Representation and Interaction of Sensorimotor Learning Processes

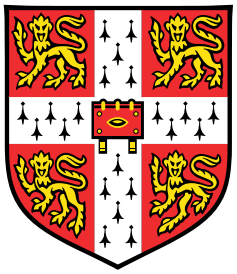

Mohsen Sadeghi

Supervisor: Prof. Daniel Wolpert

Department of Engineering

University of Cambridge

This dissertation is submitted for the degree of

Doctor of Philosophy

Robinson College July 2018

This thesis is dedicated to my loving parents and my beloved brother.

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and contains nothing which is the outcome of work done in collaboration with others, except as specified in the text and Acknowledgements. This dissertation contains fewer than 65,000 words including appendices, bibliography, footnotes, tables and equations and has fewer than 150 figures.

Mohsen Sadeghi

July 2018

Acknowledgements

I am immensely grateful to Daniel Wolpert for giving me this opportunity and for his excellent supervision. My special thanks to James Ingram for his encouraging collaboration and the enlightening discussions. James prepared the 3BOT and the WristBOT manipulanda for the experimental studies, and designed the object manipulation task which was used in this work. I am thankful to my friends and colleagues, Hannah Sheahan and James Heald for their support and helpful discussions. I would also like to thank Diane Unwin and Goda Žalalytė for their significant administrative support. This research was funded by the Wellcome Trust and the Yousef Jameel Scholarship via Cambridge Trust to which I am sincerely grateful. Finally, my thanks and gratitude go to my beloved family and friends who provided me with their unlimited support and encouragement throughout this course.

Abstract

Human sensorimotor control is remarkably adept at utilising contextual information to learn and recall systematic sensorimotor transformations. Here, we investigate the motor representations that underlie such learning, and examine how motor memories acquired based on different contextual information interact. Using a novel three-dimensional robotic manipulandum, the 3BOT, we examined the spatial transfer of learning across various movement directions in a 3D environment, while human subjects performed reaching movements under velocity-dependent force field. The obtained pattern of generalisation suggested that the representation of dynamic learning was most likely defined in a target-based, rather than an extrinsic, coordinate system. We further examined how motor memories interact when subjects adapt to force fields applied in orthogonal dimensions. We found that, unlike opposing fields, learning two spatially orthogonal force fields led to the formation of separate motor memories, which neither interfered with nor facilitated each other. Moreover, we demonstrated a novel, more general aspect of the spontaneous recovery phenomenon using a two-dimensional force field task: when subjects learned two orthogonal force fields consecutively, in the following phase of clamped error feedback, the expression of adaptation spontaneously rotated from the direction of the second force field, towards the direction of the first force field. Finally, we examined the interaction of sensorimotor memories formed based on separate contextual information. Subjects performed reciprocating reaching and object manipulation tasks under two alternating contexts (movement directions), while we manipulated the dynamics of the task in each context separately. The results suggested that separate motor memories were formed for the dynamics of the task in different contexts, and that these motor memories interacted by sharing error signals to enhance learning. Importantly, the extent of interaction was not fixed between the context-dependent motor memories, but adaptively changed according to the task dynamics to potentially improve overall performance. Together, our experimental and theoretical results add to the understanding of mechanisms that underlie sensorimotor learning, and the way these mechanisms interact under various tasks and different dynamics.

Table of contents

List of figures **xv**

List of tables **xxv**

I INTRODUCTION **1**

1 Introduction **3**

1.1 Internal models 4
1.2 Learning mechanisms 6
1.3 Generalisation of learning 11
1.4 Learning and memory 15
1.5 The organisation of the dissertation 22
1.5.1 Representation of learning in 3D environment 22
1.5.2 Interaction between the learning processes 26

II SENSORIMOTOR LEARNING IN 3D ENVIRONMENTS **29**

2 Spatial generalisation of adaptation in 3D **31**

2.1 Introduction 31
2.2 Methods 34
2.2.1 Experimental procedure 34
2.2.2 Data analysis 37
2.3 Results 39
2.4 Discussion 45

3 Spontaneous Recovery **51**

3.1 Introduction 51

3.2 Methods . . . 53
3.2.1 Experimental procedure . . . 53
3.2.2 Analysis . . . 54
3.2.3 Modelling . . . 56
3.3 Results . . . 57
3.3.1 Kinematic error . . . 58
3.3.2 Adaptation . . . 60
3.3.3 Model fits . . . 63
3.3.4 One-dimensional spontaneous recovery . . . 65
3.4 Discussion . . . 68

4 Interference and facilitation in sensorimotor learning 71
4.1 Introduction . . . 71
4.2 Methods . . . 73
4.2.1 Experimental procedure . . . 73
4.2.2 Analysis . . . 76
4.3 Results . . . 77
4.4 Discussion . . . 83

III INTERACTION BETWEEN LEARNING PROCESSES 87

5 Adaptation to novel dynamics 89
5.1 Introduction . . . 89
5.2 Methods . . . 91
5.2.1 Trial Type . . . 93
5.2.2 Conditions . . . 93
5.2.3 Analysis . . . 96
5.3 Results . . . 97
5.3.1 Model free analysis . . . 97
5.3.2 Modelling . . . 105
5.4 Discussion . . . 120

6 Adaptation to familiar dynamics 125
6.1 Introduction . . . 125
6.2 Methods . . . 126
6.2.1 Task . . . 127
6.2.2 Analysis . . . 129

6.2.3 Experiments 130
6.2.4 Model fitting 134
6.3 Results 135
6.3.1 Experiment 1 135
6.3.2 Experiment 2 139
6.4 Discussion 146

IV SUMMARY AND DISCUSSION 151

7 General discussion 153
7.1 Summary 153
7.2 Current work and future direction 156
7.3 Conclusion 165

References 167

List of figures

1.1 Spontaneous recovery induced by adaptation to two successive opposing force fields (Smith et al., 2006). A dual-rate model with slow and fast adaptation processes accounted for the empirical data. 21

1.2 **A**. The planar manipulandum. **B**. Three types of trials can be implemented by the planar robot: the null trials in which the arm can freely move in the workspace; the force field trials in which a velocity/position dependent perturbation is applied causing large deviations in the hand trajectory; and the channel trials in which the forces generated by subjects are measured. **C**. The 3D capable robotic manipulandum (3BOT). The 3BOT is a serial robotic arm with three degrees of freedom in the joint space (θ_1, θ_2 and θ_3) that can control translational movements of the end-effector in 3D. The virtual reality environment is provided by a 3D display paired with a mirror, and a pair of 3D active glasses. **D**. The 3BOT can apply force fields in different spatial directions. It can also implement channel tubes that measure subject's generated forces orthogonal to the movement direction (i.e., the X-Z plane). 23

2.1 **A**. The pattern of generalisation to different target directions (data from Howard and Franklin (2015)). Adaptation is maximum at the training target ($\phi = 0^\circ$) and decays for targets farther away. **B**. Schematic illustration of the transfer of force field learning from one movement direction (Y axis of the main frame) to another (y axis of the target-based frame). The learned adaptive force in the training direction (green arrow) can be transferred to the new orientation (ϕ) in two possible frame representations: the target-based frame (purple arrow) or the extrinsic frame (blue arrow). The force measurement by channel trials is limited to the lateral dimension of the movement direction (that is, the x axis of the target-based frame). Thus, if learning is transferred in the extrinsic frame (blue arrow) only its lateral component can be measured by channel trials. 32

2.2 **A.** Target positions. Targets were distributed on the surface of an imaginary sphere (radius of 12 cm), with the starting position at the centre of the sphere (yellow circle). Three different views of the target locations is shown: the frontal view (left panel), the sagittal view (top right) and the transverse view (bottom right). Targets were lined up on three different planes in the 3D space (transverse, sagittal and diagonal planes) with the angular distance of 22.5° with each other. **B.** The experiment consisted of four successive phases: pre-exposure phase, exposure phase, generalisation, and post exposure phase. **C.** The target based (local) frame for each target (*xyz*-red frame) was obtained by two consecutive rotations of the main (extrinsic) frame (XYZ-grey frame), first around the Z axis for ϕ_z degrees, and then along the resultant x axis for ϕ_x degrees. The y-axis of the resultant local frame pointed towards the target location, and the x and z axes were used to represent the force and adaptation measurements on the channel trials towards that target. Note that the x axis of the local frame always remained on the transverse (XY) plane, in which also the force field was applied. Therefore, compensation of the force field by subjects was reflected mainly in the x component of the measured forces. **D.** The rotation angles (in degrees) associated with the local frame of each target. 36

2.3 **A**. The average (±SE) time course of maximum perpendicular error (MPE) across 8 subjects during reaching to the central target. Data shows the MPE over pre-exposure (green), exposure (red), and post-exposure (blue) trials. **B**. The average (±SE) of MPE across the final 10 trials of pre-exposure, exposure, generalisation, and post-exposure phase. **C**. The mean (±SE) of adaptation index for the central target over the channel trials during different phases of the experiment. The bar plots show the mean adaptation in the last five channel trials of the exposure phase and generalisation phase. The error bars represent the SE across subjects, with p-values determined by paired t-test (details in the text). 40

2.4 **A**. Force vectors exerted by individual subjects in response to the force field (yellow arrow) measured on channel trials at maximum hand velocity. Blue vectors are the mean forces calculated for each individual across channel trials of the generalisation phase. Orange vectors are the overall mean of the forces across subjects for each target. **B**. The time course of generated forces in the x-direction of each target's local frame, averaged (±SE) across subjects. Each plot is associated with one target location. Force trajectories are presented for pre-exposure (black line) and generalisation (red line) phase. 43

2.5 Generalisation patterns based on the x component of adaptation across different sets of targets during pre-exposure (blue) and generalisation (orange) phases. Each error bar indicates the overall average of adaptation criterion across subjects and channel trials for each target. The patterns are generated upon the targets on the **A**. transverse plane, **B**. sagittal plane, **C**. diagonal planes, and **D**. the inner and **E**. the outer rings around the central target. . . 46

2.6 Generalisation patterns based on the z component of adaptation for pre-exposure (blue) and generalisation (orange) phases. The panels are the same as in Fig. 2.5. 47

3.1 Experimental Procedure. **A** Subjects performed a top-down reaching movement from a start position (yellow icosahedron) located at the origin of the reference frame (xyz), to a target location 12 cm below the start position on the y axis. **B** The experiment consisted of four phases: pre-exposure phase, Field A (force field in z dimension), Field B (force field in x dimension), and error clamp phase (channel tubes connecting the start position to the target position). 55

3.2 Maximum perpendicular error(MPE). **A** and **B** show the mean ($\pm$SE) time course of MPE across subjects for two conditions ($N_B = 16$ and $N_B = 32$) and separated for x and z components (blue and red lines, respectively). **C** represents the summary of MPE for x (right) and z (left) components in condition 1. The bar plots illustrate the mean MPE across the null trials (Baseline), the first 10 trials of field A (Early A), the last 10 trials of field A (Late A), and the first 10 trials of field B (Field B). **D** illustrates the same bar plots for condition 2. (t-test: $^{*}p < 0.05$, $^{***}p < 0.001$). 59

3.3 The time course of adaptation decomposed to its x (blue) and z (red) components, and shown for condition 1 (**A**) and condition 2 (**B**). The bar plots in **C** and **D** show the summary of adaptive behaviour for the z (left) and x (right) components and for each condition. Each bar plot represents the mean adaptation in the final 10 channel trials of field A (Late A), the last 2 trials of field B (Late B), the first 10 trials, and the last 10 trials of error-clamp phase (Early C and Late C). (paired t-test:$^{*}p < 0.05$, $^{***}p < 0.001$, ns: not significant). 61

3.4 **A** and **C** show the time course of adaptation angle averaged across subjects (circular mean $\pm$ circular standard error) for condition 1 and condition 2, respectively. **B** and **D** also illustrate the adaptation magnitude for each condition. The p-values are shown for the paired t-test between the adaptation levels in field B and the final trials of field A. 63

3.5 Model fit. The model was fit to the x and z components of adaptation for both conditions at once, and then used to predict the angle and the magnitude of adaptation. Different rows illustrates the z component of adaptation (**A**), the x component of adaptation (**B**), the angle of adaptation (**C**), and the magnitude of adaptation (**D**). The parameters of fit were obtained as: $\alpha_f = 0.921$, $\beta_f = 0.092$, $\alpha_s = 0.995$, and $\beta_f = 0.014$. 66

3.6 **A.** An example simulation for the one-dimensional spontaneous recovery paradigm, with fields A and B in opposite directions. **B.** Comparing the simulation results with the superposed data resulted from the subtraction of the x component of adaptation (memory of field B) from the z component of adaptation (memory of field A). 68

4.1 **A.** The 3BOT was used to conduct the experiment. Three-dimensional visual feedback was generated by an Oculus Rift virtual reality system (refresh rate: 75 Hz, resolution of 960×1080 per eye). The task was to perform vertical reaching movements along the *y*-axis from the start position to the target position. **B.** The experiment consisted of a sequence of blocks in which subjects were exposed to an initial force field (A_1), followed by a secondary force field (B), and followed by re-exposure to the initial force field (A_2; the grey background patch indicates the exposure blocks). A washout period of 50 trials (w) separated the exposure blocks. Four different conditions were used in this experiment. In each condition, we manipulated the direction of field B with respect to the initial force field (the small arrows show the direction of each force field in the *xz* plane). In condition 1, field B was in the same direction as the initial field (B=A); in condition 2, field B was in the opposite direction as the initial field (B=-A); in condition 3, field B was orthogonal to the initial field; and in condition 4, field B consisted of null trials (no force field applied). 74

4.2 Adaptation and MPE. The time course of adaptation index (left column) is shown for both inline (orange) and orthogonal (blue) components for each condition. On the right column, the magnitude of MPE is also illustrated over the trials. The small arrows on top of each panel show the relative direction of force field between the blocks A_1, B and A_2. In condition 1, all force fields were in the same direction; in condition 2, field B was in the opposite direction as field A; in condition 3, field B was directed orthogonal to field A; and in condition 4, no force field was introduced in block B. The background grey patch shows the exposure to force field, and the white patch indicates the washout trials. 79

4.3 **A.** The adaptation level prior to the onset of A_1 and A_2 (the last channel trial in the preceding washout period). Before the start of each block, the adaptation level was close to the baseline for all conditions with no significant difference between the blocks. **B.** The final adaptation level (averaged across the last five channel trials) in A_1 and A_2 blocks shown for all conditions. All error bars indicate standard error (SE) across subjects. The *p*-values are obtained by paired t-test between the blocks. 80

4.4 **A.** The inline component of adaptation in blocks A_1 and A_2 shown with the fitting results (solid lines) of the single-rate state-space model. The small arrows on top of each panel indicates the relative direction of force fields in blocks A_1, B, and A_2 (respectively from left to right). **B.** The learning rate parameter obtained from fitting the single-rate model is shown for A_1 and A_2 blocks. The difference between the learning rates of A_2 and A_1 (A_2-A_1) represents the amount of retention of learning. Error bars in all plots indicate the standard error of the mean across subjects. The *p*-values are obtained based on paired t-test between the blocks. 81

5.1 Experimental paradigm. Subjects performed reaching movements back and forth between two targets (T1 and T2) as Out and Return movements. The experiment consisted of three main phases, including Probe 1 (P1), the Training phase, and Probe 2 (P2). Each phase was followed by a washout period, mainly consisted of null trials. The experiment involved five different conditions (G1 to G5) based on the dynamics applied in the Out and Return movements. Three types of task dynamics were used in the experiments, including, exposure (shown by lateral red arrows), error-clamp (shown by grey rectangles) and null dynamics (shown by thin black arrows connecting the two targets). Subjects received rest breaks within each phase shown by the black dots (associated numbers show the trial just before the rest break). 92

5.2 The performance measures for adaptation (left column) and kinematic error (right column) for all groups (G1 to G5). Measures are the average across subjects (with error bars as $\pm SE$) separated for Out (red) and Return (blue) movements. The negative Out adaptation during P2 phase in G4 means adaptation to a reversed (opposite) force field. 98

5.3 **A.** Adaptation level at the end of Probe 1 (P1) for Out direction (last two channel trials; top panel) and Return direction (last 12 channel trials; bottom panel) presented for each group (1 to 5) and for the overall average across all groups (All). The error bars indicate the standard error of the mean (SE) across subjects. **B.** The adaptation level of Return movements at the end of the Training phase (last 5 channel trials; top panel) and after washing out (last channel trial of the washout) prior to the start of P2 (bottom panel). **C.** Adaptation behaviour during the Return trials of P2 presented for each group. The data is binned across the trials (12 trials in each bin) and each data point represents the mean adaptation, first over the trials in each bin, and then across subjects in each group (error bars show $\pm SE$ across subjects). **D.** The initial (first bin) and final (last bin) adaptation level of Probe 2 for the Return direction, compared across groups. 100

5.4 **A.** The allocation of error to the Out and Return processes within each state vector. For a given error in a given context (for example context Out), the error is fully allocated to the corresponding adaptive process (here, x_o), but is only partially used for the alternative process (i.e., x_r) through a coupling factor ($c_k^{(n)}$, $k = 1, 2$). **B.** Different variations of the model can be considered based on: 1. whether the parallel mechanisms have the same or different coupling factors (columns), and 2. the structure of the coupling factor (rows). This results in 9 different models that feature the coupling factor in various ways. **C.** The joint activity of the primitives in the 2-dimensional error space. Each dimension represents the error space in one movement direction (subscript o for Out, and subscript r for Return directions). For a given error in each movement direction, a pattern of activity is formed by the primitives along the corresponding axis ($\mathbf{A}(e_o)$ and $\mathbf{A}(e_r)$). The product of these patterns represent the joint activity of the primitives in each pair of Out and Return trials. **D.** The memory matrix is formed across trials by accumulating the joint activity of the primitives on each pair of consecutive trials. The lighter areas represent stronger connection between the corresponding errors in the Out and Return axes. On a given trial n, the product of the memory matrix and the array of primitives results in a coupling vector. Each element of the coupling vector represents the association that is formed between the error in the current context and all possible errors in the alternative context. The mean value of the coupling vector is a scalar that represents the adaptive component of the coupling factor for that trial. . . . 108

5.5 Results from fitting the model M8 to the adaptation (left column) and MPE (right column). The model was fit to both measures and across all groups at once. The solid lines represent the model fits (red lines for Out movements and blue lines for Return movements) and the data is represented by circular markers and shading error bars. 115

5.6 Left column: the predicted time course of coupling factor (model M8) shown for all five groups (G1 to G5) and separated for Out (red) and Return (black) movements. The coupling factor for Out movements means the fraction of error in Out trials that is used to update the Return process. Likewise, the coupling factor of Return movements is the fraction of error in Return trials used to update the Out process. The error bars represent the 95% confidence intervals obtained by fitting the model M8 to the bootstrapped samples. Right column: the predicted error behaviour shown for the Out and Return trials. . 119

5.7 The memory matrix shown for model M8 in some selected trials of the Training phase for group 1 (G1) and group 5 (G5). The lightened areas show the strength of the connections formed between the errors in Out and Return movements in the 2-dimensional error space. **G1**. The errors are large in the initial trials of the Training phase, resulting in connections that are formed away from the origin. As errors reduce towards zero through adaptation, the connection weights also start to form around the origin. The strength of the connections near the origin (i.e., for small errors) is higher than the connections formed away from the origin (i.e., for large errors) because small errors are more frequently experienced throughout the adaptation. **G2**. Since there was no concurrent adaptation in the Training phase of group 5, no connection was formed between the errors from Out and Return direction. . 120

6.1 **A.** Experimental setup. Subjects grasped the handle of the WristBot and rotated a virtual hammer-like object back and forth between two angular targets. **B.** The object consisted of a mass attached to a stick. Subjects were asked to rotate the object around its axis (grasp point) while trying to prevent the translational movement of the handle. The targets were two rectangle bars separated 40°. The dynamics of the object consisted of a torque (τ) that resisted the rotation, as well as translational forces (centripetal F_C and tangential F_T) that would perturb the handle of the object. Subjects had to learn to compensate for the torque to rotate the object withing the accepted speed, and the perturbing forces to keep the handle stationary during the rotations. **C.** Three types of trials were used during the experiment. In exposure trials (left) subjects experienced both the translational forces and the torque. In null trials (middle) the forces were switched off and only the torque was applied. And, in error-clamp trials the handle of the object was fixed in the home position a using simulated spring that prevented translational movement. Subjects also experienced the inertial torque during error-clamp trials. 128

6.2 Experimental paradigm. **A.** Experiment 1: the experiment consisted of three main phases: P1, the Training phase and P2, each followed by a period of washout (null trials). In P1 and P2, subjects experience exposure trials (yellow patch) in Out movements, and error-clamp trials (grey patch) in Return trials. These phases are followed by washout trials (null; white patch) w.o.1 and w.o.2. During the Training phase, subjects perform under exposure trials in both Out and Return rotations. **B.** Experiment 2: subjects perform 13 blocks of trials with different dynamic manipulations between Out and Return movements. Multiple rest breaks were provided for the subjects (60-90 seconds) during each experiment as shown with the dashed lines. **C.** different combinations of blocks between the Out and Return movements used in experiment 2. **D.** Two groups of subjects participated in this experiment. In group 1, the visual information about the orientation of object was provided, whereas in group 2, the object was ambiguous as to orientation. The dynamics of the task for both groups was the same. 132

6.3 Experiment 1. **A** and **B** show the overall performance of the subjects in terms of adaptation ratio and PD, respectively. The data is shown as the mean across subjects (the blue and red dots for Out and Return movements, respectively), with the error bars (shaded areas) as the standard error of the mean. The solid lines represent the best model fit based on model M2. **C** illustrates the adaptation behaviour during the Return movements of P1 and P2 phases. The solid lines show the single-rate state-space model fit to estimate the rate of change of the adaptation. The adaptation rate is shown in panel **D** with the 95% confidence intervals. Panel **E** also shows the after effects in the form of PD following the P1 and P2 phase (i.e., w.o.1 and w.o.2). The solid lines show the single-rate state-space model fit with the PD reduction rate shown in panel **F** (the error bars show the 95% confidence intervals). 137

6.4 Experiment 2. The adaptation and PD results are shown for visual and ambiguous objects. Data is separated for Out (red) and Return (blue) trials and is shown as the averaged across subjects (small dots) with the error bars as the standard error of the mean. The adaptation data is shown for error-clamp trials, and the PD is shown only for the exposure and null trials. The solid red and blue lines represent the best fit results of model M2. . . . 141

6.5 Experiment 2. **A** The value of the shared coupling factor (fixed component) from Model M2 shown for visual (Vis.) and ambiguous (Amb.) conditions. **B** The coupling factors associated with the fast (left panel) and slow (right panel) processes from model M3. 142

6.6 Experiment 2. The adaptation behaviour in blocks 5, 7 and 12. **A** and **C** show the decay of adaptation in blocks 5 and 7 for the visual and ambiguous conditions, respectively. The solid lines represent the best fit single rate model from equation (6.2). **B** and **D** represent the retention factors resulted from the single rate model fit, with the 95% confidence intervals shown as error bars. **E** illustrates the induced adaptation behaviour in block 12 for both visual (left panel) and ambiguous (right panel) conditions. **F**. The rate of increase in adaptation (learning rate) for each condition. 145

7.1 **A**. Model fits for the single-rate (SRM) and dual-rate (DRM) models with fixed coupling factor. The data is taken from experiment 2 in chapter 6 (visual group). For clarity, the model fits are shown only for the Return trials, though the models were fit to all trials at the same time. **B**. The sum-squared error of fit for SRM and DRM. **C**. The BIC criterion obtained for each model fit. 163

List of tables

5.1 A summary of different variations of the model 113
5.2 Model comparison: the goodness of fit for each model is shown in terms of *BIC*, the sum of squared error of fitting (*SSE*) and the coefficient of determination R^2. For each model, we took the difference between the *BIC* of that model and the *BIC* of the winning model, M8, as reported with ΔBIC. 116
5.3 Best fit parameter values for M8 with 95% confidence intervals. 117

6.1 Experiment 1. The goodness of fit for each model is shown in terms of *BIC*, the sum of squared error of fitting (*SSE*) and the coefficient of determination R^2. For each model, we took the difference between the *BIC* of that model and the *BIC* of the model M2, as reported with ΔBIC. 139
6.2 Experiment 2. The model comparison and the goodness of fit shown for each model and each group (visual and ambiguous). The ΔBIC represents the difference between the *BIC* of each model and the model with the smallest *BIC* (here, M2) as the reference model. 140

Part I

INTRODUCTION

Chapter 1

Introduction

Humans have an extensive capacity to learn a variety of skills and adapt to different environmental or internal (musculoskeletal) alterations. This capacity has been actively explored in sensorimotor learning studies in the past few decades, being addressed from different perspectives, at multiple levels of the sensorimotor hierarchy, and through diverse methodologies. In particular, a large number of studies have focused on sensorimotor learning at the level of behaviour, showing regularities in movement patterns and suggesting organisational principles in learning. Behavioural experimental paradigms have been designed around a variety of tasks from arm movements (e.g., reaching, object manipulation, etc) to saccadic and locomotor adaptation (for review, see Pélisson et al., 2010; Reisman et al., 2010; Wolpert et al., 2011). In this thesis, we focus exclusively on goal-directed arm movements in the study of sensorimotor learning, as they provide insight into the intermediate level of behaviour, encompassing both low-level motor execution and high-level cognition (Krakauer and Mazzoni, 2011).

Typically, motor learning studies of arm movements have pursued two conventional approaches to experimentally examine the modality of learning and adaptation in humans. In the first approach, visuomotor transformations are used to perturb the visual sensory information about the world, while, for instance, performing reaching movements to a target. In this approach, experimenters utilise different techniques, such as prisms (e.g., Harris, 1965; Inoue et al., 2015; Kitazawa et al., 1997) or virtual reality systems (e.g.,

Krakauer et al., 2005, 2000; Mazzoni and Krakauer, 2006; Zhou et al., 2017), to displace or rotate the visual field and measure the subjects' adaptive response to such perturbations. In a second approach, experimenters impose dynamic perturbations on the arm during task performance, and examine how subjects adapt their motor commands to compensate for such perturbations. This is done by having subjects perform reaching tasks inside a rotating room (exposure to coriolis force perturbations; Dizio and Lackner, 1995; DiZio and Lackner, 2000), or while grasping the handle of a robotic manipulandum (exposure to state-dependent force perturbations; Davidson and Wolpert, 2004; Donchin et al., 2003; Ingram et al., 2010; Shadmehr and Mussa-Ivaldi, 1994; Thoroughman and Shadmehr, 2000, and many others). Both approaches have contributed immensely to the understanding of the underlying mechanisms in human sensorimotor learning. In the following, we review some of the important and influential studies in the field, and how they relate to the studies in this thesis.

1.1 Internal models

As an influential theory of motor control, it has been suggested that the central nervous system develops internal models of the body and its surrounding world, based on which it generates appropriate motor commands, and predicts their consequences (Miall and Wolpert, 1996; Wolpert and Kawato, 1998; Wolpert et al., 1995, 1998). Internal models have been considered as pairs of inverse and forward models contributing to the predictive (feedforward) control of actions (Wolpert and Kawato, 1998). The inverse models are neural processes that map the desired movement plans into motor commands, whereas the forward models predict the sensory consequences of executing the motor commands. Evidence for the existence of inverse models can be found in adaptation studies, where the feedforward command gradually changes (adapts) to account for the new perturbations (Kawato, 1999). For example, when reaching under a force field, subjects learn to produce compensatory forces to counteract the perturbations. When the force field is switched off, subjects keep producing the compensatory forces for a few trials afterwards, causing self-generated errors in the opposite direction to the

force field. These self-generated errors, termed after-effects, are the product of feedforward motor commands that are generated by inverse models, and that can adapt to novel conditions.

Evidence for forward models came from psychophysical studies of object manipulation (Flanagan et al., 2003; Flanagan and Wing, 1997). When subjects grasp an object with a precision grip, the grasp force is typically sightly larger than the force required to prevent the object from slipping (which depends on the load force and friction between the object and fingers) . When the object load increases due to a self-generated movement (e.g., rapid movement of the arm), subjects also increase their grasp force simultaneously to prevent slippage. This simultaneous adjustment of the grasp force rules out the involvement of feedback control mechanisms in the response, as such mechanisms rely on delayed sensory feedback. Instead, the adjustments are made by a mechanism that can predict the increase in the load force in advance, and use that prediction to adjust the grasp force even before the sensory feedback is available. This predictive process is conducted by a forward model.

Sensorimotor learning involves adaptation of both forward and inverse models to novel conditions. The rate of adaptation, however, is different for each model. It has been shown that in manipulation of objects with novel dynamics, subjects learn to predict the behaviour of the object much faster than they learn how to control it (Flanagan et al., 2003). This suggests that forward models, which have a predictive role in the motor system, are adapted faster than the inverse models, which involve in the control of the object. The ability to rapidly learn a predictive model seems crucial for the motor system as it stabilises the limb and plays an important role in training the inverse model (Flanagan et al., 2003).

The difference in the learning rate between the forward and inverse models can also be explained by the theoretical representation of these models. Forward models are faster to form or adapt because they represent a unique mapping from a set of motor commands to the sensory outcomes (Jordan, 1996). For example, for a given initial configuration of the arm and a set of motor commands (joint torques), the behaviour of the arm can be uniquely described. Whereas in the inverse models, the mapping from the desired plan to the motor command is not necessarily unique. For instance, holding the index finger in a fixed position in space (desired plan) can be achieved by many different sets of joint angles, each requiring

a different set of motor commands (redundancy in solution). To solve for a single set of motor commands among infinitely possible sets, the motor system might apply optimality principles, which are challenging and computationally demanding (for review, see Todorov, 2004).

1.2 Learning mechanisms

Error-based learning

The sensory predictions made by the internal forward models are essential in detecting environmental or body-related changes (Shergill et al., 2003). The discrepancy between the predictions of the forward model and the actual sensory feedback is referred to as the sensory prediction error, which informs the motor system of a change in the dynamics of the world. How does the motor system respond to the sensory prediction error and utilise them to adapt the internal models? It has been shown that properties such as the error size, the source of error or its reliability majorly affect the way we adapt to or learn new sensorimotor transformations (Berniker and Kording, 2008; Criscimagna-Hemminger et al., 2010; Izawa and Shadmehr, 2011).

The effect of error size on adaptation has been typically examined by comparing abrupt versus gradual exposure to force field perturbations. For instance, it was shown that when errors were small (due to gradual perturbations) adaptation led to longer lasting aftereffects (Hatada et al., 2006), and smaller rate of decay (Huang and Shadmehr, 2009). Also, studying patients with severe degeneration of cerebellum, it was shown that when patients were exposed to sudden force field perturbations (which led to large errors), they showed deficits in adaptation (Rabe et al., 2009). However, when the force field was applied gradually over many trials (causing small errors), significant improvement in adaptation was observed (Criscimagna-Hemminger et al., 2010). It was suggested that large versus small error sizes in force field learning involves different adaptive mechanisms in the motor system. These findings, however, seemed to be specific to dynamic learning, as visuomotor rotation studies

in which patients were exposed to gradual versus abrupt visuomotor transformations failed to show similar effects (Schlerf et al., 2013).

Studies have also examined how adaptation scales with the size of error. In theoretical models of sensorimotor learning, it is usually assumed that adaptation increases linearly as a function of error magnitude (Donchin et al., 2003; Lee and Schweighofer, 2009; Smith et al., 2006; Thoroughman and Shadmehr, 2000). However, existing behavioural data suggests that the relationship between the error size and the extent of learning is nonlinear (Marko et al., 2012; Robinson et al., 2003; Wei and Körding, 2009). For instance, it has been shown that the amount of learning is relatively larger for small errors, and rapidly saturates as the error size increases (Marko et al., 2012; Wei and Körding, 2009). A recent study further proposed that even for a fixed error size the amount of learning could vary depending on the consistency of the environment (Herzfeld et al., 2014). It was shown that when the environment was inconsistent (error direction changed frequently due to random perturbations), the amount of learning from a given error value decreased. Whereas, in a consistent environment (i.e., the perturbation direction was consistent), the same value of error led to larger amount of learning. These results indicated that learning is modulated not only by the size of error, but also by its reliability.

Learning from error may also take place at different rates. There has been a great debate as to what determines the speed with which we learn a sensorimotor transformation (Burge et al., 2008; Gonzalez Castro et al., 2014; van Beers, 2012; Wei and Körding, 2010). A popular approach in this regard has been the use of Bayesian principles, based on which the variation of learning rate is attributed to two different sources of uncertainty in the environment: the feedback uncertainty, and the state uncertainty. For example, in visuomotor rotation studies, it was found that humans learn the task at a slower rate when the visual sensory feedback of the hand location is blurry (feedback uncertainty). In contrast, when the visual feedback is sharp, but systematically deviated from the actual hand location (state estimation uncertainty) learning takes place at a faster rate (Burge et al., 2008; Wei and Körding, 2010). This behaviour could be explained by a Kalman filter model, in which the learning rate was formulated by the Kalman gain as a function of uncertainty in feedback

and state estimation (Burge et al., 2008; van Beers, 2012):

$$K \approx \frac{\sigma_x^2}{(\sigma_x^2 + \sigma_u^2)} \tag{1.1}$$

where, K was the Kalman gain (the learning rate), and σ_x and σ_u represented the uncertainty (noise) in state estimation and visual feedback, respectively. Predictions of the Kalman filter were in agreement with some of the experimental observations, but also failed to explain some others. For instance, in the presence of random perturbation in visual feedback (which increased the feedback uncertainty), Kalman filter predicted slower adaptation rate, whereas no significant change in adaptation rate was observed in human behaviour (Burge et al., 2008).

Recently, Gonzalez Castro et al. (2014) suggested that the rate of adaptation is primarily determined, not by the estimation of state or feedback uncertainty in the current trial (i.e., the Kalman gain in equation 1.1), but by the predictability of the environmental changes in the future (how likely is it that the current experienced perturbation be repeated in the future). Experimental data from force field learning showed that when the perturbations were persistent (predictable), learning rate increased, whereas, under rapidly changing (unpredictable) perturbations, the rate of learning decreased. Similar findings were reported in a study in which the learning rate was attributed to the predictability of error signals in the future (prospective errors; Takiyama et al., 2015). As such, in an environment where errors were consistent and predictable, learning took place at a faster rate than in environments in which error direction changed randomly and frequently. Taken together, these studies shed light on how error signals, as the difference between the predicted and perceived sensory information, are used to adapt internal representation of the body and its surrounding world.

Reward-based learning

Motor learning, besides sensory prediction error, is influenced by motivational feedback such as reward and punishment. Subjects have shown the ability to learn visuomotor rotation tasks purely based on reward-based feedback (success or failure in hitting the target), whether

the rotation was imposed gradually (Izawa and Shadmehr, 2011) or abruptly (Nikooyan and Ahmed, 2015). Reward-based feedback can have a strong complementary effect on learning when provided along with the prediction error. Studies show that when error and reward feedback are both available, the learning process is significantly accelerated (Kojima and Soetedjo, 2017; Nikooyan and Ahmed, 2015; Quattrocchi et al., 2017). In addition, reward has been shown to affect exploratory features of motor variability. In a reaching task towards an invisible rewarded target, subjects increased their reach variability as the probability of reward decreased (Pekny et al., 2015). This was interpreted as a search for target locations which would be rewarded. In contrast, when the average reward increased around a target position, the variability decreased.

Recent studies have also examined the effects of punishment in sensorimotor adaptation. Galea et al. (2015) demonstrated that punishment and reward had distinct effects on the learning process during a visuomotor rotation task. They found that punishment (negative feedback) increased the speed of learning more than reward (positive feedback), but that reward led to larger retention of learning. It was suggested that reward and punishment influenced sensorimotor learning via separate mechanisms. The same group of researchers further studied the effects of reward and punishment on patients with stroke during a force field adaptation task (Quattrocchi et al., 2017). Similarly, they found that reward and punishment both increased the rate of learning in stroke patients to the level that was comparable to healthy age-matched subjects who performed the same task without reward or punishment. Interestingly, the retention of learning in the patients who received reward feedback exceeded that of healthy control subjects.

The learning patterns observed based on reward versus error-based feedback have shown to be fundamentally distinct. For instance, it has been shown that reward-based learning generalises only locally, leads to larger movement variability, and does not update the sensorimotor mapping between the hand and the cursor position (Izawa and Shadmehr, 2011). In a recent study, Cashaback et al. (2017) dissociated the effects of error-based and reward-based feedback in a reaching task. Subjects were exposed to lateral shifts of the cursor with respect to the hand position, where the magnitude of shift was sampled randomly

from a distribution whose mean and mode were separated. It was shown that when provided with reward feedback, subjects learned to compensate for the shifts based on the mode of the distribution; when error-based feedback was provided, subjects learned the task based on the mean of the distribution; and, when both error and reward feedback were available, subjects continued to adapt based on the mean of the distribution. These results suggested that reward-based and error-based learning mechanisms are separate, and that motor learning is predominantly driven by the error-based processes (Cashaback et al., 2017; Izawa and Shadmehr, 2011).

Use-dependent learning

In error-based and reward-based learning, motor commands are adjusted in a direction that would reduce error and increase reward. Recent studies have revealed a new phenomenon in which motor commands have a tendency to generate movements that are similar to previous movements. That is, when reaching movements are made repeatedly in a particular direction, the future movements will be biased towards that direction (Huang et al., 2011; Verstynen and Sabes, 2011). This induced bias towards previously performed movements is called use-dependent plasticity or use-dependent learning. Diedrichsen et al. (2010) demonstrated that use-dependent and error-based learning can simultaneously affect behaviour. Subjects performed reaching movements towards a horizontally elongated target, while their movements were constrained by channel trials (simulated one-dimensional spring walls that guided the movements in a straight line). The channel trials guided the movements in a direction that was slightly deviated from a straight movement subjects normally performed without the channels. This caused an adaptive response by the subjects due to the imposed directional error, such that when the channels were removed, subjects showed marked after-effects in the opposite direction to the channels (error-based adaptation). However, after a few trials, their movements started to converge back towards the initial direction imposed by the channel trials. This tendency towards the repeated movement direction caused by channel trials indicated a use-dependent after-effect.

Use-dependent learning is considered as a model-free process in sensorimotor learning (i.e., it does not depend on the state of the internal models), that has distinct features from the error-based learning. For example, it has been shown that unlike error-based adaptation (Malfait and Ostry, 2004), use-dependent directional biases do not transfer between the limbs (Marinovic et al., 2017). However, it is also shown that use-dependent, error-based and reward-based processes are not totally independent. In a visuomotor rotation task, subjects adapted to visual perturbations towards a training target (forward movement direction), while concurrently experiencing repetitive passive movements (guided by a robot) in a different direction associated with a probe target (backward movement direction). It was shown that the use-dependent biases caused by the passive repeated movements enhanced the generalisation of visuomotor adaptation from the training target to the probe target by 50% (Lei et al., 2017). Similarly, it has also been shown that use-dependent plasticity is modulated by reward-based learning. Mawase et al. (2017) demonstrated that when repeated performance of a motor action (which induces use-dependent plasticity) is joined with reward-based learning of a skill, it leads to stronger motor biases. Collectively, these studies illustrated that separate, yet interacting mechanisms are responsible for error-based, reward-based and use-dependent learning.

1.3 Generalisation of learning

Learning a novel sensorimotor transformation leads to adjustments in the motor system that affects behaviour, not only in the context of learning, but also in nearby contexts or conditions. Understanding the underlying mechanisms of such generalisation is of fundamental importance in answering questions like, what is learned, and how the learning is represented in the motor system. Studies of reaching movements have examined generalisation of learning in tasks such as visuomotor rotation (Brayanov et al., 2012; Ghahramani et al., 1996; Krakauer et al., 2006, 2000; Tanaka et al., 2009; Wu and Smith, 2013), and force field learning (Berniker et al., 2014; Dizio and Lackner, 1995; Donchin et al., 2003; Kadiallah et al., 2012; Malfait et al., 2002; Shadmehr and Mussa-Ivaldi, 1994; Thoroughman

and Shadmehr, 2000). From a different perspective, some studies investigated generalisation of learning on a trial-by-trial basis, in which transfer of learning was tracked as the adaptation developed (Donchin et al., 2003; Francis, 2008; Thoroughman and Shadmehr, 2000). While, others examined the generalisation in a post-adaptation phase, after learning was complete (Berniker et al., 2014; Brayanov et al., 2012; Hwang et al., 2006; Mattar and Ostry, 2007; Shadmehr and Moussavi, 2000; Shadmehr and Mussa-Ivaldi, 1994).

This dissociation highlighted the differences between the adaptation processes underlying visuomotor rotation learning (as kinematic transformations) and force field learning (as dynamic transformations). For example, Tanaka et al. (2012) demonstrated that for visuomotor rotation, both trial-by-trial and post-adaptation approaches led to narrow generalisation patterns across movement directions, the width of which increased by increasing the number of training targets. In contrast, for force field learning, they found wide generalisation patterns that spanned up to 180 degrees from the training direction in the trial-by-trial approach (also, see Donchin et al., 2003; Thoroughman and Shadmehr, 2000), while, in post-adaptation approach, the generalisation remained within 90-degree distance from the training target (also, Mattar and Ostry, 2007).

Generalisation and internal representation of learning

Studies have examined generalisation of learning across various spatial or temporal movement features, such as, movement directions (Berniker et al., 2014; Wu and Smith, 2013), limb configurations (Berniker et al., 2014; Brayanov et al., 2012; Malfait et al., 2005, 2002; Shadmehr and Moussavi, 2000), movement speeds or amplitudes (Francis, 2008; Goodbody and Wolpert, 1998), spatial movement paths (Conditt et al., 1997; Conditt and Mussa-Ivaldi, 1999), or across the arms (Burgess et al., 2007; Criscimagna-Hemminger et al., 2003; Joiner et al., 2013; Malfait and Ostry, 2004; Malfait et al., 2002; Witney and Wolpert, 2003). The outcome of these studies has provided answers to a fundamental question regarding internal representation of movement features in the brain. That is, what is the coordinate system through which the brain views the body and the environment during performing a particular task? In the context of reaching movements, there has been a wide range of studies assessing

coordinate frames such as Cartesian, joint-based, and object-based coordinate systems to describe the internal representation of information in the brain.

Shadmehr and Mussa-Ivaldi (1994) trained subjects to perform reaching tasks under velocity dependent force field within a given workspace, and investigated the transfer of learning to a new workspace by changing the arm configuration (i.e., shoulder angle). They found that the generalisation of learning is better explained within the intrinsic (joint-based) rather than extrinsic (Cartesian) coordinate frame. Subsequent studies, however, provided rather conflicting evidence regarding intrinsic versus extrinsic representations in the sensorimotor system. Studies that examined generalisation of adaptation across arm configurations found evidence in favour of intrinsic representation of generalisation (Malfait et al., 2005, 2002; Shadmehr and Moussavi, 2000). On the contrary, studies regarding inter-limb generalisation (transfer of adaptation to the contralateral arm) provided evidence in support of extrinsic representation of generalisation (Burgess et al., 2007; Criscimagna-Hemminger et al., 2003; Malfait and Ostry, 2004).

There have also been studies postulating that the internal representation of movement in the internal model is neither exclusively intrinsic nor extrinsic. For instance, Bays and Wolpert (2006) designed a bimanual reaching task, in which movements of the right hand caused force field perturbations acting on the left hand. They found interesting dissociation of kinematic and dynamic representation of learning, in which the trajectories of the right hand were represented in the extrinsic coordinates, whereas the forces produced by the left hand were better fit with the intrinsic coordinates. In an object manipulation study, Ahmed et al. (2008) suggested that objects with simpler dynamics are, to a great extent, represented in object-centred coordinates, whereas the objects with more complex dynamics are represented in arm-centred coordinates. They also reported that in the absence of visual feedback of the object an intermediate coordinate system better accounts for representation of simple object dynamics.

In a recent study, Brayanov et al. (2012) developed a computational framework to describe the internal representation of learning during visuomotor rotation task. They trained subjects in one movement direction within a given workspace, W1, and tested the generalisation

of learning to other directions both in W1 and in a novel workspace, W2. After obtaining generalisation patterns for both W1 and W2, they examined the ability of four Gaussian based models to fit the data, including, fully intrinsic, fully extrinsic, additive intrinsic-extrinsic, and composite intrinsic-extrinsic models. They showed that the transfer of learning from one workspace to the other could not be described by fully extrinsic or fully intrinsic coordinate frames. Instead, it was best accounted for by a composite intrinsic-extrinsic model, wherein the combined distance across intrinsic-extrinsic space determined the amount of generalisation (equation 1.2).

$$\begin{aligned} z(\theta_E, \theta_I) &= k \cdot e^{-D^2/2\sigma^2} \\ D &= \sqrt{s_E(\theta_E - \theta_{E0})^2 + s_I(\theta_I - \theta_{I0})^2} \end{aligned} \tag{1.2}$$

where, z was the generalisation quantity, θ_E and θ_I were the direction of movement in extrinsic and intrinsic coordinates, respectively, θ_{E0} and θ_{I0} were the training direction represented, respectively, in extrinsic and intrinsic coordinates, and k, s_I, s_E and σ were the free parameters. The proposed model predicted a unimodal generalisation pattern that could capture the generalisation peak and rate of decay across the movement directions in the test workspace, W2.

In a most recent study, Berniker et al. (2014) used a similar approach to explore the coordinate systems responsible for the generalisation of learning in force field adaptation tasks. They reassessed the previous findings regarding the intrinsic representation of generalisation observed in force field experiments (e.g., Malfait et al., 2005, 2002; Shadmehr and Mussa-Ivaldi, 1994). They argued that previous work only examined generalisation across changes in the shoulder angle, whereas, if intrinsic representation of generalisation were the case, changes to other joint angles would have elicited the same generalisation patterns. Consequently, they assessed the generalisation of learning across the total of 15 arm configurations involving changes in shoulder, elbow and wrist angles. The generalisation patterns were not accounted for by any of the joint-based, Cartesian or object-centred coordinates, individually. Rather, a mixture of coordinate frames with local learning properties explained the generalisation patterns across different arm configurations. The authors also attempted

to reproduce the findings reported in the original study by Shadmehr and Mussa-Ivaldi (1994). Surprisingly, despite using the identical experimental paradigm, the reproduced results failed to support the original findings: no clear evidence was found in support of intrinsic representation of dynamics as reported by Shadmehr and Mussa-Ivaldi (1994).

Overall, these studies indicated that the motor system does not use a single representation, but instead has access to multiple coordinate frames that can be combined and used in different tasks and contexts. It is not clear, however, what determines the weighting of such combinations.

1.4 Learning and memory

When a sensorimotor transformation is learned, various mechanisms are involved in the memory formation, retention, and recall of the learned motor commands. In this section, we review some of the experimental and theoretical work that has investigated different properties of motor memories during and after learning. In particular, we focus on the phenomena of savings, interference, and spontaneous recovery.

Savings

Savings is the phenomenon in which relearning a task after it was forgotten occurs at a faster rate. (Kojima et al., 2004; Mawase et al., 2014; Zarahn et al., 2008). This is experimentally shown when subjects learn a visuomotor or force field transformation on day 1, and are then tested for relearning after several hours, days or even months (e.g., Bock et al., 2001; Caithness et al., 2004; Krakauer et al., 2005). Similarly, it can be shown in a short-term learning process when subjects adapt to a sensorimotor transformation (task A) followed by washout of learning (forgetting), and followed by re-exposure to the same transformation (e.g., Kojima et al., 2004; Smith et al., 2006; Zarahn et al., 2008). In both scenarios, subjects showed faster or more complete learning in their second encounter to the task.

Why savings occurs and what mechanisms are involved in this phenomenon has been hotly debated in the sensorimotor learning literature. Early studies of savings proposed that

savings occurs due to the interaction of two adaptive processes that operate in parallel, but with different time-scales (Joiner and Smith, 2008; Smith et al., 2006; Zarahn et al., 2008). One process learns quickly and forgets quickly (fast process), and the other learns slowly but also forgets slowly (slow process). When a sensorimotor transformation is learned, both processes contribute to the adaptation, with the slow process taking over in the later stages of training. During the washout of adaptation (when the visuomotor or force field exposure is switched off), the fast process rapidly unlearns the task, while the slow process retains some learning (due to gradual forgetting). The net outcome of the two processes, however, emerges as an apparent forgetting of the task. Critically, when the transformation is re-exposed, the retained learning in the slow process causes faster relearning (savings). This theoretical approach, although accounting for savings in short term memory formation, can not explain savings observed over longer periods after which the adaptation was unlearned (Zarahn et al., 2008).

