8-2-2022 1:00 PM

Durability of Motor Learning by Observation

Natalia Mangos, *The University of Western Ontario*

Supervisor: Gribble, Paul L., *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Master of Science degree in Neuroscience

Abstract

Recent evidence suggests that neural representations of novel movement dynamics can be acquired by observing someone else experiencing them first-hand. Visual information about another person's movement kinematics can be transformed into an adaptation of feedforward limb control for the observer; however, little is known about the durability of this adaptation. Despite the longevity of changes in the motor system being a defining characteristic of motor learning, studies to date have only examined observation-related effects shortly after observation has occurred, leaving unknown whether such effects are transient phenomena or products of learned, durable changes in neural systems. We measured human participants' force generation patterns before and at various time points (1 minute – 24 hours) after they had either performed or observed movements that were perturbed by novel, robot-generated forces (i.e., a velocity-dependent force field). Like participants who had physically practiced, observers learned to predictively generate directionally- and temporally-specific compensatory forces during reaching. Although retention generally decayed with time, we found no evidence of an interaction between the effects of the passage of time and whether participants had performed or observed reaches in a force field, suggesting that the adaptation decayed similarly regardless of whether it was induced by observing someone else's physical force field learning or feeling the force field for oneself. Notably, the adaptation of predictive limb control induced by observation was still detectable 24 hours later, demonstrating that visually-acquired representations of movement dynamics can be retained, and continue to influence behaviour, long after the initial training period is over. Our results suggest that observing can have lasting effects on the brain that are similar to those seen for physical practice.

Keywords

Summary for Lay Audience

Motor learning involves the brain learning what forces are required for movement. The forces that our muscles need to produce in order for us to move properly will inevitably change with growth, aging, injury, and in different environments. To preserve our ability to move despite these changes, sensory and motor areas of the brain must constantly be adapting. This adaptation is critical for maintaining motor control, but also underlies the learning of new and complex motor skills.

When we learn, the brain makes new connections that can change how it controls our bodies during movement. Although motor learning typically involves extensive physical practice, recent evidence suggests that the brain can also learn what forces are required for movement by observing the movements of others. This is interesting because forces are not things that can be seen. Information about the parts of another person's movements that *can* be seen— like visible errors in where our bodies actually are, compared to where they are supposed to be—must be transformed into new connections in the brain that can then be used in the control of the observer's movements. A defining characteristic of learning is that the changes in the brain that come with it need to be long-lasting. Research suggesting that observing can teach the brain about the forces required for movement has only ever looked at the effects of observation right after observing. As a result, it was not known whether previously described observation-driven changes in the brain were transient, or whether they might form stable memories that could influence motor control well past the end of the observation period. The present thesis addresses this question for the first time. Here we provide evidence that visually-induced changes in the neural circuitry involved in force control can be retained and influence behaviour 24 hours after observing someone else's movement errors. We also find similarities in the stability of observation-related changes and physical practice-related changes over time. Our results suggest that observing can have lasting effects on the brain that are similar to those seen for physical practice.

Acknowledgments

I would first like to thank my supervisor, Paul Gribble, for his unwavering support. He welcomed me into an environment that has allowed me to be curious, make mistakes, teach and mentor others, and explore for exploration's sake. I am grateful to Paul not only for his guidance, but for making space for me to be myself. I would also like to thank the members of my advisory committee—Jessica Grahn, Andrew Pruszynski, and Blake Butler—for contributing their insights and unique perspectives to this work and the work that may follow. For his particularly constructive feedback, and for all he's done to support me, I would like to extend a special thank you to Andrew Pruszynski. For their guidance and friendship, I would like to thank the current and previous members of the Sensorimotor Superlab—especially Chris Forgaard, whose mentorship and direct contributions to this work have been invaluable. Lastly, thank you to my family, for supporting my love for science and embracing that I will probably always be an adult who is wearing a backpack.

Table of Contents

List of Figures

Chapter 1

1 Introduction

1.1 Preface

Motor learning involves a number of different components. For example, when a dancer is first learning and practicing a skill, they might learn about the direction in which their arms should move (spatial information) and with what points in the music their movements should coincide (temporal information). The dancer also learns about lower-level aspects of the skill—the mechanical requirements for movement—such as what patterns and magnitudes of muscle forces are required to generate the correct amounts of torque at particular joints. This latter element of motor learning, which provides the brain with the information it uses during movement execution, is the focus of this thesis.

1.2 Motor Adaptation and Force Field Learning

Skilled action critically depends on the brain's ability to predict what muscle forces are required for movement. Because growth, aging, injury, and changing environmental factors all affect the dynamics of movement, the brain must constantly update its predictions about what forces are needed to move successfully.

When learning a new motor skill, or regaining the ability to perform an existing skill in the face of changes that affect movement, the brain learns new mappings between sensory and motor variables, and so acquires a neural representation of the mechanical requirements for movement. The subsequent use of this representation by the neural systems involved in movement control results in changes in motor output, which improves (or restores) task performance by modifying muscle activation patterns according to predictions about movement dynamics. This process is known as motor adaptation, a form of motor learning.

1.2.1 Force Field Learning

How the brain learns what forces are required for movement has been explored through studies of motor adaptation in humans and other animals. In these studies, subjects must adapt their motor commands to correct for experimentally imposed sensory perturbations that would otherwise cause performance errors during a movement task. A common laboratory model of motor adaptation in humans involves using a robotic device to generate a novel pattern of motion-dependent forces—a *force field*—that perturbs the arm during reaching tasks. In such a paradigm, participants are tasked with performing straight-line reaching movements in the presence of robot-imposed forces that alter reach dynamics and cause movements to deviate from their intended trajectories (Shadmehr and Mussa-Ivaldi 1994). Over successive reaches, the brain learns a representation of the new dynamics, and adapts the motor commands it sends to the upper limb to generate a novel, time-varying pattern of muscle forces (Thoroughman and Shadmehr 1999) that counteracts the force field in a predictive manner (Shadmehr and Mussa-Ivaldi 1994). This adaptation of participants' force output results in the gradual return of reaching movements back to their intended trajectories, despite the continued application of perturbing forces by the robot. Upon removal of the force field, participants display *after-effects* (i.e., large trajectory curvature in the direction opposite the force field, driven by the predictive generation of compensatory forces by the participant when there is no longer a force field to compensate for). With continued reaching in the absence of any perturbing forces, after-effects wash out and baseline levels of force output are restored (Shadmehr and Mussa-Ivaldi 1994).

1.2.2 Durability of Force Field Learning

The effects of force field adaptation on the human motor system are not transient. A defining feature of motor learning is that it requires lasting change in the mechanisms underlying skilled action. Learning, therefore, must involve change that persists beyond the period of training (Schmidt and Lee 2014). Adaptive changes in predictive force output patterns (i.e., the source of after-effects) have been shown to be retained for at least 24 hours following initial force field adaptation (Joiner and Smith 2008). Force field adaptation also produces savings (Mathew et al. 2021; Shadmehr and Brashers-Krug

1997)—a phenomenon in which re-learning occurs at a faster rate than initial learning, even if after-effects of the initial learning have been washed out. Force field adaptation, therefore, gives rise to durable changes in the mechanisms underlying skilled action, which improves future performance on the same task, as well as on previously untrained but similar tasks (Conditt et al. 1997; Gandolfo et al. 1996; Goodbody and Wolpert 1998; Huang and Shadmehr 2007; Hwang et al. 2003; Malfait et al. 2002; Sainburg et al. 1999; Shadmehr and Moussavi 2000). In this way, force field adaptation is a form of temporally stable learning in humans.

1.3   Force Field Learning by Observing

In the type of force field learning paradigm discussed thus far, the brain learns what forces are required for movement through extensive physical practice, wherein changes in cortical representations of movement dynamics are thought to be driven by discrepancies between the observed and predicted sensory consequences of motor commands (Miall and Wolpert 1996; Shadmehr et al. 2010). But recent evidence suggests that motor adaptation can also occur in human observers who watch as someone else undergoes force field adaptation. That is, the brain can learn about what forces are required for movement without motor commands or sensory feedback associated with one's own movement. Instead, adaptation of the neural systems involved in limb control is driven by visual input.

In a pioneering study of force field learning by observing, Mattar and Gribble (2005) provided the first clues that novel patterns of muscle forces could be learned through visual observation. They tested performance in a force field after participants had observed one of three videos—each of which depicted another individual using a robotic device to reach in a novel mechanical environment. Observers had watched as the individual adapted to a force field that was either (1) the same as, or (2) different (opposite direction) from, the one they would later be tested in; or, they watched a control video (3) in which the individual experienced randomly varying forces (an unlearnable field; Takahashi et al. 2001). Mattar and Gribble found that relative to group 3, which did not observe learning, participants who observed another individual learning to reach in a

force field performed better (or worse) when subsequently tested in the same (or different) force field. They also found that this effect did not depend on the use of conscious, explicit strategies for movement, but rather, on the implicit engagement of the motor system. They proposed that, like physically reaching in a force field, observing had taught the brain a new representation of reach dynamics that was used to predictively generate a novel, time-varying pattern of muscle forces that mirrored the force field. More simply put: Mattar and Gribble reasoned that observers had learned something about the forces required to move in the field they observed, and that this learning was transformed into an adaptation of the forces the observers generated as soon as they began reaching, thereby improving performance when facing the force field to which their muscle activation patterns had adapted, and impairing performance when facing a force field that was opposite the one they had adapted to.

