JOINT DYADIC PRACTICE STRATEGIES WHEN LEARNING A NOVEL BODY-MACHINE INTERFACE TASK

By

Brian Fox

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Kinesiology – Doctor of Philosophy

2024

ABSTRACT

Body-Machine Interfaces can restore motion and independence to individuals with movement impairments, but often require extensive time and effort to learn due to complexities in their control scheme. Dyadic practice can help to mitigate this difficulty by allowing the learner to practice the task with a partner. Here, we examined the question of "joint dyadic" practice where the learner practices the task simultaneously with an expert partner, where the control that the expert and learner have over the task could be altered. We collected five groups of participants to explore the effectiveness of various implementations of dyadic practice: one solo group (n = 20) and four dyad groups, in which participants were paired with an expert who had prior practice on the task. In our first study, we explored how observation of one's partner benefits the learning of a cursor control task by evaluating the solo group, a constant control dyad, in which the novice and the expert shared equal control of the task and were seated in view of one another (n = 16), and a visually-separated dyad, identical to the constant control dyad, except individuals in the dyad were visually separated from their partner (n = 8). In our second study, we explored different methods of allocating control to the learner during dyadic practice by examining a gradual control dyad, in which control of the task was gradually given to the novice as training progresses (n = 12), and an adaptive control dyad, in which control of the task was given to the novice based on their prior performance (n = 12). While all groups were able to learn the task, none of the dyad groups were able to achieve the level of performance reached by the solo group. These findings provide insight on how to structure dyadic practice to encourage learning of a Body-Machine Interface task, as well as how dyadic practice may benefit more complex high degree of freedom tasks.

This dissertation is dedicated to my brother, Nicholas Raymond Bassett, and my father, Kevin John Fox: two great men whose lives inspired this research.

ACKNOWLEDGEMENTS

I would like to first thank my wife, Katherine Fox, for her unwavering love and support throughout my entire graduate school career, and for putting up with the long hours and weekends that this dissertation required. I would also like to thank my mother, Julie Fox, for all the years she spent raising me and preparing me for my PhD career, for all her kind words and support, and for bragging about me to every stranger who would listen. Thanks go out to my best friends Chris, Luke, Melissa, Kyle, Jon, Rogo, and Julieanne, for the countless D&D sessions, road trips, weddings, get-togethers, and late-night heart-to-hearts over voice chat that got me through these last four years. To my second family Ann, Vick, and Heather, thank you for being a second home for me all throughout middle school, high school, college and beyond, and for caring for me like I was one of your own. And to my wife's family Ned, Melinda, Lindsay, and Barrett, thank you for being so welcoming, generous, and thoughtful as I joined your family in the throes of graduate school. I love each and every one of you more than words can express.

I extend my thanks to my advisor, Dr Rajiv Ranganathan, for being a constant source of insight and knowledge throughout my four years at Michigan State, and for his kindness, empathy, and understanding when it mattered most. I thank my dissertation committee, Drs Mei-Hua Lee, Matthew Harkey, and Ranjan Mukherjee, for providing me with opportunities to become a better scientist through discussion and collaboration. Thanks to all my colleagues in the MOTRE/SD Labs: Jennifer, Simon, Narae, and Joey for your support and encouragement over the years. And to my undergrads, the rest of my graduate cohort, and professors in Michigan State's Department of Kinesiology, I will always value our work sessions, coffee talks, and hallway discussions. Go Green!

TABLE OF CONTENTS

Chapter 1 – Introduction

Section 1.1. Human-Machine Interfaces for controlling complex machines and assistive devices