Subsequent studies provided conflicting evidence as to the underlying mechanisms of savings. One group of studies suggested that savings was attributed to error-based learning processes (Berniker and Kording, 2011; Herzfeld et al., 2014; Leow et al., 2016; Takiyama et al., 2015). As such, when an internal model is formed or modified during error-based adaptation, it can be stored and quickly recalled when similar errors are encountered in the future. In this case, the motor system might keep a memory of previously experienced errors (Herzfeld et al., 2014) or form an estimation of errors that might occur in the future (Takiyama et al., 2015), to quickly recall appropriate motor commands.

On the other hand, observations from other studies examining savings in visuomotor rotation paradigms have shown inconsistencies with the error-based representation of learning. For example, Haith et al. (2015) showed that when subjects adapted to a visuomotor rotation under limited preparation time, the adaptation rate was slow and did not show savings when retested on the next day. Whereas, when subjects had ample time to prepare for each movement, adaptation was faster and savings was observed. It was argued that increasing preparation time allowed explicit or strategic processes to come into play, which were independent of error-based mechanisms. In a different study, Huang et al. (2011) dissociated

the error-based adaptation and model-free (use-dependent) learning in a visuomotor rotation task, and showed that savings was attributed to the memory of previous successful actions formed based on movement repetitions (use-dependent mechanism), rather than error-based adaptation. Similarly, Morehead et al. (2015) argued that savings was achieved through the recall of aiming strategies learned in the previous encounter to the visuomotor rotation, even when the true error-based adaptation was absent. In a recent study, Leow et al. (2016) also examined the effects of previous successful actions and error-based adaptation on savings. Contrary to the studies mentioned above, they proposed that the history of errors, as opposed to the history of successful actions, is necessary and sufficient for savings.

Overall, although substantial amount of experimental and theoretical work has been done towards understanding of saving phenomenon, a unifying framework to reconcile the existing conflicting results is yet to be found.

Interference

Interference in the context of sensorimotor learning refers to the phenomenon in which the motor memories of two particular tasks interfere with the formation or modification of one another. In reaching movements, this phenomenon can be experimentally examined using two types of paradigms. In the first paradigm, one learns a reaching task under two successive force fields, A then B, in which B is usually an opposite field to A (B=$-$A). In this case, the question is how adaptation to field B, after having learned A, differs from adaptation to field B without any prior learning of A. Observations showed that field B was learned remarkably slowly when preceded by field A, than when learned without any prior learning of A (Brashers-Krug et al., 1996; Shadmehr and Brashers-Krug, 1997; Shadmehr and Holcomb, 1999). This type of interference is termed anterograde, or proactive interference, which refers to the effects of initially learned motor memory (A) on the formation of a new motor memory (B). It has been shown that increasing the number of training trials in field A increases the amount of interference in the learning of field B (Sing and Smith, 2010). However, the amount of interference asymptotes in cases where the number of trials in A exceeds 40.

A second paradigm examines the relearning of field A after the sequence of A-B was performed. An A-B-A schedule is normally used in which subjects perform under field A, followed by field B, and then back to field A (Brashers-Krug et al., 1996; Caithness et al., 2004; Pekny et al., 2011; Shadmehr and Brashers-Krug, 1997; Tong et al., 2002). In this case, the adaptation behaviour is compared between the first and second exposure to field A (i.e., before, and after the learning of B). It has been shown that re-adaptation to field A after being exposed to field B takes place in a slower rate compared to the initial adaptation to field A. This type of interference is termed retrograde or retroactive interference, which refers to the damage caused by the learning of the second force field (B) on the memory of the initial force field (A).

It was initially proposed that interference is caused due to the limitations in the capacity of short-term working memory (Brashers-Krug et al., 1996; Shadmehr and Brashers-Krug, 1997; Shadmehr and Holcomb, 1999). As such, when the working memory is occupied by the learning of the initial transformation (e.g., field A), attempting to learn a second transformation not only slows down the learning process for the latter (anterograde interference), but also disrupts the memory of the former (retrograde interference). Later, Bock et al. (2001) demonstrated that competing for a limited short-term memory was not the main factor for interference; subjects were able to learn two transformations that were seemingly independent, simultaneously. Instead, when two transformations required incompatible adjustments in the motor commands (e.g., opposing force fields) the interference was observed.

Studies have also highlighted the effects of contextual cues on the level of interference (Howard et al., 2013). When two opposing transformations are associated with a single context (i.e., reaching to a single target), the corresponding motor memories interfere with one another. However, when opposite transformations are associated with different contexts, the interference can be reduced. For example, it has been shown that two opposing force fields can be learned simultaneously if each is associated with a different workspace location, either visually or proprioceptively, or object orientation (Howard et al., 2013). In fact, such contextual cues encourage the brain to allocate separate motor memories to the learning of each force field.

Other studies have also demonstrated the effects of the temporal state of the arm, right before or after the reaching movement, on the reduction of interference. It has been shown that subjects can learn opposing force fields in a single reach direction, if each force field is cued with a different lead-in (Howard et al., 2012) or follow-through (Howard et al., 2015) movement. Recently, studies on visuomotor rotation and force field perturbation (Hirashima and Nozaki, 2012; Sheahan et al., 2016) have demonstrated that motor memories of two opposing transformations can be separated by associating different movement plans to each one, even when the actual performed movements were identical. Interestingly, Sheahan et al. (2016) showed that subjects were able to reduce interference in learning two opposing force fields, solely based on *planning* different follow-through movements associated with each force field, and without the need to actually execute the planned follow-through. Overall, these studies provided evidence as to how the sensorimotor system forms separate representations for learning.

Spontaneous recovery

Spontaneous recovery is a basic phenomenon of Pavlovian conditioning, based on which a previously extinct (forgotten) memory re-emerges spontaneously (Myers and Davis, 2002; Rescorla, 2004; Robbins, 1990). In memory research, a typical experimental paradigm to encourage spontaneous recovery consists of three phases: the learning phase, in which the animal learns an association between a stimulus and a response; the extinction phase, in which the animal unlearns the association between the stimulus and response (stimulus is no longer followed by a response); and the post-extinction phase, in which the animal spontaneously recalls the previously forgotten association, such that the presentation of the stimulus induces a response.

In recent years, researchers have observed similar behaviour in sensorimotor learning (Criscimagna-Hemminger and Shadmehr, 2008; Kojima et al., 2004; Smith et al., 2006). For instance, Kojima et al. (2004) demonstrated in nonhuman primates that when saccadic adaptation (repeatedly displacing a target during a saccade to change the gain of the initial saccade) is followed by a rapid deadaptation, in the following phase of complete darkness

(no feedback of saccadic error) animals re-exhibited their initially learned adaptive behaviour. Later, Smith et al. (2006) reproduced the phenomenon of spontaneous recovery in force field learning (Fig. 1.1). In this study, learning of a force field perturbation (adaptation) was followed by a short period of reversed force field exposure (deadaptation), and then a long period of error clamp trials (trials in which the movement was constrained in a straight line using a simulated channel; Fig. 1.1). It was shown that although adaptation to initial force field was reduced to baseline by the end of deadaptation, subjects exhibited a sudden increase in adaptation towards the initial learning, gradually decaying to the baseline in the following trials. They proposed that this behaviour was the outcome of two interactive adaptive processes in the form of a linear time-independent state-space model: one process that learned quickly and forgot quickly, and the other that was slow to learn but also forgot slowly:

$$\begin{aligned} x_f^{n+1} &= A_f x_f^n + B_f e^n \\ x_s^{n+1} &= A_s x_s^n + B_s e^n \\ x^n &= x_f^n + x_s^n \\ B_f &> B_s \\ A_f &< A_s \end{aligned} \tag{1.3}$$

where, the net adaptation, x^n, was the sum of the fast process (x_f^n) and the slow process (x_s^n) on a given trial n. Also, the parameters A_s and A_f represented the retention factors, and B_s and B_f represented the learning rates of the slow (subscript s) and fast (subscript f) processes. The adaptation was driven by the error signal $e^n = F - x^n$, as the difference between the external perturbation F and the adaptation state x^n. Fig. 1.1 illustrates how fast and slow processes contribute to the adaptation behaviour and the observed spontaneous recovery.

Recently, an alternative theoretical approach based on prospective errors has been proposed to account for the phenomenon of spontaneous recovery (Takiyama et al., 2015). According to this approach, spontaneous recovery is caused by the interaction of motor

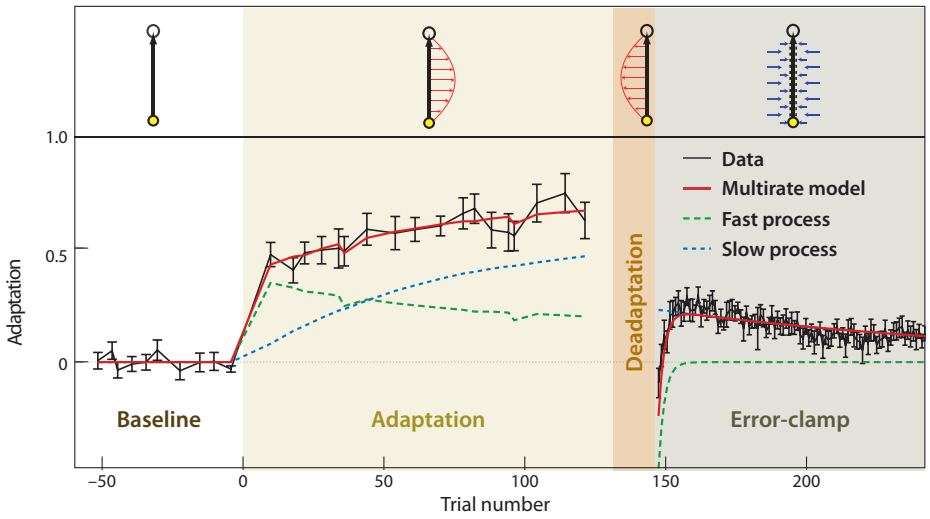

Fig. 1.1 Spontaneous recovery induced by adaptation to two successive opposing force fields (Smith et al., 2006). A dual-rate model with slow and fast adaptation processes accounted for the empirical data.

primitives that encode error signals. As such, during the initial adaptation, the primitives associated with the experienced errors (e.g., in the positive direction) are activated, and their activity strengthens as the training continues (leading to increase in adaptation in the positive direction). During the deadaptation phase, errors change sign (e.g., to negative direction) and thus the primitives that encode negative errors are activated. At the end of this phase, the activity of all primitives (those that encode positive and negative errors) sums to zero, indicating full deadaptation. However, during the error-clamp phase, the activity of primitives that encode negative errors decays quickly (due to shorter training) while the activity of primitives encoding positive errors lasts longer (due to longer training in the initial adaptation). The sum of these primitives, therefore, leads to a rebound of positive adaptation (i.e., spontaneous recovery).

Although both the dual-rate and the prospective error models accounted for spontaneous recovery, they differed in that the dual-rate model assumed a fixed value for retention factor and learning rate in the motor system (i.e., A and B in equation 1.3), whereas, in the prospective error model, these parameters were temporally dynamic and changed depending on the dynamics of the environment. This flexibility in the prospective error model provided

an explanation of adaptation behaviour under frequently changing environments (Takiyama et al., 2015).

1.5 The organisation of the dissertation

In this thesis, we focus on the representation and interaction of error-based sensorimotor learning processes. The thesis is organised in two experimental parts. In the first part, we examine the representation of learning in three-dimensional (3D) reaching tasks, and extend the notions of generalisation, interference, and spontaneous recovery to 3D movements. In the second part, we examine the interaction between multiple context-dependent sensorimotor memories, and propose a theoretical model to describe this interaction.

1.5.1 Representation of learning in 3D environment

The 3D robotic manipulandum (3BOT)

Sensorimotor learning studies mainly rely on experimental measurements of human performance in the laboratory environment. Robotic devices, in this case, are of essential apparatus that are widely used to carefully manipulate the visual or dynamic stimuli in the task environment, and measure various movement features in response to such stimuli. One of the widely utilised robotic devices in motor control research is the planar robotic manipulandum (Adelstein, 1989; de Vlugt et al., 2003; Fayé, 1986; Howard et al., 2009). These robots are mainly used in arm movement paradigms and range in studies from impedance control of the limb (Burdet et al., 2001; Franklin and Milner, 2003) to learning novel dynamics (Shadmehr and Mussa-Ivaldi, 1994) or virtual object manipulation (Ingram et al., 2010). In particular, planar robots are very common in adaptation studies in which subjects learn to perform reaching movements under artificially generated force fields.

These studies typically involve subjects grasping the handle of the manipulandum while performing a reaching task from a given start position to a target position (Fig. 1.2A). During the movement, subjects are exposed to perturbing forces applied to their hand by

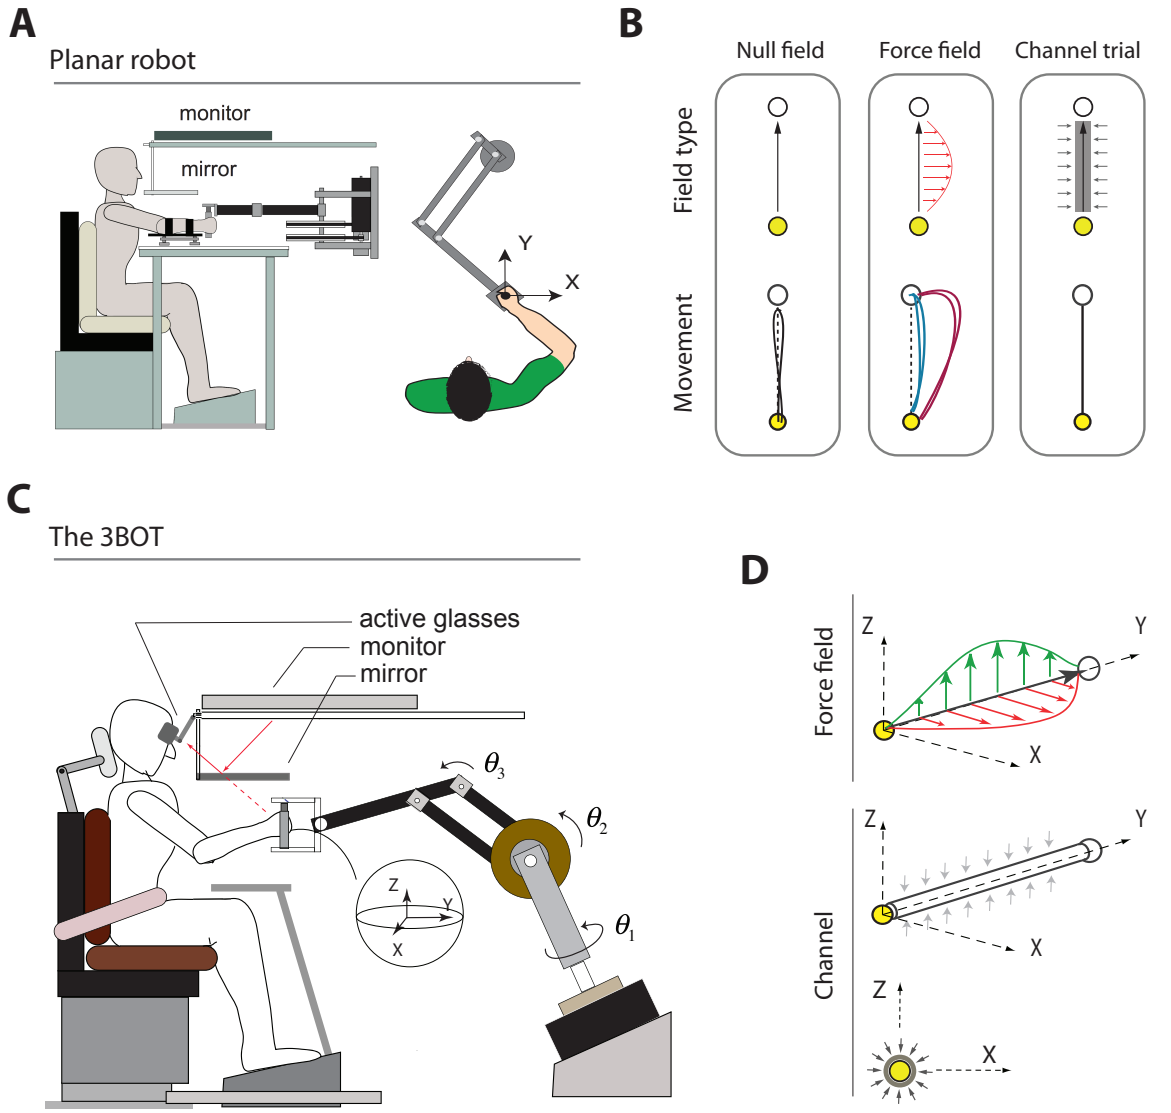

Fig. 1.2 **A**. The planar manipulandum. **B**. Three types of trials can be implemented by the planar robot: the null trials in which the arm can freely move in the workspace; the force field trials in which a velocity/position dependent perturbation is applied causing large deviations in the hand trajectory; and the channel trials in which the forces generated by subjects are measured. **C**. The 3D capable robotic manipulandum (3BOT). The 3BOT is a serial robotic arm with three degrees of freedom in the joint space (θ_1, θ_2 and θ_3) that can control translational movements of the end-effector in 3D. The virtual reality environment is provided by a 3D display paired with a mirror, and a pair of 3D active glasses. **D**. The 3BOT can apply force fields in different spatial directions. It can also implement channel tubes that measure subject's generated forces orthogonal to the movement direction (i.e., the X-Z plane).

the manipulandum, pushing their arm away from the target and causing large deviations from a normal straight movement. It is shown that subjects gradually learn to reduce the trajectory deviations by producing compensatory forces against the perturbation (Shadmehr and Moussavi, 2000; Shadmehr and Mussa-Ivaldi, 1994; Thoroughman and Shadmehr, 2000). Forces produced by subjects can be measured during the movement using a technique called channel trials (Scheidt et al., 2000). In this technique, a virtual mechanical constraint is implemented by the manipulandum that only allows the movement in one dimension (that is, from the start position to the target position) and prevents lateral movements (Fig. 1.2A). The constraint is applied by simulating a stiff spring that acts orthogonal to the target direction. When subjects produce a force against the channel wall, the force can be measured based on the displacement of the hand in the lateral dimension and the stiffness of the virtual spring. The measured force is then compared to the perturbation force as a proxy of adaptation to the force field perturbation.

This setup has been the basis for many studies in sensorimotor adaptation and has contributed to our understanding of human motor control and learning. However, the limitations associated with manipulation and measurement of movement in a two-dimensional (2D) environment has prevented researchers from studying more naturalistic movements. These limitations have encouraged researchers to develop robotic devices with the ability to perform and measure movements in the 3D space. In this case, the Phantom robots originally designed by Massie et al. (1994) are well known for their application in 3D object manipulation studies. Another 3D robotic device is the exoskeleton arm designed by Mistry et al. (2005) which can be used to directly measure joint angles and joint torques from human arm movements, and impose perturbations in the joint space.

Here, we introduce a novel three-dimensional robotic manipulandum (3BOT) developed in the Computational and Biological Learning Lab at the University of Cambridge (schematically shown in Fig. 1.2C). The 3BOT is a serial robotic arm with three degrees of freedom in the joint space that allows for 3D translational movements of the end-effector. Three DC motors are used for the control of each joint, resulting in a full control over the end-effector movement. Each joint is also equipped with an optical encoder that measures the

position of each joint angle at 2 KHz, which is then translated to the 3D position and velocity of the end-effector. The virtual reality environment is provided by a 3D display mounted horizontally (face down) on top of the workspace (Fig. 1.2C). The visual stimuli produced by the display is reflected by a mirror installed in parallel to the monitor and underneath it. Subjects receive the stimuli from the mirror through a pair of active 3D glasses that are fixed in an appropriate distance from the mirror. Visual stimuli in the virtual reality workspace were updated at 60 Hz.

The 3BOT is programmed to counterbalance the weight of the robotic arm at all times during an experiment, such that subjects could easily move the arm without carrying the weight of the robot. The 3BOT can also produce variety of force fields during a movement. These force fields can be applied in different directions in the 3D work space, and could vary in both magnitude and direction as a function of movement states (i.e., position and/or velocity; Fig. 1.2D, top panel). The performance measurements in the 3BOT can be performed at both the kinematic level (i.e., position and velocity of movement) and the dynamic level (generated forces by subjects). The kinematic measures are collected by the encoders in the robot joints, while the force measures are obtained using the channel trial technique that is extended to the 3D workspace. In channel trials, we simulate a stiff tube that only allows the movement in one direction (Y-axis; Fig. 1.2D, bottom panel) and can measure the force vector anywhere orthogonal to the movement direction (X-Z plane; Fig. 1.2D). This provides an additional degree of freedom in examining the adaptation and learning. As such, for a given force field direction we can measure how subjects adapt to both the magnitude and the direction of the force field. Note that in planar robots, adaptation could only be evaluated for the magnitude of the force field as the channel trials could only measure the force in one dimension.

Overall, the unique features of the 3BOT allow us to study more naturalistic movements without the severe constraints imposed by the planar robots.

The organisation of chapters

In the first experimental part of the thesis, we use the 3BOT to create a 3D virtual reality environment, in which we manipulate the dynamics of the movements and measure the subjects' adaptive response to such manipulations.

In chapter 2, we examine the generalisation of learning across various reach directions in a force field adaptation task. Specifically, we focus on the issue of coordinate frame representation in the motor system, and illustrate how this issue makes the interpretation of generalisation patterns obtained in previous studies so difficult. We demonstrate that generalisation patterns obtained across reach directions in a 3D environment provide insight as to the coordinate frame representations used by the motor system.

In chapter 3, we study the phenomenon of spontaneous recovery. In previous studies, the typical paradigm for spontaneous recovery consisted of adaptation to an initial force field (field A), followed by a short introduction of an opposing force field (field B), and proceeding to the error-clamp trials (Smith et al., 2006). In the current study, we use the similar paradigm, in which fields A and B are introduced in orthogonal dimensions. In this case, we demonstrate a new form of spontaneous recovery in 3D, and highlight the contribution of each memory (i.e., field A and B) to the the observed phenomenon.

Finally, in chapter 4 we examine how the motor memories associated with different force field directions interact. In this chapter, we used the 3BOT to generate force fields in various directions in a 3D environment, and measured the effects of savings and interference when the force fields were in the same direction, in opposite directions or in orthogonal directions. We demonstrate that the interaction of motor memories formed based on different force field directions is tuned to the relative angle between the force fields.

1.5.2 Interaction between the learning processes

In the second part of the thesis we investigated how motor memories of a dynamic task formed in association with different contexts interacted with one another. Previous studies have extensively examined the context dependent properties of motor memories, illustrating

how contexts (e.g., movement direction, workspace location, background colour, etc) affect the allocation of motor memories to different tasks (e.g., Howard et al., 2012, 2013). However, very few have examined the interaction of motor memories associated with different contexts. Theoretical studies of motor learning that aim to describe multi-context adaptation mostly assume a fixed interaction between context-dependent adaptive processes (Donchin et al., 2003; Ingram et al., 2013; Kim et al., 2015b; Lee and Schweighofer, 2009). This is normally done by assuming a fixed tuning function of adaptive memories across different contexts. As such, when a given context is presented, the adaptive memory associated with that context is updated the most, and other adaptive memories are updated to a lesser extent depending on the distance between their corresponding context and the presented context.

In the current study, we look more closely into the interaction of context-dependent motor memories. In chapter 5, we show that when a novel dynamic task is learned under two distinct contexts, the interaction between the associated motor memories can increase or decrease depending on the environmental condition. This behaviour is accounted for by a novel state-space model that features an adaptive coupling factor to capture the dynamics of the interaction between two motor memories. In chapter 6, we examine the properties of such interaction when subjects learn a task with a familiar dynamics (as opposed to a novel dynamics) under two different contexts. We illustrate that in the case of familiar dynamics, the learned motor memories exhibit a consistent interaction.

Part II

SENSORIMOTOR LEARNING IN 3D ENVIRONMENTS

Chapter 2

Spatial generalisation of adaptation in 3D

2.1 Introduction

A fundamental question in sensorimotor control is how the brain generalises an acquired skill or experience from one context to another. This is particularly important for the brain to be able to perform as well in similar situations without having to learn specific sensorimotor commands for every single state of the body or environment. Generalisation has been studied in various ways in sensorimotor control, from perception to action (for example, Liu and Weinshall, 1999; Roach et al., 2017; Sarwary et al., 2015; Shadmehr, 2004). In particular, there are many studies investigating the extent to which learning a new skill or a new sensorimotor mapping in one context transfers to others. In this case, various experimental paradigms, such as force field learning, visuomotor adaptation, or object manipulation are used to examine the generalisation behaviour in various contexts (Donchin et al., 2003; Ingram et al., 2010; Wu and Smith, 2013). A common aspect of generalisation in such paradigms has been the tuning of a learned dynamics to movement direction in a 2D environment. For example, many studies have looked into how learning a force field in a particular movement direction generalises to other directions (Berniker et al., 2014; Donchin et al., 2003; Mattar and Ostry, 2007; Mussa-Ivaldi et al., 1994; Shadmehr, 2004;

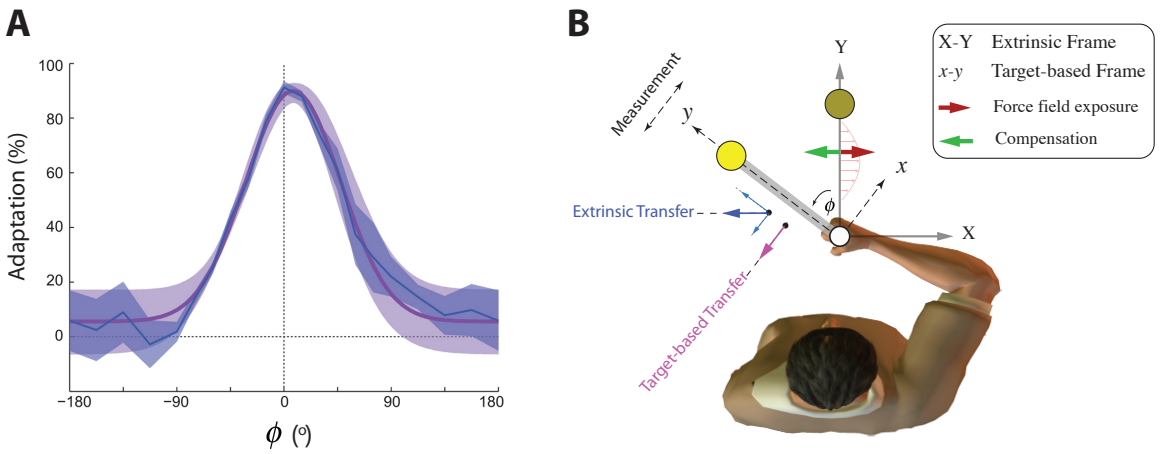

Fig. 2.1 **A**. The pattern of generalisation to different target directions (data from Howard and Franklin (2015)). Adaptation is maximum at the training target ($\phi = 0°$) and decays for targets farther away. **B**. Schematic illustration of the transfer of force field learning from one movement direction (Y axis of the main frame) to another (y axis of the target-based frame). The learned adaptive force in the training direction (green arrow) can be transferred to the new orientation (ϕ) in two possible frame representations: the target-based frame (purple arrow) or the extrinsic frame (blue arrow). The force measurement by channel trials is limited to the lateral dimension of the movement direction (that is, the x axis of the target-based frame). Thus, if learning is transferred in the extrinsic frame (blue arrow) only its lateral component can be measured by channel trials.

Thoroughman and Shadmehr, 2000). It has been shown that subjects exhibit the maximum adaptation in the training direction, and their adaptive response decays, typically in a Gaussian manner, as the movement direction deviates from the training direction (Fig. 2.1A; Brayanov et al., 2012; Howard and Franklin, 2015). This typical decaying generalisation is also observed in other contexts such as different movement speeds or target distances (Goodbody and Wolpert, 1998; Mattar and Ostry, 2010).

Most studies on generalisation have been conducted by the use of planar robots, which helped to obtain the tuning curves of adaptation transfer over different movement directions, or arm configurations on the horizontal plane (Berniker et al., 2014; Howard and Franklin, 2015). The interpretation of such tuning patterns, however, has not been easy as two major factors were involved in obtaining these patterns. First, the internal representation of the coordinate systems in which the generalisation occurs is still an open question (Berniker et al., 2014; McDougle et al., 2017). For instance, when force field adaptation is transferred

from one direction to another, it could be transferred in an extrinsic world based frame, or in an intrinsic target-based frame, or a mixture of two (Fig. 2.1B). Each of these possibilities suggests a different way of interpretation of the obtained tuning curves. Second, the measurement of adaptation in planar robots is limited to a single dimension (that is, the lateral dimension of the movement direction), which may not match the actual direction in which generalisation happens (the direction of generalisation is not clear due to unknown coordinate frame of representation). For example, in Fig. 2.1B, two different coordinate frames are compared for the transfer of learning from one target direction to another. As shown, for the target-based frame (*x*-*y* frame), the direction of generalisation (purple arrow) is orthogonal to the direction of movement, and thus matches the direction of measurement by channel trials. For the extrinsic frame (XY frame), however, the direction of transfer (blue arrow) is aligned with the direction of compensation in the original target direction (the green arrow). In this case, only the lateral component of the transferred force can be measured (that is, the component that is perpendicular to the movement direction), and thus the generalisation cannot be fully examined.

In this study, we used the 3BOT to address this issue by examining the target-based versus extrinsic (world-based) transfer of learning. Subjects adapted to a velocity dependent force field while performing reaching movements towards a training target at the intersection between the sagittal and transverse planes (Fig. 2.2A). We probed the transfer of learning when subjects reached to several targets distributed in the 3D space with different reach angles compared to the training target (Fig. 2.2A). The 3D distribution of probe targets was designed such that for one group of targets, both extrinsic and target-based representations predicted the same pattern of generalisation. Whereas, for another group of targets, different representations predicted different generalisation patterns. Comparing the measured transfer of learning between these groups determined the coordinate frame in which the generalisation had occurred.

2.2 Methods

This study was the first experiment to be designed and conducted by the 3BOT. The preparation of the 3BOT before being used in our experiments required several months of fine-tuning and validation of the generated and measured quantities, including the end-effector forces and the kinematic measures of movement. Particularly, a great deal of work was done to optimise the stability of the 3BOT during force measurements on channel trials. Here, we present the first experimental work and the first results obtained using the 3BOT.

Eight subjects were recruited to take part in this experiment (4 male, 4 female; mean age ±SD: 25.9±1.5 years). All subjects were right-handed according to the Edinburgh handedness inventory (Oldfield, 1971), with no reported neurological disorders. The experiment was approved by the institutional ethics committee, and subject provided informed consent before participation. A short diagnostic paradigm was conducted before the onset of the experiment to test for the subjects' ability to perceive the virtual workspace through the 3D display and active glasses. Subjects reported no difficulties in 3D perception of the task.

2.2.1 Experimental procedure

Subjects were seated on a sturdy chair and grasped the handle of the 3BOT (Fig. 1.2C). The visual feedback of the virtual environment included a cursor (small red sphere, 0.5 cm radius) representing the veridical hand position in the 3D workspace; a fixed starting position (yellow icosahedron, 1.25 cm radius) at the centre of the workspace; and 13 different target positions distributed around the starting point with the distance of 12 cm, located on the sagittal, transverse and diagonal planes (Fig. 2.2A). Targets on each plane were distributed with the angular distance of 22.5° with respect to the start position.

On each trial, the task was to initially maintain the cursor in the displayed start (home) position until one of the targets appeared. The movement was cued by a tone, after which subjects made a rapid reaching movement towards the displayed target. The movement finished when the cursor was stationary within the target for 100 ms (positional tolerance 0.75 cm). At the end of the movement, the robot passively returned the hand back to the

home position for the next trial. Subjects were required to finish the movement within 850 ± 250 ms, otherwise they received appropriate message indicating "too slow" or "too fast".

Three types of trials were used during the experiment. On null trials, subjects could freely move without any interference from the robot. On exposure trials, a velocity dependent force field was applied to the subject's hand while making the reaching movement. The force field was generated as shown in equation 2.1:

$$\mathbf{F} = b \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{bmatrix} \tag{2.1}$$

where, $\dot{x}$, $\dot{y}$ and $\dot{z}$ represented the components of the hand velocity vector in Cartesian coordinate frame (the X-Y-Z coordinate frame; Fig. 2.2A), and b was the field constant, which was either $b = 15$ Ns/m for clockwise force field, or $b = -15$ Ns/m for counter-clockwise force field (counter balanced between the subjects). Note that the force field was only applied in the transverse plane (X-Y plane; Fig. 2.2A) and was independent of the Z-component of the velocity. This was chosen so that we could compare the generalisation of learning across the targets in the transverse plane, where both movement direction and the force field direction were placed on the same plane (similar to previous studies with planar reaching movements), and the targets across the sagittal plane, where the movement direction was orthogonal to the force field direction (this was done for the first time in this study).

Finally, we used channel trials to assess the state of learning and generalisation to different targets. Channels were simulated as stiff tubes allowing movement only in one dimension (along a straight line from the home position to the target position), while measuring the exerted forces in the lateral dimensions. A linear spring model was used for the simulation of channel trials with the stiffness factor of 40 N/cm. On each channel trial, we could measure both the magnitude and the direction of forces produced by subjects in compensation for the force field perturbation.

Before the onset of the experiment, subjects went through a familiarisation session in which they made a few reaching movements under null trials towards all targets to get familiar

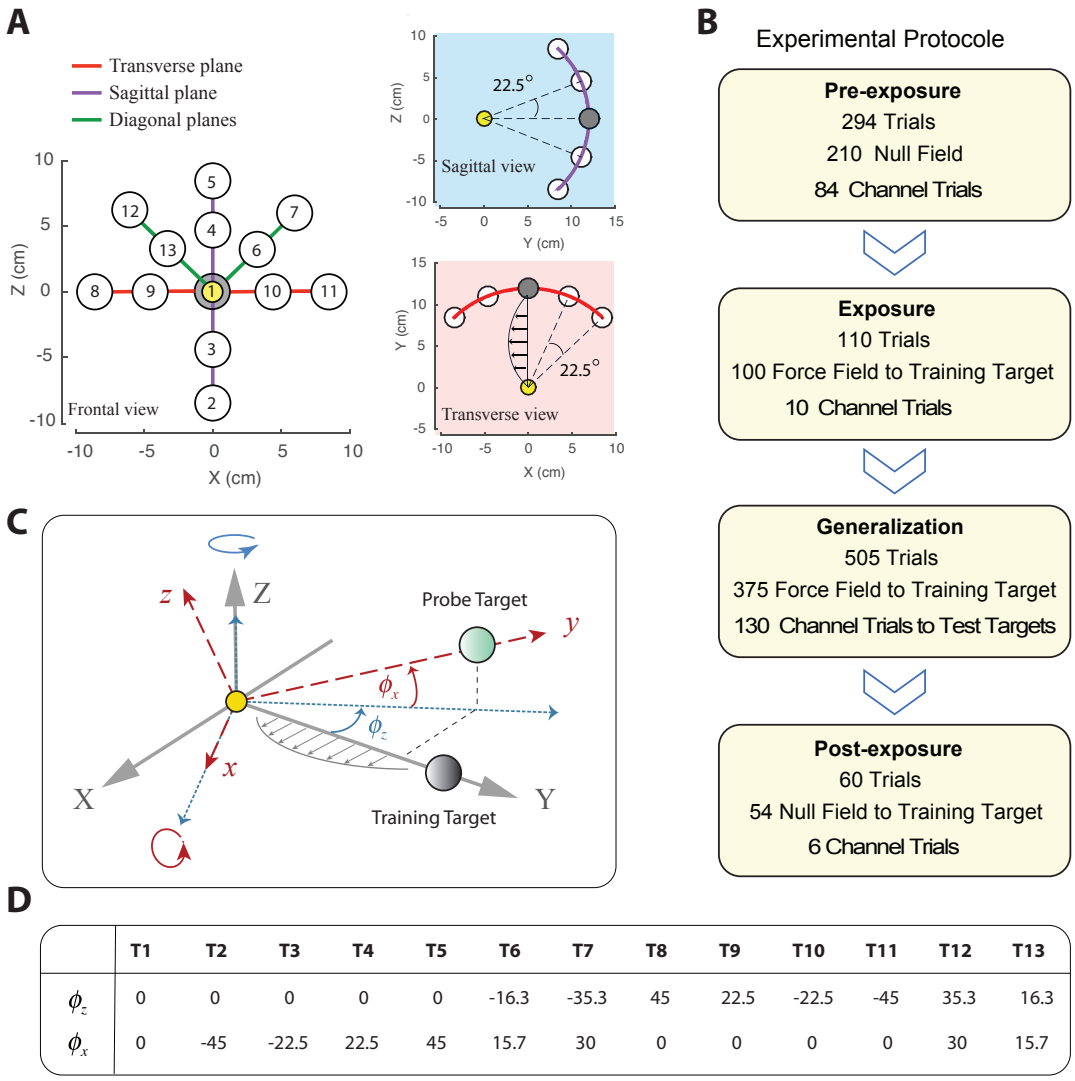

	T1	T2	T3	T4	T5	T6	T7	T8	T9	T10	T11	T12	T13
ϕ_z	0	0	0	0	0	-16.3	-35.3	45	22.5	-22.5	-45	35.3	16.3
ϕ_x	0	-45	-22.5	22.5	45	15.7	30	0	0	0	0	30	15.7

Fig. 2.2 **A.** Target positions. Targets were distributed on the surface of an imaginary sphere (radius of 12 cm), with the starting position at the centre of the sphere (yellow circle). Three different views of the target locations is shown: the frontal view (left panel), the sagittal view (top right) and the transverse view (bottom right). Targets were lined up on three different planes in the 3D space (transverse, sagittal and diagonal planes) with the angular distance of 22.5° with each other. **B.** The experiment consisted of four successive phases: pre-exposure phase, exposure phase, generalisation, and post exposure phase. **C.** The target based (local) frame for each target (*xyz*-red frame) was obtained by two consecutive rotations of the main (extrinsic) frame (XYZ-grey frame), first around the Z axis for ϕ_z degrees, and then along the resultant x axis for ϕ_x degrees. The y-axis of the resultant local frame pointed towards the target location, and the x and z axes were used to represent the force and adaptation measurements on the channel trials towards that target. Note that the x axis of the local frame always remained on the transverse (XY) plane, in which also the force field was applied. Therefore, compensation of the force field by subjects was reflected mainly in the x component of the measured forces. **D.** The rotation angles (in degrees) associated with the local frame of each target.

with the setup. The experiment consisted of four different phases: the pre-exposure phase, the exposure phase, the generalisation phase, and the post-exposure phase. In the pre-exposure phase, subjects performed reaching movements mainly under null trials towards all targets in a pseudorandom order. This consisted of 15 null trials and 6 channel trials per target, with the exception of the central (training) target (the target in the intersection of transverse and sagittal plane; Fig. 2.2A, the grey target), to which subjects performed 30 null trials and 12 channel trials. This phase provided the baseline performance during reaching movements to each target (Fig. 2.2B).

In the exposure phase, subjects performed 100 movements under the exposure trials, exclusively towards the central target. In this phase subjects adapted to the force field only in one movement direction. Ten additional channel trials were used in random locations throughout the phase to measure adaptation. During the generalisation phase, subjects continued to perform under exposure trials towards the central target, while we occasionally assessed their generalisation of learning towards other targets. This phase consisted of 5 blocks. In each block subjects performed 75 exposure trials towards the central target, as well as 26 pseudorandomly distributed channel trials towards all 13 targets (2 channel trials per target, including the central target). Finally, in the post-exposure phase, subjects performed 60 reaching movements (including 6 channel trials) towards the central target in null field to washout their adaptation. Subjects received rest breaks roughly every 100 trials throughout the experiment.

2.2.2 Data analysis

Half of the subjects in this experiment were tested under clockwise force field ($b = 15$ Ns/m), and the other half under counter clockwise ($b = -15$ Ns/m). The data collected from each group was properly combined such that the directions of error and force measurements were comparable for both groups. Two performance measures were calculated. First, on the exposure and null trials we measured the position of the hand during each movement (at 1000 Hz) and calculated the maximum spatial deviation of the hand from the straight line that connected the start and target positions (maximum perpendicular error or MPE). This

measure was considered as the kinematic error of performance (the MPE of zero meant moving in a straight line as an ideal performance).

Second, on each channel trial and for each target, we measured the lateral forces generated by subjects against the channel walls (at 1000 Hz) in compensation for the force field. The measured force at each time point during the movement was a 2D vector represented in a plane that was orthogonal to the movement direction in the 3D space. The main focus in this study was to examine whether the transfer of learning to different target directions was performed in the target-based frame or the extrinsic (world-based) frame. To this end, we first defined a target-specific (local) frame for each target, in which the measured force vectors on channel trials could be represented. In Fig. 2.2C, we schematically show a local frame for an arbitrary probe target (the *xyz*-red frame) as well as the extrinsic frame (XYZ) as the main frame of reference.

The local frame for each target was obtained by two consecutive rotations with respect to the main frame, first along the Z axis for ϕ_z degrees (the blue frame in Fig. 2.2C), and then along the new x axis for ϕ_x degrees (the red frame in Fig. 2.2C). In Fig. 2.2D the rotation angles are presented for all targets. The resultant local frame captured the target location in its y direction, and the measured force vectors (on channel trials) were described in the x and z directions. Note that for the training target, the local and the main frames were coincided. Importantly, the x axis of the local frame for all targets remained on the transverse (XY) plane, in which the force field was applied. Therefore, the compensation of the force field by subjects was mainly reflected in the x component of the measured force.

To obtain the percentage adaptation based on the measured forces on each channel trial, we first took the time course of x and z components of the measured force during the movement, and trimmed the trajectories based on the velocity of movement. That is, from the time point when the velocity exceeded 5% of its maximum value, onwards to the time when it was reduced to less than 5% of its maximum value. We then regressed (without intercept) each component of the trimmed force trajectory onto the ideal force trajectory that would fully compensate for the force field (obtained based on equation 2.1). The resultant regression coefficients were then considered as the adaptation index for each component.

2.3 Results

In this experiment, we examined adaptation to a velocity dependent force field while reaching to the central (training) target, and assessed the generalisation of adaptation to other target locations distributed in 3D virtual workspace. The baseline performance of the subjects was quantified using channel trials during the pre-exposure phase. Subjects were then exposed to the force field in the exposure phase, and learned to compensate for the perturbations in the training movement direction (central target). Exposure continued in the generalisation phase towards the central target, while we assessed generalisation to other targets using channel trials.

Force field learning

To investigate the learning process, we first examined the MPE as the kinematic measure of performance. Fig. 2.3A shows the average ($\pm$SE) time course of MPE across subjects, when reaching to the central target. During the pre-exposure phase, subjects made roughly straight reaching movements with small MPE. The mean ($\pm$SE) of MPE across the last 10 trials in this phase was 1.00 ± 0.075 cm. In the first trial of the exposure phase, the MPE showed a sudden increase up to 9.7 ± 1.45 cm (mean $\pm$SE across subjects) due to the unexpected introduction of the force field. In the subsequent trials, subjects learned to reduce the error progressively over trials down to 1.69 ± 0.80 cm (mean $\pm$SE) by the last 10 trials of the exposure phase. The amount of MPE, however, was still significantly larger than the baseline performance during the pre-exposure phase (paired t-test: $t(7) = -2.781$, $p = 0.027$). During the generalisation phase, force field exposure continued at the central target. The mean ($\pm$SE) MPE across the last 10 trials in generalisation phase was 1.44 ± 0.14 cm, not significantly different from the MPE during the late exposure phase (paired t-test, $t(7) = 1.24$, $p = 0.298$; Fig. 2.3B). In the post-exposure trials, when the force field was turned off, subjects showed an increase in the MPE up to 4.37 ± 0.26 cm due to the force field after effects, falling rapidly in the subsequent trials as they deadaptated (mean $\pm$SE across last 10 trials: 0.85 ± 0.04 cm).

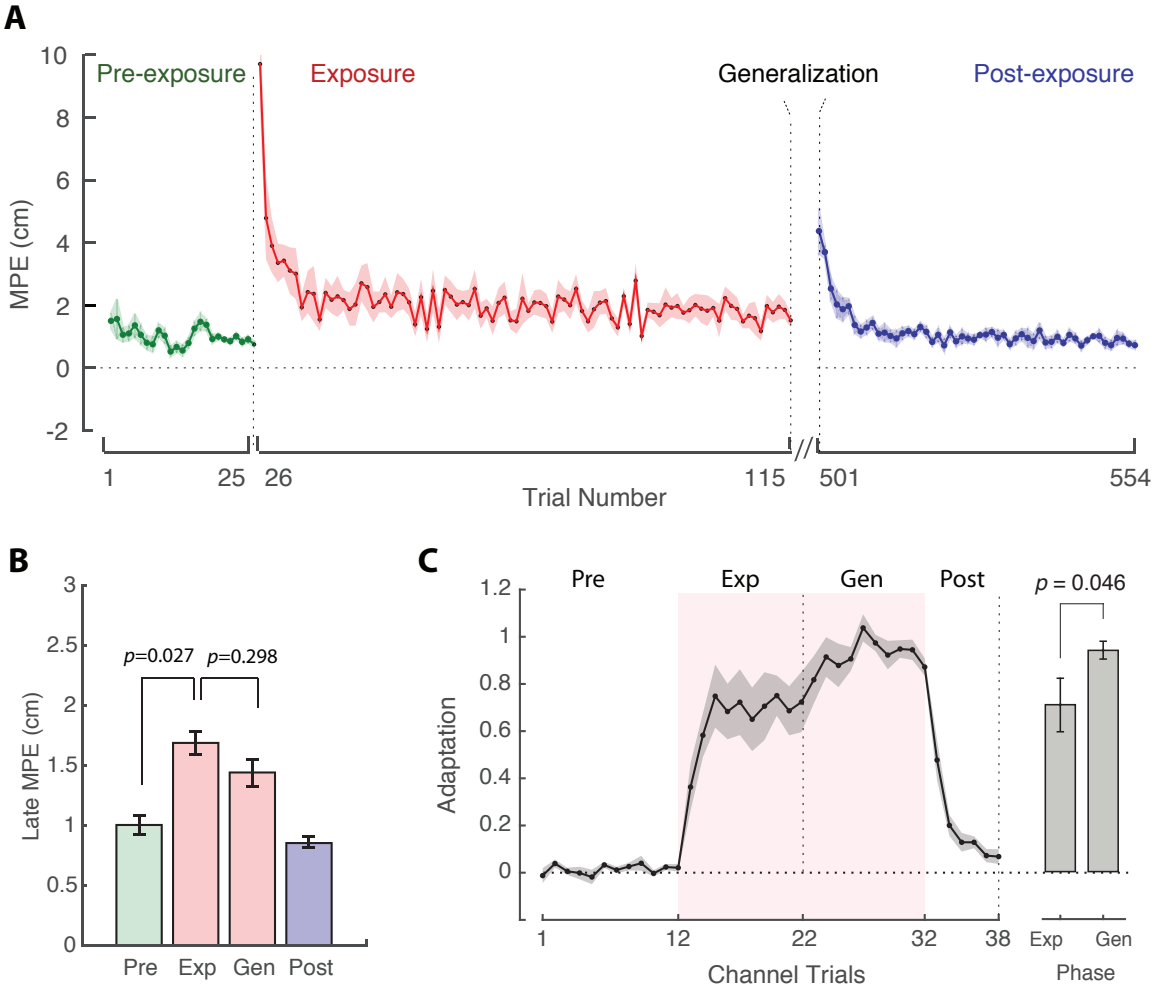

Fig. 2.3 **A**. The average (±SE) time course of maximum perpendicular error (MPE) across 8 subjects during reaching to the central target. Data shows the MPE over pre-exposure (green), exposure (red), and post-exposure (blue) trials. **B**. The average (±SE) of MPE across the final 10 trials of pre-exposure, exposure, generalisation, and post-exposure phase. **C**. The mean (±SE) of adaptation index for the central target over the channel trials during different phases of the experiment. The bar plots show the mean adaptation in the last five channel trials of the exposure phase and generalisation phase. The error bars represent the SE across subjects, with p-values determined by paired t-test (details in the text).

