3.3).

The results for the ANCOVA of **Block 3** showed that there was a significant effect of Group ($F(1, 2696) = 10.77$, $p = .001$) and a significant effect of the Post–Discovery Speed as a covariate ($F(1, 2696) = 358.55$, $p < .001$). No significant interaction of Group*Post–Discovery Speed was found $F(1, 2696) = .64$, $p = .423$).

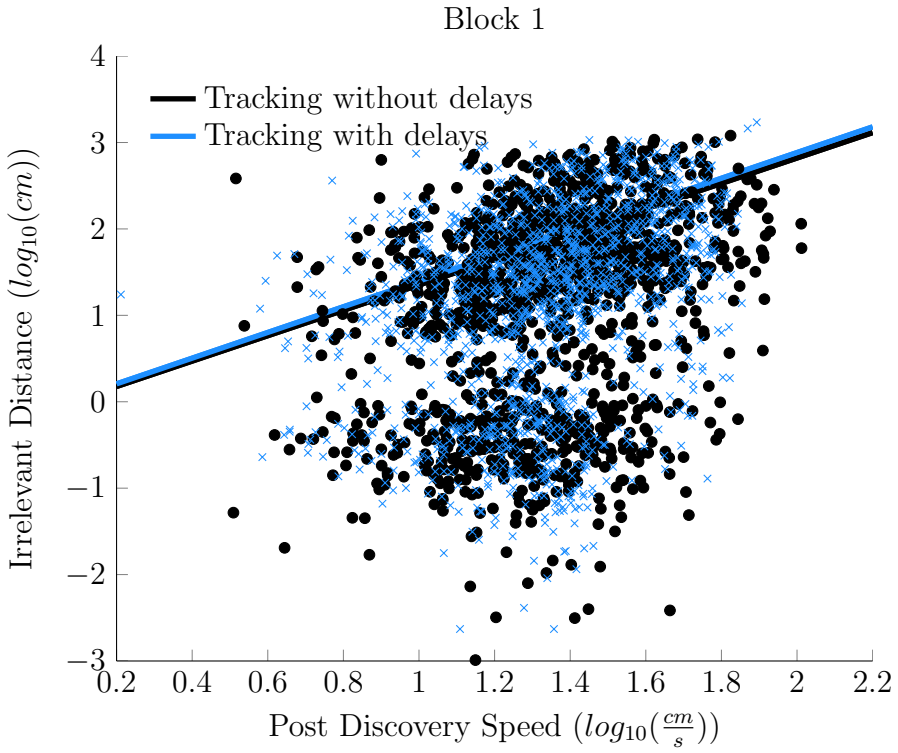

Figure 3.14: The Irrelevant Distance in Trial n is plotted against the Post–Discovery Distance in Trial n-1 for the Group A (black dots) and Group B (blue crosses) in Block 1. Superimposed are the regression lines for each group.

Figure 3.16 shows the distributions of the groups and the regression lines in Block 3. The coefficient estimates for each group (see Table 3.4) were not significant.The R^2 values were found to be very low again.

Finally, Figure 3.17 presents the data distribution and regression lines for **Block**

Table 3.2: Regression Lines Coefficients for **Block 1**.

	Coefficients	p–values	R^2
Intercept (β)	-0.8676	$< .001$	
β_{GroupA}	-0.0086	0.9412	
β_{GroupB}	0.0086	0.9412	
Slope (α)	1.478	$< .001$	
α_{GroupA}	-0.0091	0.9157	0.09498
α_{GroupB}	0.0091	0.9157	0.1049

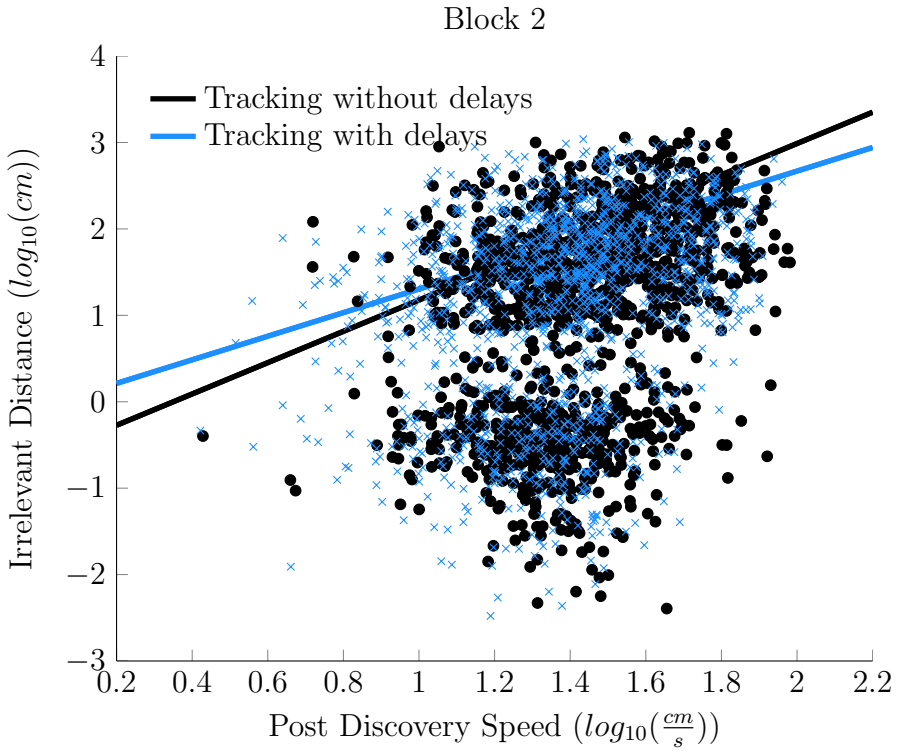

Figure 3.15: The Irrelevant Distance in Trial n is plotted against the Post–Discovery Distance in Trial n-1 for the Group A (black dots) and Group B (blue crosses) in Block 2. Superimposed are the regression lines for each group.

4. There was a significant effect of Group ($F(1, 2696) = 22.99$, $p < .001$) and a significant effect of the Post–Discovery Speed as a covariate ($F(1, 2696) = 374.65p < .001$). The interaction of Group*Post–Discovery Speed was marginally not significant ($F(1, 2696) = 3.62, p = .0573$). As a result the parameter estimates for the regression

Table 3.3: Regression Lines Coefficients for **Block 2**.

	Coefficients	p–values	R^2
Intercept (β)	-1.178	$< .001$	
β_{GroupA}	-0.4037	0.002	
β_{GroupB}	0.4037	0.002	
Slope(α)	1.5882	$< .001$	
α_{GroupA}	0.2232	0.0147	0.1123
α_{GroupB}	-0.2232	0.0147	0.08829

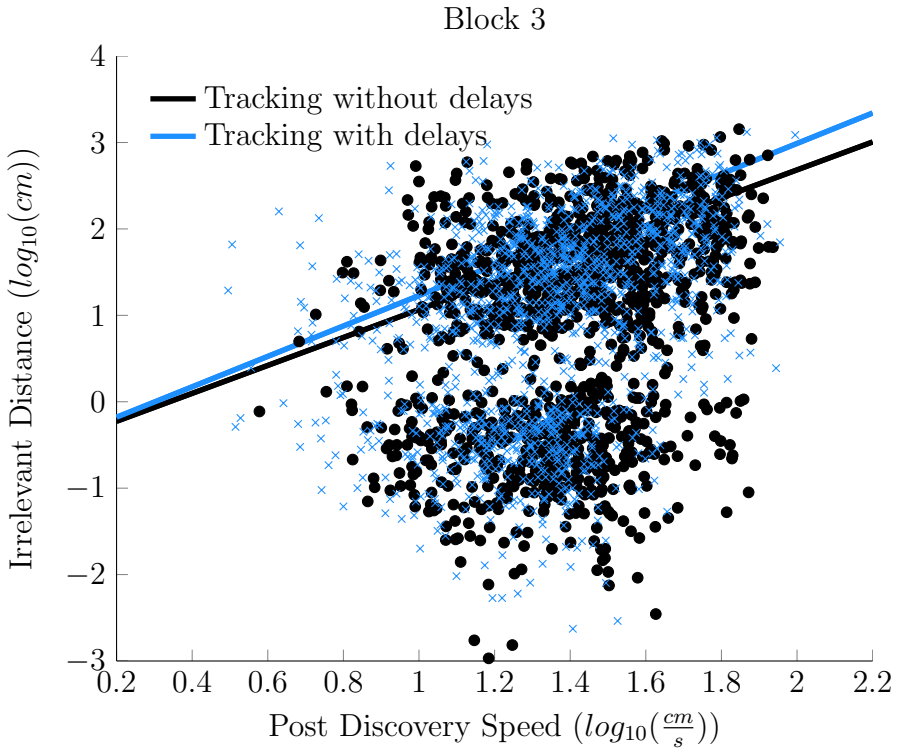

Figure 3.16: The Irrelevant Distance in Trial n is plotted against the Post–Discovery Distance in Trial n-1 for the Group A (black dots) and Group B (blue crosses) in Block 3. Superimposed are the regression lines for each group.

lines for each group could be zero (see Table 3.5). Consistent with the previous findings, the R^2 values in this case confirmed again that the model using the Post–Discovery Speed as a covariate could only explain a small part of the variance in the data.

Table 3.4: Regression Lines Coefficients for **Block 3**.

	Coefficients	p–values	R^2
Intercept (β)	-1.4231	$< .001$	
β_{GroupA}	0.0284	0.8213	
β_{GroupB}	-0.0284	0.8213	
Slope(α)	1.6883	$< .001$	
α_{GroupA}	-0.0719	0.423	0.09502
α_{GroupB}	0.0719	0.423	0.142

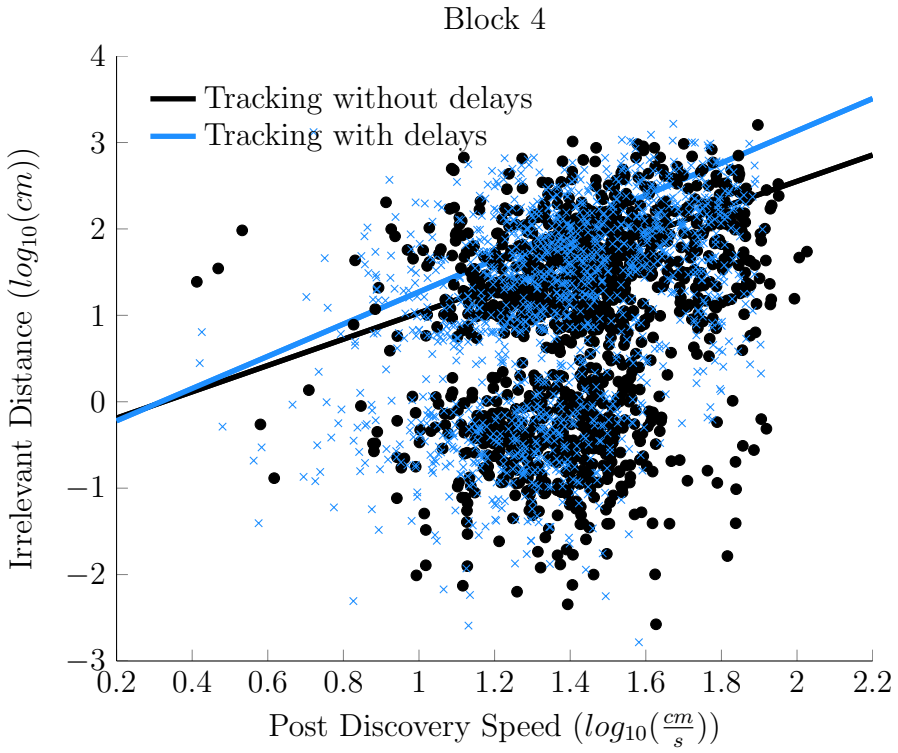

Figure 3.17: The Irrelevant Distance in Trial n is plotted against the Post–Discovery Distance in Trial n-1 for the Group A (black dots) and Group B (blue crosses) in Block 4. Superimposed are the regression lines for each group.

At this point it is worth making some general observations about the data presented in the Post–Discovery Speed section. In Figures 3.14–3.17 we can clearly see the bimodality in our dataset. It seems from the plots that a significant percentage of the higher Irrelevant Distances lobe occurs because of high Post–Discovery Speeds

Table 3.5: Regression Lines Coefficients for **Block 4**.

	Coefficients	p–values	R^2
Intercept (β)	-1.4288	$< .001$	
β_{GroupA}	0.1385	0.2788	
β_{GroupB}	-0.1385	0.2788	
Slope(α)	1.6933	$< .001$	
α_{GroupA}	-0.1706	0.0573	0.08282
α_{GroupB}	0.1706	0.0573	0.1623

– see how Irrelevant Distances are clustered with a mean Post–Discovery Speed of about 1.4 to 1.6 whereas the lower cluster has a mean of about 1.3.

Given the significant effect of Group in Blocks 2–4, we ran the following analysis to explore the effect a bit further. The ratio of the Irrelevant Distance in Trial *n* over the Post–Discovery Distance in Trial *n-1* was taken as a dependent variable. The data were log10 transformed as they were positively skewed and the resulting distribution appeared to be normal. The mean of each Batch was then calculated, as well as the Block average for each participant. Figure 3.18 shows how the ratio of these two values varied for each group across the four Exploration Task Blocks. The pattern observed in Figure 3.18 is very similar to that of Figure 3.11, which shows the Irrelevant Distance without taking into account the Post –Discovery Speed. However, the Group variance is greatly reduced when using the ratios as a variable. A 2x2x2 ANOVA, with the same factors as described in previous paragraphs, was run. As previously (p.100) we found that there was a significant effect of Day ($F(1, 28) = 18.47$, $p < .001$, $\eta^2 = .397$) but no significant effect of Block ($F(1, 28) = 2.151$, $p = .154$, $\eta^2 = .071$) or Group ($F(1, 28) = 1.005$, $p = .325$, $\eta^2 = .035$) or of any of the interactions: Block*Day($F(1, 28) = 2.873$, $p = .101$, $\eta^2 = .093$), Block*Group ($F(1, 28) = 1.776$, $p = .193$, $\eta^2 = .060$), Day*Group ($F(1, 28) = .209\, p = .651$, $\eta^2 = .007$) and Block*Day*Group ($F(1, 28) = .165$, $p = .687$, $\eta^2 = .006$). A simple main effects analysis showed that both groups got significantly better from day to day when comparing Block 1 to Block3 (Group A, $p = .002$ -Sidak adjusted and Group B, $p = 0.015$ - Sidak adjusted). Moreover, when exploring the within each day effects we found that only Group A improved significantly during Day 1 ($p = .047$ - Sidak adjusted).

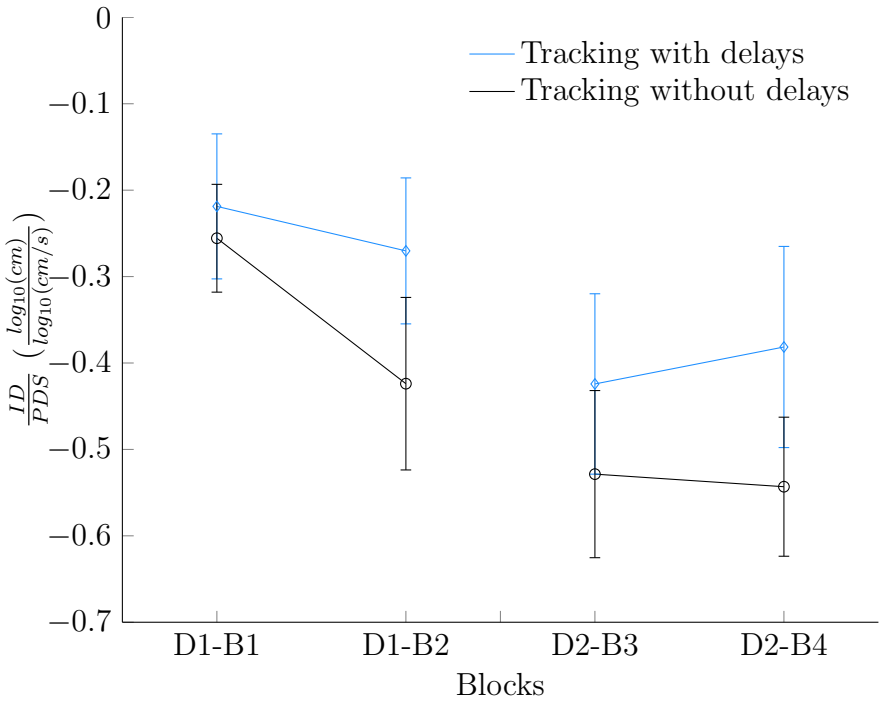

Figure 3.18: The ratio of Irrelevant Distance over the Post–Discovery Speed in Day 1 pre (D1-B1) and post (D1-B2) tracking and in Day 2 pre (D1-B3) and post (D1-B4) tracking for the Group that tracked without visual feedback delays (black) and in the Group that tracked with visual feedback delays (blue). The error bars show the standard error of the mean.

Exploration–Exploitation trade–off

In order to study the Exploration–Exploitation trade–off, we studied each Exploration Block separately. For each participant, we used the data from all the Batches in each Block. We decided not to demean each participant's data as that could obscure important information for our research question. We visually inspected the data to make sure that the data points of each participant were distributed across the entire dataset space. Figures 3.19 and 3.20 show the individual confidence ellipses (black ellipses) at Block 1 for Group A and Group B accordingly. Each ellipse encloses the 10 data points of a participant (one for each Batch). We observed that the individual ellipses overlap and that the group confidence ellipse (in blue) almost

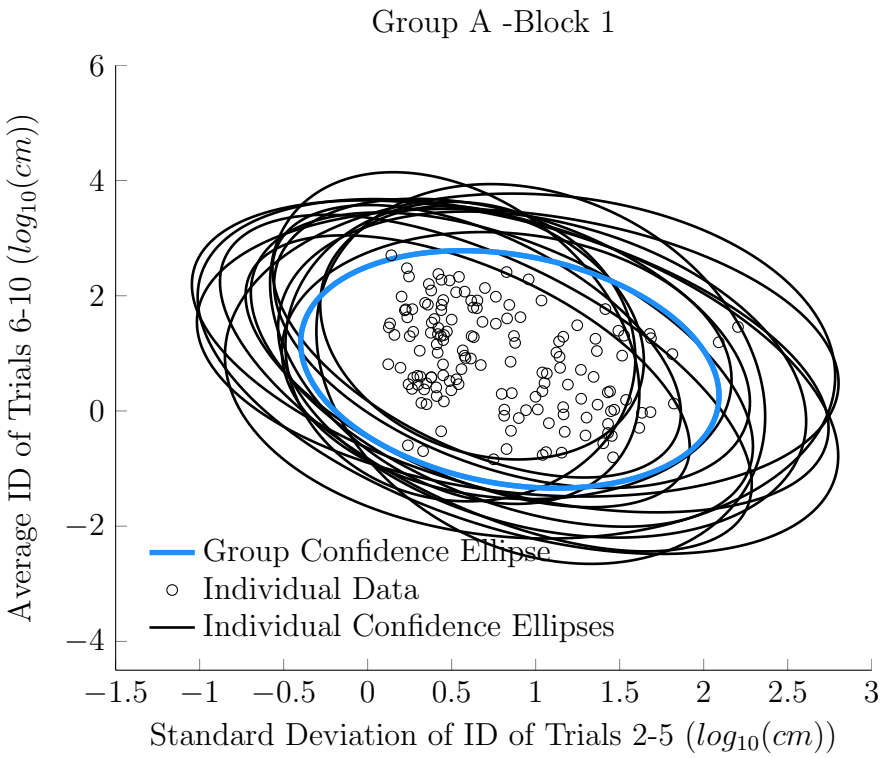

Figure 3.19: The average Irrelevant Distance in Trials 6 –10 is plotted against the Standard Deviation of the Irrelevant Distance in Trials 2–5. The black dots are the data points from all batches and participants in Group A (no tracking delays) in Block 1. The black ellipses are the individual confidence ellipses (each ellipse encircles 10 data points). The blue ellipse is the group confidence ellipse.

encircles the mutual space of the individual ellipses. We repeated the same procedure for all the Blocks and Groups and the results were similar. Therefore, the overall relationship between exploration and exploitation is not driven by between–participant differences.

Figure 3.21 shows the data points in each Block, the group confidence ellipses and the regression lines for each Group across the four Blocks–Group A (no tracking delays) is represented with black and Group B (tracking delays) with blue. The dashed line in the subplots of Block 2, 3 and 4 is the regression line of Block 1. We observed that the two groups are identical in Block 1 and almost the same in Block

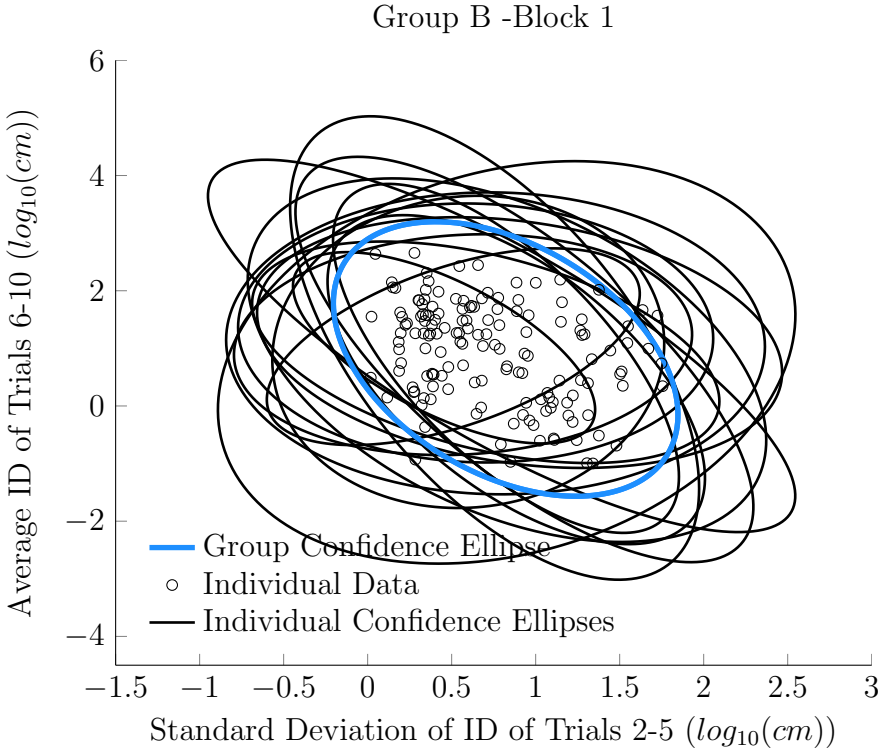

Figure 3.20: The average Irrelevant Distance in Trials 6 –10 is plotted against the Standard Deviation of the Irrelevant Distance in Trials 2–5. The black dots are the data points from all batches and participants in Group B (tracking with delays) in Block 1. The black ellipses are the individual confidence ellipses (each ellipse encircles 10 data points). The blue ellipse is the group confidence ellipse.

2 and 3. However, in Block 4 the two groups are differentiated. Group B appeared to have obtained a steeper Exploration–Exploitation trade–off regression line.

The results of the ANCOVA showed that in Block 1, there was a significant effect of the covariate (standard deviation in trials 2–5) on the average Irrelevant distance of trials 6-10 ($F(1, 296) = 36.59$, $p < .001$) but there was no Group effect ($F(1, 296) = 0.03$, $p = .8566$) and no interaction between the Group and the covariate ($F(1, 296) = 0.001$, $p = .9973$). The lines relating the average Irrelevant Distance in trials 6–10 to the variance of the Irrelevant Distance in trials 2 –5 had an *Intercept* close to $\beta = 1.387$ and a *Slope* close to $\alpha = -0.6$. Each group's coefficients

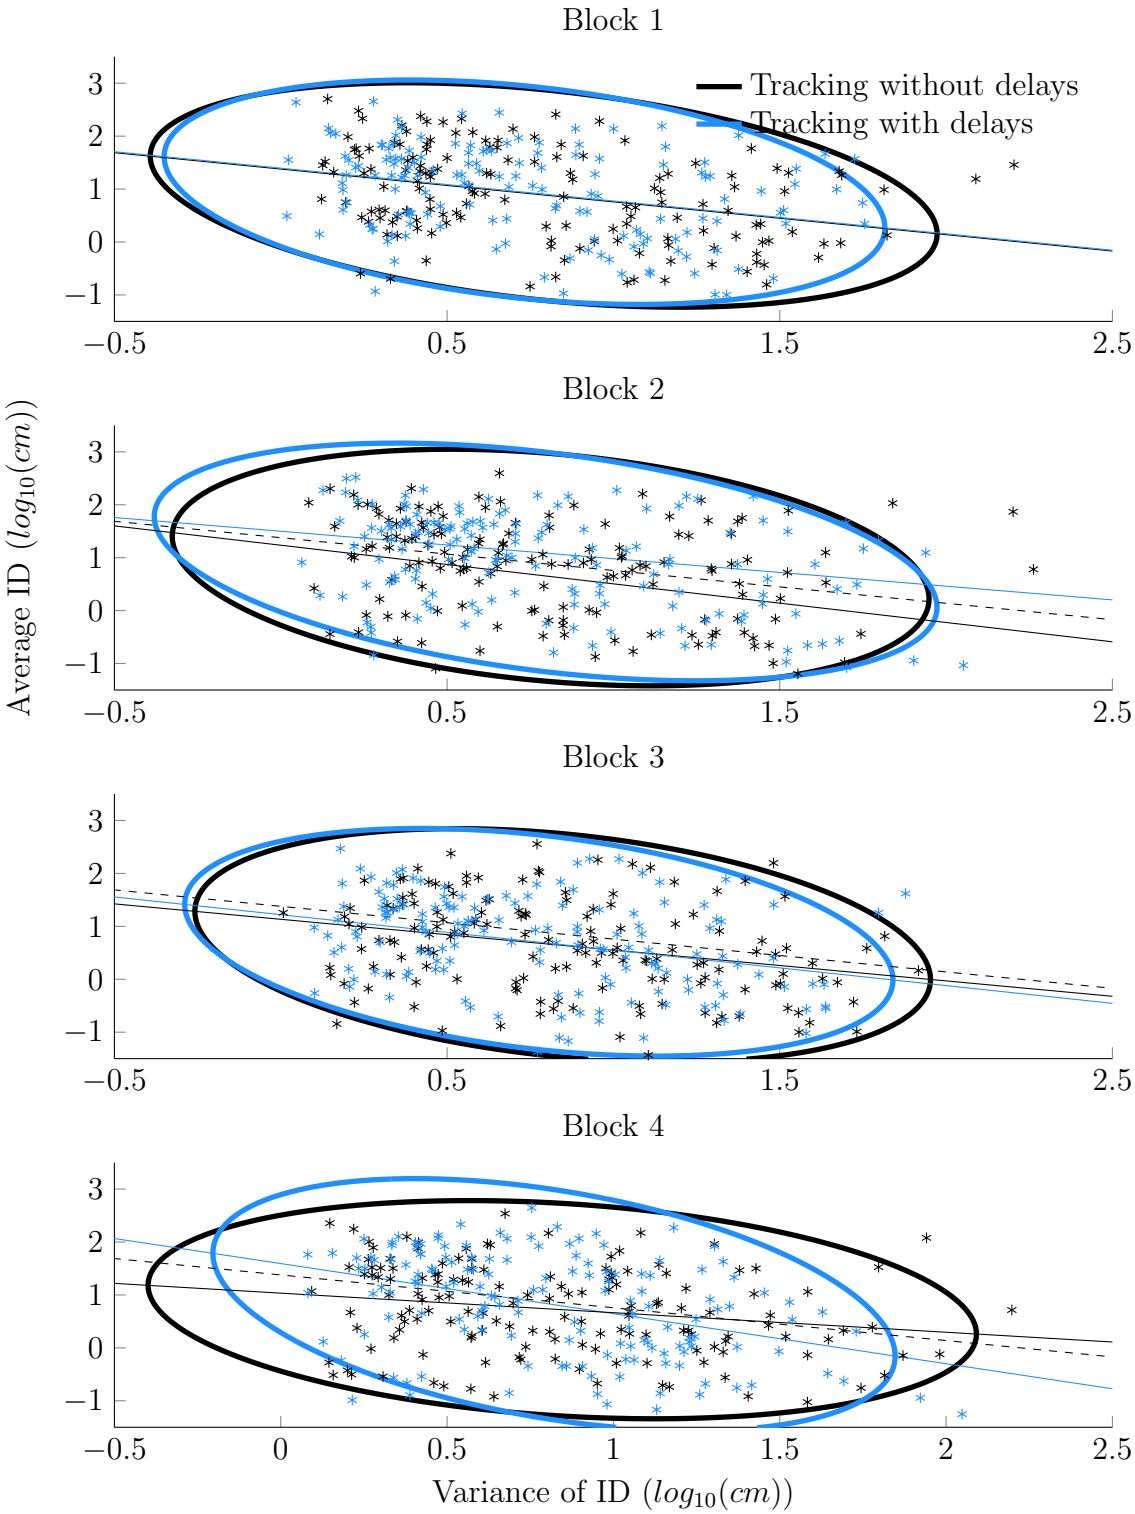

Figure 3.21: Average Irrelevant Distance (ID) of Trials 6 –10 vs Standard Deviation of the ID in Trials 2–5 across the Blocks. The data, group confidence ellipses and the regression lines are presented with black for Group A (no tracking Delays) and blue for Group B (tracking delays).

were offset from these values somewhat but the probability that they were zero was very high, meaning that they were not different form each other (see Table 3.6). Moreover, the R^2 values for each regression line were very low (Group A: 0.1195, Group B:0.1005). Thus, only a small percentage in the data variability could be explained by the variance in the first trials in each Batch. The results of the analysis for Block 2 and 3 were very similar to those of Block 1 and we do not report them for being redundant.

