VIRTUAL PARTNERS DIFFERENTIALLY INFLUENCE REACH CONTRIBUTION

REACH CONTRIBUTIONS DURING VISUOMOTOR ADAPTATION ARE DIFFERENTIALLY INFLUENCED BY ONE'S VIRTUAL PARTNER

By NOUR AL AFIF, B.Sc. (Hons)

A thesis submitted to the School of Graduate Studies in partial fulfilment of the requirements for the Degree Master of Science

Master of Science (2024) McMaster University
Department of Kinesiology Hamilton, ON, Canada

TITLE:	Reach contributions during visuomotor adaptation are differentially influenced by one's virtual partner
AUTHOR:	Nour Al Afif, Degrees (Institutions)
SUPERVISOR:	Michael Carter
NUMBER OF PAGES:	ix, 50

Lay abstract

Working with a physiotherapist is the gold-standard in rehabilitating many injuries, but this can be very time consuming and repetitive in nature. This makes it worthwhile to explore other alternatives to supplement standard rehabilitation, such as working with a virtual partner. In our experiment, we tested two partners based on human models. Participants were paired with one of the virtual partners and had to reach a target using a handle, adjusting their reach to a rotation. The partners differ in how fast they help the participant adjust for the rotation. It was found that those who completed the task with a fast-learning partner corrected less error initially and more later on, while those with the slow-learning partner corrected more error initially and less later on. These results suggest that we can influence participant behaviour with different virtual partners.

Abstract

From a parent guiding their toddler when learning to brush their teeth to a physical therapist assisting a client with their range of motion, physically interacting with other people is ubiquitous in our daily life. While some researchers have shown that haptic human-human interaction benefits performance during training as well as later individual performance (Takagi et al. 2017), others have failed to replicate these benefits (Beckers et al. 2018). Participants in these interaction groups were not aware they were haptically linked to a partner and each participant had independent control over their own virtual cursor when tracking the target. Yet, we are typically aware when we are interacting with others and often do so with tasks where we have shared control over the same control point (e.g., a toothbrush). Here, we tested the effectiveness of training alone versus training with a virtual partner when individuals were made aware of their interaction in a redundant reaching task. Participants (N = 100) completed 50 baseline trials followed by 200 trials with a clockwise cursor rotation in one of four randomly assigned groups. Two of the groups performed the adaptation trials with a virtual partner that represented either the fast (Human + Fast Agent Group) or slow (Human + Slow Agent Group) state of the two-state model (Smith et al. 2006) with 30-deg rotation. The two remaining groups performed the task alone with either the 30-deg rotation (Solo full rotation) or a 15-deg rotation (Solo half rotation). Results showed that participants in the fast agent group contributed less to correcting the rotational error early in the adaptation block, but were responsible for

most of the correction later in this block, with performance most similar to the solo full rotation group. Conversely, participants in the slow agent group corrected for a greater proportion of the initial errors, but their contribution began to drift during adaptation, with performance resembling that of the solo half rotation group. This pattern of results were consistent with our theory-driven simulations.

Acknowledgements

I would like to start off by thanking my supervisor Mike Carter for his guidance and support throughout my Masters degree. I have learned so much from you, and look forward to learning so much more over the next 4-6(?) years with you as my mentor. Thank you for letting us bring our lab visions alive — the ducks and lab instagram - and always joining in on the inside jokes.

Thank you to my committee members Dr. Jim Lyons and Dr. Keith Lohse, and my external examiner Dr. Dylan Kobsar for taking time out of their schedules to provide me with valuable input over the past year, and being a part of my journey.

To my lab mates, Mikayla, Jack, Daniel and Laura- Thank you for always listening to my crazy ideas (research or otherwise), proofreading things, and being my coding ducks when Duck Bryan is far. I enjoy our dynamic and that we have become good friends. Thanks to my roommates and friends in the department for your constant advice and support. Depali, Jacob and Ali - I would not have gotten here without your belief in my abilities.To my siblings, Hazel, Hams and Faiz — thanks for letting me ramble, even when I don't fully know what I'm saying. Finally, a special thanks to my parents, Nader and Suzan, who believed in me and supported me as I found my way. Thank you for always celebrating my accomplishments.

Table of contents

List of Figures

Declaration of academic achievement

Format and organization of thesis

This thesis is prepared in the standard format as outlined in the Guide for the preparation of master's and doctoral theses provided by McMaster University's School of Graduate Studies.

Contributions to content of thesis

Nour Al Afif; Joshua G.A. Cashaback; Michael J. Carter

Contributions using CRediT (Contributor Roles Taxonomy)

N. Al Afif: Data curation, Formal analysis, Investigation, Methodology, Project administration, Validation, Visualization, Writing - Original Draft Preparation, Writing - Review & Editing *M. Lalli:* Data curation, Investigation, Validation, Writing - Review & Editing *D. Deletsu:* Writing - Review & Editing *O.Schwarzenberg; L. Barbera:* Investigation *R. Lokesh:* Software, Methodology *J.G.A. Cashaback:* Conceptualization, Methodology, Visualization, Supervision *M.J. Carter:* Conceptualization, Data curation, Funding acquisition, Methodology, Supervision, Writing - Review & Editing

Chapter 1

Literature review

Throughout our lives, we constantly have to learn how to move our bodies, and other objects around us. This applies to many activities, from learning how to play with toys, to how to type and even to learning the movements we need for different sports and musical instruments. These same tasks can have variations when we do need to perform them. For example, we could go to a friend's house to play with their slightly different toys, use a computer at the library to type up an assignment instead of our own, or try to keep up in a soccer game with sore muscles. These events all require slight changes in the way we complete the task, without learning a new skill entirely. Such processes all fall under the umbrella of motor learning, which is broadly defined by Krakauer and colleagues as any "experience-dependent improvement in performance" (2019). Improvements in performance lead us to complete more effective movements. However, this does not only apply to the learning of completely new skills, but also the adjustment of those previously learned. With this distinction in mind, two main facets emerge encompassing motor learning; skill acquisition and skill maintenance.

1.1 Skill Acquisition

Skill Acquisition is the process in which new skills are learned through practice, such as a novice learning how to serve a ball with a tennis racquet (Krakauer et al., 2019). Skill Acquisition begins with one attempting the skill, usually for the first time, with a very high degree of error. With consistent practice on the skill, there is a decrease in error (Sternad, 2018). After an extended period of practice, there is longer term retention of this skill, such that when one wants to apply it, serving a ball at their first game, they do so without (or with very little) error (Krakauer et al., 2019). An extensive area of study under skill acquisition is dedicated to finding and implementing interventions that enhance the process of acquisition and identifying factors that limit it.

Some factors that have been found to improve skill acquisition include: providing participants with knowledge of results with error estimation information (Guadagnoli & Kohl, 2001), learning in distributed practice rather than massed practice (Mackay et al., 2002) and increasing the variability of the tasks practiced (Goodwin et al., 1998). Skill acquisition studies usually compare participants in two (or more) groups, either with the intervention(s) or without them. Participants undergo a block of skill acquisition, followed by a retention or transfer block where they are tested on the same, or a similar task (Sternad, 2018). The difference in performance between the groups is assessed to evaluate the effectiveness of the intervention. The group with less error or variability in completing the retention block is considered to have learned more effectively. From this, a researcher can conclude whether the intervention is considered to improve skill acquisition in the task.

1.2 Skill Maintenance

Skill maintenance, on the other hand, is the continuation of practice on a skill that is already known to maintain one's current level of performance in different conditions (Krakauer et al., 2019). For example, when the same novice in tennis attempts using a racquet that feels heavier than their own, many errors occur in the beginning. Those errors, however,

dissipate after a short practice session and they return to their baseline skill level.

A type of skill maintenance that is commonly studied in the field of motor learning is *adaptation*, which is adjusting a movement in order to successfully counter a perturbation. Perturbations can be an external change in the environment, such as the new racquet, or an internal change, muscle fatigue after a tough workout (Schmidt et al., 2019). When encountering a perturbation, individuals usually make transient changes, which allow them to complete the task with a similar rate of error to what they had at baseline. Adaptation can be studied in various ways in a lab setting, with tasks that include reaching, throwing and walking (Fernández-Ruiz & Díaz, 1999; Reisman et al., 2005; Shadmehr & Mussa-Ivaldi, 1994).

1.3 Studying Adaptation

1.3.1 Study Design

In a basic study setup investigating adaptation there are two main blocks; the baseline block and the adaptation block. In the baseline block, participants complete the task devoid of any perturbations, and researchers use this block to obtain baseline error rates . Following this is the adaptation block, where the perturbation is introduced. Perturbations in these experiments can be visual or proprioceptive, impacting the participant in different ways. Typical perturbations used in a laboratory setting include prism goggles, force fields, split belt walking and visuomotor rotation (Krakauer et al., 2019). In the beginning of the adaptation block, known as the early adaptation phase, participants will complete the task as they normally would, unaware of the perturbation and having a high degree of error (Shadmehr et al., 2010). As the adaptation block progresses, they begin adjusting their movement to counter the perturbation. This makes error rates in the late adaptation phase similar to the levels participants had at baseline (Krakauer et al., 2019). These phases alone are enough to study the process of adaptation, however, other phenomena such as aftereffects, washout, spontaneous recovery and savings can be examined by adjusting the

study protocol.

Aftereffects can be examined by adding a washout block following adaptation in which there is no perturbation (Krakauer et al., 2019). Aftereffects are comprised of the movement the participant used to counter the perturbation, usually in the opposite direction of the perturbation, and are an indication that one has adapted to the perturbation (Krakauer, 2009; Moore & Cluff, 2021). For example, if the perturbation shifts the participant's arm to the right while they reach towards a target, this would require the participant to move their arm to the left in order to counter this perturbation. The aftereffect seen by the researcher would be a leftward movement in the washout block. Aftereffects persist for a few trials, but dissipate quickly as the person returns to their baseline state, a phenomenon known as washout (Krakauer et al., 2019). In the aforementioned example, this would be the participant's reach slowly reverting back to the straight reach they completed in baseline. Washout does not mean that the adaptation has fully dissipated, the adaptive mechanism can manifest as spontaneous recovery or savings after washout has occurred.

Spontaneous recovery is a drift to the behaviour the participant used to counter the perturbation, that occurs after the behaviour is thought to have washed out (H. E. Kim et al., 2021). This usually occurs when adding an error-clamp block after the washout block (Criscimagna-Hemminger & Shadmehr, 2008). In an error clamp block, participants receive either no feedback or feedback that indicates that they are completing the task with no error.

As for *savings*, this is when previous adaptation manifests in the faster relearning of the perturbation upon a second introduction (Huberdeau et al., 2015; Krakauer, 2009). Savings are examined experimentally through the introduction of another perturbation block either immediately after the washout block (Morehead et al., 2015), or on the following day (Haith et al., 2015). Huberdeau and colleagues aimed to examine the smallest amount of adaptation trials that would manifest in savings after a one day delay in a reaching task with a visuomotor rotation (2015). It was found that even after a minimal adaptation block,

ranging from 2 to 5 trials, savings occurred in the block on the next day (Huberdeau et al., 2015). In another experiment, the size of the perturbation determined whether or not there would be savings, with larger perturbations causing a higher extent of savings (Morehead et al., 2015).

Savings were also examined in a study by Coltman and colleagues where participants completed the same adaptation paradigm twice, separated by a five minute break (2019a). Participants in the study adapted at a faster rate in the second adaptation block, similar to tasks that only included a savings block following washout. Various tasks and perturbations can be used to gain insight into the process of adaptation and its associated phenomena. Some examples will be discussed, in the approximate chronological order of their emergence in motor learning research.