Human-Machine Interfaces (HMIs) encompass a large array of technologies that allow humans to control or operate some type of machine. From common examples such as a computer mouse or a joystick that are used on a daily basis, to highly specialized ones such as robotic prosthetic arms and exoskeletons, HMIs allow humans to interface with machines to perform several tasks. HMIs function, in a very broad sense, by capturing signals associated with movement (whether that be the physiological signals that control the movement, or the characteristics of the movement itself) and translating that into control of the machine so the operator can successfully complete the task. These HMIs also have clinical significance in terms of assistive devices. For individuals with movement impairments, such as stroke survivors, amputees, individuals with spinal cord injuries, or persons with congenital limb loss or reduction, HMIs can restore a sense of independence and aid in their activities of daily living by allowing them to interact with the environment using assistive devices, such as a robotic prosthetic limb, on-screen cursor, or powered wheelchair. Given that more than one in four people living in the United States are living with some type of disability as of 2021 (Centers for Disease Control and Prevention, accessed 5-6-2024), that surveys among amputees list "ease of use" and/or "operability" of the prosthetic as one of the key factors influencing device abandonment (Biddiss & Chau, 2007), and that there is tremendous variation in the types of impairments from disabilities, there is a critical need to develop an understanding of how humans learn to control these interfaces efficiently.

Section 1.2. The challenge of learning to control HMIs

A key design challenge for HMIs is that the control should be intuitive - it should be relatively clear to an untrained individual what is required of them to operate the device successfully. This is typically the case in everyday devices or technologies: when driving a car, it makes sense for us to turn the steering wheel to the left if we desire the car to turn left; when using a mouse and keyboard, it makes sense that moving the mouse upward on the desk plane causes the cursor to move upward on the screen plane (interestingly, this does not come easily to kids, so one could argue it is not as intuitive as you might think!). The point of this intuitive design is that, given how regular and essential many of these tasks are for activities of daily living or work, it minimizes training time and allows these processes to be performed while minimizing cognitive effort, freeing up capacity for more complex tasks (imagine how difficult it would be to operate a vehicle at interstate speeds if you had to constantly be thinking about how to make the car turn left every time you had to merge).

Addressing the “intuitiveness” challenge becomes even more critical when applied to assistive devices, as we now have to consider (i) that the devices may have more degrees of freedom (DoFs) which require more control commands, and (ii) how the movement abilities of each individual affect their ability to control the device. In the instance of an individual with a transhumeral robotic prosthetic, the individual may need to reach and grasp an object in front of them, but lack the musculature typically used to open and close their hand. Thus, they must now find some alternative way to signal the prosthetic to open and close. Moreover, because movement abilities may differ across individuals, not only is control of these interfaces no longer intuitive, but relying solely on any single method of controlling the device may not work for all. For example, while one stroke survivor may be able to elevate and depress their affected arm, another may not, so an assistive device that relies on elevation and depression of the arm would not work for the latter individual. Therefore, a good design should account for the user's unique movement capabilities when developing the control scheme, which may require learning of "non-intuitive" ways to control the interface.

Section 1.3. BoMIs as a context to study learning

Here, we examine this challenge of learning the control of HMIs specifically in the context of a type of HMI called the Body-Machine Interface (BoMI). In contrast to Brain-Machine Interfaces in which control of devices is based on neural signals from the brain (recorded either invasively or non-invasively), a BoMI collects movement signals from the body and transforms them into control commands for the device. These signals could therefore measure variables related to an individual's kinematics, kinetics, or muscle activity (Casadio, Ranganathan, & Mussa-Ivaldi, 2012). A key advantage of BoMIs over Brain-Machine Interfaces is that they are non-invasive, robust, cheap, and have high signal-to-noise ratios, which can be useful in a wide range of motor impairments where individuals have the ability to control some residual movements.

The intuitiveness of a BoMI depends on the "map" that transforms the movement signals into device commands. This map can be used in different ways depending on the context. For assistive device contexts, the map can be customized to harness an individual's existing movement capabilities, and in the event that this individual's capabilities change over time, the map can be easily modified to account for this change in ability. For rehabilitation contexts, the map can instead be customized so that it encourages the user to use less-preferred movements that are currently "outside" their existing movement repertoire, which has been shown to improve movement capabilities/function in this population (Pierella et al., 2015). Since this map can be controlled by the experimenter which can be used to make the task difficult even for unimpaired participants, BoMIs provide an ideal context to understand the learning challenge.