This idea is consistent with the findings of a later study conducted by Wanda and colleagues (2013), in which force generation patterns after observation were measured directly. Rather than probing adaptation by testing observers' performance in a force field—which makes it difficult to assess the extent to which observers learned the novel dynamics through observation alone (i.e., without entering the perturbing environment)—Wanda *et al.* used an error clamp assay to determine whether observing force field learning produced dynamic after-effects. An error clamp (also called a 'force channel') refers to a type of mechanical environment that constrains the hand path to a straight line and produces a direct readout of the forces generated during reaching (Scheidt et al. 2000). Thus, by having participants perform mechanically clamped reaches before and after observing force field learning, one can obtain direct measurements of the forces generated by observers who have never moved in a force field, and so can quantify any adaptation of predictive force output that results solely from observing. With this approach, Wanda and colleagues confirmed that observing force field learning induced an adaptation of predictive force output in observers, and that, like with physical practice, this adaptation mirrored the timing and direction of the force field. Previous work has therefore provided convincing evidence that kinematic information obtained through observing force field learning can be transformed into a novel adaptation of reach dynamics for the observer.

1.3.1 Durability of Force Field Learning by Observing

Both observing force field learning and physically practicing in a force field can elicit temporally- and directionally-specific changes in predictive limb control; however, while force field adaptation by physical practice has been shown to have lasting effects on the human motor system (see section 1.2.2 above), no study to date has assessed the temporal stability of observation-induced adaptation. In fact, to our knowledge, there is no evidence to suggest that observing force field learning has any effect on the brain outside of a narrow time window (~1 hour) following the observation period. Indeed, most data on the effects of observing force field learning (Brown et al. 2009, 2010; Malfait et al. 2010; Mattar and Gribble 2005; McGregor et al. 2018b, 2018a; McGregor and Gribble 2017; Williams and Gribble 2012) were obtained during or only minutes after observation had occurred, except for two studies in which performance in a force field (McGregor and Gribble 2015) and on a somatosensory perception task (Bernardi et al. 2013) were assessed approximately one hour after observing. Thus, while force field adaptation by physical practice is known to elicit durable effects on the brain, the durability of the effects of observation remains to be characterized.

Investigating the durability of the adaptation that occurs through observing force field learning is primarily useful for two reasons. First, it would allow us to draw comparisons to the temporal stability of the learning that occurs through physically reaching in a force field. There is a body of literature examining neural correlates of motor learning through physical practice, including for force field learning (Diedrichsen et al. 2005; Li et al. 2001; Maschke et al. 2004; Shadmehr and Holcomb 1997; Smith and Shadmehr 2005). While there is evidence that learning movement dynamics through observation involves some of the same cortical regions as learning through physical practice (Brown et al. 2009; McGregor et al. 2016), the mechanisms underlying motor learning by observing remain largely unknown. If learning that occurs through observation and learning that occurs through physical practice were found to be similarly durable under the same conditions, this would be consistent with the idea that these two processes share common neural mechanisms. Second, a tool that can elicit temporally stable changes in the neural control of movement without physical practice could be useful in a clinical setting,

particularly for promoting the recovery of motor function after stroke. Accordingly, information about the durability of the learning that occurs through observation is relevant to the design of observation-related approaches to neurorehabilitation.

1.4   Present Study

The purpose of the present study was to investigate the durability of the learning that occurs through observing force field learning. We sought to address the overarching question: Is the adaptation of predictive limb control that follows observation a transient phenomenon, or a product of durable, learned changes in the neural systems controlling movement? Given that the latter is true in the case of physical force field learning, a comparison of the effects of observing to those of physical practice can provide insight into this question. The present study tests the hypothesis that adaptation driven by observing force field learning, and adaptation driven by physically reaching in a force field, are similarly durable.

As in Wanda *et al.* (2013), we obtained direct measurements of human participants' force generation patterns at two points in the experiment: before and after they had either observed force field learning or physically undergone force field learning. Following the period of observational or physical force field exposure, participants experienced a temporal delay lasting between one minute and 24 hours before the second set of force output measurements were obtained. We probed adaptation as changes in predictive force generation patterns from before to after observational or physical force field exposure and the variable-length delay.

In support of its purpose—to characterize the durability of the learning that occurs through observing force field learning—the objectives of this study were as follows:

> (1) Generate a video of an individual undergoing force field learning, and verify that the video induces in human observers an adaptation of predictive force output that is detectable immediately following observation (i.e., reproduce Wanda and colleagues' (2013) main finding).

(2) Determine whether observation-induced changes in predictive limb control are still detectable beyond the observation period (i.e., after 24 hours), as has been demonstrated for physical force field learning (see section 1.2.2).

(3) Compare, through quantifying adaptation at various time points after force field reaching or observing, the temporal stability of changes in limb control induced by each of the two types of force field exposure.

(4) Assess the resistance of observation- or physical practice-induced adaptation to washout, and determine if the rate of washout changes with increasing time since observing or physically practicing. After-effects are a hallmark of force field adaptation (see section 1.2.2 above); however, they wash out rapidly upon entering a conventional perturbation-absent (null) environment. In experiments where participants perform reaches in a force field, the rate at which after-effects are extinguished following the force field's removal can be sharply reduced by the presence of an error clamp during the washout period (Scheidt et al. 2000). In the present study, all reaches completed following observation or force field reaching were mechanically clamped, extending the washout phase. Slower washout allows us to better characterize the decay of the adaptation across a larger number of trials, which enables comparisons of the rate at which observation- or physical practice-induced adaptation is extinguished under the same conditions.

Chapter 2

2   Methods
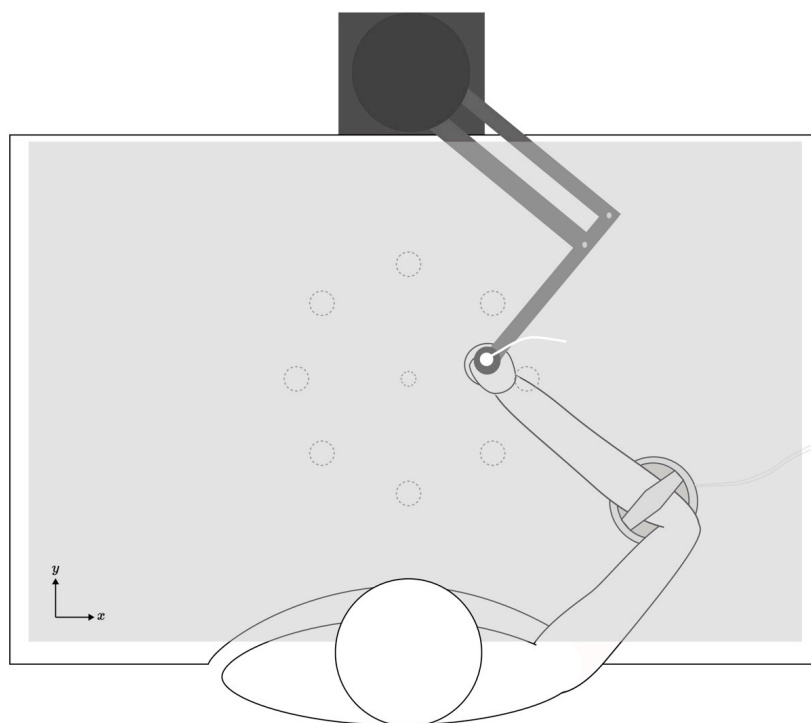
2.1   Participants

One hundred seventy-eight healthy individuals participated in this study. Eligibility requirements included right-handedness; no prior exposure to a force field learning paradigm; and no visual, neurological, or musculoskeletal disorders. Eighteen participants were excluded for failure to complete the testing session, failure to meet the originally stated eligibility criteria (disclosed after completing the testing session), failure to follow instructions during the testing session, or disruption of the testing session by technical difficulties (e.g., a piece of equipment broke). In each case, participants were recorded as needing to be excluded during or immediately following their testing sessions. New participants were recruited until each experimental group contained 16 participants (after exclusions), for a total of 160 participants (mean age: 20 ± 4 years). All study procedures were approved by Western University's Nonmedical Research Ethics Board.