1.3.2 Common Tasks in Adaptation Studies

Completing Tasks with Prism Goggles

The idea of visuomotor perturbation first came about through the study of reaching with prism goggles. Prism goggles cause vision to shift at a certain angle, and distort the participant's visual feedback of their arm, causing reaching errors (Welch, 1971). This was first studied by von Helmholtz (von Helmholtz, 1924), providing much insight into the way humans process and learn from visual information. In this work, it was found that reaching while wearing prism goggles caused large amounts of error, however, such error is reduced after a few trials with the perturbation. Hamilton and colleagues expanded on those results by finding that aftereffects occurred in the participants' opposite arms when they were allowed to move their trunk with the arm they used to reach (Hamilton, 1964). These results dissipated when participants' trunks were fixed in place.

Reaching tasks were not the only application of prism goggles, Fernandez Ruiz and colleagues made participants complete a ball-throwing task with prism goggles, where the same adaptive mechanism was shown (Fernández-Ruiz & Díaz, 1999). They also found correlations between the extent of perturbation (diopter of the prism) and the length of the

adaptation block to the persistence of the after effect. Increasing either the magnitude of the perturbation or length of time exposed to the perturbation causes the after effect to persist for longer after the perturbation is removed (Fernández-Ruiz & Díaz, 1999).

Reaching in force fields

With improvements in technology and the introduction of robotic manipulanda, reaches could be perturbed with force fields. A common type of force field used in research is a velocity dependent force field, where the amount of force is dependent on the speed the participants move their arm. Shadmehr and Mussa-Ivaldi had participants complete a reaching movement with a perpendicular velocity-dependent force field (1994). In a later block, when the force field was removed, an aftereffect was observed, indicating adaptation to the force field in the perturbed trials. In another study by Crevecoeur and colleagues, participants even showed reduced errors, as well as aftereffects to a force field even when the force field was applied in random trials (2020).

Split belt walking

Split belt treadmills are another commonly used task to study adaptation. In a split belt treadmill, there are two belts, one for each foot, which are able to move at different rates. In the baseline phase, the belts are tied together and move at the same speed. The perturbation includes the belts being unsynced and moving at different speeds. In order to adapt to this perturbation, participants would have to change their gait pattern (Schmidt et al., 2019). Aftereffects are also seen when the belts are re-synced after adaptation, with the step length of the previously faster leg being larger (Reisman et al., 2005). This result is consistent even in a post-stroke population, where both healthy controls and the post-stroke group showed after-effects, indicating that they were both able to adapt to this perturbation (Reisman et al., 2009).

Reaching in visuomotor rotation

In a visuomotor rotation task, participants are usually unable to see their hand when performing the task, and are only able to see a cursor that represents their unseen hand. In the baseline block, participants' hand movements are reflected accurately by the cursor

they see on screen (as shown in Figure 1). When the adaptation block begins, however, the cursor is rotated at an angle and no longer reflects the person's hand movement (Krakauer et al., 2019). The discrepancy between the cursor position and hand position will cause a high degree of error at the beginning of the adaptation block (or early adaptation phase in Figure 1.1).

Figure 1.1

Typical design of a visuomotor rotation paradigm with only a baseline and adaptation block.

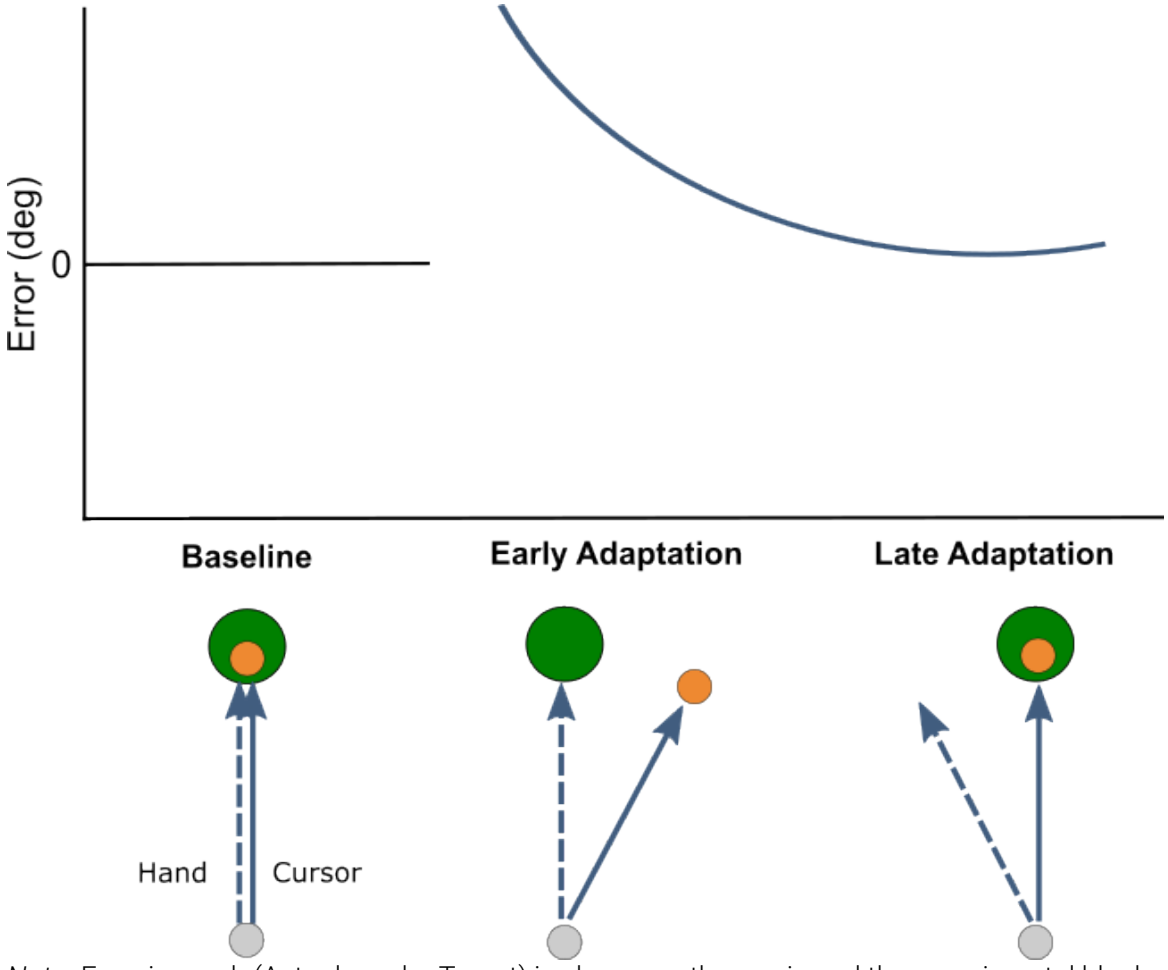

Note. Error in reach (Actual reach - Target) is shown on the y-axis and the experimental blocks are shown on the x-axis. On the graph, the baseline condition is shown in black, and the adaptation block is divided into early and late adaptation, shown in blue. The participant's reach in all phases is shown above, where in baseline they reach directly for the target with the cursor directly overlapping with their hand movement, but their reach happens at an angle in the adaptation phase to counter the perturbation angle.

This causes the participants to adapt their reach to counter the rotation and hit the target in the late adaptation phase (Shadmehr et al., 2010). Sakamoto and Kondo found that even a passive adaptation block, where participants are not actively engaging in moving their arm, improved errors from early adaptation (2012). Participants in the passive group had their arm moved by a robotic manipulandum. However, active movement and adaptation still improved error and reaction time significantly more than the passive experiment. This pattern of behaviour typically seen in an adaptation experiment is consistent across ages (E. R. Buch et al., 2003). However, older adults adapt to a lesser extent in comparison to younger participants.

Another factor that may affect the extent of adaptation is the number of targets used to practice countering the perturbation. In a study by Neva and Henriques, it was found that complete adaptation to a visuomotor rotation only occurred in participants who completed their adaptation block with a single target rather than multiple different ones (2013). Although both groups showed adaptation to the perturbation, the single target group adapted faster and more completely (Neva & Henriques, 2013).

This thesis will employ a visuomotor rotation task with a single target in the study of adaptation. As discussed, the process of adaptation presents consistently between individuals, and across different tasks and perturbations. Due to this consistency, a line of research emerged in an effort to model adaptive behaviours.

1.4 Models of Adaptation

1.4.1 State Space Models

Over the years, there has been interest in quantifying the way movements and learning occur. This was done through the use of forward models, which are models that use the current state of an organism (or effector) in order to learn what their following state will be (Miall & Wolpert, 1996). These models usually come in the form of a defined equation that can simulate the end state of a movement. A sub-type of forward model used often in

many disciplines, from engineering to ecology to motor learning, is a state space model. State space models use the current state of an organism alongside observations of error or likelihood of movement, to learn the next state that their effector will be in (Patterson et al., 2008). As more research has emerged, state-space models used for adaptation have evolved to model the change in error correction from trial to trial (Smith et al., 2006). Adaptation was first quantified with a one-state model, but developed into considerations for two and multi-state models due to gaps in the implementation of one-state model (Krakauer et al., 2019; Smith et al., 2006).

1.4.2 Evolution of the State Space Models of Adaptation

Due to the nature of adaptation, it was thought to be a solely implicit process, with only one state defining it and no awareness from the participant on how it is occurring (H. E. Kim et al., 2021). This one state model is quantified by:

$$x(t+1) = A_x \cdot x(t) + B_x \cdot e(t) \qquad (1)$$

,where $x(t)$ = position at trial t, A= retention rate, B = learning rate, $e(t)$ = error at trial t (Smith et al., 2006). The two parameters, retention rate and learning rate, in this equation work together to adjust trial-by-trial error correction. Specifically, retention rate (A) indicates the exent of the state estimate that is retained from the previous trial and learning rate (B) indicates the proportion of error that will be corrected from the previous trial (H. E. Kim et al., 2021). Although this model was able to characterize many aspects of adaptation, it did not account for savings or washout when an adaptation paradigm was simulated. In another model, Krakauer and colleagues attempted to fit the time course of adaptation with one exponential function, which was inadequate to model the changes seen in the experimental data as well (2000). Two exponential functions yielded a better fit, suggesting the presence of two different processes working on adaptation. This is explored in various

studies after this point.

Kojima and colleagues (2004) completed saccadic adaptation tasks in monkeys, where they had to adapt to a target shift during their saccade. The monkeys completed two adaptation blocks with a washout block in between, and it was found that upon the second introduction of the perturbation, that they exhibited savings. Therefore, Kojima and colleagues suggested a gain-specific model, where they defined specific mechanisms to either increase or decrease gain (Kojima et al., 2004; Smith et al., 2006). Although this model successfully modeled adaptation and savings, it still failed to account for the spontaneous recovery that the monkeys showed when completing the task without feedback.