We also examined the adaptation index for the central target during the experiment. Fig. 2.3C shows the percentage adaptation on channel trials during all four phases. As illustrated, the baseline adaptation level was close to zero during the pre-exposure phase with mean±SE across subjects equal to 0.013 ± 0.014 (t-test: $t(7) = 0.986$, $p = 0.357$). During the exposure and generalisation phase, adaptation increased, implying that subjects learned to produce compensatory forces to resist against the force field perturbations. The mean±SE of adaptation across the last 5 channel trials of exposure phase and generalisation phase were, respectively, 0.70 ± 0.11 and 0.93 ± 0.03, which showed a significant difference (paired t-test: $t(7) = -2.425$, $p = 0.046$; Fig. 2.3C). As shown, adaptation was still in progress when transitioning from the exposure to the generalisation phase, leading to larger adaptation values by the end of the generalisation phase. This was not clearly evident from the MPE data, as the difference between the late MPE values of exposure and generalisation phase was not significant (Fig. 2.3B). It could be that the MPE is usually affected by stiffening of the arm in the early stages of learning, which reduces the amount of error, but does not reflect adaptation. Due to long-run training, however, the effect of arm stiffening reduces as the adaptation increases. As a result, the error remains the same while the adaptation improves. Finally in the post-exposure phase, the adaptation dropped rapidly due to deadaptation to the force field.

Generalisation of learning

Next, we examined the transfer of learning from the central (training) target to the probe targets. During the generalisation phase, we randomly introduced channel trials towards different targets and measured the generated forces by subjects. The measured forces were all represented in the local frame of each target, and were regarded as 2D vectors with x and z components. Fig. 2.4A shows the force vectors produced by subjects on channel trials at the time of hand maximum velocity. Each force vector represents the data of an individual averaged across the channel trials of the generalisation phase (dark blue vectors). The overall averaged force across subjects is also shown by the orange vector. As shown, the direction of the force is mainly in the x dimension (the task relevant dimension) for all target locations.

The magnitude of the force is maximal for the central target and decreases for targets that are more distant.

Fig. 2.4B also illustrates the time course of the generated force in the x-direction, for the pre-exposure phase and the generalisation phase. As shown, the force trajectory during the pre-exposure phase (black line) is close to zero for all targets and across subjects (baseline performance). After learning the force field, however, subjects show a bell-shaped force profile, akin to the velocity profile during the reaching movements, which indicates the learning of the velocity dependent force field. It is observed that the shape of the force profile is similar across all target positions. This is important as it shows that subjects actually learned the mapping between the force and the velocity of movement during force field exposure. The magnitude of the force profile, however, is different across the targets, and it is maximum for the central target, while decays for probe targets as they get farther apart from the training target.

To further explore generalisation, we looked at the x component of adaptation index across different target positions. For each target, we obtained the average adaptation across its channel trials over the generalisation phase, as well as the pre-exposure phase (baseline performance) as shown in Fig. 2.5. Specifically, we grouped the targets in five different sets and examined each group separately: targets in the transverse plane (Fig. 2.5A), targets in the sagittal plane (Fig. 2.5B), targets in the diagonal plane (Fig. 2.5C), targets on the inner ring (angular distance from the central target: 22.5°; Fig. 2.5D), and targets on the outer ring (angular distance from the central target: 45°; Fig. 2.5E). In all cases, we observed the highest adaptation level for the central target (T1: mean±SE, 0.92 ± 0.05; t-test: $t(7) = 19.2$, $p < 0.0001$) and lower adaptation for other targets depending on their angular distance from the centre. The transfer of learning to the targets in the inner ring with angular distance of 22.5° (mean±SE across targets and subjects: 0.71 ± 0.043) was significantly larger than the transfer of learning to the targets in the outer ring with angular distance of 45° (mean±SE across targets and subjects: 0.49 ± 0.042; paired t-test: $t(7) = 12.92$, $p < 0.001$).

In the following, we examine the pattern of generalisation in different group of targets in more detail. In the transverse plane (Fig. 2.5A), all targets were on the same plane as the

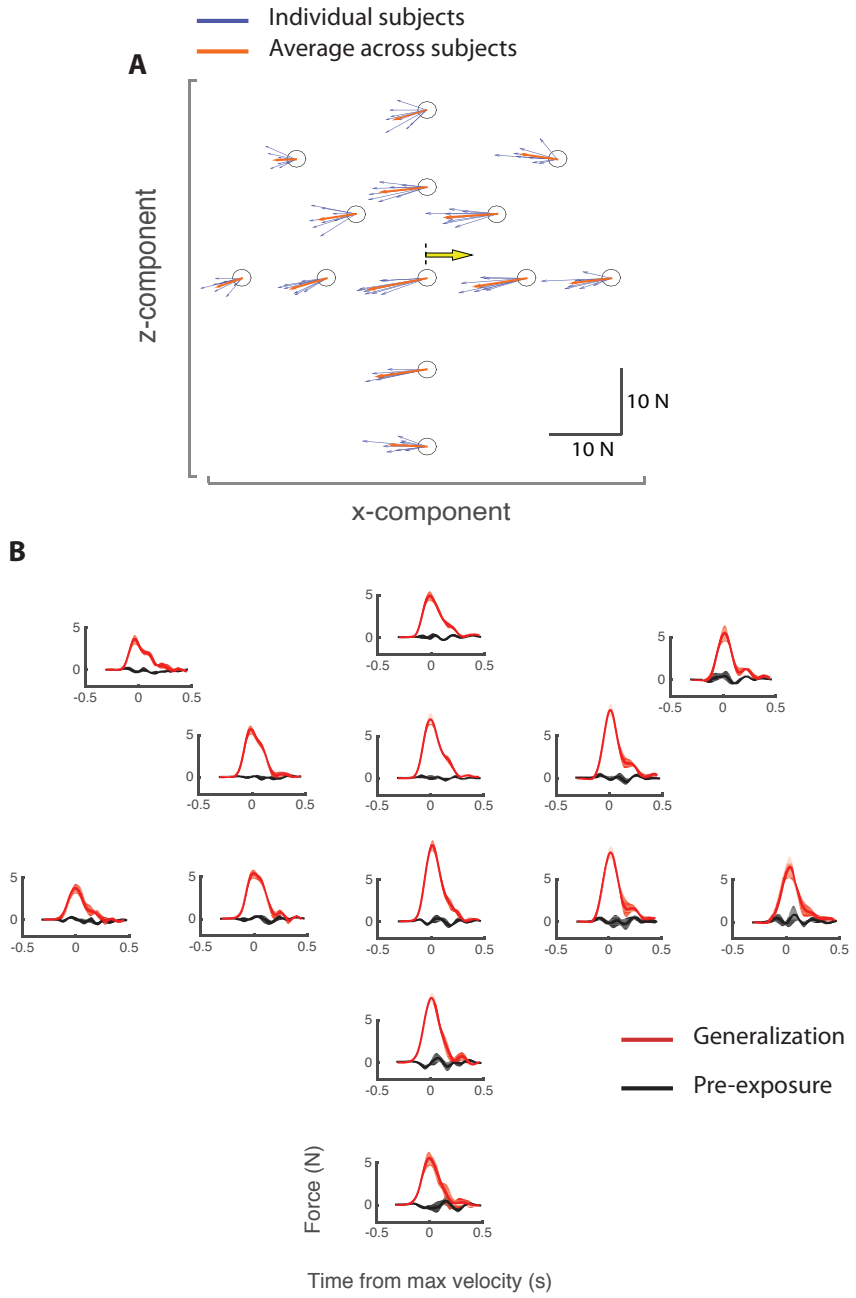

Fig. 2.4 **A**. Force vectors exerted by individual subjects in response to the force field (yellow arrow) measured on channel trials at maximum hand velocity. Blue vectors are the mean forces calculated for each individual across channel trials of the generalisation phase. Orange vectors are the overall mean of the forces across subjects for each target. **B**. The time course of generated forces in the x-direction of each target's local frame, averaged ($\pm$SE) across subjects. Each plot is associated with one target location. Force trajectories are presented for pre-exposure (black line) and generalisation (red line) phase.

force field perturbation. This was similar to the planar generalisation studies as described earlier in Fig. 2.1B (e.g., Howard and Franklin, 2015). In this case, the measurement of adaptation was performed in the x dimension of the target-based (local) frame for each target. However, depending on the coordinate frame in which the adaptation was actually transferred (extrinsic or target-based), the interpretation of the results would be different. If the transfer of learning was mediated by the target-based frame, the adaptation patterns in Fig. 2.5A would be the true measure of generalisation (transfer of adaptation in the target-based frame requires forces generated in the x dimension, which is the same as the direction of measurement; Fig. 2.1B, purple arrow). Whereas, if the transfer of learning was performed in the extrinsic frame (that is, in X-dimension of the main frame), the measurements would only reflect a fraction of the actual generalisation (that is, only the component that was aligned with the measurement direction; 2.1B, blue arrow). This meant that the generalisation pattern obtained in Fig. 2.5A for the transverse plane would be smaller than the true generalisation.

In order to examine this issue, we compared the generalisation pattern of the transverse plane with that of the sagittal plane. Crucially, in the sagittal plane, the x-axis of the local frame for all targets was aligned with the X-axis of the main (extrinsic) frame (in Fig. 2.2C, the angle ϕ_z between the x-axis and X-axis was zero for all targets in the sagittal plane). This meant that both extrinsic and target-based representation of generalisation in this case would lead to the same predictions about the pattern of generalisation (note that in extrinsic representation, the adaptation would transfer in X-dimension of the main frame, and in target-based representation, it would transfer in the x-dimension of the local frame; these two representations are identical for targets in the sagittal plane).

Contrasting the generalisation results between the transverse and sagittal plane in Fig. 2.5A and B, we could identify the coordinate frame in which transfer of learning was represented. As such, if learning was transferred in the target-based (local) frame, the pattern of generalisation between the transverse and sagittal plane would be similar with no systematic difference. However, if the transfer was mediated by an extrinsic frame, we would expect to see larger transfer of learning in the sagittal plane compared to the transverse plane, as in the transverse plane only a fraction of generalisation could be measured.

Fig. 2.5A and B shows the generalisation patterns in transverse and sagittal planes. As shown, both patterns are fairly similar, with the largest adaptation at the central target, and smaller adaptation for distant targets. We did not observe any significant difference in the amount of transfer between the targets in transverse and sagittal planes. Particularly, comparing the transfer of learning among the targets with the same angular distance from the centre (inner and outer ring; Fig. 2.5D and E), regardless of the plane in which they laid, the amount of adaptation was not significantly different (Repeated measures ANOVA: $F_{5,35} = 0.94$, $p = 0.465$ for the inner ring, and, $F_{5,35} = 1.92$, $p = 0.115$ for the outer ring). These results implied that the transfer of learning was most likely represented in the target-based coordinate frame.

In Fig. 2.5 we examined the x component of adaptation as the task-relevant measure of adaptation (the x-axis of the local frame was on the same plane in which the force field was applied). It was observed that adaptation decayed as the angular distance between the training target and the probe target increased. Here, we further examined whether the z component of adaptation also changed when transferred from one movement direction to another. In other words, we asked whether the direction of adaptation vector altered when transferred from the training target to the probe targets. Fig. 2.6 illustrates the generalisation patterns obtained based on the z component of adaptation. As shown, the level of adaptation in this component was close to the baseline for all targets (in different spatial planes) and was not significantly different between the pre-exposure and generalisation phases. This indicated that subjects did not develop any adaptive response in the z direction, which was particularly irrelevant to the goal of the task for the targets on the transverse and sagittal planes (for targets on the transverse and sagittal planes the z component of the local frame was orthogonal to the force field perturbation, which was applied in the X-direction of the main frame).

2.4 Discussion

In this study, we assessed the generalisation of force field learning across various movement directions in a 3D space. Subjects learned the force field for a training target direction at

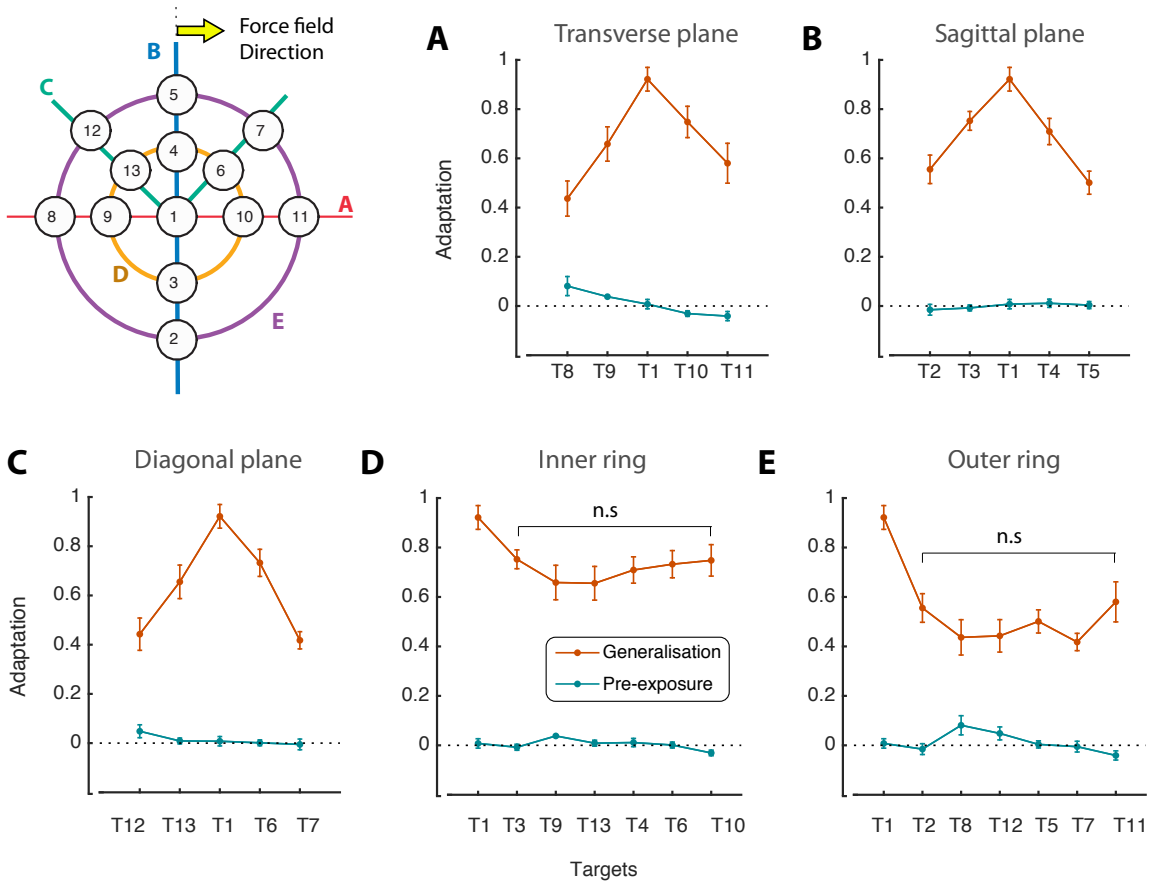

Fig. 2.5 Generalisation patterns based on the x component of adaptation across different sets of targets during pre-exposure (blue) and generalisation (orange) phases. Each error bar indicates the overall average of adaptation criterion across subjects and channel trials for each target. The patterns are generated upon the targets on the **A**. transverse plane, **B**. sagittal plane, **C**. diagonal planes, and **D**. the inner and **E**. the outer rings around the central target.

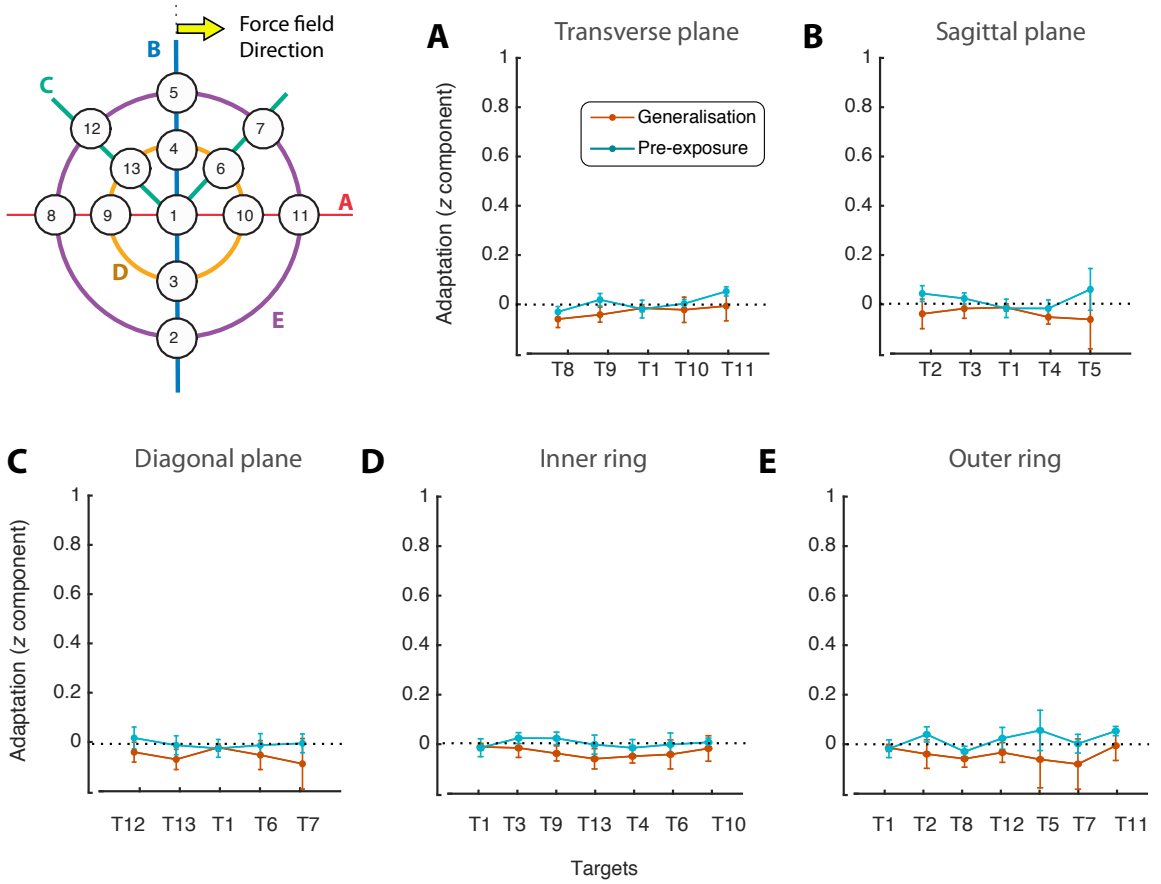

Fig. 2.6 Generalisation patterns based on the z component of adaptation for pre-exposure (blue) and generalisation (orange) phases. The panels are the same as in Fig. 2.5.

the intersection of transverse and sagittal planes, and were probed for other targets located on the transverse, sagittal, and diagonal planes (Fig. 2.2A). We found similar pattern of generalisation in the transverse and sagittal planes, and showed that targets with the same angular distance from the training target (inner and outer ring; Fig. 2.5D and E) exhibited similar amount of learning transfer.

In the context of force field learning, a large number of studies have examined the generalisation of learning to novel workspace or movement directions (e.g., Berniker et al., 2014; Howard and Franklin, 2015; Malfait et al., 2002; Mattar and Ostry, 2007, 2010; Shadmehr and Mussa-Ivaldi, 1994; Thoroughman and Shadmehr, 2000). The findings from these studies provided evidence that learning was transferred only locally, and generalisation to very different conditions was severely limited. Our results were consistent with these findings showing that adaptation level dropped very quickly as the angular distance between the training target and the probe targets increased (decay of adaptation form $\sim 92\%$ in the training target to $\sim 55\%$ in the targets with the angular distance of 45°; Fig. 2.5).

Mainly, we were interested in how generalisation of learning across different movement directions provided evidence about the coordinate frame representation of learning in the motor system. In recent years, meany studies have addressed this question by looking at force field or visuomotor adaptation generalisation (for example, Berniker et al., 2014; Brayanov et al., 2012). Various coordinate frames have been examined, and evidence has been provided in favour of extrinsic frame (Burgess et al., 2007; Krakauer et al., 2000), intrinsic or joint-based frame (Malfait et al., 2005, 2002; Shadmehr and Mussa-Ivaldi, 1994), or a mixture of several coordinate frames (Berniker et al., 2014; Brayanov et al., 2012; Parmar et al., 2015). Most of these studies examined the generalisation across different workspaces or different arm configurations (e.g., by varying joint angles), which allowed to test for joint-based versus world-based coordinate frames. However, due to the limitations of measurement in planar robots (Fig. 2.1B), very few studies have addressed the coordinate frame representation of learning when the workspace is fixed, and only the direction of movement varies (Brayanov et al., 2012). In addition, the generalisation patterns obtained across different movement directions have been difficult to interpret. For example, it has been

shown that generalisation decays when the movement direction deviates from the training direction (Donchin et al., 2003; Howard and Franklin, 2015). However, it was not clear whether the decay of generalisation was purely due to the angular distance between the movement directions, or was affected by the incomplete measurement of adaptation (e.g., in case of representations in the extrinsic coordinate frame; Fig. 2.1B, blue arrow).

Here, we examined the internal representation of generalisation when the workspace was fixed, and only movement direction changed within the 3D workspace. Two possible coordinate frames were tested: the extrinsic frame and the target-based frame. We first showed that for a group of targets laid on the sagittal plane (orthogonal to the force field perturbation), both coordinate frames led to the same representation of generalisation. That is, the direction of force compensation (X-axis in the extrinsic frame) matched the direction of measurement (the x-axis of the target-based frame). This arrangement of targets with respect to the force field direction was only possible using the 3BOT. As a result, the measured transfer of learning obtained in the sagittal plane represented the true generalisation of learning. Second, by comparing the pattern of generalisation in the sagittal plane and that of the transverse plane (akin to previous planar studies of generalisation), we found similar results. This indicated that the generalisation of learning was most likely represented in the target-based frame rather than the extrinsic frame (the extrinsic frame predicted less amount of generalisation for the transverse plane due to incomplete measurement of adaptation).

Taken together, we showed that when the workspace is fixed, movement to various directions is mainly represented by a frame that is target specific. This, however, does not mean that the extrinsic coordinate frame is not used in the motor system, but that when the workspace is fixed, the effect of extrinsic representation is limited. In fact, it has been shown that when the workspace or arm configuration changes, the extrinsic coordinate frame contributes substantially in the representation of learning (Berniker et al., 2014; Brayanov et al., 2012). It is possible that the motor system uses a mixture of target-based, world-based (extrinsic) and joint-based (intrinsic) coordinate frames, and that the contribution of each one varies depending on the task demands or the complexity of the movements.

Chapter 3

Spontaneous Recovery

3.1 Introduction

Spontaneous recovery has been extensively studied in memory research and is regarded as one of the most fundamental observations in classical conditioning (Myers and Davis, 2002; Rescorla, 2004; Robbins, 1990). The experimental design used to study spontaneous recovery normally involves three main phases. In the acquisition phase, the animal learns an association between a conditioned stimulus (CS; e.g., a tone) and an unconditioned stimulus (US; e.g., a puff of air aimed at animal's eye). This learning involves exhibiting an appropriate response to the CS, which is blinking after hearing a tone. In the extinction phase the US no longer follows the CS, leading to unlearning of the association between CS and US. As a result, the animal shows no response to the CS (in this example, the tone is no longer followed by the puff, thus the animal learns to show no blinking after the tone). Finally, the post-extinction phase in which the animal spontaneously re-exhibits some of the previously learned responses to the CS, despite the apparent extinction of such responses.

In recent years, studies have shown similar behaviour in sensorimotor learning (Criscimagna-Hemminger and Shadmehr, 2008; Kojima et al., 2004; Smith et al., 2006). For instance, Kojima et al. (2004) demonstrated in nonhuman primates that when saccadic adaptation (repeatedly displacing a target during a saccade to change the gain of the initial saccade) is followed by a rapid deadaptation, in the following phase of complete darkness (no feedback

of saccadic error) animals re-exhibit their initial adaptive behaviour. Later, spontaneous recovery was observed in a force field learning task (Smith et al., 2006). Human subjects learned a force field perturbation (field A) during reaching to a single target, and were immediately exposed to a brief force field reversal (field B) before going through a long period of error-clamp (channel) trials. It was observed that adaptation to field A rapidly reduced to baseline by the end of field B, and then showed a sudden rebound (spontaneous recovery) over the course of channel trials. It was suggested that such behaviour was the product of two interactive adaptive processes: one that learned quickly and forgot quickly, and the other that was slow to learn but also forgot slowly (Fig. 1.1).

In these studies, the exposure to the reversed force field (field B) has been considered as the essential part of the paradigm, causing the rapid deadaptation of the initial learning (equivalent to the extinction phase in classical conditioning). It has been unclear, however, how the motor memories associated with field A and field B interact and contribute to the observed spontaneous recovery. The main challenge has been that both motor memories act in the same spatial dimension (in opposite directions), and thus only the net adaptive behaviour could be measured, while the individual state of adaptation for each motor memory remains hidden. In addition, it is not clear, though generally assumed, whether spontaneous recovery observed in force field learning paradigms only happens in a one-dimensional setup, where fields A and B need to be counteracting.

Here, by utilising the 3BOT, we designed a novel experimental paradigm, akin to the spontaneous recovery experiment, in which the fields A and B were exposed in two orthogonal dimensions (as opposed to a single dimension with opposing directions). In this case, field B did not directly contribute to the deadaptation (extinction) of the memory of field A. We examined whether the interaction of these orthogonal force fields induced spontaneous recovery. Furthermore, using the 3BOT we were able to track the memories of field A and field B, independently, by measuring the adaptation to each field in its corresponding dimension. This way, adaptation was defined as a two-dimensional vector with each element representing the adaptive response to one of the force fields. The contribution of the associated motor memories to the overall adaptive behaviour, therefore, could be examined.

We demonstrate that spontaneous recovery occurs, not as a rebound of adaptation in the memory of field A, but as the rotation of the adaptation vector from the recently formed motor memory (field B), towards the initially learned more robust motor memory (field A). We argue that this behaviour is a more general form of spontaneous recovery, which reveals itself as the adaptation rebound in the case of one-dimensional task. We also test the ability of the dual-rate model (Smith et al., 2006) in explaining the spontaneous recovery observed in our two-dimensional task.

3.2 Methods

Twelve subjects participated in this experiment (6 female; age: 27.2 ± 3.6 years). All subjects were right-handed according to the Edinburgh handedness inventory (Oldfield, 1971), with no reported record of neurological disorders. Informed consent was provided by all subjects prior to the experiment. The experiment was approved by the institutional ethics committee. The procedure and apparatus was similar to that of the previous chapter, and here we only highlight the differences.

3.2.1 Experimental procedure

Subjects performed repeated reaching movements from a given start position in the 3D workspace (represented by a yellow icosahedron, 1.25 cm radius; Fig. 3.1A), to a single target position (represented by white icosahedrons, 1 cm radius) located 12 cm below the start position along the vertical y-axis. Each trial started with holding the cursor (a red sphere 0.5 cm radius representing the veridical hand position in the virtual environment) within the start icosahedron (positional tolerance 0.75 cm). The movement initiated with a tone, after which subjects were required to make a rapid reaching movement towards the target location. The movements were required to be finished within $600 - 1100$ ms, otherwise an appropriate message “too slow” or “too fast” appeared on the screen. The movement finished when subjects remained stationary within the target for 100 ms (positional tolerance 0.75

cm). Subjects then received a "Relax Arm" message, after which the robot passively returned subject's hand to the home position (using a minimum jerk motion) for the next trial.

The experiment consisted of four phases (Fig. 3.1B): the pre-exposure phase (30 null trials and 10 pseudorandomly interleaved channel trials), exposure to field A (135 exposure trials as well as 45 pseudorandomly ordered channel trials), exposure to field B (either 16 or 32 trials, including 4 or 8 channel trials, respectively), and error clamp phase (75 channel trials). The force field perturbations during field A and B were velocity dependent force fields with the same strength, but in orthogonal directions as shown in Fig. 3.1B (field A was applied in the z dimension, and field B was applied in the x dimension). The forces generated in field A and B are:

$$\mathbf{F}_A = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & b & 0 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{bmatrix}, \mathbf{F}_B = \begin{bmatrix} 0 & b & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{bmatrix}, \tag{3.1}$$

where, $\dot{x}$, $\dot{y}$ and $\dot{z}$ represent the components of hand velocity in the main coordinate frame (XYZ) shown in Fig. 3.1A. The constant $b = \pm 15$ Ns/m represents the field strength, and $\mathbf{F}_A$ and $\mathbf{F}_B$ indicate the force vectors associated with fields A and B, respectively.

The sign of the field constant ($b = 15$ or $b = -15$ Ns/m) was counterbalanced between the subjects, and the measurements for each field constant were combined to be comparable across all subjects. Each subject performed the task twice in separate weeks: once with $N_B = 16$ (number of trials in field B) and once with $N_B = 32$ (the order was counterbalanced between the subjects). We aimed to examine how the duration of adaptation to field B (N_B) affected the spontaneous recovery during the error clamp phase.

3.2.2 Analysis

We used two main measures to evaluate learning: the kinematic error of movement and the percentage adaptation to the force field. The kinematic error on each trial was measured as the maximum distance of the hand trajectory (measured at 1000 Hz) from the straight line that connected the home position to the target position. We refer to this measurement as the

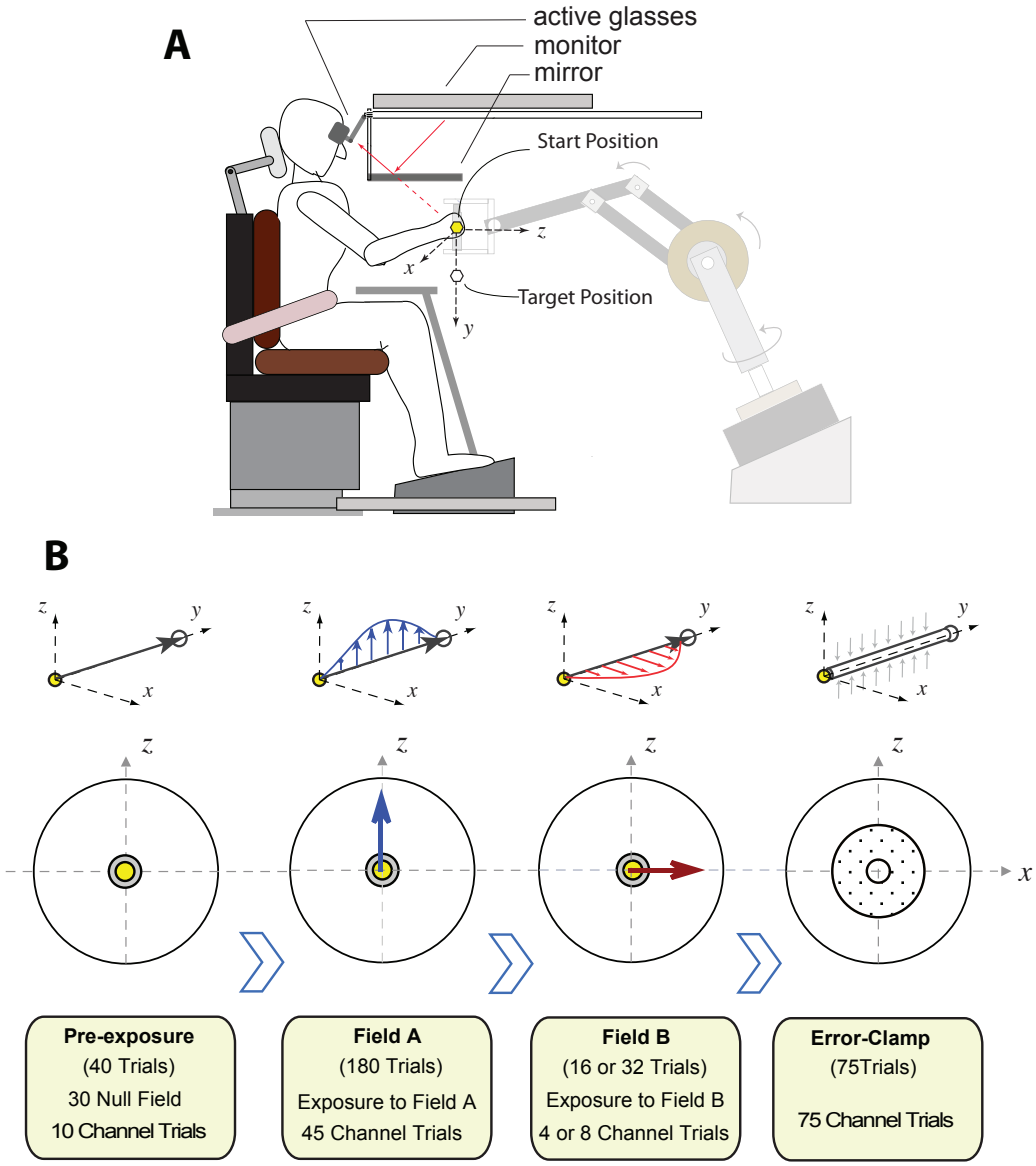

Fig. 3.1 Experimental Procedure. **A** Subjects performed a top-down reaching movement from a start position (yellow icosahedron) located at the origin of the reference frame (*xyz*), to a target location 12 cm below the start position on the *y* axis. **B** The experiment consisted of four phases: pre-exposure phase, Field A (force field in *z* dimension), Field B (force field in *x* dimension), and error clamp phase (channel tubes connecting the start position to the target position).

maximum perpendicular error (MPE), which is a two-dimensional vector with components in the x and z direction (Fig. 3.1B).

On channel trials, we measured the time course of generated forces by subjects (at 1000 Hz) in compensation for the force field perturbation. Like MPE, the measured forces were represented as two-dimensional force vectors with x and z components at each time during the movement. To obtain adaptation, we first trimmed the measured force trajectories on each channel trial based on the speed of movement, from when the speed exceeded 5% of its maximum value, onward to the time when it reduced to less than 5% of its maximum value. We then regressed (without intercept) the resultant force trajectory of each component onto the ideal force trajectory that would fully compensate for the force field. The regression coefficient was used as the adaptation index, ranging from zero (no adaptation) to one (full adaptation).

3.2.3 Modelling

Various state-space models have recently been proposed to describe and predict the adaptation process during force field learning (Donchin et al., 2003; Smith et al., 2006; Thoroughman and Shadmehr, 2000), visuomotor learning (Lee and Schweighofer, 2009; Tanaka et al., 2009), and object manipulation (Ingram et al., 2013, 2011). Recently, it has been shown that learning is driven by multiple adaptive processes that work in parallel, and their sum activity accounts for observed adaptive behaviour (Lee and Schweighofer, 2009; Smith et al., 2006). In this case, a dual-rate state-space model was shown to be able to account for the phenomena of savings and spontaneous recovery during force field learning. Here, we use a similar approach by extending the dual-rate model to a vector format to examine the adaptation behaviour in a two-dimensional force field task. Equation 3.2 illustrates the model formulation.

$$\begin{aligned}
\mathbf{x}_f^{(n+1)} &= \alpha_f \cdot \mathbf{x}_f^{(n)} + \beta_f \cdot \mathbf{e}^{(n)} \\
\mathbf{x}_s^{(n+1)} &= \alpha_s \cdot \mathbf{x}_s^{(n)} + \beta_s \cdot \mathbf{e}^{(n)} \\
\mathbf{x}^n &= \mathbf{x}_f^{(n)} + \mathbf{x}_s^{(n)} \\
\beta_f &> \beta_s \quad \alpha_f < \alpha_s
\end{aligned} \tag{3.2}$$

In the model, $\mathbf{x}^n$ represents the net vector of adaptation (with components in x and z directions), which is the sum of two adaptive processes with fast ($\mathbf{x}_f^n$) and slow ($\mathbf{x}_s^n$) learning rates on a given trial n. The parameters α_s and α_f represent the retention factors associated with the slow and fast processes, respectively. Likewise, the parameters β_s and β_f are, respectively, the learning rates in the slow and fast processes. Adaptation process is driven by the error vector $\mathbf{e}^n$ defined as the vector difference between the net adaptation ($\mathbf{x}^n$), and the external perturbation $\mathbf{f}$:

$$\mathbf{e}^n = \mathbf{f}^n - \mathbf{x}^n \tag{3.3}$$

where, $\mathbf{f}$ is a vector with binary values of 1 (if the perturbation in the corresponding dimension is on), and 0 (if the perturbation in the corresponding dimension is off).

The above model was fit to the average time course of adaptation data from both conditions ($N_B = 16$ and $N_B = 32$) simultaneously. In this case, the parameters of fit (α_s, α_f, β_s, and β_f) were obtained by minimising the overall sum of squared errors between each component of the net adaptation model $\mathbf{x}^{(n)}$ and the corresponding components of the measured adaptation data from both conditions.

3.3 Results

Subjects performed vertical reaching movements under an initial force field direction (field A), followed by a short introduction of an orthogonal force field (field B), and continued

under error-clamp trials. Each subject performed the task in two sessions (one week apart) in which we varied the duration of field B ($N_B = 16$ and $N_B = 32$). Here, we look at the behavioural performance of subjects in each group in terms of MPE and adaptation index.

3.3.1 Kinematic error

Fig. 3.2A and B show the time course of MPE separated into the x and z components and averaged ($\pm$SE) across subjects for each condition. In the pre-exposure phase (null field), subjects performed nearly straight reaching movements with the MPE close to baseline for both components and across both conditions (overall mean $\pm$SE: 0.028 ± 0.058, t-test: $t(11) = 0.299$, $p = 0.771$).

During field A, the perturbation was in the z dimension, causing the z component of MPE to increase in the first trial of field A (2.7 ± 0.21 cm for condition 1, and 2.3 ± 0.24 cm for condition 2). In the following trials, the MPE progressively decreased as subjects adapted to the force field. Fig. 3.2C and D (left panels) show the MPE in z dimension for the early and late trials of field A (window of 10 trials). Meanwhile, the x component of MPE showed a small, yet significant increase at the beginning of field A (Fig. 3.2C and D, right panels; t-test on early trials: $t(11) > 6.98$, $p < 0.001$ for both conditions), despite the fact that the force field did not have any component in the x dimension during this phase. This induced component of error could arise from the nonlinear dynamics of the arm, based on which the direction of hand movement due to the force field perturbation (error) slightly deviated from the direction in which the force was actually applied (Hogan, 1985). Because of this biomechanical property of the arm, the hand motion under a given perturbation direction could have two components: one inline with the perturbation direction (here, z component), and one orthogonal to it (i.e., x component). This induced component of MPE, however, is reduced quickly to the baseline in the following trials of field A (Fig. 3.2C and D, right panels, late trials of A). Later, we will discuss that such induced component of error has minimal effect on the adaptation behaviour.

During field B, the perturbation was switched to the x dimension. This caused a negative MPE in the z dimension due to the after-effects of field A (Fig. 3.2C and D), while at the

same time increased the MPE in x dimension due to the new force field direction. Over the course of trials, both components of MPE concurrently reduced towards the baseline, indicating rapid deadaptation in the z dimension as well as adaptation in the x dimension. In the next section, we examine the behaviour of adaptation index for each dimension and show its consistency with the MPE behaviour.

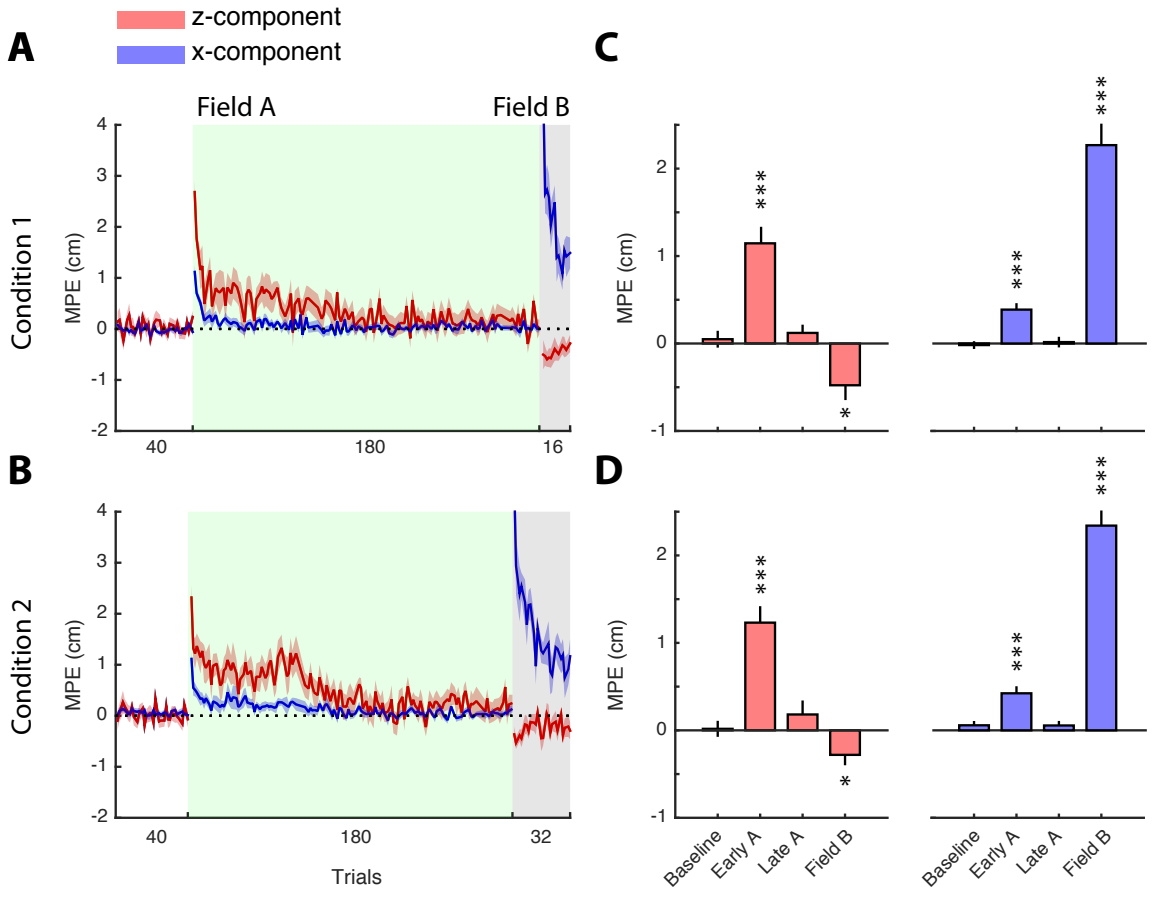

Fig. 3.2 Maximum perpendicular error(MPE). **A** and **B** show the mean ($\pm$SE) time course of MPE across subjects for two conditions ($N_B = 16$ and $N_B = 32$) and separated for x and z components (blue and red lines, respectively). **C** represents the summary of MPE for x (right) and z (left) components in condition 1. The bar plots illustrate the mean MPE across the null trials (Baseline), the first 10 trials of field A (Early A), the last 10 trials of field A (Late A), and the first 10 trials of field B (Field B). **D** illustrates the same bar plots for condition 2. (t-test: $^{*}p < 0.05$, $^{***}p < 0.001$).

3.3.2 Adaptation

As shown in Fig. 3.3A and B (blue lines), during field A the x component of adaptation remained close to baseline throughout the phase in both conditions ($N_B = 16$ and 32). This indicates an appropriate lack of adaptive response, as during field A no perturbation was applied in the x dimension. However, we earlier observed that the x component of MPE showed an increase in the beginning of this phase without any perturbations involved (Fig. 3.2, blue line). To assess whether this induced MPE had any effects on the x component of adaptation, we looked at the average adaptation across the first 10 channel trials of field A in the x dimension. The values were obtained as 0.056 ± 0.024 for condition 1 (t-test: $t(11) = 2.41,\ p = 0.04$), and 0.019 ± 0.020 for condition 2 ($t(11) = 0.844,\ p = 0.416$), indicating minimal effects by the x-component of MPE on the adaptation behaviour. This implied that the MPE in the x dimension was irrelevant to the task and thus resulted in minimal adaptive response. Previous studies have also shown that the nervous system estimates the relevance of observed errors, and only adapts to the errors that are relevant to the goal of the task (Wei and Körding, 2009).

During field B, the force field was in the x dimension, resulting in a rapid increase in the x component of adaptation up to $\sim 55\%$ for both conditions. This rapid adaptation was then followed by a rapid decay in the subsequent error-clamp phase (Fig. 3.3C and D, right panels: Early vs Late trials of error-clamp phase), indicating a fragile memory formed for the learning of field B.

Examining the z component of adaptation in Fig. 3.3A and B (red line), we observed a progressive increase in the adaptation level over the course of field A, reaching up to $\sim 80\%$ by the end of the field for both conditions. At the beginning of field B, the force field was turned off in the z dimension, resembling a washout period for the z component of adaptation. As a result, subjects deadapted to the force field (Fig. 3.3C and D, left panels), and by the end of field B, the adaptation dropped to a lower, yet still significantly positive level. This suggested an incomplete deadaptation process before the onset of error-clamp phase (Fig. 3.3C and D, left panels, Late B: t-test: $t(11) = 3.16$, $p = 0.009$ for condition 1, and $t(11) = 2.60$, $p = 0.024$ for condition 2). Finally during the error-clamp phase, the adaptation

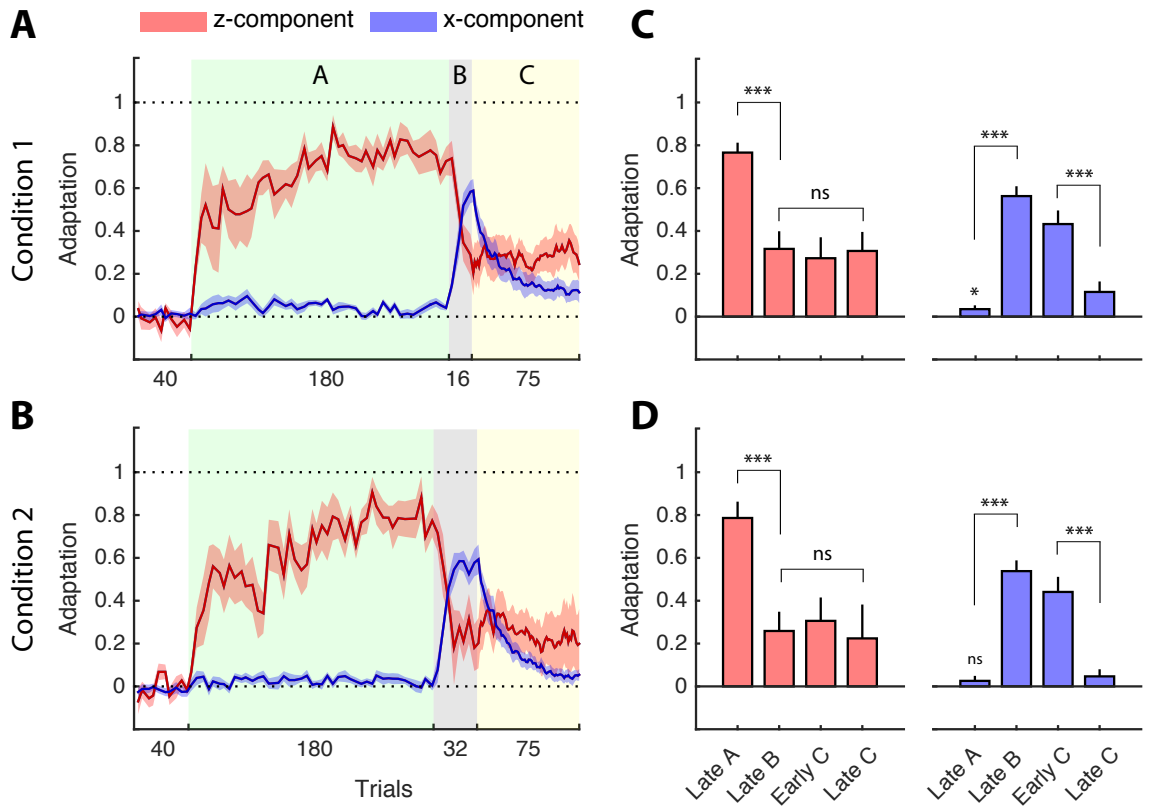

Fig. 3.3 The time course of adaptation decomposed to its x (blue) and z (red) components, and shown for condition 1 (**A**) and condition 2 (**B**). The bar plots in **C** and **D** show the summary of adaptive behaviour for the z (left) and x (right) components and for each condition. Each bar plot represents the mean adaptation in the final 10 channel trials of field A (Late A), the last 2 trials of field B (Late B), the first 10 trials, and the last 10 trials of error-clamp phase (Early C and Late C).
(paired t-test:$^{*}p < 0.05$, $^{***}p < 0.001$, ns: not significant).

level remained almost flat, with no significant difference between the early and late trials in this phase (paired t-test: $t(11) = -0.592$, $p = 0.565$ for condition 1, and $t(11) = 0.673$, $p = 0.515$ for condition 2; Fig. 3.3C and D, left panels).