Section 1.4. Providing learning assistance when learning BoMIs

What kind of training schedules can optimize learning of such non-intuitive interfaces? There are typically two distinct approaches – (i) no guidance – i.e., trial-and-error exploration, and (ii) guided practice from an expert. In no guidance conditions, the person performs a movement pattern or technique to observe its outcome and determine if the performed action contributes to successful task performance, or if that movement should be avoided in future attempts. The advantage of the no-guidance approach is that it forces the participant to learn on their own through trial-and-error. However, exploring the task in this way can be time-consuming, as the individual could theoretically explore several suboptimal options of completing the task before arriving at the desired movement pattern. Relatedly, in instances where the interface is not immediately intuitive or contains many manipulable DoFs, this process could be difficult, frustrating and/or discouraging for the individual.

Guidance from an expert aims to expedite exploration by having the individual practice the task while under the supervision of a trained expert who understands the task and how to perform it. The advantage of this method is that the expert can provide the individual with demonstrations and/or feedback, thus potentially guiding the learner quickly toward the solution. In a task with many manipulable DoFs, it can rule out ineffective or inefficient movement patterns early, allowing the learner to focus more attention on more efficient coordination patterns. However, a disadvantage of guided practice is that the feedback provided by the expert can create “information-overload” especially in cases where the task is complex. In addition, excessive guidance could make the learner “dependent” on the expert such that performance drops when the learner has to eventually do the task on their own (known as the “guidance hypothesis”; Salmoni, Schmidt, & Walter, 1984).

As the goal for these tasks is to ultimately perform them without external assistance, engaging in the learning process in such a way that causes the novice to become reliant on the presence/feedback of an expert is less than ideal. Instead, it may be more beneficial to “fade” the amount of feedback given to the novice as their training on the task goes on, in which this feedback is given more frequently during the early stages of learning, and then gradually decreases as learning progresses. So how can practice be structured in such a way to encourage the novice to learn the task without overwhelming them or causing them to become reliant on their partner?

One approach that has the potential for combining the advantages of the methods is “dyadic” practice, in which two individuals work together to learn a task that is intended to be performed on an individual basis (Crook & Beier, 2010). In dyadic practice, the learner interacts with a partner either sequentially (through a combination of observation and performance), or jointly performing the task simultaneously with the partner. Dyadic practice provides individuals with an opportunity to engage in observational learning (in which they observe successful completion of the task before performing it themselves), or to share simultaneous control of the task with another user in order to reduce cognitive effort and error magnitude (often referred to as “shared control” in the field of human-machine interaction). Engaging in dyadic practice using either of these methods (observational learning or shared control) has been previously demonstrated to aid in learning complex tasks such as surgical techniques (Bjerrum et al., 2014; Tolsgaard et al, 2015) and teleoperation of machines (Griffin, Provancher, & Cutkosky, 2003;

Nudehi, Mukherjee, & Ghodoussi, 2005), and we believe that it could also benefit the learning of complex BoMIs.

When engaging in sequential dyadic practice, learning can be optimized by taking advantage of rest periods. In tasks that follow this sequential ("one on, one off") practice schedule, participants engage in physical practice of the task, and then are given the opportunity to rest and observe the actions of their partner on the following trial. This appears to not only allow for the participant to consolidate their own movements on the prior trial, but simultaneously observe the movements of their partner and gain insight into how their partner's movements affected task performance. In the event that a novel coordination pattern must be found, this style of sequential practice can help rule out ineffective patterns: as one person performs the task, their partner can observe how well their strategy works, and decide to further pursue that pattern, or rule it out on future attempts. These benefits to learning efficiency do not appear to manifest in instances in which one participant performs all trials before allowing their partner to perform theirs (Shea, Wulf, & Whitacre, 1999); it appears that alternating between physical practice and observational practice during rest is what makes this method of dyadic practice effective, as doing so otherwise does not give the partner the opportunity to implement movement strategies they may have observed.