Smith and colleagues (2006) simulated multiple different models, including Kojima's two-state gain specific model (2004), as well as a single state model and a novel two-state gain independent multi-rate model. The two-state gain independent multi-rate model will be referred to as the two state model in this work. Smith and colleagues also had participants complete a force field task, with baseline, adaptation, counter adaptation and error clamp blocks (2006). The simulated models were then fit to the data collected, with the two-state model having the best fit. This suggests that within the behaviour that is seen from a participant in an adaptation task, there are two gain-independent models operating on different timescales (Shown graphically in Figure 1.2). There is a slow state (in blue) that takes long to learn how to counter perturbations but retains these error correction mechanisms for a longer time. It is the slow state that counters the bulk of the perturbation near the end of the adaptation block. This slow state happens in parallel with a fast state (in orange) that is highly sensitive to errors in the beginning, when they are bigger and more frequent. At this point, the fast state adapts to counter the errors quickly, but does not retain those error correction mechanisms and stops correcting near the end of the adaptation block. It is the slow model that retains most of the error correction mechanisms near the end of the adaptation block. These two models are quantified with the equations below, ,where t represents the current trial, with $t + 1$ being the next trial. x_f and x_s

represent the fast and slow states, respectively:

$$x_f(t+1) = A_f \cdot x(t) + B_f \cdot e(t) \tag{2a}$$

$$x_s(t+1) = A_s \cdot x(t) + B_s \cdot e(t) \tag{2b}$$

with net behaviour seen from the participant represented as:

$$x_{net}(t+1) = x_f(t+1) + x_s(t+1) \tag{2c}$$

, the sum of the fast and slow models. A_f and A_s represent the different retention rates of the fast and slow states respectively, with $A_s > A_f$. B_f and B_s are learning rates of the fast and slow states, with $B_f > B_s$, and $e(t)$ represents error in the given trial t.

Figure 1.2

The Two State Model of adaptation as described in the equations in Smith et al., 2006 in a visuomotor rotation task.

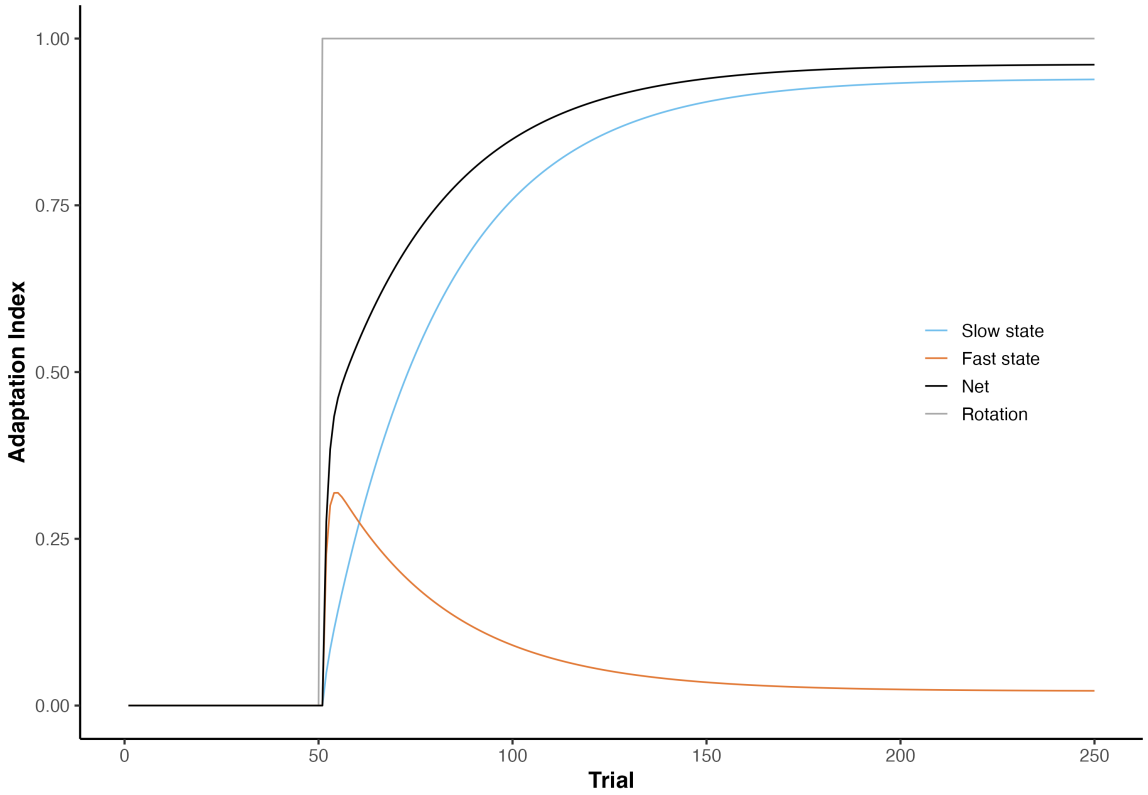

Note. The first 50 trials representing a baseline block and trials 50-250 representing the adaptation block with a visuomotor perturbation of 30 degrees. Rotation as an adaptation index is shown in grey, fast state model (with values Af = 0.597 , Bf = 0.227) in orange, slow state model (with values As = 0.998 , Bs = 0.049) in blue and the combination of the two models as net movement in black.

Smith and colleagues used the two-state model to simulate Kojima and colleagues' data and were able to fit the two-state model to the data successfully (2004). Most notably, it was able to model the spontaneous recovery that occurred in the beginning of the second relearning block, which was lost in the previous models of adaptation. Experimentally, the two state model has been tied to the explicit and implicit strategies used in adaptation. This has been tested in a visuomotor rotation task completed by McDougle and colleagues (McDougle et al., 2015). Participants completed outward reaching motions with a visuomotor rotation, while actively reporting the angle they were heading to. The reported value serves as the explicit aiming value, while implicit learning is measured by subtracting the reported value from the participants' actual aim. These explicit aiming values showed a good fit to the fast model, while the slow model resembled that of the implicit aiming values. Similarly, the two state model accounted for error correctly in a split-belt treadmill task (Mawase et al., 2013) and a force field perturbation task (Coltman et al., 2019a). The two-state model and other models of adaptation, although accurate, only represent adaptation in one individual. However, humans complete many motor tasks with others, making it important to examine how adaptation occurs during an interaction.

1.5 Interacting with others

In our day-to-day lives, there are many motor tasks that we complete together, from team sports such as soccer and hockey, to simply helping a friend move a couch. In all of these instances, there is a need to coordinate and collaborate with other people to complete these tasks effectively. This is studied in the field of joint action, defined as a social interaction where individuals coordinate their actions in order to make a change in their environment (Sebanz et al., 2006). During an interaction, partners do not necessarily have to share auditory or visual information. Haptic information plays an important role in interaction as it has less sensory delay than other forms of feedback, and can still convey information about one's partner (Sawers & Ting, 2014). In moving a couch, one can infer their upcoming direction of movement from where their partner nudges them to, even with the couch

blocking their view. Undeniably, not all joint actions are the same, as there are many ways two (or more) people can interact. These interactions can be classified either through the role each participant takes or with the nature of the task they are completing together.

1.5.1 Classifying Joint Interactions

Jarrasse and colleagues created categories for joint interactions, based on the roles the co-actors embody in the interaction (Jarrassé et al., 2012a). This framework was made with both human-human and human-robot interactions in mind. From these categories, only two will be briefly discussed in this work, collaboration and cooperation. Collaboration is an interaction where one takes into account their partner's as well as their own effort and error (Sawers & Ting, 2014). In this type of interaction, the roles that participants have in the interaction are not predefined, so they work together to achieve their common goal. For example, when carrying a couch with a friend, these roles are constantly changing with the various turns and paths that need to be taken, so they are not predefined. On the other hand, in cooperative interactions, the roles are predefined, with one person usually being the person receiving education or assistance completing the task (Sawers & Ting, 2014). In this scenario, the person receiving assistance does not take into account the error of their partner, and is mostly focusing on their own effort and error (Jarrassé et al., 2012a). An example of this would be physiotherapy, where a patient receives help from the physiotherapist, without expecting to have to reciprocate the help and error correction given to them.

Tasks can also be classified by the nature of the task itself, whether the outcome is achieved through a single point or needs more than one point to be completed. In a non-redundant task, participants work on independent tasks but are haptically connected with a force during task completion (Beckers et al., 2020; eg. Takagi et al., 2017a). Tasks of this nature have been used in many recent experiments studying joint action. However, non-redundant tasks are not reflective of real world applications. Many tasks that we complete with others align more closely to the classification of a redundant task. In a redundant task, the task

can be completed successfully with either one or more people, making changes on the same control point (Peña-Pérez et al., 2023). Examples of this would be akin to a physiotherapist helping a patient move their knee, or experimentally, when participants share a single cursor (Reed et al., 2006). Redundancy has previously been studied in bimanual tasks — where error is usually attributed by the participant to the most likely cause of error, the non-dominant hand (Reichenbach et al., 2013; White & Diedrichsen, 2010). Since the study of redundancy usually formed theories based on one person completing a task individually, it is hard to understand how these processes work in a joint interaction. This is due to numerous reasons, most importantly that in a joint action, there is the additional challenge in that each person does not know much, if anything at all, about their partner's plan of movement. In some joint laboratory tasks, the participants are not even aware of their partner's presence or movement location (Beckers et al., 2018; Ganesh et al., 2014). This makes it more difficult to attribute and assign error to the most likely cause as there are a lot of unknown factors to each participant regarding their partner. Despite these gaps in the study of redundancy, there has been an abundant line of research examining the effects of joint interactions on both group and individual performance.

1.5.2 Effects of interaction on group performance

Wegner and Zeaman completed some of the first work examining performance in a joint action, finding that those completing a pursuit rotor task in a group of two or four had better performance during the task than those completing it on their own (Wegner & Zeaman, 1956). In a later study, Reed and colleagues found that performing a redundant task with a partner improved performance in comparison to performing it alone (2006). Participants completed a target acquisition task in pairs seated across from each other with no view of their partner, where they shared a single rigid crank. The majority of the dyads had better performance than either group member working alone (Reed et al., 2006). Therefore, working with another person can improve performance during task completion. This same result was found during the completion of other target tracking tasks while adapting to visuomotor rotations (Beckers et al., 2020; Ganesh et al., 2014; Takagi et al., 2017a) and force

fields (Beckers et al., 2018). These results also extend to working with a robotic partner, where robot-assisted participants improved more in reaching performance compared to free-reaching participants (Iwamuro et al., 2008). However, there is no indication from these results specifically whether completing a task as a dyad improves solo performance later on.

1.5.3 Effects of interaction on solo performance

To examine the effects of group interaction on solo performance, studies are conducted so that participants experience the task in both connected and solo trials. In these experiments, participants were seated across from each other, each holding a robotic handle, with their view of their hand and their partner occluded (Beckers et al., 2018, 2020; Ganesh et al., 2014; Takagi et al., 2017a). A key difference to note between this paradigm and that of Reed and colleagues is that rather than performing a redundant task together, participants were performing a non-redundant task, each with control over an individual point. The participants were haptically connected to their partner by a virtual elastic band, which would produce a force, nudging them towards their partner's cursor if they were far away (Ganesh et al., 2014). Ganesh and colleagues had participants complete a target tracking task with a visuomotor rotation introduced in the middle block of trials (2014). Participants alternated between solo and connected trials. Subjects intermittently connected to a partner had less error than those who never interacted with a partner. Takagi and colleagues corroborated the findings from Ganesh and colleagues, finding the same results with a virtual (robotic) partner as well (2017a). Therefore, participants who collaborated with either a human or virtual partner on a non-redundant tracking task had less error in their solo trials than participants who did not interact with a partner.

However, Beckers and colleagues did not replicate these results in a velocity dependent force field learning task (2018). In this study, participants completed 3 blocks consisting of a baseline, force field and washout block, in either a solo group or a connected group. Those in the solo group completed the task without interacting with a partner on any trials, while those in the connected group completed the task haptically connected with a partner

on alternating trials. For the connected group, task performance improved during the interaction trials but not on solo trials. A limitation in the study, however, was that both the perturbation and the interaction forces were haptic, making it difficult for participants to discriminate between the forces. In a follow up study, Beckers and colleagues sought out to rectify this limitation, by replicating the task with visuomotor rotation. Participants completed one baseline block and 3 visuomotor rotation blocks (80 degrees). In contrast to the previously mentioned studies, no difference was found between the solo and connected groups in performance on the solo trials (Ganesh et al., 2014; Reed et al., 2006; Takagi et al., 2017a). Therefore, Beckers and colleagues did not find a benefit to solo performance from a dyadic interaction. In conclusion, although the benefits of joint action on task completion have been established across various studies, there are still inconsistencies as to whether there are benefits to individual performance. Considering that the aim of many cooperative scenarios is to improve individual performance, forming a better understanding of these tasks is important for the design of programs and virtual partners applicable to the real world.