2.2   Experimental Setup

With their right arm, participants used the handle of a two-joint robotic manipulandum (KINARM Laboratories, Kingston, Canada) to perform planar reaching movements to eight visual targets. A custom air-sled was placed beneath the right arm to prevent it from getting tired during the experiment. Targets were displayed on a semisilvered mirror that was mounted horizontally between eye-level and the workspace. This surface occluded participants' direct vision of the handle and right arm but showed the position of the hand as a circular, white cursor, which dragged behind it a transient tracing of the hand's recent trajectory (Figure 1). Targets were placed equally around the circumference of a circle with a 10 cm radius and a 'home' target at its center. On a given trial, participants moved the robot handle to the 'home' target until prompted, by the appearance of one of the eight possible movement targets, to perform a straight-line reaching movement to that target. Shortly after the end-target was reached, the robot moved the participant's passive

arm so that the handle was placed back at the home position. Targets appeared in random order within bins of eight trials, such that each target appeared once per bin.

To regulate reach velocity, participants received colour feedback after each reach to indicate whether the reach was completed too slowly (end-target turned blue), too quickly (end-target turned red), or within the desired time window of 400 ± 50 milliseconds (end-target disappeared rather than changing colour). These time constraints did not include reaction time; participants were told that they could wait as long as they wanted before starting their reach, and that the colour feedback would reflect the amount of time elapsed between leaving the 'home' target and reaching the end-target. Participants were instructed to reach to the end-target in as straight of a line as possible and to do so within the correct time constraints (i.e., such that the end-target would disappear rather than changing colour), to stop on the end-target rather than reaching past or through it, and to wait in the end-target until the robot initiated a return back to the 'home' target.

Figure 1. Experimental setup. Participants were seated in front of a two-joint robotic manipulandum. A custom air-sled was placed on top of the table to support the arm. Direct vision of the arm was completely occluded by the opaque, horizontal display (depicted here as translucent) onto which targets were projected.

2.3   Task Design

The experiment was divided into four blocks (Figure 2A). Participants were assigned to one of 10 experimental groups (*N*=16 participants per group). Each group designation referred to which of two possible force field learning protocols (Movement or Observation) and which of five possible delay period protocols (1-minute, 10-minute, 30-minute, 60-minute, or 24-hour delay) a participant completed.
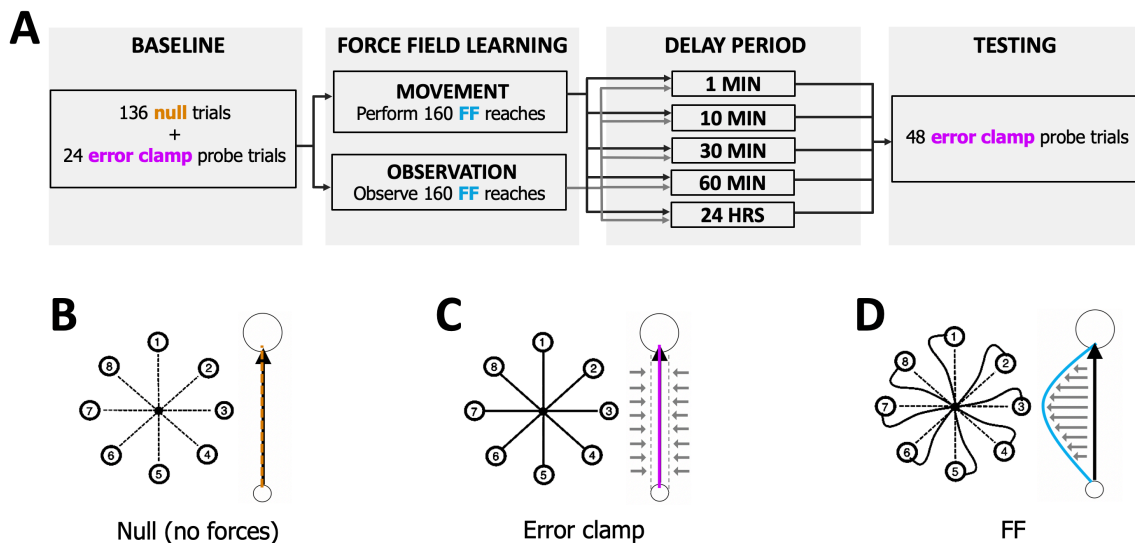
Figure 2. Task design. *A.* Experimental blocks. Ten groups of participants completed a Baseline block, Force Field (FF) Learning block, Delay Period block, and Testing block. All participants completed the same Baseline and Testing blocks; which FF Learning protocol (Movement or Observation) and Delay Period protocol (1-minute, 10-minute, 30-minute, 60-minute, or 24-hour delay) participants completed was manipulated. Force generation patterns were probed during error clamp trials in Baseline and Testing blocks. *B–D.* Hand paths during null trials, error clamp trials, and early FF trials—for reaches to all eight targets (left) or for one reach to an individual target (right). The large black arrow indicates the intended reach direction; small grey arrows depict robot-imposed forces.

Baseline Block. Participants began by performing 160 reaches in the absence of any opposing forces applied by the robot ('null' environment; Figure 2B). The latter half of trials in this block included 24 pseudorandomly interspersed error clamp trials (Scheidt et al. 2000)—probe trials in which 'walls' of force imposed by the robot restricted, or clamped, the position of the handle to a straight-line path between the home position and the end-target (Figure 2C). Such trials allowed us to quantify participants' force output over the time course of a reach, prior to any force field exposure.

Force Field Learning Block. Following the Baseline block, participants completed one of two possible force field learning protocols. Participants assigned to a Movement group performed 160 reaches in a counter-clockwise velocity-dependent force field environment. The force field was applied by the robot in accordance with the following equation,

(1)
$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} 0 & -k \\ k & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix}$$

where F_x and F_y are the commanded forces, v_x and v_y are hand velocities in the x- and y-axes (left-right and front-back) respectively, and $k=14$ Ns/m (Figure 2D). Participants assigned to an Observation group did not perform reaches in a force field environment; rather, they watched a video of a tutor performing 160 reaches in the counter-clockwise force field environment given by Equation 1. The video depicted a top-down view of the tutor's right arm as they learned to make straight movements in the presence of perturbing forces (Figure 3).
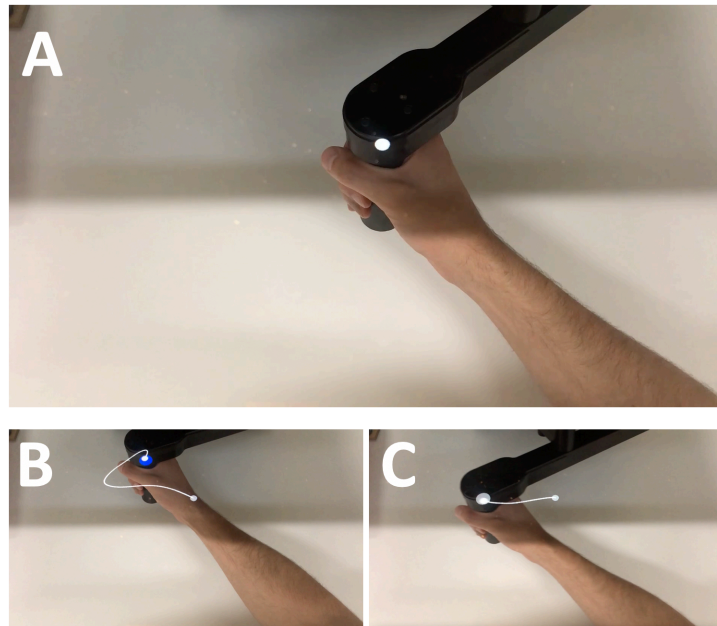
Figure 3. Sample frames from the video used for Observation groups. *A.* Tutor resting at the 'home' position. *B.* Reach completed during the early stages of force field adaptation. The tutor experiences reaching errors (large trajectory curvature). *C.* Reach completed during the late stages of force field adaptation. The force field causes minimal error (trajectory curvature).

Participants assigned to an Observation group remained seated in front of the manipulandum while the video was projected onto their horizontal display. Participants were not informed of the presence of robot-generated forces in the video they observed. They were only told that the video showed "someone doing reaches", and were instructed to keep count of the number of times the actor in the video reached the end-target in the desired time window (i.e., the number of times the end-target disappeared rather than changing colour). To verify that participants were paying attention during the video, they were asked to report their counts (verbally, to be documented by the researcher) at each of four checkpoints throughout the video. Any participant who reported a number not within 20% of the correct value at any of the four checkpoints was excluded from data analysis. Only one participant—who disclosed that they had fallen asleep during the video—reported a count that fell outside of this range and was dismissed with full compensation after the video had finished playing. Because moving while observing force field learning has been shown to interfere with the effects of observing (Mattar and Gribble 2005), participants were instructed to keep their hands resting flat on the table beneath the display, and to remain as still as possible, throughout the duration of the video.

Delay Period Block. After the Force Field Learning block, participants completed one of five possible delay period protocols, in which they took a 1-minute, 10-minute, 30-minute, 60-minute, or 24-hour break from using the robot. During delay periods ≤ 60 minutes, participants remained seated in the robot chair and kept their hands resting flat on the table. Because we were unsure of whether moving during the delay period would retroactively interfere with the effects of learning by observing, participants were instructed to remain as still as possible throughout the duration of the delay period. The researcher was present to monitor participants' motion, and any movement observed by the researcher was accompanied by a verbal reminder that participants should restrict their movement. Participants were permitted to listen to a podcast or speak with the researcher during the delay period, and throughout the delay the horizontal display showed only a solid black screen.