While sequential practice can benefit the learning of tasks intended to be performed by a single user, it may not be appropriate in all circumstances. For high DOF tasks (which may require coordination of many joints simultaneously), or tasks containing redundancy (i.e., the same tasks can be performed with different movement solutions), the abundance of movement options can cause the learner to either underexplore (Lee and Ranganathan, 2019) or overexplore (Ranganathan et al., 2019) their available movement options. In the case of Body-Machine

Interfaces, in which the control scheme will be unique to the individual based on their residual movements (i.e., movements that persist following a movement impairment), the movements that one individual may use to control the interface may not be comfortable or even possible for another. If this is the case, it may instead be beneficial to implement a joint dyadic practice paradigm and examine how the magnitude of each person's contribution benefits the learning of these more complex tasks.

Shared control tasks pair a novice with a trained expert with the goal of completing a given task simultaneously. Control of the task can be represented by the following equation (Nudehi, Mukherjee, & Ghodoussi, 2005):

$$X_{task} = \lambda X_{expert} + (1 - \lambda) X_{novice}$$

where X is a vector that describes some task variable that is being controlled (say the position of a cursor). In these paradigms, the key variable that determines the type of practice is the "responsibility" λ, which can take values between 0-1. When λ is 1, the task is completely controlled by the expert (i.e., the novice has no impact on task performance), and when λ is 0, the task is completely controlled by the novice (i.e., the expert has no impact on task performance). Because λ can take values in between as well (a value of 0.5 would allow both participants to contribute equally to task performance), this variable allows us to parametrically vary the "assistance" provided by the expert in relation to the task during dyadic practice.

Therefore, manipulating this factor across practice provides us with a method of evaluating the effect of dyadic practice with a focus on both the observational aspect as well as the shared control aspect. While previous research has discussed using this method of shared control to allocate responsibility to the novice based on their skill level (Khademian & Hashtrudi-Zaad, 2011), the means of determining how much control should be allocated based on the learner's skill has not been evaluated. It is possible that the expert's contribution to task performance could benefit from the same "fading" idea used to determine how feedback is given to a novice throughout their learning of the task.

Section 1.5. Focus of dissertation

Thus, the purpose of this dissertation is to assess the effectiveness of a dyadic practice paradigm on learning a novel Body-Machine Interface task. Our goal is to examine how varying the interaction between the two participants during dyadic practice can facilitate the learning of the interface, where "learning" is inferred based on performance when participants eventually perform the task on their own (i.e., without the expert). Specifically, we address the following questions – (i) how does observation of a trained expert during dyadic practice compare to solo motor exploration when learning a novel motor task, and (ii) can fading strategies that reduce expert contribution throughout dyadic practice improve learning?

In chapter 2, we begin by exploring the difficulties that arise when learning these Body-Machine Interface tasks, the drawbacks to current methods of training, and the benefits that dyadic practice could provide. In chapter 3, we discuss the first of two experiments, in which we examine the benefits of observational learning within a dyad when learning a BoMI task and how dyadic practice compares to learning the task entirely on one's own. In chapter 4, we discuss the second experiment that explores how to best allocate control to a novice throughout the learning of a BoMI task. Finally, in chapter 5, we discuss common themes observed in chapters 3 and 4, as well as future directions for this line of work.

The general outline of the experiments is as follows: The task we examine is one where participants were asked to learn a two-dimensional cursor control task, controlled via inertial measurement units placed on their shoulders. In the dyadic practice groups, participants were paired with an "expert" who has had prior practice on the task. Five total groups were collected: a constant control dyad, in which the novice and the expert share equal control of the task and are seated in view of one another; a visually-separated dyad, identical to the constant control dyad, except individuals in the dyad are visually separated from their partner; a gradual control dyad, in which control of the task is gradually given to the novice as training progresses; an adaptive control dyad, in which control of the task is given to the novice based on their prior performance; and a solo control group, that learned an identical cursor control task without the presence of a partner.

In evaluating both the observational learning and shared control aspects of dyadic practice, this dissertation hopes to bring insight into effective methods of learning Body-Machine interface tasks. As the present paradigm has previously seen success in restoring motion to individuals following spinal cord injury, as well as having applications in 2D, virtual 3D, and true 3D space, we anticipate that the findings from this research will go on to inform the effectiveness of dyadic practice in BoMI tasks of various complexities and applications. By expediting the learning of these devices for individuals who need them, we can work to reduce individual attrition from the task, allowing these individuals to resume a more independent, self-sustaining lifestyle.