Chapter 2

Introduction

Adaptation is an error-driven form of motor learning that has important implications for the development of evidence-based rehabilitation interventions (Cashaback et al., 2024; Roemmich & Bastian, 2018). Adaptation allows individuals to adjust their movements over time in response to changes in their environment or their own body. This process is thought to be driven by sensory prediction errors, which result from differences between expected and actual sensory consequences (Krakauer et al., 2019; Shadmehr et al., 2010; Wolpert et al., 2011; Wolpert & Flanagan, 2016). Following a neurological disorder or injury, interacting with a device that augments performance or promotes relearning through adaptation can improve the ability to perform functional tasks (Huang & Krakauer, 2009; Reinkensmeyer et al., 2004). These devices range from brain-machine interfaces that help control a cursor on a screen to neurorehabilitation robots that promote recovery (E. Buch et al., 2008; Fasoli et al., 2003; Gulati et al., 2015; López-Larraz et al., 2018; Yap et al., 2017). More recently, researchers have begun to investigate how human performance can be further improved by using adaptive agent algorithms that incorporate human learning rates (e.g., Chasnov et al., 2023; De Santis, 2021; Madduri et al., 2023). Despite these important advances, there remains limited use of adaptive algorithms to augment rehabilitation.

Co-adaptation is an important aspect of interactive behaviours within a rehabilitation setting. For instance, a physiotherapist or a rehabilitation robotic device must consider and adapt to the client's current functional state and recovery rate when providing treatment. Understanding the principles that govern such co-adaptation would provide valuable insights on how two intelligent entities (e.g., human or synthetic agent) interact [*], which can then be leveraged when developing robot-guided neurorehabilitation (Cashaback et al., 2024). Interestingly, several human-human motor tasks have consistently revealed that an individual will display relatively greater (e.g., Batson et al., 2020; Feth et al., 2009; Ganesh et al., 2014; Takagi et al., 2017b) or similar (e.g., Beckers et al., 2020; Che et al., 2016) motor performance in a task after physically interacting with another person compared to those that only ever performed the same task alone. Similar benefits have also been reported when an individual interacts with a robot (i.e., virtual) partner (e.g., Takagi et al., 2017b). However, in these experiments the interactions have often been intermittent and the participants are not typically made aware that they will be interacting with another agent. Despite these positive effects of human-human and human-agent interactions, we still have a limited understanding of the co-adaptation principles that lead to coordinated and seamless interactions between intelligent agents.

A popular paradigm to assess adaptation in a single individual is with a visuomotor rotation task (for reviews see Krakauer, 2009; Krakauer et al., 2019). In these tasks, an individual reaches from a start position to a target. The position of the cursor is rotated relative to the *unseen* hand, about the start position. The individual must learn to counter this rotation (e.g., counterclockwise) by adjusting their reach in the opposite direction (e.g., clockwise). A hallmark of this behaviour is that it is the result of multiple learning processes, each with distinct timescales (S. Kim et al., 2015; Kording et al., 2007; Lee & Schweighofer, 2009; Shadmehr et al., 2010; Smith et al., 2006). In their seminal work, Smith and colleagues (2006) modelled how a single individual adapts using a state-space approach that included both a fast and slow learning process. The fast process learns quickly but forgets easily, whereas the slow process learns more gradually but has better retention. Thus, key assumptions

in Smith and colleagues' (2006) two-state model is that the learning rate is higher for the fast state compared to the slow state and the retention factor is greater for the slow state compared to the fast state. Crucially, this model can account for various motor learning phenomena, including spontaneous recovery (i.e., the reappearance of a previously learned behaviour following extinction trials), anterograde interference (i.e., learning of a new task interferes with the recall of a previously learned task), and savings (i.e., faster relearning of a previously learned task) (Smith et al., 2006).

The goal of the present experiment was to test the effectiveness of adaptive agent algorithms as virtual partners during sensorimotor co-adaptation. Participants completed a visuomotor rotation task that consisted of one baseline block with veridical cursor feedback and one adaptation block where cursor feedback was rotated 30 degrees. Participants in the co-adaptation groups completed the adaptation block where the cursor was jointly controlled by them and an agent partner that embodied either the fast process or the slow process of the two-state model (Smith et al., 2006). Based on model simulations (see Section *A priori* model simulations below), we made the following predictions:

Interacting with the fast state agent results in greater individual early and late adaptation (i.e., human reach angle only). Interacting with the slow state agent results in the human gradually reducing their individual reach contribution over adaptation trials.

Chapter 3

Methods

3.1 Participants

One hundred participants (M_{age} = 20.63, SD = 3.62, 70 females, and 91 right-handed) from the undergraduate and graduate student population at McMaster University volunteered to participate in the experiment. Participants were randomly assigned to one of the Human+Fast Agent (n = 25), Human+Slow Agent (n = 25), Solo Full Rotation (n = 25), or the Solo Half Rotation (n = 25) groups. The Solo Full Rotation and Solo Half Rotation were included as control groups for the Human+Fast Agent and Human+Slow Agent groups, respectively, based on our model simulations. Although sample size was determined based on resource (e.g., time and money) constraints (Lakens, 2022), it is larger than that typically used in motor learning research (median n per group of 11 reported by K. R. Lohse et al., 2016). All participants self-reported that they were free of sensorimotor impairments and had normal or corrected-to-normal vision. Participants were compensated with $10 or with course-credit for their time. All participants provided written informed consent and the experiment was approved by the McMaster Research Ethics Board.

3.2 Task and apparatus

Participants grasped the handle of a robotic manipulandum of a Kinarm End-Point Lab (Kinarm; Kingston, ON, Canada) with their preferred hand and made reaching movements in the horizontal plane (Figure 3.1A). Participants were seated in an adjustable chair at a height such that their forehead could rest comfortably on the system's padded headrest. A semi-silvered mirror blocked vision of both the participant's upper-limb and the robotic manipulandum, and projected images from an LCD screen onto the horizontal plane of motion. A circular cursor (4 mm radius) was displayed to represent the position of the center of the handle that was grasped by the participants. On each trial, participants were presented with a circular start position (10 mm diameter) and a circular target (10 mm diameter) located 15 cm forward of the start position. Participants were instructed to reach and "slice" through the virtually displayed target without vision of their hand or cursor feedback during the movement. Hand position was recorded at 1000 Hz and stored for offline data analysis.

Figure 3.1

Experimental Setup.

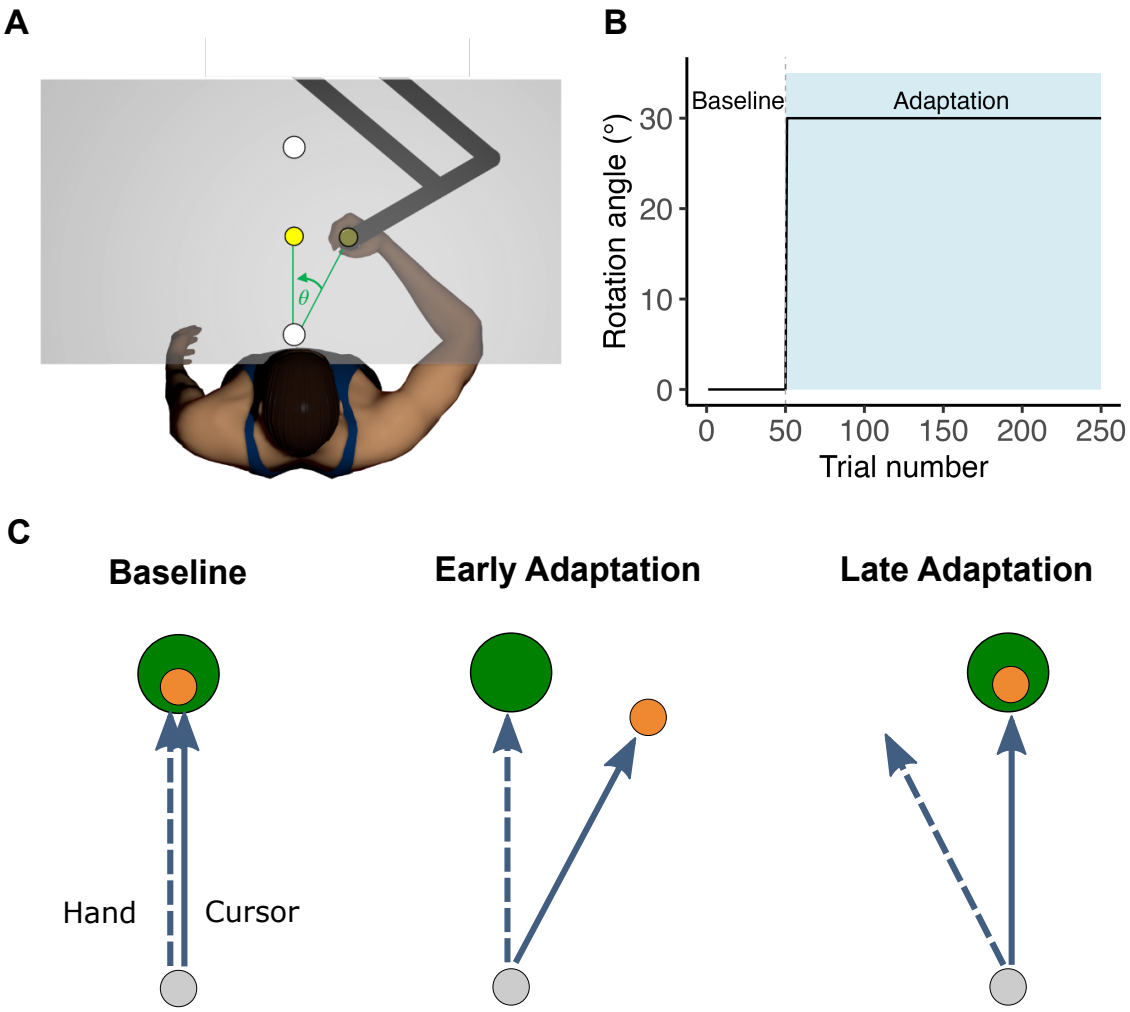

Note. **(A)** Participants grasped the handle of a Kinarm Endpoint robot and controlled a virtual cursor (yellow circle) to hit a target (white circle) positioned in front of them. **(B)** Participants completed one testing session consisting of a baseline phase (50 trials) and an adaptation phase (200 trials). The baseline phase was completed alone, while the adaptation phase was completed according to their experimental group. **(C)** During the baseline phase, the virtual cursor matched the location of the participant's unseen hand. During the adaptation phase, the virtual cursor was rotated 30 degrees about the start position.