Table 3.6: Exploration–Exploitation trade–off analysis: Regression Lines Coefficients for **Block 1**.

	Coefficients	p–values	R^2
Intercept (β)	1.3869	$< .001$	
β_{GroupA}	-0.0087	0.9237	
β_{GroupB}	0.0087	0.9237	
Slope(α)	-0.6201	$< .001$	
α_{GroupA}	0.0003	0.9973	0.1195
α_{GroupB}	-0.0003	0.9973	0.1005

Table 3.7: Exploration–Exploitation trade–off analysis: Regression Lines Coefficients for **Block 4**.

	Coefficients	p–values	R^2
Intercept (β)	1.3133	$< .001$	
β_{GroupA}	-0.2794	0.007	
β_{GroupB}	0.2794	0.007	
Slope(α)	-0.6579	$< .001$	
α_{GroupA}	0.2885	0.0084	0.04977
α_{GroupB}	-0.2885	0.0084	0.1659

In Block 4, we found that there was a significant effect of the covariate (standard deviation in trials 2–5) on the average Irrelevant distance of trials 6-10 ($F(1, 296) = 31.88$, $p < .001$) but there was no Group effect ($F(1, 296) = 0.65$, $p = .4199$). However, there was a significant interaction between the Group and the covariate

($F(1, 296) = 7.04$, $p = .0084$). The coefficients of the regression lines for Block 4 are found in Table 3.7. Group B (tracking delays) is described by a steeper slope ($\alpha + \alpha_{GroupB} = -0.9464$) compared to Block A ($\alpha + \alpha_{GroupA} = -0.3694$). The slope and intercept estimates for both groups were significant ($p = .0084$ and $p = .007$, accordingly). The R^2 value for Group A was extremely small (0.04977) but higher for Group B (0.1659).

The difference in the exploration–exploitation regression lines must have been driven by a steeper increase in the variance of the first trials in Group B that tracked with delays (see Figure 3.22) rather than in differences in the average performance towards the end of a Batch of trials (see Figure 3.23).

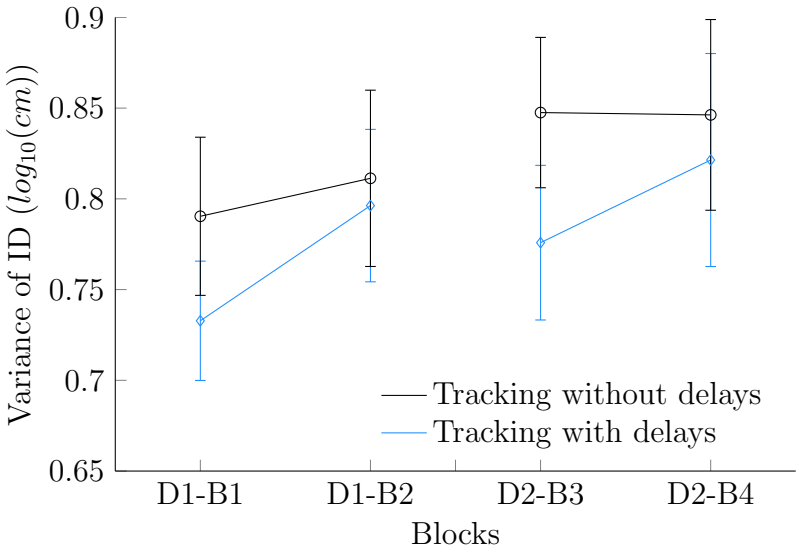

Figure 3.22: Standard Deviation of the Irrelevant Distance (ID) in Trials 2–5 across the Blocks for Group A (no tracking Delays) and Group B (tracking delays). The error bars show the standard error of the mean.

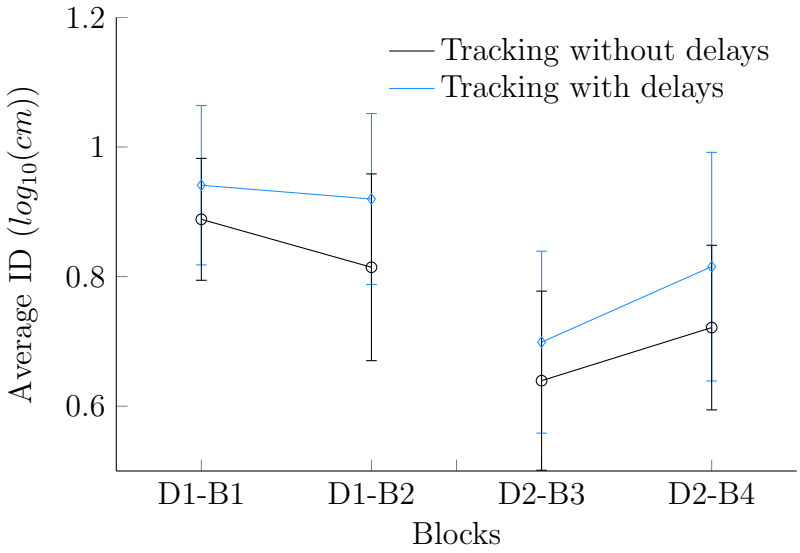

Figure 3.23: Average Irrelevant Distance (ID) in Trials 6 –10 across the Blocks for Group A (no tracking Delays) and Group B (tracking delays). The error bars show the standard error of the mean.

3.4 Discussion

Our hypothesis in the present study was that as participants adapt to tracking delays and perhaps attribute them as a property of the vBots, they would be able to better associate their actions with the brought outcome (receiving the reinforcement flash) during the exploration task. The predictions we made based on the dependent variables used were the following. If participants were able to "model" the delay during the tracking task and better understand where the target area was during exploration, we then expected that the group that tracked with delays:

- would be able to find a more efficient path to reach the target – depicted in shorter Irrelevant Distances and/or
- participants would stop on the target and wait for the delivery of the reinforcement – depicted in shorter Post–Discovery Distances and/or

- they would slow down on purpose to compensate for the delays – depicted in lower speeds and as a result shorter Post–Discovery Distances.

However, none of these hypotheses were confirmed by the results.

Irrelevant Distance. Contrary to our expectation, tracking experience without delays had a positive impact on the Irrelevant Distances during exploration. On the other hand, tracking with delays seemed to impede improvement in the exploration task. It could be argued that shorter Irrelevant Distances in Group A (no tracking delays) are related not to tracking experience but to familiarisation with the exploration task. This could be an explanation but it needs to be further tested. However, the different direction in the effect between Group A and B implies that at least part of this behaviour is related to practising the tracking task. Finally, the group–independent significant reduction in Irrelevant Distance from Day 1 to Day 2 is most probably related to exploration experience; participants developed and used more efficient strategies to explore the workspace over practice.

Overall Speed and Post–Discovery Distance. Tracking experience did not have an impact on the Overall Speed during the exploration task. Firstly, participants did not adopt the average tracking speeds during the exploration task; exploration speeds were much higher. Moreover, the adaptation to the delays during tracking did not lead to lower speeds. One could have expected that if participants had attributed the delay as a characteristic of the manipulandum would slow down to perform better in the exploration task. The fact that both groups showed gradual increase in speed indicates that these speed changes should be attributed to the extended exploration practice and maybe to familiarisation of the workspace during both the tracking and the exploration task.

It is apparent from the analysis of the Overall Speed that the changes seen in the Post–Discovery Distance must be a byproduct of the speed used by the participants. So, there is no sign that participants adapted to the delay in the hypothesised way, by stopping on the target and waiting until reinforcement delivery.

The effect of Post–Discovery Speed to the Irrelevant Distance. Aside from the hypotheses stated above, we were also interested in looking at how Post–Discovery Speed on a trial interacts with the Irrelevant Distance in the next trial. The results of this analysis indicated that there is a correlation between Post–Discovery Speed and Irrelevant Distance.

However, the extremely low R^2 values in all fits indicate that the model used to describe the data was poor. This was to be expected because we ignored the effect of the individual within–participant variance because of the repeated measures within each Batch and within each Block. A mixed effects model could have been used instead to take into account the random effects that arise from the repeated measures. However, the degrees of freedom of such a model would make it extremely difficult to systematically interpret the results. Manipulation of the participants' speed and exploration technique experimentally could perhaps make more sense.

Overall, we could say that the Post–Discovery Speed can explain some percentage of the Irrelevant Distance variability. This was also true when averaging across the trials of each Batch (this analysis was not presented) and actually the R^2 values were a bit higher (0.17–0.32). However, with exception of the Block 2 results, there was not any significant interaction of Group and Post–Discovery Distance.

Finally, an interesting finding of this analysis was the significant effect of Group in Blocks 2–4. Did this mean that if we took into account the variance induced by the Post–Discovery Speed, we would observe a Group effect on the Irrelevant Distance?

We simplified the analysis by calculating the ratio of the Irrelevant Distance in Trial n over the Post–Discovery Distance in Trial n-1. The results were similar to those of the Irrelevant Distance, suggesting that tracking experience without delays enhanced exploration performance whereas tracking with delays did not.

Tracking Performance The absence of any evidence in favour of our initial hypotheses is not related to lack of adaptation during the tracking task. Similarly to what has been previously observed in this task (Miall and Jackson, 2006), participants showed an improvement in performance in terms of executing smoother movements in both groups. Moreover, the group that tracked with visual feedback delays significantly reduced tracking errors over practice. Finally, although not explicitly tested, improvements in tracking were observed during both days.

Since participants adapted to the delays in the expected way in the tracking task, why did we observe exploration behaviour opposite to the predicted effects? Our hypotheses implied that participants in Group B (tracking with visual feedback delays)were aware of the delays. However, this was not the case for the majority of the participants as they reported in the debriefing session after the study. In previous studies (Miall and Jackson, 2006; Foulkes and Miall, 2000), participants practised baseline trials (without delays) before being exposed to tracking with delays. So, the unexpected behaviour they encountered during the trials with delays could only be attributed to a experimental manipulation. However, in our study participants were not told about the delays and they did not practise baseline trials because we did not want to give them any clue of the delays in any stage of the study. In a sense, they had to infer that the errors in the performance arise from delays attributed to the vBot.

We propose that, at least in the context of our set–up, where participants do not

have direct visual feedback of their arm, tracking with delays might have increased proprioceptive uncertainty. This could be because participants were not able to accurately attribute (at least during the tracking experience gained during the two days practise) the errors in the performance to the themselves or to the vBots. An increase in proprioceptive uncertainty could have resulted in higher Irrelevant Distances (or at least no improvement in the exploration task behaviour) and could also explain why non–contingent output had a more severe effect in the group that tracked with delays. We will return to this hypothesis in the following chapter.

Exploration–Exploitation trade–off Increase in proprioceptive uncertainty could also lead to more exploration and so the variance of the Irrelevant Distances would be higher. Any differences in variance would be depicted in the Exploration–Exploitation trade–off analysis. The results of this analysis showed that the exploration–exploitation slopes of the two groups were significantly different in the last exploration block. This was explained partly by an increased variance in the first trials for the group that tracked with delays. However, we should mention that the exploration–exploitation analysis was also characterised by low R^2 values in all fits. Again this means that this model was also poor to explain the behaviour.

Distribution of the Irrelevant Distance The distribution of the data was expected to be positively skewed, as found before (Walton, 2011). Heavy tailed (positively skewed) distributions of human travels in virtual environments have been suggested to describe exploration–exploitation behaviour (Volchenkov et al., 2013). So a heavy tailed distribution is composed by a gaussian distribution representing exploitative behaviour whereas the power law tail distribution emerges from explorative trials.

Our data of the Irrelevant distance were positively skewed but not as expected.

There was a component in the distribution, representing trials with Irrelevant Distances below 1.1cm, which was giving a high initial bin in the histogram of raw data that concentrated 25% of the total trials (excluding these data would lead to a distribution as expected). We tested if there was any relation between exploration techniques that the participants used and the individual distributions, but we failed to observe anything systematic.

Lack of a consistent relationship between exploration techniques and these trials does not exclude the hypothesis that this bias in the data emerges from specific techniques. This is because each participant did not always use the same technique (even within the same trial). Observation of the Batches indicated that the majority of these trials arise because of the technique but also because of accidental encounter of the target through an almost straight line (either as a lucky strike since the first trial or later in the Batch).

A difference between our study and previous ones (Walton, 2011) was the number of Batches each participant experienced in each delay condition. In our case, each participant completed 40 Batches in total compared to 2 in (Walton, 2011). This experience in the exploration task could have contributed to the development of strategies to accomplish the task that are responsible for this percentage of trials with very short distances. However, looking into the first couple of batches in our data indicated that there was still a populated early bin in the histogram, and log–transformed data were still falling into a bimodal distribution.

Another difference in our study is that we used a robotic manipulandum rather than a joystick. Perhaps the practice of the task with the vBots makes the task more vulnerable to the exploration techniques.

Of course the distribution of the data can be attributed to processes that we can not understand at the moment and would have to investigate further in the future.

A first step could be to confine the participants to using a random–only exploration technique and see if the same effect persists.

Excluding the trials that gave rise to lefthand lobe of the distribution of the data was not an option in the current analysis. Firstly, because they represented 25% of the total trials. Moreover, we could not filter the data using a consistent criterion and also within these short trials there were trials representing desired behaviour (finding a optimal path to the target).

Conclusions To sum up, tracking experience with delays in visual feedback of the target did not provide a mean to alleviate superstitious behaviour due to delayed visual reinforcement. This is because there must be mechanisms occurring during tracking that we had not accounted for in the first place. On the other hand, the duration of tracking training provided during our study might have been insufficient (see Appendix A.2). Improvements in tracking performance have been observed over days (Miall and Jackson, 2006). It could be that extended practice in the tracking task might lead to the expected effect and this will require further experiments. The remaining preferred explanation for the effect we observed is that tracking influences the certainty of localising the hand within the workspace. In the next chapter we present a study conducted to investigate the effects of tracking with delays on proprioceptive uncertainty.

Chapter 4
THE EFFECT OF TRACKING ON PROPRIOCEPTIVE UNCERTAINTY

4.1 Introduction

In the precious chapter, tracking experience with delays in visual feedback of the target did not produce the expected results. A possible explanation is that tracking with delays influences the certainty of localising the hand within the workspace. The aim of the study in this chapter was to investigate this hypothesis, i.e. that tracking with delays has an impact on proprioceptive uncertainty.

Adaptation to visuomotor rotations induces a spatial error between vision and proprioception. However, participants are able to adapt their movements by counteracting the effects of the imposed perturbation via an implicit adaptation of an internal model (Taylor and Ivry, 2011). In addition to changes in motor output during visuomotor adaptation (but also force–field adaptation), there is also a sensory recalibration component (Haith et al., 2008; Simani et al., 2007; Henriques and Cressman, 2012).

Delays are inherit in sensorimotor control in terms of afferent and efferent transmission delays and central processing. Disruption of temporal coherence between vision and proprioception impairs performance. However, it has been shown that participants are able to adapt to externally imposed delays in several tasks, such as: manual tracking (Foulkes and Miall, 2000; Miall and Jackson, 2006), guiding a computer mouse through obstacles projected on a monitor (Cunningham et al., 2001b) driving in a high–fidelity simulator (Cunningham et al., 2001a), prism adaptation (Kitazawa et al., 1995; Kitazawa and Yin, 2002; Tanaka et al., 2011), visuomotor adaptation (Honda et al., 2012a,b) and force field adaptation (Levy et al., 2010).

Temporal motor adaptation has been suggested to be a similar process to adapta-

tion visuomotor rotation. Specifically, it has been demonstrated to have after–effects (Cunningham et al., 2001b) and generalise in untrained environments (Cunningham et al., 2001a). This provides evidence that during temporal motor adaptation there is a new mapping constructed between the actions and the consequences of them. One might expect that temporal adaptation, similarly to spatial adaptation, would lead to recalibration in estimates of simultaneity. It has been shown that delays between the presentation of two different modalities of stimuli lead to temporal recalibration that also transfers between modalities (Di Luca et al., 2009). In another example, Tanaka et al. (2011) showed that exposure to a delay between reaching to a target and receiving visual feedback of the reaching shifted the perception of asynchrony between the two events. Moreover, (Cunningham et al., 2001a) mentioned in their results that several subjects reported that they felt that visual and haptic feedback appeared simultaneous at the end of the training.

The paradigm of Cunningham et al. (2001a), as all the paradigms mentioned before, involved movement and not discrete events and so temporal delay causes a spatial error. Are temporal delays perceived as spatial perturbations and perhaps vice versa? Unless explicitly informed about the delays, one should solve a spatiotemporal assignment problem. Are the errors experienced because of when or where? One would then expect that adaptation to delays could cause both temporal perceptual changes and intersensory recalibration.

In the present study we aimed to explore if there are any sensory recalibration effects occurring during a manual tracking paradigm. In this task, adaptation is exposed by improvement in performance in terms of a reduction in the error between the hand and the delayed hand representation but also in terms of performing smoother movements (Foulkes and Miall, 2000; Miall and Jackson, 2006). These improvements of motor performance imply an adaptation of a predictive forward

model, probably implemented in the cerebellum (Miall and Jenkinson, 2005).

We used two different conditions of the tracking task. In first the target was moving with variable speeds, as used before in (Foulkes and Miall, 2000). In a second condition, the target was moving with a constant speed. We hypothesised that if participants in both conditions were adapting to the delays per se then we should fail to find any changes in the perception of hand position after tracking or any changes found would be similar for both groups. Alternatively, if spatial adaptation was induced we would expect that the variable versus constant speed conditions would lead to differential effects in proprioception because of the variability in spatial errors in each case.

4.2 Materials and Methods

To address the hypothesis participants performed reach–out movements to targets in the workspace without being given any visual feedback of their movements, before and after a tracking task. We compared the accuracy of the movements pre and post tracking by assessing the mean End Point Error (EPE) to the target and the EPE variance.

4.2.1 Participants

72 right handed subjects (age range: 18-36 , mean=20, female: 48) participated in the study. 36 participants took part in Condition 1 and 36 participants took part in Condition 2. All of them had normal or corrected to normal vision. They did not have any restricted mobility or suffer from any neurological condition. They were informed about all the aspects of the experiment and gave informed written consent. The experimental protocol was approved by the Science, Technology, Engineering and Mathematics Ethical Review Committee of the University of Birmingham. Par-

ticipants received either cash or credits upon the completion of the study.

4.2.2 Apparatus

A vBOT robotic manipulandum (Howard et al., 2009) was used as an input device. An 30-inch Apple Cinema HD Display monitor was used to display visual objects viewed via the projection mirror surface (Figure 4.1). Subjects were seated on a stool and held the robotic arm with their right arm. The height of the stool was adjusted so that the shoulders of the subject were at the same height as the horizontal projection surface they were looking at. This allowed an almost fixed body posture throughout the study. The vision of the arm was obscured during the study.

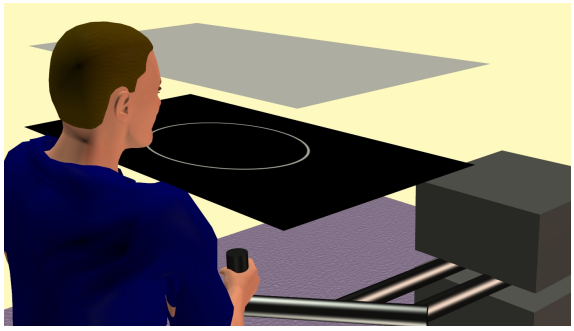

Figure 4.1: The Set – up. Subjects held the robotic arm with their right arm. A monitor (grey surface) displayed visual objects viewed via the projection mirror surface (black surface).

The background of the experimental area displayed on the monitor was black. In all the tasks of both experiments a white circle of 13 cm radius marked the area where the visual objects appeared. When there were no delays imposed on the visual feedback (see below Section 4.2.3) participants were physically constrained in within an area bounded by the circle by a high resistive position sensitive force field with added damping: $F[i] = 30 * (13 - |x|) * \frac{x[i]}{|x|} - 0.015 * v[i]\ N/m$, where i is the x and y components, F is the Force, p the position and v the velocity. This gave the feeling of bumping into a wall colocated with the circle. When there

were visual delays, the forces developed on the circle were those of a viscous force field:$F[i] = -0.015 * v[i]\ N/m$, where i is the x and y components, F is the Force and v the velocity. In this way the delayed visual and proprioceptive sensation were approximately aligned. However, the participants were not supposed to experience these forces as they were expected to remain within the workspace.

4.2.3 Procedure

The study consisted of two different tasks; a continuous tracking task and a discrete centre–out pointing movements task. In total, participants completed two blocks of 24 tracking trials each. Each tracking trial lasted for 45 seconds. At the end of each trial participants returned to the centre of the workspace and they were given three seconds of rest. Moreover, they could take a longer break between the two tracking blocks. In this task participants did have visual feedback of their hand and they were instructed to track the target as accurately as possible. The centre–out movements task was performed before and after the tracking task. Each time four blocks of 60 trials were completed (for design details see Table 4.1 below).

Tracking Task

In the tracking task, participants had to pursuit a circle target (4mm diameter) that was moving smoothly in an unpredictable 2-D path. An example of a tracking trajectory is shown in Figure 4.2. Participants were tracking the target with a cursor (6 mm diameter) that reflected the position of the vBot handle.

There were two tracking task conditions. In Condition 1, the pseudorandom target trajectories were generated in each axis independently as the sum of four non – harmonic sinusoids (0.11, 0.23, 0.35 and 0.42Hz) whose relative phases were randomised for each trial (Foulkes and Miall, 2000). This means that the target was moving with *variable speeds*. The speeds at which the target was moving were

between 1-20 cm/sec, with an average speed at around 8.5cm/sec.

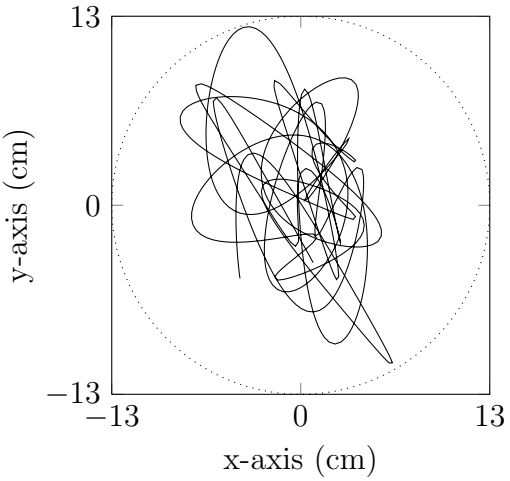

Figure 4.2: Example of a tracking target trajectory. The circle encloses the workspace.

In Condition 2, the target was moving with *constant speed*. 144 trajectories that were generated for Condition 1 were spatially resampled so that speed of the target motion would be constant, and on the average value of the speeds presented in Condition 1. The constant speed was therefore 8.5cm/sec.

In both Condition 1 and 2 participants were assigned in two groups. Participants in Group A tracked without any visual feedback delays imposed on the hand cursor. On the contrary, participants in Group B were presented a cursor that represented a 300ms delayed version of their hand position.

Participants were given feedback about their performance at the end of each trial in the form of a green upwards arrow in trials where the average error between the target and the cursor in the current trial was smaller compared to the previous trial. Otherwise, they received a red downwards arrow. However, they were instructed that the arrows showed their performance in the last trial compared to their overall performance up to that moment. Each green arrow was rewarded with extra 10 pence. The reason for introducing this feedback was to avoid subjective verbal instructions

for encouranging maximum performance.

Centre–Out Movements (COM) Task

In the centre–out movements task participants had to place the vBot handle in the centre of the workspace circle, which was marked by a crosshair. When they were within 0.5cm of the centre they saw a black dot of 0.33cm in diameter. The crosshair disappeared when the hand cursor was within a circle of 1mm radius around the centre of the crosshair. At that time a target appeared in one of 12 possible positions. Eight targets were in a circle of 9.75cm radius (75% of workspace radius) and four of them in a circle of 3.75cm radius (25% of workspace radius). The exact position of the targets is shown in Figure 4.3.

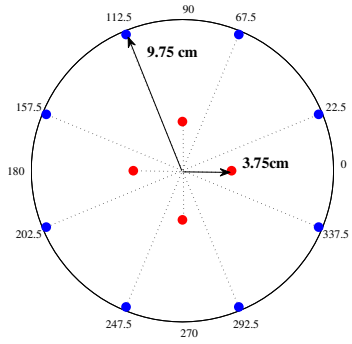

Figure 4.3: The 12 targets used in the Centre – Out Task.

Participants were instructed to make a fast, uncorrected movement to the target and conclude their movement on the target (van Beers et al., 2004). During the movement participants had no visual information about their movement. When the speed of their movement dropped below 0.4cm/s and the distance travelled from the centre was higher than that between the hand and the target, the target disappeared and the central crosshair re–appeared. Following that, participants had to go back

to the centre. The robotic arm would further assist them to return to the centre by applying a spring force drawing the vBot handle along a minimum jerk trajectory toward the central crosshair. All participants completed nine blocks of centre out movements. Each block entailed 60 centre–out pointing movements, five at each target. The targets appeared in a random order. The first block was a training block and was excluded from any further analysis. Four blocks were completed before the tracking task and four after the tracking task. The design is summarised in Table 4.1.

Table 4.1: Experimental Design. Participants performed a Centre – Out Movements (COM) task before and after a Tracking task.

COM task	Tracking task	COM task
1+ 4 Blocks x 5 Trials /Target ($\sim 15min$)	2 Blocks x 24 Trials ($\sim 40min$)	4 Blocks x 5 Trials /Target ($\sim 12min$)

4.2.4 Data Analysis

Tracking task

Matlab R2011b was used to analyse the data recorded. For each tracking trial, the RMS error between the target and the cursor positions was calculated after excluding the first 2 seconds of each trial. Moreover, a power analysis was performed. The positional data were filtered using a 4th order zero phase low pass Butterworth filter with a cut-off frequency at 10 Hz. A Hanning filter was used to smooth the data. We then calculated the velocity vector by taking the gradient of the position data independently in each axis. The power spectral density was taken using a Fast Fourier Transform in each axis. The mean power spectral density across the two axis was then calculated in the frequencies between 0.1Hz and 3Hz.