Participants assigned to a group with a 24-hour delay left the testing facility immediately following completion of the Force Field Learning block and returned 24 ± 2 hours later.

Testing Block. Following the Delay block, participants were instructed to perform a final set of reaches with the same rules as in the Baseline block. The Testing block consisted of 48 error clamp trials, which allowed us to obtain a second set of measurements of participants' force output during reaching.

2.4   Analyses

The force field used in this experiment perturbed the arm laterally relative to the reach direction (Figure 2). Adaptation could therefore be measured as the change in lateral forces generated by participants during reaching. For each error clamp trial in the Baseline and Testing blocks, the mechanical clamp that restricted the hand's trajectory to a straight line was generated by the application of robot-generated forces that mirrored the lateral forces generated by the participant. Accordingly, measurements of participants' lateral force output over the time course of a given reach were obtained by taking the sign-flipped time series of robot-generated clamp forces. All kinematic and force data were digitally sampled at 1000 Hz and low-pass filtered (10 Hz, double-pass, third-order Butterworth filter). Time-series force and velocity data were aligned on peak hand speed and sliced from 400 ms before to 400 ms after the time of the peak. Although the desired reach time was approximately 400 ms, and the peak velocity ought to coincide with the spatial and temporal midpoint of the reach, we sliced with a margin of 200 ms on either side of the optimal reach time to ensure all data from all reaches would be included. Reaches were only excluded in the infrequent event that they were jerky rather than smooth—that is, the participant sped up and slowed down multiple times between leaving the 'home' target and reaching the end-target, as determined by the presence of two or more acceleration phases separated by a deceleration phase.

Adapted and ideal force output profiles. For each participant, sliced lateral force output data from error clamp trials in the Baseline block were collapsed across same-target trials to generate one average baseline force output profile per target. Then, for any single trial in the Testing block, an *adapted lateral force output profile* was generated by time-series

subtraction of the target-matched baseline profile from the sliced lateral force output profile for that Testing block trial. Adapted lateral force output profiles therefore represented the change, from pre- to post-force field exposure, in the magnitude and temporal pattern of lateral forces produced by participants during reaching. For each trial in the Testing block, we also generated an *ideal lateral force output profile*, which gave the magnitude and temporal pattern of lateral forces that would have been required to perfectly oppose the force field if the force field had been applied during that trial. The ideal profile for a given reach was computed according to Equation 1, taking as input the instantaneous velocity of a participant's hand over the time course of the reach. Because the force field was designed to push the arm counter-clockwise relative to the direction of the reach, compensatory forces produced by participants were those exerted in the clockwise direction. Accordingly, for simplicity we re-signed all lateral force data to make the sign convention for clockwise lateral forces positive.

Adaptation metric. In previous studies, force field adaptation has traditionally been quantified as the reduction in peak lateral error (i.e., maximum perpendicular deviation from the straight-line path during a reach) over successive reaches in a force field—a proxy for measuring the compensatory changes in force output that underlie kinematic performance gains. In our study, which measured lateral forces (using error clamp probe trials) rather than kinematics, we used an analogous metric: peak adapted lateral force (i.e., the maximum change in lateral force output from baseline on a given trial in the Testing block). For each trial in the Testing block, we quantified the amount of adaptation as the peak adapted lateral force divided by peak ideal lateral force. The fraction of ideal force exerted provided a simple adaptation metric that reflected how much of the required compensatory force participants had learned to produce after force field exposure. Peak force values for any given trial were obtained from within a 400 ms epoch centered on peak velocity; however, retrieving peak values from the entire adapted or ideal profiles produced very similar results as those reported here.

Washout rate. Linear regression of the fraction of ideal force exerted (peak adapted / peak ideal) onto trial number was used to determine the rate at which the adaptation washed out over successive Testing block trials. A washout rate for each participant was

given by the slope of the line that best fit (in a least-squares sense) the fraction of ideal force exerted per trial over 48 trials.

Velocity dependence. Because during force field reaching the magnitude of robot-imposed perturbing forces scaled with reach velocity, we also assessed the extent to which participants' adapted force output exhibited velocity dependence. We fit the following linear model to individual participant data across the first eight trials (one reach to each target) in the Testing block:

$$(2) \qquad F_{ADAPTED}(t) = \beta_0 + \beta_1 v(t) + \beta_2 T + \varepsilon(t)$$

where $F_{ADAPTED}(t)$ denotes adapted lateral force output over the time course of a reach, $v(t)$ denotes hand velocity over the time course of a reach, and T denotes trial number. Trial number was included as a covariate to account for any potential trial-to-trial washout. For each participant, we solved for the regression coefficients that would best fit (in a least-squares sense) the model to the data. The β_1 coefficient, the velocity-dependent gain, was used to determine the sensitivity of each participant's adapted force output to reach velocity. We then normalized the modeled velocity-dependent gain by the actual velocity-dependent gain ($k = 14$ Ns/m) to compute a metric of learned velocity dependence, where perfect velocity dependence would be reflected by $\beta_1 / k = 1$, and no velocity dependence would be reflected by $\beta_1 / k = 0$.

Statistical analyses. Adaptation metrics, washout rates, and peak velocities were compared between groups by two-way analysis of variance (ANOVA), with type of force field exposure (Movement or Observation), and delay length (1 min 10 min, 30 min, 60 min, or 24 h), as factors. ANOVA and Bonferroni-corrected post-hoc tests were carried out in JASP v.0.16.2. Regressions and related statistical analyses were carried out in MATLAB R2021b. Modeled β_1 parameters for each participant, and mean learned velocity dependence parameters (β_1 / k) for each group, were compared to zero using multiple one-sample, two-tailed t tests with Bonferroni corrections.

Chapter 3

3 Results

3.1 Adaptation of Force Output After Force Field Exposure

Participants were assigned to one of 10 experimental groups, each involving one of two types of force field exposure (Movement or Observation) followed by one of five possible delay lengths (1 min 10 min, 30 min, 60 min, or 24 h). For each participant, the adapted force profile for any trial in the Testing block reflected the change in their lateral force output from their baseline measurements (see Methods). Adapted force profiles for each group, averaged over the first eight reaches in the Testing block (one to each target) across participants, are shown in Figure 4A. Positive values represent a change from baseline in the compensatory (clockwise) direction.

To determine whether there were group-level differences in movement speed over the trials depicted in Figure 4A, differences in mean peak velocities across the first eight trials were assessed using two-way ANOVA [$F(4,150)=1.66$, $P=0.163$ for main effect of delay length; $F(1,150)=5.92$, $P=0.016$ for main effect of type of force field exposure; $F(4,150)=3.27$, $P=0.013$ for interaction]. Bonferroni-corrected post-hoc tests returned non-significant ($P>0.05$) results for all but one pairwise comparison, which found mean peak velocity to be significantly higher for the Movement group with the 10-minute delay than for the Observation group with the 10-minute delay ($P=0.007$; Figure 4B). To visualize the adaptation of participants' force output in a way that would account for group-level differences in movement speed, we generated velocity-normalized adapted force profiles by dividing the profiles in Figure 4A by the average peak velocity for each respective group. Normalized profiles are shown in Figure 4C.

We used the (non-normalized) adapted and ideal force data for each trial to calculate the fraction of the peak ideal force participants had learned to produce after observing or physically reaching in a force field. This fraction served as our adaptation metric, and accounted for variation in velocity through its use of peak ideal force (a product of velocity; Equation 1) as the denominator. Across the first eight trials in the Testing block