3.3 Procedure

Each trial began with the start position (white circle, 10 mm diameter) and target (grey circle, 10 mm diameter) displayed. The robotic manipulandum would then move the participant's hand such that the virtual cursor (orange circle, 4 mm radius) representing their unseen hand was positioned in the center of the start position. After maintaining this position for 500 ms, the grey target turned green and served as the "go" signal for participants to initiate their movement to the target. Participants were instructed to hit the target with their cursor by making a fast reaching movement. To minimize online corrections during the reach, participants were not only instructed to "slice" through the target with the cursor, but also made their reach in the absence of any visual feedback of the cursor. If the movement was initiated (i.e., the cursor had moved fully out of the start position) before the "go" signal, the trial was aborted and a feedback text message "too early" appeared centrally on the screen. If participants did not complete their reach within 750 ms following the "go" signal, a feedback text message "too slow" appeared centrally on the screen. Visual feedback of their endpoint reach error (i.e., the orange cursor) was displayed on the screen for 1 s before the robotic manipulandum returned the participant's hand to the starting position.

Data collection was completed in a single session lasting approximately 35 minutes. This session (Figure 3.1B) consisted of a baseline block (50 trials) and an adaptation block (200 trials). During baseline, all participants performed their reaches alone and received veridical endpoint cursor feedback after each trial (Figure 3.1C). During the adaptation trials, participants performed the task either alone or with an agent partner based on their randomly assigned group. For the Human+Fast Agent and Human+Slow Agent groups, the jointly controlled cursor was rotated 30 degrees clockwise about the start position. That is, the terminal cursor feedback received at the end of each trial represented the sum of the human participant's reach and the agent's contribution. For the solo groups, the solely controlled cursor was rotated either 30 degrees (Solo Full Rotation group) or 15 degrees

(Solo Half Rotation group) about the start position. Participants in the two agent groups were made aware that the endpoint cursor feedback they received after each trial in the adaptation block represented the contribution of their reach plus that of the agent.

To enable an agent that would adapt in a similar manner to a human, we leveraged the following state-space equations to update the aim (x, $\circ$) of an agent on each trial (t) according to fast (f) and slow (s) processes that are governed by their respective learning (B) and retention (A) rates:

$$x_f(t+1) = A_f \cdot x(t) + B_f \cdot e(t) \tag{3a}$$

$$x_s(t+1) = A_s \cdot x(t) + B_s \cdot e(t) \tag{3b}$$

$$x_{net}(t+1) = x_f(t+1) + x_s(t+1) \tag{3c}$$

$$e(n) = T(n) - c^{a_1,h_1}(n) \tag{3d}$$

where, e is the difference between the target (T) and cursor (c) position. Critically, the cursor position is a function of position changes by the human (h_1) and agent (a_1) that enables co-adaptation. In the Human+Fast Agent group, we set B_f, A_f, B_s, A_s to 0.227, 0.597, 0.0, 0.0, promoting fast adaptation and low retention by the agent. In the Human+Slow Agent group, we set B_f, A_f, B_s, A_s to 0.0, 0.0, 0.049, 0.998, promoting slow adaptation and high retention by the agent.

3.4 *A priori* model simulations

We novelly extended upon the state-space approach to capture co-adaptation of human-agent pairs, which has previously only been used to predict how a *single* individual adapts (e.g., Smith et al., 2006). We simulated *a priori* predictions to update the aim (x, $\circ$) of an agent (a) and human (h) on each trial (t). This four-state model accounts for fast (f) and slow (s) processes governed by learning (B) and retention (A) rates according to:

$$x_f^{h_1}(t+1) = A_f^{h_1} \cdot x(t) + B_f^{h_1} \cdot e(t) \tag{4a}$$

$$x_s^{h_1}(t+1) = A_s^{h_1} \cdot x(t) + B_s^{h_1} \cdot e(t) \tag{4b}$$

$$x_f^{a_1}(t+1) = A_f^{a_1} \cdot x(t) + B_f^{a_1} \cdot e(t) \tag{4c}$$

$$x_s^{a_1}(t+1) = A_s^{a_1} \cdot x(t) + B_s^{a_1} \cdot e(t) \tag{4d}$$

$$c^{h_1,a_1}(t+1) = x_f^{h_1}(t+1) + x_s^{h_1}(t+1) + x_f^{a_1}(t+1) + x_s^{a_1}(t+1) \tag{4e}$$

$$e(t) = T(t) - c^{h_1,a_1}(t) \tag{4f}$$

where, e is the difference between the target (T) and cursor (c) position. In the Human+Fast Agent group, we set A_f^a, B_f^a, A_s^a, and B_s^a respectively to 0.597, 0.227, 0.0, and 0.0, which corresponds with the known coefficients in **Eq. 1**. In the Human+Slow Agent group, we set A_f^a, B_f^a, A_s^a, and B_s^a to 0.0, 0.0, 0.998, and 0.049, respectively. To simulate the human, we used previously found coefficients used for a two-state model (Coltman et al., 2019b), where A_f^a, B_f^a, A_s^a, and B_s^a were set to 0.597, 0.227, 0.998, and 0.049, respectively. *A priori* predictions based on the simulations for the Human+Fast Agent and Human+Slow Agent groups are shown in Figure 3.2.

Figure 3.2

Simulation of human and virtual partner contribution to shared cursor.

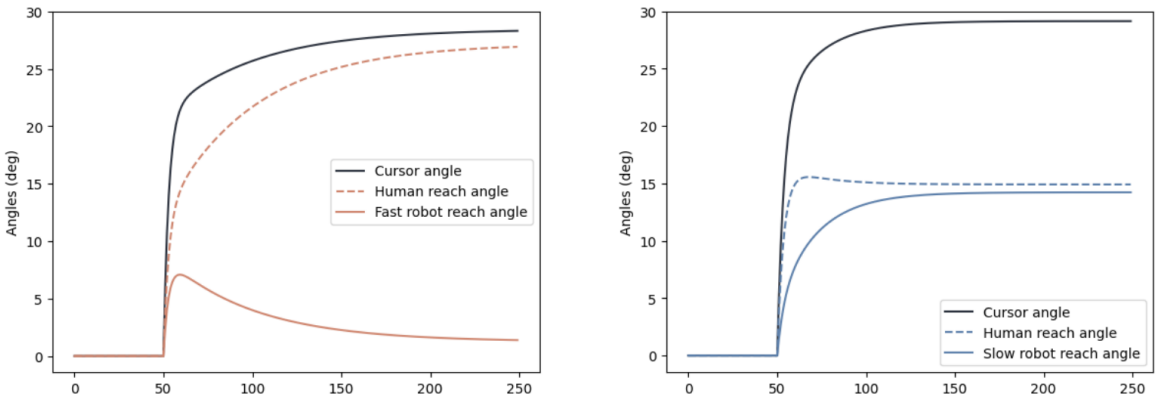

Note. In both graphs, the human angle is shown in the dashed line and the virtual partner angle is shown in the solid line. **(A)** Reach angle is shown for a simulated participant (As = 0.998, Bs = 0.049, Af = 0.597, Bf= 0.227) interacting with the fast agent (As = 0, Bs = 0, Af = 0.597, Bf= 0.227), where the virtual partner will counter more of the rotation in the early trials. Participants interacting with a fast agent will do less of the correction in the beginning then the bulk of it in late adaptation. **(B)** Reach angle is shown for a simulated participant (As = 0.998, Bs = 0.049, Af = 0.597, Bf= 0.227) interacting with the slow agent (As = 0.998, Bs = 0.049, Af = 0, Bf= 0), where the virtual partner will counter less of the rotation in the early trials. Participants interacting with a slow agent will do more of the correction in the beginning then split it equally with the virtual partner.

3.5 Data analysis

To assess motor performance on each trial, we calculated the endpoint hand angle with respect to the target at the moment the cursor reached the target distance of 15 cm from the start position. Prior to statistical analyses, we calculated the 20% trimmed mean human endpoint hand angles for early and late adaptation epochs based on trials 50 to 75 and trials 225 to 250, respectively. Alpha was set to .05 for all statistical tests, which are described below. Post-hoc comparisons were performed using the percentile bootstrap (2000 bootstrap samples) on the 20% trimmed means and we adjusted for multiple comparisons with the Holm-Bonferroni approach.

Chapter 4

Results

Figure 4.1 shows the trial-by-trial endpoint reach angles for the Human + Agent (A and D), the Human only (B and E), and the Agent only (C and F) for a representative participant from the Human+Fast Agent group (left side) and Human+Slow Agent group (right side). Higher reach angles correspond to greater adaptation. Consistent with our *a priori* model predictions (see Figure 3.2), we see that the participant in the Human+Slow Agent group showed faster overall (i.e., Human + Agent) early adaptation compared to the participant in the Human+Fast Agent group. However, both of these participants showed similar overall late adaptation. As seen in our simulations, we see that the participant in the Human+Fast Agent group has a larger individual hand reach angle compared to the participant in the Human+Slow Agent group, who gradually decreased their individual hand reach angle across adaptation trials.

Figure 4.1

Individual participant reach data for the Human + Agent groups.

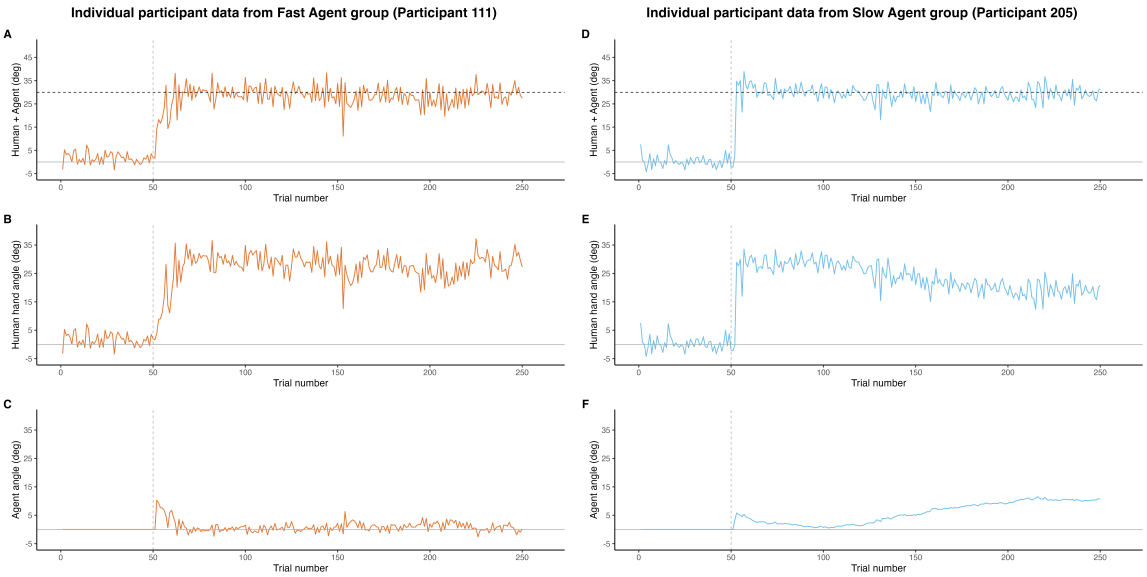

Note. Individual participant data is shown for participant 111 from the fast agent group in orange (left) and participant 205 from the slow agent group in light blue (right). The x-axis on all graphs indicates the trial number. The y-axes are individually described below. The vertical dashed line indicates the separation between the baseline and adaptation blocks on all 3 graphs. (A & D) The sum of the human and agent angle in degrees, is shown for all trials. (B & E) Human hand angle, the angle of the participant's contribution to the displayed cursor in degrees is shown for all trials. Participant 111 (with Fast Agent) shows less contribution to the cursor in the beginning but counters most of the perturbation near the end, while participant 205 (with Slow Agent) shows more contribution in the beginning trials and less towards the end of the adaptation block. (C & F) The agent contribution in degrees to the displayed cursor is shown for all trials.