For each participant we calculated the mean RMS position error and average mean power spectral density of the velocity in the first 24 trials (Block 1) and the

last 24 trials (Block 2). If participants showed no adaptation across the two blocks, they were excluded from any further analysis. In total, three participants in each group were excluded. So, any further analysis is based on 15 participants in each group. The mean error and the average mean power spectral density were analysed separately for each delay group, using a 2x2 mixed ANOVA with a) one between participant factor *Tracking* condition (variable speed – Condition 1 or constant speed – Condition 2) and b) one within participant factor of *Block* (Block 1 or Block 2).

Centre–out movements task

Matlab R2011b was used to analyse the data recorded. For each centre out movement the absolute peak velocity was found. The end point of each movement was taken as the point at which the velocity had dropped below 5% of the absolute peak velocity. The End Point Error (EPE) was defined as the euclidean distance between the end point of the movement and the target at which each movement was directed to. In order to exclude any outliers, a Principal Component Analysis (PCA) was performed (using princomp Matlab function) for each target and for each participant (separately for the before and after tracking movements). The reason for performing a PCA was to further compute Hotteling's T^2 (one of the output measures of princomp function), which is a measure of the multivariate distance of each observation from the centre of the data set. Any point that had Hotteling's T^2 above two standard deviations from the mean Hotteling's T^2 (the centre of the dataset) was excluded. Finally we calculated the signal to ratio (SNR), defined as the mean error over the standard deviation of the error, for the four inner and the eight outermost targets. We found that the SNR of the four inner targets was much lower ($meanSNR = 2.01, StDevSNR = 0.15$) compared to that of eight outer targets ($meanSNR = 3.94, StDevSNR = 0.299$). We thus excluded the four inner

targets from further analysis.

The mean EPE and the standard deviation of the EPE for each of the eight outer targets were used as dependent variables for the statistical analysis. For each dependent variable a 2x2x2x8 mixed ANOVA was used with a) one between participant factor of *Tracking* condition (variable speed- Condition 1 or constant speed -Condition 2), b) one between participant factor of *Delay* in the visual feedback (two conditions of with or without delay), c) one within participant factor of *Block* (Block 1 – pre tracking or Block 2 – post tracking) and d) one within participant factor of Targets (8 targets).

4.3 Results

4.3.1 Tracking

General Obervations

Typical tracking behaviour of individual trials in all four groups is presented in the graphs of Figure 4.4. Tracking errors are larger in early trials compared to late trials, especially when participants tracked with visual feedback delays. Over practise, participants improved their tracking behaviour by minimising the error between the target and the cursor. As expected, the individuals that tracked with visual feedback delays did not achieve the same small errors as the groups that tracked without visual feedback delays (Foulkes and Miall, 2000).

The improvement of performance is also depicted in the intermittency of the tracking behaviour in late trials compared to the early ones. Participants managed to perform smoother movements over practice. Again, there is a noticeable difference between the groups that tracked without and with visual feedback delays. The improvement in the intermittent behaviour is mirrored in the power spectra in Figure 4.5. Overall, the spectra have less power in late trials as a result of the

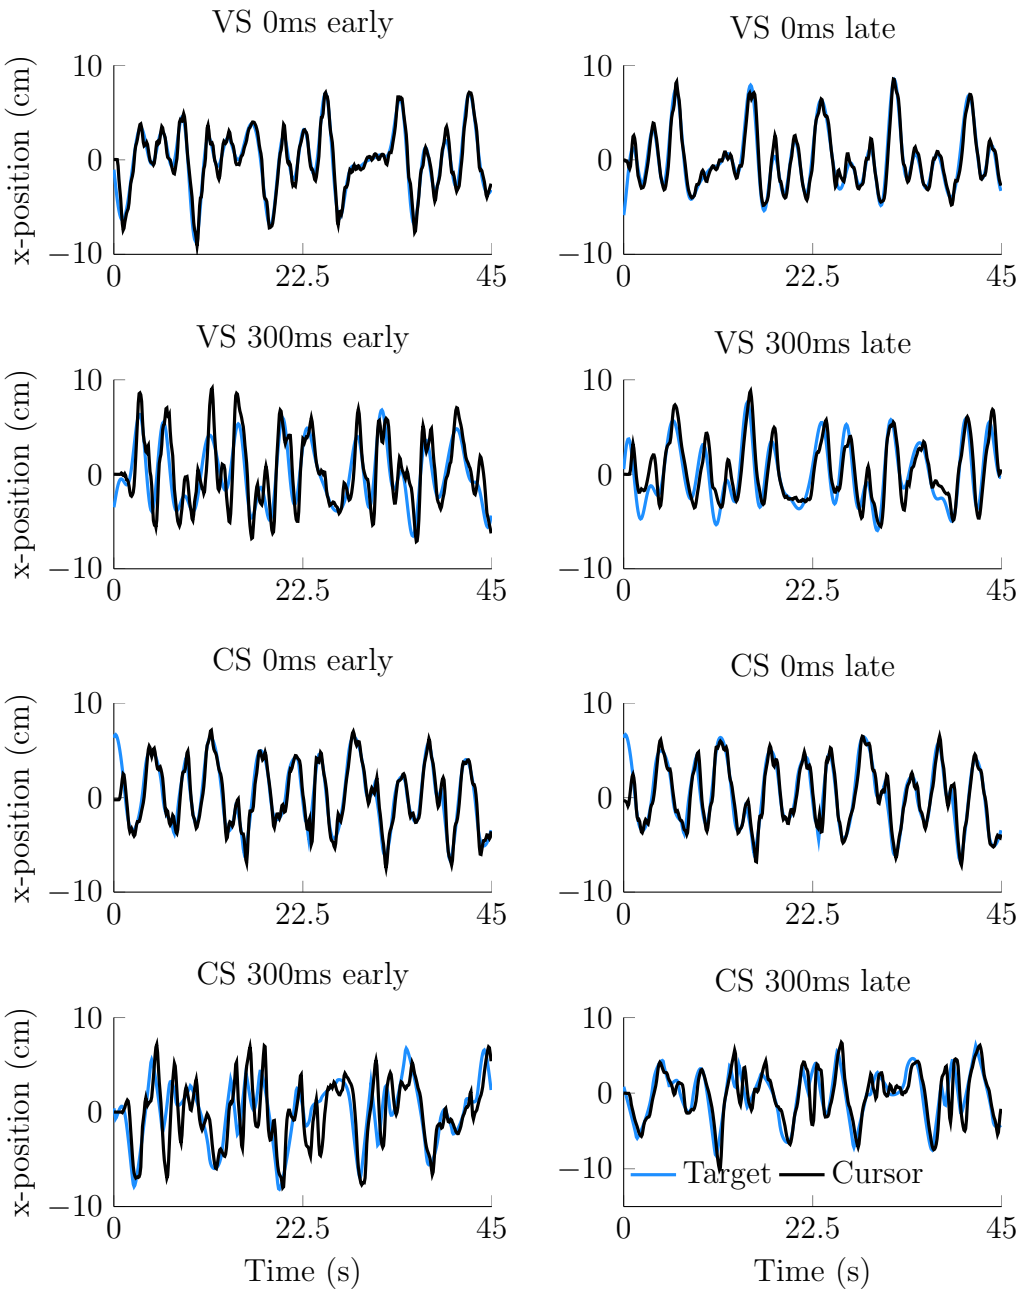

Figure 4.4: Error between the target (blue line) and the cursor (black line) in early and late trials, for individuals that tracked without (0ms) and with (300ms) visual feedback delays in the Variable Speed Group (VS) and the Constant Speed Group (CS).

smoother movements achieved towards the end of training. In line with previous studies (Foulkes and Miall, 2000), during early trials there is a power component

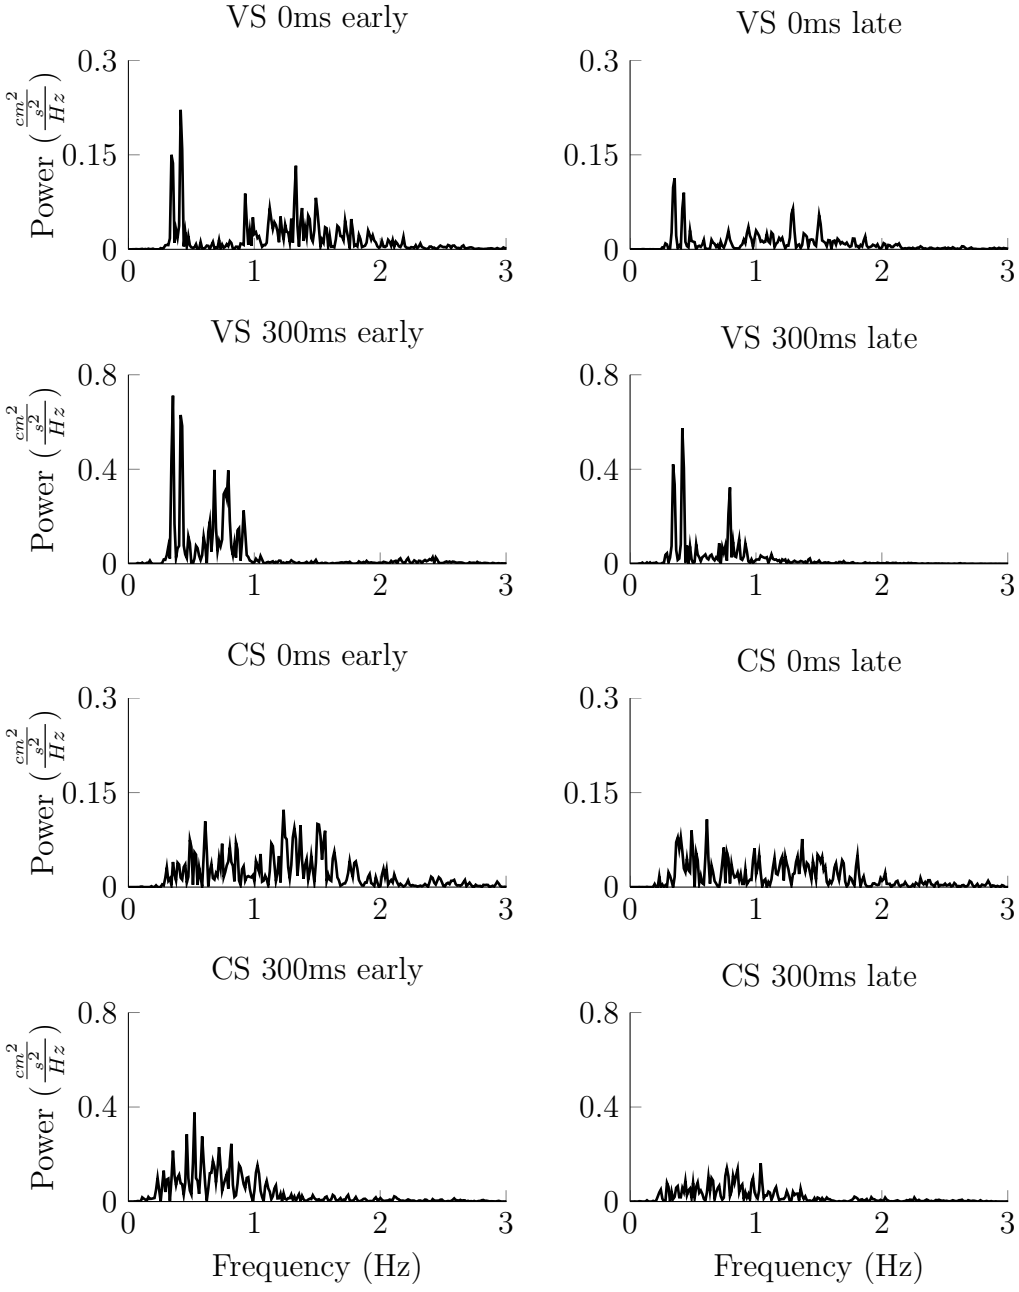

Figure 4.5: Power Spectral Densities of individual early and late trials of participants that tracked without (0ms) and with (300ms) visual feedback delays in the Variable Speed Group (VS) and the Constant Speed Group (CS). Each spectrum is the mean power of the velocity error calculated separately in horizontal and vertical axes. Note the change in scale for the vertical axis in the 300ms delay condition.

around the frequency of 1.2Hz for participants that tracked without visual feedback delays and around the frequency of 0.7Hz for participants that tracked with visual feedback delays. These components are less dominant in late trials. Moreover, the spectra of the participants that tracked in the variable speed condition compared to those of the participants that tracked with constant speed differ in the power of the components around the frequencies of 0.1–0.5Hz. This is the range of target frequencies (0.11, 0.23, 0.35 and 0.42Hz). In case of the constant speed condition, spatial resampling of the trajectories distributed the power across a broader range of frequencies than the constituent sinusoids.

Improvement in tracking performance

The RMS error between the target and the cursor was the first measure used to assess the adaptation of the participants in the tracking task. The groups that tracked without visual feedback delays showed no significant improvement in terms of minimising their error (see Figure 4.6.B). There was no main effect of Block ($F(1, 28) = .001$, $p = .979$, $\eta^2 < .001$) or of Group ($F(1, 28) = 2.356$, $p = .136$, $\eta^2 = .078$) and there was no significant interaction of Block*Group ($F(1, 28) = .446$, $p = .510$, $\eta^2 = .016$).

Both groups that tracked with visual feedback delays improved their performance significantly across the two experimental blocks (see Figure 4.6.A). There was a significant main effect of Block ($F(1, 28) = 48.697$, $p < .001$, $\eta^2 = .635$) but no significant effect of Group ($F(1, 28) = .919, p = .346, \eta^2 = .032$). Moreover, there was a significant interaction of Block and Group ($F(1, 28) = 10.410, p = .003, \eta^2 = .271$). This interaction was a result of: 1) a significant drop ($t = .0013$, adjusted for multiple comparisons Sidak) in the mean RMS tracking error by Δmean=.091 in the group that tracked in the Constant Speed condition and 2) a significant difference ($t < .001$,

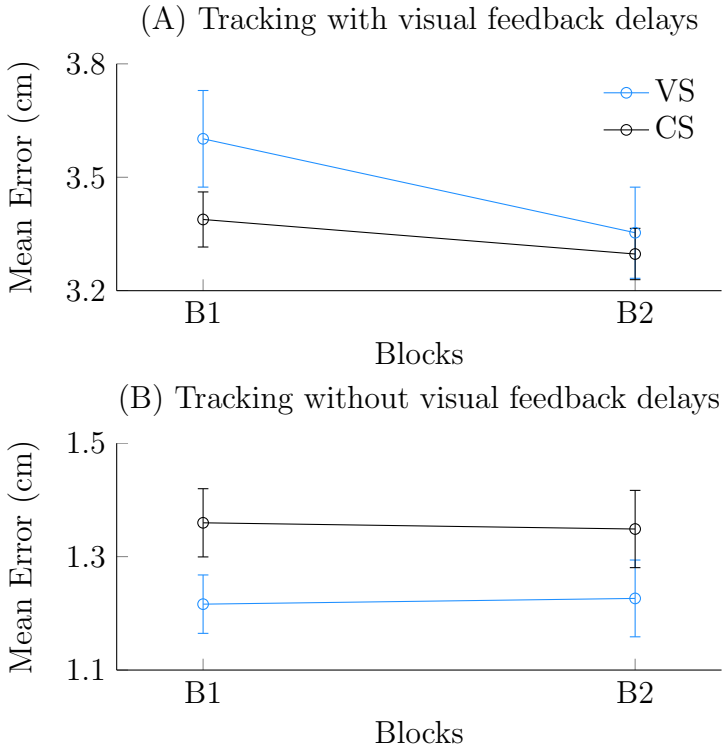

Figure 4.6: Mean RMS error between the target and the cursor in the Groups that tracked with visual feedback delays (A) and the groups that tracked without visual feedback delays (B). VS= Variable Speed Group, CS= Constant Speed Group. The error bars show the standard error of the mean.

adjusted for multiple comparisons Sidak) in the mean RMS tracking error by Δ mean=.248 across the two Blocks in the group that tracked in the Variable Speed condition. Thus, the groups exposed to the delay showed significant learning, whereas the groups without delay did not (see Table B.1 in Appendix B.1).

The mean power spectral density was the second variable used to study the adaptation to the tracking task. Participants were expected to achieve lower mean power spectral density with practice as a result of adapting to the task and executing smoother movements. Figure 4.7 shows that all four experimental groups did improve over the practice of the task.

The mixed 2x2 ANOVA in the groups that tracked without visual feedback delays

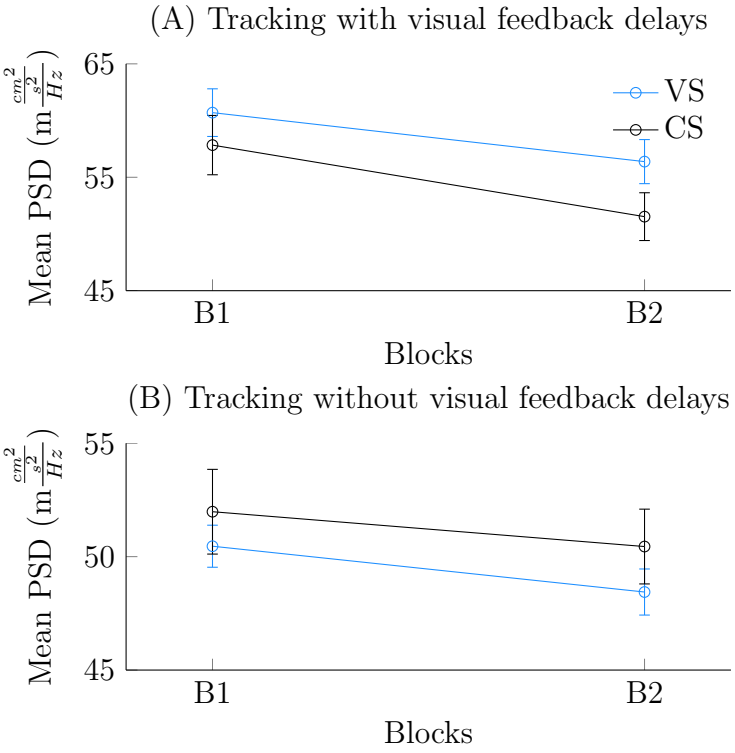

Figure 4.7: Mean Power Spectral Density in the Groups that tracked with visual feedback delays (A) and the groups that tracked without visual feedback delays (B). VS= Variable Speed Group, CS= Constant Speed Group. The error bars show the standard error of the mean.

showed that there was a significant main effect of Block ($F(1, 28) = 41.81$, $p < .001$, $\eta^2 = .599$) but no significant effect of Group ($F(1, 28) = .783$, $p = .384$, $\eta^2 = .027$) or of the Block*Group interaction ($F(1, 28) = .775, p = .386, \eta^2 = .027$). In the groups that tracked with visual feedback delays the results of the 2x2 ANOVA were similar. There was a significant main effect of Block ($F(1, 28) = 19.09, p < .001, \eta^2 = .405$) but there was no main effect of Group ($F(1, 28) = 1.803, p <= .190, \eta^2 = .060$) or any interaction of Block*Group ($F(1, 28) = .669$, $p = .420$, $\eta^2 = .023$).

4.3.2 Centre–Out Movements

General Observations

Figure 4.8.A1 shows the end point positions for the centre–out movements of a typical subject that participated in Condition 1 (variable tracking speed) in Group A (no visual feedback delays). There seems to be an enhanced mean accuracy of the pointing movements after tracking. On the other hand, a typical subject from Group B (visual feedback delays) shows an increased variance of her movements after tracking (see Figure 4.8.B1).

However, adaptation to tracking appears to have no effect in pointing accuracy in Condition 2 (constant speed tracking) in both Group A and Group B, as shown in Figures 4.8.A2 and 4.8.B2 accordingly.

No spatial shift after adaption to tracking

Figure 4.9 shows the mean End Point Errors (EPE) in the centre–out movements before and after tracking (collapsed across the 8 targets) in all four groups. There is a tendency towards smaller errors in the groups that tracked without delays in the variable speed condition (Figure 4.9.A1) and towards larger errors in all other conditions (Figures 4.9. A2/B1-2).

However, none of these differences were significant as exposed by the analysis. The results of the 2x2x2x8 mixed ANOVA of the mean EPE showed that there was no effect of the adaptation to tracking on this variable. Specifically, there were no main effects of the within participant factor of Block ($F(1, 56) = .105$, $p = .748$, $eta^2 = .002$), neither for the between participant factors of Tracking condition ($F(1, 56) = .041$, $p = .841$, $\eta^2 = .001$) and Delay ($F(1, 56) = 2.097$, $p = .153$, $\eta^2 = .036$). There were also no interactions among the aforementioned factors (refer to Table B.2 in Appendix B.1 for details).

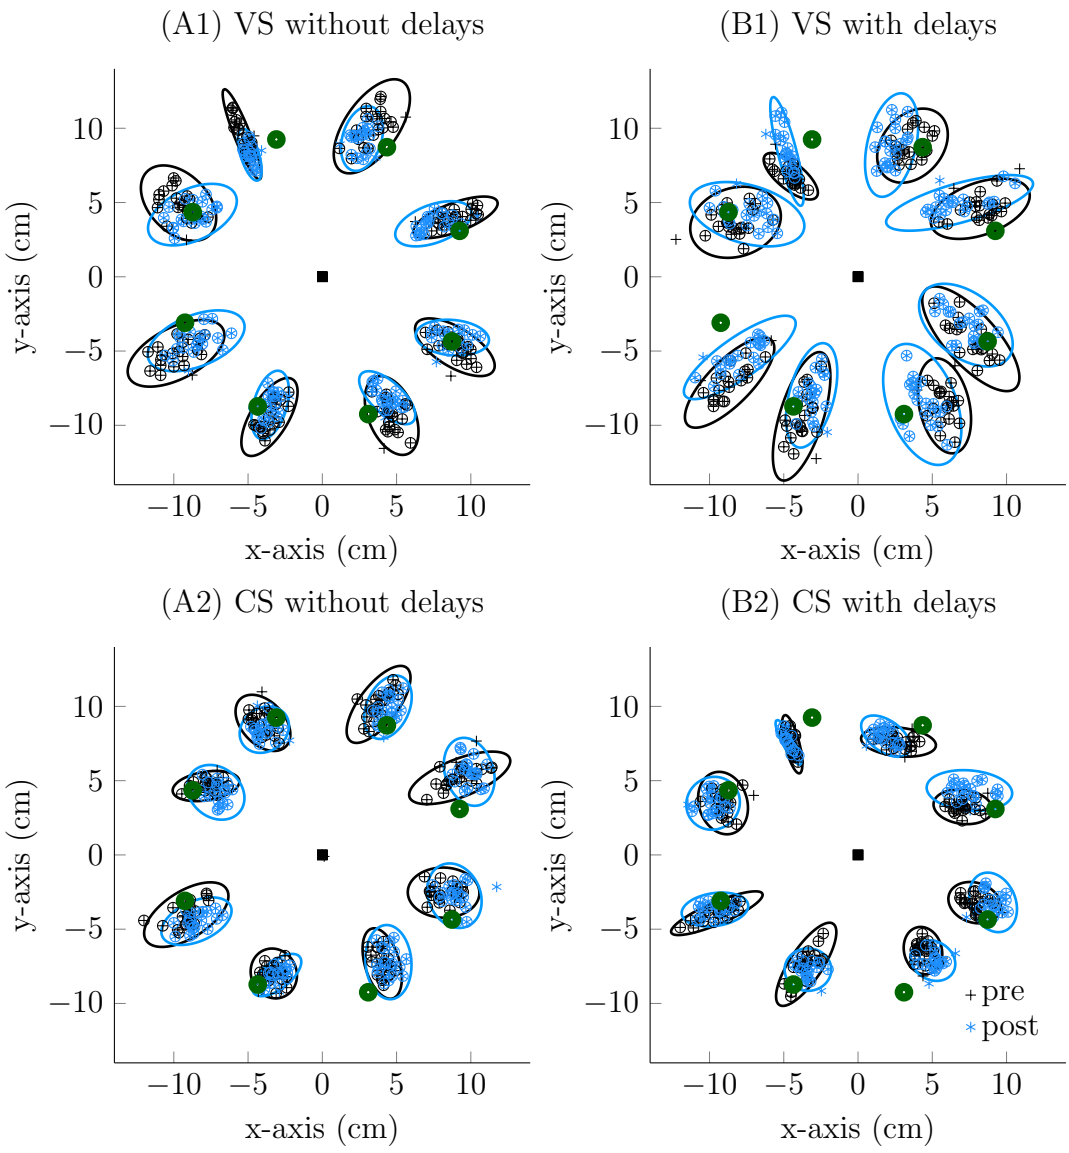

Figure 4.8: Examples of the end–points of the centre–out movements of individuals before (blue stars) and after (black crosses) exposure to the tracking task: in Condition 1 (Variable Speed condition–VS) without visual feedback delays (A1) and with visual feedback delays (B1) and in Condition 2 (Constant Speed condition–CS) without visual feedback delays (A2) and with visual feedback delays (B2). The ellipses around the distributions of the end–points are the 95% confidence ellipses. The green dots represent the actual positions of the targets. The black point in the middle is the centre of the workspace. The points (stars and crosses) that are without circles are outliers.

However, there was a significant effect of the within – participant factor of Targets ($F(7, 392) = 10.153$, $p < .001$, $\eta 2 = .153$). This result was expected because of the

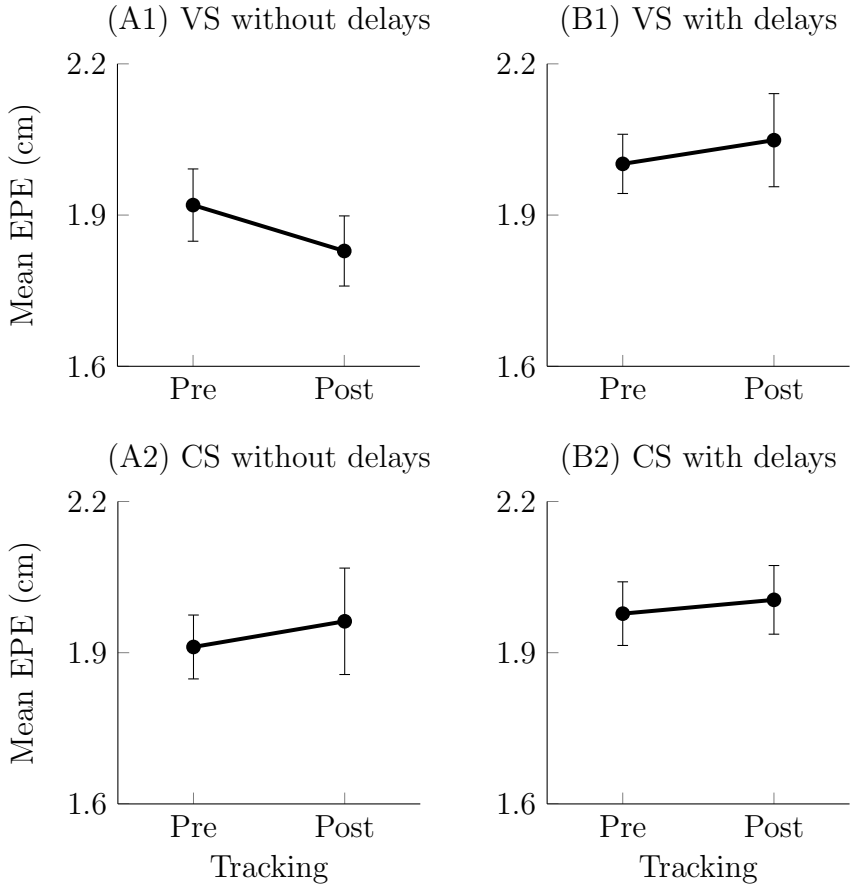

Figure 4.9: Mean End Point Errors of the Centre–Out Movements before and after Tracking in Condition 1 (Variable Speed condition–VS) without visual feedback delays (A1) and with visual feedback delays (B1) and in Condition 2 (Constant Speed condition–CS) without visual feedback delays (A2) and with visual feedback delays (B2). The error bars show the standard error of the mean.

differences in accuracy when pointing to different locations of space (van Beers et al., 2004). Moreover, there was a significant interaction between the Targets and the Block ($F(5.4, 302.8) = 4.544$, $p < .001$, $\eta 2 = .075$). This significant interaction was due to changes in the mean EPE in 3 of the 8 targets (Table B.5 in Appendix B.1). Moreover, there were some significant differences in mean EPE between the targets, both in Block 1 and in Block 2 (Tables B.3 and B.4 in Appendix B.1).