Chapter 2 – Literature Review

Section 2.1. Learning in BoMIs: Theoretical considerations

Individuals learning to use a BoMI to operate an assistive device, such as an upper body prosthetic, will undergo weeks or months of training in order to effectively control the device and restore their movement capabilities to a self-sustaining capacity (Wake Forest Baptist, accessed 11-9-23). As these devices typically rely on "residual movements" of the individual – movements that the individual is capable of following their injury or impairment – for control, individuals are tasked with developing a new coordination pattern in order to control the device, not unlike learning to play an instrument or a sport. The development and learning of these coordination patterns is referred to as "de novo learning" (Telgen, Parvin, & Diedrichsen, 2014; Sternad, 2018; Krakauer et al., 2019; Haith et al., 2022), and requires significantly more time spent learning before the task is learned adequately compared to adapting an existing movement pattern.

De novo learning requires the development of novel coordination patterns in such a way that previous knowledge of other tasks do not benefit the learning of the new task. Learning to shoot a free throw, for instance, does not benefit from prior knowledge of how to play the cello, as the two tasks require two fundamentally different coordination patterns. As it could be argued that no task is truly "novel" in this way, as they will consist of movements with which we are familiar, the definition of de novo learning can be expanded to include learning to (i) select an existing action under unique circumstances, or (ii) generate novel motor output via less explored coordinations (Krakauer et al., 2019; Haith et al., 2022). Given our focus on learning to operate an assistive device via residual movements, discussion regarding the difficulties of de novo learning will primarily favor the generation of these novel motor outputs. The timescale of de novo learning will generally take place over the span of weeks or months, making it difficult to study in a laboratory context.

This stands in contrast to adaptation learning, in which the learner can utilize a previously learned coordination pattern and modify it to accommodate the constraints of the task. Adaptation learning is typically characterized by two distinct outcomes: (i) relatively quick adjustments to movement strategies in the face of a perturbation (such as a visuomotor rotation, force-field, or gain modification), and (ii) persistent aftereffects, in which upon removal of the perturbation, the adapted movement pattern persists (Shadmehr & Mussa-Ivaldi, 1994). These quick adjustments to movement strategies can typically be elicited in the span of just a few trials (Krakauer et al., 2000; Bastian, 2008; Liu et al., 2011) as existing coordination strategies are adjusted to fit the new task requirements. These adjustments will then persist once the perturbation is removed as the individual "de-adapts" to the task (Kitago et al., 2013), even if the individual is explicitly told that the perturbation has been removed (Kluzik et al., 2008; Taylor & Ivry, 2011). These characteristics make adaptation learning an attractive subject for scientific testing, as they are quick to observe and (relatively) simple to predict.

De novo learning does not get this luxury, and there appears to be several reasons as to why. For one, the presence of de novo learning is generally determined by ruling out adaptation. If a response to a particular perturbation does not happen in the span of a few trials, and does not persist upon removal of the perturbation, it is assumed that any existing movement patterns cannot be adapted to complete the present task, and a new pattern must be developed (Haith et al., 2022). This has previously been observed when comparing visuomotor rotation (VM) tasks to mirror rotation tasks; common adaptation effects observed in the VM rotation task, such as a lack of aftereffects once the perturbation has been removed (Yang, Cowan, & Haith, 2021) or adaptations occurring within a handful of trials (Telgen, Parvin, & Diedrichsen, 2014), are absent in the mirror rotation task, suggesting that the individual cannot quickly compensate for the mirror rotation, instead having to develop a new coordination pattern to accommodate.