Trial-by-trial endpoint reach angles as a function of group are shown in Figure 4.2 and the pattern of results for the Human + Agent and the Human only reach angles for the two agent groups are strikingly similar to those of the representative participants. In terms of the Human only reach angle, it can be seen that the Human+Fast Agent group and the Solo Full Rotation group were most similar in late adaptation, but the Solo Full Rotation group adapted their reach angle more quickly in early adaptation. The Human+Slow Agent group and the Solo Half Rotation group were quite similar in late adaptation, but the Human+Slow Agent group adapted their reach angle more quickly in early adaptation; albeit the extent of adaptation was less than that observed in either the Solo Full Rotation group or the Human+Fast Agent group.

Figure 4.2

Trial by trial reaches for all groups.

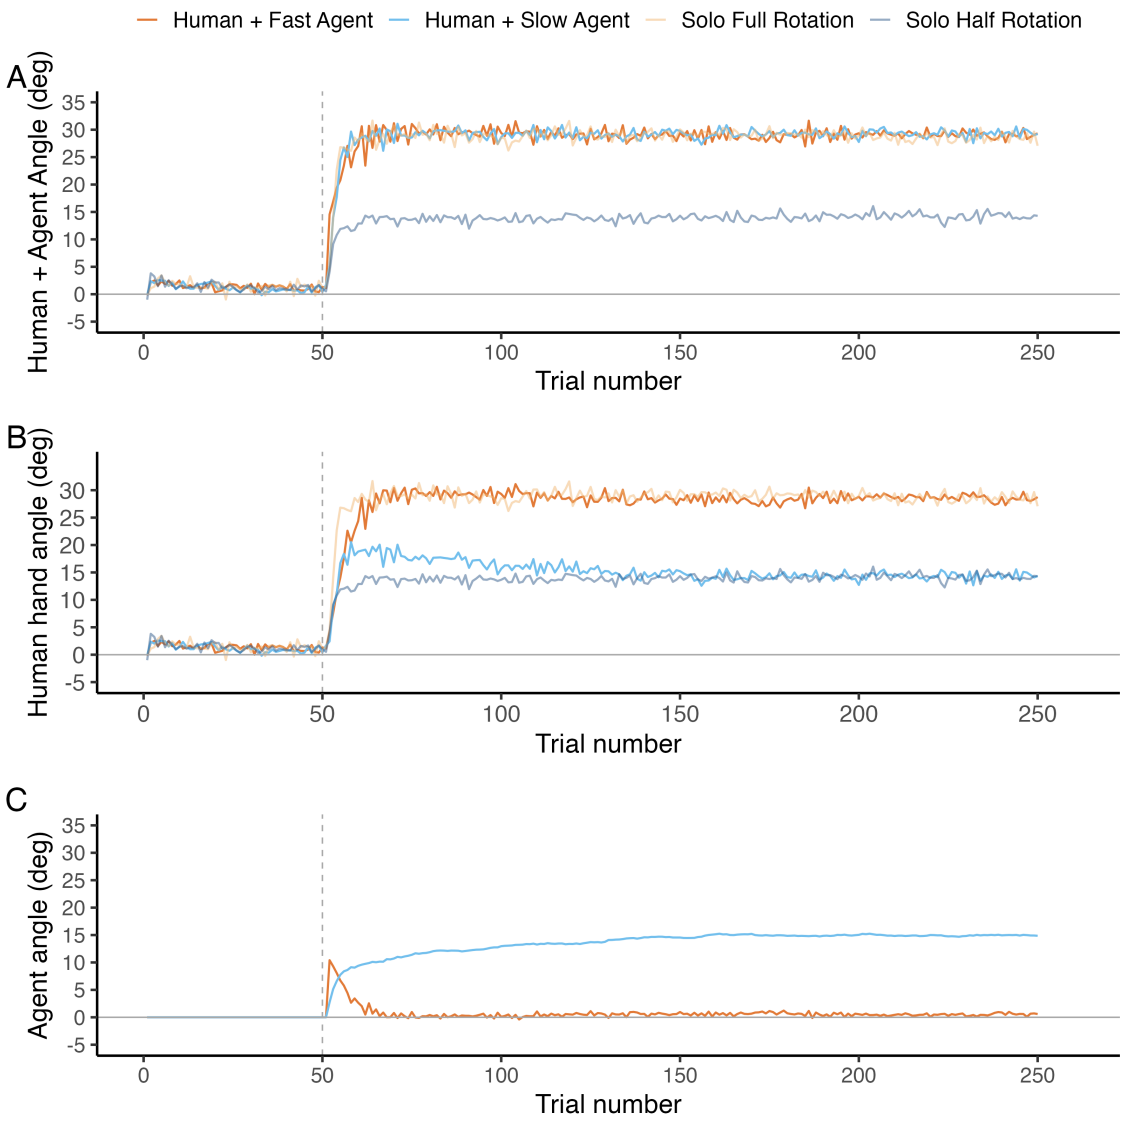

Note. The x-axes on all graphs indicates the trial number. The y-axes are individually described below. The vertical dashed line indicates the separation between the baseline and adaptation blocks on all 3 graphs. Humans with the fast agent are shown in orange, while those working with the slow agent are shown in light blue. The solo full rotation (30 degrees) group is shown in yellow and the solo half rotation (15 degrees) group is shown in dark blue. (A) The sum of the human and agent angle in degrees, (if applicable) is shown as the trimmed mean for each group on all trials. (B) Human reach angle, the angle of the participant's contribution to the displayed cursor in degrees is shown as the trimmed mean for each group across all trials. (C) The agent contribution in degrees to the displayed cursor is shown for each of the agents as the trimmed mean for each of the groups.

To test for any group differences in early and late adaptation (see Figure 4.3), we analyzed human reach angle using a 4 Group (Human+Fast Agent, Human+Slow Agent, Solo Full Rotation, Solo Half Rotation) x 2 Epoch (Early Adaptation, Late Adaptation) robust mixed ANOVA on 20% trimmed means. The significant Group main effect, $F(3, 28.26) = 316.85$, $p < .0001$, and the significant Epoch main effect, $F(1, 36.46) = 17.98$, $p = .0001$, were superseded by a significant Group x Epoch interaction, $F(3, 27.86) = 5.34$, $p = .0049$. Post-hoc analyses revealed that the Human+Fast Agent group had a significantly larger human reach angle compared to the Human+Slow Agent group in both early adaptation (*adj.* $p = .006$) and late adaptation (*adj.* $p < .001$) epochs. The Human+Slow Agent group had a significantly larger human reach angle in early adaptation compared to the Solo Half Rotation group (*adj.* $p = .006$). All groups except for the Human+Slow Agent group had significantly larger human reach angles in late adaptation compared to early adaptation (all *adj.* p's $< .001$).

Figure 4.3

Box plots showing human reach angle in degrees for all groups.

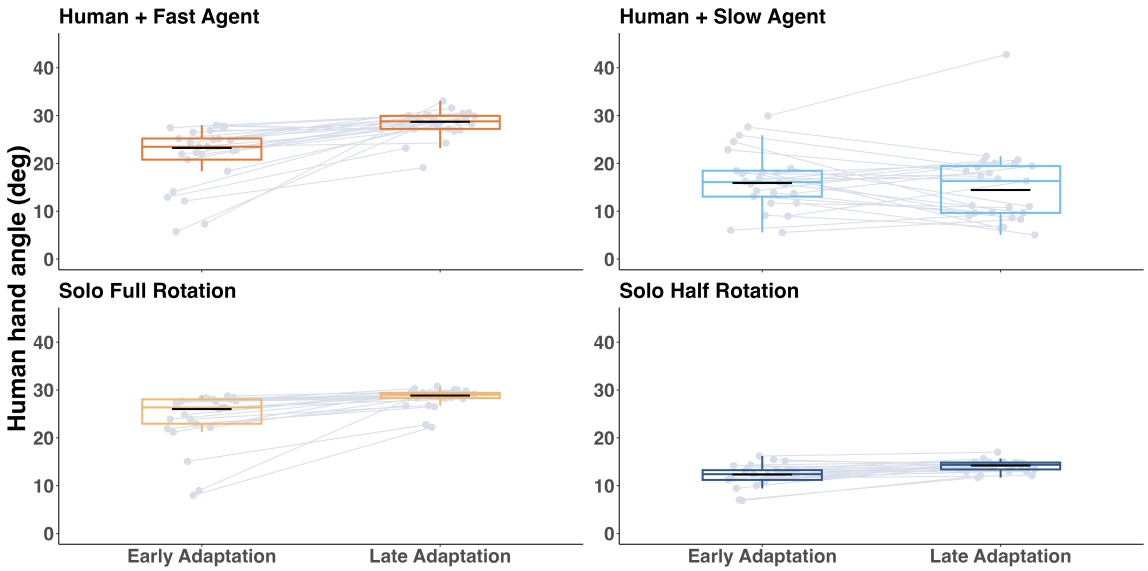

Note. Each of the epochs indicate the mean of 25 trials, with the early epoch being the beginning of the adaptation block and late indicating the end. Each box plot indicates the interquartile range of the data, with the median indicated with the horizontal line on each box and the 20% trimmed mean indicated with the black horizontal line. Each circular point indicates the trimmed mean of the epoch for one individual participant, with lines connecting each participant's behaviour across the two epochs. The human + fast agent group is shown in orange, while human + slow agent are shown in light blue. The solo full rotation (30 degrees) group is shown in yellow and the solo half rotation (15 degrees) group is shown in dark blue.

Given the noticeable difference in the spread of scores between the Human+Fast Agent and Human+Slow Agent groups, we performed an exploratory quantile comparison between the two groups (Mair & Wilcox, 2020; Wilcox et al., 2014) using the quantile estimator proposed by Harrell & Davis (1982). This test provides a more detailed understanding of where and how distributions differ (Mair & Wilcox, 2020; Wilcox, 2023). Separate analyses based on the 25th, 50th, and 75th quantiles were performed for the early and late adaptation epochs and *p*-values were adjusted using Hochberg's method. No significant group differences were found in any of the quantiles for the early adaptation epoch. In the late adaptation epoch, however, the Human+Fast Agent and Human+Slow Agent groups were significantly different at all quantiles (all *adj. p*'s < .001).

Chapter 5

Discussion

The purpose of the present experiment was to investigate the effectiveness of adaptive agent algorithms as virtual partners during sensorimotor co-adaptation in a visuomotor rotation task. In line with our *a priori* model simulations, we found that interacting with the fast state agent led to a greater individual contribution by the human partner when counteracting the visuomotor rotation compared to interacting with the slow state agent in both early and late adaptation epochs. Also consistent with our simulations, we found that interacting with the slow state agent resulted in the human gradually decreasing their reach contribution over trials wherein the human and slow agent were essentially making an equal contribution to counteracting the visuomotor rotation. However, this gradual decrease was modest in size and failed to reach statistical significance. In our exploratory analyses, we did not find that the distribution of scores differed between the Human+Fast Agent and Human+Slow Agent groups in the early adaptation epoch; however, we did find significant group differences at the 25th, 50th, and 75th quantiles in the late adaptation epoch. Overall, our results show that interacting with either a virtual agent that embodies the fast state or slow state of the prominent two-state model (Smith et al., 2006) results in similar overall (i.e., Human + Agent) adaptation. Yet, the human's individual reach contribution is differentially impacted by these virtual agents.