(Figure 4D), the peak adapted force generated one minute after physically reaching in the force field averaged ~75% of the peak ideal force, while observers' adapted force output after the same delay averaged ~30% of the ideal (~40% of the adaptation induced by physical reaching). The mean fraction of peak ideal force exerted across the first eight Testing block trials was significantly different from zero for all combinations of delay length and type of force field exposure (Bonferroni-corrected P<0.001 for all 10 groups), showing that adaptation was still detectable after all delay lengths for both observers and participants who had physically reached in the force field. The effects of type of force field exposure and delay length on the fraction of ideal force generated were assessed by two-way ANOVA. We found a significant main effect of type of force field exposure [$F(1,150)$=430.10, P=6.6x10^{-46}; Figure 4E] and a significant main effect of delay length [$F(4,150)$=8.18, P=5.4x10^{-6}] on the fraction of ideal force exerted, with no significant interaction between the two factors [$F(4,150)$=2.04, P=0.092], suggesting that the effect of delay on the fraction of ideal lateral force exerted was not different for participants who observed versus physically practiced in the force field. The average fraction of ideal force generated by observers consistently fell between 35–45% that of the Movement group with the same delay (Figure 4D, E). Bonferroni-corrected post-hoc tests found that, averaged across both types of force field exposure, participants produced a significantly lower fraction of the ideal force after 24 hours than after 1-, 10-, or 30-minute delays (P=6.5x10^{-5}, P=3.7x10^{-6}, P=0.001, respectively); however, there was no significant difference (P=0.103) between the fraction of ideal force generated after 60-minute and 24-hour delays (Figure 4F).
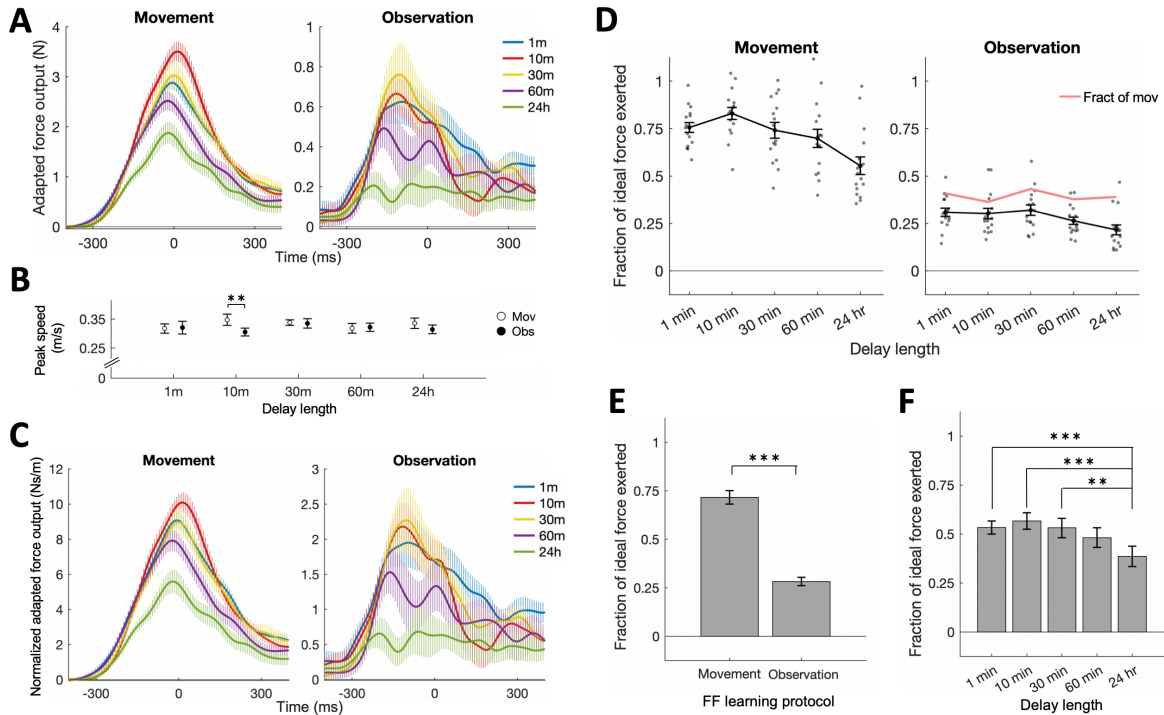
Figure 4. Adaptation of lateral force output after force field exposure. *A.* Average (± SE) adapted lateral force output profiles for Movement (left) and Observation (right) groups with different delay lengths between force field exposure and the Testing block. Profiles represent the difference in Testing-block lateral force output from baseline lateral force output, averaged over the first eight reaches (one to each target) in the Testing block. Positive values represent force output in the clockwise (compensatory) direction. *B.* Mean (± 95% CI) peak reach velocity over the first eight trials in the Testing block, for Movement (open circle) and Observation (closed circle) groups with different delays. No group's peak velocity differed significantly from that of any other group, except for the Movement group with the 10-minute delay with respect to the Observation group with the 10-minute delay. *C.* Velocity-normalized adapted force output profiles, computed by dividing the profiles shown in 4A by each group's mean peak velocity over the first eight reaches in the Testing block. *D.* Amount of adaptation, quantified as the fraction of peak ideal lateral force exerted, averaged across the first eight Testing block trials. Values equal to 1 reflect perfect adaptation and values equal to 0 reflect no adaptation. Means (± SE) are shown for Movement (left) and Observation (right) groups with different delays. Closed circles are individual participant means. The pink line represents the fraction of the delay-matched Movement group's adaptation that was achieved after Observation. *E.* Mean (± 95% CI) amount of adaptation across delay lengths, for Movement and Observation groups. *F.* Amount of adaptation, averaged (± 95% CI) across types of force field exposure, by delay length. **$P<0.01$, ***$P<0.001$

3.2   Washout of Adaptation

Next, we assessed the effect of delay length on the rate at which reaching- or observation-related adaptation washed out during error clamp probe trials after the delay period. The mean fraction of ideal lateral force exerted per trial in the Testing block is shown for each of the 10 experimental groups in Figure 5A.

For each individual participant, the fraction of ideal lateral force generated across the 48 Testing block trials was modeled as a linear function of trial number, where the rate of washout was given by the slope of the least-squares line of best fit. Average washout rates were compared between groups by two-way ANOVA, with type of force field exposure (Movement, Observation) and delay length (1 min, 10 min, 30 min, 60 min, 24 h) as factors. There were significant main effects of type of force field exposure [$F(1,150)=138.25$, $P=4.9 \times 10^{-23}$] and delay length [$F(4,150)=6.46$, $P=7.9 \times 10^{-5}$] on washout rate; however, these effects were qualified by a significant interaction [$F(4,150)=5.24$, $P=5.6 \times 10^{-4}$]. A simple main effects analysis showed that delay length affected washout rate after physical force field exposure ($P=1.8 \times 10^{-5}$), but not after observation ($P=0.252$). Bonferroni-corrected post-hoc tests were conducted for all possible pairwise comparisons except for pairs of Observation groups with different delay lengths. Adaptation related to physical force field exposure was found to have washed out significantly faster after 1- or 10-minute delays than after 24 hours ($P=0.038$, $P=2 \times 10^{-5}$ respectively); however, there were no significant differences in washout rates at 30 minutes, 60 minutes, and 24 hours after physical force field exposure ($P>0.05$ for all pairwise comparisons). Among post-hoc comparisons of washout rate for Movement and Observation groups with the same delay length, adaptation was found to have washed out significantly faster ($P<0.001$) for Movement than Observation groups at all delays except 24 hours ($P=1.000$), at which point washout rates for Movement and Observation groups were not statistically different (Figure 5B). Additionally, the washout rate for the Movement group with 24-hour delay did not differ significantly from the washout rate of any Observation group ($P>0.05$ for all pairwise comparisons). Washout rates for all other Movement groups (i.e., delays ≤ 60 min) were significantly different from all Observation groups ($P<0.05$ for all pairwise comparisons).
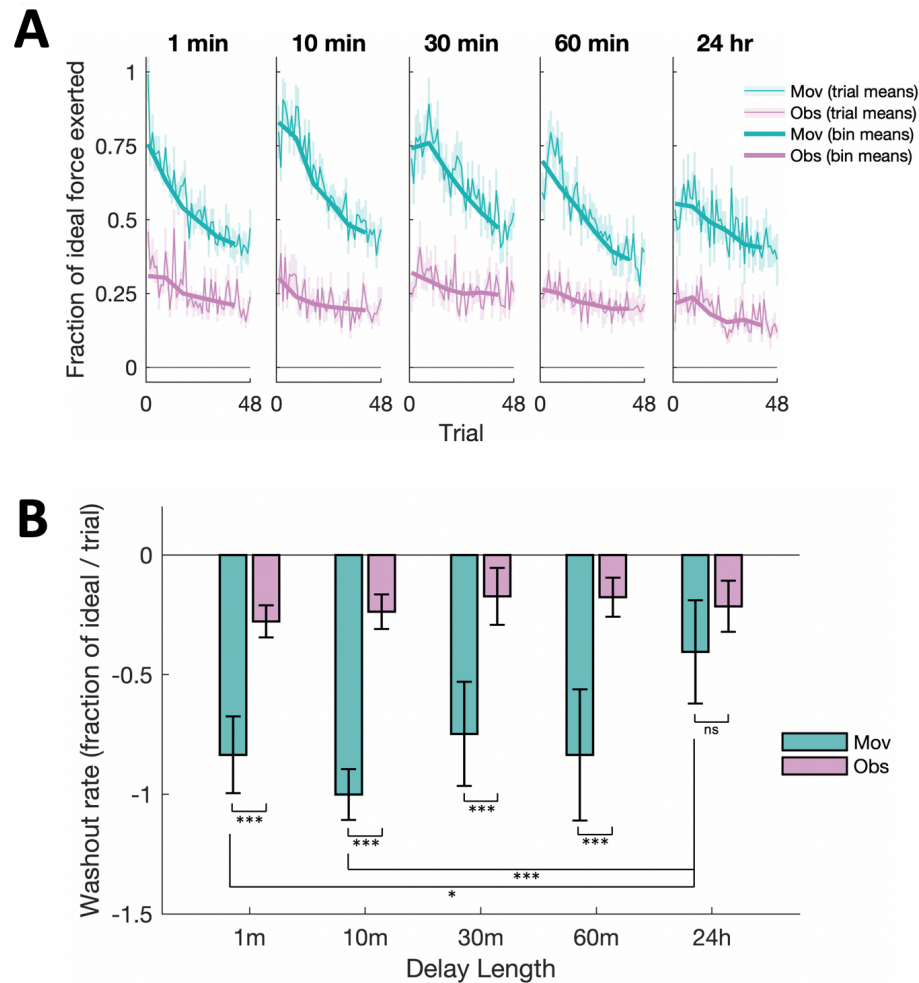
Figure 5. Washout of adaptation. *A.* Fraction of ideal force exerted across all 48 Testing block trials, for each group. Thin lines and surrounding shading represent single trial means ± SE across participants in a given group; thick lines represent means for bins of eight trials (one reach to each target per bin). *B.* Mean (± 95% CI) washout rates for each group. Washout rates were computed for individual participants as the slope of the least-squares line of best fit to the fraction of ideal force exerted across the 48 Testing block trials. *$P<0.05$, ***$P<0.001$, ns $P>0.05$.