Comparing the adaptive behaviour in z dimension with what is normally observed in conventional spontaneous recovery paradigms (Criscimagna-Hemminger and Shadmehr, 2008; Pekny et al., 2011; Smith et al., 2006), it can be seen that in previous paradigms the memory of field A was seemingly fully washed out by an opposing field (B), but was reactivated (in the form of adaptation rebound) early in the following error-clamp phase. In our paradigm, however, we did not observe such behaviour in the memory of field A (z component of adaptation). As shown, the adaptation to field A was never totally washed out before the error-clamp phase, and more importantly, preserved some learning in a low yet consistent level throughout channel trials (i.e., no significant rebound was observed).

We further examined the adaptation behaviour from a different perspective, by looking at the magnitude and angle of the adaptation vector in the xz plane. Fig. 3.4A and C show the angle of adaptation (measured from the x axis) averaged across subjects (circular mean) for each condition. As shown, the adaptation angle remained close to optimal (i.e., 90°) during field A, consistent with the angle of perturbation in this field. When the force field was switched to field B (x dimension), the vector of adaptation also rotated towards 0°, inline with the direction of the second force field. Interestingly, during this rotation the magnitude of adaptation remained unchanged (no significant difference was observed in the magnitude of adaptation between field B and the final trials of field A; Fig. 3.4B and D). This indicated that when the field direction switched, subjects did not reduce the strength of their compensatory forces, but just rotated the force direction in order to compensate for the perturbations in the new direction (field B). The amount of rotation was slightly larger in the second condition (due to longer exposure in field B), but the effect was not significant (last two trials of field B: $65.0° \pm 7.2°$ for condition 1, and $67.7° \pm 9.5°$ for condition 2; Watson-Williams test: $F_{1,22} = 0.05$, $p = 0.82$).

During the error-clamp phase, we observed that the adaptation vector spontaneously rotated back towards the initially learned force field direction (i.e., 90°) while also decaying

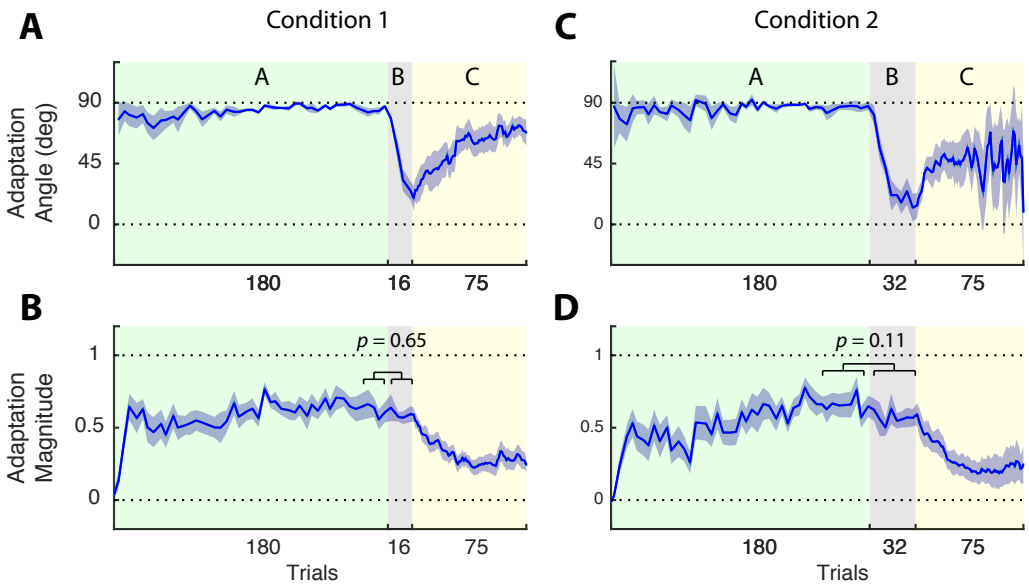

Fig. 3.4 **A** and **C** show the time course of adaptation angle averaged across subjects (circular mean $\pm$ circular standard error) for condition 1 and condition 2, respectively. **B** and **D** also illustrate the adaptation magnitude for each condition. The p-values are shown for the paired t-test between the adaptation levels in field B and the final trials of field A.

progressively in strength (Fig. 3.4). This behaviour was mainly attributed to the x component of adaptation (that is, the memory of field B) during the error-clamp phase. As shown earlier in Fig. 3.3A and B, the adaptation in x dimension showed a fast decay over the error-clamp trials, while the z component of adaptation preserved a somewhat constant level throughout. This led to an effect on the adaptation angle, whereby the direction of adaptation showed a rebound (spontaneous recovery) towards the direction of the initially learned force field (field A). This behaviour suggested a more general form of spontaneous recovery that could not be recognised in conventional paradigms, in which fields A and B were in opposite directions. In the following, we test whether a dual-rate model can account for this observation.

3.3.3 Model fits

The model was fit to the x and z components of adaptation from both $N_B = 16$ and $N_B = 32$ conditions at once. The fitting parameters were then used to simulate the adaptation angel and magnitude to be compared with the data. Fig. 3.5 illustrates the model fits, including the

net adaptation along with the slow and fast processes for x and z components, as well as the angle and magnitude of adaptation. The model was able to produce the overall behaviour observed in the data with decent accuracy.

According to the model fits, it was observed that the adaptive behaviour in the z dimension (memory of field A) was mainly driven by the slow process over the course of the error-clamp phase (Fig. 3.5A). Due to the long period of training during field A, the slow process gradually took over the adaptation, and thus largely contributed to the behaviour in the final phase. The fast process, on the other hand, only caused a small rebound in the net adaptation due to its quick decay in the beginning of this phase. In the data, this slight rebound was observed only in condition 2, yet the significance of the rebound was not achieved: comparing the final trial of field B (Late B) and the initial trials of the error-clamp phase (Early C) did not show a significant increase (rebound) in the level of adaptation (Fig. 3.3C and D, left panels).

The x component of adaptation (memory of field B) on the other hand, was predominantly driven by the fast process (Fig. 3.5B). As shown, during field B, due to the shortness of the exposure period, the fast process takes over the adaptation very quickly, whereas the slow process contributes to the adaptation only marginally.

The interaction of x and z components of adaptation, one with the properties of a fast process, and the other driven mainly by a slow process, resulted in an interesting form of spontaneous recovery in the level of adaptation angle. As shown in Fig. 3.5C, both slow and fast processes initially matched the direction of force field during field A (i.e., 90°). By introducing field B (in 0°), both processes made a rotation towards 0°, but the extend of rotation was larger for the fast process than the slow process. This caused the net adaptation to rotate appropriately mainly under the influence of the fast process. During the error-clamp phase, the slow and fast processes preserved their angle throughout the phase, however, the decay of their magnitude caused the net adaptation angle to rotate back towards the initially learned angle (90°).

In this case, the model makes two noteworthy predictions. First, the angle of the net adaptation in the final trials of error-clamp phase is determined by the slow process (Fig. 3.5C,

the red line converges the magenta line), which in turn means that the final net adaptation angle never fully returns to the angle of the initial learning (i.e., 90°). Second, longer exposure during field B allows more involvement of the slow process in this field, leading to larger rotation of this process by the end of the field (Fig. 3.5C). This means that in the following error-clamp phase, the angle of the slow process is farther away from the direction of the initial learning for the $N_B = 32$ condition, compared to the $N_B = 16$ condition (Fig. 3.5C, magenta line). Since the final net adaptation angle in error-clamp phase is determined by the slow process, the amount of rebound in adaptation angle would, hence, be smaller for the $N_B = 32$ condition.

We examined both these predictions on the data. First, we tested the final level of adaptation angle against the angle of the initial force field (i.e., 90°). To this end, we took the circular mean of adaptation angle across the last 20 trials of the error-clamp phase for each condition. This was $69.87° \pm 7.9°$ for condition 1, and $45.44° \pm 32.33°$ for condition 2. As shown, on average, the final adaptation angle was smaller than the initial force field angle for both conditions. Although, due to the large amount of noise, especially in condition 2, the significance of the difference for this condition was compromised (one sample t-test against 90°: $t(11) = 6.18$, $p < 0.001$ for condition 1, and $t(11) = 1.49$, $p = 0.162$ for condition 2).

Second, as shown above, the final adaptation angle for the second condition was (on average) smaller than that in the first condition, showing smaller effect of rebound when subjects were exposed to longer period of field B. Again, due to the noise in the data, especially for the second condition, the significance of the difference was not found (Watson-Williams test: $F_{1,22} = 0.410$, $p = 0.528$).

Overall, the model accounted for the angular spontaneous recovery of adaptation and its predictions on the adaptation angle were consistent with the averaged data.

3.3.4 One-dimensional spontaneous recovery

In our experiments, we were able to directly measure the behaviour of the memories for field A and B by projecting each field on a different spatial dimension. In this case, the adaptive memory was represented in a two-dimensional space, in which the adaptation started in

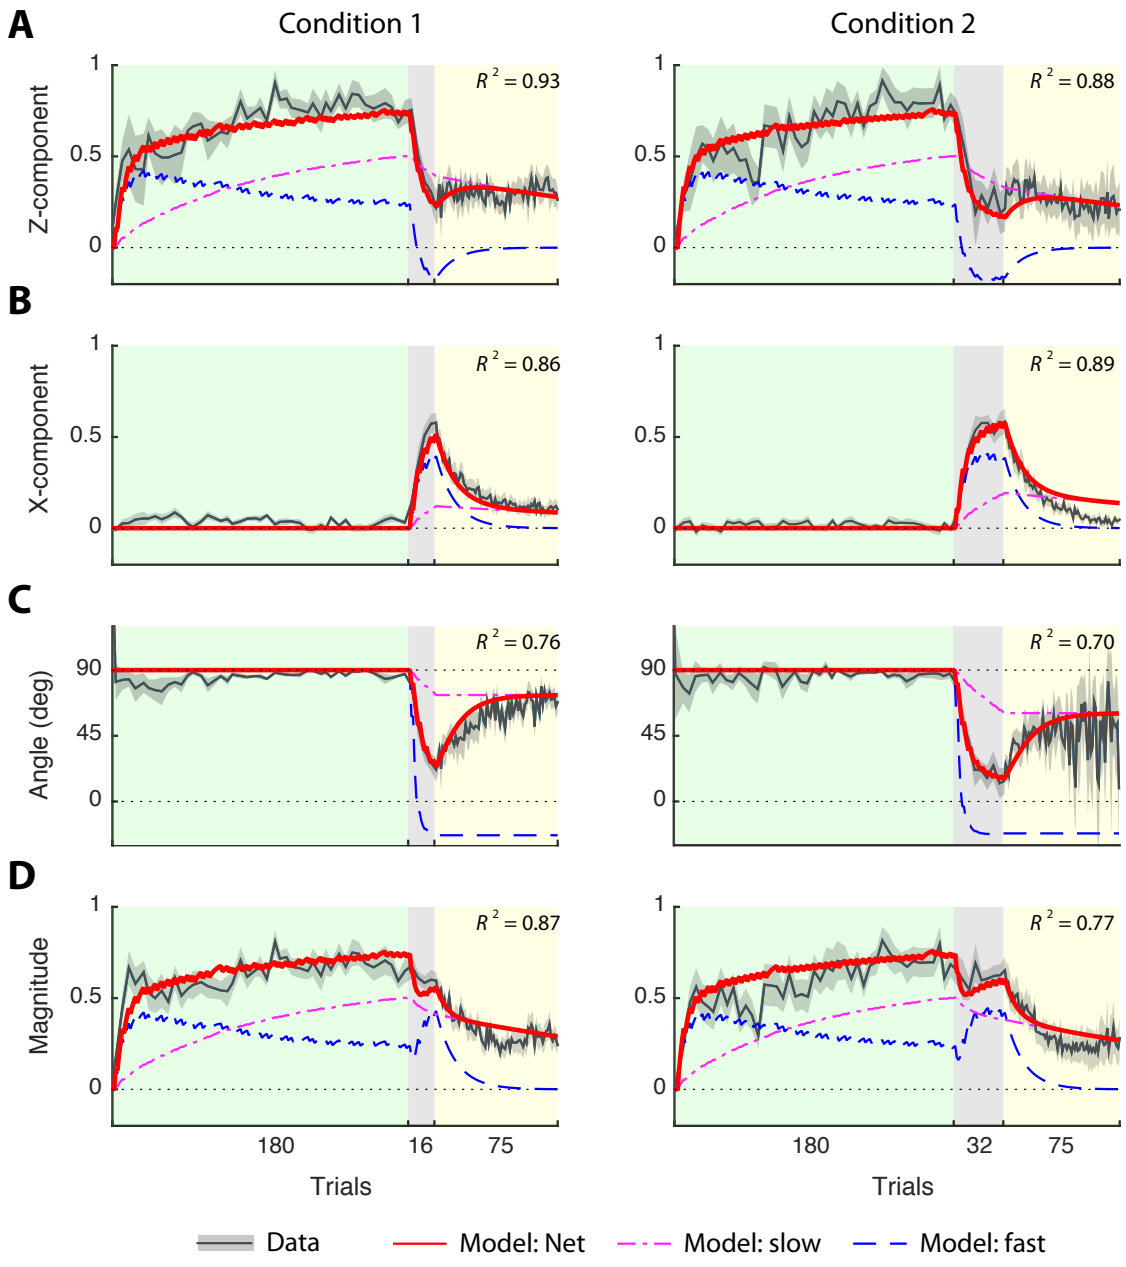

Fig. 3.5 Model fit. The model was fit to the x and z components of adaptation for both conditions at once, and then used to predict the angle and the magnitude of adaptation. Different rows illustrates the z component of adaptation (**A**), the x component of adaptation (**B**), the angle of adaptation (**C**), and the magnitude of adaptation (**D**). The parameters of fit were obtained as: $\alpha_f = 0.921$, $\beta_f = 0.092$, $\alpha_s = 0.995$, and $\beta_f = 0.014$.

the first dimension during field A, rotated towards the second dimension by introducing field B, and rotated back towards the initial dimension when the error was removed. This behaviour showed a more general form of spontaneous recovery that could not be observed in conventional paradigms in one dimension (Criscimagna-Hemminger and Shadmehr, 2008; Smith et al., 2006).

Here, we examine whether this notion of spontaneous recovery could generate similar behaviour observed in the case of one-dimensional force field learning. To this end, we first used the model to simulated a one-dimensional spontaneous recovery paradigm, in which field B was in the opposite direction of field A. We used the exact same model parameters as obtained from the model fitting in the previous section. We also used the same number of trials in the simulation as used in the experiment for both conditions. Fig. 3.6A shows the result of the simulation for condition 1 as an example. Next, as an analogy to opposing learning memories, we subtracted the x component of adaptation in our data (memory of field B), from the z component of adaptation (memory of field A). This is considered as an approximation of the case where field B is in the opposite direction of field A, and thus the overall adaptation behaviour is approximated as the signed superposition of both fields (that is, $A + (-B)$). Fig. 3.6B illustrates the data along with the model simulation for both conditions. As shown, the behaviour in this synthetic case is similar to the behaviour normally observed in one-dimensional spontaneous recovery; rapid drop of adaptation level (to the negative direction) during field B, followed by a quick rebound of adaptation towards the positive direction during the error-clamp phase. This suggests that the spontaneous rebound in one-dimensional tasks could be the result of the interaction of two separate memories: one with a more stable behaviour and driven by a slow process (memory of field A; z component of adaptation), and one with a fragile and transient state which is driven by a fast process (memory of field B; x component of adaptation). We were able to dissociate these memories by projecting them into two orthogonal dimensions.

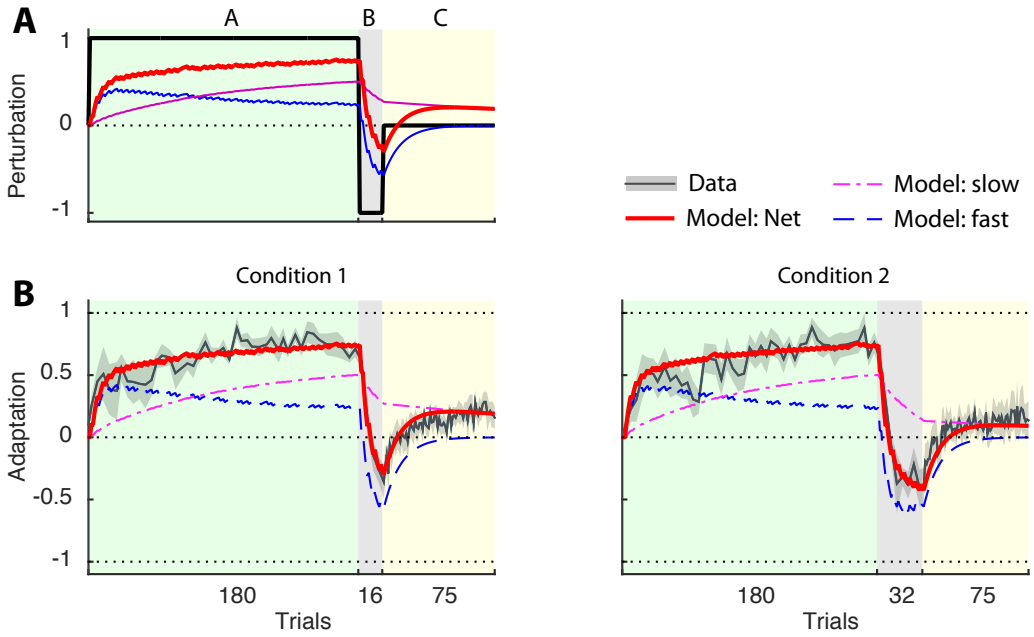

Fig. 3.6 **A.** An example simulation for the one-dimensional spontaneous recovery paradigm, with fields A and B in opposite directions. **B.** Comparing the simulation results with the superposed data resulted from the subtraction of the x component of adaptation (memory of field B) from the z component of adaptation (memory of field A).

3.4 Discussion

In this study we examined the phenomenon of spontaneous recovery in a two-dimensional force field learning task. We showed that spontaneous recovery happens even when the secondary force field (B) is not counteracting the memory of the initial learning. Instead, it seems as though the spontaneous recovery is the tendency of the motor system to transition from the unstable, fragile motor memory, towards the more stable and robust motor memory. This was observed in our data as the rotation of adaptation vector from the direction of the second force field towards the direction of the initial force field.

In one-dimensional force field studies, spontaneous recovery could be explained based on two different hypotheses. One suggesting that there is a single motor memory formed for the adaptation of the initial force field (field A), while the exposure to the opposing force field (B) only accelerates the unlearning of this memory. The spontaneous recovery then appears as the sudden increase in the activity of the initial memory, after it was washed out (apparent

extinction). This hypothesis mainly relies on interference studies in which the memory of an initial learning is washed out by the learning of the following opposite force field (retrograde interference; e.g., Caithness et al., 2004). By contrast, other studies suggest that the memory of the initial learning is protected, and only masked by a competing memory formed for the force field perturbation in the opposite direction (Criscimagna-Hemminger and Shadmehr, 2008; Pekny et al., 2011). According to this view, the effects of the initial memory could be revealed if enough time is allowed for the competing memory to decay, and thus unmask the initial memory.

Our results in the two-dimensional force field task are mainly consistent with the latter hypothesis. It seems that the overall adaptation behaviour is the result of the competition between two motor memories: one that is more robust due to longer training (field A), and one that is formed in a short time and thus is more fragile (field B). We could distinguish the behaviour of each motor memory by projecting each into separate (orthogonal) dimensions. During the error-clamp trials, the memory of field A was almost constant throughout and was never fully washed out. Its contribution to the adaptation vector, however, was undermined by the memory of the most recent adaptation (field B) in the beginning of the error-clamp phase. Over the trials, the memory of field B underwent a rapid decay, unmasking the memory of field A. This revealed itself as the rotation of the adaptation vector from the direction of field B towards the direction of field A. In other words, the adaptation behaviour shifted from a fragile state towards a more robust state. This might as well be the case in the one-dimensional spontaneous recovery. That is, the observed rebound of adaptation may not be due to the sudden increase in the activity of the initial memory (memory of field A), but due to the rapid decay of the competing motor memory (i.e., memory of field B).

The aforementioned view implies that the strength of spontaneous recovery (that is, the rebound of adaptation angle) might depend on the fragility (rate of decay) of the recently formed motor memory. According to the predictions of the dual-rate model, this fragility is modulated by the duration of training under field B. As such, longer training would increase the contribution of the slow process in the memory of field B, which in turn reduces the speed of decay for this memory during the error clamp phase. As a result, spontaneous recovery

will occur slower or to a lesser extend. We tested this prediction by comparing the rebound of adaptation angle between condition 1 ($N_B = 16$) and condition 2 ($N_B = 32$). The data, when averaged across subjects, seemed to be in agreement with this prediction, although, its significance was not quite found due to the large noise in the data.

In summary, our results showed a more general form of spontaneous recovery that had not been shown before, and its effect was accounted for by an extended (two-dimensional) dual-rate model.

Chapter 4

Interference and facilitation in sensorimotor learning

4.1 Introduction

Adaptation to novel sensorimotor transformations causes adjustments in the motor memory that are retained for weeks or even months after learning (Bock et al., 2001; Caithness et al., 2004; Krakauer et al., 2005). How these motor memories are formed, maintained or interfered with has been of significant interest in the sensorimotor learning research. In particular, two important properties of motor memories have been extensively investigated. First, it has been shown that relearning a sensorimotor transformation after it is forgotten occurs in a faster rate; a phenomenon known as savings (Bock et al., 2001; Brashers-Krug et al., 1996; Huang et al., 2011; Leow et al., 2016; Roemmich and Bastian, 2015; Zarahn et al., 2008). Also, when adaptation to a new transformation is followed, immediately or after a break, by an opposing transformation, the memory of the initial learning is prone to disruption; a phenomenon referred to as interference (Bock et al., 2001; Brashers-Krug et al., 1996; Caithness et al., 2004; Krakauer et al., 2005).

Studies have investigated various factors that might constitute the underlying mechanisms for interference or savings. The early studies suggested that interference occurs if two transformations share the same kinematic features (e.g., position or velocity), but each

associated with opposite motor outputs. For example, learning a position-dependent force field could interfere with the learning of a position-dependent visuomotor rotation (in the opposite direction), because both transformations, despite being structurally different, depend on the same kinematic parameter (i.e., position; Tong et al., 2002). However, an acceleration dependent dynamic transformation may not interfere with a position dependent visuomotor rotation (Krakauer et al., 1999), as they involve different kinematic variables in the task. Further studies, however, showed that interference happened even when two transformations involved different kinematic variables (Bays et al., 2005): the learning of a position dependent force field was disrupted by the learning of an oppositely directed velocity dependent force field.

Recently, studies on error-based learning have proposed that the similarity of performance errors caused by two successive transformations, regardless of their kinematic dependencies in the task, determine whether interference or facilitation would be seen. Computational approaches suggested that the similarity of observed errors in the second transformation with that of the first could lead to facilitation in learning (Herzfeld et al., 2014; Takiyama et al., 2015). Whereas, discrepancy between the observed error in the second transformation and the expected error (based on the first transformation) would cause interference (Takiyama et al., 2015). The representation of error in the motor system and its role in the formation and interaction of motor memories, however, is poorly understood. In previous studies conducted mainly by a planar manipulandum, error was normally measured as a scalar defined in a single dimension (perpendicular to the reach direction). However, naturalistic human movements in a 3D space usually lead to errors that are characterised not just by their magnitude, but also by their spatial direction. For example, reaching to a target in the front in a 3D space could be disturbed horizontally (left-right) or vertically (up-down), leading to error vectors in different directions. How different error vectors, experienced during different transformations, modulate the interaction of motor memories (that is, facilitation or interference) is still unknown.

Here, by using the 3BOT, we are able to introduce force field transformations in various directions, causing error vectors in a 2D space (perpendicular to the movement direction). We

examined how the relative angle between the error vectors experienced in two different force field perturbations predicted savings or interference in the corresponding motor memories. Specifically, we examined three main cases, in which two force fields were applied either in the same direction, in the opposite directions, or in the orthogonal directions. The first two cases have been tested extensively in previous studies to address savings and interference. The last case, however, is a novel paradigm tested for the first time in this study and is only possible in a 3D setup. Here we discuss how facilitation and interference between two successive transformations are tuned to the relative angle between the errors caused by each transformation.

4.2 Methods

32 subjects were recruited to participate in this experiment (14 female; age: 24.2 ± 5.7 years). All subjects were right-handed with no reported record of neurological disorders. Informed consent was provided by all subjects prior to the experiment. The experiment was approved by the institutional ethics committee.

4.2.1 Experimental procedure

Apparatus

The 3BOT was used in this study (Fig. 4.1A). we upgraded the virtual reality system (previously a 3D screen and active 3D glasses) with an Oculus rift (Development Kit 2) which was installed in a fixed location above the robot workspace (Fig. 4.1A). Subjects were seated comfortably on an adjustable chair and their head was fixed according to the location and orientation of the Oculus rift. Visual feedback of the task was provided at 75 Hz with the resolution of 960×1080 per eye, and subjects had a 100° field of view which fully encompassed the task workspace.

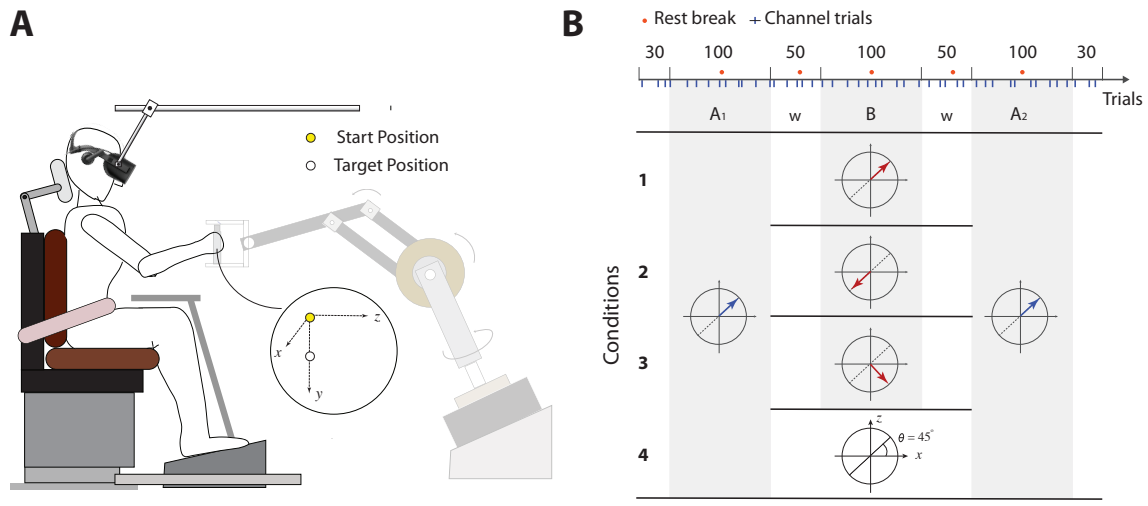

Fig. 4.1 **A.** The 3BOT was used to conduct the experiment. Three-dimensional visual feedback was generated by an Oculus Rift virtual reality system (refresh rate: 75 Hz, resolution of 960×1080 per eye). The task was to perform vertical reaching movements along the y-axis from the start position to the target position. **B.** The experiment consisted of a sequence of blocks in which subjects were exposed to an initial force field (A_1), followed by a secondary force field (B), and followed by re-exposure to the initial force field (A_2; the grey background patch indicates the exposure blocks). A washout period of 50 trials (w) separated the exposure blocks. Four different conditions were used in this experiment. In each condition, we manipulated the direction of field B with respect to the initial force field (the small arrows show the direction of each force field in the xz plane). In condition 1, field B was in the same direction as the initial field (B=A); in condition 2, field B was in the opposite direction as the initial field (B=-A); in condition 3, field B was orthogonal to the initial field; and in condition 4, field B consisted of null trials (no force field applied).

Task

Subjects grasped the handle of the 3BOT and performed vertical reaching movements (along the y-axis; Fig. 4.1A) from a specified start position in the 3D workspace (represented by a yellow sphere, 0.75 cm radius), to a single target position (represented by a white sphere, 0.75 cm radius) located 12 cm below the start position. Each trial started with maintaining the cursor (a red sphere 0.5 cm radius representing the veridical hand position in the virtual environment) within the start position. The movement initiated with the appearance of the target followed by a tone, after which subjects were required to make a rapid reaching movement towards the target location. The movements were required to be finished within $600 - 1100$ ms, otherwise an appropriate message indicating "too slow"

or “too fast” appeared on the screen. The movement was finished when subjects remained stationary within the target area for 100 ms (positional tolerance 0.75 cm). Subjects then received a “Relax Arm” message, after which the robot passively returned subject’s hand back to the home position (using a minimum jerk motion) for the next trial.

Three types of trials were used in our experiment: null trials (no perturbations applied), error-clamp (channel) trials (movement was constrained by a simulated channel tube that connected the start position to the target position), and exposure trials (velocity dependent force field perturbations applied ; Fig. 1.2D). The force field was generated perpendicular to the movement direction in the *xz* plane as shown below:

$$\begin{aligned} \mathbf{f}(\theta) &= b \cdot \mathbf{R}(\theta) \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{bmatrix}, \\ \mathbf{R}(\theta) &= \begin{bmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{bmatrix} \end{aligned} \tag{4.1}$$

where, $b = 15$ Ns/m was the field constant, θ represented the angle of the force field in the *xz* plane (measured from the *x*-axis), and $\dot{x}$, $\dot{y}$, and $\dot{z}$ were the components of the velocity vector represented in the main frame (Fig. 4.1A). Here, a rotation matrix $\mathbf{R}(\theta)$ was used to adjust the angle of the force field in the *xz* plane. For example, for $\theta = 0^\circ$ the force field was applied in *x* direction, and for $\theta = 90^\circ$ it was applied in the *z* direction.

Experimental paradigm

The experiment consisted of four different conditions in which we examined savings and interference when adapting to force fields in different directions. Separate groups of subjects were used for each condition (n=8 per group). In each group, a paradigm A-w-B-w-A was used in which subjects performed under three exposure blocks, starting with an initial force field A, followed by a secondary force field B, and back to the initial force field A. From

here on, we denote the first and second exposure to field A as A_1 and A_2, respectively. The exposure blocks were separated by washout periods (w) in between (Fig. 4.1B). The direction of field B was manipulated in each condition with respect to the direction of field A in the following ways: B was in the same direction as A (condition 1), B was in the opposite direction as A (condition 2), B was orthogonal to the direction of A (condition 3), and B was a null field (condition 4).

In the first and second conditions we examined the effects of savings and retrograde interference, respectively. In condition 3 we tested a novel case in which we examined how the memory of field A was influenced by the following adaptation to an orthogonal field (B). Condition 4 was also used as a control condition in which no secondary force field interfered with the memory of field A (B=Null).

In each condition, the direction of field A was counterbalance between $\theta_A = 45^\circ$ (Fig. 4.1B) and $\theta_A = -135^\circ$ across subjects (half of the subjects in the group performed under $\theta_A = 45^\circ$ and the other half under $\theta_A = -135^\circ$; the performance measures were then properly combined across both cases for the analysis). Accordingly, the angle of field B was determined as: $\theta_B = \theta_A$ for condition 1, $\theta_B = \theta_A - 180^\circ$ for condition 2, and $\theta_B = \theta_A - 90^\circ$ for condition 3 (field B was a null field in condition 4).

Prior to the experiment, each subject performed 40 null trials to get familiar with the task and the virtual reality environment (these trials were not analysed). The experiment in each condition then started with 30 baseline trials, including 27 null trials as well as 3 interleaved channel trials. Subjects then performed the sequence of exposure and washout blocks, A_1-w-B-w-A_2. The exposure blocks (A_1, B, and A_2) each consisted of 90 exposure trials as well as 10 pseudo-randomly interleaved channel trials. The washout periods (w) also included 45 null trials and 5 pseudo-randomly interleaved channel trials (Fig. 4.1B).

4.2.2 Analysis

Two performance measures were used in the analysis of behaviour during this experiment. First, in each trial we recorded the time course of the hand position (at 1000 Hz) and obtained the maximum perpendicular error (MPE) between the hand trajectory and the straight line

that connected the start and target positions. The value of MPE in each trial was regarded as the kinematic error of performance in that trial.

For each channel trial, we measured the time course of the force vector generated during the movement (in the *xz* plane) in compensation for the force field (collection rate 1000 Hz). The measured force vector at each time point was projected into two dimensions, one inline with the direction of field A (as a reference direction), and one orthogonal to it. We also measured the velocity of movement on that trial ($\dot{x}$, $\dot{y}$, and $\dot{z}$) and obtained the force trajectory that subjects would have experienced had they moved with the same velocity in an exposure trial (ideal force trajectory; equation 4.1). The ideal and measured force trajectories in each dimension were then trimmed based on the speed of movement, from when the movement speed first exceeded 5% of its maximum value, until it reduced to less than 5% of its maximum value (that is after the peak velocity). Finally, we regressed (without intercept) the trajectory of the measured force onto the trajectory of the ideal force and used the regression coefficient as an index of adaptation to the force field.

4.3 Results

The MPE and adaptation index were analysed for each condition separately. In each condition, we mainly focused on the learning behaviour during the first and second exposure to field A (A_1: initial learning, and A_2: relearning), as a function of adaptation to an intermediate force field (B). More specifically, in the first two conditions we looked at the effects of savings and retrograde interference on the learning memory of field A, while in condition 3, we examined the characteristics of the interaction between two orthogonal force fields. Condition 4 was also used as a control condition in which the memory of field A was not interfered with any intermediate force field.

Fig. 4.2 illustrates the learning behaviour in each condition in terms of the adaptation index (inline and orthogonal components; left column) as well as the magnitude of MPE (right column). As shown, during A_1 and A_2 blocks, for all conditions, the inline component of adaptation (orange line) showed an increase in the positive direction, indicating adaptation

to field A. The orthogonal component of adaptation in these blocks (blue line) remained unchanged and close to baseline, implying near ideal performance in this dimension (no force field was applied in the orthogonal dimension in these blocks). During field B, different adaptive behaviour was observed in each condition depending on the type of force field: in condition 1, adaptation was similar to blocks A_1 and A_2; in condition 2, the inline component of adaptation increased in the negative direction, indicating adaptation to a reversed (opposing) force field; in condition 3, the inline component remained close to baseline, and instead the orthogonal component increased in the positive direction (adaptation to orthogonal field); and in the final condition, both inline and orthogonal components of adaptation remained flat and close to zero (null field).

We examined the adaptive behaviour in the A_1 and A_2 blocks more closely to highlight the effects of the intermediate field B on the memory of field A. First, to make sure that the adaptation behaviour is correctly comparable across blocks A_1 and A_2, it is important that the adaptation value in each block starts from a similar level; thus, we examined the value of adaptation prior to the onset of each block. To this end, we took the last channel trial from the washout periods preceding A_1 and A_2, and compared the level of adaptation in these trials. As shown in Fig. 4.3A, in each condition, the level of adaptation right before the start of the A_1 and A_2 was close to the baseline (t-test against zero for every block and every condition: $|t(7)| < 1.4$, $p > 0.204$). Next, we compared the adaptation process between A_1 and A_2 blocks in two different ways: 1. the completeness of adaptation (how much have been learned), and, 2. the rate of adaptation (how fast the learning happened). In the following, we explain each case in detail.

Completeness of adaptation

For each condition, we examined the final level of adaptation acquired in blocks A_1 and A_2 to see how the level of learning differed between the blocks as a function of field B. In Fig. 4.3B the average of inline adaptation over the last five channel trials (the second half) of each block is shown for A_1 and A_2. In the first condition, where the force field in block B was the same as A_1 and A_2, we observed that the adaptation was larger for the A_2 (relearning)

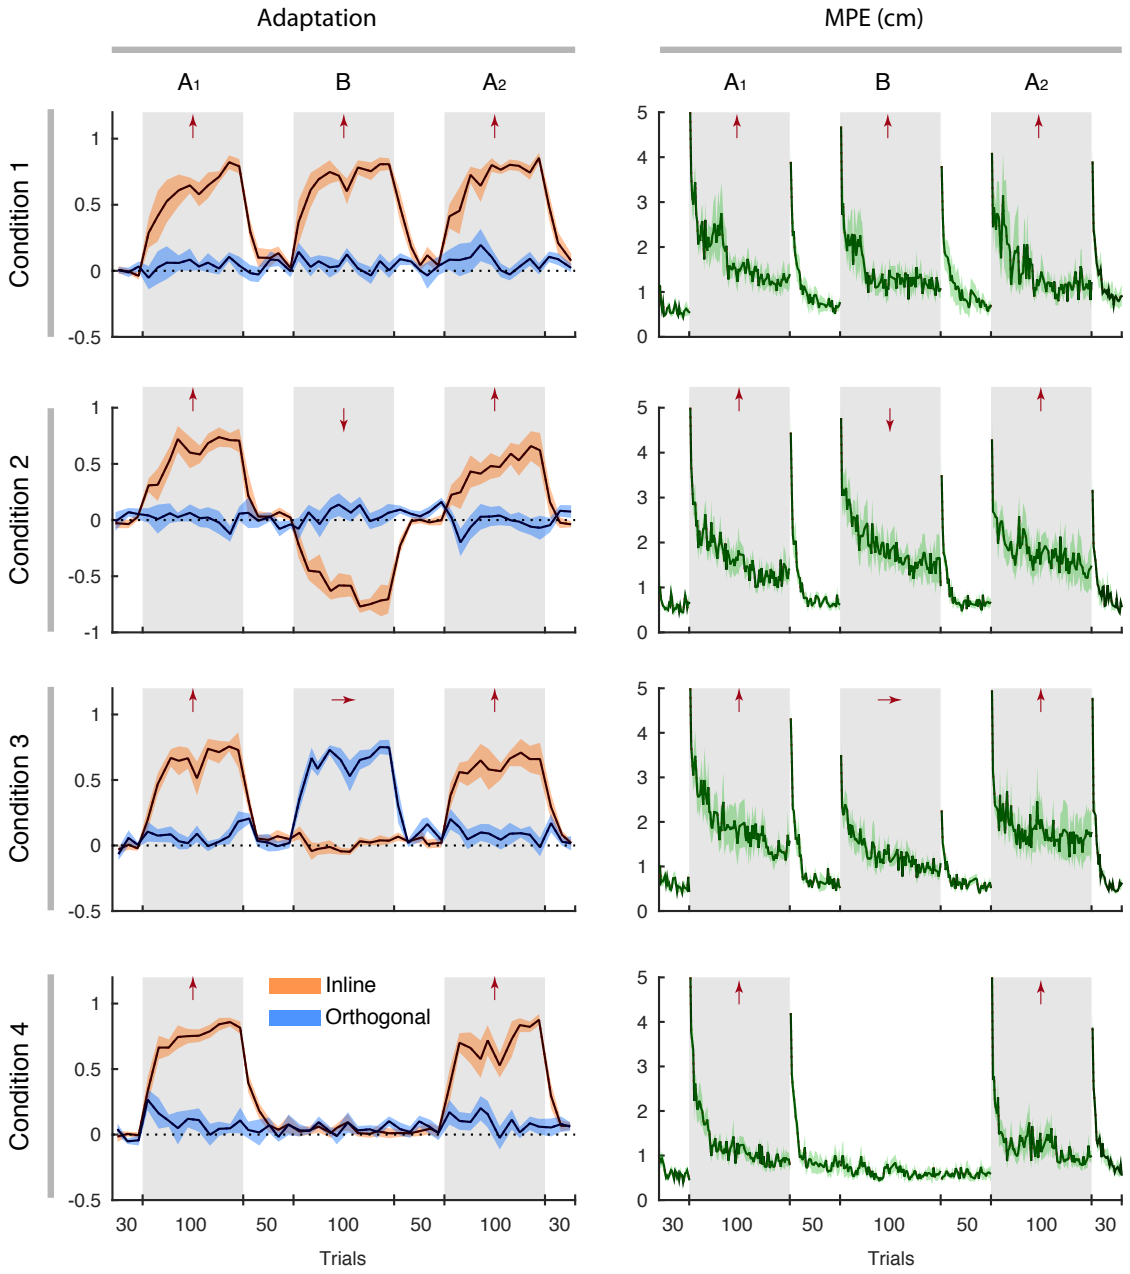

Fig. 4.2 Adaptation and MPE. The time course of adaptation index (left column) is shown for both inline (orange) and orthogonal (blue) components for each condition. On the right column, the magnitude of MPE is also illustrated over the trials. The small arrows on top of each panel show the relative direction of force field between the blocks A_1, B and A_2. In condition 1, all force fields were in the same direction; in condition 2, field B was in the opposite direction as field A; in condition 3, field B was directed orthogonal to field A; and in condition 4, no force field was introduced in block B. The background grey patch shows the exposure to force field, and the white patch indicates the washout trials.

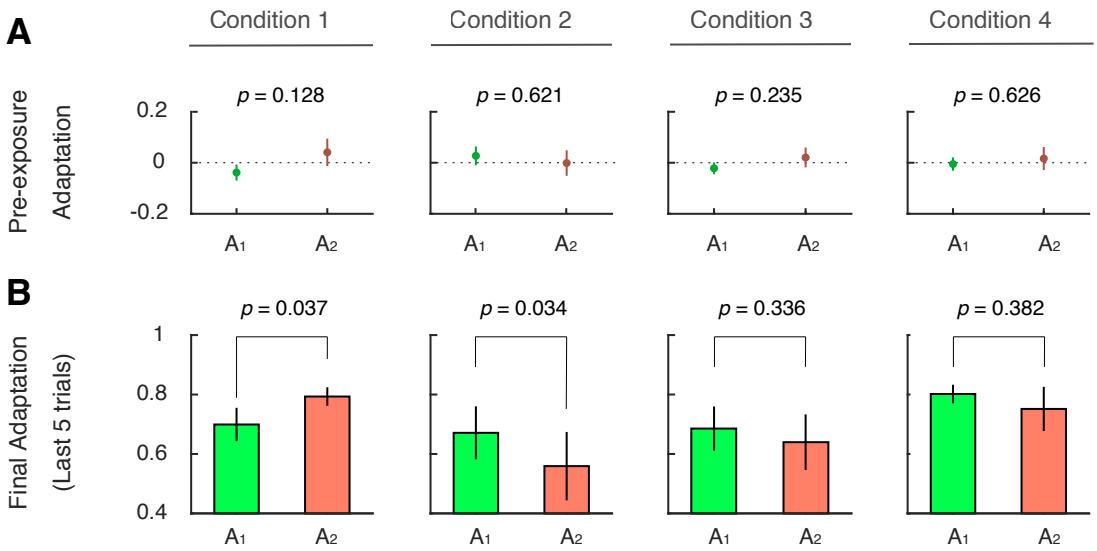

Fig. 4.3 **A.** The adaptation level prior to the onset of A_1 and A_2 (the last channel trial in the preceding washout period). Before the start of each block, the adaptation level was close to the baseline for all conditions with no significant difference between the blocks. **B.** The final adaptation level (averaged across the last five channel trials) in A_1 and A_2 blocks shown for all conditions. All error bars indicate standard error (SE) across subjects. The p-values are obtained by paired t-test between the blocks.

compared to A_1 (initial learning; paired t-test: $t(7) = -2.58$, $p = 0.037$). This larger amount of relearning could be attributed to the effects of savings in this condition. In the second condition, subjects were exposed to an opposing force field during field B, a paradigm that invoked retrograde interference. As expected, the level of adaptation in A_2 was observed to be smaller than the level of adaptation in the initial learning, A_1 (Fig. 4.3B; paired t-test: $t(7) = 2.62$, $p = 0.034$), indicating a clear effect of interference caused by the opposing force field, B, on the memory of the initial force field, A.

In the third group, we designed a novel condition to test the effects of an orthogonal force field (B) on the memory of the initial learning (A). We found that the final level of adaptation between A_1 and A_2 in this condition was roughly the same (paired t-test: $t(7) = 1.03$, $p = 0.336$, Fig. 4.3B). This implied that the orthogonal force field in block B neither caused interference nor facilitation in the memory of field A. However, to better examine the interference or facilitation effects of orthogonal force fields, we contrasted the results in group 3 with that of group 4, in which no force field exposure was introduced during

block B. As illustrated in Fig. 4.3B, quite similar to condition 3, we found no significant difference between the final level of adaptation in A_1 and A_2 for the forth condition (paired t-test: $t(7) = 0.934$, $p = 0.382$). These observations indicated that performing under the orthogonal force field during block B had no different effect on the memory of the initial learning (field A) than performing under null trials with no transformation exposed. This could mean that when two force fields are exposed in (spatially) orthogonal dimensions, the resultant motor memories might also be formed orthogonal to one another (that is, with no interaction or overlap between the memories).

Adaptation rate

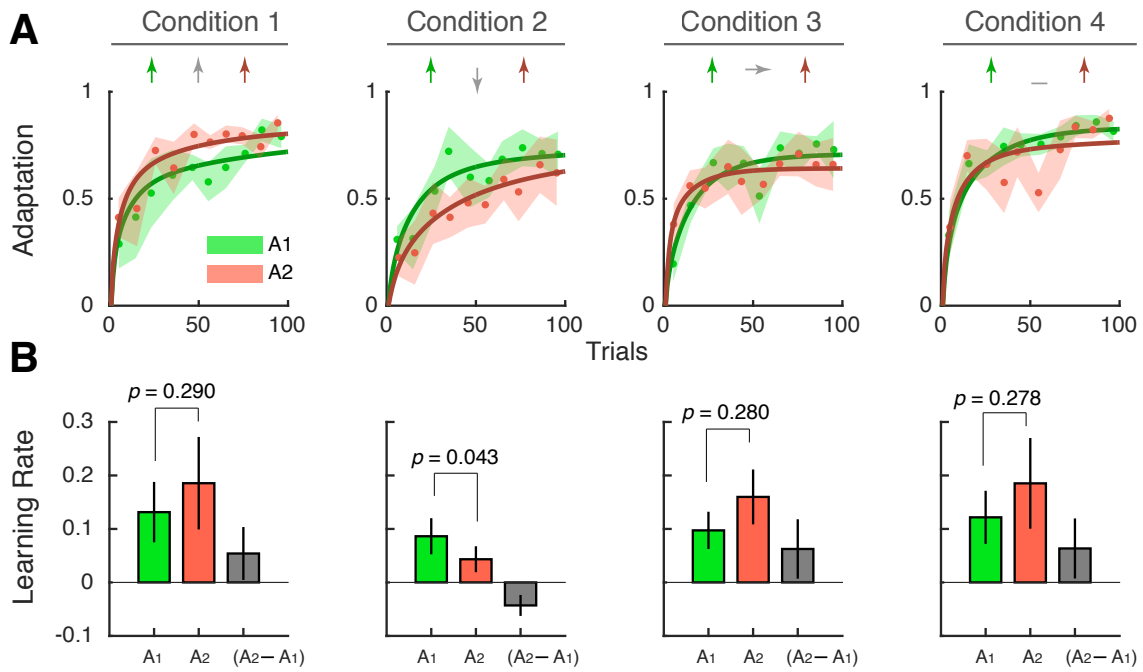

Fig. 4.4 **A.** The inline component of adaptation in blocks A_1 and A_2 shown with the fitting results (solid lines) of the single-rate state-space model. The small arrows on top of each panel indicates the relative direction of force fields in blocks A_1, B, and A_2 (respectively from left to right). **B.** The learning rate parameter obtained from fitting the single-rate model is shown for A_1 and A_2 blocks. The difference between the learning rates of A_2 and A_1 (A_2-A_1) represents the amount of retention of learning. Error bars in all plots indicate the standard error of the mean across subjects. The p-values are obtained based on paired t-test between the blocks.

In the second analysis, we examined the rate of learning between A_1 and A_2 blocks as a function of the intermediate force field (B) in different conditions. To obtain the learning rate, for each subject and for each block (A_1 and A_2), we fitted a single-rate state-space model to the inline component of adaptation:

$$\begin{aligned} z^{(n+1)} &= \alpha \cdot z^{(n)} + \beta \cdot (1 - z^{(n)}) \\ z^{(1)} &= 0 \end{aligned} \tag{4.2}$$

where, $z^{(n)}$ represented the adaptation index on trial n within each block, and α and β were, respectively, the forgetting and learning rates considered as the free parameters in the model. The parameters of fit were then obtained by minimising the sum of squared error between $z^{(n)}$ and the adaptation level in the data.

Fig. 4.4 shows the fitting results for A_1 and A_2, including the obtained learning rates for each one (parameter β). In each condition, we took the difference between the learning rates of blocks A_2 and A_1 for each individual ($\beta_{A_2} - \beta_{A_1}$), and considered the outcome as the amount of retention of learning (Fig. 4.4B): a positive value for this difference represented savings, and a negative value indicated interference. As shown, in condition 2, we observed clear interference between A_2 and A_1 as the difference between the corresponding learning rates was significantly negative (paired t-test: $t(7) = 2.49$, $p = 0.043$; Fig. 4.4B).