Despite the novelty of our approach to sensorimotor co-adaptation, our results for overall adaptation in these two groups was similar to that of the Solo Full Rotation group (see Figure 3.2A), and thus, the extant motor adaptation literature that has primarily focused on how a single individual adapts to a systematic perturbation (for reviews see Krakauer et al., 2019; Wolpert et al., 2011; Wolpert & Flanagan, 2016). The improved performance that we observed during the adaptation trials when the interaction with the virtual agents was present is consistent with previous sensorimotor interaction research (e.g., Beckers et al., 2018, 2020; Ganesh et al., 2014; Takagi et al., 2017b). In these experiments, target tracking performance was more accurate on trials where participants were intermittently connected to another person through a virtual elastic that permitted them to feel the haptic forces generated by one another. A key difference between these experiments and our experiment is that participants always had sole control over their virtual cursor and were misled about the haptic interaction. Thus, our work extends the benefit of sensorimotor interactions to a redundant reaching task and one where participants are made explicitly aware of the nature of the interaction. Based on the interaction framework forwarded by Jarrassé et al. (2012b), we suggest that the successful co-adaptation of the participants in our virtual agent groups resulted from participants considering both their own and virtual partner's errors and efforts to adjust their reaching behaviours accordingly. Future research could further test this hypothesis by having virtual partners that become more assistive or resistive on a series of probe trials and examine how the human adjusts their reaches across these probe trials.

Our finding that the human only endpoint reach angle (see Figure 3.2B and Figure 4.1) was heavily impacted by one's virtual agent partner has important practical significance. That is, the potential application of these virtual agents may largely depend on the rehabilitative goals of the therapist and/or the client. Given the fairly equal task distribution that resulted in the Human+Slow Agent group, this type of agent algorithm characterized by a high retention rate might be particularly useful in the early stages of a rehabilitation program when a client may be experiencing limited functional skills, as is common following a stroke.

Importantly, this slow state agent helps to ensure that the overall task performance is accurate relative to the task goal, which could be important for motivational reasons linked to task engagement (K. R. Lohse et al., 2016), perceptions of competence (Ryan & Deci, 2000, 2017), and/or task enjoyment (Deci & Ryan, 2012; Ryan & Deci, 2020). Conversely, if the clinical goal is to have an individual maximize their own performance to promote recovery during rehabilitation while still providing some level of initial assistance, then it would appear to be best if a client interacts with a fast state agent that has a low retention factor. These initial results highlight a potentially promising approach to tailor one's training or rehabilitation to their individualized needs (Cashaback et al., 2024). Along these lines, it would be fruitful for future research to investigate the effects of changing the learning rate parameter or the retention parameter, while holding the other one constant, of the fast and slow agents to potentially identify the most suitable algorithms for individuals based on their disease or injury characteristics.

Although our data suggests promising within-session performance gains, a limitation of the current design is we are unable to discern whether these gains are temporary versus relatively permanent (Kantak & Winstein, 2012; e.g., Schmidt & Bjork, 1992). To address this limitation, we are currently investigating the effects of these same fast and slow state agents in a visuomotor paradigm that includes a counter rotation block, an error clamp block, and an identical second session following a period of no-practice. These additional phases will respectively provide important insights regarding spontaneous recovery, retention, and savings during co-adaptation. Another potential limitation is that we relied on a visual perturbation whereas learning novel task dynamics via a mechanical perturbation (e.g., velocity-dependent force field) may have greater "out-of-the-lab" validity and applicability. However, a key benefit of using a visual perturbation is that it allowed us to isolate behavioural changes based on the joint goal of hitting the target with the shared virtual cursor. Lastly, given the "proof-of-concept" and exploratory nature of this project, our experiment and analysis plan was not pre-registered. While we have reported our results in a transparent and open manner and generated our predictions based on *a*

priori model simulations, there have been several recent calls within motor learning (e.g., K. Lohse et al., 2016; McKay et al., 2023) for greater use of pre-registration to help combat various forms of questionable research practices (e.g., Munafò et al., 2017).

Chapter 6

Conclusion

In the present experiment, we investigated the effects of knowingly completing a task with a virtual agent partner on co-adaptation in a visuomotor rotation task. Participants were paired with one of two virtual partners based on the fast process or slow process of Smith and colleagues' (2006) two-state model of adaptation. Although our *a priori* model simulations indicated that both agent groups should learn to successfully counteract the visuomotor rotation, the simulations also revealed that the relative contributions of the human partner in these groups should differ. This pattern of results were found not only at the individual participant level, but also at the group level. Thus, this novel co-adaptation paradigm may be a useful approach to robot-guided neurorehabilitation; however, further studies are still needed to identify the principles underlying sensorimotor co-adaptation.

References

Batson, J. P., Kato, Y., Shuster, K., Patton, J. L., Reed, K. B., Tsuji, T., & Novak, D. (2020). Haptic coupling in dyads improves motor learning in a simple force field. *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 4795–4798.

Beckers, N., Asseldonk, E. H. van, & Kooij, H. van der. (2020). Haptic human–human interaction does not improve individual visuomotor adaptation. *Scientific Reports*, *10*(1), 19902.

Beckers, N., Keemink, A., Asseldonk, E. van, & Kooij, H. van der. (2018). Haptic human-human interaction through a compliant connection does not improve motor learning in a force field. *Haptics: Science, Technology, and Applications: 11th International Conference, EuroHaptics 2018, Pisa, Italy, June 13-16, 2018, Proceedings, Part i 11*, 333–344.

Buch, E. R., Young, S., & Contreras-Vidal, J. L. (2003). Visuomotor Adaptation in Normal Aging. *Learning & Memory*, *10*(1), 55–63. https://doi.org/10.1101/lm.50303

Buch, E., Weber, C., Cohen, L. G., Braun, C., Dimyan, M. A., Ard, T., Mellinger, J., Caria, A., Soekadar, S., Fourkas, A., et al. (2008). Think to move: A neuromagnetic brain-computer interface (BCI) system for chronic stroke. *Stroke*, *39*(3), 910–917.

Cashaback, J. G., Allen, J. L., Chou, A. H.-Y., Lin, D. J., Price, M. A., Secerovic, N. K., Song, S., Zhang, H., & Miller, H. L. (2024). NSF DARE—transforming modeling in neurorehabilitation: A patient-in-the-loop framework. *Journal of NeuroEngineering and Rehabilitation*, *21*(1), 23.

Chasnov, B. J., Ratliff, L. J., & Burden, S. A. (2023). Human adaptation to adaptive machines

converges to game-theoretic equilibria. *arXiv Preprint arXiv:2305.01124*.

Che, Y., Haro, G. M., & Okamura, A. M. (2016). Two is not always better than one: Effects of teleoperation and haptic coupling. *2016 6th IEEE International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, 1290–1295.

Coltman, S. K., Cashaback, J. G. A., & Gribble, P. L. (2019a). Both fast and slow learning processes contribute to savings following sensorimotor adaptation. *Journal of Neurophysiology*, *121*(4), 1575–1583. https://doi.org/10.1152/jn.00794.2018

Coltman, S. K., Cashaback, J. G., & Gribble, P. L. (2019b). Both fast and slow learning processes contribute to savings following sensorimotor adaptation. *Journal of Neurophysiology*, *121*(4), 1575–1583.

Crevecoeur, F., Thonnard, J.-L., & Lefèvre, P. (2020). A Very Fast Time Scale of Human Motor Adaptation: Within Movement Adjustments of Internal Representations during Reaching. *eNeuro*, *7*(1), ENEURO.0149–19.2019. https://doi.org/10.1523/ENEURO.0149-19.2019

Criscimagna-Hemminger, S. E., & Shadmehr, R. (2008). Consolidation Patterns of Human Motor Memory. *Journal of Neuroscience*, *28*(39), 9610–9618. https://doi.org/10.1523/JNEUROSCI.3071-08.2008

De Santis, D. (2021). A framework for optimizing co-adaptation in body-machine interfaces. *Frontiers in Neurorobotics*, *15*, 662181.

Deci, E. L., & Ryan, R. M. (2012). Self-determination theory. In *Handbook of theories of social psychology* (pp. 416–436). Sage Publications Ltd.

Fasoli, S. E., Krebs, H. I., Stein, J., Frontera, W. R., & Hogan, N. (2003). Effects of robotic therapy on motor impairment and recovery in chronic stroke. *Archives of Physical Medicine and Rehabilitation*, *84*(4), 477–482.

Fernández-Ruiz, J., & Díaz, R. (1999). Prism Adaptation and Aftereffect: Specifying the Properties of a Procedural Memory System. *Learning & Memory*, *6*(1), 47–53. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC311278/

Feth, D., Groten, R., Peer, A., Hirche, S., & Buss, M. (2009). Performance related energy

exchange in haptic human-human interaction in a shared virtual object manipulation task. *World Haptics 2009-Third Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 338–343.

Ganesh, G., Takagi, A., Osu, R., Yoshioka, T., Kawato, M., & Burdet, E. (2014). Two is better than one: Physical interactions improve motor performance in humans. *Scientific Reports*, *4*(1), 1–7.

Goodwin, J. E., Grimes, C. R., Eckerson, J. M., & Gordon, P. M. (1998). Effect of different quantities of variable practice on acquisition, retention, and transfer of an applied motor skill. *Perceptual and Motor Skills*, *87*(1), 147–151. https://doi.org/10.2466/pms.1998.87.1.147

Guadagnoli, M. A., & Kohl, R. M. (2001). Knowledge of Results for Motor Learning: Relationship Between Error Estimation and Knowledge of Results Frequency. *Journal of Motor Behavior*. https://doi.org/10.1080/00222890109603152

Gulati, T., Won, S. J., Ramanathan, D. S., Wong, C. C., Bodepudi, A., Swanson, R. A., & Ganguly, K. (2015). Robust neuroprosthetic control from the stroke perilesional cortex. *Journal of Neuroscience*, *35*(22), 8653–8661.

Haith, A. M., Huberdeau, D. M., & Krakauer, J. W. (2015). The Influence of Movement Preparation Time on the Expression of Visuomotor Learning and Savings. *Journal of Neuroscience*, *35*(13), 5109–5117. https://doi.org/10.1523/JNEUROSCI.3869-14.2015

Hamilton, C. R. (1964). Intermanual Transfer of Adaptation to Prisms. *The American Journal of Psychology*, *77*(3), 457–462. https://doi.org/10.2307/1421017

Harrell, F. E., & Davis, C. (1982). A new distribution-free quantile estimator. *Biometrika*, *69*(3), 635–640.

Huang, V. S., & Krakauer, J. W. (2009). Robotic neurorehabilitation: A computational motor learning perspective. *Journal of Neuroengineering and Rehabilitation*, *6*, 1–13.

Huberdeau, D. M., Haith, A. M., & Krakauer, J. W. (2015). Formation of a long-term memory for visuomotor adaptation following only a few trials of practice. *Journal of Neurophysiology*, *114*(2), 969–977. https://doi.org/10.1152/jn.00369.2015

Iwamuro, B. T., Cruz, E. G., Connelly, L. L., Fischer, H. C., & Kamper, D. G. (2008). Effect of a gravity-compensating orthosis on reaching after stroke: Evaluation of the Therapy Assistant WREX. *Archives of Physical Medicine and Rehabilitation*, *89*(11), 2121–2128. https://doi.org/10.1016/j.apmr.2008.04.022

Jarrassé, N., Charalambous, T., & Burdet, E. (2012a). A Framework to Describe, Analyze and Generate Interactive Motor Behaviors. *PLOS ONE*, *7*(11), e49945. https://doi.org/10.1371/journal.pone.0049945

Jarrassé, N., Charalambous, T., & Burdet, E. (2012b). A framework to describe, analyze and generate interactive motor behaviors. *PloS One*, *7*(11), e49945.

Kantak, S. S., & Winstein, C. J. (2012). Learning–performance distinction and memory processes for motor skills: A focused review and perspective. *Behavioural Brain Research*, *228*(1), 219–231.