3.3   Velocity Dependence

We next determined whether or not individual participants' adapted lateral force output exhibited dependence on reach velocity. Fitting the linear model in Equation 2 to the adapted force and velocity time-series data from the first eight Testing block trials (one reach to each target), we estimated the β_1 coefficient for each participant and compared its value to zero using two-tailed Bonferroni-corrected t tests. A β_1 value significantly larger than zero indicated that the adapted forces generated in the compensatory direction scaled with instantaneous velocity. β_1 was positive and significantly different from zero for all 80 participants who physically performed reaches in the force field, regardless of delay length ($P<0.05$ for all participants). Of the 80 participants who observed rather than performed perturbed reaches, 61 had a β_1 parameter that was positive and significantly different from zero ($P<0.05$ for all), 10 had a β_1 parameter that was negative and significantly different from zero ($P<0.05$ for all), and the remaining nine had a β_1 parameter that did not significantly differ from zero ($P>0.05$ for all). The number of participants in each group who learned appropriate velocity dependence ($\beta_1>0$), inappropriate velocity-dependence ($\beta_1<0$), or no velocity dependence ($\beta_1=0$) is shown in Figure 6A, accompanied by adapted force output profiles (averaged over the first eight trials) from sample observers whose β_1 parameter fell into each category (Figure 6B).

We normalized modeled β_1 parameters by the real velocity-dependent gain of the force field, k, to generate a metric of learned velocity dependence: β_1 / k. The mean ($\pm$ 95% CI) learned velocity dependence for each group, with or without participants who did not learn any velocity dependence, is shown in Figure 6C. Including all participants' parameters, the mean learned velocity dependence was significantly different from zero for all groups (multiple two-tailed t tests, all Bonferroni-corrected $P<0.05$) except the Observation group with a 24-hour delay ($P=0.922$).
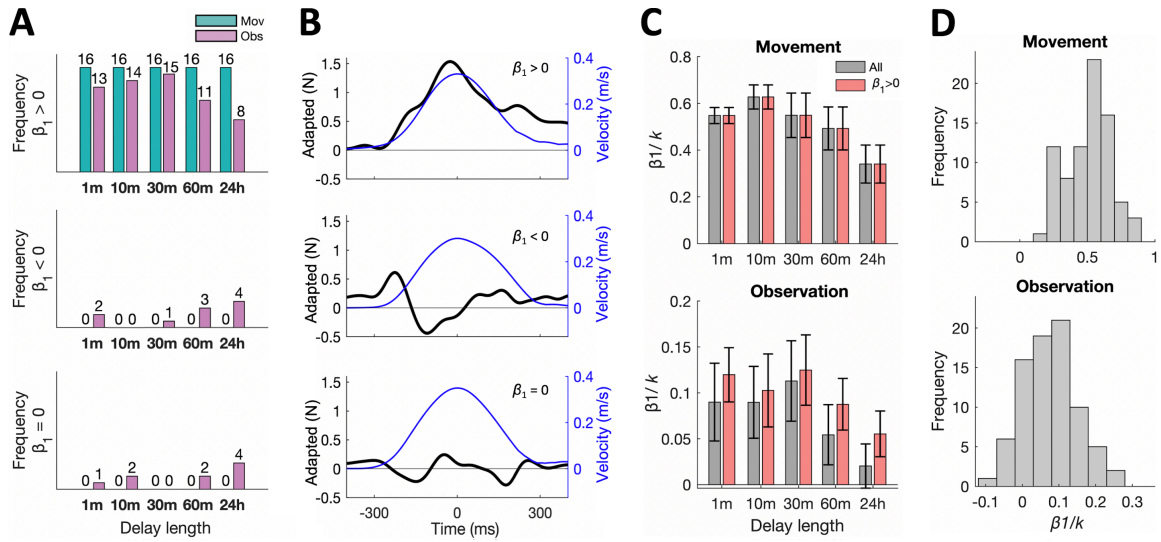
Figure 6. Velocity dependence of adapted force output. *A.* Number of participants in each group (*N*=16/group) with a modeled β_1 parameter significantly greater than zero (top), significantly less than zero (middle), or not significantly different from zero (bottom). These categories separate participants based on whether they learned appropriate velocity dependence, inappropriate velocity dependence, or no velocity dependence, respectively. Differences from zero were evaluated by Bonferroni-corrected two-tailed *t* tests (significant if adjusted *P*<0.05). *B.* Adapted (i.e., baseline-subtracted) force output and velocity data, averaged over the first eight reaches in the Testing block, from a sample participant whose modeled β_1 parameter was significantly greater than zero (top), a sample participant whose modeled β_1 parameter was significantly less than zero (middle), and a sample participant whose modeled β_1 parameter was not significantly different from zero (bottom). All three sample participants were in Observation conditions. *C.* Mean normalized velocity dependence parameter (β_1 / k) across all members of a group (grey) or across members whose adapted force output exhibited appropriate velocity dependence (β_1 significantly greater than zero; pink). Error bars represent 95% confidence intervals. *D.* Density of normalized velocity dependence parameters (β_1 / k) across participants in Movement (top) or Observation (bottom) groups (*N*=80/group).

Chapter 4

4   Discussion

The present study sought to characterize the durability of the learning that occurs through observing another individual undergoing force field learning. Previous studies have reported changes in motor behaviour, somatosensory perception, and functional connectivity in the brain, shortly after human participants observed another individual undergoing force field learning. But this body of work left unknown whether the effects of observing are transient, or whether they might reflect learned, durable changes in neural systems. We addressed this question in an experiment that probed adaptation at time points up to 24 hours after participants had observed another individual undergoing force field learning.

4.1   Temporal Stability of Observation- or Practice-Related Adaptation

Consistent with the findings of Wanda and colleagues (2013), we found that, like participants who learned to reach in a force field through physical practice, observers also learned to predictively generate directionally- and temporally-specific compensatory forces during reaching. Novel to the present study, we found that observation induced an adaptation of predictive limb control that was still detectable 24 hours later, demonstrating that new representations of movement dynamics acquired through observing can be retained, and continue to influence behaviour, well past the duration of the initial training period.

Further, although observing drove less adaptation than physical practice, the passage of time following the training period did not have differential effects on the amount of adaptation measured for participants who observed versus physically performed reaches in a force field. The adaptation had similar temporal stability regardless of whether it was acquired by observing someone else's physical force field learning or feeling the force field for oneself.

Previous studies of physical force field learning have provided evidence that neural representations learned during a single training session can form long-term motor memories that persist for several months after the initial session (Shadmehr and Brashers-Krug 1997). Our finding that observing or physically undergoing force field learning can elicit long-term (>24 hours) effects that are similarly stable over time lends support to the idea that observing may also drive the formation of such long-lived motor memories as those described for physical practice.

4.1.1   Consolidation

Newly acquired representations of movement dynamics are gradually consolidated into stable motor memories over several hours following the end of a physical training period (Brashers-Krug et al. 1996). Although the present study provides evidence that representations acquired by observing force field learning can form long-term motor memories, it does not necessarily demonstrate that such memories have been consolidated. Consolidation refers specifically to a process in which a newly acquired memory trace becomes resistant to interference through the progressive stabilization of a once labile representation (Dudai 2004). That such consolidation occurs in the case of physical force field learning has been demonstrated by the existence of a 5-6 hour consolidation window during which the learned representation is susceptible to behavioural interference but becomes resistant to interference thereafter (Shadmehr and Brashers-Krug 1997).

In the present study, activities undertaken by observers during their 24-hour delay periods were unsupervised and not controlled for. While during this time participants might have carried out activities that could have interfered with the consolidation of the newly acquired representation, we did not specifically assess the susceptibility of observation-induced adaptation to interference by learning of another skill in close temporal proximity to the initial observation period. Future studies may wish to investigate whether a consolidation period such as that which has previously been described in the force field learning literature also exists for representations acquired by observing. If so, this would provide conclusive evidence that observational learning of movement

dynamics is consolidated into a stable motor memory following completion of the observation period (Dudai 2004).