Adaptation to tracking with variable speeds affects proprioceptive uncertainty

The effect of the adaptation to tracking in the proprioceptive uncertainty is shown in Figure 4.10. There is a prominent decrease of the mean standard deviation in the group that tracked without visual delays in the variable tracking speed condition (Group A, Condition 1), as shown in Figure 4.10.A1.

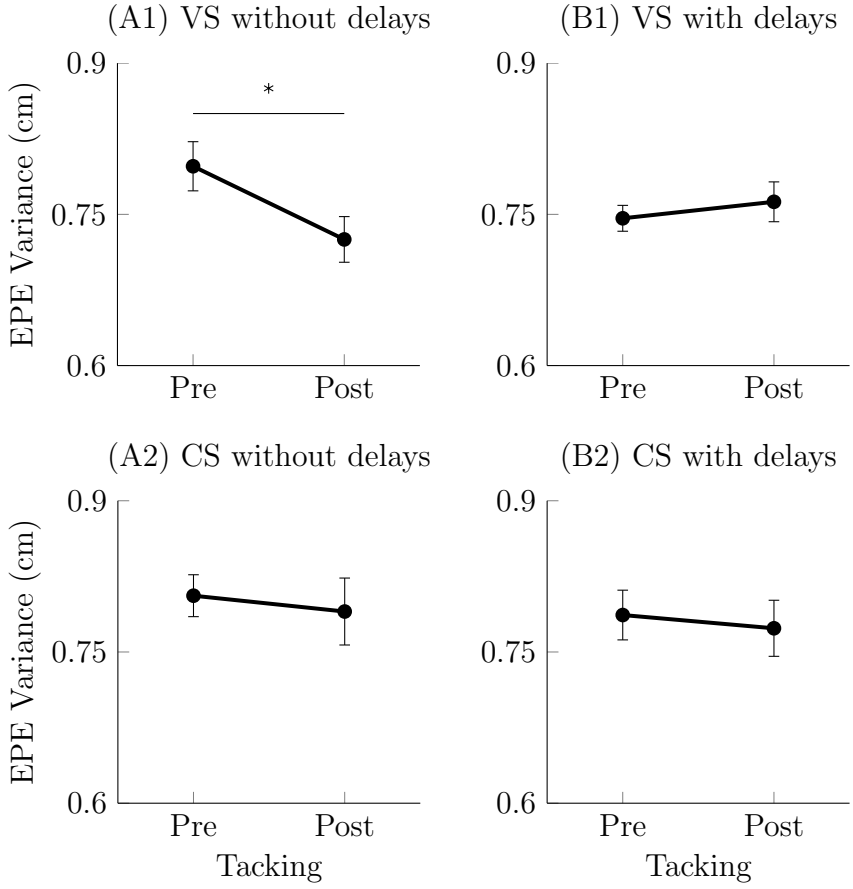

Figure 4.10: Variance of the End Point Errors of the centre–out Movements before and after Tracking in Condition 1 (Variable Speed condition–VS) without visual feedback delays (A1) and with visual feedback delays (B1) and in Condition 2 (Constant Speed condition–CS) without visual feedback delays (A2) and with visual feedback delays (B2). The error bars show the standard error of the mean.

On the other hand, the group that tracked with delayed visual feedback at Con-

dition 1 shows an increase in its variance (Figure 4.10.B1). Finally, both groups that tracked with constant speed appear to have a minor decrease of their variance (Figure 4.10.A2-B2).

The mixed 2x2x2x8 ANOVA of the standard deviation of the End Point Errors showed that there was a significant main effect of Block ($F(1, 56) = 4.153$, $p = .046$, $\eta^2 = .069$), a significant interaction of Block*Delays ($F(1, 56) = 4.805$, $p = .033$, $\eta^2 = .079$), and of Block*Tracking*Delays ($F(1, 56) = 4.284$, $p = .043$, $\eta^2 = .071$) (details of this analysis are presented in Table B.6 in Appendix B.1). A simple main effects analysis was run to further explore these findings. It was found that these interactions arise because of a significant ($p = .001$, adjusted for multiple comparisons Sidak) decrease in the average variance by Δmean=.073 in the group that tracked without visual feedback delays in the variable speed tracking condition (Group A, Condition 1). On the other hand, the groups that tracked with visual feedback delays at the same tracking condition showed an overall increase of the average variance (Δmean=-.016). Both groups that tracked with constant speed velocity showed a decrease of their variance by Δmean=.016 in case of no visual feedback delays and by Δmean=.013 in case of visual feedback delays. None of these three results were statistically significant (the results of the pairwise comparisons are presented in Table B.7 in Appendix B.1). Moreover, there were no significant baseline differences among the four groups.

As expected, there was a main effect of the Target position ($F(6.145, 344.104) = 22.395, p < .001$, $\eta^2 = .291$) due to the different movement characteristics across the workspace. No interactions between the Targets and the other factors were observed (see Table B.6 in Appendix B.1).

4.4 Discussion

The results of the present study showed that training in a manual tracking task, where the speed of the target was variable, significantly increased movement accuracy in a subsequent centre–out reaching task. On the other hand, practice and adaptation to delays in the same tracking task caused a smaller, statistically insignificant decrease in accuracy. Moreover, adaptation to a tracking task where the target moved with a constant speed, both with and without the presence of visual feedback delays, did not cause any significant effects in reaching performance.

Improvement in tracking performance

In line with previous studies (Foulkes and Miall, 2000; Miall and Jackson, 2006), participants that tracked with variable target speed (Condition 1) both in with and without visual feedback conditions showed an improved performance in tracking. A significant decrease in mean power spectra was found in both groups. This result indicates that participants adapted to the task in terms of performing smoother movements and were managing to appropriately intercept the target, suggesting a modification of predictive feedforward movements (Miall and Jackson, 2006). Moreover, adaptation to tracking delays (and variable target speeds) was further exposed by the significant decrease in end–point tracking error.

The findings for the groups that tracked a target with constant speed (Condition 2) in terms of improvement in performance were the same. A difference between the two groups with delays in Condition 1 and 2 was found in terms of end–point tracking error decrease. There was a higher drop in error in Condition 1 and since there were no significant differences between the groups in Block 1 and 2, we conclude that learning with variable speeds might have driven a faster learning rate.

However, it is important to remember that these observations about the effect

of practise in the tracking task are based on only 40 minutes of exposure to delays. Adapting to visuomotor tracking, especially with delays, can continue over several days (Miall and Jackson, 2006), so unless tested we cannot know how the current results will generalise to longer exposure to these delays (see Appendix A.2).

What was measured in the reaching task? Reaching movements, before any tracking, present baseline variability attributed to inherit execution noise (van Beers et al., 2004) and planning accuracy (Gordon et al., 1994). During the reaching task there was no visual feedback of the arm. Thus, part of the variance seen in the EPE must be because of inherit proprioceptive uncertainty and perhaps a further increase in it induced by the unfamiliar to the participants set–up.

By using the reaching task before and after the tracking task, we intended to measure proprioception changes. It has been argued that using voluntary reaching movements to assess proprioceptive recalibration could be a misleading way to do so (Henriques and Cressman, 2012). This is because any changes observed could be due to both sensory and motor recalibration. However, Simani et al. (2007) showed that realignment tests and a reaching task after tracking were correlated, providing evidence that reach after–effects are due to sensory recalibration.

No proprioceptive shifts after tracking

The first measure we used to assess changes in proprioception was mean End–Point Error (EPE) between the targets and the end–point of the reaching movement. Here, we expected that since the temporal delays in the tracking would induce a positional error between vision and proprioception, adaptation to the delays would lead to sensory realignment. Perhaps no differences would be expected between the two tracking conditions since the average positional error in the two cases was the same and what was changing (due to variable target speeds in Condition 1) was the

variance of the error.

Subjects that participated in the variable speed condition (Condition 1) in the no delays group showed a tendency towards smaller errors, compared to the other three groups. However, based on the statistical analysis, we failed to find any changes in mean EPE before and after tracking, across the groups and the conditions. This result indicates that tracking experience (with and without delays) did not induce any direction–specific shifts in proprioception. Perhaps this should have been expected as the visual errors were always occurring along the path of the tracking trajectory and thus they were varying in direction. So any corrective movements to counteract the effects of errors were not towards the same direction, in contrast to visuomotor adaptation paradigms where sensory realignment is observed (Henriques and Cressman, 2012). Moreover, in the tracking task movements are continuous and variable, so that the same error is observed under different circumstances. Differences in reach after–effects after exposure to either a tracking or to a reaching task were observed before (Simani et al., 2007). Simani et al. attributed this finding in differences in error–corrective learning rule exposure which is less obvious and intense in tracking compared to reaching.

A final point to consider is the amount of the expected shift. The average tracking speed of participants was around 10cm/sec. This is 15% higher than the average speed because of overshooting at target turning points and because of directional errors (see Figure 4.4). If the adaptation to tracking was translated to a pure spatiotemporal transformation, the expected spatial shift after tracking, in the visual feedback condition would be on average 3cm. This value is almost four times greater than the variance of the centre – out movements before any tracking. Thus, we can exclude the scenario that any spatial shift would have been occluded by the inherit movement variance.

Changes of proprioceptive uncertainty after tracking

The second variable chosen to examine differences in proprioception was the variance of EPE, that would depict any changes in proprioceptive uncertainty. We hypothesised that the effect of the tracking condition would lead to different results in terms of reaching variance, given that in Condition 1 (variable target speed) the positional errors were variable and in Condition 2 (constant target speed) constant. The main finding here was a significant decrease in EPE variance in the group that tracked without delays.

When comparing the two tracking conditions (variable versus constant target speed) we might expect a similar change in proprioception since both groups showed the same tracking performance over the blocks. The decrease of the mean power spectral density during tracking showed that in both cases participants adapted in intercepting the target in case of sudden direction changes. However, in the variable speed tracking condition participants had to intercept the target by performing corrective movements during variable speed changes. So, we can conclude that by being exposed to a more variable task participants became more accurate. Since the tracking performance in both groups is comparable that means that both groups adapted in a similar way in terms of motor performance. This observation adds to the argument that the differences in EPE variance in the case of tracking a target with variable speeds is because of changes in proprioception rather than motor adaptation.

The delays in visual feedback introduced an additional component of error. While in Condition 2 (constant target speed) this error was constant along the tracking direction, in Condition 1 (variable target speeds) it was frequently changing. This variability in the perturbations perhaps made it more difficult for the participants to differentiate between self–uncertainty and task-dependent errors, leading to an

increased uncertainty overall. Although, the increase in EPE variance was not statistically significant, given the sign direction of the change in EPE variance compared to all other groups and conditions, we should consider the possibility that a significant effect might have been concealed by order effects.

Differences between the targets

Baseline differences in variance among the movements towards the various locations of the task targets were expected because of differences in the execution movement noise in various directions of the workspace (van Beers et al., 2004) and differences in planning accuracy (Gordon et al., 1994). This is why we found the Target factor to be significant in the statistical analyses. Moreover, in case of the mean EPE we also found an interaction between Targets and Block. The lack of any other interactions between Targets*Block and the other between participant factors (Speed or Delay) underlines that this result can not be reliably related to the tracking experience. It could instead be an effect of the experience gained through practicing the centre–out movements, in case of a decreased mean EPE, or an effect of fatigue in case of an increased mean EPE.

Underlying neural structures

The cerebellum has been related to visuomotor co–ordination and adaptation to delays between vision and proprioception (Miall et al., 2001; Miall and Jenkinson, 2005), possibly by retaining internal forward models that are able to predict the sensory consequences of our actions and allow the computation of sensory prediction errors. So, to a large extent changes seen in the tracking task are expected to be cerebellar–dependent. Sensory realignment on the other hand has been suggested not to be related with cerebellum and possibly to depend on posterior parietal cortex (Block and Bastian, 2012).

Conclusions

To sum up, we found that adaptation of movements during a tracking task improves proprioceptive accuracy when the tracking target moves with variable speeds but not with constant speeds. Temporal delays between seen and felt hand position disrupt tracking performance and, in the case of variable speeds, lead to a small increase of proprioceptive uncertainty in the early stages of adaptation to delays.

Chapter 5
THE EFFECT OF TDCS IN A NOVEL EXPLORATION TASK

5.1 Introduction

There can be several different components during novel action acquisition that contribute to learning. Short–latency phasic activity of dopamine neurones in basal ganglia could serve to associate the delivery of rewards with recent motor output (Redgrave et al., 2008). On the other hand, action learning can be influenced by high–level cognitive/declarative areas. For example, Hikosaka and Wurtz (1983) showed that lateral prefrontal cortex can hold spatial information used to guide motor behaviour.

The issue of non–declarative and declarative components in the exploration task has been raised by Walton (2011). In a variant of the exploration task where participants had to explore and learn performing gestures, it was shown that exposure to spatial information had little impact on learning.

The aim of the study in this chapter was to investigate the involvement of different brain areas in the exploration task. To do so, we used transcranial direct current stimulation (tDCS) over the motor cortex (M1), dorsolater prefrontal cortex (dlPFC) and cerebellum.

Motor cortex projects directly to the striatum (putamen) (Alexander et al., 1986). The cortico–basal–ganglia connections have been proposed to represent state–action value for possible actions (Samejima and Doya, 2007) and changes to pre– or post– synaptic weights to the striatum could be one of the mechanisms mediating action selection (Redgrave et al., 2008) in the basal ganglia. M1 stimulation could potentially bias action selection by affecting the input salience (magnitude signals) of pre–synaptic weights to the striatum (Redgrave et al., 2008). It could be further

speculated that anodal M1 stimulation, which enhances excitability (Jacobson et al., 2012), would increase the probability of a more favourable movement to be selected, via the aforementioned mechanism. Selection of more favourable movements would perhaps be depicted in lower irrelevant distances.

The dorsolateral prefrontal cortex projects to the dorsolateral head of the caudate nucleus (Tekin and Cummings, 2002). dlPFC is related to "the maintenance and manipulation of belief states" (Samejima and Doya, 2007), for example by processing spatial information (Hikosaka and Wurtz, 1983). We should note here that this state estimation can be quite different from the cerebellar state estimation, in that in case of the dlPFC we can be talking about models of the world whereas in case of cerebellum the focus is on state estimation of own body state. Here we hypothesised that anodal dlPFC stimulation, which enhances working memory (Fregni et al., 2005), could lead to more certainty about the task state, perhaps the movement leading to successful behaviour. In other words, anodal dlPFC stimulation could lead to the enhancement of use of spatial information. That would not necessarily lead to optimal action selection. On the contrary, it could encourage a more superstitious behaviour (Skinner, 1948).

Finally, we were interested in further exploring the functional interaction between cerebellum and basal ganglia. Cerebellar tDCS could mediate changes in the basal ganglia in three (at least) different ways. As discussed in previous chapters, the cerebellum has been shown to have direct anatomical connections to basal ganglia (Bostan et al., 2013). So the first way that cerebellum could be affecting basal ganglia function would be through the direct cerebello–basal ganglia anatomical pathway. However, we could make no predictions about the functional nature of this interaction (whether motor or cognitive). The second way that cerebellar stimulation could affect the basal ganglia would be via the motor cortex. Stimulation of the

cerebellar cortex leads to changes in the cerebellar–brain inhibition (CBI) (Galea et al., 2009). In particular, cathodal stimulation results in higher excitability of motor cortex by reducing the CBI. In this case we would expect that cerebellar stimulation will lead to similar effects to M1 stimulation. Finally, a third way that cerebellar stimulation could have an effect on the exploration task is via the cerebello–thalamo–dlPFC loop (Middleton and Strick, 1994). Cathodal cerebellar stimulation has been shown to induce similar cognitive effects to anodal dlPFC stimulation (Pope and Miall, 2012). So, maybe cerebellar stimulation would lead to changes in behaviour similar to anodal dlPFC stimulation.

5.2 Materials and Methods

5.2.1 Participants

80 right handed subjects[1] (mean:23 age range:18-55, 24 male) participated in the study. All of them had normal or corrected to normal vision. They did not have any restricted mobility or suffer form any neurological condition. They were informed about all the aspects of the experiment and gave informed written consent. The experimental protocol was approved by the Science, Technology, Engineering and Mathematics Ethical Review Committee of the University of Birmingham Ethics Committee. Participants received either cash or credits upon the completion of the study.

[1]The data in this study were collected by Sam Westwood and Alexandra Doherty as part of their undergraduate dissertation. I provided to them the experimental and analysis codes and contributed to their supervision. The data analysis presented in this chapter is entirely conducted by me.

5.2.2 Experimental Set–up

Participants sat around 50 cm away from a 19–inch LG Flatron monitor that was used to deliver visual feedback. A custom–made contactless joystick (6.5cm in length, of light weight and low friction, with the self–centring spring removed) was used as an input device in this study. They held the joystick while keeping their forearm rested on a flat surface. The joystick was calibrated so that the limits of the joystick's excursion were in correspondence with the edges of the monitor.

5.2.3 Exploration Task

Participants had to explore the workspace, which was extended to the whole screen surface, and find a hidden target area (Stafford et al., 2012). No visual feedback of the cursor position was provided. At the beginning of each trial participants had first to locate a circular area 100 pixels ($\simeq$2.65 cm) in diameter that would serve as their starting point (start area). When they reached the start area, a blue circular disk flashed with a 380 ms delay. Participants were never told explicitly about the delay. The disk had the actual size and location of the start area. They were then instructed that they had to explore the workspace and locate another circular area that was defined as the target area. The target area had a diameter of 250 pixels ($\simeq$6.6cm) and it occupied 4.7% of the workspace. 380 ms after participants reached the target area they received a whole–screen yellow flash. Again, participants were not instructed that they would receive the 'reinforcement' flash with delay. Following that, they had to return to the start area and then immediately try to relocate the target area, as efficiently and as quickly as possible. After 10 trials, the start area and the target area were randomly assigned new locations, set at least 525 pixels apart (from centre to centre). Each set of 10 trials (with the same start and target area) will be referred to as a Batch. Moreover, the centre of start and target areas could

not be found within 200 pixels of the border of the workspace (screen). Finally, a break of three seconds was given between the Batches and a 15 seconds break every ten Batches.

The reason for having the start area located to a different point at each batch (as opposed to having it at the same location, for example the centre of the workspace) was to avoid the adoption of a specific exploration strategies by the participants (see Section 3.3.2 in Chapter 3 for more details). The 380ms of delay was chosen as close to the 450ms delay shown to be effective to impair performance in the exploration when using the joystick (Walton, 2011; Thirkettle et al., 2013).

Another major change in the paradigm of this chapter is the travelling from the target area to the start area. We decided not to present the start area until participants actually reached it. In this way we hope to promote a continuous movement between the target and the start – participants were looking for two targets in a sense. This manipulation perhaps would leave way for the exposure of declarative components of learning, in a similar way to the gesture version of the exploration paradigm (Walton, 2011). Moreover, the introduction of a delay between the start area and its presentation could add an extra spatial/cognitive component to the task. In a sense, participants had the opportunity to infer the spatiotemporal characteristics of the task and so improve their performance.

Finally, we chose a larger target area compared to that in the example of Chapter 3 as we wished to make the encounter of the target area relatively easier in order to achieve greater number of successful trials.

5.2.4 tDCS

A Magstim Eldith DC–Stimulator was used to deliver the direct current via a pair of conductive rubber electrodes (thickness (mm): $d \simeq 1$, surface (cm^2): A$\simeq 3.7 \times 3.7$)

that were inserted in NaCl-moistened sponge pockets (surface (cm^2): $A_{sponge} \sim 4.7 \times 5.72$). Depending on the brain area stimulated (see below) the current intensities used were $1.5mA$ and $2mA$ corresponding to 0.056 and 0.075 mA/cm^2 current density accordingly. These current density values were within the safety limits (Liebetanz et al., 2009). The duration of the stimulation was 20 minutes. When real stimulation was delivered, the current intensity was raised in a ramp-like way over the first 10 seconds until the desired level and faded out the same way at the end of stimulation. In case of pseudo stimulation (control condition) a current pulse of 110μA of intensity and 15 ms duration was delivered every 550 ms, giving the participants a similar sensation to that of the real stimulation but without eliciting the same neuromodulating effects.

Depending on the stimulation area, participants were randomly assigned in four different groups in a double–blind procedure. Participants in the first group received $2mA$ (Galea et al., 2009; Ferrucci et al., 2008) of cathodal cerebellar stimulation (cerebellar group). The active electrode was placed in the right cerebellar cortex 1cm below and 4cm laterally to the inion (Pope and Miall, 2012) and the reference electrode was placed on the ipsilateral buscinator muscle (Galea et al. (2009)). In a second group, participants were administered $1.5mA$ of anodal left motor cortical stimulation (M1 group), with the active electrode being placed over C3 (according to the EEG 10/20 standard system) and the reference electrode over the contralateral supraorbital area. In a third group, left dorsolateral pre–frontal cortex (dlPFC group) was stimulated by administering $1.5mA$ of anodal stimulation. In this case, the active electrode was placed over F3 (according to the EEG 10/20 standard system) and the reference electrode over the contralateral supraorbital area. Finally, 20 participants received pseudo stimulation (sham group); 8 of them were prepared for cerebellar stimulation, 6 for M1 stimulation and 6 for dlPFC stimulation. Table 5.1

summarises the stimulation parameters of the four experimental groups.

Table 5.1: Experimental Groups: Participants were randomly assigned in one of four experimental groups depending on the brain area stimulated and the kind of stimulation (real or pseudo)

Group	Active Electrode	Reference Electrode	Current Intensity	Stimulation Polarity
Cerebellar (right)	1cm below & 4cm across the inion	Left Buscinator Muscle	$2mA$	Cathodal
M1 (left)	C3	Right Supraorbital Area	$1.5mA$	Anodal
dlPFC (left)	F3	Right Supraorbital Area	$1.5mA$	Anodal
Sham	6 Cerebellar, 8 M1 8 dlPFC		110μA for15 ms every 550 ms	Pseudo

5.2.5 Design

Prior to stimulation participants had first to complete five tracking trials during which they tracked via the joystick a target moving in a pseudorandom trajectory (see Section 3.2.4) for details of a tracking task trial). In these trials participants had visual feedback of a cursor corresponding to the joystick movement. No visual feedback delays were imposed on the cursor. The reason for introducing these tracking trials was to familiarise the participants with the workspace and instruct them on following a reasonable speed in the exploration task. Following that, participants completed five Batches of the exploration task (*baseline phase* Phase).

Upon completion of the baseline exploration Batches, the DC–stimulator was turned on. Participants rested for the first 10 minutes of the stimulation and then they started performing the exploration task. For the next ten minutes they performed the exploration task while still receiving the stimulation. After that the stimulation was turned off and participants continued performing the task for another ten minutes. We will call the practise of the exploration task during and post

tDCS as *test phase*. The design is summarised in Table 5.2.

Table 5.2: Experimental Design. During the Baseline Phase, participants first completed 5 tracking trials and 5 baseline exploration Batches. After that, they received 20 minutes of stimulation. They performed the exploration task for 20 minutes starting in the middle of the stimulation period.

		tDCS		
	PRE	(10 min)	PERI (10 min)	POST (10 min)
Tracking	**Baseline Phase** Exploration task		**Test Phase** Exploration Task	

5.2.6 Data Analysis

Matlab R2011b was used to process the data recorded from the joystick. SPSS was used for running the statistical analysis.

The first dependent variable used to analyse the present dataset was the *Number of Batches* completed during the 20 minutes of exploration. Increase in the Number of Batches could be a result of either a decrease of Irrelevant Distance or of an increase in Speed.

Changes in *Speed* could be expected and could be related to several parameters. For example, stimulation of the motor areas (cerebellum and motor cortex) could drive participants to move faster, thus change kinematic parameters of the movement. On the other, hand, enhancing the working memory (via dlPFC stimulation) could result in better registering successful movements and repeating them fast afterwards. Generally, familiarisation to the task could lead to an increased speed (in all groups). So, to safely interpret any changes in the Number of Batches we first looked at the average Speed of participants. The average Speed was calculated by taking the ratio of the Irrelevant Distance (as defined below) to the time interval the distance was travelled. We expected to observe no difference during the baseline Batches as participants were encouraged to adopt a reasonable speed similar to that

of the tracking task.

Similar to Chapter 3, any improvements in exploration performance were expected to be revealed by shorter Irrelevant Distances between the start and target area. Again, the *Irrelevant Distance* was defined as the distance travelled between encountering the start area and entering the target area minus the euclidean distance between the two areas (from periphery to periphery).

Since we encouraged a continuous movement from the start to the target area and back, we also looked at the *Total Irrelevant Distance*, defined as the distance travelled between two sequential encounters of the start area minus double the euclidean distance between the two areas (from periphery to periphery).

Embedded in the metric of the Irrelevant Distances, there was a second non–contingent motor output; the output produced between entering the start area and the hand position 380ms later when the blue disk revealing the start area was presented. Since participants were given potential access about the information of the delays during this phase, we were interested in looking if participants learned to compensate for the effect of delays within stopping within the start area and waiting until the blue disk flashed. We calculate the *Hand–Start Distance*, defined as the distance travelled between entering the start area to the hand position 380 later minus the radius of the start area. Evidence of discontinuous way moving would be revealed by the analysis of the Hand–Start Distance and might suggest awareness of delay.

Finally, we looked at changes in the *Target Post–Discovery Motor Output*; the non–contingent output between reaching the target area and receiving the reinforcement flash. Again, we looked at this variable in order to investigate whether there were any signs of adaptation to delays or disruption of the continuous movement.

All but the Batches variable were log–10 transformed. The Number of Batches

dependent variable was analysed with a one–way ANOVA, with four levels of the independent variable of Group, one for each experimental group. The rest of the variables, unless otherwise stated in the results section were analysed with a 4x10 (or 9) x2 ANOVA, with one between–subject factor of Group (4 levels: Sham, dlPFC, Cerebellar and M1), one within–subject factor of Trials (with either 10 or 9 levels, depending on the variable) and one within–subject factor of Phase (Baseline or Test Phase). The degrees of freedom were corrected using Greenhouse–Geisser estimates of sphericity.

Also we should mention that in the figures presented in the next section, unless otherwise stated, the error bars show the standard error of the mean.

5.3 Results

5.3.1 General Observations

Figure 5.1 shows the behaviour of a typical participant in a Batch. The first trial of each Batch started with the participant trying to find the new start area (the small circle in the plots). This exploration phase is omitted from the plot of the first Trial, as it was not taken into account in the behavioural metrics. Once the start area was revealed to the participants they would go on to explore the workspace until they found the target area (big black circle). In the subplots of Figure 5.1 the black line shows the trace of the path travelled between reaching the start area and the target area. This is the path used for the calculation of the Irrelevant Distance from start to target. The beginning and the end of this trajectory is marked with black dots. However, the position of the of the hand (non delayed) at the presentation of the start area and the reinforcement flash of finding the target area were in reality 380ms ahead. The real positions of the hand at stimuli presentation are not presented in the graph. When the participant received the reinforcement flash on reaching the

target area she was then free to return to the start area. The orange line presents the entire path from reaching the start area in one trial and returning to it. This would trigger the next trial.