Second, it is difficult to quantify the development of a novel coordination pattern. Developing a new coordination pattern is not a linear process, nor will one person's learning resemble that of another. In adaptation learning, an individual's behavior can be quantified rather easily by looking at outcome metrics of the individual's performance, such as the initial direction of a reach or an applied force (Haith et al., 2022), offering that learning can at least be suggested by these metrics (i.e. an initial reaching trajectory closer to that of the presented target suggests that the learner is "getting better" at the adjusting to the current perturbation). In de novo learning, however, it is less clear what metrics could be consistently measured in order to quantify the development of this new coordination pattern at any given time (Ranganathan & Scheidt, 2016), or how the learning of this pattern progresses between trials (Haith et al., 2022). Essentially, it is difficult to quantify whether adjustments made to a coordination pattern during the learning process benefit the task, or are otherwise irrelevant, and becomes even more difficult in tasks with many degrees of freedom.

Section 2.2. Current methods of training

Much of the previous research examining de novo learning in BoMI tasks have done so using a 2D cursor control task (Farshchiansadegh et al., 2014; Lee et al., 2016; Abdollahi et al., 2017; Lee & Ranganathan, 2019; Rizzoglio et al., 2020; Pierella et al., 2021) or a 3D robotic arm end effector task (Ranganathan et al., 2019; Aspelund et al., 2020; Rizzoglio et al., 2023), both of which will typically require reduction of the degrees of freedom (DoF) of which an individual is capable into the DoF of the task being performed (referred to as "dimensionality reduction"). For instance, previous research examining 2D cursor control will outfit participants with IMU sensors on their trunk or upper extremities, and while the number of signals collected from these sensors may provide an eight DoF space (2 signals by 4 sensors, for example), these eight signals will be condensed into control of a two-dimensional task. This will present the user with some degree of redundancy, in that there is now more than one way an individual could move in order to successfully complete the given task.

For individuals with movement impairments, this redundancy is usually a good thing: if multiple methods exist to control the BoMI, it is likely that some combination of existing residual movements from that individual will serve to successfully complete the task. Or, in an ideal scenario, the BoMI will allow individuals to calibrate the interface to fit their own residual movements, and then transform these body-space movements into a map that corresponds to 2D or 3D control of a cursor or end effector. In most cases, while this results in a map with orthogonal directional control, the coordinate system of the body-space may not directly align with the coordinate system of the task space (e.g., moving the cursor upward on the screen may require the person to engage in lateral leaning of the trunk), thus providing us with an opportunity to explore different movement synergies.

However, the difficulty in learning these tasks comes from asking these participants to explore this redundancy. Individuals with movement impairments may have some residual movements left of their distal limb segments, but prefer to rely on the more gross movements that remain more proximal to the body. And yet, this approach of exploring redundant movement options to force use of more distal segments has been proven effective for not only learning the task, but also for restoring general motion to these individuals (Pierella et al., 2015; Seáñez-González et al., 2016; Pierella et al., 2017; Pierella et al., 2021). Similar difficulties in exploiting this redundancy are seen in children (Ranganathan et al., 2019) and older adults (Lee and Ranganathan, 2019) when exploring these tasks, as children tend to explore all options available to them (i.e. not narrowing down viable options, or spending too much time exploring a movement pattern that is ineffective), while older adults tend to do the opposite and explore only the options that are familiar or comfortable to them (akin to Bernstein's "freezing degrees of freedom" problem [1967]).

This difficulty in learning becomes even more salient once the BoMI is used to control an object in 3D space. As is the case with human movement, control of an object in 3D space can require control of more than three degrees of freedom, due to the body's near-infinite combinations of joint postures that can give rise to a proper solution (Bernstein, 1967). Many BoMI studies that have evaluated task performance in 3D space have used a 7 DoF robot arm (2 DoF at the shoulder, 1 DoF at the elbow, 3 DoF at the wrist, and 1 DoF to open/close the grasp) (Ranganathan et al., 2019; Aspelund et al., 2020). As the number of DoF in an apparatus increases, the more difficult it becomes to operate all of its DoF simultaneously, instead resorting to controlling a set of constrained dimensions and then toggling the interface to control a different set of dimensions as they become relevant (e.g. a participant may be able to control the 3 DoF of the robotic wrist using a single IMU, then toggle a button that allows them to open/close the grasp by using the same movements used to control the wrist). This can present difficulties for individuals with a limited movement repertoire, as the same coordination pattern must now be associated with control of multiple degrees of freedom.