4.2 Implicit Contributions to Observational Learning

Interestingly, the adaptation of predictive limb control after observing force field learning did not seem to depend on observers' conscious awareness of perturbing forces in the video they observed. In our study, participants were never informed of any external influence on the tutor's reaching. Many participants in fact audibly remarked on how "bad" the tutor's reaches were as soon the video began to play, attributing reaching errors to poor control on the part of the tutor rather than suspecting the presence of environmental disturbances. More than half of observers requested information about the purpose of the study—and specifically, the purpose of the video—after they had completed the experiment. Among them, 100% of those who learned that the tutor was being perturbed said they had not previously been aware of any external forces acting on the arm they observed. All but one participant said they thought the tutor was simply poor at reaching straight (the remaining participant reported suspecting that the arm in the video had been their own, covertly recorded during the Baseline block minutes earlier). In any case, each participant said they did not suspect that the observed errors were due to altered dynamics in the reaching environment. That observing can induce an adaptation of force output in the absence of participants' explicit knowledge of the force field is consistent with Mattar and Gribble's (2005) finding that the effects of observation are mediated through the implicit engagement of the motor system, rather than the use of conscious strategy.

Outside of facilitating the development of explicit strategies for reaching, conscious awareness of the presence of perturbing forces during observation could potentially influence the amount and/or temporal stability of the learning that occurs by observing. As much is known to be true for the influence of attentional processes on sensory perception (Logan et al. 1999) and implicit learning in contextual cueing tasks (Jiang and Chun 2001). In their informal comments made after the experiment had ended, many observers shared that during the observation period they were focused (as instructed) on

counting the number of times the individual in the video had reached the end-target in the desired time window. Future research should investigate whether instructing observers to attend to the curvature of the tutor's reach trajectories, or explicitly informing them of the presence of a force field in the video, can modulate the effects of observing force field learning. Similar methods as those used in the present study can be employed to determine how changes in the instructions that participants receive prior to observing might affect subsequent adaptation. In future studies in which observers might be told of the presence of perturbing forces in the video they will observe, the contributions of explicit versus implicit control to the adaptation of predictive force output after observing can be dissociated by manipulating the instructions participants receive prior to starting the Testing block. For example, whether or not participants are told they will begin reaching in the same environment as they felt during the Baseline block would presumably influence their subsequent reaching strategy.

4.3   Resistance to Washout

In our study, the adaptation induced by observing force field learning was more resistant to washout during repeated error clamp trials than was the adaptation induced by physically reaching in a force field—with one exception. After 24 hours, the rate at which physical practice-related adaptation washed out was equivalent to the rate at which observation-related adaptation washed out after a post-training interval of any length (see Figure 5B). We note two potential explanations for our findings. First, it is possible that observing force field learning engages different neural mechanisms than physically reaching in a force field, leading to differences in the rates at which observation- and physical practice-related adaptation are extinguished under the same conditions. In this case, perhaps representations acquired through physical practice undergo a consolidation process that results in their functional migration to the same areas as representations acquired through observing, yielding similar robustness to washout after a 24-hour delay but not in the acute post-training period. Migration of learned representations used for predictive control of reaching has previously been documented to occur in association with consolidation of physical force field learning (Shadmehr and Holcomb 1997).

A second option, which does not preclude the possibility that representations acquired by observing or physical practice both use the same neural substrates and are equally robust to washout, is that the differences in washout rates reported in the present study could be attributable to differences in the amount of adaptation that was available to be washed out in the first place. At the beginning of the washout period (i.e., early Testing block), the amount of adaptation measured after observing was only about 40% of that measured after physical practice. Only at the end of the washout period did levels of adaptation measured for Movement groups approach the amount of adaptation that was initially available to be washed out among observers. In studies of motor adaptation, washout following the removal of a perturbing stimulus characteristically follows a falling exponential curve (see Taylor and Ivry 2014 for example). Washout therefore occurs at a faster rate in more highly adapted states, and at a slower rate in less adapted states. One possibility is that washout rate is causally related to how much adaptation remains to be washed out, rather than being determined as a function of the progression of the de-adaptation process. The washout rates reported for Observation and Movement groups could therefore represent different phases of the same washout process, differentiated by the amount of adaptation left over at the time of the probe. The washout rates we reported for observers might approximate what would have been seen for Movement groups if there had been enough trials for them to continue reaching after meeting the level of adaptation that observers started with. Future studies should investigate this possibility by first determining whether the amount of adaptation detectable at the start of a washout period is causally related to washout rate. This could be done by quantifying washout rate after participants have performed a variable number of reaches in a force field and so have undergone variable amounts of adaptation.

4.4   Inter-Individual Differences in Learned Velocity Dependence

The present study found that changes in lateral forces measured shortly after observing force field learning exhibited a dependence on velocity signals to an extent comparable to what has been documented previously (Wanda et al. 2013); however, there were some participants who did not learn velocity dependence. Although the subset of observers

who failed to produce appropriate velocity-dependent forces in the early Testing block spanned every delay group, the proportion of participants whose adapted forces were statistically unrelated to reach velocity, or negatively correlated with reach velocity, was notably higher 24 hours after observing than within the minutes that followed the end of the observation period. We cannot explain why some participants learned to produce forces that scaled with reach velocity while others did not (or otherwise did not retain this learning for long enough to be detected). Differences in baseline resting-state functional connectivity in sensorimotor brain areas have previously been reported to predict inter-individual differences in the effects of observing force field learning (McGregor and Gribble 2017). Whether such differences might also influence the temporal stability of the learning that occurs through observing may be of interest for future research.

4.5   Additional Future Directions

Some directions for future study have been proposed above; however, future research should also explore the mechanisms underlying the acquisition, retention, and use of neural representations acquired by observing force field learning. There is evidence that some neural regions—namely, primary somatosensory (McGregor et al. 2016) and motor (Brown et al. 2009) cortices—play necessary roles in acquiring new representations of reach dynamics both while observing force field learning and while carrying out physical practice in a force field. Regions of the cerebellum, intraparietal sulcus, and dorsal premotor cortex are also engaged both when physically experiencing reaching errors and when observing another individual undergoing force field learning (Malfait et al. 2010). Observing force field learning is also associated with changes in functional connectivity in a network involving the middle temporal visual area (V5/MT) and cerebellar, primary somatosensory, and primary motor cortices, which could provide a neuroanatomical basis for how visual information about others' movements might reach sensorimotor circuits to facilitate learning (see McGregor and Gribble 2015).

One possible direction for future study is to explore whether, like physically reaching in a force field, learning by observing may occur through sensory prediction error, where discrepancies between the sensed and predicted sensory consequences of motion drive

motor adaptation (Shadmehr et al. 2010). There is evidence that during observation, the observer's central nervous system uses neural representations of movement to predict the tutor's movement kinematics. For example, Flanagan and Johansson (2003) showed that while observing a tutor perform a motor task, observers' eye movements were predictively, rather than reactively, coordinated with movement of the tutor's hand, suggesting that a feedforward process was used to predict the tutor's hand kinematics and appropriately direct the observers' gaze. In the context of observing force field learning, we speculate that during observation, visual information about the tutor's reach kinematics may be compared against the central nervous system's predictions of what reach kinematics ought to look like. Presumably, such predictions are informed by movement representations acquired during unperturbed reaching in the Baseline block. Future studies could delve deeper into this line of thinking by examining adapted force output profiles generated by observers who were trained to perform Baseline reaches along a different trajectory than would be provided by a straight line between 'home' and end targets.

Outside of further investigation into the mechanisms by which new representations are acquired by observing, future studies should also look into methods for improving the efficacy of videos generated to induce learning of neural representations of movement dynamics. The present study has demonstrated that representations of dynamics can be learned by observing, and that such representations can continue to influence behaviour beyond the duration of the initial training period. Here, watching only a single, 12-minute video induced about 40% of the adaptation gained through physically performing reaches in a novel dynamic environment. We suspect that our video and observation protocol could be modified to produce stronger effects—for example, by increasing the size and number of reaching errors participants observe (Brown et al. 2010), or playing the observation video twice (Wanda et al. 2013). Carrying out multiple training sessions might also be helpful, since this has been shown to improve the learning and retention of skills learned through physical practice (Shea et al. 2000). It would be beneficial for future work to determine whether such modifications might improve learning and retention after observing.