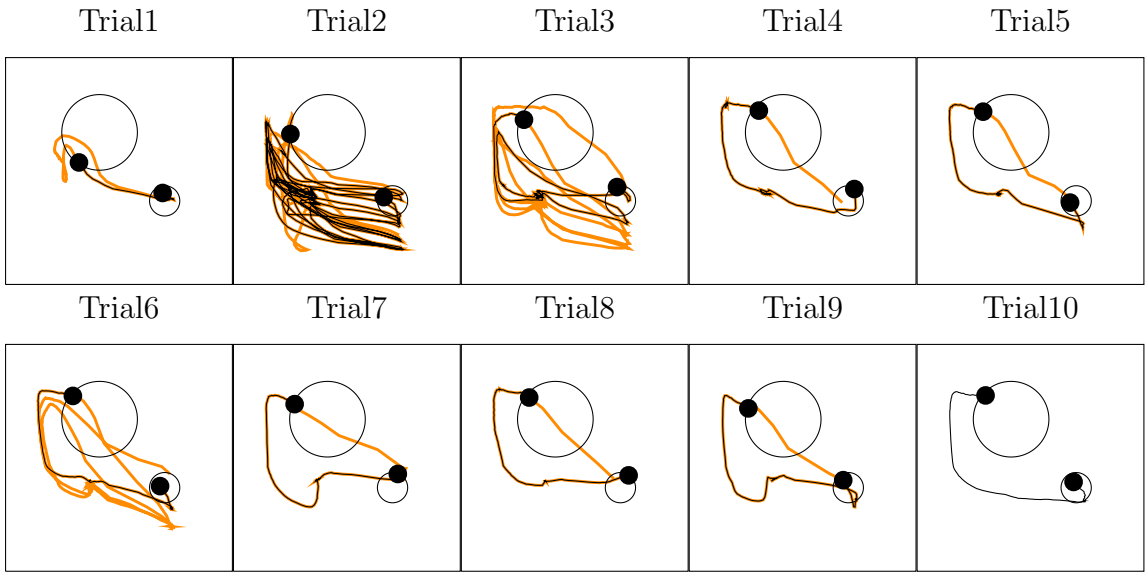

Figure 5.1: The *black line* shows the distance travelled between the Start Area (small black circle) and the Target Area (big black circle). The position of interception of the cursor with the start area and the target area is symbolised with small black filled circles. The *orange line* in each subplot shows the path from encountering the start area in the current trial till returning to it (which signals the beginning of the next trial). The Batch finishes when reaching to the Target for the 10^{th} time. So, for this trial there is no from Start to Start path. Over the trials the participant adopted a continuous movement that brought him through the starting and the target area. Note that the start and area sizes are plotted proportionally to its real sizes.

In the example presented, although the participant was quite lucky in finding the target area during the first trial, she still had to spend trial two and three in order to find a movement that brought her successfully to the target and all–in–all finding an action that was successful in terms of achieving the delivery of both stimuli. Moreover, the example of this participant was used to show the effect of both reinforcement delivery delay and start area presentation delay in action discovery. In trial two the participant attempts to go back where she was when the reinforcement

stimuli was delivered but she fails to receive the same stimulus. That makes her to start exploring the workspace again, which is depicted by a long black trace. Somewhat obscured by the black trace there is a long orange line which shows that the participant had to explore for some time as well to return to the starting point. The long orange line trace of trial three is produced as the participant missed the start area in first place and had to explore again to find it (or else to explore until the start area was presented). This behaviour was caused mainly as result of the small size of the start area.

To sum up, on top of the delay in the delivery of the reinforcement for reaching the target, the delayed presentation of the start area also induced a second source of non–contingency. Over the trials participants found a (continuous) movement that brought the delivery of both stimuli.

5.3.2 Number of Batches

The first variable we examined was the Number of Batches completed during the test phase of the study. Figure 5.2 shows that participants in the group that received M1 stimulation completed on average more Batches compared to all the other groups.

A one–way ANOVA revealed that the effect of the between–participants factor Group was significant ($F(3, 79) = 6.185$, $p = .001$). Tukey's HSD post–hoc test showed that this result was because of a significant difference between the following groups: M1 and Sham ($p = .003$), M1 and dlPFC ($p = .001$) and finally M1 and Cerebellar group ($p = .039$). No other significant multiple comparisons were observed.

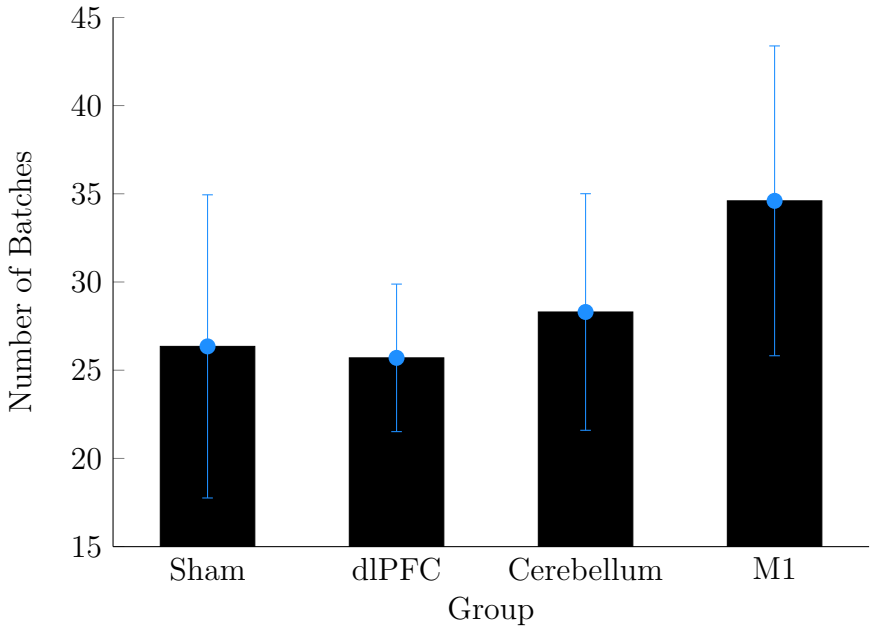

Figure 5.2: Number of Batches across Groups. The error bars show the standard deviation and not the standard error of the mean. Participants that received M1 stimulation completed on average a higher number of Batches.

5.3.3 Speed

Figure 5.3 shows the baseline Speed across the Groups. We observe a small but steady increase in the average Baseline Speed across the trials. Sham, Cerebellar and M1 Groups appear to have similar baseline Speeds but the dlPFC group starts on average about 20% slower. All the groups though show an increase of Speed across trials about 20% (from first to last trial). Figure 5.4 presents the average Speed across Trials and Groups during the test phase. All Groups seem to keep an almost steady average Speed across Trials (starting from Trial 2), which is higher compared to the average Speed during the last Trials of baseline phase. Similar to what we observed in the baseline Trials, the dlPFC Group on average adopted a lower Speed.

The 4x10x2 ANOVA showed that there was a significant effect of the following

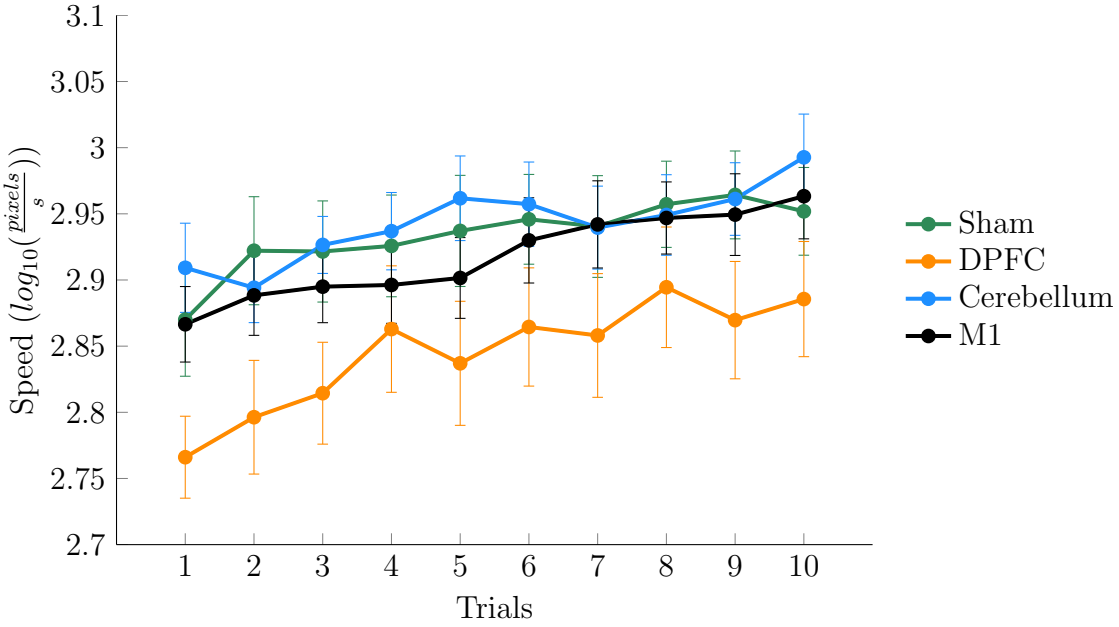

Figure 5.3: Average Speed during baseline phase exploration Batches across Groups. The average Speed increases over the Trials. The error bars show the standard error of the mean.

within–subject factors on Speed: Phase ($F(1, 76) = 28.054, p < .001, \eta^2 = .270$), Trials ($F(5.24, 398.6) = 25.481, p < .001, \eta^2 = .251$) and the interaction of Phase*Trials ($F(6.27, 477.2) = 10.649,\ p < .001,\ \eta^2 = .123$). A simple main effect analysis (not presented here in details) showed that this interaction was because of a significant difference when comparing each Trial during the baseline and test phase but also because of differences among Trials within each phase of the study. This interaction also exposes the difference in the rate of change of trials' speed within each phase.

Although no Phase*Group interaction was observed ($F(3, 76) = .399, p = .754, \eta^2 = .015$), we searched if the Phase effect was found in all Groups. A simple main effect analysis of the effect of Phase (tDCS) on each Group revealed that all Groups but dlPFC significantly increased their speed from Baseline to Test Phase: Sham ($p = .008$), dlPFC ($p = .068$), Cerebellar ($p = .001$), M1($p = .01$). All these p–values

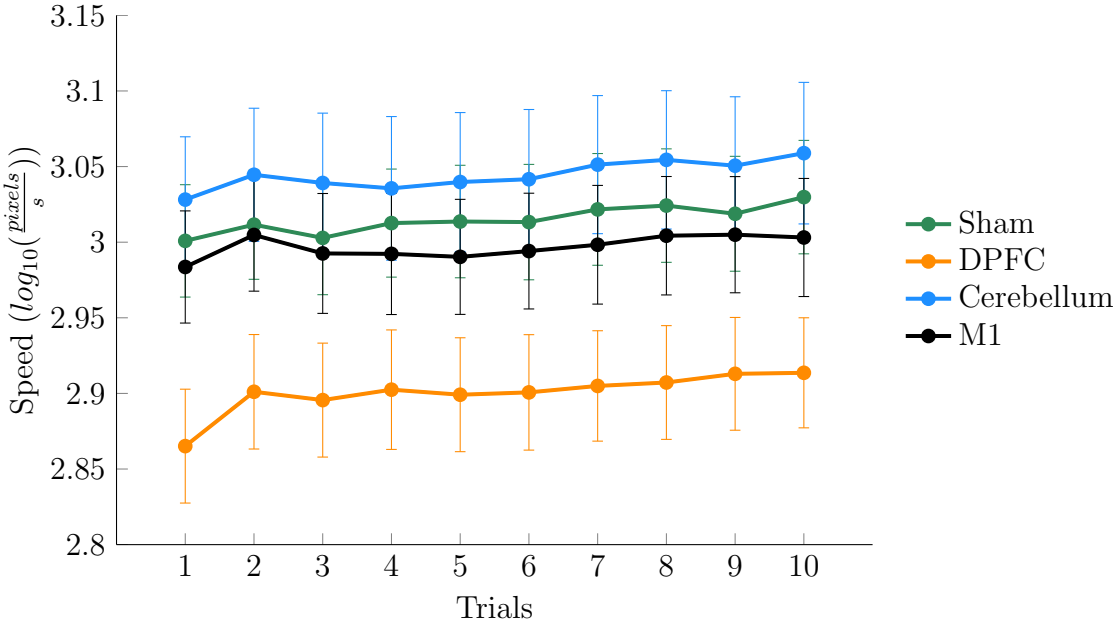

Figure 5.4: Average Speed during test (during and post tDCS) phase exploration Batches across Groups. The average Speed remains steady from Trial 2 on. Average Speeds during the test phase are increased compared to baseline phase. Compare this figure to previous one and note the slight difference in the values of the y–axis. The error bars show the standard error of the mean.

were adjusted for multiple comparisons (Sidak).

No significant effects on the Speed were observed as a result of the Trials*Group interaction ($F(15.73, 398.6) = 1.128, p = .299, \eta^2 = .043$) and the Phase*Trials*Group interaction ($F(18.83, 477.2) = 1.165$, $p = .259$, $\eta^2 = .044$). Finally, the between–subject factor of Group did not have a significant effect ($F(3, 76) = 2.55$, $p = .062$, $\eta^2 = .092$) on the Speed.

In summary, Speed increased in all groups, but there was no evidence of a tDCS effect either between groups or between Phases.

5.3.4 Irrelevant Distance from Start to Target

All four experimental Groups showed similar behaviour in terms of average start to target Irrelevant Distance across Trials in the first five baseline Batches, as shown

in Figure 5.5. The data in the figure are in real space units (pixels). Participants appear to be able to find a shorter pathway from start to target area over the 10 trials given at a Batch.

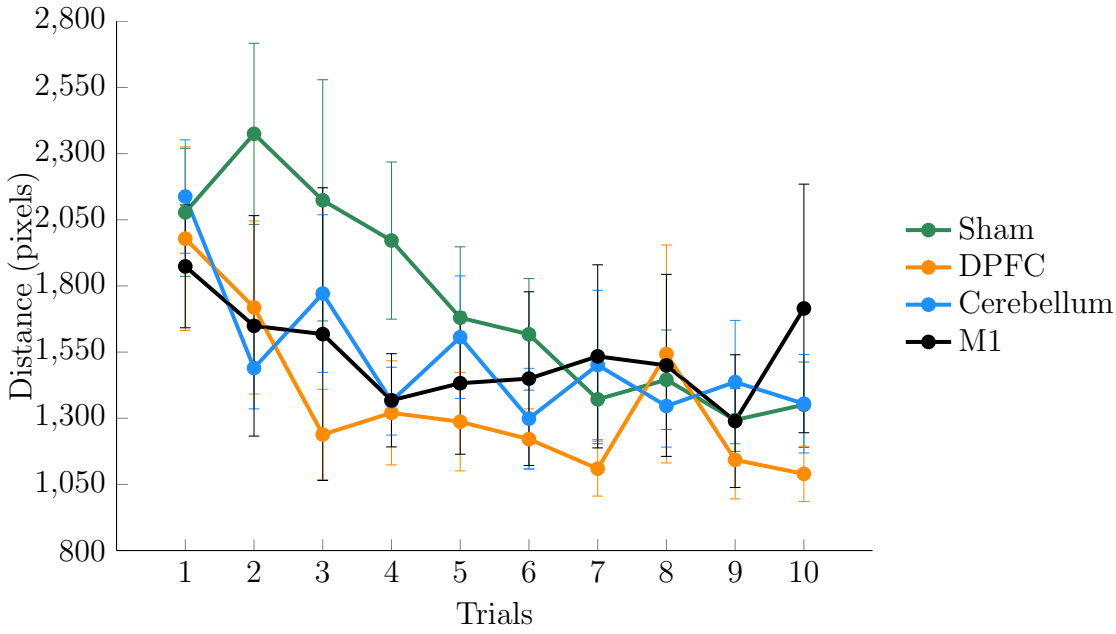

Figure 5.5: Average Irrelevant Distance from Start to Target Area during the Baseline Phase. All groups showed similar behaviour. The data in this graph are real size (pixel) data. The error bars show the standard error of the mean.

The behaviour during the test phase is shown in Figure 5.6.Note that in this figure we plotted the transformed (log_{10}) values of Irrelevant Distances. Here, we observe a differentiation among the Groups. The most pronounced effects appear to be those between Sham and M1 and between Sham and dlPFC Groups. Also, all Groups have a smoother learning curve in this phase compared to Baseline. This is attributed to the fact that the average Irrelevant Distance in each Trial is the result of taking the mean of more Batches per participant (compared to only five Batches during the baseline phase). Refer to Section 5.3.2 for details on the average numbers on Batches per Group during the test phase.

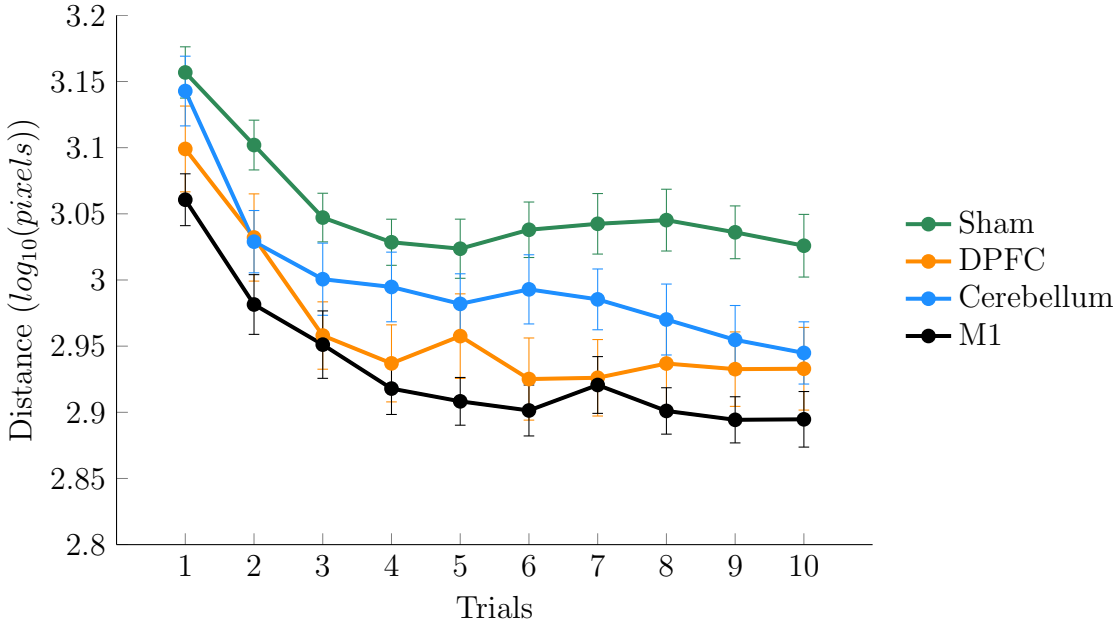

Figure 5.6: Average Irrelevant Distance from Start to Target Area during the Test Phase. The four Groups appear to have developed differences in behaviour. The error bars show the standard error of the mean.

The 4x10x2 ANOVA showed that there was a significant effect of Trials on the Irrelevant Distance ($F(6.33, 481.7) = 33.135$, $p < .001$, $\eta^2 = .304$). On the other hand, the within–subject factor of Phase had no effect on the dependent variable ($F(1, 76) = .707$, $p = .403$, $\eta^2 = .009$). In addition, none of the interactions between the various factors was significant: Phase*Group ($F(3, 76) = .771$, $p = .514$, $\eta^2 = .030$), Trials*Group ($F(19.015, 481.7) = .831$, $p = .670$, $\eta^2 = .032$), Phase*Trials ($F(7.15, 543.4) = .501$, $p = .874$, $\eta^2 = .007$) and Phase*Trials*Group ($F(21.45, 543.4) = .842$, $p = .697$, $\eta^2 = .032$).

Finally, and most importantly, there was a significant effect of the between–subject factor of Group ($F(3, 76) = 4.59$, $p = .005$, $\eta^2 = .153$). Tukey's HSD post–hoc test reveal that this significant effect was because of a significant difference between Sham and dlPFC Groups ($p = .015$) and between Sham and M1 Groups

($p = .007$). There was no difference between the Cerebellar and any of the other Group (dlPFC: $p = .787$, M1: $p = .638$, Sham: $p = .148$). Moreover, there was no significant difference between the dlPFC and M1 Groups ($p = .994$).

5.3.5 Irrelevant Distance from Start to Start

Figures 5.7 and 5.8 show the performance of participants in terms of Irrelevant Distance from Start to Start during Baseline and Test Phase, accordingly. Similarly to the Irrelevant Distance between Start and Target (Figure 5.5) there are no Group differences during the Baseline Phase. However, during the Test Phase the differences among the Groups have a different pattern compared to that of Irrelevant Distances between the Start and the Target (Figure 5.6). Sham and Cerebellar Groups seem to have a similar behaviour, which is different from the behaviour of the dlPFC and M1 Groups.

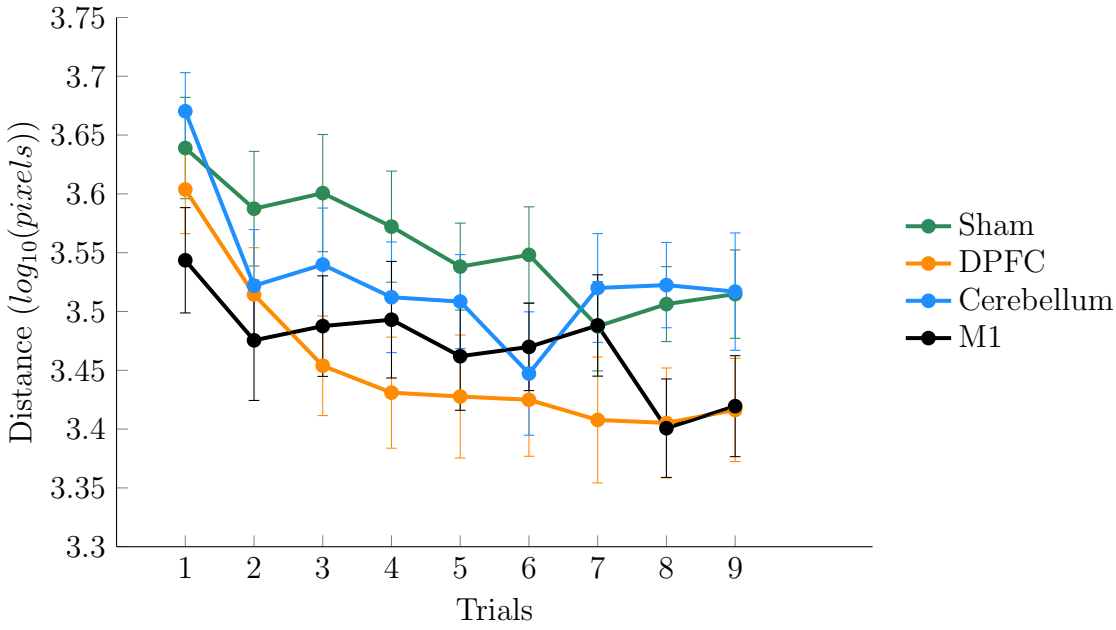

Figure 5.7: Average Irrelevant Distance from Start to Start Area during the Baseline Phase. All groups showed similar behaviour. The error bars show the standard error of the mean.

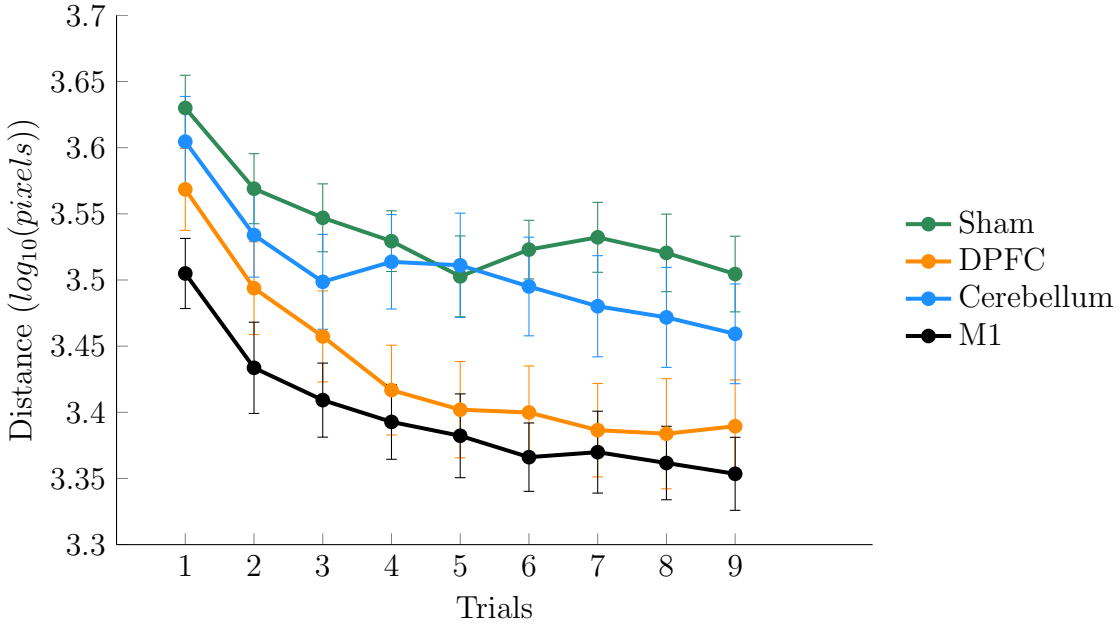

Figure 5.8: Average Irrelevant Distance from Start to Start Area during the Test Phase. The four Groups appear to have developed differences in behaviour. The error bars show the standard error of the mean.

The result of the 4x9x2 ANOVA showed that there was a significant effect of Phase ($F(1, 76) = 5.540$ $p = .021$, $\eta^2 = .068$) and of Trials ($F(5.808, 441.4) = 37.085$, $p < .001$, $\eta^2 = .328$) on the Start to Start Irrelevant Distance. However, none of the interactions between the various factors was significant: Phase* Group ($F(3, 76) = .986$, $p = .404$, $\eta^2 = .037$), Trials*Group ($F(17.42, 441.4) = 1.238$, $p = .229$, $\eta^2 = .047$), Phase*Trials ($F(6.56, 498.8) = .309$, $p = .943$, $\eta^2 = .004$) and Phase*Trials*Group ($F(19.69, 498.8) = 1.143$, $p = .302$, $\eta^2 = .043$).

Although there was not a significant Phase*Group we ran a simple main effects analysis to look if there was any group specific Phase effect. We found that the only group that had a significant phase effect was M1 ($p = .01$, Sidak–adjusted for multiple comparisons). The other groups showed no such an effect (Sham: $p = .590$, Cerebellar: $p = .449$, dlPFC: $p = .458$).

Finally, again importantly, there was a significant effect of the between–subject factor of Group ($F(3, 76) = 3.485$, $p = .020$, $\eta^2 = .121$). Tukey's HSD post–hoc test revealed that this significant effect was because of a significant difference between Sham and M1 Groups ($p = .043$) There was no difference between the Cerebellar and any of the other Groups (dlPFC: $p = .295$, M1: $p = .199$, Sham: $p = .900$). Moreover, there was no significant difference between the dlPFC and M1 ($p = .996$) and between Sham and dlPFC Groups ($p = .074$).