For the conditions 1, 3 and 4, the difference in learning rate was generally positive, indicating faster relearning in all three conditions, although the significance of increase was not found (paired t-test: $t(7) = 1.14$, $p = 0.29$ for condition 1, $t(7) = 1.17$, $p = 0.28$ for condition 3, and $t(7) = 1.24$, $p = 278$ for condition 4). This showed only a weak effect of savings in terms of faster relearning, which was not modulated by the intermediate field (B) across these conditions. It might be that the similarity of results between condition 1 (B=A) and condition 4 (B=Null) was due to a ceiling effect on the learning rate, such that it could not be further improved in condition 1 (compared to condition 4), despite experiencing the same force field (A) in three consecutive blocks.

On the other hand, comparing condition 3 (B orthogonal to A) and condition 4 (B=Null), the similarity of increase in learning rate could imply that the orthogonal force field did

not cause interference with the memory of field A, as any interference would have caused a smaller improvement in learning rate compared to condition 4. Also, it seems unlikely that the slight increase in the learning rate in condition 3 was due to the experience of the orthogonal field during block B, as the same amount of increase is also observed for condition 4 with only null trials in block B.

Overall, the results from both analyses (completeness of adaptation and rate of adaptation) illustrated that, consistent with previous studies, savings and interference happened when field B was in the same direction, or opposite to field A, respectively (conditions 1 and 2). Also, found in this study for the first time, we observed that learning an orthogonal force field during block B had no marked effect of facilitation nor interference on the memory of the initial learning.

4.4 Discussion

In this chapter, we examined the facilitation and interference effects when adapting to successive force fields (A-w-B-w-A) with different spatial directions. Four conditions were considered in which the direction of field B was either the same, opposite, orthogonal or neutral (null) to the direction of field A. We examined the difference between the first and second encounter to field A (A_1 and A_2) and used two different measures to quantify interference and facilitation in each condition: the final level of adaptation, and the learning rate.

In previous studies, the savings and interference phenomena were extensively examined when two transformations were, respectively, in the same or in the opposite directions (for example, Caithness et al., 2004; Krakauer et al., 2005). Here, using the 3BOT, we tested the similar cases in our experiments in conditions 1 and 2, to compare the results with those of other studies. We found clear effects of interference in condition 2 for both measures, consistent with previous findings in literature. We also found significant effect of savings in condition 1 based on the completeness of adaptation, but less so based on the rate of adaptation. This lack of savings effect on the learning rate may rise from various

differences between the current study and the studies that have shown faster relearning in similar paradigms. Mainly, in previous studies movements were constrained in the horizontal plain with no degrees of freedom orthogonal to the plane of movement. In this study, we used the 3BOT for conducting the experiment and subjects performed vertical movements with more degrees of freedom in the 3D space. This could have caused more variability in the performance (the adaptation measures for individuals were quite noisy in our data), and this might have masked the significance of the increase in the learning rate.

In condition 4 of our study (B=Null), we found no significant difference between the adaptive behaviour in block A_1 and A_2, indicating no effect of savings in this case. Different studies in the past have reported different results as to the effect of washout trials in between two exposures to the same sensorimotor transformation. In some studies of visuomotor rotation, people have shown that savings happen even if relearning of the task is preceded by a washout period (Berniker and Kording, 2011; Zarahn et al., 2008). Whereas, in some other studies, the washout periods between the initial and second encounter to the task have shown substantial disruption to the retention of learning (i.e., no savings; Caithness et al., 2004; Roemmich and Bastian, 2015). Our results in condition 4 are more consistent with the latter group of studies. It seems that a long period of washout (200 null trials in our case) between the first and second encounter to field A has caused complete forgetting of the initial learning (note that the washout period was twice as long as the initial learning session with 100 trials).

The main focus of the current study was to examine the facilitation/interference effects of two orthogonal force fields (condition 3). Two different hypotheses could be considered regarding how the memory of orthogonal force fields might interact. First, it has been argued that fundamental phenomena such as savings and interference are not merely the product of error-based adaptation, but instead rise from model-free processes that involve action selection (Huang et al., 2011). As such, savings is attributed to the recall of an aiming strategy rather than remembering familiar errors (Morehead et al., 2015). This view is also supported by visual tracking studies in which learning two independent transformations (that is, neither opposing nor similar in action) have produced facilitation (Bock et al., 2001). According to this view, one could argue that in the case of orthogonal force fields (where required actions

are neither opposing nor of the same direction), once a strategy is learned during adaptation to an initial force field, it can be recalled and applied during adaptation to the orthogonal force field. Therefore, it is expected that two orthogonal force fields produce savings.

Our results, however, were not consistent with this hypothesis as we did not find any evident effects of savings or interference in the learning of two orthogonal force fields. Importantly, we observed that learning an orthogonal field in block B showed no different effects on the memory of field A than performing under a session of null trials. Instead, it seems that our results are more compatible with a different hypothesis which emphasises the role of experienced errors in the savings and interference phenomena (Berniker and Kording, 2011; Herzfeld et al., 2014; Leow et al., 2016; Smith et al., 2006; Takiyama et al., 2015). It has been shown that the prior experience of error is necessary and sufficient for savings (Herzfeld et al., 2014; Leow et al., 2016); that is, savings happens when errors caused by two transformations occur in the same direction. Also, interference has shown to be the product of the discrepancy between the previously experienced errors and the currently observed errors (Takiyama et al., 2015). This indicates that interference happens when two successive transformations lead to errors in opposite directions. Our results in conditions 1 and 2, were indeed consistent with these findings.

However, most of these studies examined only a special case in which the errors experienced by different transformations were either in the same direction, or in opposite directions. As such, no representation or predictions were provided for the vectorised concept of error, where two error vectors might neither be in the same direction nor in opposite directions. Here, we examined this novel case in our experiments. Incorporating our findings with previous error-based studies, one could propose a framework based on which when two error vectors (caused by two different transformations) form an acute angle (i.e., their dot product is positive), it results in savings. Likewise, if two error vectors form an obtuse angle (their dot product is negative), it would result in interference. And finally, if two error vectors are orthogonal or one of the error vectors is null, no effect of savings or interference would be observed.

This framework suggests a particular tuning of interference/facilitation behaviour to the relative direction of error vectors experienced in successive transformations. Recently, we proposed a cosine-shaped tuning of adaptation behaviour in an object manipulation study, whereby interference and facilitation were modulated based on the cosine of the relative angle between the experienced error vectors in two transformations (Ingram et al., 2017). We showed that a cosine tuning provided an optimal representation of how two successive transformations lead to savings, interference, or remain independent. It is plausible that savings and interference in force field learning is also modulated by the cosine of the relative angle between experienced error vectors. As such, the largest amount of savings happens when two error vectors are aligned in the same direction (relative angle of zero), and it transitions towards interference as the relative angle increases towards 180°, while passing through the neutral angle of 90° in which no savings or interference occurs. This framework can be further investigated in future studies.

Part III

INTERACTION BETWEEN LEARNING PROCESSES

Chapter 5

Adaptation to novel dynamics

5.1 Introduction

The sensorimotor system is adept at using contextual information to adapt to systematic perturbations (Wolpert and Kawato, 1998). This enables the motor system to learn multiple (sometimes opposing) tasks simultaneously by associating each with a separate context. For example, subjects have shown the ability to learn two opposing force fields when cued with separate contextual information for each one, such as different hand positions or appropriate visual cues (Howard et al., 2013). Many studies have investigated how multiple motor memories are tuned to different contextual information and how context is used to activate appropriate motor commands (e.g., Imamizu et al., 2007; Ingram et al., 2011; Lee and Schweighofer, 2009; Thoroughman and Shadmehr, 2000; Wolpert and Kawato, 1998). Little is known, however, about how motor memories (or internal models) associated with different contexts interact with one another. In other words, how the adaptive changes in the motor memory of one context influences the motor behaviour in a different context.

In general, the interactive properties of context-dependent motor memories can be found in two main groups of studies. In generalisation studies, it has been found that when a sensorimotor transformation is learned in one context (e.g., one movement direction), it locally affects the behaviour in other nearby contexts (Donchin et al., 2003; Howard and Franklin, 2015; Ingram et al., 2010; Sarwary et al., 2015). The extent of this behavioural change

depends on how different the contexts are, and indicates an interaction in the form of transfer of learning between the learning mechanisms associated with each context. A different group of studies, mainly examine the interaction of context-dependent motor memories through the study of interference. It has been shown that motor memories for two opposing force fields can be separated using appropriate contextual cues, and that the interaction between these motor memories can affect the amount of interference between them (Howard et al., 2012, 2013, 2015; Sheahan et al., 2016). For example, motor memories separated based on two different visual cues have stronger interaction (larger interference) than motor memories that are separated based on different hand positions (Howard et al., 2013). The underlying mechanisms governing such interactions, however, are less understood.

In this study, we examine the interaction of context-dependent motor memories during a force field learning task. We designed a novel experimental paradigm based on reciprocating movements. These are back and forth movement cycles of the same magnitude but in opposite directions, wherein each cycle might be associated with a different dynamics in the task or the environment. For example, when rowing or swimming, the movement of the oar or hand consists of two cycles: one that is performed under the water thrusting the boat or body forward (drive), and the other that happens above the water to relocate the oar or arm for the next cycle (recovery). Each movement cycle is performed in a different environmental condition (e.g., under the water versus above the water) and with different dynamic loads applied on the movement. In this case, separate motor memories are responsible to implement appropriate motor commands based on the dynamics of each movement cycle. The interaction between such motor memories can be examined by manipulating the dynamics of movement across the two movement cycles.

In our experiments subjects performed successive back and forth reach actions between two specified targets. In this case, each movement direction was considered as a separate context. We manipulated the dynamics of the reach using force field perturbations in one movement direction (context 1) and tracked the behavioural changes in the other (context 2). This allowed us to examine the interaction of the corresponding motor memories under various dynamic conditions. We found that motor memories associated with each movement

direction were initially independent with no evident interaction. However, when both movements were subject to force field perturbations simultaneously, an interaction developed between the corresponding motor memories, so that they shared error signals in their learning so as to potentially improve behaviour in future movements.

5.2 Methods

The study was approved by the local ethics committee and subjects provided written consent before participation. Fifty healthy right-handed subjects (26.9±3.9 years old; 29 female) took part in the experiment in five groups.

Subjects were seated at a virtual reality system while holding the handle of a robotic manipulandum (vBOT) which allowed them to move their arm on the horizontal plane (Howard et al., 2009). The manipulandum could impose haptic perturbations and constraints on the arm movement and measure movement kinematics and generated forces. The virtual reality environment was generated and displayed by a horizontal screen and was reflected via a horizontal mirror to provide subjects with the visual information associated with the task. The visual feedback about the task included a red circular cursor (5 mm diameter) representing the veridical position of the hand, as well as two circular targets (10 mm diameter) aligned 12 cm apart in the virtual workspace (T1 and T2 in Fig. 5.1). The task was to make reaching movements back and forth between the two targets.

Each trial started with one of the targets illuminated on the screen (home target), in which subjects were required to hold the cursor stationary. The second target (away target) then appeared on the screen, towards which subjects were cued (by a beep) to make a fast reaching movement; this was considered as the "Out" movement cycle. In the following trial, subjects reversed the reaching movement from the away target back to the home target; this was regarded as the "Return" movement cycle. In each movement cycle, subjects had to initiate the movement within one second of hearing the beep, otherwise the trial would start over. They were also required to finish the trial within 800 ms or else they would receive a warning message ‘Too slow’. After the completion of each trial, subjects remained stationary inside

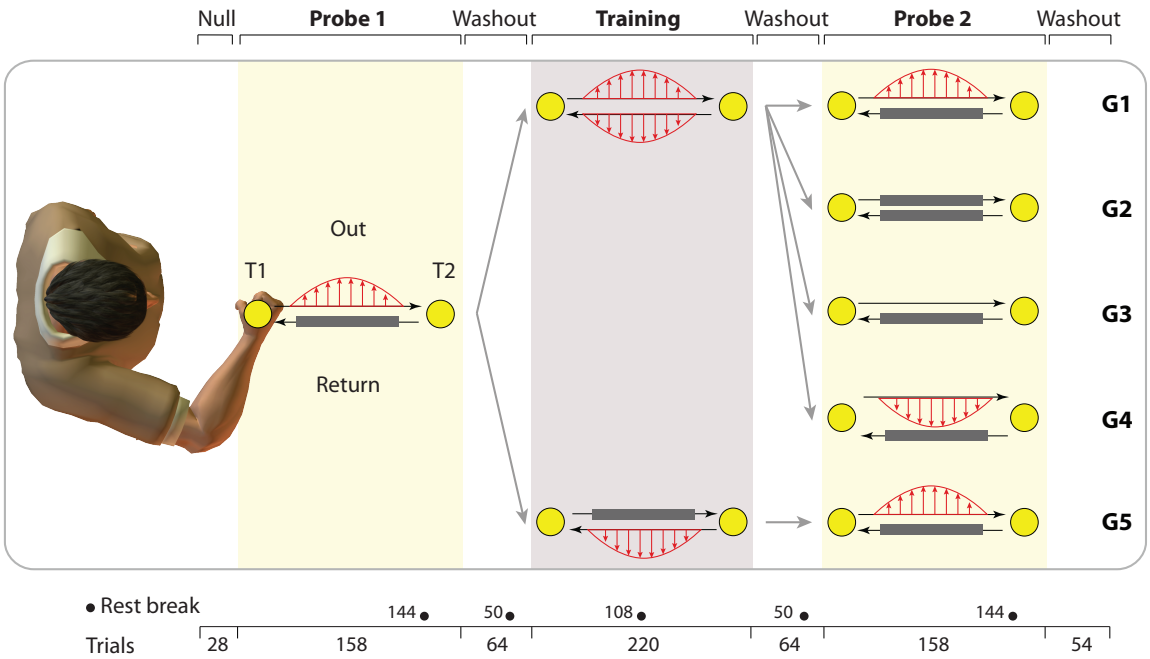

Fig. 5.1 Experimental paradigm. Subjects performed reaching movements back and forth between two targets (T1 and T2) as Out and Return movements. The experiment consisted of three main phases, including Probe 1 (P1), the Training phase, and Probe 2 (P2). Each phase was followed by a washout period, mainly consisted of null trials. The experiment involved five different conditions (G1 to G5) based on the dynamics applied in the Out and Return movements. Three types of task dynamics were used in the experiments, including, exposure (shown by lateral red arrows), error-clamp (shown by grey rectangles) and null dynamics (shown by thin black arrows connecting the two targets). Subjects received rest breaks within each phase shown by the black dots (associated numbers show the trial just before the rest break).

the target for 100 ms for the trial to finish, before receiving the cue for the next trial (note that the Out and Return movement cycles were discrete trials).

The trials continued as subjects made reciprocating Out and Return movements between the targets. The choice of the home and away targets was counterbalanced across the subjects in each experimental group, with half of the subjects starting with T1 (near target), and the other half starting with T2 (far target), as the home target.

5.2.1 Trial Type

Using the vBOT, we could impose various dynamics on subjects' hand during the movements. We used three different trial types; null trials, exposure trials, and error-clamp (channel) trials. In a null trial, subjects could freely move their arm in the workspace without any intervention from the robot. In exposure trials, we introduced a velocity dependent curl-field in which a force perturbation was applied on the subject's hand orthogonal to the direction of the hand's velocity. Specifically, the force vector was obtained as:

$$\mathbf{f} = b \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} \tag{5.1}$$

where, $\dot{x}$ and $\dot{y}$ represented the horizontal and vertical components of velocity, respectively, and the constant b determined the intensity of the field with $b = 15$ Ns/m for a CW field and $b = -15$ Ns/m for a CCW field. In an error-clamp trial, a virtual mechanical channel was applied that constrained the hand movement to a straight line between the two targets by simulating a stiff one dimensional spring and damper (spring coefficient, 6 kN/m; damping coefficient, 5 Ns/m). On these trials, the forces generated by subjects against the channel wall were measured as a proxy of adaptation to the curl-field.

We manipulated the trial type separately for Out and Return movements, so that they were either the same in both movements (i.e., both null, both exposure, or both error-clamp), or different for each movement.

5.2.2 Conditions

The aim of the experiment was to investigate the interaction of learning processes between Out and Return movement cycles. We focused on two main questions in designing each experiment. First, whether there was any a priori transfer of learning between Out and Return movements. And, second, whether the transfer of learning depended on the history of dynamic exposure in each movement.

Five different experimental conditions were designed. Each condition consisted of three main phases: 1. the initial probe phase (P1), in which we measured the extent of interaction that existed a priori between the Out and Return motor memories; 2. the Training phase, in which we tested how simultaneous dynamic exposure in both Out and Return movements altered the interaction between the corresponding motor memories; and 3. a second probe phase (P2), in which we tested the properties of such interaction under various dynamic manipulations of the task (Fig. 5.1).

Five groups of subjects participated in the experiment (n=10/group). All groups underwent a familiarisation session in which they performed 40 null trials in both movement directions prior to the onset of the main experiments. These trials were not analysed. The main experiments started with 14 pairs of Out-Return null trials followed by the three phases of P1, the Training phase and P2 as follows.

During P1, all groups of subjects performed 79 pairs of Out-Return trials (158 trials in total), in which subjects adapted to the force field exposure during Out movements (64 exposure trials with 15 error-clamp trials pseudo-randomly located in between), while the Return movements were all error-clamps (Fig. 5.1. Probe 1). During this phase we assessed whether the adaptation process to the force field in the Out movement direction affected the adaptive behaviour in the Return direction. Note that in P1 subjects never experienced a force field during the Return trials. This phase was followed by a short washout period consisting of 30 pairs of null trials followed by 2 pairs of error-clamp trials in both cycles (64 trials in total).

In the Training phase, subjects from groups 1 to 4 were exposed to force field perturbation simultaneously in both Out and Return movements. This phase consisted of 90 pairs of exposure trials (in both Out and Return directions) with 20 pairs of error-clamp trials pseudo-randomly placed in between. In contrast to the groups 1 to 4, in group 5, instead of simultaneous adaptation in both movement directions subjects only adapted to the force field in the Return movements, while the Out movements were all under error-clamp trials (Fig. 5.1G5). This way the adaptive memory for the Return trials was formed in isolation without any concurrent experience of force field for both movements. We could therefore

contrast the effects of simultaneous adaptation observed in groups 1 to 4, with the condition in which subjects never experienced the concurrent adaptation (group 5). The Training phase was followed by a short deadaptation period including 30 pairs of null trials and two pairs of error-clamp trials to bring the adaptive behaviour back to the baseline.

Finally, during P2 we examined the interactive properties of Out and Return learning memories after having experienced the Training phase. For groups 1-4 this allowed us to examine whether concurrent adaptation process affected the subsequent interaction between the underlying adaptive mechanisms in Out and Return movements. To do so, we manipulated the task dynamics in the Out movements of P2 for each group, and measured the adaptive response in the Return movements using error-clamp trials.

In group 1, P2 was identical to P1 with exposure trials in the Out and error-clamp trials in the Return directions (Fig. 5.1G1). In this group we could examine the effects of adaptation during Out trials on the adaptive behaviour during Return trials, both before the Training phase and after (i.e., in P1 and P2).

In groups 2 and 3, we aimed to dissociate the effects of interaction observed in the Return trials of P2, from other adaptive features in sensorimotor learning, such as spontaneous recovery. Smith et al. (2006) demonstrated that when subjects performed under error-clamp trials immediately after a phase of adaptation and deadaptation to a force field exposure, they exhibited spontaneous increase in their adaptation level, despite the absence of error during the error-clamp trials. This could also be the case in the adaptation behaviour during the Return trials of P2, where subjects perform under error-clamp trials after adaptation and deadaptation to force field during the Training phase. In order to tease apart such effect from the effects of interaction, in groups 2 and 3, the Out trials of P2 consisted of either error-clamp trials (Fig. 5.1G2), or null trials (Fig. 5.1G3), in which no adaptation was occurred. In this case, the adaptive behaviour during the Return trials of P2 could not be affected by any adaptation process in the Out trials, and thus, the possible effects of spontaneous recovery could be detected.

In group 4, we exposed subjects to a reversed force field perturbation during the Out trials of P2, to see how reversing the sign of the force field altered the adaptive behaviour

in the Return movements (Fig. 5.1G4). This could show whether the interaction effects were specific to the force field direction experienced over the Training phase, or could be generalised to force fields in the opposite direction.

Finally in group 5, subjects performed identical P1 and P2 phases as in group 1, except that the Training phase was different between the two groups. In group 1, subjects experienced simultaneous adaptation in both Out and Return movements, while in group 5, subjects only adapted to the force field in the Return trials (no concurrent adaptation; Fig. 5.1G5). We compared the behaviour observed during the P2 phase between the two groups to highlight the effects of concurrent adaptation on the interaction of Out and Return motor memories.

In all groups, subjects received rest breaks (90 seconds) at specific trials during the experiment as shown in Fig. 5.1.

5.2.3 Analysis

Two different performance measures were considered in the analysis of behaviour. On each trial we recorded the time course of position and velocity of the hand during movement (collection rate 1000 Hz). In null and exposure trials, we measured the maximum lateral deviation of the hand position from the straight line between the two targets as the kinematic error of performance. We refer to this measure as the maximum perpendicular error (MPE).

On error-clamp trials, we also collected the lateral force profiles generated by subjects against the channel walls (collection rate 1000 Hz) and compared it to the ideal force patterns required for complete force field compensation on that trial. To this end, the ideal force on each channel trial was calculated based on equation 5.1 and the velocity of movement on that trial. We then trimmed the resultant ideal and measured force trajectories based on the velocity profile as follows: we took the time points from when the velocity first exceeded 5% of its maximum value, forward to the time when the velocity was reduced again to less than 5% of its maximum value. Finally, we regressed the resultant measured force trajectory onto the ideal force trajectory (without intercepts), and considered the regression coefficient as the adaptation index. The adaptation index represented the progress in learning the force field from zero (no learning) to one (complete learning).

The statistical tests across groups were performed using one-way ANOVA. For pair-wise comparisons between groups (that is, comparing group 1 with groups 2–4), the Tukey–Kramer method Tukey (1949) was applied on the ANOVA test to correct for family-wise error rate due to multiple comparisons. This was implemented using Matlab functions *anova1* and *multcompare*. To test for the significance of quantities withing each group the two-tailed student's t-test was applied.

During the experiment, subjects were given multiple rest breaks (Fig. 5.1). We excluded from our analysis the data in the first 10 trials following each rest break.

5.3 Results

5.3.1 Model free analysis

We investigated the interaction of learning processes in a reciprocating adaptation task. Subjects performed successive Out and Return reaching movements between two targets, while we manipulated the dynamics of the movement in each cycle. Five groups of subjects participated in the experiment and each group underwent a different scenario regarding the dynamics of the task (Fig. 5.1). The experiment involved three main phases: the first probe phase (P1), the Training phase, and the second probe phase (P2). After each phase, subjects performed a short period of null trials on both directions to bring the adaptation level to the baseline before proceeding to the next phase. The overall performance measures, including adaptation and kinematic error, are shown as averages across subjects for each group in Fig. 5.2.

A priori interaction between Out and Return motor memories

We first examined P1, where subjects from all groups underwent a force field exposure during Out movements, while the Return movements were error-clamp trials. We asked whether adapting to the force field during Out movements for the first time affected the adaptive behaviour during Return movements. As shown in Fig. 5.2 (left column), throughout the first

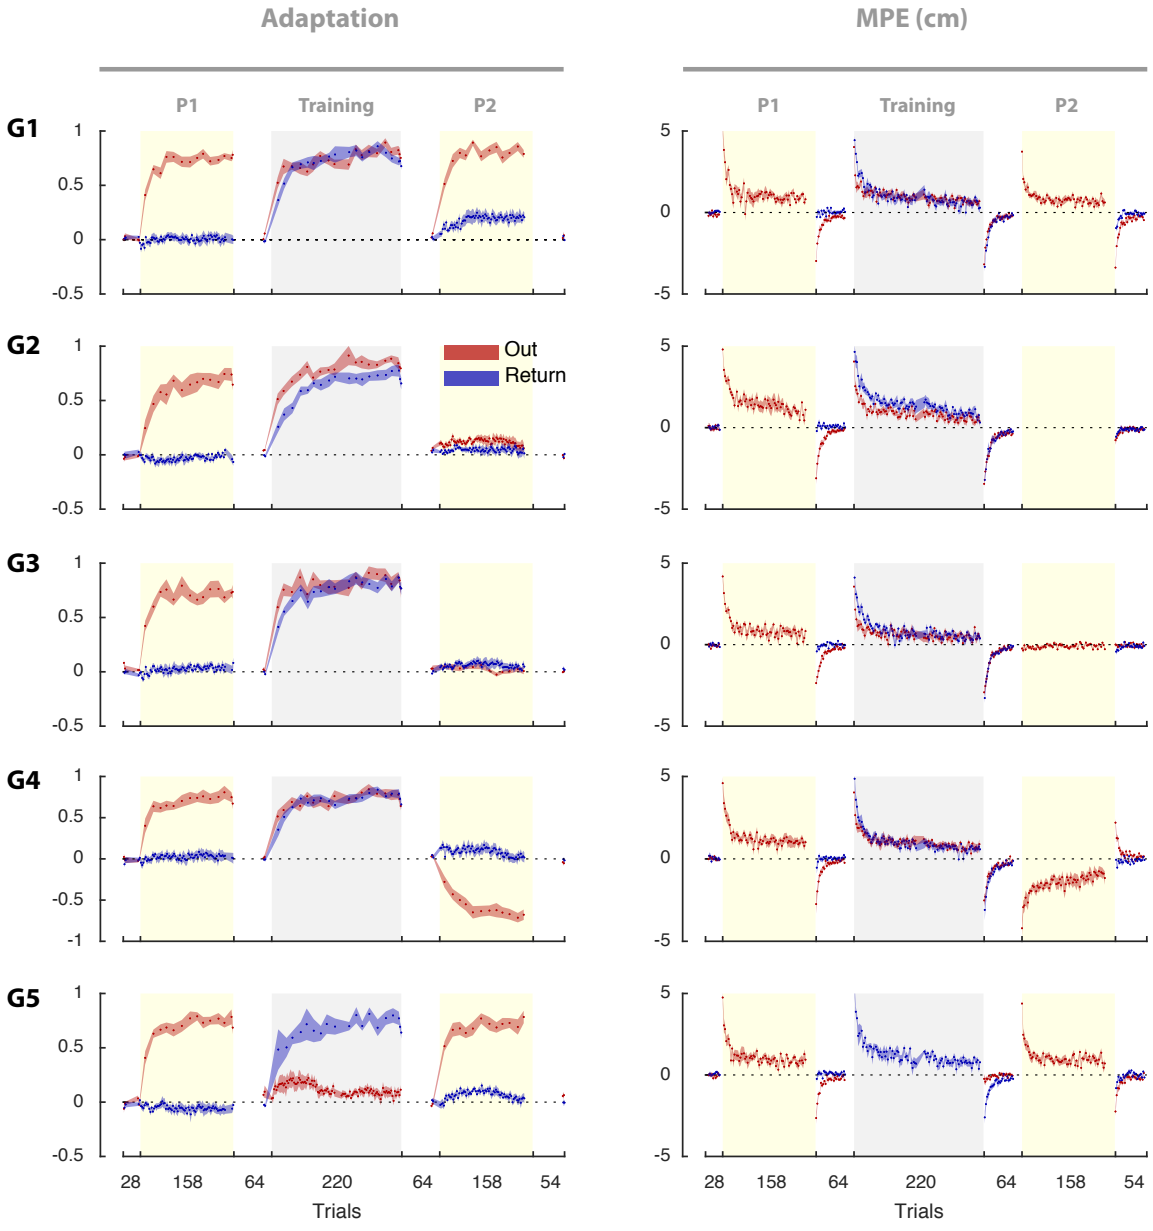

Fig. 5.2 The performance measures for adaptation (left column) and kinematic error (right column) for all groups (G1 to G5). Measures are the average across subjects (with error bars as $\pm SE$) separated for Out (red) and Return (blue) movements. The negative Out adaptation during P2 phase in G4 means adaptation to a reversed (opposite) force field.

probe phase and across all groups, adaptation in Out direction (red) increased progressively, whereas the adaptation level in Return movements (blue) seemed to be unaffected over trials. We particularly examined the final trials of P1, when the adaptation level in the Out movements nearly plateaued. Fig. 5.3A illustrates the mean adaptation level in the last 12 trials in the Return direction, as well as the last 2 error-clamp trials in the Out direction for each group. As shown, while the adaptation in the Out movement had reached $\sim 75\%$ by the end of P1 for all groups (significant effect of learning, t-test: $t(9) > 11.14$, $p < 0.0001$ for each group; no main effect of the group: one way ANOVA: $F_{4,45} = 0.278$, $p = 0.89$), the adaptation in Return movements remained mainly unaffected and very close to the baseline (t-test: $|t(9)| < 1.66$, $p > 0.13$ for each group and $t(49) = 0.073$, $p = 0.94$ overall). This indicated minimal transfer of learning from Out movement direction to the Return, implying that the sensorimotor memories associated with Out and Return movements were initially separate with minimal interaction during learning.

Adaptive coupling

We further examined how the dynamic exposure during the Training phase affected the interaction between Out and Return trials in different groups. Before doing so, to allow comparisons across conditions it was important that similar learning was seen during the Return trials of the Training phase, and that prior to P2 all groups showed similar washout of adaptation. To this end, for each group we first looked at the adaptation level across the last five error-clamp trials of the Training phase for Return movements. As illustrated in Fig. 5.3B (top panel), the adaptation was almost at the same level for all groups with no main effect of the group (one way ANOVA: $F_{4,45} = 0.976$, $p = 0.42$). Also, we tested whether the adaptation level after the washout of the Training phase (prior to P2) was at the same level across different groups and ideally close to the baseline. To do so, we analysed the adaptation level in the last error-clamp trial of the washout period preceding the P2. As shown in Fig. 5.3B (bottom panel), the adaptation level was near baseline for all groups with no significant difference found across the groups before entering the P2 phase (ANOVA: $F_{4,45} = 0.677$, $p = 0.61$).

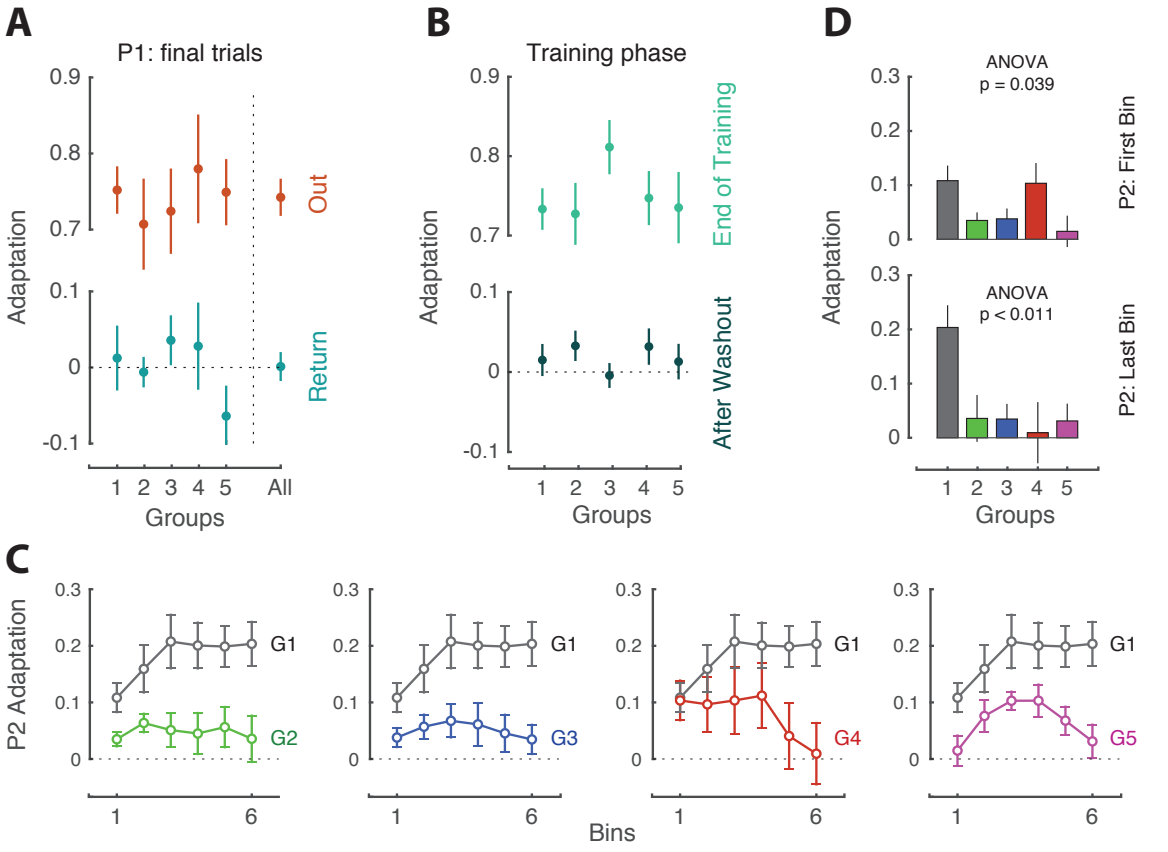

Fig. 5.3 **A.** Adaptation level at the end of Probe 1 (P1) for Out direction (last two channel trials; top panel) and Return direction (last 12 channel trials; bottom panel) presented for each group (1 to 5) and for the overall average across all groups (All). The error bars indicate the standard error of the mean (SE) across subjects. **B.** The adaptation level of Return movements at the end of the Training phase (last 5 channel trials; top panel) and after washing out (last channel trial of the washout) prior to the start of P2 (bottom panel). **C.** Adaptation behaviour during the Return trials of P2 presented for each group. The data is binned across the trials (12 trials in each bin) and each data point represents the mean adaptation, first over the trials in each bin, and then across subjects in each group (error bars show $\pm SE$ across subjects). **D.** The initial (first bin) and final (last bin) adaptation level of Probe 2 for the Return direction, compared across groups.

In the following, we examine the adaptive behaviour observed in each group, separately. In the first group we asked how simultaneous adaptation during the Training phase affected the interaction between the Out and Return motor memories. During P2, similar to P1, the force field was applied only in Out movements, while the Return movements were error-clamps. We observed that unlike P1 (in which no transfer of learning was observed), subjects showed an increase in the adaptation level during the Return trials of P2 (Fig. 5.2G1). This indicated that while subjects adapted to the force field during the Out trials, the adaptation transferred to the Return trials.

For better analysis of this behaviour and for statistical testing across groups, we binned the trials of Return movements in P2 into 6 bins of 12 trials each, and calculated the mean adaptation across the trials of each bin. Fig. 5.3C illustrates the binned adaptation averaged across all subjects in each group. Also, Fig. 5.3D directly compares the adaptation for the initial (top panel) and final (bottom panel) bin across all groups. Adaptation in the initial bin represents the rapid recall of the memory, while adaptation in the final bin represents the persistence of adaptation throughout the block. One-way ANOVA showed significant effect of group for both initial ($F_{4,45} = 2.76$, $p = 0.039$) and final bins ($F_{4,45} = 3.69$, $p = 0.011$).

As shown for group 1 (G1, black), the adaptation level of Return movements increased up to ~ 0.20 by the end of P2, even though all these movements were performed under error-clamp trials. Note that prior to P2, adaptation in Return movements was washed out to baseline (adaptation level after washout and prior to P2 for Return movements: $0.015\% \pm 0.019$; t-test: $t(9) = 0.862$, $p = 0.41$). We hypothesised that this adaptive behaviour was due to a coupling between the Out and Return motor memories that was developed during the Training phase.

To test whether the adaptive behaviour observed during the Return trials of P2 was in fact a coupling effect (induced by force field exposure in the Out trials), or merely an effect of spontaneous recovery (Kojima et al., 2004; Smith et al., 2006), we examined groups 2 and 3 who did not experience exposure trials during P2. Instead, they either experienced error-clamp trials (group 2) or null trials (group 3) during the Out movements (the Return movements were still error-clamps for all groups). In Fig. 5.3C, the adaptation pattern

of group 2 and 3 (G2 blue and G3 red) are compared with group 1. As shown, due to spontaneous recovery in group 2 and 3, adaptation showed a slight rebound that decayed quickly to the baseline by the end of P2. The magnitude of this rebound, however, was far from the adaptation level observed in group 1. Comparing the last bin of P2 between the groups, we observed significant difference between group 1 and the groups 2 and 3 in the final level of adaptation (ANOVA with Tukey-Kramer correction, G1 versus G2: p=0.048, and G1 versus G3: p=0.037; Fig. 5.3D, bottom panel).

Overall, the results showed that the induced adaptation pattern observed in group 1 was not merely a spontaneous rebound of learning in the Return trials, but was influenced by the ongoing adaptation process in the Out trials.

Specificity in coupling of learning memories

In group 1, subjects experienced the same force field in the Out trials of P2, as they did during the Out trials of the Training phase. This led to an induced adaptive response in the Return trials of P2, which was also in the same direction as the adaptation in the Return trials of the Training phase. We further examined whether the coupling effect observed in this group was specific to the force field direction between the training phase and P2. Subjects in group 4 performed a similar paradigm as group 1 with identical P1 and Training blocks. During P2, however, the direction of the force field in the Out trials was reversed to the opposite direction compared to the force field during the Training phase (Fig. 5.1G4). We examined whether reversing the direction of force field in the Out trials would reverse the induced adaptive behaviour in the Return trials.

Fig. 5.2G4 illustrates the overall performance of subjects in group 4. As shown, during the Training phase the adaptation is in the positive direction in both Out and Return movements. In P2, the adaptation direction in Out movement is negative due to the reversed force field direction in this phase. However, the adaptive behaviour in the Return movements did not reverse and remained positive (Fig. 5.3C), although, by the end of P2 the adaptation decayed to baseline. This shows that the coupling that is formed during the Training phase under a particular force field direction is not generalizable to an opposite force field direction.

Also, comparing the strength of the adaptive behaviour in group 4 and 1 in Fig. 5.3C, it is shown that the expression of adaptation in group 4 is not as strong as in group 1. In fact, the adaptation level in the last trial bin of P2 shows a significantly weaker adaptation in group 4 compared to group 1 (Tukey-Kramer corrected ANOVA on the final bin: p=0.019; Fig. 5.3D, bottom panel). In brief, it appears that the expression of adaptive behaviour during Return movements of P2 is specific to the force field direction under which the Training phase was performed.

Coupling requires simultaneous exposure in both movement directions

We examined the importance of simultaneous adaptation during the Training phase in the formation of coupling between Out and Return learning memories. Subjects in group 5 performed a similar experiment to group 1 with identical P1 and P2 phases. However, in the Training phase, instead of simultaneous adaptation in both Out and Return directions, they only experienced force field adaptation in the Return movements (the Out movements were all error-clamps; Fig. 5.1G5). This way, subjects formed an adaptive memory for the Return movements separately, without experiencing the force field exposure in both movement directions at the same time. We asked whether separate formation of the learning memory for Return movements during the Training phase could still lead to induced adaptive behaviour in the Return trials of P2.

In Fig. 5.2G5 the overall adaptation performance in group 5 is shown. Subjects separately adapted to the force field in Return direction during the Training phase (final adaptation level: 73%; Fig. 5.3D, top panel), while the Out direction was clamped. In this phase, we also observed a rebound of adaptation in the Out movements of the Training phase which progressively decayed to the baseline over the trials. During P2, subject were re-exposed to the force field in the Out movement, in response to which we examined the adaptive behaviour in the Return movement. Fig. 5.3C (rightmost panel) illustrates the adaptation behaviour of Return movements in group 5 compared to group 1. As shown for group 5 (G5), a strong rebound is observed, starting from the baseline in the beginning of the P2, reaching to almost 10%, and decaying back to the baseline by the end of the phase. This

behaviour was qualitatively different from the induced adaptation observed in group 1, both in the vigour of response throughout the phase, and the persistence by the end of the phase. Fig. 5.3D, shows a significant difference between the adaptation level in group 1 and 5 for the final bin of the P2 phase (Tukey-Kramer corrected ANOVA: p=0.0282; bottom panel). These results indicate that the induced adaptation observed in group 1 requires a coupling formed between the Out and Return learning memories via simultaneous adaptation in both movement directions.

Instant memory recall

Comparing different groups in our experiments and focusing on the adaptive behaviour in the Return movements of P2, we found that coupling of learning memories observed in group 1 caused a rather strong induced adaptive behaviour up to 20% adaptation and lasted until the end of P2. This feature was not observed in other groups as in all other cases, the adaptation was close to baseline towards the end of P2 (Fig. 5.3D, bottom panel). The persistence of this induced adaptation relied on simultaneous adaptation process during the Training phase as well as re-exposure to the same force field direction in the Out movement during the P2 phase. Here, we focus on another aspect of the adaptive behaviour observed during the Return movements of P2, that is, the adaptation level in the beginning of the phase (first bin of P2). This measure could represent the instant recall of the learning memory of the Return movements (formed previously in the Training phase) due to a sudden force field exposure in the Out movement. Fig. 5.3D (top panel) shows the adaptation level observed in the first bin of trials during P2 for all groups. As illustrated, there is a significant effect of group (One way ANOVA: $F_{4,45} = 2.76$, $p = 0.039$) with larger adaptation level for group 1 and 4 and smaller for the rest of the groups. Here, we discuss this observation in more details.

The difference between the experimental paradigm in group 1 and that of groups 2 and 3 was that there was no re-exposure to force field introduced during the Out movements of P2 for groups 2 and 3. This led to smaller adaptive response in groups 2 and 3 compared to group 1 in the beginning of P2, which implied no instant recall of learning memory due to lack of dynamic exposure in the Out movements (Fig. 5.3D). On the other hand, comparing

group 5 and group 1, the re-exposure to the force field was present for both groups during the Out movements of P2 (dynamic cue), but there was no simultaneous adaptation involved during the Training phase for group 5. This led to weaker adaptive response in the beginning of P2 for group 5, indicating that simultaneous training is also necessary for instant recall of the motor memory in the Return movements.

In group 4, subjects experienced simultaneous adaptation in the Training phase, just like in group 1, and also were introduced to re-exposure during the Out trials of P2. However, unlike group 1, the re-exposure in group 4 was in the reverse direction compared to the Training phase. This difference, although led to different adaptive behaviour between the two groups in the late trials of P2 (Fig. 5.3D; bottom panel), it showed almost the same adaptation level in the beginning of P2 for both groups (5.3D; top panel). These results indicated that simultaneous adaptation during Training phase as well as a dynamic cue (regardless of its direction) is sufficient to encourage the motor system to express instant recall of the learning memory during the Return movements.

5.3.2 Modelling

Before describing the model in detail, we will give a brief overview to provide some intuition about the idea behind the model. The model posits that the Out and Return movements are each represented by a separate adaptive process, which is mainly updated based on the errors experienced in its corresponding movement (that is, the Out process learns from errors in the Out movements, and the Return process learns from errors in the Return movements). We define the interaction between these processes as the extent to which each process learns from the errors that occur in the alternative movement (how much the Out process learns from an error in the Return movement, and vise verse). This interaction is modulated by a "coupling factor" which ranges from zero, meaning Out and Return processes are independent (they only learn from their own corresponding movement errors), to one, meaning each process learns from errors in both movements equally. The coupling factor is, itself, a function of error. Mainly, the coupling factor increases when errors are experienced in both movements simultaneously (error in one movement is always followed by an error in the alternative

movement). In this case, the coupling factor learns an association between the errors, such that any error occurred in the current movement might be followed by an error in the following, alternative movement. Therefore, it uses a portion of the error in the current movement to modify the behaviour in the following movement in a predictive way (e.g., uses a portion of the error from an Out movement, to modify the Return process). In the following, we describe the implementation of this idea in more detail.

Assume $x_o^{(n)}$ and $x_r^{(n)}$ represent the adaptive states on trial n associated with the Out and Return movements, respectively. We represent these states within a state vector ensemble $\mathbf{x}^{(n)}$ as described below:

$$\mathbf{x}^{(n)} = \begin{bmatrix} x_o \\ x_r \end{bmatrix}^{(n)} \tag{5.2}$$

In the motor system, there might be multiple (at least two) of such state vectors that operate in parallel (Lee and Schweighofer, 2009; Smith et al., 2006). Accordingly, akin to fast and slow adaptive mechanisms, one can imagine two parallel state vectors $\mathbf{x}_1^{(n)}$ and $\mathbf{x}_2^{(n)}$ which differ in learning and forgetting rate and their sum contributes to the motor output. Note that within each state vector, the Out and Return processes have the same rate of learning and forgetting. The net output in this case is defined as:

$$y^{(n)} = \mathbf{q}^T \cdot (\mathbf{x}_1^{(n)} + \mathbf{x}_2^{(n)}) \tag{5.3}$$

where, $y^{(n)}$ represents the motor output, and $\mathbf{q}$ is a binary selection vector that depending on the movement direction (Out or Return) determines whether the Out or Return process should be responsible for the current trial, n ($\mathbf{q}$ equals to $[1,0]^T$ for Out trials, and $[0,1]^T$ for Return trials). The difference between the resultant motor output and the external perturbation leads to an error as obtained below:

$$e^{(n)} = f^{(n)} - y^{(n)} \tag{5.4}$$

where, $f^{(n)}$ represents the external perturbation (force field) and is either equal to $+1$ (for CW force field), -1 (for CCW force field), or 0 (for null field). Note that for error-clamp trials error is set to $e^{(n)} = 0$ by definition. When the error is observed, the Out and Return processes within each state vector update their state according to their relevance to the trial in which the error occurred:

$$\begin{aligned} \mathbf{x}_k^{(n+1)} &= \alpha_k \cdot \mathbf{x}_k^{(n)} + \beta_k \cdot e^{(n)} \cdot \mathbf{G}_k^{(n)} \\ k &= 1,2 \end{aligned} \tag{5.5}$$

where, k represents the number of multiple parallel mechanisms (here, two). Within each mechanism, α_k and β_k represent the retention factor and the learning rate, respectively. Also $\mathbf{G}_k^{(n)}$ represents a gating vector that determines how much of the observed error in the current trial should be used to update each of the Out and Return processes. As such, for an Out trial, the gating vector is defined as $[1, c_k^{(n)}]^T$, and for a Return trial, it is defined as $[c_k^{(n)}, 1]^T$. According to this definition, when an error occurs, for example, in an Out trial, $[1, c_k^{(n)}]^T$ implies that the Out process should be fully updated (with the weight of 1), whereas the Return process should be partially updated via the coupling factor $c_k^{(n)}$. This is schematically shown in Fig. 5.4A, where an error from an Out trial (e_o) is appropriately apportioned between the Out and Return processes ($x_o^{(n)}$ and $x_r^{(n)}$) within each of the parallel mechanisms $\mathbf{x}_k^{(n)}$. In the following section, we describe the coupling factor in more detail and explain how it develops trial by trial as a function of error.

The structure of the coupling factor

The coupling factor ($c_k^{(n)}$) is the essential part of the model that determines the interaction between the Out and Return processes. The value of zero for the coupling factor means no interaction between the Out and Return processes (each process is only updated by the errors from its own corresponding movements), while, a value of 1 means full interaction between Out and Return processes (each process is equally updated by errors from both Out and Return movements).