Kim, H. E., Avraham, G., & Ivry, R. B. (2021). The Psychology of Reaching: Action Selection, Movement Implementation, and Sensorimotor Learning. *Annual Review of Psychology*, *72*(1), 61–95. https://doi.org/10.1146/annurev-psych-010419-051053

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

Kojima, Y., Iwamoto, Y., & Yoshida, K. (2004). Memory of Learning Facilitates Saccadic Adaptation in the Monkey. *Journal of Neuroscience*, *24*(34), 7531–7539. https://doi.org/10.1523/JNEUROSCI.1741-04.2004

Kording, K. P., Tenenbaum, J. B., & Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience*, *10*(6), 779–786.

Krakauer, J. W. (2009). Motor learning and consolidation: The case of visuomotor rotation. *Progress in Motor Control: A Multidisciplinary Perspective*, 405–421.

Krakauer, J. W., Hadjiosif, A. M., Xu, J., Wong, A. L., & Haith, A. M. (2019). Motor learning. *Compr Physiol*, *9*(2), 613–663.

Krakauer, J. W., Pine, Z. M., Ghilardi, M.-F., & Ghez, C. (2000). Learning of Visuomotor Transformations for Vectorial Planning of Reaching Trajectories. *Journal of Neuroscience*, *20*(23), 8916–8924. https://doi.org/10.1523/JNEUROSCI.20-23-08916.2000

Lakens, D. (2022). Sample size justification. *Collabra: Psychology*, *8*(1), 33267.

Lee, J.-Y., & Schweighofer, N. (2009). Dual adaptation supports a parallel architecture of motor memory. *Journal of Neuroscience*, *29*(33), 10396–10404.

Lohse, K. R., Boyd, L. A., & Hodges, N. J. (2016). Engaging environments enhance motor skill learning in a computer gaming task. *Journal of Motor Behavior*, *48*(2), 172–182.

Lohse, K., Buchanan, T., & Miller, M. (2016). Underpowered and overworked: Problems with data analysis in motor learning studies. *Journal of Motor Learning and Development*, *4*(1), 37–58.

López-Larraz, E., Sarasola-Sanz, A., Irastorza-Landa, N., Birbaumer, N., & Ramos-Murguialday, A. (2018). Brain-machine interfaces for rehabilitation in stroke: A review. *NeuroRehabilitation*, *43*(1), 77–97.

Mackay, S., Morgan, P., Datta, V., Chang, A., & Darzi, A. (2002). Practice distribution in procedural skills training. *Surgical Endoscopy And Other Interventional Techniques*, *16*(6), 957–961. https://doi.org/10.1007/s00464-001-9132-4

Madduri, M. M., Burden, S. A., & Orsborn, A. L. (2023). Biosignal-based co-adaptive user-machine interfaces for motor control. *Current Opinion in Biomedical Engineering*, *27*, 100462.

Mair, P., & Wilcox, R. (2020). Robust statistical methods in r using the WRS2 package. *Behavior Research Methods*, *52*, 464–488.

Mawase, F., Haizler, T., Bar-Haim, S., & Karniel, A. (2013). Kinetic adaptation during locomotion on a split-belt treadmill. *Journal of Neurophysiology*, *109*(8), 2216–2227. https://doi.org/10.1152/jn.00938.2012

McDougle, S. D., Bond, K. M., & Taylor, J. A. (2015). Explicit and Implicit Processes Constitute the Fast and Slow Processes of Sensorimotor Learning. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *35*(26), 9568–9579.

https://doi.org/10.1523/JNEUROSCI.5061-14.2015

McKay, B., Bacelar, M. F., Parma, J. O., Miller, M. W., & Carter, M. J. (2023). The combination of reporting bias and underpowered study designs has substantially exaggerated the motor learning benefits of self-controlled practice and enhanced expectancies: A meta-analysis. *International Review of Sport and Exercise Psychology*, 1–21.

Miall, R. C., & Wolpert, D. M. (1996). Forward Models for Physiological Motor Control. *Neural Networks*, *9*(8), 1265–1279. https://doi.org/10.1016/S0893-6080(96)00035-4

Moore, R. T., & Cluff, T. (2021). Individual Differences in Sensorimotor Adaptation Are Conserved Over Time and Across Force-Field Tasks. *Frontiers in Human Neuroscience*, *15*. https://www.frontiersin.org/articles/10.3389/fnhum.2021.692181

Morehead, J. R., Qasim, S. E., Crossley, M. J., & Ivry, R. (2015). Savings upon Re-Aiming in Visuomotor Adaptation. *Journal of Neuroscience*, *35*(42), 14386–14396. https://doi.org/10.1523/JNEUROSCI.1046-15.2015

Munafò, M. R., Nosek, B. A., Bishop, D. V., Button, K. S., Chambers, C. D., Percie du Sert, N., Simonsohn, U., Wagenmakers, E.-J., Ware, J. J., & Ioannidis, J. (2017). A manifesto for reproducible science. *Nature Human Behaviour*, *1*(1), 1–9.

Neva, J. L., & Henriques, D. Y. P. (2013). Visuomotor adaptation and generalization with repeated and varied training. *Experimental Brain Research*, *226*(3), 363–372. https://doi.org/10.1007/s00221-013-3444-1

Patterson, T. A., Thomas, L., Wilcox, C., Ovaskainen, O., & Matthiopoulos, J. (2008). State–space models of individual animal movement. *Trends in Ecology & Evolution*, *23*(2), 87–94. https://doi.org/10.1016/j.tree.2007.10.009

Peña-Pérez, N., Eden, J., Ivanova, E., Farkhatdinov, I., & Burdet, E. (2023). How virtual and mechanical coupling impact bimanual tracking. *Journal of Neurophysiology*, *129*(1), 102–114. https://doi.org/10.1152/jn.00057.2022

Reed, K., Peshkin, M., Hartmann, M. J., Grabowecky, M., Patton, J., & Vishton, P. M. (2006). Haptically Linked Dyads: Are Two Motor-Control Systems Better Than One? *Psychological Science*, *17*(5), 365–366. https://doi.org/10.1111/j.1467-9280.2006.01712.x

Reichenbach, A., Costello, A., Zatka-Haas, P., & Diedrichsen, J. (2013). Mechanisms of responsibility assignment during redundant reaching movements. *Journal of Neurophysiology*, *109*(8), 2021–2028. https://doi.org/10.1152/jn.01052.2012

Reinkensmeyer, D. J., Emken, J. L., & Cramer, S. C. (2004). Robotics, motor learning, and neurologic recovery. *Annu. Rev. Biomed. Eng.*, *6*(1), 497–525.

Reisman, D. S., Block, H. J., & Bastian, A. J. (2005). Interlimb Coordination During Locomotion: What Can be Adapted and Stored? *Journal of Neurophysiology*, *94*(4), 2403–2415. https://doi.org/10.1152/jn.00089.2005

Reisman, D. S., Wityk, R., Silver, K., & Bastian, A. J. (2009). Split-Belt Treadmill Adaptation Transfers to Overground Walking in Persons Poststroke. *Neurorehabilitation and Neural Repair*, *23*(7), 735–744. https://doi.org/10.1177/1545968309332880

Roemmich, R. T., & Bastian, A. J. (2018). Closing the loop: From motor neuroscience to neurorehabilitation. *Annual Review of Neuroscience*, *41*(1), 415–429.

Ryan, R. M., & Deci, E. L. (2000). Self-Determination Theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, *55*(1), 68–78. https://doi.org/10.1037110003-066X.55.1.68

Ryan, R. M., & Deci, E. L. (2017). *Self-determination theory: Basic psychological needs in motivation, development, and wellness*. Guilford Press.

Ryan, R. M., & Deci, E. L. (2020). Intrinsic and extrinsic motivation from a self-determination theory perspective: Definitions, theory, practices, and future directions. *Contemporary Educational Psychology*, *61*, 101860. https://doi.org/10.1016/j.cedpsych.2020.101860

Sakamoto, T., & Kondo, T. (2012). Can passive arm movement affect adaptation to visuomotor rotation? *2012 IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL)*, 1–6. https://doi.org/10.1109/DevLrn.2012.6400824

Sawers, A., & Ting, L. H. (2014). Perspectives on human-human sensorimotor interactions for the design of rehabilitation robots. *Journal of NeuroEngineering and Rehabilitation*, *11*(1), 142. https://doi.org/10.1186/1743-0003-11-142

Schmidt, R. A., & Bjork, R. A. (1992). New conceptualizations of practice: Common

principles in three paradigms suggest new concepts for training. *Psychological Science*, *3*(4), 207–218.

Schmidt, R. A., Lee, T. D., Winstein, C. J., Wulf, G., & Zelaznik, H. N. (2019). Conditions of Practice. In *Motor Control and Learning: A Behavioural Emphasis* (Sixth Edition).

Sebanz, N., Bekkering, H., & Knoblich, G. (2006). Joint action: Bodies and minds moving together. *Trends in Cognitive Sciences*, *10*(2), 70–76. https://doi.org/10.1016/j.tics.2005.12.009

Shadmehr, R., & Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience*, *14*(5), 3208–3224. https://doi.org/10.1523/JNEUROSCI.14-05-03208.1994

Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience*, *33*(1), 89–108.

Smith, M. A., Ghazizadeh, A., & Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biology*, *4*(6), e179.

Sternad, D. (2018). It's not (only) the mean that matters: Variability, noise and exploration in skill learning. *Current Opinion in Behavioral Sciences*, *20*, 183–195. https://doi.org/10.1016/j.cobeha.2018.01.004

Takagi, A., Ganesh, G., Yoshioka, T., Kawato, M., & Burdet, E. (2017b). Physically interacting individuals estimate the partner's goal to enhance their movements. *Nature Human Behaviour*, *1*(3, 3), 1–6. https://doi.org/10.1038/s41562-017-0054

Takagi, A., Ganesh, G., Yoshioka, T., Kawato, M., & Burdet, E. (2017a). Physically interacting individuals estimate the partner's goal to enhance their movements. *Nature Human Behaviour*, *1*(3), 0054.

von Helmholtz, H. (1924). *Helmholtz's treatise on physiological optics, Vol. 1, Trans. From the 3rd German ed* (J. P. C. Southall, Ed.; pp. xxi, 482). Optical Society of America. https://doi.org/10.1037/13536-000

Wegner, N., & Zeaman, D. (1956). Team and Individual Performances on a Motor Learning Task. *The Journal of General Psychology*, *55*(1), 127–142. https://doi.org/10.1080/002213

09.1956.9920301

Welch, R. B. (1971). Discriminative conditioning of prism adaptation. *Perception & Psychophysics*, *10*(2), 90–92. https://doi.org/10.3758/BF03214321

White, O., & Diedrichsen, J. (2010). Responsibility assignment in redundant systems. *Current Biology: CB*, *20*(14), 1290–1295. https://doi.org/10.1016/j.cub.2010.05.069

Wilcox, R. R. (2023). *A guide to robust statistical methods*. Springer.

Wilcox, R. R., Erceg-Hurn, D. M., Clark, F., & Carlson, M. (2014). Comparing two independent groups via the lower and upper quantiles. *Journal of Statistical Computation and Simulation*, *84*(7), 1543–1551.

Wolpert, D. M., Diedrichsen, J., & Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nature Reviews Neuroscience*, *12*(12), 739–751.

Wolpert, D. M., & Flanagan, J. R. (2016). Computations underlying sensorimotor learning. *Current Opinion in Neurobiology*, *37*, 7–11.

Yap, H. K., Lim, J. H., Nasrallah, F., & Yeow, C.-H. (2017). Design and preliminary feasibility study of a soft robotic glove for hand function assistance in stroke survivors. *Frontiers in Neuroscience*, *11*, 547.