5.3.6 Distance between Hand position and Start Area upon Start Area presentation

The distance travelled by the hand between reaching the Start Area and Start Area presentation in Baseline and Test phases is shown in Figures 5.9 and 5.10. No differences are observed both during Baseline and Test Phases in terms of average Distance across Trials. Moreover, there seems to be only a small increase on average between the two Phases. The behaviour across the Groups appears to be similar. Note that the values in the graphs are in real size units. Moreover, we have subtracted the start area radius from our data presented in the figures. So, it is clear that participants were always outside the start area when the blue disk showing the start area was presented to them.

The 4x10x2 ANOVA showed that there was no significant effect of Trials on the Distance ($F(6.95, 528.8) = 1.26$, $p = .265$, $\eta^2 = .016$). For this reason we collapsed the dataset across Trials and ran a simpler 4x2 ANOVA with a between–subject factor of four Groups and a within–subject factor of two levels of tDCS (Baseline and Test Phase).

The results of the 4x2 ANOVA showed that there was no significant effect of the within subject factor of Phase ($F(1, 76) = .022$, $p = .883$, $\eta^2 < .001$) and no

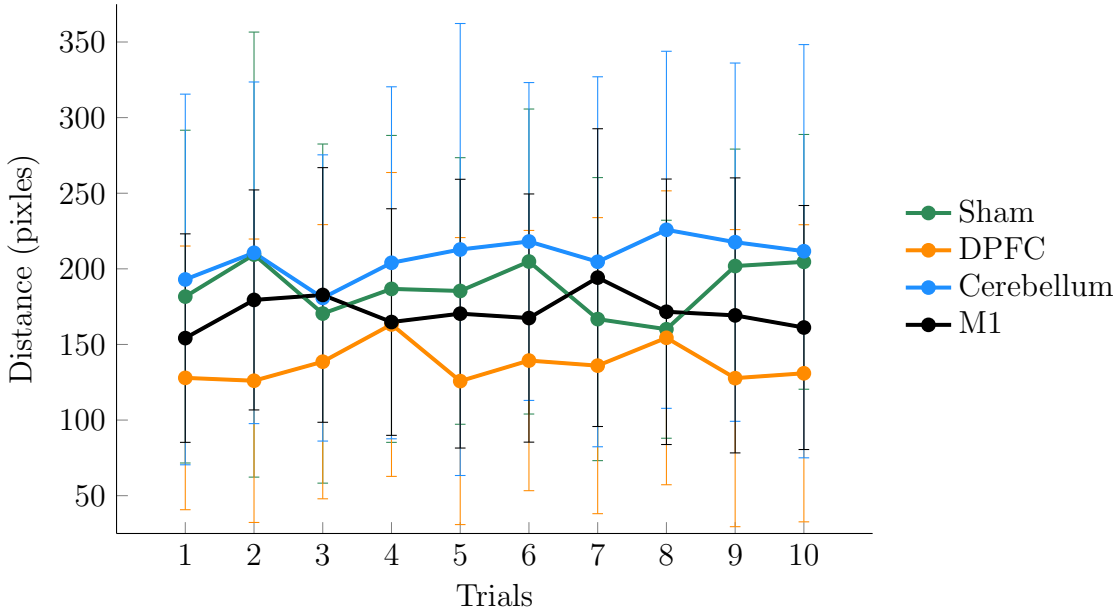

Figure 5.9: Distance between Hand position and Start Area upon Start Area presentation during Baseline Phase. There are no significant differences across Trials. All four Groups seem to have a similar behaviour. The error bars present the standard deviation of the mean.

Phase*Group interaction ($F(3, 76) = .166$, $p = .919$, $\eta^2 = .007$). Moreover, there was no significant Group effect ($F(3, 76) = 2.389\, p = .075$, $\eta^2 = .086$).

5.3.7 Target Post–Discovery Motor Output

Figures 5.11 and 5.12 show the Target Post–Discovery motor output across Trials and among Groups during the baseline and test phase accordingly. When comparing these figures to Figures 5.3 and 5.4, it becomes clear that the target post–discovery motor output is driven by the average group speeds.

Similar to the Speed analysis, the 4x10x2 ANOVA showed that there was a significant effect of Phase ($F(1, 76) = 15.450$, $p < .001$, $\eta^2 = .169$), of Trials ($F(6.86, 521.9) = 4.844$, $p < .001$, $\eta^2 = .060$) and of the Phase*Trials interaction ($F(7.01, 533.1) = 4.641$, $p < .001$, $\eta^2 = .058$) on the Post–Discovery Motor Output.

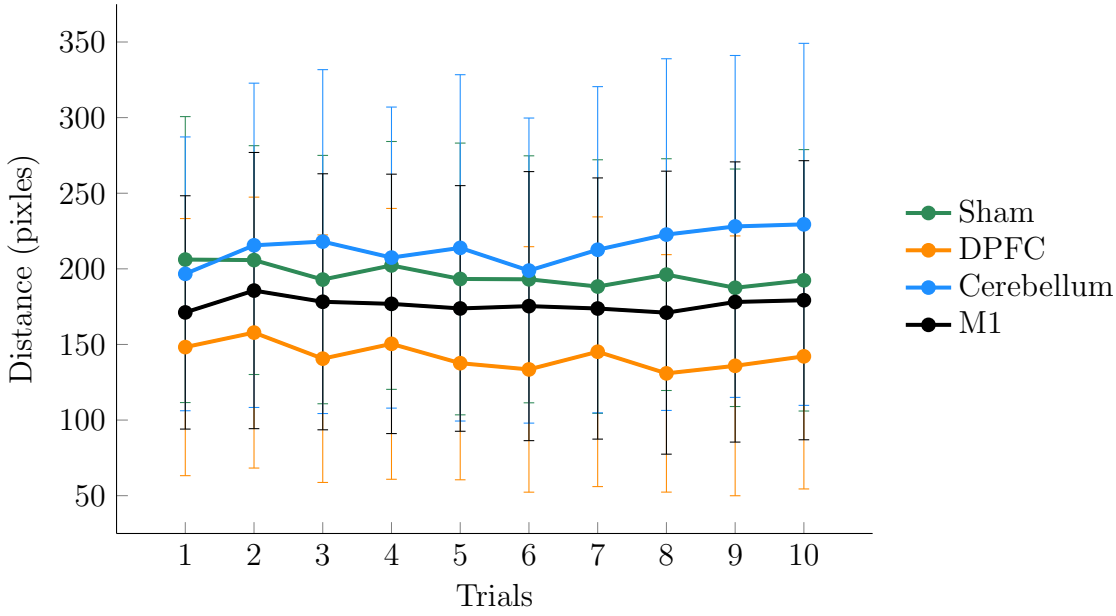

Figure 5.10: Distance between Hand position and Start Area upon Start Area presentation during Test Phase. There are no significant differences across Trials. All four Groups seem to have a similar behaviour. The error bars present the standard deviation of the mean.

The effect of Trials and of the Phase*Trials interaction is a result of the gradual increase of distances during the Baseline Phase and the difference in the average distances between the Baseline Phase and the Test Phase (simple main effect analysis not presented here).

We also ran a simple main effect analysis to look of any group–specific effects before and during/after the tDCS. Although no Phase*Group interaction was observed ($F(3, 76) = .872$, $p = .459$, $\eta^2 = .033$), the simple main effect analysis showed that there was a significant effect of Phase in the Post–Discovery Distance for the Cerebellar ($p = .006$) and Sham ($p = .012$) Groups but not for the dlPFC ($p = .428$) and M1($p = .104$) Groups. All these p–values were adjusted for multiple comparisons Sidak.

No other significant effects in the Post–Discovery Motor Output were observed as

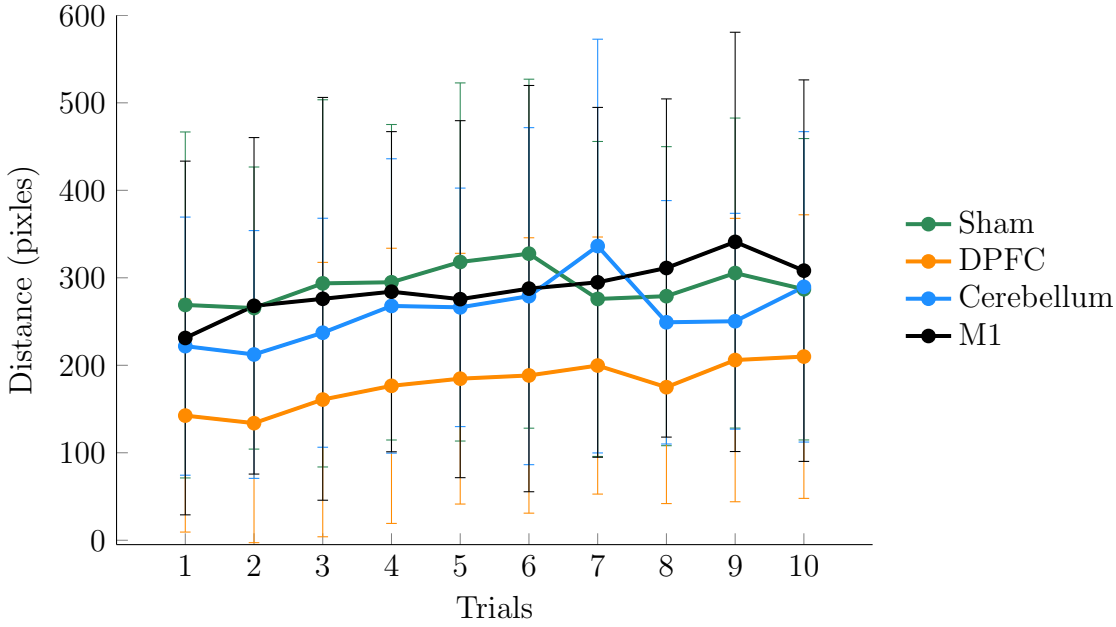

Figure 5.11: Target Post–Discovery Distance during Baseline Phase. Participants produced more non–contingent over the Trials but there are no obvious between Group differences. The error bars present the standard deviation of the mean.

a result of the Trials*Group interaction ($F(20.60, 521.9) = .889$, $p = .605$, $\eta^2 = .034$) or the Phase*Trials*Group interaction ($F(21.04, 533.1) = 1.319$, $p = .155$, $\eta^2 = .049$). Finally, the between–subject factor of Group did not have a significant effect ($F(3, 76) = 2.356$, $p = .078$, $\eta^2 = .085$) on the dependent variable.

5.4 Discussion

In the present study we aimed to explore the involvement of M1, dlPFC and cerebellum in the exploration task presented in Chapter 3 and as a consequence their role in action acquisition. The results showed that anodal M1 stimulation resulted in an increased number of Batches completed during and post tDCS stimulation. Speed differences among the groups could not account for this result. Moreover, the M1 group was the only group that showed changes in Total Irrelevant Distance from

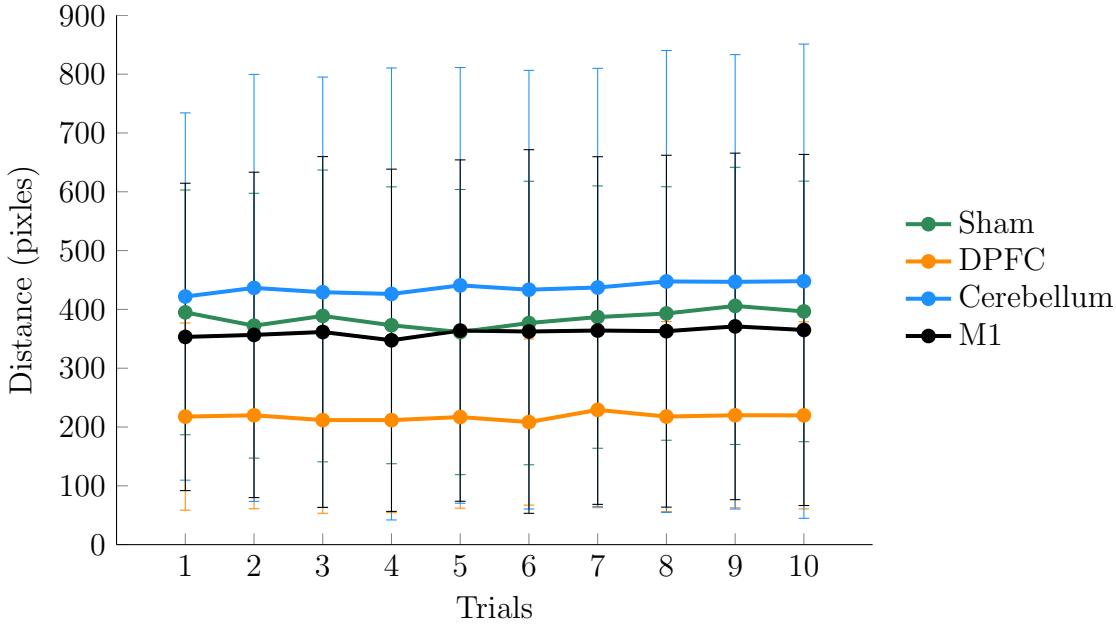

Figure 5.12: Target Post–Discovery Distance during the Test Phase. No changes over the Trials and no significant between Group differences are observed. The error bars present the standard deviation of the mean.

baseline to test phase.

Similar Speeds Across Groups

High or low speeds in tasks with delays can augment or diminish, accordingly, the effect of delay and introduce unwanted between subjects differences (Miall et al., 1986). This is why, before starting the exploration task participants performed five tracking trials to familiarise with the workspace and to adopt a reasonable exploration speed, avoiding very fast or very slow movements. However, we should note that we did not select tracking speeds taking into consideration the delays involved in the exploration task. We just aimed to ensure that participants would tend to explore with similar speeds. This is why we were not interested in analysing the tracking data.

The results of the exploration data showed that there were no significant group

differences in speed. All four groups adopted a faster speed during the test phase compared to baseline. This change is most probably attributed to familiarisation to the task. However, the dlPFC groups appeared to have a slower average speed, both during baseline and test phases. Finally, stimulation did not induce any significant effects in any of the groups.

Effect of tDCS on the Irrelevant Distances

Irrelevant Distance from Start to Target Area. Similarly to the analysis in Chapter 3, we were interested in looking at the Irrelevant Distance between the start and the target areas. Overall, participants managed to perform (on average) movements with less Irrelevant Distance across the 10 trials of a Batch. No baseline differences were observed among the groups. Although we observed group differences when plotting the data, the statistical analysis did not reveal any significant effect of phase (tDCS). We attributed this result in lack of power and large individual differences in this unconstrained exploration task.

Irrelevant Distance from Start to Start. In Chapter 3, participants made discrete movements to the target area. Once they received the reinforcement flash, the start point appeared and they had to go back to it while being further assisted by the robotic manipulandum. In this Chapter, the trials were not discrete. Once participants received the reinforcement flash they were free to go back to the start area that only appeared once the hand reached the area. Even then the start area was presented to the participants with 380ms delay. This set–up encouraged a continuous movement from start to target and back. So, we were interested not only in the Irrelevant Distance from start to target but actually from start to start.

The analysis of the start–to–start Irrelevant Distance revealed that participants were actually learning a continuous movement that brought them to pass thought

the start and target areas and receive the delayed stimuli (presentation of the target area and reinforcement flash). Comparing the four groups we found that there was a significant effect of Phase on the start–to–start Irrelevant distance only in the M1 group. Moreover, the group effect was driven by the difference between the Sham and M1 groups. The fact that the dlPFC group was not significantly different from the Sham group, as might have been expected by observing Figure 5.8, indicates that the behaviour of this group might have been a result of the slower average group speed (although not significant). Generally, the fact that the dlPFC and the cerebellar groups did not show any significant differences compared to the Sham or the M1 group could imply that there have been some effects of tDCS on these groups that either were not strong enough to reach significance or could not be exposed by the highly variable Irrelevant Distance metric.

tDCS effect on the Number of Batches

Based on the Speed and Irrelevant Distance results we can conclude that the Number of Batches effect is driven (partly at least) by shorter Irrelevant Distances and not by between speed differences on a group level. However, there must be other factors affecting this result. This is because if the Batch effect was driven only by the Irrelevant Distances then we should expect to see no significant Batch effect between the M1 and Cerebellar and between M1 and dlPFC groups, like in the case of the Irrelevant Distances.

Delayed presentation of the start area and delayed reinforcement

The delayed presentation of the start area provided potential information about the delay, if participants were able to assess the start position (across several trials) independently of the start circle "flash" onset. We looked at the Distance travelled by the hand from the moment the hand crossed the start area until the participants

were presented the actual location of the start area. Participants showed no signs of stopping on the start area and waiting for its presentation, although the spatial information about the start area made them quite consistent across the trials. However, the behaviour of this variable does not seem to correlate to the speed behaviour.

On the other hand, the non–contingent output from crossing the target area till the reinforcement delivery was driven apparently by the speed. Again, there is no evidence that participants integrated the information of the delay in a different way across the groups.

All in all, our observations on the delayed presentation of stimuli show that tDCS did not have any effect on the way participants processed this information, in any of the groups.

Differential tDCS effects over stimulated brain areas

It is clear from the results presented that anodal tDCS over M1 lead to a significant effect in the exploration task. It was discussed in Chapter 2 that the effects of tDCS can be related either to increased excitability and/or to interference with homeostatic plasticity mechanisms. However, it would be very difficult to tightly bind any of these theories to the behaviour observed. What we can speculate is that M1 stimulation must have affected the cortico–basal–ganglia connections. A proposed mechanism in line with the theory of Redgrave et al. (2008) is that M1 stimulation could have biased action selection by introducing synaptic weight changes between the motor cortex and the striatum or by influencing the "time stamp" or the eligibility traces. By enhancing the cortico–striatal system, it is therefore possible that the eligibility of each exploratory action is more easily assessed, and therefore a more efficient action was reached more quickly. Again these are only hypotheses that can

be further pursued in future research. Finally, M1 did not appear to affect kinematic parameters of the movements (e.g. speed).

The fact that we failed to observe any cerebellar and dlPFC tDCS effects might indicate that these two areas are not heavily involved in the exploration task. This could imply that the exploration task does not have a strong declarative component such that anodal dLPFC or cathodal cerebellar stimulation could have an effect on cognitive counterparts of the task. Stimulation of these two areas has been shown to drive effects only if the tasks are demanding enough (Pope and Miall, 2012). This observation can add evidence to the fact that what is learnt in the exploration task is movement trajectory and not spatial characteristics of the space (Thirkettle et al., 2013) and also support the idea that the exploration task is a suitable task for studying reinforcement learning. We hypothesised that stimulation of dlPFC might result in greater superstitious learning effects. Our analysis is not well designed to detect this, however, an additional spatial analysis (for example with Procrustes analysis) might be needed.

Perhaps we could have expected that cathodal cerebellar tDCS could have lead to similar effects to anodal M1 stimulation via the cerebello–cortical pathway. Based on our result, we can conclude that most probably cerebellar tDCS did not have the power to drive the same effect compared to direct M1 stimulation.

Finally, we can exclude the hypothesis that the tDCS did not have any effect on cerebellum or dlPFC. There are several studies to have shown cerebellar effects in motor tasks (Galea et al., 2011; Jayaram et al., 2012) and, dlPFC and cerebellar effects in cognitive tasks (Pope and Miall, 2012). Moreover, the number of participants was comparable to previous studies. Therefore, while we would expect to have modulated those areas, there was no behavioural effect.

In the analysis used in this chapter we chose not to differentiate between during

and post stimulation effects. The main reason for that was that we wanted to look at the overall effect of tDCS over each brain area. However, we could pursuit such analysis in the future. Although, integrating an extra time component in the analysis could be complicated to interpret, it could reveal possible interactions between non–declarative and declarative components during the time course on learning new actions.

Conclusions

Taken together, these results suggest that only stimulation over M1 affected action acquisition, possibly by modulating cortico–basal–ganglia connections. This could provide evidence supporting the crucial involvement of low–level process rather than cognitive ones to the initial phase of action acquisition.

Chapter 6
GENERAL DISCUSSION

"**What is the involvement of cerebellum in processes of motor learning beyond error–based learning, which is traditionally linked to the structure?**"

This has been the common ground question of the studies presented in my PhD thesis. In this General Discussion chapter, I would like to summarise the main findings and conclusions of the studies. In addition, I will present several points to consider when designing an experiment that involves transcranial Direct Current Stimulation (tDCS). Finally, I will briefly mention the limitations imposed by the analysis tools when studying about learning.

In Chapter 2, we asked what is the role of cerebellum in motor skill learning. A keypress task was used that involved learning coordinated finger movements, similar to the movements one need to learn when practising the piano. Anodal and cathodal tDCS was applied over the cerebellum and after several minutes of stimulation participants started practising the task. The hypothesis was that tDCS will differentially affect learning in each stimulation group. The results of this study were inconclusive both because of insufficient understanding of the effects of tDCS on brain activity and mechanisms, but also because of certain design flaws as discussed in Section 2.4. However, the results of this study, along with a couple of other tDCS studies that I ran during my PhD (not presented in this thesis), might suggest that cerebellar tDCS could be affecting motor variance.

In Chapter 3, we were interested in investigating the role of cerebellum in reinforcement learning and particularly in novel action acquisition. The goal of this chapter was to work towards a paradigm that involved both the basal ganglia and

the cerebellum, and that could be perhaps translated into an imaging study. We combined the exploration task of Stafford et al. (2012) with a visuomotor tracking task, based on the paradigm used by (Foulkes and Miall, 2000). Delays in the delivery of the reinforcement signal during the exploration task lead to acquisition of actions that are biased by the motor output non–contingent to the reinforcement (Walton, 2011; Stafford et al., 2012). This is the motor output produced between reaching to a target and receiving the reinforcement signal. We hypothesised that if participants are exposed to the same delay during a visuomotor tracking task and adapt to the delays, they might then be able to acquire movements that are less contaminated by the non–contingent motor output.

Based on the results of Chapter 3, we can reach the following conclusions. Firstly, participants were not able to attribute the delay as a property of the vBots and use it to perform better in the exploration task. Attributing the delay to the vBots has in a sense two dimensions. The first is a cognitive one; participants become aware of the delay, attribute it to the vBots and then apply this information to the exploration task. In this case, they infer the target location and stop to it until delivery of reinforcement. There were no signs of this sort of behaviour. Based on what participants that tracked with delays reported during a debriefing session, they were not able to make this connection between the two tasks.

The second dimension of "attributing the delay to the vBots" is equivalent to low-level motor adaptation. Here we expected that adaptation to the delays could influence the decision about which part of the movement was related to the delivery of the reinforcement. This behaviour would be depicted in shorter Irrelevant Distances performed by the group that tracked with delays. Although participants adapted to the tracking delays they did not show any signs confirming our hypothesis. Contrary to that, the results of this study showed that the subjects that practised the tracking

task with delays tended to perform worse in terms of Irrelevant Distance compared to the group that practised the tracking task without delays. Further analysis indicated that this might be because of increased proprioceptive uncertainty experienced by the group that tracked with delays. So, the second conclusion that was drawn by this study was that adaptation to tracking delays might have interacted in a way that we had not accounted for in the first place.

This observation lead to the study presented in Chapter 4. Here, we asked if there are any changes in proprioceptive uncertainty occurring during the tracking task that could explain the result in the exploration and tracking task study. We assessed proprioceptive variance before and after four different conditions of the tracking task. The results in this study indicated that tracking without delays and with variable speeds increases proprioceptive accuracy. On the other hand, adaptation to the delays does not lead to a similar effect. These observations bring forward the role of variance during learning.

In the final study presented in this thesis (Chapter 5), we returned to the question of the involvement of cerebellum in action acquisition and extended the question to other brain areas, as well. We used tDCS to perturb the motor cortex (M1), dorsolateral prefrontal cortex (dlPFC) and cerebellum during the exploration task. The results showed that only anodal M1 stimulation affected action acquisition, in terms of lower Irrelevant Distances. This finding suggests that action acquisition might be primarily driven by low–level mechanisms and not by cognitive ones, as dlPFC stimulation did not bring any significant effect. It is more difficult to explain why cerebellar stimulation did not affect task performance. It could be that the cerebellum simply is not involved heavily in the task. Another possible explanation is that the cerebellum could be keeping the balance between cognitive and motor functions. So, tDCS over the cerebellum would be equivalent to just shifting a baseline

brain state. Currently, there are no experimental evidence supporting this idea but it would be an interesting question to pursuit the future.

To sum up, the results presented in this thesis might indicate that the cerebellum is not only responsible for providing state estimations for the motor apparatus but also for providing the level of certainty or uncertainty related to the estimates. This is supported by the marginally significant increase in the mean and standard deviation of the intra–tap interval in the cerebellar stimulation groups during the keypress study in Chapter 2. The same conclusion could be inferred implicitly based on the behavioural data of Chapter 4, given the involvement of the cerebellum the tracking task. Increasing uncertainty about our own body state can lead to more variance in the movements and this is necessary during any kind of motor learning. Increased variance will allow for more exploration and exploitation of new motor patterns. So, again it would be interesting to re–analyse the data in Chapter 5 in a way that could expose greater within–batch variability in the movements and then compare across groups to investigate any potential tDCS effects. In the same framework, managing the contribution of cognitive and motor resources, perhaps based on the level of certainty attributed to state estimates, can further influence learning. All–in–all, the results presented bring forward the important role of variance during learning and raise the question about whether the conventional analysis techniques are appropriate to expose aspects of learning.

Design consideration when using tDCS

A big proportion of my PhD time has been spent in exploring tDCS and I would like to summarise the knowledge gained during using this neuromodulating technique, especially in motor learning studies.

In the study in Chapter 2, we did not have any *baseline measure* of performance.

We reached this decision because we wanted to avoid exposing the participants to the task, since our aim was to study the effect of tDCS in a task where participants would learn a task having no prior experience. However, between groups differences at the first block can not be solely attributed to tDCS. In this case, another measure could have been used, like for example practising another finger sequence. However, this solution could be problematic as a person might have more or less difficulty in performing different sequences. Nonetheless, having a baseline measure of performance is vital to accurately interpret the results. So, in order to be able to make a consistent comparison before and after the tDCS manipulation, the same baseline measure of performance is perhaps necessary even at the cost of exposing the participants to the task.

Learning a motor skill takes days, months or even year to be learnt. That means that a single session studies cannot give a sufficient degree of training and might not be enough to reveal group differences. *Multiple day sessions* could be preferable and essential to reveal the effect of tDCS over a brain area (see for example Reis et al., 2009).

In Pope and Miall (2012), cerebellar tDCS was effective only in the task that was more demanding (although the tasks were cognitive). Moreover, as discussed in Section 2.4, afferent signals during a motor task impose changes to cortical excitability (Nitsche et al., 2007) that are potentially different form task to task and lead to different susceptibility of each area to tDCS. These observations outline the importance of considering *task difficulty* and the engagement of an area in a task when designing a study and interpreting the results.

When to apply tDCS (whether to apply it before or while practising a task) is also important. A better understanding of how tDCS interacts with brain activity and cellular mechanisms must be gained in order to be able to predict when to apply

tDCS in order to observe maximum differences is behaviour and perhaps when to apply tDCS to increase the beneficial effects of tDCS in case of a treatment.

Another important factor to consider is the susceptibility of individuals to stimulation due to *genetic factors* (Paulus, 2011). If no information of the underlying individual genetics is provided, testing a large *number of participants* could provide more power to expose behavioural differences. In any case, when studying about motor learning, a large number of participants might be necessary to overcome the difficulty of variance inherent to learning processes.

The *size* of electrodes and the *position* of both the stimulating and reference electrodes are also crucial parameters as they influence the focality of the tDCS effect. Last but not least, the intensity of the stimulation os one of the most important parameters that will determine the effect of the stimulation. Usually, the higher intensity, within the safety limits is chosen, to achieve a maximum effect.

To sum up, there is a number of considerations to be taken into account when designing a tDCS study. However, it is difficult to predict with safety what is the optimal parameter for each of the aforementioned factors mainly because of insufficient knowledge about tDCS mechanisms.