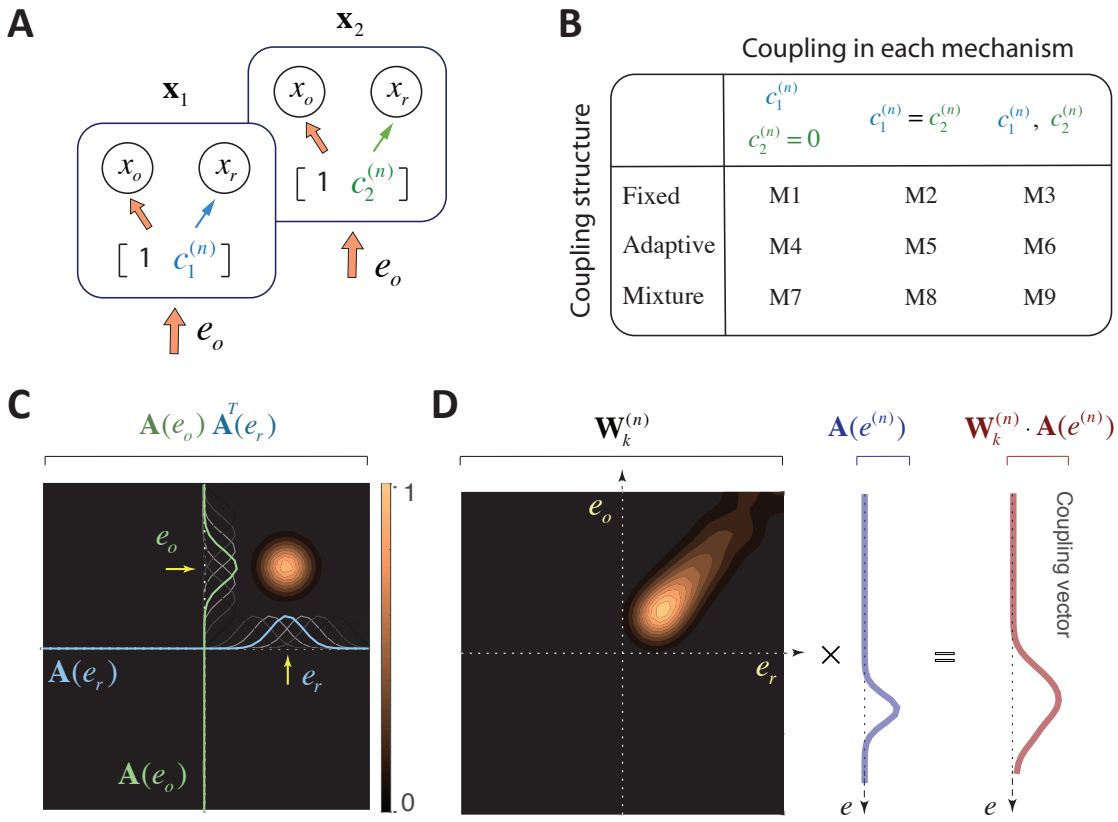

Fig. 5.4 **A.** The allocation of error to the Out and Return processes within each state vector. For a given error in a given context (for example context Out), the error is fully allocated to the corresponding adaptive process (here, x_o), but is only partially used for the alternative process (i.e., x_r) through a coupling factor ($c_k^{(n)}$, $k = 1, 2$). **B.** Different variations of the model can be considered based on: 1. whether the parallel mechanisms have the same or different coupling factors (columns), and 2. the structure of the coupling factor (rows). This results in 9 different models that feature the coupling factor in various ways. **C.** The joint activity of the primitives in the 2-dimensional error space. Each dimension represents the error space in one movement direction (subscript o for Out, and subscript r for Return directions). For a given error in each movement direction, a pattern of activity is formed by the primitives along the corresponding axis ($\mathbf{A}(e_o)$ and $\mathbf{A}(e_r)$). The product of these patterns represent the joint activity of the primitives in each pair of Out and Return trials. **D.** The memory matrix is formed across trials by accumulating the joint activity of the primitives on each pair of consecutive trials. The lighter areas represent stronger connection between the corresponding errors in the Out and Return axes. On a given trial n, the product of the memory matrix and the array of primitives results in a coupling vector. Each element of the coupling vector represents the association that is formed between the error in the current context and all possible errors in the alternative context. The mean value of the coupling vector is a scalar that represents the adaptive component of the coupling factor for that trial.

Here, we consider three possible structures that the coupling factor might adopt. First, the coupling factor could be a fixed parameter across all trials, representing a fixed interaction between the Out and Return processes:

$$c_k^{(n)} = c_k^0 \tag{5.6}$$

where, c_k^0 is a free parameter representing a fixed value for the coupling factor. Second, the coupling factor could involve an adaptive process through which it changes across the trials and can be increased or reduced depending on the dynamics of the task in each movement context:

$$c_k^{(n)} = \tilde{C}_k^{(n)} \tag{5.7}$$

where, $\tilde{C}_k^{(n)}$ represents an adaptive dynamics for the coupling factor (see below). Finally, the coupling factor could be considered as a mixture of fixed and adaptive coupling, in which the fixed component describes the baseline interaction, and the adaptive component captures how the interaction evolves or decays around the baseline:

$$c_k^{(n)} = c_k^0 + \tilde{C}_k^{(n)} \tag{5.8}$$

On the other hand, each of the parallel mechanisms ($\mathbf{x}_1^{(n)}$ and $\mathbf{x}_2^{(n)}$; Fig. 5.4A), could have a separate coupling factor of its own ($c_1^{(n)}, c_2^{(n)}$), or they both could share the same coupling factor ($c_1^{(n)} = c_2^{(n)}$), or only one of the mechanisms involves a coupling factor (only $c_1^{(n)}$, and $c_2^{(n)} = 0$). The combination of different coupling structures and their allocation to each parallel mechanism results in 9 different variations of the model (M1 to M9) as described in Fig. 5.4B.

Here, we describe the adaptive structure of the coupling factor in detail. The adaptive coupling is assumed to be a function of errors that are observed on each pair of Out-Return trials. When errors from the Out movements are repeatedly observed along with errors from the Return movements, an association (coupling) develops between these errors, such

that the occurrence of error in one movement predicts the occurrence of error in the other. Therefore, when an error is experienced in the current trial, it not only improves the behaviour in its corresponding process, but the coupling factor allows the alternative process also to be partially updated to reduce the expected error in the following trial. The more frequently the errors in Out and Return trials happen together, the stronger the coupling factor becomes (meaning that it becomes more likely that an error in an Out movement be followed by an error in the following Return movement).

To implement this notion, we assume that the errors in Out and Return movements are each encoded by a set of primitives that form a two dimensional error space as illustrated in Fig. 5.4C. As shown in the figure, the horizontal axis represents the errors in the Return movements (e_r) and the vertical axis represents the errors in the Out movement (e_o). When a pair of consecutive Out and Return errors occur, a pattern of activity is formed for each error in its corresponding dimension (i.e., $\mathbf{A}(e_r)$ and $\mathbf{A}(e_o)$ represented by Gaussian like patterns). The activity of primitives is described as Gaussian functions with preferred error sizes as shown in equation 5.9:

$$\begin{aligned} \mathbf{A}(e^{(n)}) &= [A_1(e^{(n)}),\ldots,A_m(e^{(n)})]^T \\ A_i(e^{(n)}) &= \exp\frac{-\left(\bar{E}_i - e^{(n)}\right)^2}{2\sigma^2} \end{aligned} \tag{5.9}$$

where, m is the number of primitives in each movement direction, $\bar{E}_i$ is the preferred error size for primitive i, and σ is the width of each primitive. The joint activity of primitives in Out and Return movements is defined as $\mathbf{A}(e_o)\mathbf{A}^T(e_r)$ which is shown as the lightened area in the two-dimensional error space in Fig. 5.4C. Over the trials, this joint activity is accumulated and stored in a "memory matrix" $\mathbf{W}_k^{(n)}$, whose elements represent the connection that is formed between the errors from the Out movements and the errors from the Return movements. Fig. 5.4D ($\mathbf{W}_k^{(n)}$) illustrates how the memory matrix would look like after multiple trials. As shown, the lightened areas in this matrix indicate the strength of the coupling that is formed between the corresponding errors from Out and Return movements. This strength is determined by how frequently a particular pair of errors is repeated (and thus

accumulated) across trials. The mechanism by which the memory matrix is accumulated is shown in equation 5.10.

$$\begin{aligned} \mathbf{W}_k^{(n+1)} &= \lambda_k \cdot \mathbf{W}_k^{(n)} + \eta_k \cdot \text{sign}(e_o \cdot e_r) \cdot \mathbf{A}(e_o)\mathbf{A}^T(e_r) \\ \begin{bmatrix} e_o \\ e_r \end{bmatrix} &= \mathbf{D} \begin{bmatrix} e^{(n)} \\ e^{(n-1)} \end{bmatrix} \end{aligned} \tag{5.10}$$

where, λ_k is the decay ratio and η_k is the accumulation rate of the memory matrix across trials. The errors e_o and e_r denote the errors experienced in the last pair of Out and Return movements. Therefore, in the past consecutive trials (n) and $(n-1)$, the associated errors $e^{(n)}$ and $e^{(n-1)}$ are related to the e_o and e_r via a context matrix $\mathbf{D}$:

$$\mathbf{D} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, n \equiv \text{Out} \qquad ; \qquad \mathbf{D} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, n \equiv \text{Return} \tag{5.11}$$

The context matrix in fact determines whether the error in a given trial n is from an Out movement, or a Return movement. The $\text{sign}(e_o \cdot e_r)$ term in the update equation 5.10 also determines whether the pair of Out and Return errors were of the same or the opposite signs. This was to take into account the direction of the errors that were coupled, in addition to their magnitude.

Given the memory matrix in each trial, we could determine the value of the adaptive coupling in the following way. For a given error on trial n with the pattern of activity $\mathbf{A}(e^{(n)})$, the product $\mathbf{W}_k^{(n)T}\mathbf{A}(e^{(n)})$ (if n is an Out trial) or $\mathbf{W}_k^{(n)}\mathbf{A}(e^{(n)})$ (if n is a Return trial) is an array whose elements represent the connection that was formed between the observed error $e^{(n)}$ in the current movement (Out or Return), and all possible error values ($\bar{E}_i$, $i = 1, ..., m$) that could occur in the alternative movement (Fig. 5.4D). The average across the elements of this vector determines the overall value of the adaptive coupling for the given error on trial n:

$$\tilde{C}_k^{(n)} = \begin{cases} \left\langle \mathbf{W}_k^{(n)T} \mathbf{A}(e^{(n)}) \right\rangle & n \equiv \text{Out} \\ \\ \left\langle \mathbf{W}_k^{(n)} \mathbf{A}(e^{(n)}) \right\rangle & n \equiv \text{Return} \end{cases} \tag{5.12}$$

where, $\langle\rangle$ represents the averaging operator that returns the average across the elements of its argument. In summary, the pattern of activation caused by error in a given trial ($\mathbf{A}(e^{(n)})$) is acted upon by the memory matrix to determine how the observed error is associated with different possible errors in the following trial, and thus how much of the current error should be used to update the alternative process to reduce the possible errors in the following trial.

Model selection and fitting

The model in its full description with mixture coupling has the following free parameters: the retention factor and the learning rate for each of the parallel mechanisms (α_k, β_k), the fixed component of coupling (c_k^0), the rate of decay and accumulation in the memory matrix(λ_k, η_k), and the width of each primitive (σ). Except for σ, the rest of the parameters are defined separately for each mechanism (i.e., $k = 1, 2$). Therefore, the model in its full form includes 11 free parameters. Note that in our model, we did not constrain any of the mechanisms ($\mathbf{x}_1^{(n)}$ and $\mathbf{x}_2^{(n)}$) to be specifically fast or slow, but allowed that to be specified during the model fitting. We used different variations of the model, each with different set of parameters, by manipulating two model features as shown in Fig. (5.4B). First, the structure of the coupling factor could take either of the forms as fixed, adaptive or a mixture of both (equation 5.6 to 5.8). Second, we could vary the way the coupling factor was used in each mechanism (Fig. 5.4B). A summary of all the variations of the model is shown in Table 5.1. In addition, we also considered the case in which none of the mechanisms were coupled, that is $c_1^{(n)} = c_2^{(n)} = 0$. This meant that the learning memories of Out and Return movements were fully independent with no interaction whatsoever.

All these variations led to ten different models that we fitted to the data. Each model was fitted to both the MPE and the adaptation measures of all groups simultaneously. We

Table 5.1 A summary of different variations of the model

#	**Model Description**	**Parameters**
M1	Fixed coupling: single adaptive mechanism	α_1, β_1, α_2, β_2, c_1^0
M2	Fixed coupling: shared between two adaptive mechanisms	α_1, β_1, α_2, β_2, $c_1^0 = c_2^0$
M3	Fixed coupling: separate for each adaptive mechanism	α_1, β_1, α_2, β_2, c_1^0, c_2^0
M4	Adaptive coupling: single adaptive mechanism	α_1, β_1, α_2, β_2, λ_1, η_1 , σ
M5	Adaptive coupling: shared between two adaptive mechanisms	α_1, β_1, α_2, β_2, $\lambda_1 = \lambda_2$, $\eta_1 = \eta_2$, σ
M6	Adaptive coupling: separate for each adaptive mechanism	α_1, β_1, α_2, β_2, λ_1, λ_2, η_1, η_2 , σ
M7	Mixture coupling: single adaptive mechanism	α_1, β_1, α_2, β_2, c_1^0, λ_1, η_1 , σ
M8	Mixture coupling: shared between two adaptive mechanisms	α_1, β_1, α_2, β_2, $c_1^0 = c_2^0$, $\lambda_1 = \lambda_2$, $\eta_1 = \eta_2$, σ
M9	Mixture coupling: separate for each adaptive mechanism	α_1, β_1, α_2, β_2, c_1^0, c_2^0, λ_1, λ_2, η_1, η_2 , σ
M10	No coupling	α_1, β_1, α_2, β_2

chose to fit both MPE and adaptation measures so that the number of fitting data points for all groups were the same (the number of adaptation data points varied between the groups depending on the paradigm, but the total number of data points including adaptation and MPE was the same for all groups). Bayesian information criterion (*BIC*) was used to select the best model among all the variations:

$$BIC = N \cdot \ln(\frac{SSE}{N}) + p \cdot \ln(N) \tag{5.13}$$

where, N is the number of data points including both the MPE and adaptation data points, *SSE* is the sum of squared errors of prediction for MPE and adaptation together, and p represents the number of fitting parameters in the model. Since the measured kinematic error and adaptation were of different mathematical units, we used an additional set of mapping parameters to linearly map adaptation and MPE measures into the dimensionless model space (where adaptation and MPE values are normalised within the range of 0 and 1). This added four additional parameters to the model fitting process (one scaling and one offset parameter for each of the MPE and adaptation measures). After model fitting, the mapping parameters were inverted to bring the model predictions of adaptation and kinematic error back to the measured space.

In order to find the confidence intervals for the parameters of each model, we prepared 1,000 bootstrap samples to which we fitted the models. Each sample was prepared by randomly choosing 10 subjects (with replacement) from each group, making it 50 selected subjects across all five groups for each sample. In order for the model comparison, for each given sample we calculated the difference in *BIC* (ΔBIC) between each model and a reference model (the model that had the smallest *BIC* overall). This provided a single ΔBIC for each model and for each data sample. We then took the average of ΔBIC across all data samples for each model, and used that as the criterion for model comparison. Other quantities used to determine the goodness of fit, such as sum-squared error of fit (*SSE*) and the coefficient of determination (R^2), were also reported as average across all samples for any given model.

Model fitting results

Table (5.2) summarises the quality of fit for each model in terms of the *BIC* difference (compared to the preferred model, M8), the sum squared error (*SSE*) and the coefficient of determination R^2. As shown, the model with no coupling (M10) shows the poorest performance in explaining the data. The models with fixed coupling (M1 to M3) are also far from the best model with the *BIC* difference of more than 230. However, models that include adaptive coupling (either adaptive only or a mixture of adaptive and fixed; M4 to M9) show substantial improvement in fitting the data. Among these models M8, which includes a mixture coupling shared between the adaptive processes $\mathbf{x}_1^{(n)}$ and $\mathbf{x}_2^{(n)}$, shows the best performance in terms of *BIC*. Note that although the R^2 values are similar across all models, these are independent of the number of data points in the fitting process. Whereas the BIC values scale with the number of data points. Therefore, the small improvement in R2 for a model can translate into a large difference in likelihood and hence the BIC.

In Fig. 5.5 the model predictions of adaptation and MPE after fitting M8 to the data is illustrated for all groups (solid lines). Table 5.3 also presents the best fit parameter values for model M8 with the 95% confidence intervals for each parameter. In the following, we examine the properties of M8 in more details. In particular, we focus on the behaviour of the

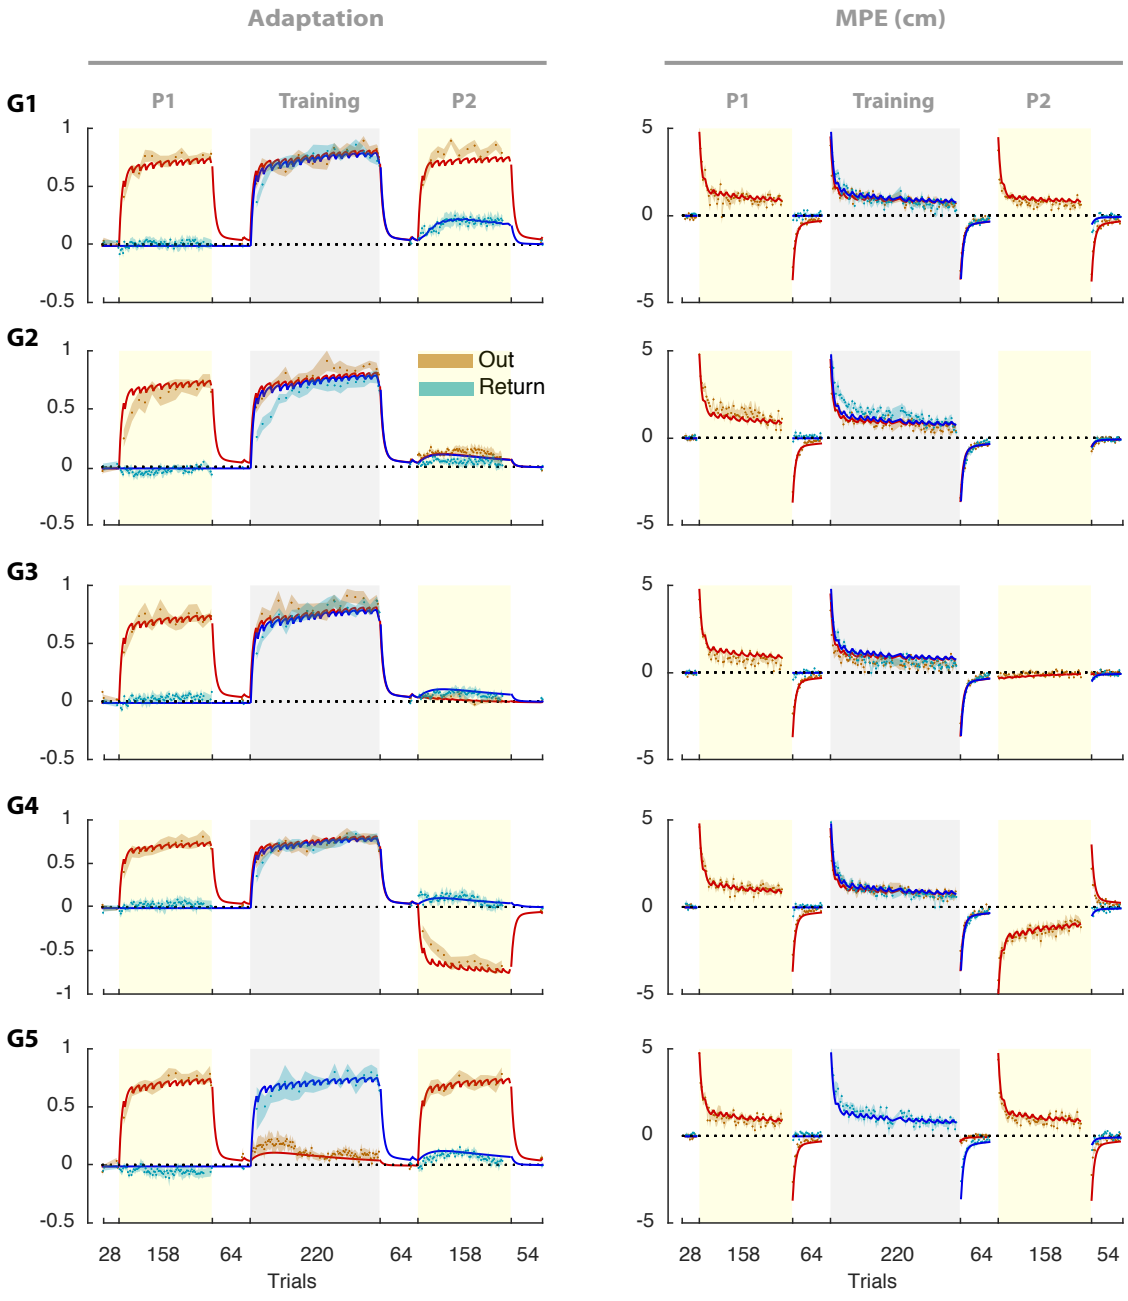

Fig. 5.5 Results from fitting the model M8 to the adaptation (left column) and MPE (right column). The model was fit to both measures and across all groups at once. The solid lines represent the model fits (red lines for Out movements and blue lines for Return movements) and the data is represented by circular markers and shading error bars.

Table 5.2 Model comparison: the goodness of fit for each model is shown in terms of *BIC*, the sum of squared error of fitting (*SSE*) and the coefficient of determination R^2. For each model, we took the difference between the *BIC* of that model and the *BIC* of the winning model, M8, as reported with ΔBIC.

Models	Description	ΔBIC	*SSE*	$R^2(\%)$
M1	Fixed coupling: single adaptive mechanism	287.5	13.074	94.11
M2	Fixed coupling: shared between the two adaptive mechanisms	246.7	12.922	94.18
M3	Fixed coupling: separate for each adaptive mechanism	232.0	12.838	94.22
M4	Adaptive coupling: single adaptive mechanism	19.2	12.047	94.57
M5	Adaptive coupling: shared between the two adaptive mechanisms	30.1	12.090	94.55
M6	Adaptive coupling: separate for each adaptive mechanism	29.9	12.027	94.58
M7	Mixture coupling: single adaptive mechanism	17.2	12.011	94.59
M8	Mixture coupling: shared between the two adaptive mechanisms	0	11.952	94.61
M9	Mixture coupling: separate for each adaptive mechanism	12.9	11.913	94.63
M10	No coupling	380.6	13.464	93.93

coupling factor and its components (i.e., fixed and adaptive coupling; equation 5.8) after fitting M8 to the data.

The coupling factor acts on the error that is experienced on each trial, and represents the amount of learning (from error) that is transferred from the current context (movement direction) to the other. The fixed component of the coupling factor (c_k^0) is shown in Table 5.3 and represents the baseline interaction that exists between the learning memories of Out and Return movement directions. As shown, the fixed coupling parameter has taken the value of 0.024, indicating a limited amount of interaction (2.4% transfer of learning) that initially exists between the Out and Return learning memories. This value mainly captures the adaptive behaviour during the P1 phase where we examined the initial transfer of learning from Out movements to Return movements. As shown in our model-free analysis of P1 (Fig. 5.3A; Return adaptation), the effect of adaptation during Out movements was limited on the adaptive behaviour during Return movements, which is reflected in the value of the fixed coupling parameter in our model-based analysis. In the following, we show that the adaptive component of the coupling factor plays the major role in explaining the data in our experiments.

We first compare the model performance between the models M2, M5 and M8 (Table 5.2; Fig. 5.4B, second column). In all these models, the coupling factor is shared between the adaptive processes ($\mathbf{x}_1^{(n)}$ and $\mathbf{x}_2^{(n)}$). However, the structure of the coupling factor in each

Table 5.3 Best fit parameter values for M8 with 95% confidence intervals.

Parameters	Description	Mean	Lower bound	Upper bound
α_1	Fast retention factor	0.947	0.941	0.953
β_1	Fast learning rate	0.237	0.220	0.260
α_2	Slow retention factor	0.994	0.993	0.995
β_2	Slow learning rate	0.037	0.028	0.045
$c_1^0 = c_2^0$	Fixed Coupling	0.024	0.005	0.043
$\lambda_1 = \lambda_2$	Coupling retention factor	0.995	0.990	0.997
$\eta_1 = \eta_2$	Coupling rate of change	0.087	0.060	0.112
σ	Standard deviation of primitives	0.059	0.043	0.078

model is different: M2 uses only the fixed coupling, M5 uses only the adaptive coupling, and M8 uses a mixture of fixed and adaptive coupling. The $\Delta BIC = 30.1$ between M5 and M8 shows that the inclusion of the fixed coupling in M8 has reasonably improved its performance over M5. However, this improvement is much less profound than the influence that the adaptive coupling has in differentiating the models. The $\Delta BIC = 246.7$ between M2 and M8 shows that the adaptive coupling in M8 has given this model a remarkable advantage over M2 in fitting the data, indicating the importance of an adaptive mechanism that connects the learning memories of Out and Return movements.

In Fig. 5.6 the time course of the coupling factor (mixture of fixed and adaptive) for M8 is illustrated for both Out and Return trials. As shown, the value of the coupling factor during P1 is very small and is only captured by the fixed component of coupling. In the Training phase, as shown for groups 1 to 4, due to the concurrent force field exposure the errors were experienced in both movement directions. This led to the formation of adaptive coupling, increasing up to $\sim 53\%$ by the end of the Training phase. Note that in the beginning of the Training phase, the coupling factor shows a very small growth due to the rapid reduction of error in the first few trials. However, in the following trials, the errors change very slowly, and thus the same error values are repeated between Out and Return movements for a longer period (Fig. 5.6, right column). This means that the same primitives are activated regularly,

causing a stronger connection between the corresponding error values from each movement direction.

During P2, we observed that the adaptive coupling revealed itself during the Out trials of group 1, but remained close to the baseline for other groups. The reason is that in group 1, subjects experienced exposure trials in the Out movements of P2, which resulted in similar errors that were previously experienced during the Training phase. Therefore, the coupling factor that was formed based on these errors reappears during P2. As shown in Fig. 5.6G1 (left column), the value of the coupling factor stands around 0.15 throughout P2, indicating roughly 15% transfer of learning from the Out movements to the Return movements. This, therefore, causes an increase in the adaptation level during the Return trials of P2 in group 1 (Fig. 5.3C, G1).

This is not the case for the other groups. In groups 2 and 3, the Out trials are either error-clamp (G2) or null trials (G3), neither resulting in errors that would induce coupling (Fig. 5.6G2 and G3, Right column). In group 4, subjects experienced a reverse force field during P2, which resulted in new errors that were in the opposite direction as the ones experienced during the Training phase (Fig. 5.6G4, Right column). Since these errors were not observed previously, no coupling was formed and observed around them. Finally in group 5, subjects never experienced concurrent exposure throughout the experiment, and thus no adaptive coupling was developed in this group.

We can also look at the behaviour of the memory matrix, particularly during the Training phase, to show how it evolves across trials. We specifically compare the memory matrix for group 1 (with concurrent adaptation during the Training phase) and group 5 (without concurrent adaptation). Fig. 5.7 illustrates the memory matrix for a few selected trials from the Training phase for each group. As shown for group 5 (G5), there were no connections formed between the errors in the Training phase, as subjects never experienced errors in both movement directions simultaneously. Therefore, no coupling was observed in this group. In group 1, however, due to concurrent adaptation during the Training phase, the connections start to form between the errors from Out and Return movements (Fig. 5.7G1). Since the memory matrix accumulates the joint activity of primitives on each pair of Out-Return trials,

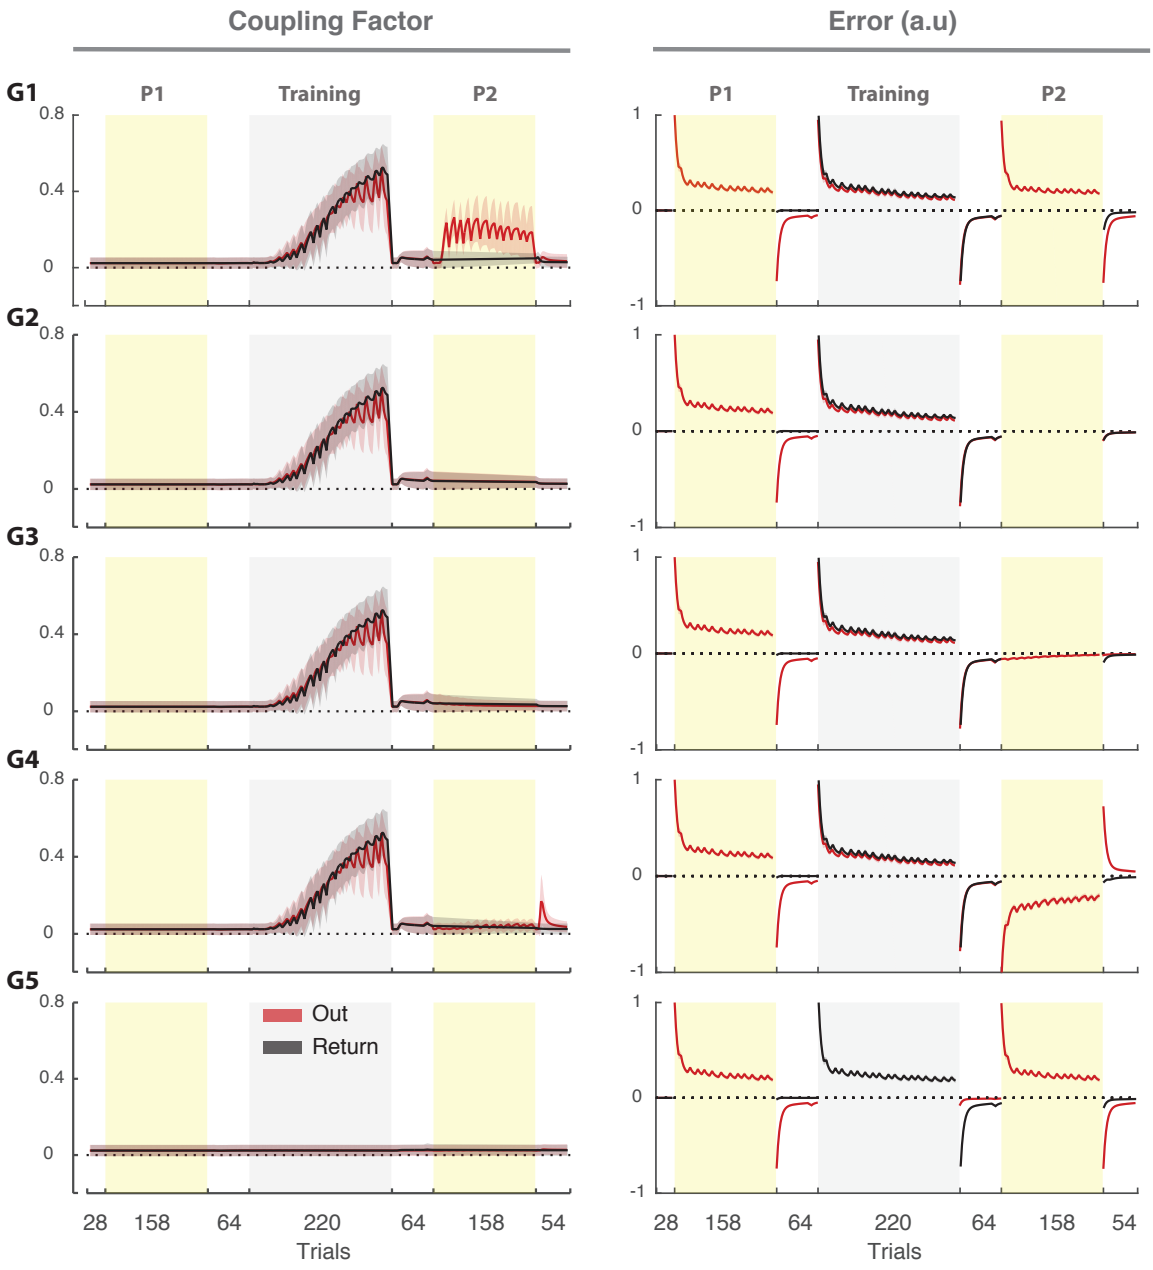

Fig. 5.6 Left column: the predicted time course of coupling factor (model M8) shown for all five groups (G1 to G5) and separated for Out (red) and Return (black) movements. The coupling factor for Out movements means the fraction of error in Out trials that is used to update the Return process. Likewise, the coupling factor of Return movements is the fraction of error in Return trials used to update the Out process. The error bars represent the 95% confidence intervals obtained by fitting the model M8 to the bootstrapped samples. Right column: the predicted error behaviour shown for the Out and Return trials.

the more a particular pair of errors is repeated, the stronger the connections will be in the memory matrix. Therefore, for the early trials of the Training phase, where the errors reduce rapidly, there is weaker connections formed between the large errors. Whereas in the later trials where the error plateaus in small values, the connections become stronger around the origin in the error space.

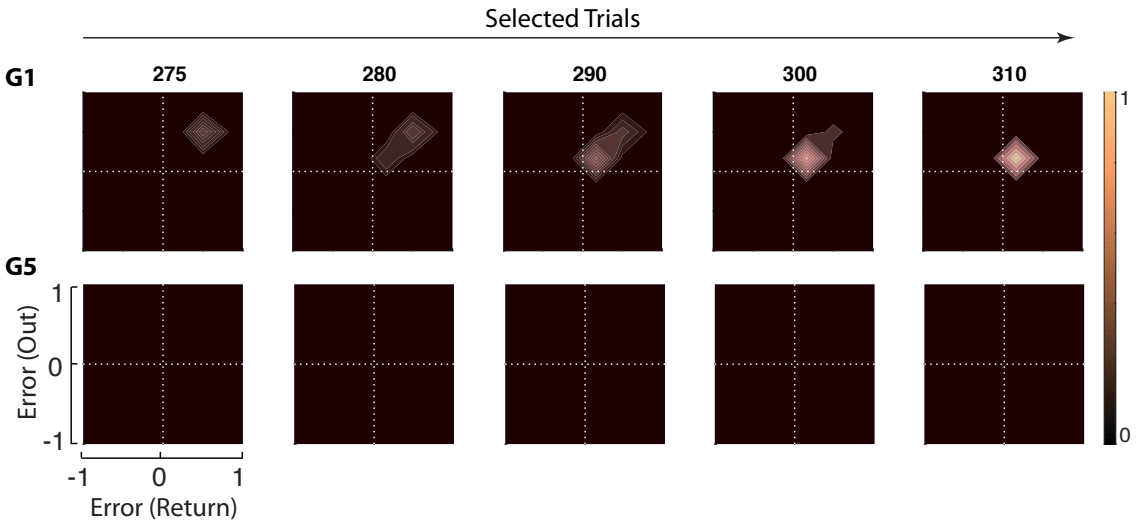

Fig. 5.7 The memory matrix shown for model M8 in some selected trials of the Training phase for group 1 (G1) and group 5 (G5). The lightened areas show the strength of the connections formed between the errors in Out and Return movements in the 2-dimensional error space. **G1**. The errors are large in the initial trials of the Training phase, resulting in connections that are formed away from the origin. As errors reduce towards zero through adaptation, the connection weights also start to form around the origin. The strength of the connections near the origin (i.e., for small errors) is higher than the connections formed away from the origin (i.e., for large errors) because small errors are more frequently experienced throughout the adaptation. **G2**. Since there was no concurrent adaptation in the Training phase of group 5, no connection was formed between the errors from Out and Return direction.

5.4 Discussion

We examined how motor memories of novel dynamics, formed and separated based on kinematic contexts (movement directions) interacted with one another. Examining reciprocal reaching movements under force fields, and manipulating the dynamics of the reach in each movement cycle (Out and Return), we showed that associated motor memories interacted,

and the interaction (coupling) could increase or decrease depending on the dynamic condition of the task. We observed that coupling evolved when both movement cycles underwent a force field perturbation simultaneously, and was specific to the error signals experienced during this simultaneous exposure. Our behavioural observations were accounted for by a computational model in which separate processes were responsible for the adaptive state of each movement cycle. These processes interacted via an adaptive coupling factor that distributed the error signal between the processes based on the current context.

Previous state-space models of sensorimotor learning have addressed the context-specific features of motor memories when learning a novel task (Ingram et al., 2013, 2011; Kim et al., 2015b; Lee and Schweighofer, 2009; Shadmehr and Moussavi, 2000; Thoroughman and Shadmehr, 2000; Voss et al., 2013). In these models, typically a context-dependent learning-rate or retention factor is considered to determine how acquisition or decay of adaptation is transferred across contexts (e.g., movement directions, object orientations, etc). For instance, it has been shown that both learning and decay of motor memories are largest for the current active context, and diminish in a Gaussian manner for nonactive distant contexts (Ingram et al., 2013). However, such patterns of interaction have been assumed to remain unchanged through the course of learning, indicating a fixed generalisation behaviour in the motor system. Our results challenged this assumption and showed that even for distant contexts with minimal transfer of learning (i.e., two opposite movement directions; Donchin et al., 2003; Howard and Franklin, 2015; Ingram et al., 2013), the interaction is dynamic and can increase or decrease depending on the dynamic condition of the environment.

In the current study, we used error-based primitives that separately encoded error signals for each movement direction (Out and Return). The notion of primitives used in this study is different from the state-dependent motor primitives used to directly contribute to the adaptation state (that is, the adaptive state is obtained as the weighted sum of state-dependent motor primitives; Donchin et al., 2003; Sing et al., 2009; Thoroughman and Shadmehr, 2000). It is, instead, consistent with previous theoretical (Herzfeld et al., 2014; Takiyama et al., 2015) and neurophysiological (Ferrera and Barborica, 2010; Popa et al., 2012) studies suggesting the existence of primitives (or neurons) that respond directly to the error signals. In this

case, our model might seem similar in formulation to two previous models of sensorimotor learning: the memory of error (Herzfeld et al., 2014) and the prospective error (Takiyama et al., 2015) models. In these models, error-based primitives were used mainly to account for how the motor memory of a task was formed under varying environmental conditions, by manipulating the sensitivity to experienced errors. In these cases, formation and retrieval of memory was for a single context (e.g., one reach direction), but varied with environmental consistency.

In our study, however, we asked a fundamentally different question: how motor do memories associated with different contexts interact? In this case, primitives were used to determine the connection between the motor memories of different movement directions (Out and Return movements). The nature of the interaction was determined in a principled way: when primitives in both Out and Return movements were activated concurrently through simultaneous exposure, the connection between the activated primitives strengthened. This was specified by the weighting coefficients in the memory matrix (equation 5.10) that were updated based on the joint activity of primitives ($\mathbf{A}(e_o)\mathbf{A}^T(e_r)$), as well as the the rate of increase (η_k) and decay of connections (λ_k; equation 5.10). Such a mechanism is also consistent with the notion of associative (Hebbian) learning (Hebb, 2005), wherein the motor system learns an association between the observed error signals (stimuli) in Out and Return movements. As such, experiencing error in both contexts repeatedly, causes repeated activation of corresponding primitives, which then leads to an increase in the connection weights between those primitives. As a result, when error is experienced in one context (i.e., Out), it not only updates the motor memory of the corresponding context, but also partially modifies the memory of the alternative context (i.e., Return).

Why is developing such interaction between separate context-dependent motor memories useful? It might seem rather detrimental for the motor system to change the memory of the distant, nonactive contexts based on the errors observed in the current context. Previous findings have shown that the brain protects motor memories (low learning and high retention) that correspond to contexts deviated from the task in hand (Ingram et al., 2013). This allows the brain to preserve multiple motor memories for different contexts and prevent

them from change due to irrelevant stimuli (error signals). We suggest, however, that under specific conditions, it could be an efficient strategy to keep the interaction patterns flexible between distinct motor memories. As such, when the environment alternates between two different contexts, each followed by an external perturbation, the motor system learns that an experienced error in one movement context is highly likely to be followed by error in the following context. Therefore, it not only uses the error to improve performance in the current context, but also modifies the motor memory of the alternative context to reduce potential errors in the future in a predictive manner. In this way, the motor system is prepared for the errors in the following context before they occur.

How strongly the motor system transfers learning from one context to another also depends on how certain it is about the association between the contexts. This seems to be determined by the duration of training under simultaneous exposure in both contexts. As such, longer training under concurrent exposure increases the probability that perturbation in the current context will be followed by perturbation in the following context, thus leading to larger transfer of learning. In our results, this could be seen in the behaviour of the coupling factor during the training phase, where the coupling increased gradually over trials of simultaneous exposure (Fig. 5.6, left column). Collectively, our results suggest a flexible interactive behaviour between context-dependent motor memories that allows for sharing information (error signals) among different contexts to efficiently improve overall performance.

Chapter 6

Adaptation to familiar dynamics

6.1 Introduction

In the previous chapter, we examined the interaction of sensorimotor memories associated with different kinematic contexts when learning a novel dynamics (i.e., force field perturbation). We showed that such interaction was of an adaptive nature and increased due to simultaneous exposure in both movement contexts (Out and Return reaching movements). In the current chapter, we examine the interaction of context-dependent motor memories when manipulating objects with familiar dynamics.

Previous studies have highlighted the difference between the underlying mechanisms of adaptation in novel versus familiar dynamic conditions (Ingram et al., 2011). When a novel dynamics is learned (e.g., a state-dependent force field), the motor system needs to learn the structure of the dynamics from scratch (Braun et al., 2009, 2010), a process with a relatively long time constant (e.g., about 50 trials; see Fig. 5.2). Since the structure is initially unknown, subjects might also need to learn whether there is an association between the dynamics of the task in different contexts. For instance, as shown in our previous chapter, while there was no initial interaction between the memories of Out and Return reach contexts, over the course of the Training phase subjects learned the dynamic association (i.e., simultaneous exposure) and developed a pattern of between-context interaction.

By contrast, when dealing with familiar dynamics in the environment, such as manipulating a tennis racket or lifting a cup, the structure of the dynamics is already formed in the motor system due to everyday use of similar objects. In this case, for successful manipulation, we only need to update our knowledge of the parameters of the dynamics, such as the mass or the moment of inertia. This form of adaptation, known as parameterisation, occurs in a shorter time scale (in a few trials; Fu and Santello, 2012; Ingram et al., 2010).

In this chapter, we examine how such structural knowledge of the dynamics in object manipulation affects the interaction of motor memories across different contexts. Subjects performed back and forth (reciprocal) rotations of a hammer-like object (clockwise and counterclockwise) between two angular targets. Each rotational direction was considered as a separate kinematic context. In different conditions, we varied the dynamics of the object in one direction, and tracked how subjects modified their performance both in the current direction and in the alternative reversed direction. We examined whether the motor system held a single motor memory for the dynamics of the object in both kinematic contexts (rotational directions), or whether motor memories were separated for each context. In the latter case, we asked how the motor memories of different contexts interacted, and how such interaction differed from that of learning a novel dynamics. To answer these questions, we designed various object manipulation paradigms and used our context-dependent model (previous chapter) to analyse the behaviour. We show that in the case of familiar dynamics, the motor memories associated with different kinematic contexts are still separate, but interact via a fixed coupling factor.

6.2 Methods

Thirty-six healthy right handed subjects (20 female; age between 20 and 30 years old) participated in this study. All subjects were naive to the purpose of the experiment, and provided written consent before participation. During the experiments, subjects were seated at a virtual reality system and rotated a virtual hammer-like object in the horizontal plane by grasping and rotating a vertical handle of a robotic manipulandum (the WristBOT; Howard

et al. (2009); Fig. 6.1A). The WristBOT allowed us to simulate the dynamic behaviour of the object during rotation, including the translational forces in the horizontal plane as well as the torque about its vertical handle. The visual feedback associated with the object and the task was provided in real time using a virtual reality display system. The visual object consisted of a circular handle (radius 0.5 cm) attached to a 4 cm square mass by an 8 cm rod (width 0.2 cm; Fig. 6.1B). The position and orientation of the object were determined by the position and orientation of the WristBOT handle.

6.2.1 Task

The task was to grasp the object by the handle at the base of the rod and rotate it back and forth between two angular targets (Fig. 6.1A). Subjects were instructed to generate pure rotational movement about the object's handle while keeping the position of the handle stationary at the centre of the home region (a visually presented 1 cm radius disc). The trials consisted of alternating clockwise (CW) and counterclockwise (CCW) rotations between the targets (two 40°-apart oriented bars emanating from the home region; Fig. 6.1B). Each trial began with the handle of the object stationary within the home region and the rod of the object aligned with either the CCW target or the CW target. Movement initiation was cued by a tone and the appearance of the second target, towards which subjects rotated the object. The trial ended when the orientation of the object was aligned with the orientation of the aimed target. Subjects were required to finish the movement within 400 ms, otherwise they were warned with a 'two slow' message. Movements exceeding 500 ms had to be repeated. Rest breaks (60 s) were provided during the experiments every 3–5 minutes.

The dynamics of the object was simulated as a point mass (1% of the subject's body mass) at the end of the rigid rod (Fig. 6.1B). Rotating the object generated forces and torques at the handle which were simulated by the manipulandum. The torque was associated with the moment of inertia of the object, and the generated force was associated with the circular motion of the mass, which could be decomposed into two components: tangential and centripetal (**F** in Fig. 6.1B). The direction of force depended on the orientation of the object, while its magnitude was determined by the mass and the length of the object. Critically, the

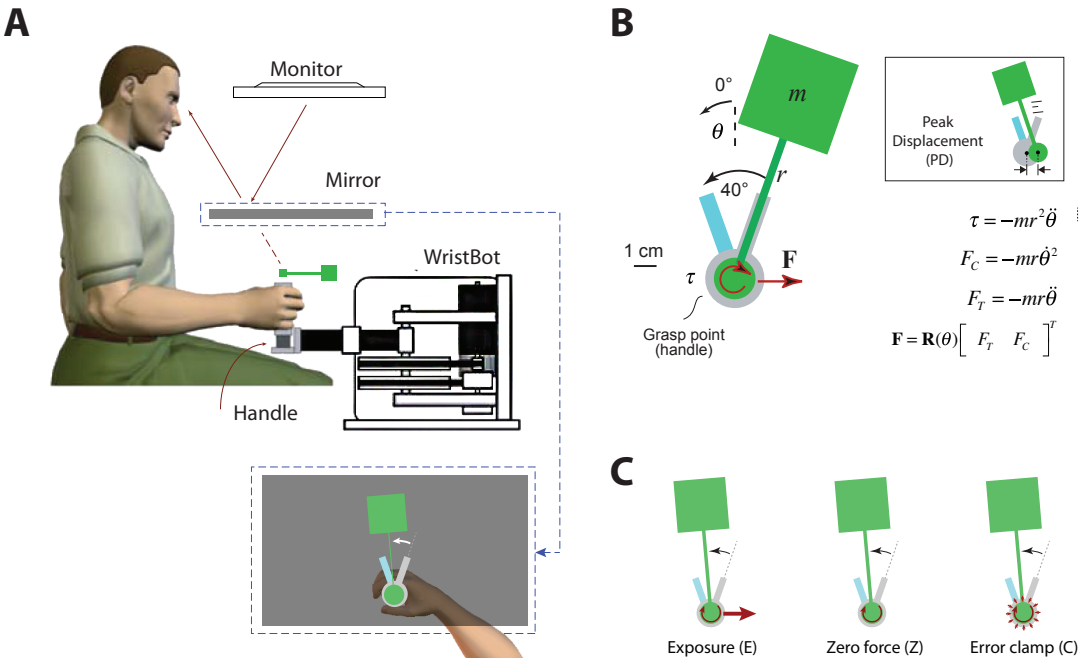

Fig. 6.1 **A.** Experimental setup. Subjects grasped the handle of the WristBot and rotated a virtual hammer-like object back and forth between two angular targets. **B.** The object consisted of a mass attached to a stick. Subjects were asked to rotate the object around its axis (grasp point) while trying to prevent the translational movement of the handle. The targets were two rectangle bars separated 40°. The dynamics of the object consisted of a torque (τ) that resisted the rotation, as well as translational forces (centripetal F_C and tangential F_T) that would perturb the handle of the object. Subjects had to learn to compensate for the torque to rotate the object withing the accepted speed, and the perturbing forces to keep the handle stationary during the rotations. **C.** Three types of trials were used during the experiment. In exposure trials (left) subjects experienced both the translational forces and the torque. In null trials (middle) the forces were switched off and only the torque was applied. And, in error-clamp trials the handle of the object was fixed in the home position a using simulated spring that prevented translational movement. Subjects also experienced the inertial torque during error-clamp trials.

force caused the handle of the object to displace unless subjects produced a compensatory force in the opposite direction.

During the experiments, we manipulated the dynamics of the object to be one of three possible trial types: exposure trial, null trial or error-clamp trial (Fig. 6.1C). On exposure trials, subjects experienced the full dynamics of the object including the force on the handle as well as the torque. In these trials, subjects had to learn to compensate for the translational forces in order to keep the handle stationary during the rotation. On null trials, the manipulandum did not produce any forces and the handle was free to move. Importantly, any forces produced by subjects on null trials would cause the handle to displace. Finally, on error-clamp trials, the manipulandum simulated a stiff two-dimensional spring, centred on the handle position at the start of the trial (the spring constant was 40 N/cm). Error-clamp trials effectively eliminated kinematic errors and prevented error-driven adaptation. They also allowed the compensatory forces produced by subjects to be measured (Ingram et al., 2011).