For this population, control of the assistive device no longer has the luxury of being intuitive. As the coordination pattern required to operate these devices is based on the movement capabilities of its user, the options for control will decrease along with the capabilities of the individual. Now not only are users developing a new movement pattern, but the coordination of joints or muscles necessary to operate the assistive device may not align with how that movement pattern may have been performed if the user had unrestricted movement capabilities. However, while these movements may be unintuitive relative to an individual with unrestricted capabilities, it is likely that the solution to which these individuals arrive is the most personally optimal given their specific movement repertoire (Latash & Anson, 1996). As such, it is important to capture and harness all residual movements of an individual so they can utilize the entirety of their movement capabilities for control of the interface.

As these tasks become more time-consuming or difficult to "figure out" quickly, the risk of attrition (dropout) increases as well. In instances in which the "correct" coordination pattern must be found, the increased time spent learning is likely attributable to simply exploring all available options. If the "correct" coordination pattern is already known (i.e. in instances in which the coordination pattern is mapped to an individual's residual movements), this increased time spent learning is instead spent associating this new coordination pattern with the given task (i.e., engaging in de novo learning). Either way, a task that takes too long to "get," or whose requirements for successful performance are not immediately clear, can result in increased time spent training (which can be tiring and/or frustrating for participants), which then can lead to attrition from the task. For individuals with movement impairments, attrition becomes much more severe: failure to learn the present task indicates a failure to restore motion and independence to those who are putting forth effort to do so. Given this risk of attrition and increased time spent learning, exploring other methods of practice that can reduce these potential barriers without compromising learning becomes appealing.

Section 2.3. Dyadic practice to expedite learning

A promising method to reduce time spent learning is through dyadic practice, in which two individuals work together to learn a task that is intended to be performed on an individual basis (Crook & Beier, 2010). Typically, dyadic practice will take one of two distinct forms: (i) sequential dyadic practice, in which the learner observes the actions of their partner during their rest periods, or (ii) joint dyadic practice, in which simultaneous control of a task is shared between the learner and their partner. Practicing a task within a dyad can offer several benefits, including improved solo performance after being trained with a partner (in the event that the task is ultimately intended to be performed alone) and increased training efficiency (as multiple participants can be trained on a task in the time it would normally take to train one person).

Sequential dyadic practice is most effective when individuals get time to reflect on their own performance, observe the performance of others, and engage in discussion of effective and ineffective movement strategies. De novo learning, as previously mentioned, can be time-consuming due to the possibility of the individual exploring every option available to them, but if individuals are first provided with some guidance on what successful task performance looks like, many of the ineffective options can be immediately ruled out. In the event that individuals are alternating practice trials (i.e. Participant A tries one trial, then Participant B tries one trial, etc.), this time spent observing their partner perform the task also allows the individual to consolidate their own previous trials and better understand what did and did not work (Shea, Wulf, & Whitacre, 1999). Alternating trials in this way tends to be more effective than having one individual perform all of their trials before the other individual has a chance to practice, as even though the other individual has observed what successful performance looks like, they have received no opportunity to practice the movement pattern themselves, and have likely lost important information that exceeds what can be stored in the individual's short term memory.

It is also not necessary for the individual that the novice is paired with to be trained expert to benefit from this sequential dyadic practice. Instead, pairing two novices together may allow them to engage in the aforementioned combination of motor exploration and observational learning by observing strategies performed by their partner and integrating those strategies into their next attempt. This type of dyad has been found effective in medical students learning medical and surgical techniques, as the combination of physical practice and observational learning proved to be just as effective as solo practice while requiring fewer trials from each participant (rather, individuals in a dyad learned just as well while receiving half as much physical practice as their solo counterparts; Bjerrum et al., 2014; Cordovani & Cordovani, 2016; Tolsgaard et al., 2015).