Dealing with variance in motor learning paradigms

As mentioned in Chapter 1, CNS motor control on the musculoskeletal system is challenged by factors like neuronal noise, delays, fatigue etc. Another issue, that was not discussed earlier, is *motor redundancy* (Bernstein, 1967). For example, a point in space can be reached with different movements, using different number of joint rotations and recruiting different muscle groups. Optimal control offers a framework to deal with motor redundancy "by choosing the best possible control law" (Todorov, 2004). An alternative view considers motor redundancy in terms of

motor abundance, where variability is prerequisite for adaptive behaviour (Latash, 2012a). In these terms, there is not a single optimal solution but movements can be both variable and optimal (Latash, 2012b). Optimal control is often challenged as being not biologically plausible and instead a "good–enough control" should be considered that allows for local rather than global optimal minima to be adopted (Loeb, 2012; de Rugy et al., 2012). Irrespectively of the motor control framework used to explain motor variance, it is indisputable that motor variability drives motor learning and it can actually predict motor learning ability (Wu et al., 2014).

All the studies presented in my PhD dealt with motor learning and a final point that I would like to make is the issue of dealing (in terms of analysis) with variance in motor learning paradigms. A large proportion of participants in Chapter 2 was excluded because they showed no learning curve. In the tracking task (Chapter 3 and 4), there were also participants that showed big variance from trial to trial and no improvement in the performance (at least in terms or cursor–target error). Similarly, in the exploration task (Chapter 3 and 5) there were batches where the participants showed big fluctuations from trial–to-trial.

Undisputedly, part of this behaviour can be explained by factors like boredom, tiredness and lack of attention. However, a considerable proportion of this variant behaviour must be an effect of learning. Thus, considering a single index of performance and averaging across trials and batches can conceal a lot of interesting aspects of learning behaviour.

For example, in case of the tracking task we did not find a correlation between the improvement in mean tracking error (averaged over a time window or a few trials) and the increase or decrease in variance in the centre–out movements (this analysis was not presented). However, it could be that there is a correlation between the variance across the tracking trials and proprioceptive certainty in the centre–out

paradigm.

The exploration task is uniquely open to individual difference in behaviour and in learning ability. A lot of important information is lost just by looking the Irrelevant Distance and not the shape of the trajectories and another analysis like the Procrustes analysis (Walton, 2011), might be a more appropriate tool to investigate behaviour in the exploration paradigm. Moreover, perhaps one should consider not transforming the positively skewed data distribution but identifying the different components of the distribution (a gaussian distribution and a power tail law distribution), which could give a better insight about the exploration–exploitation behaviour (Volchenkov et al., 2013).

To sum up, important information about learning might be lost when analysing the data during motor learning tasks with simple performance indices. Alternative analysis methods can be considered in the future that will "explore and exploit" behaviour that are hidden in the data.

Appendix A
Appendix of Chapter 3

A.1 Exploration Pilot Study

The purpose of the pilot study was two–fold. At first, we wanted to define an appropriate Target Area Size (TAS). Moreover, we wanted to find the delay that would significantly impair participants' performance compared to the no delay condition.

Twenty right handed subjects (age range: 18–35, m=24, 6 males) participated in the study. All of them had normal or corrected to normal vision. They did not have any restricted mobility or suffer form any neurological condition. They were informed about all the aspects of the experiment and gave informed written consent. The experimental protocol was approved by the Science, Technology, Engineering and Mathematics Ethical Review Committee of the University of Birmingham Ethics Committee. Participants received either cash or credits upon the completion of the study.

The apparatus was the same as the one described in Chapter 3. For a detailed description of the Exploration task please refer to Section 5.2.3.

Participants were randomly assigned in two groups. Group A had to find a target area that was 0.25% of the total workspace, whereas the target area of Group B was 0.5% of the total workspace. We used four different reinforcement delay conditions: no delay, 100ms, 200ms and 300ms. The higher exploration delay was limited by the higher tracking delay that we could use to keep the task coherent. Each participant completed 64 Batches (different target areas), 16 in each delay condition. The reinforcement delay was randomly assigned in each Batch. The mean Irrelevant Distance in Trials 2 – 10 was used as a dependent variable (see Section 3.2.6). A 4x2 mixed ANOVA was used to analyse the data, with one within subject factor of Delay (four levels) and one between subject factor of TAS Group (0.25% and 0.5%).

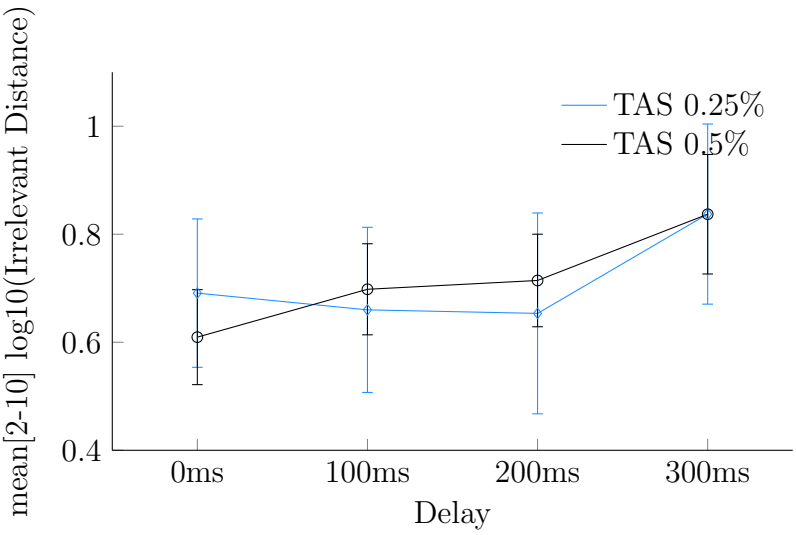

Figure A.1: Mean Irrelevant Distance in Trials 2 – 10. Participants' performance becomes worse while the reinforcement Delay increases. However, there are no differences between the two Target Area Size (TAS) Groups.

The results of the Exploration Pilot study are shown in Figure A.1. The 2x4 mixed ANOVA revealed that there was a significant result of Delay ($F(3, 54) = 5.112$, $p = .003$). The Pairwise Comparisons (Sidak corrected) showed that the following mean differences were significantly different: No Delay condition – 300ms Delay condition ($p = .018$), 200ms Delay condition – 300ms Delay condition ($p = .036$). The mean difference between the 100 ms Delay condition and the 300ms Delay condition was not significant ($p = .067$), as well as the other pairwise comparisons. Moreover, we failed to find any significant interaction of Delay and Group ($F(3, 54) = .702, p = .555$), nor any significant effect of Group ($F(3, 54) = .001, p = .981$).

Based on the aforementioned findings, we decided to use a 300ms reinforcement delay in the main Exploration Study as it was the only condition that was significantly different from other conditions. Although there was no significant Group differences, we chose the smallest Target Area Size for the main study. A final decision we took, based on analyses in the pilot study that are not mentioned here,

was that all participants would be exposed to the same pseudorandomly selected target area locations in each Block of Exploration as a way to minimise the variance between the block and the participants. For each participant, the target locations were fully unpredictable between Blocks.

A.2 Tracking Pilot Study

The main purpose of the Tracking Pilot study was to ensure that we could replicate the adaptation in tracking with visual feedback delays (Foulkes and Miall, 2000) when using the vBots instead of a joystick. Moreover, we wanted to confirm that a visual feedback delay of 300 ms would make the tracking task challenging enough while keeping it coherent at the same time. Finally, we aimed to determine how long participants would have to practice the tracking task in order to show a significant improvement in their performance.

Ten right handed subjects (age range: 20-22, m=21, 2 males) participated in the study. All of them had normal or corrected to normal vision. They did not have any restricted mobility or suffer form any neurological condition. They were informed about all the aspects of the experiment and gave informed written consent. The experimental protocol was approved by the Science, Technology, Engineering and Mathematics Ethical Review Committee of the University of Birmingham Ethics Committee. Participants received either cash or credits upon the completion of the study.

The apparatus, the characteristics of the tracking trials and the dependent variables of the Tracking Pilot Study were the same as the ones described in Chapter 3.

Participants performed 110 tracking trials. Each trial lasted 45 seconds followed by 5 seconds of break. Every 20 trials ($\sim$15 minutes) participants were allowed to take a longer break. The first 10 trials were baseline trials and there were no visual

feedback delays imposed on the hand position. In the reamaining 100 trials the hand position was displayed with a 300ms delay in cursor. Participants were told that they would perform 10 baseline trials and that then the trials would become harder without explicitly being told about the delay.

Two participants were excluded from the analysis because of problems with the program controlling the vBots during their session and one participant withdrew in the last 20 trials because of fatigue. So, the final analysis was conducted in a group of 7 participants. The 100 tracking trials with delay were grouped in five Blocks of 20 trials for the analysis. The RMS error between the target and the cursor and the Mean Power Spectral Density of speed were used as indices of performance (see Section 3.2.6 for details). A Repeated measures ANOVA with one within subject factor of *Block* with 5 levels was conducted for each dependent variable.

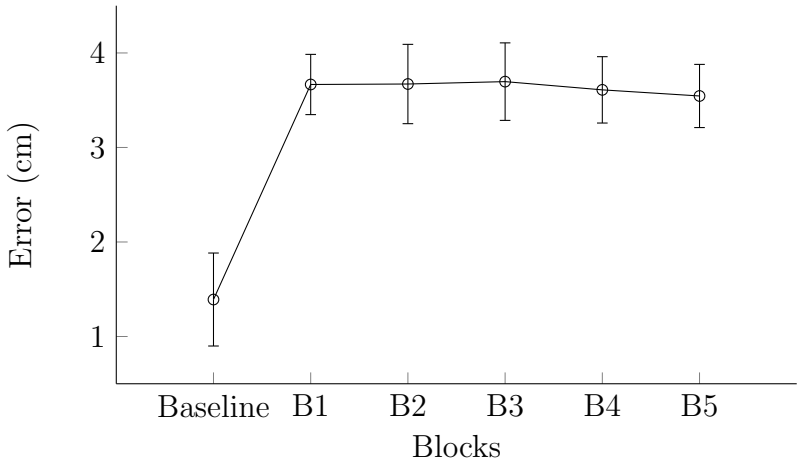

Figure A.2: Mean RMS error between the target and the cursor in the Baseline Block (10 trials) and the 5 Blocks (20 trials each) with visual feedback delays.

Figure A.2 shows how the RMS tracking error varied across the Blocks. The error in the Baseline trials was much smaller compared to the trials with delayed visual feedback (Blocks 1-5). The 1x5 ANOVA did not reveal any significant improvement in the RMS tracking error ($F(2, 13) = .518$, $p = .623$) for the five delayed blocks. As

this might have been expected by the very variable individual responses (Foulkes and Miall, 2000), we further fitted a regression line (see Figure A.3) grouping the trials in ten Blocks of ten trials. The tracking error could be predicted from time (10 Blocks) by the following formula: $Error(time) = -0.016 * time + 3.725, R^2 = 0.502$. The gradient of the regression line was significant ($p = 0.022$) and the 95% confidence intervals were $[-0.02872, -0.002984]$ for the line's slope. From this we can conclude that participants did improve in terms of the tracking error but with a very slow rate. Approximately 1000 trials would be required to return to baseline errors, if this linear rate is extrapolated.

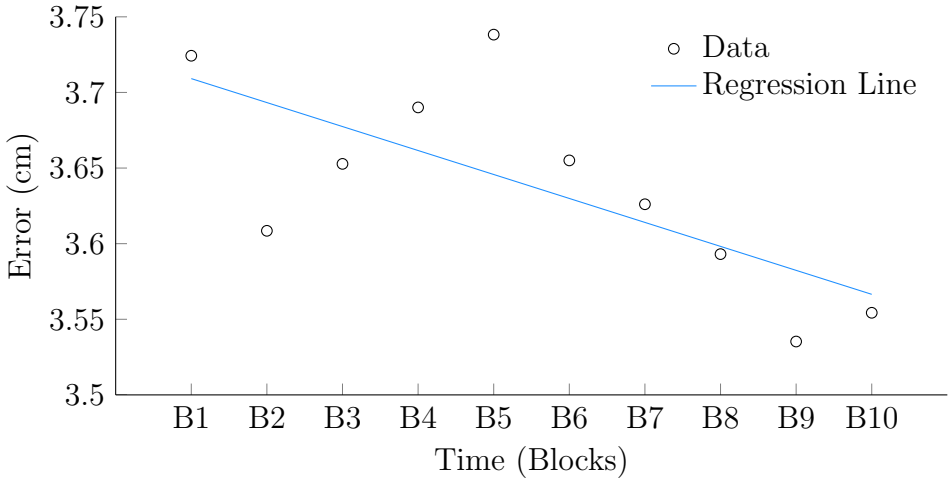

Figure A.3: Regression line of the mean RMS error between the target and the cursor (blue line) and the average group data in Blocks of 10 trials each (black circles)

Figure A.4 shows the performance of the participants in terms of the smoothness of their movements. The 1x5 ANOVA showed that there was a significant improvement in the Mean Power Spectral Density of Speed ($F(2, 12) = 5.024$, $p = .026$). This means that the participants managed to perform smoother movements over practice. Approximately, the power had reduced by 2/3 after the 5 Blocks, suggesting about 10 Blocks of 10 trials might be needed to fully restore the level of smooth

tracking seen in the baseline.

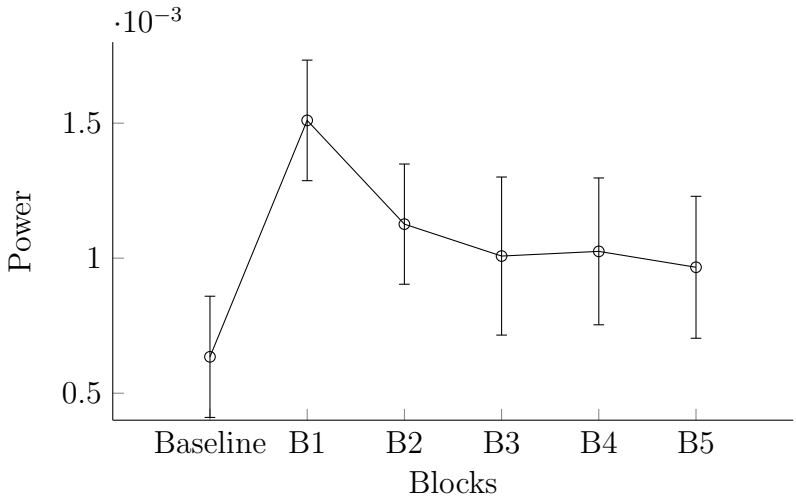

Figure A.4: Mean Power Spectral Density of Speed in the Baseline Block (10 trials) and the 5 Blocks (20 trials each) with visual feedback delays.

Based on the aforementioned results we took the following decisions about the tracking task of the main study. If we wanted to maximise the tracking error performance we would have to increase that total tracking time. However, tracking for longer times within a day would be exhausting for the participants who already reported fatigue towards the end of the pilot study. Moreover, in previous studies (Miall and Jackson, 2006) tracking over multiple days lead to an increased improvement in tracking performance. For these reasons, we chose a two–days design with one hour practice (3 Blocks of 24 trials each) of tracking on each day. This amount of tracking training was a compromise between allowing for better performance to be achieved and not extending the study across many days (more than two). Moreover, given the individual variability in the tracking task we decided to increase the number of participants in the main study.

Appendix B
Appendix of Chapter 4

B.1 Statistical Analysis of Tracking and Centre–Out movements study

Table B.1: Pairwise Comparisons of the Mean Error in tracking - *Visual feedback delays* Groups

1.Block Tracking*							
			Mean Difference			95% Confidence Interval for $Difference^b$	
Tracking	(I) Block	(J) Block	(I-J)	Std. Error	$Sig.^b$	Lower Bound	Upper Bound
CS	1	2	.091*	.034	.013	.021	.162
VS	1	2	.248*	.034	.000	.178	.319
2.Block Tracking*							
			Mean Difference			95% Confidence Interval for $Difference^b$	
Block	(I) Tracking	(J) Tracking	(I-J)	Std. Error	$Sig.^b$	Lower Bound	Upper Bound
1	CS	VS	−.213	.147	.158	−.515	.088
2	CS	VS	−.056	.138	.687	−.340	.227

Based on estimated marginal means
* . The mean difference is significant at the
b . Adjustment for multiple comparisons Sidak

Table B.2: 2x2x2x8 ANoVA of the mean End Point Error

Mauchly's Test of Sphericity

Within Subjects Effect	Mauchly's W	Approx. Chi-Square	df	Sig.	Epsilon: Greenhouse-Geisser
Block	1.000	.000	0	.	1.000
Targets	.200	85.979	27	.000	.706
Block*Targets	.295	65.292	27	.000	.778

Tests of Within-Subjects Effects

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Observed Power
Block	Greenhouse-Geisser	.018	1	.018	.105	.748	.002	.062
Block*Delay	Greenhouse-Geisser	.197	1	.197	1.168	.284	.020	.186
Block*Tracking	Greenhouse-Geisser	.224	1	.224	1.333	.253	.023	.206
Block*Delay*Tracking	Greenhouse-Geisser	.392	1	.392	2.327	.133	.040	.323
Error(Block)	Greenhouse-Geisser	9.428	56	.168				
Targets	Greenhouse-Geisser	26.586	4.941	5.380	10.153	.000	.153	1.000
Targets*Delay	Greenhouse-Geisser	3.380	4.941	.684	1.291	.268	.023	.452
Targets*Tracking	Greenhouse-Geisser	4.630	4.941	.937	1.768	.120	.031	.601
Targets*Delay*Tracking	Greenhouse-Geisser	4.635	4.941	.938	1.770	.120	.031	.601
Error(Targets)	Greenhouse-Geisser	146.644	276.717	.530				
Block*Targets	Greenhouse-Geisser	2.768	5.444	.508	4.544	.000	.075	.979
Block*Targets*Delay	Greenhouse-Geisser	.449	5.444	.083	.738	.607	.013	.277
Block*Targets*Tracking	Greenhouse-Geisser	.649	5.444	.093	1.066	.381	.019	.398
Block*Targets*Delay*Tracking	Greenhouse-Geisser	.658	5.444	.126	1.124	.348	.020	.419
Error(Block*Targets)	Greenhouse-Geisser	34.106	304.858	.112				

Tests of Between-Subjects Effects

Source	Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Observed Power
Intercept	3673.034	1	3673.034	3053.461	.000	.982	1.000
Delay	2.524	1	2.524	2.097	.153	.036	.296
Tracking	.049	1	.049	.041	.841	.001	.055
Delay*Tracking	.558	1	.558	.463	.499	.008	.103
Error	67.418	56	1.204				

Table B.3: Pairwise Comparisons of the Mean of the End Point Error (A)

*1.Block * Targets*

Block	(I) Target	(J) Target	Mean Difference (I-J)	Std. Error	*Sig.*[b]	95% Confidence Interval for *Difference*[b] Lower Bound	Upper Bound
1	1	2	.143	.067	.663	-.077	.363
		3	-.082	.092	1.000	-.382	.218
		4	.106	.075	.993	-.141	.353
		5	-.353*	.098	.018	-.672	-.033
		6	.245*	.073	.038	.007	.484
		7	-.299	.103	.139	-.636	.038
		8	.008	.086	1.000	-.274	.291
	2	3	-.225*	.069	.049	-.449	.000
		4	-.036	.063	1.000	-.243	.170
		5	-.495*	.098	.000	-.816	-.175
		6	.102	.066	.976	-.112	.317
		7	-.442*	.093	.000	-.746	-.137
		8	-.135	.090	.985	-.428	.159
	3	4	.188	.068	.192	-.034	.410
		5	-.271	.092	.121	-.571	.029
		6	.327*	.072	.001	.092	.561
		7	-.217	.076	.162	-.476	.033
		8	.090	.100	1.000	-.238	.418
	4	5	-.459*	.085	.000	-.736	-.182
		6	.139	.071	.800	-.094	.371
		7	-.405*	.092	.001	-.707	-.103
		8	-.098	.077	.998	-.349	.153
	5	6	.598*	.096	.000	.283	.912
		7	.054	.113	1.000	-.315	.423
		8	.361*	.085	.002	.082	.640
	6	7	-.544*	.079	.000	-.803	-.285
		8	-.237	.076	.079	-.486	.012
	7	8	.307	.109	.172	-.050	.664

Based on estimated marginal means
* . The mean difference is significant at the
b . Adjustment for multiple comparisons Sidak

Table B.4: Pairwise Comparisons of the Mean of the End Point Error (B)

*1.Block * Targets*

Block	(I) Target	(J) Target	Mean Difference (I-J)	Std. Error	*Sig.*[b]	95% Confidence Interval for *Difference*[b] Lower Bound	Upper Bound
2	1	2	.160	.071	.554	-.073	.393
		3	.132	.098	.997	-.188	.452
		4	.247	.096	.302	-.067	.562
		5	-.088	.111	1.000	-.452	.276
		6	.438*	.089	.000	.146	.730
		7	.038	.103	1.000	-.300	.376
		8	.173	.103	.943	-.163	.509
	2	3	-.028	.068	1.000	-.250	.193
		4	.087	.076	1.000	-.161	.336
		5	-.248	.095	.288	-.560	.065
		6	.278*	.069	.004	.053	.502
		7	-.122	.099	.999	-.446	.201
		8	.013	.095	1.000	-.297	.323
	3	4	.116	.069	.947	-.110	.342
		5	-.219	.102	.642	-.553	.115
		6	.306*	.084	.016	.031	.581
		7	-.094	.088	1.000	-.383	.195
		8	.041	.105	1.000	-.301	.384
	4	5	- .335*	.086	.008	-.617	-.052
		6	.191	.073	.277	-.048	.430
		7	-.209	.089	.468	-.501	.082
		8	-.074	.091	1.000	-.371	.223
	5	6	.526*	.091	.000	.227	.825
		7	.126	.108	1.000	-.228	.479
		8	.261*	.072	.017	.026	.495
	6	7	-.400*	.076	.000	-.649	-.152
		8	-.265	.086	.084	-.546	.016
	7	8	.135	.095	.993	-.176	.446

Based on estimated marginal means
* . The mean difference is significant at the
b . Adjustment for multiple comparisons Sidak

Table B.5: Pairwise Comparisons of the Mean of the End Point Error (C)

*1.Block * Targets*							
			Mean Difference			95% Confidence Interval for $Difference^b$	
Targets	(I) Block	(J) Block	(I-J)	Std. Error	$Sig.^b$	Lower Bound	Upper Bound
1	1	2	-.175*	.064	.009	-.304	-.046
2	1	2	-.158*	.058	.009	-.274	-.042
3	1	2	.038	.055	.488	-.072	.149
4	1	2	-.034	.057	.556	-.149	.081
5	1	2	.090	.054	.099	-.018	.198
6	1	2	.018	.057	.753	-.096	.132
7	1	2	.162*	.056	.006	.049	.275
8	1	2	-.010	.053	.849	-.116	.096

Based on estimated marginal means
* . The mean difference is significant at the
b . Adjustment for multiple comparisons Sidak

Table B.6: 2x2x2x8 ANoVA of the Standard Deviation of the End Point Error

Mauchly's Test of Sphericity

Within Subjects Effect	Mauchly's W	Approx. Chi-Square	df	Sig.	Epsilon: Greenhouse-Geisser
Block	1.000	.000	0	.	1.000
Targets	.557	31.302	27	.260	.878
Block*Targets	.608	26.589	27	.488	.887

Tests of Within-Subjects Effects

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Observed Power
Block	Greenhouse-Geisser	.109	1	.109	4.153	.046	.069	.517
Block*Delay	Greenhouse-Geisser	.126	1	.126	4.805	.033	.079	.577
Block*Tracking	Greenhouse-Geisser	.011	1	.011	.430	.515	.008	.099
Block*Delay*Tracking	Greenhouse-Geisser	.112	1	.112	4.284	.043	.071	.530
Error(Block)	Greenhouse-Geisser	1.463	56	.026				
Targets	Greenhouse-Geisser	5.335	6.145	.868	22.935	.000	.291	1.000
Targets*Delay	Greenhouse-Geisser	.212	6.145	.034	.910	.489	.016	.365
Targets*Tracking	Greenhouse-Geisser	.212	6.145	.035	.913	.487	.016	.366
Targets*Delay*Tracking	Greenhouse-Geisser	.191	6.145	.031	.820	.557	.014	.329
Error(Targets)	Greenhouse-Geisser	13.026	344.104	.038				
Block*Targets	Greenhouse-Geisser	.209	6.209	.034	1.438	.197	.025	.569
Block*Targets*Delay	Greenhouse-Geisser	.094	6.209	.015	.643	.701	.011	.261
Block*Targets*Tracking	Greenhouse-Geisser	.129	6.209	.021	.890	.505	.016	.359
Block*Targets*Delay*Tracking	Greenhouse-Geisser	.176	6.209	.028	1.211	.299	.021	.486
Error(Block*Targets)	Greenhouse-Geisser	8.142	347.705	.023				

Tests of Between-Subjects Effects

Source	Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Observed Power
Intercept	574.278	1	574.278	5164.298	.000	.989	1.000
Delay	.038	1	.038	.339	.563	.006	.088
Tracking	.233	1	.233	2.098	.153	.036	.296
Delay*Tracking	.007	1	.007	.062	.804	.001	.057
Error	6.227	56	.111				

Table B.7: Pairwise Comparisons of the Standard Deviation of the End Point Error

*1.Delay*Velocity*Block*

				Mean Difference			95% Confidence Interval for $Difference^b$	
Delay	Tracking	(I) Block	(J) Block	(I-J)	Std. Error	$Sig.^b$	Lower Bound	Upper Bound
No Delay	CS	1	2	.016	.021	.455	−.026	.057
	VS	1	2	.073*	.021	.001	.031	.114
Delay	CS	1	2	.013	.021	.531	−.029	.055
	VS	1	2	−.016	.021	.436	−.058	.025

*2.Delay*Velocity*Block*

				Mean Difference			95% Confidence Interval for $Difference^b$	
Block	Tracking	(I) Delay	(J) No Delay	(I-J)	Std. Error	$Sig.^b$	Lower Bound	Upper Bound
1	CS	No Delay	Delay	.019	.030	.525	−.041	.079
	VS	No Delay	Delay	.052	.030	.091	−.008	.112
2	CS	No Delay	Delay	.017	.037	.657	−.058	.091
	VS	No Delay	Delay	−.037	.037	.321	−.112	.037

*3.Delay*Velocity*Block*

				Mean Difference			95% Confidence Interval for $Difference^b$	
Block	Delay	(I) Tracking	(J) Tracking	(I-J)	Std. Error	$Sig.^b$	Lower Bound	Upper Bound
1	No Delay	CS	VS	.008	.030	.788	−.052	.068
	Delay	CS	VS	.041	.030	.182	−.020	.101
2	No Delay	CS	VS	.065	.037	.087	−.010	.140
	Delay	CS	VS	.011	.037	.768	−.064	.086

Based on estimated marginal means
* . The mean difference is significant at the
b . Adjustment for multiple comparisons Sidak

Bibliography

Alexander GE, DeLong MR, Strick PL (1986) Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annu Rev Neurosci* 9:357–381.

Antal A, Nitsche MA, Kincses TZ, Kruse W, Hoffmann KP, Paulus W (2004) Facilitation of visuo-motor learning by transcranial direct current stimulation of the motor and extrastriate visual areas in humans. *Eur J Neurosci* 19:2888–2892.

Barto AG, Dietterich TG (2004) Reinforcement learning and its relationship to supervised learning In Si J, Barto AP W, Wunsch D, editors, *Handbook of Learning and Approximate Dynamic Programming*. John Wiley & Sons, Inc., Hoboken, NJ, USA.

Bastian AJ (2008) Understanding sensorimotor adaptation and learning for rehabilitation. *Curr Opin Neurol* 21:628–633.