6.2.2 Analysis

Two performance measures were used for the analysis. On exposure and null trials we measured the position of the handle (at 1000 Hz) during the rotation, and took the peak displacement (PD) of the handle on that trial with respect to the starting position. The PD was considered as the kinematic error of movement (the task was to keep the handle stationary), and thus the PD of zero meant ideal performance. Also, on each error-clamp trial, we recorded the time course of the force magnitude generated by subjects on that trial (at 1000 Hz) in compensation for the object dynamics. We also obtained the ideal force trajectory that could fully compensate for the object dynamics based on the angular velocity and acceleration of movement on that trial (that is, Fig. 6.1B). The ideal and measured force trajectories were trimmed according to the angular velocity of the object, from when the velocity exceeded 5% of its maximum value, onwards until when the velocity reduced to less than 5% of its maximum value. We then regressed the measured force onto the ideal force (without intercept) and took the regression coefficient as the adaptation index.

Overall, for each error-clamp trial we had a single measure of adaptation, and for each null or exposure trial we had a single measure of PD. We excluded from our analysis the data from the first 10 trials that followed each rest break during the experiment.

6.2.3 Experiments

Experiment 1

Twelve subjects participated in the first experiment. We examined the representation of the object's dynamics during reciprocal object rotations. Specifically, we asked whether the dynamics of CW and CCW rotations were represented by separate adaptive mechanisms, and whether these mechanisms interacted. To this end, we conducted a similar paradigm as in the previous chapter (Group 1; Fig. 5.1G1). The experiment included pairs of CW and CCW rotations back and forth between two angular targets (Fig. 6.1A). The first trial of each pair was considered as the "Out" rotation, followed by the second trial as the "Return" rotation (these were discrete trials). We counterbalanced the choice of Out and Return rotations as half of the subjects performed the paired trials of CW-CCW as Out-Return rotations, while the other half performed the CCW-CW pair as Out-Return rotations. The data from these subgroups was properly combined for the analysis and presentation. Each subject had a familiarisation session in which s/he performed 20 pairs of null trials. These trials were not analysed.

The experiment started with a short block of 18 trials (9 pairs of Out-Return trials), that included four pairs of error-clamp trials as well as 5 pairs of null trials (Fig. 6.2A). This short block provided the baseline level for PD and adaptation ratio. The experiment continued with three main phases: the first probe phase (P1), the Training phase, and the second probe phase (P2). During P1 (35 pairs of trials), subjects experienced the full dynamics of the object in the Out rotations, while the Return rotations were all under error-clamp trials. At the end of P1 there was two more pairs of error-clamp trials (in both rotations) to measure the final adaptive state of Out and Return movements. Subjects then deadapted to the object dynamics in a following washout period with 22 pairs of null trials in both rotations.

The Training phase included 50 pairs of exposure trials, in which subjects adapted to the object dynamics concurrently in both Out and Return rotations. This phase was completed with two pairs of error-clamp trials in the end of the phase to measure the final adaptive state in each rotation. The Training phase was also followed by 22 pairs of null trials to deadapt subjects to the object dynamics. Finally, subjects performed the P2 phase which was identical to P1 (35 pairs of trials), with exposure trials in the Out and error-clamp trials in the Return rotation. The P2 phase was completed with two more pairs of error-clamp trials (in both rotations) in the end of the phase, followed by final washout period of 17 pairs of null trials.

Experiment 2

A new group of 12 subjects participated in the experiment. In this paradigm, we conducted a more extensive experiment in which we tested all possible combinations of task dynamics (trial types) between Out and Return rotations and examined how the interaction between the adaptive processes associated with each rotation changed throughout the experiment. Each subject had a familiarisation session in which they performed 20 pairs of null trials to get used to the task and the apparatus. During the main experiment, subjects performed 13 consecutive blocks of trials, each consisting of 13 to 18 pairs of Out-Return rotations (Fig. 6.2B). Within each block, the dynamics of the task between Out and Return rotations was of a different combination (Fig. 6.2C). In some blocks we applied the same dynamics for both rotations (i.e., the diagonal elements in Fig. 6.2C). While in the other blocks, we had different object dynamics associated with each rotation (blocks 5, 7, 9, and 12; the off-diagonal elements of Fig. 6.2C). Each block was followed by two pairs of error-clamp trials in which we measured the adaptation level at the end of that block for each rotational direction (as shown by grey patch in Fig. 6.2B).

In the first block, subjects performed 13 pairs of null trials as the baseline. After that, in blocks 2 to 10, we alternated the blocks of simultaneous exposure in both rotations (i.e., 2,4,6,8 and 10) with probe blocks in which we varied the dynamics between the two rotations (blocks 3,5,7,9). We examined all possible combinations of dynamics in probe blocks for when the Return movement was a clamp trial, while the Out movement was either a clamp

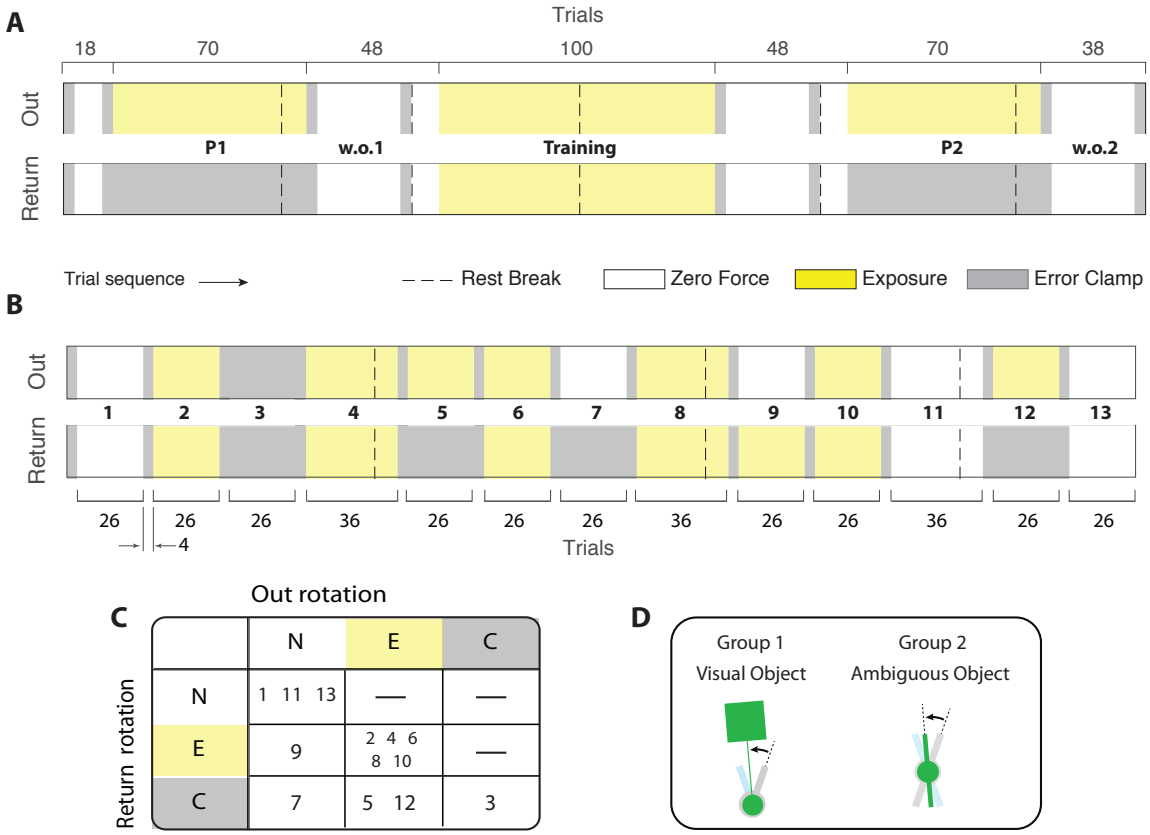

Fig. 6.2 Experimental paradigm. **A.** Experiment 1: the experiment consisted of three main phases: P1, the Training phase and P2, each followed by a period of washout (null trials). In P1 and P2, subjects experience exposure trials (yellow patch) in Out movements, and error-clamp trials (grey patch) in Return trials. These phases are followed by washout trials (null; white patch) w.o.1 and w.o.2. During the Training phase, subjects perform under exposure trials in both Out and Return rotations. **B.** Experiment 2: subjects perform 13 blocks of trials with different dynamic manipulations between Out and Return movements. Multiple rest breaks were provided for the subjects (60-90 seconds) during each experiment as shown with the dashed lines. **C.** different combinations of blocks between the Out and Return movements used in experiment 2. **D.** Two groups of subjects participated in this experiment. In group 1, the visual information about the orientation of object was provided, whereas in group 2, the object was ambiguous as to orientation. The dynamics of the task for both groups was the same.

(block 3), exposure (block 5), or null (block 7). We also examined cases where the Out movement was a null trial and the Return movement was either a clamp (block 7) or exposure (block 9) trial. In block 11, subjects had a washout period in both rotations to return their performance back to the baseline. We then examined learning after a washout in block 12 in which the Out movement was an exposure and Return movement was a clamp trial. Finally, subjects finished the experiment with a null block for both movement rotations in block 13.

All subjects underwent the same order of probe blocks so that the behaviour in each block was comparable across subjects (i.e., for each block, the preceding trials and the following trials were the same for all subjects; this is important as the behaviour in each block is highly influenced by the history of adaptation preceding that block). This way we could appropriately average the time course of adaptation and PD across all subjects, and perform the model fitting analysis on the averaged data over the whole trial sequence. Accordingly, we chose the order of probe blocks as follows. For the first probe block (block 3; Fig. 6.2B) we used error-clamp trials in both rotations so that we could measure the baseline rate of memory decay in each rotation, after subjects first adapted to simultaneous exposure in block 2 (note that all probe blocks were preceded by simultaneous exposure blocks). In blocks 5 and 7, we examined the decay rate of adaptation in the Return rotation (error clamp trials), while having exposure and null trials in the Out rotation, respectively. In this case, we chose to introduce the exposure-clamp trials in block 5, before the null-clamp trials in block 7. This was encouraged by some predictions of the model which we will discuss in detail in the discussion section. We then introduced block 9 with null trials in Out and exposure in Return. Finally, blocks 11, 12 and 13 were introduced last to resemble the P2 phase of experiment 1, where we examined the induced adaptation behaviour (i.e., exposure-clamp trials, preceded and followed by washout trials; Fig. 6.2A and B).

We further aimed to examine how the visual information about the object's shape and geometry would affect the representation of the object dynamics in Out-Return rotations. In this experiment, subjects were provided with the information about the shape of the object (i.e., the square mass attached to a rod). We recruited another group of subjects (n=12), who performed the same paradigm as the first group in this experiment, except that visual cue

about the object's actual shape was masked by an ambiguous object (Fig. 6.2D). Therefore, knowledge of the orientation of the object was only possible by the haptic perception of the object dynamics. We asked whether the interaction of Out and Return adaptive mechanisms was modulated by the visual information of the object.

6.2.4 Model fitting

We performed a model based analysis on the data to examine the interaction of Out and Return adaptive mechanisms, quantitatively. In the previous chapter, we developed a novel state space model by introducing the coupling factor that represented the amount of interaction between the learning memories of Out and Return reaching movements (Section 5.3.2). Ten different variations of the model were examined to investigate the nature of the coupling that was formed between the movement states (Table 5.1). Here, we use the same models to analyse the object manipulation data.

For each experiment and group, we fit the models to the adaptation data and the PD simultaneously. To this end, the net motor output, $y^{(n)}$ (equation 5.3) was fit to the adaptation ratio, and the absolute value of the error term, $|e^{(n)}|$, (equation 5.4) was fit to the PD. We used four mapping parameters to match the adaptation and PD values in the data with the fitting variables in the model ($y^{(n)}$ and $|e^{(n)}|$) using a linear mapping (each performance measure in the data was independently mapped to its corresponding model variable via a scaling and an offset parameter). The *BIC* was used to decide the best model for each experiment and group (equation 5.13).

In order to obtain the confidence intervals for the parameters of each model, we performed a sub-sampling bootstrap analysis, in which we produced a large number of unique samples in the following way. For each group of subjects (n=12), we first drew all possible combinations of 8 subjects (without replacement) out of 12 (i.e., 495 samples). We then went on to draw all unique combinations of 9 subjects out of 12 (220 samples), 10 subjects out of 12 (66 samples), 11 subjects out of 12 (12 samples), and finally 12 subjects out of 12 (one sample). We put all samples together (total of 794 unique samples) and fitted the models to the average data of each sample. The 95% confidence interval of each parameter for a given model was obtained

based on the 2.5th and 97.5th percentile of the parameter distribution across all samples. For the model comparison, for each given data sample we calculated the difference in *BIC* (ΔBIC) between each model and a reference model (the model that had the smallest *BIC* overall). This provided a single ΔBIC for each model and for each sample. We then took the average of ΔBIC across all samples for each model, and used that as the criterion for model comparison. Other quantities used to determine the goodness of fit, such as sum-squared error of fit (*SSE*) and the coefficient of determination (R^2), were also reported as average across all samples for any given model.

For the statistical test between two given parameters of interest in a model, we first found the difference between the bootstrap distributions of the corresponding parameters. This was done by taking the difference between all the bootstrap samples from one parameter, and every bootstrap sample from the other parameter, leading to a new distribution for parameter difference. We then performed a two tailed test on the resultant distribution against zero. The p-values in this case were obtained by calculating the ratio of samples that were smaller than zero over the total number of samples, and multiplying the resultant ratio by 2 (for the two-tailed test).

6.3 Results

6.3.1 Experiment 1

In this study, we examined the interactive properties of sensorimotor learning mechanisms associated with the Out and Return rotations in a reciprocating object manipulation task. In the first experiment, we tested whether separate representations of object dynamics are formed for Out versus Return rotations of the same object. Subjects performed the experiment in three main phases: the P1 phase (exposure in Out and error-clamp in Return), the Training phase (exposure trials in both Out and Return) and the P2 phase (the same as P1). During P1 and P2 we examined the interaction between Out and Return adaptive processes by measuring how an ongoing adaptation in the Out rotations affected the adaptive behaviour in the Return

rotations. We also asked whether this interaction was enhanced going from P1 to P2 by the concurrent adaptation during the Training phase.

Fig. 6.3A and B illustrate the overall performance of the subjects in terms of the adaptation ratio and the peak displacement (PD), separated for the Out and Return rotations. Subjects adapted to the object's dynamics during the Out rotations of P1 and P2 (Fig. 6.3A; blue data points) and they showed clear after-effects in the following washout periods (i.e., w.o.1 and w.o.2; Fig. 6.3B; blue data points). We examined how the observed behaviour in the Out rotations affected the adaptive behaviour during the Return rotations, in terms of both the adaptation ratio in P1 and P2, and the PD during w.o.1 and w.o.2.

As shown in Fig. 6.3C, the adaptation ratio for the Return movements shows a slight increase over the trials in both P1 and P2 phases, indicating a transfer of adaptation from the Out rotations to the Return rotations. This is particularly important for P1 phase, where the adaptation ratio increases in the Return movements without ever experiencing the object dynamics in this movement direction. In order to precisely quantify the amount of increase in adaptation ratio and examine its significance, we fit a simple single rate state-space model to the adaptation ratio in P1 and P2:

$$\begin{aligned} z^{(n+1)} &= A \cdot z^{(n)} + B \cdot (1 - z^{(n)}) \\ z^{(1)} &= Z_0 \end{aligned} \tag{6.1}$$

where, $z^{(n)}$ represents the adaptation ratio on trial n, and the free parameters A, B, and Z_0 are, respectively, the decay rate ($0 \leq A \leq 1$), the rate of increase of adaptation, and the initial value of adaptation in the beginning of the phase. The solid lines in Fig. 6.3C show the result of the fit for P1 (orange line) and P2 (green line). We were mainly interested in the rate of increase (B) to see if the increase of adaptation was significant. Fig. 6.3D shows the best fit value of parameter B for P1 and P2 phases, with the confidence intervals shown as error bars (see section 6.2.4 for calculation of confidence intervals). A strictly positive value for B means that the increase in adaptation is significant. As shown in the figure, the confidence intervals of B for both P1 and P2 fall within the positive values (the lower bound of B for

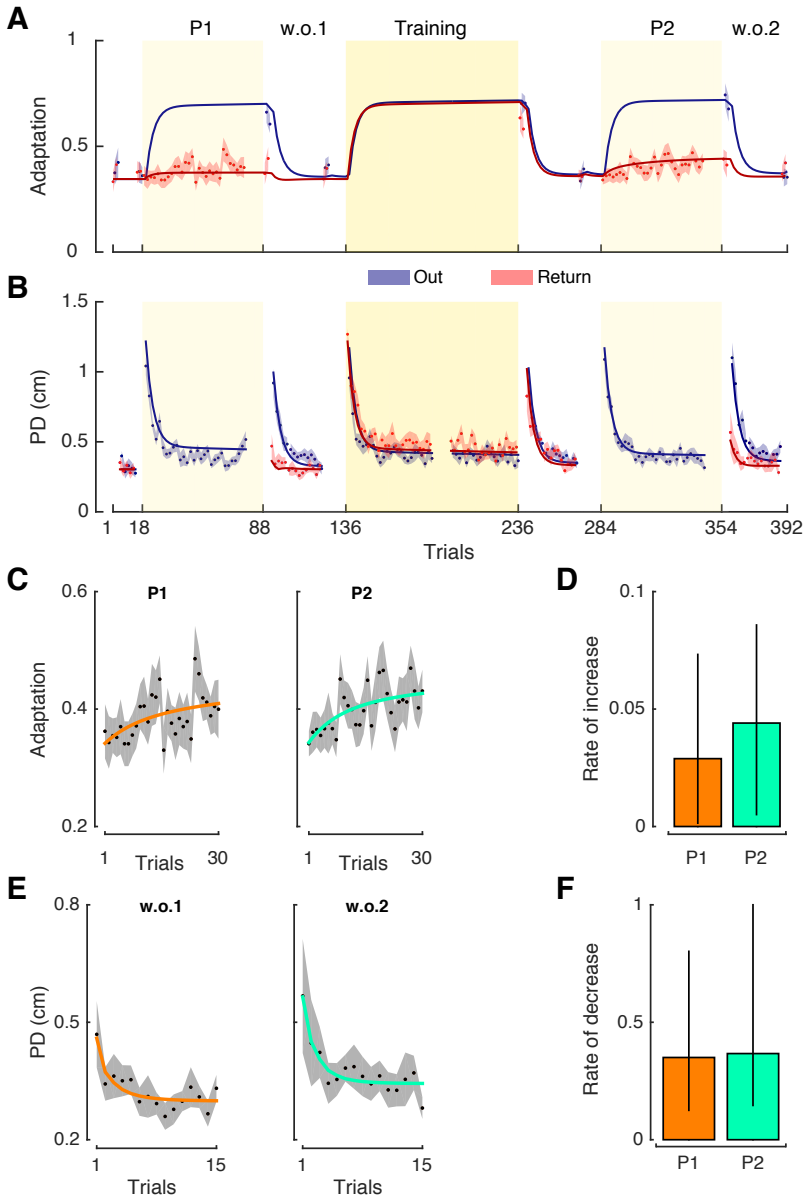

Fig. 6.3 Experiment 1. **A** and **B** show the overall performance of the subjects in terms of adaptation ratio and PD, respectively. The data is shown as the mean across subjects (the blue and red dots for Out and Return movements, respectively), with the error bars (shaded areas) as the standard error of the mean. The solid lines represent the best model fit based on model M2. **C** illustrates the adaptation behaviour during the Return movements of P1 and P2 phases. The solid lines show the single-rate state-space model fit to estimate the rate of change of the adaptation. The adaptation rate is shown in panel **D** with the 95% confidence intervals. Panel **E** also shows the after effects in the form of PD following the P1 and P2 phase (i.e., w.o.1 and w.o.2). The solid lines show the single-rate state-space model fit with the PD reduction rate shown in panel **F** (the error bars show the 95% confidence intervals).

P1 was close to zero with the value of 1.4×10^{-3}, but still positive) showing a significant, although small, increase of adaptation over trials.

We further examined the after effects of the induced adaptation during the Return trials of P1 and P2. Fig. 6.3E illustrates the PD profile during the washout periods following P1 and P2 for the Return movements. As shown, there is a clear after effect observed in the form of sudden increase in the PD in the beginning of the washout period, followed by progressive decrease towards the end of the washout. We quantified this behaviour by fitting the single rate state-space model (equation 6.1) to the PD data, in which PD was represented by the term $(1 - z^{(n)})$. In this case, the parameter B represented the rate at which the PD decreased. For a strictly decreasing PD profile (indicating the after effect), the parameter B should be strictly positive (a negative value or the value of zero for B means no reduction in the PD, and thus no after effects). Fig. 6.3F illustrates the value of B with the confidence intervals for the washout periods following P1 and P2. As shown, the whole interval of the value of B rests in the positive direction, indicating a significant after effect in the Return rotations, both for P1 and P2.

The above results show that there is an interaction, in the form of adaptation transfer, between the adaptive mechanisms of Out and Return rotations. As such, an adaptation process during the Out movements induces an adaptive response in the Return movements. This is particularly interesting about P1, which shows the interaction exists without any prior experience of object dynamics during the Return trials. Comparing the behaviour between P1 and P2, however, we did not find any significant difference, neither in the adaptation ratio, nor in the after effects. This could imply that the pre-existing interaction observed in P1 is not majorly enhanced by the simultaneous dynamic exposure between Out and Return trials during the Training phase. This can be seen from the perspective of the coupling models introduced previously (Table 5.1). We fitted all 10 variations of the model to the data and presented the goodness of fit for each model in Table (6.1). According to the ΔBIC shown in the table, the models that feature a fixed coupling factor (i.e., M1, M2 and M3) outperform other models with adaptive or mixture coupling (Fig. 6.3A illustrates the model fits based on model M2). Particularly the models with the adaptive coupling (M4, M5 and M6) show

Table 6.1 Experiment 1. The goodness of fit for each model is shown in terms of *BIC*, the sum of squared error of fitting (*SSE*) and the coefficient of determination R^2. For each model, we took the difference between the *BIC* of that model and the *BIC* of the model M2, as reported with ΔBIC.

Models	Description	ΔBIC	*SSE*	$R^2(\%)$
M1	Fixed coupling: single adaptive mechanism	0.25	2.334	75.54
M2	Fixed coupling: shared between two adaptive mechanisms	0	2.333	75.56
M3	Fixed coupling: separate for each adaptive mechanism	4.33	2.322	75.67
M4	Adaptive coupling: single adaptive mechanism	33.7	2.464	74.20
M5	Adaptive coupling: shared between two adaptive mechanisms	33.4	2.462	74.22
M6	Adaptive coupling: separate for each adaptive mechanism	45.3	2.462	74.22
M7	Mixture coupling: single adaptive mechanism	18.1	2.334	75.55
M8	Mixture coupling: shared between two adaptive mechanisms	17.6	2.330	75.58
M9	Mixture coupling: separate for each adaptive mechanism	29.8	2.295	75.96
M10	No coupling	16.1	2.476	74.18

the poorest performance in explaining the data. This suggests a somewhat settled interaction between the adaptive mechanisms of Out and Return rotations, that is not further enhanced by concurrent training. In the following sections, we conduct a more extensive experiment to better quantify the coupling effects in an object manipulation task with familiar dynamics.

6.3.2 Experiment 2

In the first experiment we mainly focused on how an adaptation process in the Out rotations affected the adaptive behaviour in the Return rotations, using a simple paradigm. In the second experiment, we aimed to explore the interactive properties of associated adaptive mechanisms more extensively by designing a richer set of dynamic manipulations between Out and Return rotations. Subjects (n=12) performed 13 blocks of Out-Return pairs of trials (Fig. 6.2). Each block consisted of either the same object dynamics for both Out and Return movements (i.e., both exposure, both error-clamp or both null trials), or different dynamics for each one (i.e., pairs of exposure and error-clamp, zero force and error-clamp, and, null and exposure).

Having a variety of different dynamic manipulations between Out and Return movements, allowed us to investigate more thoroughly the interactive behaviour of Out and Return adaptive mechanisms through our coupling models. We further investigated the effects

Table 6.2 Experiment 2. The model comparison and the goodness of fit shown for each model and each group (visual and ambiguous). The ΔBIC represents the difference between the *BIC* of each model and the model with the smallest *BIC* (here, M2) as the reference model.

Models	Visual			Ambiguous		
	ΔBIC	*SSE*	$R^2(\%)$	ΔBIC	*SSE*	$R^2(\%)$
M1	8.91	2.277	88.90	6.04	2.831	80.31
M2	0	2.229	89.13	0	2.790	80.58
M3	5.29	2.225	89.15	2.87	2.769	80.73
M4	28.7	2.320	88.69	25.7	2.883	79.95
M5	44.2	2.408	88.26	24.5	2.873	80.02
M6	44.5	2.342	88.58	35.1	2.862	80.09
M7	14.9	2.215	89.19	24.1	2.830	80.31
M8	14.2	2.209	89.22	15.8	2.774	80.69
M9	29.6	2.196	89.29	29.1	2.742	80.91
M10	120.7	3.001	85.45	12.9	2.919	79.71

of the visual presentation of the object on such interactive behaviour. A second group of subjects participated (n=12) and went through the exact same task and dynamic exposures, but without the visual information provided about the object's orientation or shape (i.e., the hammer-like object). In this group, we hid the shape of the object and replaced it with an ambiguous shape (a disk with two blades emanated from its centre; Fig. 6.2D). We asked how the visual information about the geometry of the object modulates the interaction of adaptive mechanisms associated with Out and Return rotations.

Fig. 6.4 illustrates the overall performance of the subjects based on adaptation and PD for both groups. We first perform a model-based analysis by fitting all the model variations separately to the visual and ambiguous groups. Table 6.2 compares the fitting performance between the models for each group. As shown, similar to the first experiment, the models with the fixed coupling factor (M1, M2 and M3) outperform the models with either adaptive or mixture coupling. In this case, model M2 which features a shared coupling factor between the slow and fast processes shows the smallest *BIC*, but models M1 (with a fixed coupling either for the slow or for the fast process, whichever fits the data better) and M3 (separated coupling factors for fast and slow processes) were not far behind. In Fig. 6.4 the fitting result is shown for model M2 (solid blue and red lines), with the goodness of fit of $R^2 = 0.89$ and $R^2 = 0.80$ for visual and ambiguous groups, respectively.

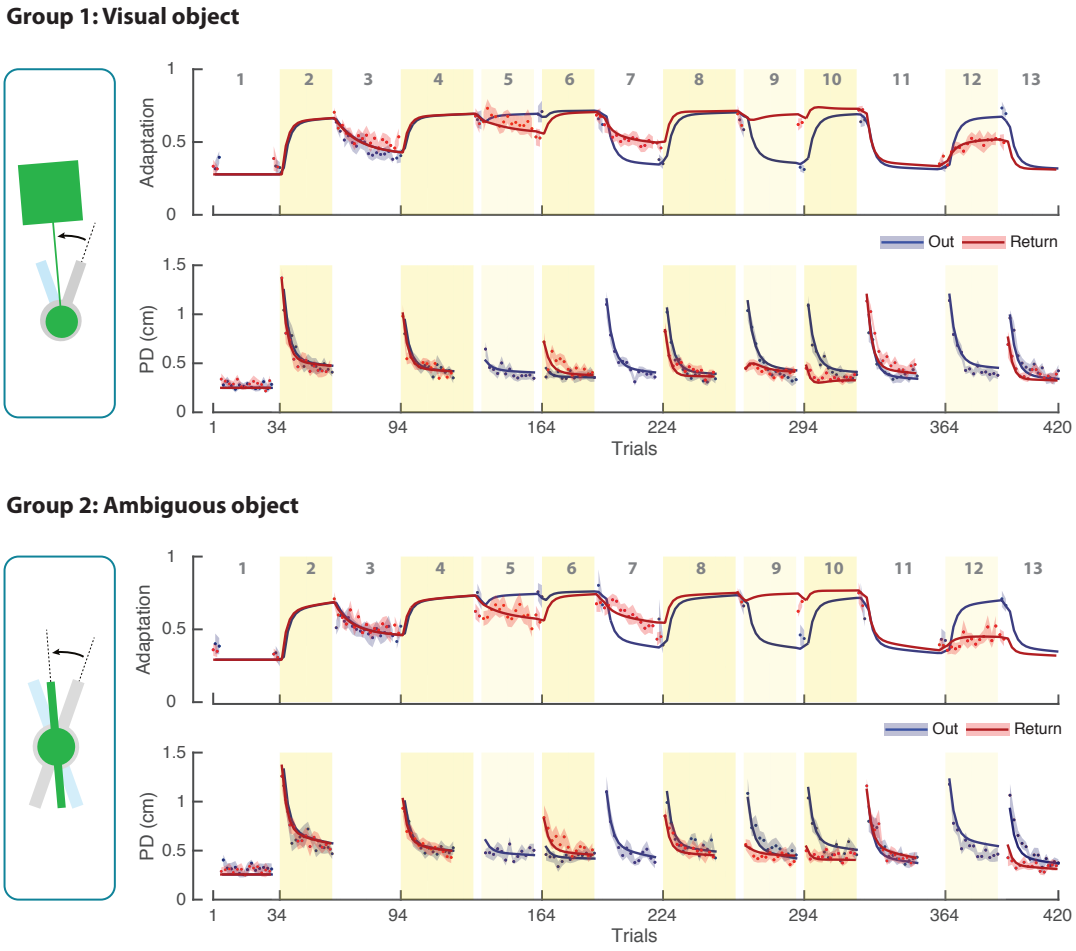

Fig. 6.4 Experiment 2. The adaptation and PD results are shown for visual and ambiguous objects. Data is separated for Out (red) and Return (blue) trials and is shown as the averaged across subjects (small dots) with the error bars as the standard error of the mean. The adaptation data is shown for error-clamp trials, and the PD is shown only for the exposure and null trials. The solid red and blue lines represent the best fit results of model M2.

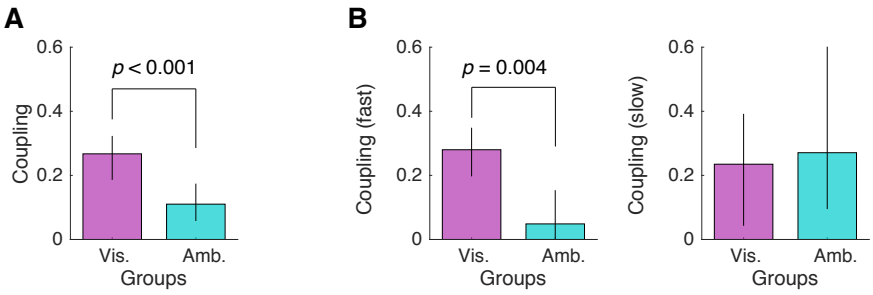

Fig. 6.5 Experiment 2. **A** The value of the shared coupling factor (fixed component) from Model M2 shown for visual (Vis.) and ambiguous (Amb.) conditions. **B** The coupling factors associated with the fast (left panel) and slow (right panel) processes from model M3.

In Fig. 6.5 the value of the coupling factor for models M2 and M3 (as the two best models) is compared between the visual and ambiguous groups. As shown, according to model M2 (with shared coupling) there is about 26% coupling between Out and Return adaptive mechanisms for the visual group. This value is dropped to about 10% for the ambiguous group, indicating a significant effect of the object's visual information on the coupling (two-tailed test against zero for the difference between bootstrap distributions of coupling factors between groups: $p < 0.001$). Note that for the second group, although no visual information is provided regarding the shape of the object, the type of coupling is the same as the visual group (that is, fixed coupling; Table 6.2 shows better performance in terms of *BIC* for the fixed coupling models). Furthermore, although the value of the coupling factor is much smaller for the ambiguous object group, its distribution (with 95% intervals) still remains larger than zero. In Fig. 6.5B we also show the coupling factors for model M3 that provides separate coupling components for fast and slow processes. As shown, the coupling factor associated with the slow process is similar for both visual and ambiguous objects ($\sim 25\%$) with no significant difference (two-tailed test on bootstrap distributions: $p = 0.923$). However, the coupling factor associated with the fast process shows a clear difference between the the two groups (about 28% for the visual group, and about 5% for the ambiguous group; two-tailed test on the corresponding bootstrap distributions $p = 0.004$). This implies that the visual information of the object mainly affects the fast process of adaptation.

Coupling in memory decay

In the previous section we showed how the overall performance of the subjects in both groups could be explained by a model that featured a fixed component for coupling factor. The amount of coupling was about 26% for the visual group, and 10% for the ambiguous group, according to model M2. Here, we examine whether the obtained coupling from the model fitting reveals itself in different parts of the data. In particular, we focus on the probe blocks 5, 7 and 12.

In blocks 5 and 7 we had error-clamp trials in Return movements while the Out movements were either exposure (block 5) or null (block 7) trials. Therefore, by comparing the Return trials between these blocks we could analyse how the adaptation behaviour in Return movements was affected by different object dynamics in the Out movements. As seen in Fig. 6.4, the adaptation in the Return trials of both blocks (red line) showed some decay. However, the amount of decay seemed to be different between the blocks. The coupling effect would predict that adaptation in the Return movements of block 5 would show slower decay (due to ongoing exposure in the Out movements), while adaptation in block 7 would show larger decay (due to deadaptation in Out movements). To test this, we fit a simple state-space model with a retention factor to the Return trials of blocks 5 and 7:

$$\begin{aligned} z^{(n+1)} &= A \cdot z^{(n)} \\ z^{(1)} &= Z_0 \end{aligned} \tag{6.2}$$

where, $z^{(n)}$ represents the adaptation ratio with the initial value Z_0, and parameter A represents the retention of adaptation. A larger retention factor means smaller decay of adaptation. We compare the retention factor between the blocks within each group and across groups to quantify the decay behaviour. Fig. 6.6A and C illustrate the decay of adaptation in blocks 5 (exposure in Out movements) and 7 (null in Out movements) for both groups with the best fit result from the equation 6.2 (solid lines). The retention factor for each block is shown in Fig. 6.6B and D with the 95% confidence intervals obtained through sub-sampling (see section 6.2.4). Comparing the blocks 5 and 7 for both groups (Fig. 6.6B and D), it is

shown that the retention factor is larger for block 5, where subjects continue on adapting to the object dynamics during the Out trials. In block 7, however, due to deadaptation during the Out trials, the decay of memory during Return trials is faster, thus showing smaller retention (two-tailed test on the bootstrap distributions between blocks: $p = 0.003$ for visual, and $p < 0.001$ for the ambiguous groups). This indicates that the decay of adaptive memory in one rotational direction (i.e., Return) is affected by the dynamics of the task in the other direction (i.e., Out) through coupling. Comparing the decay of adaptation between the visual and ambiguous conditions, however, we did not find any significant difference in the value of retention factors.

Finally, we examined the adaptive behaviour in block 12, where, following a washout of adaptation in block 11, subjects were re-exposed to the dynamics of the object in the Out movements, while performing the task under error-clamp trials during the Return movements. As shown in Fig. 6.4, in both groups subjects adapt to the dynamics of the object during Out trials of block 12 (blue data points; increase in the adaptation value and the progressive reduction of the PD). Due to this adaptation, subjects also show an increase in the adaptation level during the error-clamp trials of the Return movements (red data points in adaptation). In Fig. 6.6E and F we have a closer look at the induced adaptation in Return movements of block 12 for both visual and ambiguous conditions. To quantify the amount of induced adaptation in this block, we fitted the single-rate model from equation (6.1) to the adaptation data and obtained the rate of increase in the adaptation level (parameter B). As shown in Fig. 6.6E there is an increase in the adaptation level for both groups. However, the rate of increase is different for visual versus ambiguous conditions. According to Fig. 6.6F, the rate of increase in adaptation is about 10% for the visual condition, which is significantly larger than that in the ambiguous condition ($\sim 3\%$; two-tailed test on bootstrap distributions: $p < 0.001$). This shows that the coupling has a strong effect on the induced adaptation for the visual condition, whereas it is weaker when the the visual cue about the object is absent.

Overall, the effects of coupling were observed in both the decay of adaptation (comparing blocks 5 and 7) and the induced increase in adaptation (block 12).

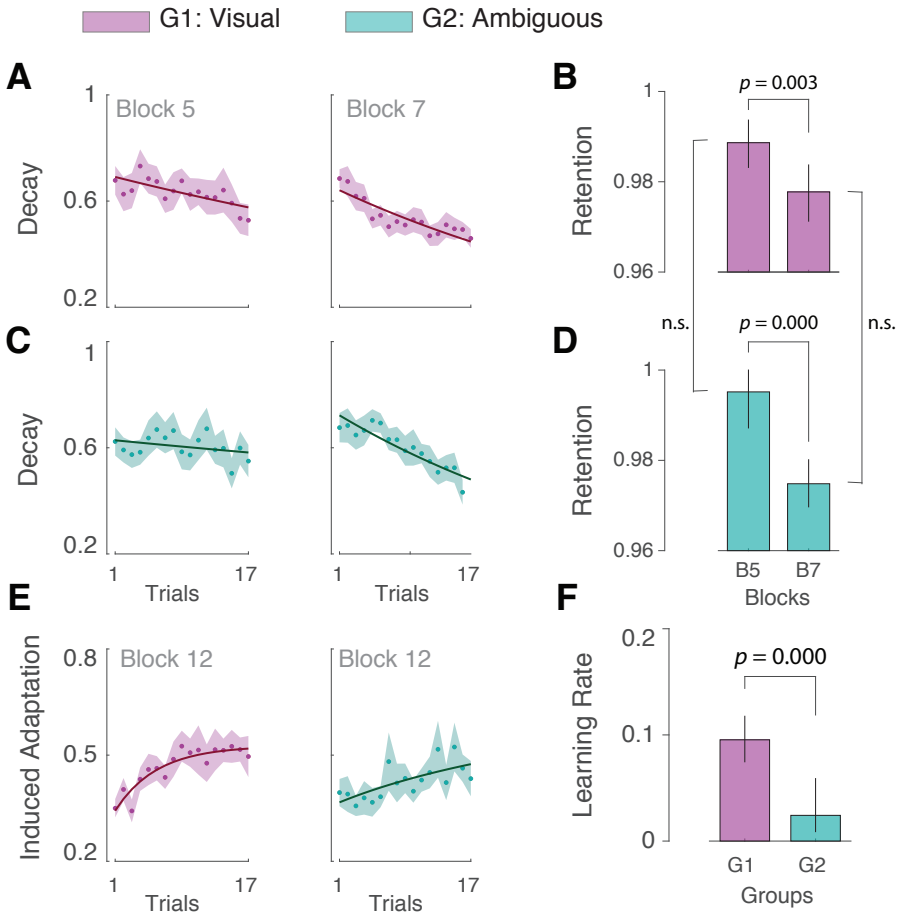

Fig. 6.6 Experiment 2. The adaptation behaviour in blocks 5, 7 and 12. **A** and **C** show the decay of adaptation in blocks 5 and 7 for the visual and ambiguous conditions, respectively. The solid lines represent the best fit single rate model from equation (6.2). **B** and **D** represent the retention factors resulted from the single rate model fit, with the 95% confidence intervals shown as error bars. **E** illustrates the induced adaptation behaviour in block 12 for both visual (left panel) and ambiguous (right panel) conditions. **F**. The rate of increase in adaptation (learning rate) for each condition.

6.4 Discussion

In this study we examined the formation of context-dependent motor memories for the dynamic manipulation of a hammer-like object, and asked how these memories interact. Our results showed that separate motor memories were responsible for the adaptation of Out versus Return rotations (contexts) when manipulating an object with familiar dynamics. This indicated that although the dynamic structure of the object was known, when the parameters of the dynamics changed in one rotation (i.e., exposure to the full dynamics of the object), the adapted behaviour did not fully transfer to the reversed rotation. Instead, we found that the motor memories associated with each rotation interacted via a fixed coupling factor which partially transferred the learning from one rotational direction to another. This fixed coupling factor indicates a constant transfer of learning that exists between the motor memories of Out and Return rotations, even before subjects experience the dynamic exposure in both rotations simultaneously (see the adaptation behaviour in P1, Fig. 6.3C and E).

The value of the fixed coupling was modulated by the visual information about the object's shape and geometry. When this information was available (group 1, experiment 2), the coupling factor was larger, showing a stronger interaction between Out and Return rotations (Fig. 6.5). Previous studies have demonstrated the importance of object's visual presentation in the selection of control commands during object manipulation. For example, in the task of lifting a U-shaped object, it has been shown that subjects generate near ideal torques on the handle of the object to prevent the tilt, only based on the shape of the object (Fu and Santello, 2012). Also, when rotating a hammer-like object, subjects have shown the ability to produce forces in the correct direction, solely based on the given visual orientation of the object, and without ever experiencing the dynamics (Ingram et al., 2010). These results indicates the importance of visual information in recalling the dynamic structure of the object. In relation to our results, it seems that when the structure of the dynamics is instantly recalled due to the available visual information about the object (group 1), it leads to stronger coupling between the memories of different rotations (Fig. 6.5A). By contrast, when the object's visual information is masked, it slows down the recall of the corresponding dynamic structure, thus leading to weaker coupling.

The difference between the coupling factors for visual versus ambiguous groups was obtained from the perspective of two different models, M2 and M3 (the models with the smallest *BIC*; Fig. 6.5). In model M2, where the coupling factor was shared between the fast and slow processes, we observed that, overall, the coupling factor for the visual group was stronger compared to ambiguous group (Fig. 6.5A). More interestingly, when we separated the coupling factors for the fast and slow processes in M3, we observed that the difference between the visual and ambiguous groups was mainly reflected in the coupling factor of the fast process, while the slow process showed similar values of coupling for both groups (Fig. 6.5B). This could imply that the fast process is more sensitive to the visual information about the object than the slow process. Indeed, previous studies have demonstrated that fast and slow processes have distinct properties and contribute differently to the learning. For instance, it is suggested that the slow process is involved in implicit learning, and is mainly error driven, whereas the fast process is mostly driven by explicit or strategic mechanisms and is sensitive to reward (for review, see Huberdeau et al., 2015). It is possible, therefore, that the object's visual information is mainly used by the explicit mechanisms to rapidly employ appropriate motor commands solely based on the shape of the object, even in the absence of error (Ingram et al., 2010). Such influence of the object's visual information is then reflected in the behaviour of the fast process, as seen in the results (Fig. 6.5B).

We further examined whether the observed coupling from the model-based analysis revealed itself in a meaningful way in the model free analysis of memory decay. We particularly looked at the adaptive behaviour in blocks 5 and 7, where the Return trials were error clamp, while the Out trials were, respectively, exposure and null. We showed that when the Out movements were exposure trials, the decay rate in the Return trials (error-clamp) was slower due to transfer of learning (coupling) from Out to Return movements. Whereas, when the Out movements were null trials, causing rapid deadaptation, the rate of decay in the Return trials was also faster (Fig. 6.6). One could argue that the difference in the decay behaviour observed in blocks 5 and 7 could be merely the effect of the order in which the blocks were introduced. In fact, all subjects performed the exposure-clamp trials (block 5)

before the null-clamp trials (block 7). We argue however, that this choice of order was in fact crucial to distinguish the effect of coupling in the decay behaviour.

If there was no coupling involved, a dual-rate model with independent Out and Return processes (i.e., model M10; table 5.1) would predict that the rate of decay in block 7 should in fact be smaller than that in block 5 (predictions not shown). This can be simply explained by the interaction between the slow and fast processes in the model. As such, by the time subjects reach to block 7, they have experienced longer period of exposure trials in the Return movements (i.e., in blocks 2, 4 and 6; Fig. 6.2B). As a result, the contribution of the slow process to adaptation behaviour increases to a larger extent by the onset of block 7, compared to block 5 (the slow process takes over the adaptation as the number of exposure trials increases; Sing and Smith, 2010; Smith et al., 2006). This leads to slower decay of adaptation in block 7 (as it is mainly driven by the slow process) compared to block 5. Accordingly, by choosing to introduce the null-clamp trial in block 7, we in fact counteracted the effect of block order (predicting slower decay) with that of coupling (predicting faster decay). The behavioural data showed that the decay behaviour in block 5 and 7 were consistent with the predictions of coupling, and not the order of blocks.

The overall results in this chapter were different from those in the previous chapter when subjects learned a novel dynamics. Specifically, when subjects performed reciprocal reaching movements under force field perturbation (novel dynamics), the coupling was mainly adaptive, evolving from almost no interaction in the beginning, to about 50% by the end of the training phase (simultaneous exposure; Fig. 5.6, left column). On the contrary, when adapting to the dynamics of an object with familiar structure, the behaviour was best described by a fixed coupling throughout the experiment. The reason for this difference might be that in the object manipulation task (familiar dynamics), the knowledge about the dynamic structure of the object is shared across all contexts. It has been shown that when manipulating a familiar object, subjects can generate forces in the correct direction for all different object orientations or movement directions, even before they experience the actual dynamics of the object (Ingram et al., 2010). This shared knowledge of dynamics across the contexts, therefore, might be the key element contributing to the fixed interaction between

the context-dependent motor memories. Whereas, when learning a novel dynamics, there is no initial knowledge about whether the dynamics applied in one movement direction (Out) will also be applied in the opposite movement direction (Return). Therefore, subjects need to experience the dynamics in both Out and Return contexts simultaneously to learn that the applied dynamics is common across both contexts. As this knowledge is gradually acquired through training, the coupling between the context-dependent memories also evolves.

Part IV

SUMMARY AND DISCUSSION

Chapter 7

General discussion

7.1 Summary

This thesis was organised in two main experimental parts in which we examined the representation and interaction of sensorimotor learning processes during reaching and object manipulation tasks. In this section, we summarise the results found in each part.

Multi-dimensional sensorimotor learning

In the first part, we introduced a novel experimental apparatus, the 3BOT, which allowed for studying force field adaptation in three-dimensional space. This feature was unique amongst usual planar robotic manipulanda used in most previous studies of motor learning. We benefited from this feature to examine the representation of learning in 3D by investigating the generalisation of adaptation (chapter 2), spontaneous recovery of learning (chapter 3) and interference and facilitation in learning (chapter 4).

In chapter 2, we demonstrated how generalisation of learning across various movement directions in a 3D environment provides insight as to the coordinate frame representation of learning in the motor system. Subjects learned a velocity-dependent force field applied in the horizontal (transverse) plane during reaching towards a training target located along the body midline. The transfer of adaptation was then probed to multiple target directions, deviated from the training target and located in the transverse plane (the plane in which the

force field was applied), the sagittal plane, and two diagonal planes. We asked whether the transfer of learning to each probe target was represented in the motor system by a coordinate frame specific to that target (target-based frame), or whether a global coordinate frame mediated the generalisation of learning across all targets. In order to examine this question, we compared the generalisation patterns obtained for two groups of probe targets. One group for which both target-based and global coordinate frames led to the same representation of generalisation (i.e., targets on the sagittal plane, where the force field was always orthogonal to the movement direction). And another group in which the representation of generalisation deviated for target-based versus global coordinate frames (i.e., targets on the transverse plane, for which the relative angle between the force field and the movement direction varied). The global representation of generalisation predicted smaller transfer of learning across targets on the transverse plane compared to the sagittal plane. Whereas, a target specific representation predicted similar patterns of generalisation for transverse and sagittal planes. The results supported the latter, indicating that generalisation of learning to each target was represented in a coordinate frame specific to that target.

In chapter 3, we examined the phenomenon of spontaneous recovery in sensorimotor learning in a novel 2-dimensional force field task. Subjects learned an initial force field applied in a specified direction (field A), followed by a short introduction of a second force field in the orthogonal direction (field B). We then probed the adaptive behaviour in a following period of error-clamp trials. The adaptation state was defined as a two-dimensional vector, in which the first component represented the adaptation to field A, and the second component represented the adaptive behaviour in field B. We found that the overall adaptation behaviour was the vector sum of adaptation in field A and field B. Spontaneous recovery appeared as the tendency to shift from the newly formed, fragile memory (B), towards the initially learned more robust memory (A). This shift revealed itself as the rotation of the adaptation vector from the most recent adaptation direction towards the direction of the initial learning.

Finally, in chapter 4, we examined the facilitation and interference effects of learning when adapting to force fields in different spatial directions. In particular, we asked

whether learning of two orthogonal force fields produced facilitation or interference. Subjects performed reaching movements to a single target in a sequence of blocks, A_1-w-B-w-A_2, including three force field blocks (A_1, B, and A_2) with short blocks of washout trials (w) in between. The initial and final blocks of force field were exactly the same (A_1=A_2), while the direction of the force field in the middle block (B) was either the same, the opposite, orthogonal, or neutral (null) with respect to the initial force field. We compared the adaptation performance between the first and second encounter to field A (that is, A_1 and A_2) as a function of the force field direction in the middle block (field B). The results implied that the effect of field B on the memory of field A depended on the similarity of errors observed between field A and field B. When both fields caused errors in the same direction, it led to facilitation (improved performance in A_2); when errors occurred in opposite directions, it caused interference (worsened performance in A_2); and when errors were in orthogonal directions, no marked effect of facilitation or interference was observed (performance in A_2 was not different from A_1). These results suggested that the interaction of motor memories associated with different force field directions might be tuned to the relative angle between the errors caused by each force field.