Future studies should also explore potential observation-based approaches to neurorehabilitation. The reorganization of motor circuits involved in force production and control has been recognized as a fundamental aspect of motor recovery after stroke (Kokotilo et al. 2009). Studies in non-human primates have demonstrated that physical force field exposure induces neuronal plasticity in primary motor cortex, reflecting the acquisition of new representations of dynamics (Li et al. 2001). The authors also found that changes in neural activity persisted even after washout of after-effects, providing one potential neural basis for how motor memories might be preserved in primary motor cortex. Other studies have also provided evidence that motor cortical representations of dynamics for simpler tasks could contribute to representations of dynamics for more complex behaviour (Gribble and Scott 2002), providing a potential basis for how learning simpler motor skills could contribute to learning more complex ones. In the present study we have demonstrated that observing force field learning can drive the formation of long-term motor memories, which influence predictive limb control in a manner similar to what is seen following physical force field learning. For patients with disorders that reduce their ability to undergo physical practice (for example, stroke patients with severe hemiparesis), an approach to rehabilitation that drives lasting cortical reorganization primarily by observing could be beneficial. A future line of research could work towards determining the extent to which observational methods may be able to reduce the amount of physical therapy required for recovery of motor function. Taking into consideration the recommendations that have been made elsewhere within this thesis (particularly with respect to assessing the influence of instructional variation on adaptation, and making modifications that could reasonably be expected to boost the efficacy of observational intervention), a good place to start might be to carry out a hybrid version of the present study that aims to determine how much physical force field reaching can be replaced by observing without reducing the total amount of learning that occurs.

References

Bernardi NF, **Darainy M**, **Bricolo E**, **Ostry DJ**. Observing motor learning produces somatosensory change. *J Neurophysiol* 110: 1804–1810, 2013.

Brashers-Krug T, **Shadmehr R**, **Bizzi E**. Consolidation in human motor memory. *Nature* 382: 252–255, 1996.

Brown LE, **Wilson ET**, **Gribble PL**. Repetitive transcranial magnetic stimulation to the primary motor cortex interferes with motor learning by observing. *J Cogn Neurosci* 21: 1013–1022, 2009.

Brown LE, **Wilson ET**, **Obhi SS**, **Gribble PL**. Effect of trial order and error magnitude on motor learning by observing. *J Neurophysiol* 104: 1409–1416, 2010.

Conditt MA, **Gandolfo F**, **Mussa-Ivaldi FA**. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *J Neurophysiol* 78: 554–560, 1997.

Diedrichsen J, **Hashambhoy Y**, **Rane T**, **Shadmehr R**. Neural Correlates of Reach Errors. *J Neurosci* 25: 9919–9931, 2005.

Dudai Y. The neurobiology of consolidations, or, how stable is the engram? *Annu Rev Psychol* 55: 51–86, 2004.

Flanagan JR, **Johansson RS**. Action plans used in action observation. *Nat 2003 4246950* 424: 769–771, 2003.

Gandolfo F, **Mussa-Ivaldi FA**, **Bizzi E**. Motor learning by field approximation. *Proc Natl Acad Sci U S A* 93: 3843–3846, 1996.

Goodbody SJ, **Wolpert DM**. Temporal and amplitude generalization in motor learning. *J Neurophysiol* 79: 1825–1838, 1998.

Gribble PL, **Scott SH**. Overlap of internal models in motor cortex for mechanical loads

during reaching. *Nat 2002 4176892* 417: 938–941, 2002.

Huang VS, **Shadmehr R**. Evolution of motor memory during the seconds after observation of motor error. *J Neurophysiol* 97: 3976–3985, 2007.

Hwang EJ, **Donchin O**, **Smith MA**, **Shadmehr R**. A Gain-Field Encoding of Limb Position and Velocity in the Internal Model of Arm Dynamics. *PLoS Biol* 1, 2003.

Jiang Y, **Chun MM**. Selective attention modulates implicit learning. *Q J Exp Psychol Sect A Hum Exp Psychol* 54: 1105–1124, 2001.

Joiner WM, **Smith MA**. Long-Term Retention Explained by a Model of Short-Term Learning in the Adaptive Control of Reaching. *J Neuro-physiol* 100: 2948–2955, 2008.

Kokotilo KJ, **Eng JJ**, **Boyd LA**. Reorganization of brain function during force production after stroke: A systematic review of the literature. *J Neurol Phys Ther* 33: 45–54, 2009.

Li CSR, **Padoa-Schioppa C**, **Bizzi E**. Neuronal Correlates of Motor Performance and Motor Learning in the Primary Motor Cortex of Monkeys Adapting to an External Force Field. *Neuron* 30: 593–607, 2001.

Logan GD, **Taylor SE**, **Etherton JL**. Attention and automaticity: Toward a theoretical integration. *Psychol Res 1999 622* 62: 165–181, 1999.

Malfait N, **Shiller DM**, **Ostry DJ**. Transfer of Motor Learning across Arm Configurations. *J Neurosci* 22: 9656–9660, 2002.

Malfait N, **Valyear KF**, **Culham JC**, **Anton JL**, **Brown LE**, **Gribble PL**. fMRI Activation during Observation of Others' Reach Errors. *J Cogn Neurosci* 22: 1493–1503, 2010.

Maschke M, **Gomez CM**, **Ebner TJ**, **Konczak J**. Hereditary cerebellar ataxia progressively impairs force adaptation during goal-directed arm movements. *J Neurophysiol* 91: 230–238, 2004.

Mathew J, **Lefèvre P**, **Crevecoeur F**. Savings in Human Force Field Learning Supported by Feedback Adaptation. *eNeuro* 8, 2021.

Mattar AAG, **Gribble PL**. Motor learning by observing. *Neuron* , 2005. doi:10.1016/j.neuron.2005.02.009.

McGregor HR, **Cashaback JGA**, **Gribble PL**. Functional Plasticity in Somatosensory Cortex Supports Motor Learning by Observing. *Curr Biol* 26: 921–927, 2016.

McGregor HR, **Cashaback JGA**, **Gribble PL**. Somatosensory perceptual training enhances motor learning by observing. *J Neuro-physiol* 120: 3017–3025, 2018a.

McGregor HR, **Gribble PL**. Changes in visual and sensory-motor resting-state functional connectivity support motor learning by observing. *J Neurophysiol* 114: 677–688, 2015.

McGregor HR, **Gribble PL**. Functional connectivity between somatosensory and motor brain areas predicts individual differences in motor learning by observing. *J Neurophysiol* 118: 1235–1243, 2017.

McGregor HR, **Vesia M**, **Rinchon C**, **Chen R**, **Gribble PL**. Changes in corticospinal excitability associated with motor learning by observing. *Exp Brain Res* 236: 2829–2838, 2018b.

Miall RC, **Wolpert DM**. Forward Models for Physiological Motor Control. *Neural Networks* 9: 1265–1279, 1996.

Sainburg RL, **Ghez C**, **Kalakanis D**. Intersegmental dynamics are controlled by sequential anticipatory, error correction, and postural mechanisms. *J Neurophysiol* 81: 1045–1056, 1999.

Scheidt RA, **Reinkensmeyer DJ**, **Conditt MA**, **Zev Rymer W**, **Mussa-Ivaldi FA**. Persistence of motor adaptation during constrained, multi-joint, arm movements. *J Neurophysiol* 84: 853–862, 2000.

Schmidt RA, **Lee TD**. *Motor Learning and Performance: From Principles to Application*. 6th ed. Champaign, IL: Human Kinetics, 2014.

Shadmehr R, **Brashers-Krug T**. Functional Stages in the Formation of Human Long-Term Motor Memory. *J Neurosci* 17: 409–419, 1997.

Shadmehr R, **Holcomb HH**. Neural Correlates of Motor Memory Consolidation. *Science (80-)* 277: 821–825, 1997.

Shadmehr R, **Moussavi ZMK**. Spatial Generalization from Learning Dynamics of Reaching Movements. *J Neurosci* 20: 7807–7815, 2000.

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

Shadmehr R, **Smith MA**, **Krakauer JW**. Error Correction, Sensory Prediction, and Adaptation in Motor Control. *Annu Rev Neurosci* 33: 89–108, 2010.

Shea CH, **Lai Q**, **Black C**, **Park JH**. Spacing practice sessions across days benefits the learning of motor skills. *Hum Mov Sci* 19: 737–760, 2000.

Smith MA, **Shadmehr R**. Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. *J Neurophysiol* 93: 2809–2821, 2005.

Takahashi CD, **Scheidt RA**, **Reinkensmeyer DJ**. Impedance control and internal model formation when reaching in a randomly varying dynamical environment. *J Neurophysiol* 86: 1047–1051, 2001.

Taylor JA, **Ivry RB**. Cerebellar and Prefrontal Cortex Contributions to Adaptation, Strategies, and Reinforcement Learning. In: *Progress in Brain Research*. Elsevier, 2014, p. 217–253.

Thoroughman KA, **Shadmehr R**. Electromyographic Correlates of Learning an Internal Model of Reaching Movements. *J Neurosci* 19: 8573–8588, 1999.

Wanda PA, **Li G**, **Thoroughman KA**. State dependence of adaptation of force output following movement observation. *J Neurophysiol* 110: 1246–1256, 2013.

Williams A, **Gribble PL**. Observed effector-independent motor learning by observing. *J Neurophysiol* 107: 1564–1570, 2012.