Bays PM, Wolpert DM (2007) Computational principles of sensorimotor control that minimize uncertainty and variability. *J Physiol* 578:387–396.

Berniker M, Kording K (2008) Estimating the sources of motor errors for adaptation and generalization. *Nat Neurosci* 11:1454–1461.

Berniker M, Kording KP (2011) Estimating the relevance of world disturbances to explain savings, interference and long-term motor adaptation effects. *PLoS Comput Biol* 7:e1002210.

Bernstein NA (1967) *The co-ordination and regulation of movements* Pergamon Press, Oxford.

Bhushan N, Shadmehr R (1999) Computational nature of human adaptive control during learning of reaching movements in force fields. *Biol Cybern* 81:39–60.

Blakemore SJ, Wolpert DM, Frith CD (1998) Central cancellation of self-produced tickle sensation. *Nat Neurosci* 1:635–640.

Block HJ, Bastian AJ (2012) Cerebellar involvement in motor but not sensory adaptation. *Neuropsychologia* 50:1766–1775.

Bostan AC, Dum RP, Strick PL (2010) The basal ganglia communicate with the cerebellum. *Proc Natl Acad Sci U S A* 107:8452–8456.

Bostan AC, Dum RP, Strick PL (2013) Cerebellar networks with the cerebral cortex and basal ganglia. *Trends Cogn Sci* 17:241–254.

Bostan AC, Strick PL (2010) The cerebellum and basal ganglia are interconnected. *Neuropsychol Rev* 20:261–270.

Boyer EO, Babayan BM, Bevilacqua F, Noisternig M, Warusfel O, Roby-Brami A, Hanneton S, Viaud-Delmon I (2013) From ear to hand: the role of the auditory-motor loop in pointing to an auditory source. *Front Comput Neurosci* 7:26.

Calderon DP, Fremont R, Kraenzlin F, Khodakhah K (2011) The neural substrates of rapid-onset dystonia-parkinsonism. *Nat Neurosci* 14:357–365.

Cerminara NL, Apps R, Marple-Horvat DE (2009) An internal model of a moving visual target in the lateral cerebellum. *J Physiol* 587:429–442.

Cohen DA, Pascual-Leone A, Press DZ, Robertson EM (2005) Off-line learning of motor skill memory: a double dissociation of goal and movement. *Proc Natl Acad Sci U S A* 102:18237–18241.

Criscimagna-Hemminger SE, Bastian AJ, Shadmehr R (2010) Size of error affects cerebellar contributions to motor learning. *J Neurophysiol* 103:2275–2284.

Cunningham DW, Billock VA, Tsou BH (2001b) Sensorimotor adaptation to violations of temporal contiguity. *Psychol Sci* 12:532–535.

Cunningham DW, Chatziastros A, von der Heyde M, Blthoff HH (2001a) Driving in the future: temporal visuomotor adaptation and generalization. *J Vis* 1:88–98.

Dam G, Kording K, Wei K (2013) Credit assignment during movement reinforcement learning. *PLoS One* 8:e55352.

Daw ND, Doya K (2006) The computational neurobiology of learning and reward. *Curr Opin Neurobiol* 16:199–204.

de Rugy A, Loeb GE, Carroll TJ (2012) Muscle coordination is habitual rather than optimal. *J Neurosci* 32:7384–7391.

Di Luca M, Machulla T, Ernst M (2009) Recalibration of multisensory simultaneity: cross-modal transfer coincides with a change in perceptual latency. *J Vis* 9:7.1–716.

Diedrichsen J, Shadmehr R, Ivry RB (2010b) The coordination of movement: optimal feedback control and beyond. *Trends Cogn Sci* 14:31–39.

Diedrichsen J, White O, Newman D, Lally N (2010a) Use-dependent and error-based learning of motor behaviors. *J Neurosci* 30:5159–5166.

Doya K (2000) Complementary roles of basal ganglia and cerebellum in learning and motor control. *Curr Opin Neurobiol* 10:732–739.

Ebner TJ, Pasalar S (2008) Cerebellum predicts the future motor state. *Cerebellum* 7:583–588.

Ethier V, Zee DS, Shadmehr R (2008) Spontaneous recovery of motor memory during saccade adaptation. *J Neurophysiol* 99:2577–2583.

Ferrucci R, Marceglia S, Vergari M, Cogiamanian F, Mrakic-Sposta S, Mameli F, Zago S, Barbieri S, Priori A (2008) Cerebellar transcranial direct current stimulation impairs the practice-dependent proficiency increase in working memory. *J Cogn Neurosci* 20:1687–1697.

Ferrucci R, Brunoni AR, Parazzini M, Vergari M, Rossi E, Fumagalli M, Mameli F, Rosa M, Giannicola G, Zago S, Priori A (2013) Modulating human procedural learning by cerebellar transcranial direct current stimulation. *Cerebellum* 12:485–492.

Foulkes AJ, Miall RC (2000) Adaptation to visual feedback delays in a human manual tracking task. *Exp Brain Res* 131:101–110.

Fregni F, Boggio PS, Nitsche M, Bermpohl F, Antal A, Feredoes E, Marcolin MA, Rigonatti SP, Silva MTA, Paulus W, Pascual-Leone A (2005) Anodal transcranial direct current stimulation of prefrontal cortex enhances working memory. *Exp Brain Res* 166:23–30.

Frens MA, Donchin O (2009) Forward models and state estimation in compensatory eye movements. *Front Cell Neurosci* 3:13.

Fricke K, Seeber AA, Thirugnanasambandam N, Paulus W, Nitsche MA, Rothwell JC (2011) Time course of the induction of homeostatic plasticity generated by repeated transcranial direct current stimulation of the human motor cortex. *J Neurophysiol* 105:1141–1149.

Galea JM, Jayaram G, Ajagbe L, Celnik P (2009) Modulation of cerebellar ex-

citability by polarity-specific noninvasive direct current stimulation. *J Neurosci* 29:9115–9122.

Galea JM, Vazquez A, Pasricha N, de Xivry JJO, Celnik P (2011) Dissociating the roles of the cerebellum and motor cortex during adaptive learning: the motor cortex retains what the cerebellum learns. *Cereb Cortex* 21:1761–1770.

Gentner R, Gorges S, Weise D, aufm Kampe K, Buttmann M, Classen J (2010) Encoding of motor skill in the corticomuscular system of musicians. *Curr Biol* 20:1869–1874.

Gentner R, Wankerl K, Reinsberger C, Zeller D, Classen J (2008) Depression of human corticospinal excitability induced by magnetic theta-burst stimulation: evidence of rapid polarity-reversing metaplasticity. *Cereb Cortex* 18:2046–2053.

Gibo TL, Criscimagna-Hemminger SE, Okamura AM, Bastian AJ (2013) Cerebellar motor learning: are environment dynamics more important than error size? *J Neurophysiol* 110:322–333.

Glaescher J, Daw N, Dayan P, O'Doherty JP (2010) States versus rewards: dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron* 66:585–595.

Gordon J, Ghilardi MF, Cooper SE, Ghez C (1994) Accuracy of planar reaching movements. ii. systematic extent errors resulting from inertial anisotropy. *Exp Brain Res* 99:112–130.

Haggard P, Wolpert DM (2005) Disorders of body scheme. In Freund H, Jeannerod M, Hallett M, Leiguarda R, editors, *Higher-order motor disorders*. Oxford University Press.

Haith A, Jackson C, Miall CV S (2008) Unifying the sensory and motor components of sensorimotor adaptation. *Proceedings of Advances in Neural Information Processing Systems (NIPS)* p. 593600.

Haith AM, Krakauer JW (2013) Model-based and model-free mechanisms of human motor learning. *Adv Exp Med Biol* 782:1–21.

Hardwick RM, Dagioglou M, Miall R (2013) State estimation and the cerebellum In Manto M, Schmahmann J, Rossi F, Gruol D, Koibuchi N, editors, *Handbook of the Cerebellum and Cerebellar Disorders*, pp. 1297–1313. Springer Netherlands.

Hartigan J, Hartigan P (1985) The dip test of unimodality. *Annals of Statistics* 13:7084.

Hasan A, Hamada M, Nitsche MA, Ruge D, Galea JM, Wobrock T, Rothwell JC (2012) Direct-current-dependent shift of theta-burst-induced plasticity in the human motor cortex. *Exp Brain Res* 217:15–23.

Hatada Y, Rossetti Y, Miall RC (2006) Long-lasting aftereffect of a single prism adaptation: shifts in vision and proprioception are independent. *Exp Brain Res* 173:415–424.

Henriques DYP, Cressman EK (2012) Visuomotor adaptation and proprioceptive recalibration. *J Mot Behav* 44:435–444.

Hikosaka O, Wurtz RH (1983) Visual and oculomotor functions of monkey substantia nigra pars reticulata. iii. memory-contingent visual and saccade responses. *J Neurophysiol* 49:1268–1284.

Honda T, Hirashima M, Nozaki D (2012a) Adaptation to visual feedback delay influences visuomotor learning. *PLoS One* 7:e37900.

Honda T, Hirashima M, Nozaki D (2012b) Habituation to feedback delay restores degraded visuomotor adaptation by altering both sensory prediction error and the sensitivity of adaptation to the error. *Front Psychol* 3:540.

Hoshi E, Tremblay L, Fger J, Carras PL, Strick PL (2005) The cerebellum communicates with the basal ganglia. *Nat Neurosci* 8:1491–1493.

Howard IS, Ingram JN, Wolpert DM (2009) A modular planar robotic manipulandum with end-point torque control. *Journal of neuroscience methods* 181:199–211.

Huang VS, Haith A, Mazzoni P, Krakauer JW (2011) Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. *Neuron* 70:787–801.

Hunter T, Sacco P, Nitsche MA, Turner DL (2009) Modulation of internal model formation during force field-induced motor learning by anodal transcranial direct current stimulation of primary motor cortex. *J Physiol* 587:2949–2961.

Izawa J, Criscimagna-Hemminger SE, Shadmehr R (2012) Cerebellar contributions to reach adaptation and learning sensory consequences of action. *J Neurosci* 32:4230–4239.

Izawa J, Shadmehr R (2011) Learning from sensory and reward prediction errors during motor adaptation. *PLoS Comput Biol* 7:e1002012.

Jacobson L, Koslowsky M, Lavidor M (2012) tdcs polarity effects in motor and cognitive domains: a meta-analytical review. *Exp Brain Res* 216:1–10.

Jayaram G, Tang B, Pallegadda R, Vasudevan EVL, Celnik P, Bastian AJ (2012) Modulating locomotor adaptation with cerebellar stimulation. *J Neurophysiol* .

Jenkinson N, Miall RC (2010) Disruption of saccadic adaptation with repetitive transcranial magnetic stimulation of the posterior cerebellum in humans. *Cerebellum* 9:548–555.

Jueptner M, Frith CD, Brooks DJ, Frackowiak RS, Passingham RE (1997c) Anatomy of motor learning. ii. subcortical structures and learning by trial and error. *J Neurophysiol* 77:1325–1337.

Jueptner M, Jenkins IH, Brooks DJ, Frackowiak RS, Passingham RE (1996) The sensory guidance of movement: a comparison of the cerebellum and basal ganglia. *Exp Brain Res* 112:462–474.

Jueptner M, Ottinger S, Fellows SJ, Adamschewski J, Flerich L, Mller SP, Diener HC, Thilmann AF, Weiller C (1997a) The relevance of sensory input for the cerebellar control of movements. *Neuroimage* 5:41–48.

Jueptner M, Stephan KM, Frith CD, Brooks DJ, Frackowiak RS, Passingham RE (1997b) Anatomy of motor learning. i. frontal cortex and attention to action. *J Neurophysiol* 77:1313–1324.

Jueptner M, Weiller C (1998) A review of differences between basal ganglia and cerebellar control of movements as revealed by functional imaging studies. *Brain* 121 (Pt 8):1437–1449.

Keisler A, Shadmehr R (2010) A shared resource between declarative memory and motor memory. *J Neurosci* 30:14817–14823.

Kitazawa S, Kohno T, Uka T (1995) Effects of delayed visual information on the rate and amount of prism adaptation in the human. *J Neurosci* 15:7644–7652.

Kitazawa S, Yin PB (2002) Prism adaptation with delayed visual error signals in the monkey. *Exp Brain Res* 144:258–261.

Kluzik J, Diedrichsen J, Shadmehr R, Bastian AJ (2008) Reach adaptation: what determines whether we learn an internal model of the tool or adapt the model of our arm? *J Neurophysiol* 100:1455–1464.

Koerding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427:244–247.

Krakauer JW, Mazzoni P (2011) Human sensorimotor learning: adaptation, skill, and beyond. *Curr Opin Neurobiol* 21:636–644.

Lackner JR, DiZio P (2005) Motor control and learning in altered dynamic environments. *Curr Opin Neurobiol* 15:653–659.

Latash ML (2012a) The bliss (not the problem) of motor abundance (not redundancy). *Exp Brain Res* 217:1–5.

Latash ML (2012b) Movements that are both variable and optimal. *J Hum Kinet* 34:5–13.

Levy N, Pressman A, Mussa-Ivaldi FA, Karniel A (2010) Adaptation to delayed force perturbations in reaching movements. *PLoS One* 5:e12128.

Liebetanz D, Koch R, Mayenfels S, Knig F, Paulus W, Nitsche MA (2009) Safety limits of cathodal transcranial direct current stimulation in rats. *Clin Neurophysiol* 120:1161–1167.

Liebetanz D, Nitsche MA, Tergau F, Paulus W (2002) Pharmacological approach to the mechanisms of transcranial dc-stimulation-induced after-effects of human motor cortex excitability. *Brain* 125:2238–2247.

Liu X, Robertson E, Miall RC (2003) Neuronal activity related to the visual representation of arm movements in the lateral cerebellar cortex. *J Neurophysiol* 89:1223–1237.

Loeb GE (2012) Optimal isn't good enough. *Biol Cybern* 106:757–765.

Martin TA, Keating JG, Goodkin HP, Bastian AJ, Thach WT (1996) Throwing while looking through prisms. i. focal olivocerebellar lesions impair adaptation. *Brain* 119 (Pt 4):1183–1198.

Mazzoni P, Krakauer JW (2006) An implicit plan overrides an explicit strategy during visuomotor adaptation. *J Neurosci* 26:3642–3645.

Miall C (2013) 10,000 hours to perfection. *Nat Neurosci* 16:1168–1169.

Miall RC, Imamizu H, Miyauchi S (2000) Activation of the cerebellum in co-ordinated eye and hand tracking movements: an fmri study. *Exp Brain Res* 135:22–33.

Miall RC, Jackson JK (2006) Adaptation to visual feedback delays in manual tracking: evidence against the smith predictor model of human visually guided action. *Exp Brain Res* 172:77–84.

Miall RC, Jenkinson EW (2005) Functional imaging of changes in cerebellar activity related to learning during a novel eye-hand tracking task. *Exp Brain Res* 166:170–183.

Miall RC, Reckess GZ, Imamizu H (2001) The cerebellum coordinates eye and hand tracking movements. *Nat Neurosci* 4:638–644.

Miall RC, Weir DJ, Stein JF (1986) Manual tracking of visual targets by trained monkeys. *Behav Brain Res* 20:185–201.

Miall RC, Weir DJ, Wolpert DM, Stein JF (1993) Is the cerebellum a smith predictor? *J Mot Behav* 25:203–216.

Miall RC, Christensen LOD, Cain O, Stanley J (2007) Disruption of state estimation in the human lateral cerebellum. *PLoS Biol* 5:e316.

Middleton FA, Strick PL (1994) Anatomical evidence for cerebellar and basal ganglia involvement in higher cognitive function. *Science* 266:458–461.

Middleton FA, Strick PL (2000) Basal ganglia and cerebellar loops: motor and cognitive circuits. *Brain Res Brain Res Rev* 31:236–250.

Monte-Silva K, Kuo MF, Liebetanz D, Paulus W, Nitsche MA (2010) Shaping the optimal repetition interval for cathodal transcranial direct current stimulation (tdcs). *J Neurophysiol* 103:1735–1740.

Nitsche MA, Paulus W (2000) Excitability changes induced in the human motor cortex by weak transcranial direct current stimulation. *J Physiol* 527 Pt 3:633–639.

Nitsche MA, Roth A, Kuo MF, Fischer AK, Liebetanz D, Lang N, Tergau F, Paulus W (2007) Timing-dependent modulation of associative plasticity by general network excitability in the human motor cortex. *J Neurosci* 27:3807–3812.

Nitsche MA, Schauenburg A, Lang N, Liebetanz D, Exner C, Paulus W, Tergau F (2003) Facilitation of implicit motor learning by weak transcranial direct current stimulation of the primary motor cortex in the human. *J Cogn Neurosci* 15:619–626.

Orban de Xivry JJ, Marko MK, Pekny SE, Pastor D, Izawa J, Celnik P, Shadmehr R (2011) Stimulation of the human motor cortex alters generalization patterns of motor learning. *J Neurosci* 31:7102–7110.

Ostry DJ, Darainy M, Mattar AAG, Wong J, Gribble PL (2010) Somatosensory plasticity and motor learning. *J Neurosci* 30:5384–5393.

Pasalar S, Roitman AV, Durfee WK, Ebner TJ (2006) Force field effects on cerebellar purkinje cell discharge with implications for internal models. *Nat Neurosci* 9:1404–1411.

Paulus W (2011) Transcranial stimulation techniques: which genetics is the best for which purpose? *Journal of Physiology* 589:1245.

Picard N, Matsuzaka Y, Strick PL (2013) Extended practice of a motor skill is associated with reduced metabolic activity in m1. *Nat Neurosci* 16:1340–1347.

Pope PA, Miall RC (2012) Task-specific facilitation of cognition by cathodal transcranial direct current stimulation of the cerebellum. *Brain Stimul* 5:84–94.

Rabe K, Livne O, Gizewski ER, Aurich V, Beck A, Timmann D, Donchin O (2009) Adaptation to visuomotor rotation and force field perturbation is correlated to different brain areas in patients with cerebellar degeneration. *J Neurophysiol* 101:1961–1971.

Rao RPN (2010) Decision making under uncertainty: a neural model based on partially observable markov decision processes. *Front Comput Neurosci* 4:146.

Redgrave P, Gurney K (2006) The short-latency dopamine signal: a role in discovering novel actions? *Nat Rev Neurosci* 7:967–975.

Redgrave P, Gurney K, Reynolds J (2008) What is reinforced by phasic dopamine signals? *Brain Res Rev* 58:322–339.

Reis J, Schambra HM, Cohen LG, Buch ER, Fritsch B, Zarahn E, Celnik PA, Krakauer JW (2009) Noninvasive cortical stimulation enhances motor skill ac-

quisition over multiple days through an effect on consolidation. *Proc Natl Acad Sci U S A* 106:1590–1595.

Robertson EM, Miall RC (1999) Visuomotor adaptation during inactivation of the dentate nucleus. *Neuroreport* 10:1029–1034.

Robertson EM, Pascual-Leone A, Miall RC (2004) Current concepts in procedural consolidation. *Nat Rev Neurosci* 5:576–582.

Roitman AV, Pasalar S, Johnson MTV, Ebner TJ (2005) Position, direction of movement, and speed tuning of cerebellar purkinje cells during circular manual tracking in monkey. *J Neurosci* 25:9244–9257.

Roth MJ, Synofzik M, Lindner A (2013) The cerebellum optimizes perceptual predictions about external sensory events. *Curr Biol* 23:930–935.

Sadnicka A, Hoffland BS, Bhatia KP, van de Warrenburg BP, Edwards MJ (2012) The cerebellum in dystonia - help or hindrance? *Clin Neurophysiol* 123:65–70.

Samejima K, Doya K (2007) Multiple representations of belief states and action values in corticobasal ganglia loops. *Ann N Y Acad Sci* 1104:213–228.

Schaefer SY, Shelly IL, Thoroughman KA (2012) Beside the point: motor adaptation without feedback-based error correction in task-irrelevant conditions. *J Neurophysiol* 107:1247–1256.

Schambra HM, Abe M, Luckenbaugh DA, Reis J, Krakauer JW, Cohen LG (2011) Probing for hemispheric specialization for motor skill learning: a transcranial direct current stimulation study. *J Neurophysiol* 106:652–661.

Schlerf JE, Xu J, Klemfuss NM, Griffiths TL, Ivry RB (2013) Individuals with cerebellar degeneration show similar adaptation deficits with large and small visuomotor errors. *J Neurophysiol* 109:1164–1173.

Schultz W (2013) Updating dopamine reward signals. *Curr Opin Neurobiol* 23:229–238.

Scott SH (2004) Optimal feedback control and the neural basis of volitional motor control. *Nat Rev Neurosci* 5:532–546.

Seidler RD, Noll DC, Chintalapati P (2006) Bilateral basal ganglia activation associated with sensorimotor adaptation. *Exp Brain Res* 175:544–555.

Seymour B, O'Doherty JP, Dayan P, Koltzenburg M, Jones AK, Dolan RJ, Friston KJ, Frackowiak RS (2004) Temporal difference models describe higher-order learning in humans. *Nature* 429:664–667.

Shabbott BA, Sainburg RL (2010) Learning a visuomotor rotation: simultaneous visual and proprioceptive information is crucial for visuomotor remapping. *Exp Brain Res* 203:75–87.

Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14:3208–3224.

Shadmehr R (2010) Control of movements and temporal discounting of reward. *Curr Opin Neurobiol* 20:726–730.

Shadmehr R, Smith MA, Krakauer JW (2010) Error correction, sensory prediction, and adaptation in motor control. *Annu Rev Neurosci* 33:89–108.

Shmuelof L, Krakauer JW (2011) Are we ready for a natural history of motor learning? *Neuron* 72:469–476.

Siebner HR, Lang N, Rizzo V, Nitsche MA, Paulus W, Lemon RN, Rothwell JC (2004) Preconditioning of low-frequency repetitive transcranial magnetic stimulation with transcranial direct current stimulation: evidence for homeostatic plasticity in the human motor cortex. *J Neurosci* 24:3379–3385.

Simani MC, McGuire LMM, Sabes PN (2007) Visual-shift adaptation is composed of separable sensory and task-dependent effects. *J Neurophysiol* 98:2827–2841.

Skinner B (1948) Superstition in the pigeon. *J Exp Psychol* 38:168–172.

Smith MA, Ghazizadeh A, Shadmehr R (2006) Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol* 4:e179.

Stafford T, Thirkettle M, Walton T, Vautrelle N, Hetherington L, Port M, Gurney K, Redgrave P (2012) A novel task for the investigation of action acquisition. *PLoS One* 7:e37749.

Stagg CJ, Jayaram G, Pastor D, Kincses ZT, Matthews PM, Johansen-Berg H (2011) Polarity and timing-dependent effects of transcranial direct current stimulation in explicit motor learning. *Neuropsychologia* 49:800–804.

Stanley J, Krakauer JW (2013) Motor skill depends on knowledge of facts. *Front Hum Neurosci* 7:503.

Sutton RS, Barto AG (1998) *Reinforcement Learning: An Introduction* MIT Press, Cambridge, MA.

Synofzik M, Lindner A, Thier P (2008) The cerebellum updates predictions about the visual consequences of one's behavior. *Curr Biol* 18:814–818.

Synofzik M, Thier P, Lindner A (2006) Internalizing agency of self-action: perception of one's own hand movements depends on an adaptable prediction about the sensory action outcome. *J Neurophysiol* 96:1592–1601.

Tanaka H, Homma K, Imamizu H (2011) Physical delay but not subjective delay determines learning rate in prism adaptation. *Exp Brain Res* 208:257–268.

Taylor JA, Ivry RB (2011) Flexible cognitive strategies during motor learning. *PLoS Comput Biol* 7:e1001096.

Taylor JA, Klemfuss NM, Ivry RB (2010) An explicit strategy prevails when the cerebellum fails to compute movement errors. *Cerebellum* 9:580–586.

Tekin S, Cummings JL (2002) Frontal-subcortical neuronal circuits and clinical neuropsychiatry: an update. *J Psychosom Res* 53:647–654.

Thirkettle M, Walton T, Redgrave P, Gurney K, Stafford T (2013) No learning where to go without first knowing where you're coming from: action discovery is trajectory, not endpoint based. *Front Psychol* 4:638.

Thoret H, Haque J, Pascual-Leone A (2001) Increased variability of paced finger tapping accuracy following repetitive magnetic stimulation of the cerebellum in humans. *Neurosci Lett* 306:29–32.

Todorov E (2004) Optimality principles in sensorimotor control. *Nat Neurosci* 7:907–915.

Tseng YW, Diedrichsen J, Krakauer JW, Shadmehr R, Bastian AJ (2007) Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *J Neurophysiol* 98:54–62.

Turrigiano GG, Nelson SB (2004) Homeostatic plasticity in the developing nervous system. *Nat Rev Neurosci* 5:97–107.

van Beers RJ, Haggard P, Wolpert DM (2004) The role of execution noise in movement variability. *J Neurophysiol* 91:1050–1063.

van Donkelaar P, Stein JF, Passingham RE, Miall RC (1999) Neuronal activity in the primate motor thalamus during visually triggered and internally generated limb movements. *J Neurophysiol* 82:934–945.

van Donkelaar P, Stein JF, Passingham RE, Miall RC (2000) Temporary inactivation in the primate motor thalamus during visually triggered and internally generated limb movements. *J Neurophysiol* 83:2780–2790.

Verstynen T, Sabes PN (2011) How each movement changes the next: an experimental and theoretical study of fast adaptive priors in reaching. *J Neurosci* 31:10050–10059.

Volchenkov D, Helbach J, Tscherepanow M, Kuhnel S (2013) Exploration exploitation trade-off features a saltatory search behaviour. *J R Soc Interface* 10:20130352.

Walton T (2011) The Discovery of Novel Actions Ph.D. diss., University of Sheffield.

Walton T, Thirkettle M, Redgrave P, Gurney KN, Stafford T (2013) The discovery of novel actions is affected by very brief reinforcement delays and reinforcement modality. *J Mot Behav* 45:351–360.

Wankerl K, Weise D, Gentner R, Rumpf JJ, Classen J (2010) L-type voltage-gated ca2+ channels: a single molecular switch for long-term potentiation/long-term depression-like plasticity and activity-dependent metaplasticity in humans. *J Neurosci* 30:6197–6204.

Wei K, Koerding K (2009) Relevance of error: what drives motor adaptation? *J Neurophysiol* 101:655–664.

Wei K, Wert D, Koerding K (2011) The nervous system uses nonspecific motor learning in response to random perturbations of varying nature. *J Neurophysiol* 104:3053–3063.

Wiestler T, Diedrichsen J (2013) Skill learning strengthens cortical representations of motor sequences. *Elife* 2:e00801.

Wilke C, Synofzik M, Lindner A (2013) Sensorimotor recalibration depends on attribution of sensory prediction errors to internal causes. *PLoS One* 8:e54925.

Wolpert DM, Ghahramani Z, Jordan MI (1995) An internal model for sensorimotor integration. *Science* 269:1880–1882.

Wolpert DM, Miall RC (1996) Forward models for physiological motor control. *Neural Netw* 9:1265–1279.

Wolpert DM, Miall RC, Kawato M (1998) Internal models in the cerebellum. *Trends Cogn Sci* 2:338–347.

Wolpert DM, Diedrichsen J, Flanagan JR (2011) Principles of sensorimotor learning. *Nat Rev Neurosci* 12:739–751.

Wolpert DM, Flanagan JR (2010) Motor learning. *Curr Biol* 20:R467–R472.

Wolpert DM, Ghahramani Z, Flanagan JR (2001) Perspectives and problems in motor learning. *Trends Cogn Sci* 5:487–494.

Wu HG, Miyamoto YR, Castro LNG, lveczky BP, Smith MA (2014) Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat Neurosci* 17:312–321.

Wu T, Hallett M (2013) The cerebellum in parkinson's disease. *Brain* 136:696–709.