Interaction between sensorimotor learning systems

In the second experimental part of the thesis, we examined how multiple context-dependent motor memories interact. Subjects performed reciprocating reaching movements (chapter 5) and object rotations (chapter 6) which consisted of two movement cycles with the same magnitude, but in opposite directions (Out and Return). Each movement cycle was considered as a separate kinematic context. We manipulated the dynamics of the movement in one movement context (i.e., Out) and measured how behaviour changed in the other (Return). We found that separate motor memories were responsible for the control of movement in each context, and that these motor memories interacted by sharing error signals for adaptation. The nature of this interaction, however, differed depending on the dynamic structure of the task. When motor memories were formed during adaptation to novel dynamics (i.e., force field exposure), the interaction was flexible and enhanced or weakened depending on the

task dynamics in each movement context (chapter 5). Whereas, when subjects adapted to the dynamics of an object with familiar dynamics (i.e., a hammer-like object), the interaction between Out and Return motor memories was mainly fixed and consistent throughout (chapter 6). This behaviour was well accounted for by a multi-context, dual-rate state-space model that featured a coupling factor to represent interaction (transfer of learning) between Out and Return adaptive states. For the learning of a novel dynamics, the coupling factor took an adaptive form, evolving across trials depending on the perturbation schedule in each context. For the case of familiar dynamics, the coupling factor took on a fixed value across all conditions.

7.2 Current work and future direction

Study of naturalistic behaviour

Behavioural studies of motor learning have been mostly focused on reaching movements performed in constrained planar workspace. Indeed, applying careful and well-motivated constraints in the experimental design is helpful and sometime necessary to control for the effects of unwanted variables, and track the behaviour of the variables of interest in question. However, it should also be noted that the motor system is, in essence, a complex system that naturally operates in a multi-dimensional space. As such, any attempt to project this multi-dimensional behaviour onto a space with reduced dimensionality could lead to the loss of valuable information. Hence, to understand such a system more accurately, it may be essential to shift the focus of our studies towards more naturalistic movement behaviour. This, of course, requires development of appropriate experimental apparatus that allows for manipulation and measurement of movements that resemble natural behaviour in everyday life.

In the current study, we introduced the 3BOT which was designed to expand the study of motor learning to multi-dimensional movements. In this case, as a simple demonstration for the first-time use of the device, we studied three simple paradigms which were widely used in the past in the planar case, but were extended here to three-dimensional space. This provided

useful information about the underlying mechanisms of sensorimotor learning which were hidden in planar paradigms. For example, the dissociation of global versus target-specific representation of generalisation was only made possible when we were able to apply the force field direction and the movement direction in two orthogonal planes (that is, force field was applied in the transverse plane and the targets were located in the sagittal plane; chapter 2). As a different example, we showed how facilitation or interference between two sensorimotor transformations was modulated by the spatial direction in which these transformations were applied (chapter 4). This was also made possible using a 3D capable device. These examples, though simple in design, indicate how studying multi-dimensional and natural movements provides a broader, more thorough view of how the motor system might work.

State-space models

Over recent years, many variations of state-space models have been used in sensorimotor learning studies as a powerful mechanistic approach to understanding the learning process (for example, Ingram et al., 2011; Lee and Schweighofer, 2009; Smith et al., 2006; Thoroughman and Shadmehr, 2000). The main feature of this class of models has been the ability to summarise the learning behaviour down to two main characteristics: the rate of learning, and the rate of decay. As such, a single adaptive process is described by how fast it learns from error, and how quickly it forgets in the absence of error. Complex human learning behaviour has been explained only based on how these parameters are determined in different conditions.

One group of studies have focused on the variation of these parameters within a single context and across different statistics of the environment. For instance it has been proposed that the learning rate varies with noise in the estimation of environmental or internal states (Burge et al., 2008; Kording et al., 2007; Wei and Körding, 2010), and adjusts according to the consistency of the perturbations (Gonzalez Castro et al., 2014; Herzfeld et al., 2014). A different group of studies have highlighted the tuning patterns of learning/decay parameters across multiple contexts. As such, each parameter is a function of the distance between the current context and other non-active contexts. For example, it has been suggested that the

rate of learning or decay is largest for the current active context, and diminishes progressively for contexts deviated farther away (e.g., Donchin et al., 2003; Ingram et al., 2013, 2011; Lee and Schweighofer, 2009). These studies mainly assume that the value of each parameter withing a given context is constant.

Adding to the large body of studies mentioned above, here we presented a novel aspect of the learning behaviour through state-space models, which had not been examined before (chapter 5). We assessed the context-dependent learning behaviour and asked how the tuning pattern of the learning rate across different contexts varied based on the variation of environmental dynamics. This was investigated by introducing a coupling factor that tuned the learning rate across multiple contexts (similar to previous context-dependent models), but with the ability to adjust the tuning to increase efficiency in learning (new feature introduced in this study). In the developing of the model, we also made a few assumptions. Notably, we assumed that the interaction between the context-dependent processes, only affected the learning rate, and not the decay rate. As such, the rate of decay is assumed to be context-independent and invariable under various dynamic changes in the environment. This is a common assumption in most of the context-dependent state-space models (Donchin et al., 2003; Lee and Schweighofer, 2009; Nozaki and Scott, 2009; Tanaka et al., 2012, 2009; Thoroughman and Shadmehr, 2000), although, it has been shown that under specific scenarios the rate of decay will vary with context and time (Ingram et al., 2017, 2013; Kitago et al., 2013). Further work is required to examine the context dependent features of the decay rate, and whether the motor system adjusts this rate across multiple contexts in an adaptive way.

Model fitting and confidence intervals

An important step in validating the predictive and generative features of a candidate model is to tune the parameters of the model to best describe the available experimental data (i.e., model fitting). Model fitting can be challenging in various aspects, from avoiding observation noise in the fitting process to obtaining the level of certainty over the estimated model parameters.

Depending on the richness of collected data, model fitting can be performed in the individual, or the group level. For example, in perceptual decision making studies where a large number of data points can be collected for each subject and condition, models are normally fit to the individual data (e.g., Resulaj et al., 2009; Selen et al., 2012; Van Den Berg et al., 2016; van den Berg et al., 2016). This enables the analysis of performance for each individual and shows how the model parameters vary from one subject to another (between-subject variability). By contrast, in sensorimotor learning studies models are normally fit to the group data (e.g., Ingram et al., 2013, 2011; Kim et al., 2015a,b; Smith et al., 2006), as the individual data is usually too noisy and leads to reduced quality of fit (Ingram et al., 2017; Kim et al., 2015a). The main source of such noise is the fact that studies of motor learning normally examine the time course of learning, where at each point in time (i.e., trial), only a single data point (adaptation or kinematic error) can be measured for each individual. Therefore, the collected data point for each trial could be highly uncertain as to how robustly it represents the actual state of learning. Furthermore, measurement of adaptation usually relies on channel trials (Scheidt et al., 2000), in which subjects sometimes show strange behaviour, especially at the early stages of learning. This also results in increased noise. In this case, data is usually averaged across subjects so as to reduce the noise and hence provide better quality of fit.

When a model is fit to the group data, only a single value is obtained for each model parameter. An important part of model fitting is then to estimate the confidence/credibility interval over each parameter. This interval is an estimated range of values in the parameter space that for any given sample from the actual (unseen) population, there is a minimum probability (e.g., 95%) that the resultant fit parameters would fall within this interval. There are two main approaches to estimating such interval, with each offering a slightly different perspective towards how such intervals should be interpreted: the non-parametric, frequentist approach, and the parametric, Bayesian approach (VanderPlas, 2014). In the current study, we used bootstrapping (in chapter 5) and subsampling (in chapter 6) techniques as non-parametric approaches to obtaining confidence intervals. Bootstrapping (Efron, 1992) is the procedure of resampling, with replacement, from the original sample (available data),

at the same sample size. In this case, samples are drawn with the right size, but from an empirical distribution. Subsampling (Politis et al., 1999), on the other hand, is described as resampling from the original sample without replacement, and at a smaller size. In this case, samples are drawn from the true distribution, but the sample size is smaller than the original size. Both techniques are simple to implement and commonly used in the literature, though recent studies have suggested that subsampling provides a more robust analysis in certain conditions (De Bin et al., 2016; Geyer, 2006). The issue with subsampling is that for it to work, the original sample size needs to be large enough, so that the number of drawn subsamples is reasonable. For example, in Chapter 6, each experimental group consisted of 12 subjects, from which all possible subsamples of sizes 8 and above constituted about 800 unique samples. Whereas, in Chapter 5, each experimental group consisted of 10 subjects, and all possible subsamples of size 8 and above in this case would result in only 56 unique samples, which is too small. As such, subsampling is suitable when the number of subjects is relatively large.

An alternative way to finding the certainty over the model parameters is the parametric, Bayesian approach. In this approach, model parameters are treated as probability distributions. In brief, a prior probability distribution is assumed for each parameter, which is then updated by the observed data, through Bayesian inference, to obtain the posterior distribution (i.e., the probability of each parameter given the observed data). The certainty over each parameter, known as the credibility interval, is then obtained as the range of parameter values that cover 95% of the posterior distribution. The Bayesian approach highly relies on the choice of the prior distribution and its hyper-parameters (that is, the parameters of the prior distribution). The hyper-parameters are in some cases adjusted in the process of model fitting so as to maximise the likelihood of the observed data. In a more rigorous approach, however, the hyper-parameters could themselves be considered as probability distributions (hyperpriors) that are integrated into the Bayesian inference. This is known as the *hierarchical* Bayesian modelling (Allenby, 2006).

There are advantages and disadvantages in using either the non-parametric or the Bayesian approach in finding confidence/credibility intervals. Namely, in the non-parametric approach,

samples are drawn from a discrete set of observations. This is problematic when the original sample size is small, because it only covers a restricted range of values, and thus the amount of variation in the actual population is severely underestimated. The Bayesian approach in this case provides a more powerful analysis, and can also be used to estimate the model parameters for noisy individual data. As such, even when the model is fit to individual data, the Bayesian inference prevents the parameters from taking impossible (e.g., negative) or nonsensical values: the probability of a parameter taking an impossible value is highly restricted through the prior probability distribution. On the other hand, although Bayesian approach provides more power than the non-parametric bootstrap, it does so on the basis of an inherently arbitrary choice of prior distribution. This makes the Bayesian inference highly susceptible to wrong outcomes if the prior distribution is badly chosen or just assumed out of thin air. This is, indeed, not an issue for non-parametric methods, as in these methods there is no assumption about how the observed data (or the parameters of the model) are distributed.

In most sensorimotor learning studies, multi-rate state-space models of adaptation are used to describe experimental data. These models, in the conventional form (Smith et al., 2006), are not formulated in the Bayesian form. Therefore, confidence intervals for the parameters of such models are more commonly obtained through non-parametric methods (e.g., Ingram et al., 2013, 2011; Kim et al., 2015a,b; Smith et al., 2006). In line with these studies, here we also developed a context-dependent state-space model in which bootstrapping and subsampling methods were used to obtain confidence intervals. In recent years, however, researchers have developed a framework in which multi-rate state-space models could also be formulated in the Bayesian form with uncertainty taken into account in the learning process (Albert and Shadmehr, 2017; Cheng and Sabes, 2006; Heald et al., 2018; Kording et al., 2007; Tanaka et al., 2012). Such framework has made it possible to use hierarchical Bayesian modelling in obtaining the credibility intervals for model parameters.

Single-rate versus dual-rate learning

From the perspective of state-space models, it is widely assumed that the learning behaviour is the product of at least two learning processes that operate in parallel, and that each is

characterised with a distinct rate of decay and learning (Lee and Schweighofer, 2009; Smith et al., 2006). In this case, one process is slow, stable, implicit and expressive in low reaction times; and the other is fast and explicit, and requires longer preparation time to be expressed (for review, see Huberdeau et al., 2015). This decomposition of behaviour has been mainly shown for tasks that require learning of a novel sensorimotor transformation (e.g., visuomotor rotation or force field perturbation). By contrast, when learning a transformation with familiar dynamics, it has been suggested that a single process is sufficient to account for the adaptive behaviour (Ingram et al., 2013, 2011).

Here, we developed a set of models based on the dual-rate representation of learning, which was used to describe the learning behaviour in the reciprocal reaching movement (i.e., learning novel dynamics; chapter 5). The same set of models were also used to examine the behaviour in our object manipulation task, which was regarded as a task with familiar dynamics (chapter 6). One might consider, on the basis of previous works (Ingram et al., 2011), whether the dual-rate model used to analyse our object manipulation task was excessive, and instead a single-rate model could similarly account for the behaviour. To address this concern, we tested a single rate model with a fixed coupling factor (similar to the model M1 in Fig. 5.4B, with $\mathbf{x}_2^{(n)} = 0$) and compared it to the winning model in chapter 6 (i.e., M2: the dual-rate model with a fixed coupling factor for both model rates; table 6.2). Fig. 7.1A illustrates the model fits for the single-rate (SRM) and dual-rate (DRM) models with a fixed coupling factor. The models were fit to the data from the visual group in experiment 2 of chapter 6 (Fig. 6.2). For clarity, only the model fits to the Return trials are presented. As shown, it is clear that the dual-rate model out-performs the single-rate model both in the quality of fit (sum of squared error between the model fit and the data; Fig. 7.1B) and the BIC criterion (Fig. 7.1C). This indicates that the adaptive behaviour in this case is governed by a dual-rate process.

This might seem inconsistent with previous findings suggesting that a single-rate process is sufficient to explain the learning behaviour during object manipulation (Ingram et al., 2011). However, there are substantial differences in methodology between the current work and the previous work. In the previous study, subjects performed clockwise (CW) and

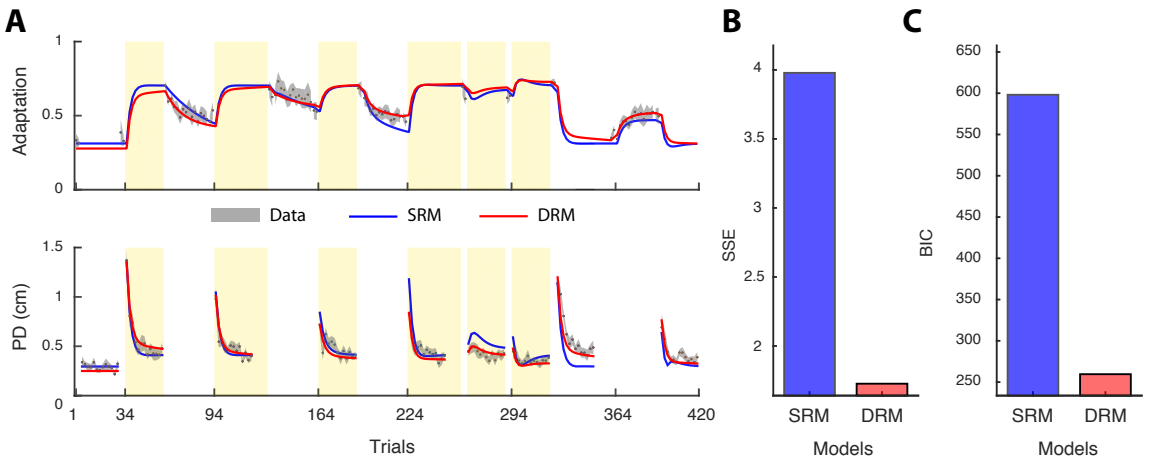

Fig. 7.1 **A**. Model fits for the single-rate (SRM) and dual-rate (DRM) models with fixed coupling factor. The data is taken from experiment 2 in chapter 6 (visual group). For clarity, the model fits are shown only for the Return trials, though the models were fit to all trials at the same time. **B**. The sum-squared error of fit for SRM and DRM. **C**. The BIC criterion obtained for each model fit.

counter-clockwise (CCW) rotations of an object (similar to Out and Return rotations in our experiment), where the dynamics of the task was always the same for both rotations throughout the experiments. Furthermore, the model fitting in the previous work was performed on the collective data from both rotations put together (that is, CW and CCW trials were combined in model fitting). In contrast, in our study we separated the Out and Return rotations and used distinct (yet interacting) learning mechanisms for each one. In addition, the dynamics of the task in each block of trials was varied between the Out and Return rotations in several different ways. These differences along with the results shown in Fig. 7.1 suggest that the behaviour in each rotation (Out or Return) might be governed separately by a dual-rate process, and this is more pronounced when the dynamics of the task is differed between the Out and Return rotations. However, when the dynamics is the same and the trials from both rotations are put together, the overall performance seems to be sufficiently accounted for by a single-rate process. Further investigation is required, both experimentally and theoretically, to determine the conditions in which a combination of two interacting dual-rate processes (i.e., for Out and Return rotations) leads to a single-rate behaviour.

Future direction

The experimental and computational results presented in this work raise novel questions and promise potential ideas for future research. One of the main subjects in motor learning studies has been the role of error in sensorimotor adaptation (Shadmehr et al., 2010). Substantial amount of work has been done regarding how error signals are encoded and compensated for in the motor system. These studies, however, mostly regard error as a one-dimensional (scalar) variable (for example, Herzfeld et al., 2014; Takiyama et al., 2015) and very few have examined error as a multi-dimensional vector in adaptation studies (e.g., Ingram et al., 2017). An important question in this case, is how the motor system encodes error signals in both magnitude and direction, and how the formation and recall of motor memories are modulated by the spatial direction of error. For instance, in chapter 4 we found that motor memories formed based on orthogonal error directions do not interfere. However, the exact modulation of such interaction by error direction is unknown. This can be examined using the 3BOT and carefully controlling for error directions in an adaptation task. For example, one could implement a planar channel, in which the movement is allowed within a defined plane in the 3D space, but is constrained orthogonal to that plane. By controlling the angle of this planar channel, one could control the direction in which errors might occur.

In the second part of this study we examined the interaction of multiple context-dependent motor memories, and how it varies depending on the dynamics of the task. This interaction was mediated by a coupling factor that scaled the learning rate across contexts (highest rate for the active context and lower for the non-active context), and could increase or decrease as a function of experienced error. A remaining question in this case is how the rate of memory decay is determined across different contexts, and whether it can alter in a meaningful way to improve retention and recall of motor memories. In other words, although we assumed that the coupling factor only affects the rate of learning, it is not clear how the rate of decay might be influenced by such adaptive coupling. This requires further experimental and computational work to be understood.

7.3 Conclusion

The sensorimotor system recruits multiple learning processes to remain constantly alert to the environmental or internal changes, and to adapt accordingly to preserve the optimal or satisfactory performance. During this complex process, accurate coordinate representations are required in which the motor system defines the task goal, calculates the deviation from the goal (error), and prepares the appropriate motor commands to achieve the goal (reducing the error). These representations are not universal, but remain specific to the context in which the task is introduced, yet also transfer properly between the contexts to potentially increase efficiency.

For many years, scientists have endeavoured to understand this fascinating process. How such coordinate representations are made, how the errors are encoded in the motor system, and how contextual information is used to serve memory formation and recall have been a few of many questions researchers attempted to answer. In the current work, we tried to highlight some of the features of the sensorimotor learning system by examining the representation and interaction of motor memories when adapting to novel and familiar dynamics. By introducing a novel experimental apparatus (the 3BOT), we investigated the coordinate representation of generalisation in force field learning, and tested the interaction of memories formed based on orthogonal force fields. Also, by examining context-dependent learning, we showed how motor memories associated with different contexts interact, and how such interaction improves depending on the dynamic characteristics of the environment. Our finding add to the understanding of the sensorimotor learning processes, and raise further questions and potential ideas for future research.

References

Adelstein, B. D. (1989). *A virtual environment system for the study of human arm tremor*. PhD thesis, Massachusetts Institute of Technology.

Ahmed, A. A., Wolpert, D. M., and Flanagan, J. R. (2008). Flexible Representations of Dynamics Are Used in Object Manipulation. *Current Biology*, 18(10):763–768.

Albert, S. T. and Shadmehr, R. (2017). Estimating properties of the fast and slow adaptive processes during sensorimotor adaptation. *Journal of Neurophysiology*, pages jn–00197.

Allenby, G. M. (2006). Hierarchical bayes models. *The handbook of marketing research: Uses, misuses, and future advances*, pages 418–440.

Bays, P. M., Flanagan, J. R., and Wolpert, D. M. (2005). Interference between velocity-dependent and position-dependent force-fields indicates that tasks depending on different kinematic parameters compete for motor working memory. *Experimental Brain Research*, 163(3):400–405.

Bays, P. M. and Wolpert, D. M. (2006). Actions and consequences in bimanual interaction are represented in different coordinate systems. *The Journal of neuroscience*, 26:7121–7126.

Berniker, M., Franklin, D. W., Flanagan, J. R., Wolpert, D. M., and Kording, K. (2014). Motor learning of novel dynamics is not represented in a single global coordinate system: evaluation of mixed coordinate representations and local learning. *Journal of neurophysiology*, 111(6):1165–82.

Berniker, M. and Kording, K. (2008). Estimating the sources of motor errors for adaptation and generalization. *Nature neuroscience*, 11(12):1454–61.

Berniker, M. and Kording, K. P. (2011). Estimating the relevance of world disturbances to explain savings, interference and long-term motor adaptation effects. *PLoS computational biology*, 7(10):e1002210.

Bock, O., Schneider, S., and Bloomberg, J. (2001). Conditions for interference versus facilitation during sequential sensorimotor adaptation. *Experimental brain research*, 138(3):359–365.

Brashers-Krug, T., Shadmehr, R., and Bizzi, E. (1996). Consolidation in human motor memory. *Nature*, 382(6588):252–5.

Braun, D. A., Aertsen, A., Wolpert, D. M., and Mehring, C. (2009). Learning optimal adaptation strategies in unpredictable motor tasks. *The Journal of neuroscience*, 29(20):6472–8.

Braun, D. A., Mehring, C., and Wolpert, D. M. (2010). Structure learning in action. *Behavioural brain research*, 206(2):157–65.

Brayanov, J. B., Press, D. Z., and Smith, M. A. (2012). Motor memory is encoded as a gain-field combination of intrinsic and extrinsic action representations. *The Journal of neuroscience*, 32(43):14951–65.

Burdet, E., Osu, R., Franklin, D. W., Milner, T. E., and Kawato, M. (2001). The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature*, 414(6862):446.

Burge, J., Ernst, M. O., and Banks, M. S. (2008). The statistical determinants of adaptation rate in human reaching. *Journal of vision*, 8(4):1–19.

Burgess, J. K., Bareither, R., and Patton, J. L. (2007). Single limb performance following contralateral bimanual limb training. *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society*, 15(3):347–55.

Caithness, G., Osu, R., Bays, P., Chase, H., Klassen, J., Kawato, M., Wolpert, D. M., and Flanagan, J. R. (2004). Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks. *The Journal of neuroscience*, 24(40):8662–71.

Cashaback, J. G., McGregor, H. R., Mohatarem, A., and Gribble, P. L. (2017). Dissociating error-based and reinforcement-based loss functions during sensorimotor learning. *PLoS computational biology*, 13(7):e1005623.

Cheng, S. and Sabes, P. N. (2006). Modeling Sensorimotor Learning with Linear Dynamical Systems. *Neural Computation*, 18(4):760–793.

Conditt, M. A., Gandolfo, F., and Mussa-Ivaldi, F. A. (1997). The Motor System Does Not Learn the Dynamics of the Arm by Rote Memorization of Past Experience. *Journal of Neurophysiol*, 78(1):554–560.

Conditt, M. A. and Mussa-Ivaldi, F. A. (1999). Central representation of time during motor learning. *Proceedings of the National Academy of Sciences*, 96(20):11625–11630.

Criscimagna-Hemminger, S. E., Bastian, A. J., and Shadmehr, R. (2010). Size of error affects cerebellar contributions to motor learning. *Journal of neurophysiology*, 103(4):2275–84.

Criscimagna-Hemminger, S. E., Donchin, O., Gazzaniga, M. S., and Shadmehr, R. (2003). Learned dynamics of reaching movements generalize from dominant to nondominant arm. *Journal of neurophysiology*, 89(1):168–76.

Criscimagna-Hemminger, S. E. and Shadmehr, R. (2008). Consolidation patterns of human motor memory. *The Journal of neuroscience*, 28(39):9610–8.

Davidson, P. R. and Wolpert, D. M. (2004). Scaling down motor memories: de-adaptation after motor learning. *Neuroscience letters*, 370(2-3):102–7.

De Bin, R., Janitza, S., Sauerbrei, W., and Boulesteix, A.-L. (2016). Subsampling versus bootstrapping in resampling-based model selection for multivariable regression. *Biometrics*, 72(1):272–280.

de Vlugt, E., Schouten, A. C., van der Helm, F. C., Teerhuis, P. C., and Brouwn, G. G. (2003). A force-controlled planar haptic device for movement control analysis of the human arm. *Journal of neuroscience methods*, 129(2):151–168.

Diedrichsen, J., White, O., Newman, D., and Lally, N. (2010). Use-dependent and error-based learning of motor behaviors. *The Journal of neuroscience*, 30(15):5159–66.

Dizio, P. and Lackner, J. R. (1995). Motor adaptation to Coriolis force perturbations of reaching movements: endpoint but not trajectory adaptation transfers to the nonexposed arm. *Journal of neurophysiology*, 74:1787–1792.

DiZio, P. and Lackner, J. R. (2000). Congenitally blind individuals rapidly adapt to coriolis force perturbations of their reaching movements. *Journal of neurophysiology*, 84:2175–2180.

Donchin, O., Francis, J. T., and Shadmehr, R. (2003). Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. *The Journal of neuroscience*, 23(27):9032–45.

Efron, B. (1992). Bootstrap methods: another look at the jackknife. In *Breakthroughs in statistics*, pages 569–593. Springer.

Fayé, I. C. (1986). *An impedance controlled manipulandum for human movement studies*. PhD thesis, Massachusetts Institute of Technology.

Ferrera, V. P. and Barborica, A. (2010). Internally generated error signals in monkey frontal eye field during an inferred motion task. *Journal of Neuroscience*, 30(35):11612–11623.

Flanagan, J. R., Vetter, P., Johansson, R. S., and Wolpert, D. M. (2003). Prediction precedes control in motor learning. *Current Biology*, 13(2):146–150.

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

Francis, J. T. (2008). Error generalization as a function of velocity and duration: human reaching movements. *Experimental brain research*, 186(1):23–37.

Franklin, D. W. and Milner, T. E. (2003). Adaptive control of stiffness to stabilize hand position with large loads. *Experimental Brain Research*, 152(2):211–220.

Fu, Q. and Santello, M. (2012). Context-dependent learning interferes with visuomotor transformations for manipulation planning. *The Journal of neuroscience*, 32(43):15086–92.

Galea, J. M., Mallia, E., Rothwell, J., and Diedrichsen, J. (2015). The dissociable effects of punishment and reward on motor learning. *Nature neuroscience*, 18(4):597–602.

Geyer, C. J. (2006). 5601 notes: The subsampling bootstrap.

Ghahramani, Z., Wolpert, D. M., and Jordan, M. I. (1996). Generalization to Local Remappings of the Visuomotor Coordinate Transformation. *The journal of neuroscience*, 16(21):7085–7096.

Gonzalez Castro, L. N., Hadjiosif, A. M., Hemphill, M. A., and Smith, M. A. (2014). Environmental consistency determines the rate of motor adaptation. *Current biology : CB*, 24(10):1050–61.

Goodbody, S. J. and Wolpert, D. M. (1998). Temporal and amplitude generalization in motor learning. *Journal of neurophysiology*, 79(4):1825–38.

Haith, A. M., Huberdeau, D. M., and Krakauer, J. W. (2015). The influence of movement preparation time on the expression of visuomotor learning and savings. *The Journal of neuroscience*, 35(13):5109–17.

Harris, C. S. (1965). Perceptual adaptation to inverted, reversed, and displaced vision. *Psychological review*, 72(6):419.

Hatada, Y., Miall, R., and Rossetti, Y. (2006). Two waves of a long-lasting aftereffect of prism adaptation measured over 7 days. *Experimental brain research*, 169(3):417–426.

Heald, J. B., Ingram, J. N., Flanagan, J. R., and Wolpert, D. M. (2018). Multiple motor memories are learned to control different points on a tool. *Nature human behaviour*, 2(4):300.

Hebb, D. O. (2005). *The organization of behavior: A neuropsychological theory*. Psychology Press.

Herzfeld, D. J., Vaswani, P. A., Marko, M., and Shadmehr, R. (2014). A memory of errors in sensorimotor learning. *Science (New York, N.Y.)*, pages science.1253138–.

Hirashima, M. and Nozaki, D. (2012). Distinct motor plans form and retrieve distinct motor memories for physically identical movements. *Current biology : CB*, 22(5):432–6.

Hogan, N. (1985). The mechanics of multi-joint posture and movement control. *Biological cybernetics*, 52(5):315–331.

Howard, I. S. and Franklin, D. W. (2015). Neural tuning functions underlie both generalization and interference. *PloS one*, 10(6):e0131268.

Howard, I. S., Ingram, J. N., Franklin, D. W., and Wolpert, D. M. (2012). Gone in 0.6 seconds: the encoding of motor memories depends on recent sensorimotor states. *Journal of Neuroscience*, 32(37):12756–12768.

Howard, I. S., Ingram, J. N., and Wolpert, D. M. (2009). A modular planar robotic manipulandum with end-point torque control. *Journal of neuroscience methods*, 181(2):199–211.

Howard, I. S., Wolpert, D. M., and Franklin, D. W. (2013). The effect of contextual cues on the encoding of motor memories. *Journal of neurophysiology*, 109(10):2632–44.

Howard, I. S., Wolpert, D. M., and Franklin, D. W. (2015). Report The Value of the Follow-Through Derives from Motor Learning Depending on Future Actions. *Current Biology*, 25(3):397–401.

Huang, V. S., Haith, A., Mazzoni, P., and Krakauer, J. W. (2011). Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. *Neuron*, 70(4):787–801.

Huang, V. S. and Shadmehr, R. (2009). Persistence of motor memories reflects statistics of the learning event. *Journal of neurophysiology*, 102(2):931–40.

Huberdeau, D. M., Krakauer, J. W., and Haith, A. M. (2015). Dual-process decomposition in human sensorimotor adaptation. *Current Opinion in Neurobiology*, 33:71–77.

Hwang, E. J., Smith, M. A., and Shadmehr, R. (2006). Adaptation and generalization in acceleration-dependent force fields. *Experimental brain research*, 169(4):496–506.

Imamizu, H., Sugimoto, N., Osu, R., Tsutsui, K., Sugiyama, K., Wada, Y., and Kawato, M. (2007). Explicit contextual information selectively contributes to predictive switching of internal models. *Experimental brain research*, 181(3):395–408.

Ingram, J., Sadeghi, M., Flanagan, J., and Wolpert, D. (2017). An error-tuned model for sensorimotor learning. *PLoS Computational Biology*.

Ingram, J. N., Flanagan, J. R., and Wolpert, D. M. (2013). Context-dependent decay of motor memories during skill acquisition. *Current biology : CB*, 23(12):1107–12.

Ingram, J. N., Howard, I. S., Flanagan, J. R., and Wolpert, D. M. (2010). Multiple grasp-specific representations of tool dynamics mediate skillful manipulation. *Current biology : CB*, 20(7):618–23.

Ingram, J. N., Howard, I. S., Flanagan, J. R., and Wolpert, D. M. (2011). A single-rate context-dependent learning process underlies rapid adaptation to familiar object dynamics. *PLoS computational biology*, 7(9):e1002196.

Inoue, M., Uchimura, M., Karibe, A., O'Shea, J., Rossetti, Y., and Kitazawa, S. (2015). Three timescales in prism adaptation. *Journal of neurophysiology*, 113(1):328–38.

Izawa, J. and Shadmehr, R. (2011). Learning from sensory and reward prediction errors during motor adaptation. *PLoS computational biology*, 7(3):e1002012.

Joiner, W. M., Brayanov, J. B., and Smith, M. A. (2013). The training schedule affects the stability, not the magnitude, of the interlimb transfer of learned dynamics. *Journal of neurophysiology*, 110(4):984–98.

Joiner, W. M. and Smith, M. A. (2008). Long-term retention explained by a model of short-term learning in the adaptive control of reaching. *Journal of neurophysiology*, 100(5):2948–55.

Jordan, M. I. (1996). Computational aspects of motor control and motor learning. *Handbook of Perception and Action*.

Kadiallah, A., Franklin, D. W., and Burdet, E. (2012). Generalization in adaptation to stable and unstable dynamics. *PloS one*, 7(10):e45075.

Kawato, M. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9(6):718–727.

Kim, S., Ogawa, K., Lv, J., Schweighofer, N., and Imamizu, H. (2015a). Neural substrates related to motor memory with multiple timescales in sensorimotor adaptation. *PLoS biology*, 13(12):e1002312.

Kim, S., Oh, Y., and Schweighofer, N. (2015b). Between-Trial Forgetting Due to Interference and Time in Motor Adaptation. *PloS one*, 10(11):e0142963.

Kitago, T., Ryan, S. L., Mazzoni, P., Krakauer, J. W., and Haith, A. M. (2013). Unlearning versus savings in visuomotor adaptation: comparing effects of washout, passage of time, and removal of errors on motor memory. *Frontiers in human neuroscience*, 7.

Kitazawa, S., Kimura, T., and Uka, T. (1997). Prism adaptation of reaching movements: specificity for the velocity of reaching. *The Journal of neuroscience*, 17:1481–1492.

Kojima, Y., Iwamoto, Y., and Yoshida, K. (2004). Memory of learning facilitates saccadic adaptation in the monkey. *The Journal of neuroscience*, 24(34):7531–9.

Kojima, Y. and Soetedjo, R. (2017). Selective reward affects the rate of saccade adaptation. *Neuroscience*, 355:113–125.

Kording, K. P., Tenenbaum, J. B., and Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature neuroscience*, 10(6):779–86.

Krakauer, J. W., Ghez, C., and Ghilardi, M. F. (2005). Adaptation to visuomotor transformations: consolidation, interference, and forgetting. *The Journal of neuroscience*, 25(2):473–8.

Krakauer, J. W., Ghilardi, M. F., and Ghez, C. (1999). Independent learning of internal models for kinematic and dynamic control of reaching. *Nature neuroscience*, 2(11):1026–31.

Krakauer, J. W. and Mazzoni, P. (2011). Human sensorimotor learning: adaptation, skill, and beyond. *Current opinion in neurobiology*, 21(4):636–44.

Krakauer, J. W., Mazzoni, P., Ghazizadeh, A., Ravindran, R., and Shadmehr, R. (2006). Generalization of motor learning depends on the history of prior action. *PLoS biology*, 4(10):e316.

Krakauer, J. W., Pine, Z. M., Ghilardi, M.-F., and Ghez, C. (2000). Learning of Visuomotor Transformations for Vectorial Planning of Reaching Trajectories. *The journal of neuroscience*, 20(23):8916–8924.

Lee, J.-Y. and Schweighofer, N. (2009). Dual adaptation supports a parallel architecture of motor memory. *The Journal of neuroscience*, 29(33):10396–404.

Lei, Y., Bao, S., Perez, M. A., and Wang, J. (2017). Enhancing generalization of visuomotor adaptation by inducing use-dependent learning. *Neuroscience*, 366:184–195.

Leow, L.-A., De Rugy, A., Marinovic, W., Riek, S., and Carroll, T. J. (2016). Savings for visuomotor adaptation require prior history of error, not prior repetition of successful actions. *Journal of neurophysiology*, 116(4):1603–1614.

Liu, Z. and Weinshall, D. (1999). Mechanisms of generalization in perceptual learning. In *Advances in Neural Information Processing Systems*, pages 45–51.

Malfait, N., Gribble, P. L., and Ostry, D. J. (2005). Generalization of motor learning based on multiple field exposures and local adaptation. *Journal of neurophysiology*, 93:3327–3338.

Malfait, N. and Ostry, D. J. (2004). Is interlimb transfer of force-field adaptation a cognitive response to the sudden introduction of load? *The Journal of neuroscience*, 24(37):8084–9.

Malfait, N., Shiller, D. M., and Ostry, D. J. (2002). Transfer of Motor Learning across Arm Configurations. *The journal of neuroscience*, 22(22):9656–9660.

Marinovic, W., Homan, M., and Carroll, T. J. (2017). Use dependent directional bias does not transfer to the untrained limb during bimanual contractions. *European Journal of Neuroscience*.

Marko, M. K., Haith, A. M., Harran, M. D., and Shadmehr, R. (2012). Sensitivity to prediction error in reach adaptation. *Journal of neurophysiology*, 108(6):1752–63.

Massie, T. H., Salisbury, J. K., et al. (1994). The phantom haptic interface: A device for probing virtual objects. In *Proceedings of the ASME winter annual meeting, symposium on haptic interfaces for virtual environment and teleoperator systems*, volume 55, pages 295–300. Chicago, IL.

Mattar, A. A. G. and Ostry, D. J. (2007). Modifiability of generalization in dynamics learning. *Journal of neurophysiology*, 98(6):3321–9.

Mattar, A. A. G. and Ostry, D. J. (2010). Generalization of dynamics learning across changes in movement amplitude. *Journal of neurophysiology*, 104(1):426–38.

Mawase, F., Shmuelof, L., Bar-Haim, S., and Karniel, A. (2014). Savings in locomotor adaptation explained by changes in learning parameters following initial adaptation. *Journal of neurophysiology*, 111(7):1444–54.

Mawase, F., Uehara, S., Bastian, A. J., and Celnik, P. (2017). Motor learning enhances use-dependent plasticity. *Journal of Neuroscience*, 37(10):2673–2685.

Mazzoni, P. and Krakauer, J. W. (2006). An implicit plan overrides an explicit strategy during visuomotor adaptation. *The Journal of neuroscience*, 26(14):3642–3645.

McDougle, S. D., Bond, K. M., and Taylor, J. A. (2017). Implications of plan-based generalization in sensorimotor adaptation. *Journal of Neurophysiology*, 118(1):383–393.

Miall, R. and Wolpert, D. (1996). Forward Models for Physiological Motor Control. *Neural Networks*, 9(8):1265–1279.

Mistry, M., Mohajerian, P., and Schaal, S. (2005). An exoskeleton robot for human arm movement study. In *Intelligent Robots and Systems, 2005.(IROS 2005). 2005 IEEE/RSJ International Conference on*, pages 4071–4076. IEEE.

Morehead, J. R., Qasim, S. E., Crossley, M. J., and Ivry, R. (2015). Savings upon Re-Aiming in Visuomotor Adaptation. *The Journal of Neuroscience*, 35(42):14386–14396.

Mussa-Ivaldi, F. A., Giszter, S. F., and Bizzi, E. (1994). Linear combinations of primitives in vertebrate motor control. *Proceedings of the National Academy of Sciences of the United States of America*, 91(16):7534–8.

Myers, K. M. and Davis, M. (2002). Behavioral and neural analysis of extinction. *Neuron*, 36(4):567–584.

Nikooyan, A. A. and Ahmed, A. A. (2015). Reward feedback accelerates motor learning. *Journal of Neurophysiology*, 113(2):633–646.

Nozaki, D. and Scott, S. H. (2009). Multi-compartment model can explain partial transfer of learning within the same limb between unimanual and bimanual reaching. *Experimental brain research*, 194(3):451–63.

Oldfield, R. C. (1971). The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia*, 9(1):97–113.

Parmar, P. N., Huang, F. C., and Patton, J. L. (2015). Evidence of multiple coordinate representations during generalization of motor learning. *Experimental brain research*, 233(1):1–13.

Pekny, S. E., Criscimagna-Hemminger, S. E., and Shadmehr, R. (2011). Protection and expression of human motor memories. *The Journal of neuroscience*, 31(39):13829–39.

Pekny, S. E., Izawa, J., and Shadmehr, R. (2015). Reward-dependent modulation of movement variability. *The Journal of neuroscience*, 35(9):4015–24.

Pélisson, D., Alahyane, N., Panouilleres, M., and Tilikete, C. (2010). Sensorimotor adaptation of saccadic eye movements. *Neuroscience & Biobehavioral Reviews*, 34(8):1103–1120.

Politis, D. N., Romano, J. P., and Wolf, M. (1999). *Subsampling*. Springer-Verlag, New York.

Popa, L. S., Hewitt, A. L., and Ebner, T. J. (2012). Predictive and feedback performance errors are signaled in the simple spike discharge of individual purkinje cells. *Journal of Neuroscience*, 32(44):15345–15358.

Quattrocchi, G., Greenwood, R., Rothwell, J. C., Galea, J. M., and Bestmann, S. (2017). Reward and punishment enhance motor adaptation in stroke. *J Neurol Neurosurg Psychiatry*, pages jnnp–2016.

Rabe, K., Livne, O., Gizewski, E. R., Aurich, V., Beck, A., Timmann, D., and Donchin, O. (2009). Adaptation to visuomotor rotation and force field perturbation is correlated to different brain areas in patients with cerebellar degeneration. *Journal of neurophysiology*, 101(4):1961–1971.

Reisman, D. S., Bastian, A. J., and Morton, S. M. (2010). Neurophysiologic and rehabilitation insights from the split-belt and other locomotor adaptation paradigms. *Physical Therapy*, 90(2):187–195.

Rescorla, R. A. (2004). Spontaneous recovery. *Learning & memory (Cold Spring Harbor, N.Y.)*, 11(5):501–9.

Resulaj, A., Kiani, R., Wolpert, D. M., and Shadlen, M. N. (2009). Changes of mind in decision-making. *Nature*, 461(7261):263–266.

Roach, N. W., McGraw, P. V., Whitaker, D. J., and Heron, J. (2017). Generalization of prior information for rapid Bayesian time estimation. *Proceedings of the National Academy of Sciences of the United States of America*, 114(2):412–417.

Robbins, S. J. (1990). Mechanisms underlying spontaneous recovery in autoshaping. *Journal of Experimental Psychology: Animal Behavior Processes*, 16(3):235.

Robinson, F. R., Noto, C. T., and Bevans, S. E. (2003). Effect of visual error size on saccade adaptation in monkey. *Journal of neurophysiology*, 90(2):1235–44.

Roemmich, R. T. and Bastian, A. J. (2015). Two ways to save a newly learned motor pattern. *Journal of neurophysiology*, 113(10):3519–3530.

Sarwary, A. M. E., Stegeman, D. F., Selen, L. P. J., and Medendorp, W. P. (2015). Generalization and transfer of contextual cues in motor learning. *Journal of neurophysiology*, 114(3):1565–76.

Scheidt, R. A., Reinkensmeyer, D. J., Conditt, M. A., Rymer, W. Z., and Mussa-Ivaldi, F. A. (2000). Persistence of Motor Adaptation During Constrained, Multi-Joint, Arm Movements. *Journal of Neurophysiol*, 84(2):853–862.

Schlerf, J. E., Xu, J., Klemfuss, N. M., Griffiths, T. L., and Ivry, R. B. (2013). Individuals with cerebellar degeneration show similar adaptation deficits with large and small visuomotor errors. *Journal of Neurophysiol*, 109:1164–1173.

Selen, L. P. J., Shadlen, M. N., and Wolpert, D. M. (2012). Deliberation in the Motor System: Reflex Gains Track Evolving Evidence Leading to a Decision. *Journal of Neuroscience*, 32(7).

Shadmehr, R. (2004). Generalization as a behavioral window to the neural mechanisms of learning internal models. *Human movement science*, 23(5):543–68.

Shadmehr, R. and Brashers-Krug, T. (1997). Functional Stages in the Formation of Human Long-Term Motor Memory. *The journal of neuroscience*, 17(1):409–419.

Shadmehr, R. and Holcomb, H. H. (1999). Inhibitory control of competing motor memories. *Experimental Brain Research*, 126(2):235–251.

Shadmehr, R. and Moussavi, Z. M. K. (2000). Spatial Generalization from Learning Dynamics of Reaching Movements. *The journal of neuroscience*, 20(20):7807–7815.

Shadmehr, R. and Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *The Journal of neuroscience*, 14(5 Pt 2):3208–3224.

Shadmehr, R., Smith, M. A., and Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual review of neuroscience*, 33:89–108.

Sheahan, H. R., Franklin, D. W., and Wolpert, D. M. (2016). Motor planning, not execution, separates motor memories. *Neuron*, 92(4):773–779.

Shergill, S. S., Bays, P. M., Frith, C. D., and Wolpert, D. M. (2003). Two eyes for an eye: the neuroscience of force escalation. *Science*, 301(5630):187–187.

Sing, G. C., Joiner, W. M., Nanayakkara, T., Brayanov, J. B., and Smith, M. A. (2009). Primitives for motor adaptation reflect correlated neural tuning to position and velocity. *Neuron*, 64(4):575–89.

Sing, G. C. and Smith, M. A. (2010). Reduction in learning rates associated with anterograde interference results from interactions between different timescales in motor adaptation. *PLoS computational biology*, 6(8):14.

Smith, M. A., Ghazizadeh, A., and Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS biology*, 4(6):1035–1043.

Takiyama, K., Hirashima, M., and Nozaki, D. (2015). Prospective errors determine motor learning. *Nature communications*, 6:5925.

Tanaka, H., Krakauer, J. W., and Sejnowski, T. J. (2012). Generalization and multirate models of motor adaptation. *Neural computation*, 24(4):939–66.

Tanaka, H., Sejnowski, T. J., and Krakauer, J. W. (2009). Adaptation to visuomotor rotation through interaction between posterior parietal and motor cortical areas. *Journal of neurophysiology*, 102(5):2921–32.

Thoroughman, K. A. and Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature*, 407(6805):742–747.

Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature neuroscience*, 7(9):907–915.

Tong, C., Wolpert, D. M., and Flanagan, J. R. (2002). Kinematics and dynamics are not represented independently in motor working memory: evidence from an interference study. Technical Report 3, Department of Psychology and Centre for Neuroscience Studies, Queen's University, Kingston, Ontario, K7L 3N6, Canada.

Tukey, J. W. (1949). Comparing individual means in the analysis of variance. *Biometrics*, pages 99–114.

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

Van Den Berg, R., Anandalingam, K., Zylberberg, A., Kiani, R., Shadlen, M. N., and Wolpert, D. M. (2016). A common mechanism underlies changes of mind about decisions and confidence. *Elife*, 5:e12192.

van den Berg, R., Zylberberg, A., Kiani, R., Shadlen, M. N., and Wolpert, D. M. (2016). Confidence Is the Bridge between Multi-stage Decisions. *Current Biology*, 26(23):3157–3168.

VanderPlas, J. (2014). Frequentism and bayesianism: a python-driven primer. *arXiv preprint arXiv:1411.5018*.

Verstynen, T. and Sabes, P. N. (2011). How each movement changes the next: an experimental and theoretical study of fast adaptive priors in reaching. *Journal of Neuroscience*, 31(27):10050–10059.

Voss, A., Nagler, M., and Lerche, V. (2013). Diffusion Models in Experimental Psychology. *Experimental Psychology*, 60(6):385–402.

Wei, K. and Körding, K. (2009). Relevance of error: what drives motor adaptation? *Journal of neurophysiology*, 101(2):655–64.

Wei, K. and Körding, K. (2010). Uncertainty of feedback and state estimation determines the speed of motor adaptation. *Frontiers in computational neuroscience*, 4:11.

Witney, A. G. and Wolpert, D. M. (2003). Spatial representation of predictive motor learning. *Journal of neurophysiology*, 89(4):1837–43.

Wolpert, D. and Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks*, 11(7-8):1317–1329.

Wolpert, D. M., Diedrichsen, J., and Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nature reviews. Neuroscience*, 12(12):739–51.

Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science (New York, N.Y.)*, 269(5232):1880–2.

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

Wu, H. G. and Smith, M. A. (2013). The generalization of visuomotor learning to untrained movements and movement sequences based on movement vector and goal location remapping. *The Journal of neuroscience*, 33(26):10772–89.

Zarahn, E., Weston, G. D., Liang, J., Mazzoni, P., and Krakauer, J. W. (2008). Explaining savings for visuomotor adaptation: linear time-invariant state-space models are not sufficient. *Journal of neurophysiology*, 100(5):2537–48.

Zhou, W., Fitzgerald, J., Colucci-Chang, K., Murthy, K. G., and Joiner, W. M. (2017). The temporal stability of visuomotor adaptation generalization. *Journal of Neurophysiology*, 118(4):2435–2447.