If instead two individuals are paired together and tasked with simultaneously completing the task (i.e., joint dyadic practice), responsibility (i.e., how much of the task each individual controls) can either be shared by splitting the task into discrete components that each individual is solely responsible for (referred to as the Active Interlocking Model; Shebilske, 1992), or by providing the individual with a percentage of control of the entire task based on their skill level (Nudehi, Mukherjee, & Ghodoussi, 2005). Regarding the former, by splitting the task into subtasks and allowing participants to alternate control of the subtasks through training, participants are afforded similar benefits to those observed in sequential dyadic practice, as the time spent performing one subtask can simultaneously be used to consolidate the information acquired when performing the other subtask (i.e., participants get physical practice on Subtask A while observing and consolidating information on Subtask B). This method allows for increased efficiency as two people can be trained on the task in the same amount of time it would have originally taken to train one, and both participants be able to perform the task in its entirety once trained. This interlocked model also reduces the risk of loafing (a reduction in performance or effort observed when an individual participates in a dyadic or group task; Latané, Williams, Harkins, 1979) as by interlocking the dyad in this way, the task is rendered impossible to complete without the assistance of both participants. However, splitting a task in this way may not always be feasible or appropriate; instead, participants may benefit from a paradigm that allows the learner to control all aspects of the task, but varies the magnitude of that control to reduce errors as learning progresses.

This alternative shares control between participants so that the learner can experience the full environment and feedback of the task without immediately having full control. In surgical simulators, this paradigm is used to pair a trained surgeon with a student learning the procedure: this allows the student to explore the task while allowing the trained surgeon to correct for any errors incurred by the novice in the early stages of learning (Shamaei, Kim, & Okamura, 2015). For complex tasks with many degrees of freedom, this "expert-novice" pairing may be beneficial as the expert can demonstrate successful task performance or mitigate excess degrees of freedom. If the goal of the expert is to expedite learning for their novice partner, their demonstration can help guide novices to a correct coordination pattern, potentially allowing the novice to "get" the task quicker and reduce time spent exploring the entirety of the space.

The concept of having an expert present to mitigate excess degrees of freedom has been explored in robotics and teleoperation, in which a human and a robot act together to perform a complex task, providing the human with control of the more complex/detail-oriented parameters of the task (i.e. end-effector position), leaving the robot to control the lower order parameters to aid in task performance (i.e. force production of the various degrees of freedom to allow the end effector to reach the desired position). In our case, the use of an expert within a dyad can provide the same benefits as the robot – the expert can serve to reduce the cognitive load on the novice by taking over extraneous degrees of freedom, leaving the novice to explore the space without having to control every degree of freedom of the task. If the goal is for the novice to eventually perform the entirety of the task on their own, the expert may gradually concede control as the novice becomes more comfortable with the task (i.e. the "fading" idea mentioned previously). This allows the novice to initially focus more on the degrees of freedom that directly relate to successful task performance (as we see in shared control research) and limits the effect that their exploration of the space has on the completion of the task (i.e., errors that a novice makes during task performance are less dire as they are not controlling the entirety of the task).

However, this expert-novice pairing may not be beneficial in all circumstances; previous research suggests that while the performance of the dyad as a whole increases when it contains an expert, the learner's performance within that dyad actually suffers (Kager et al., 2019). This is likely due to the expertise of the expert completely taking over the task, and effectively "dragging the novice through the mud". The novice gets no chance to explore the space of their own accord and is more or less just "along for the ride." If the goal is for the novice to be able to perform the task entirely on their own, novices may benefit more from exploring the task on their own before being paired with their expert partner (Avila Mireles et al., 2017), or by being paired with another novice instead of an expert throughout the entirety of the learning process (Kager et al., 2019; Avila Mireles et al., 2017; Saracbasi et al., 2021).

Section 2.4. Literature Gaps

Existing laboratory paradigms implementing joint dyadic practice will usually test adaptation learning in low-dimensional tasks. Most often, dyads are asked to complete a cursor or target tracking task in one (Takagi, Beckers, & Burdet, 2016; Ivanova et al., 2022; Kim et al., 2023) or two dimensions (Ganesh et al., 2014; Avila Mireles et al., 2017; Kager et al., 2019; Batson et al., 2020; Beckers, van Asseldonk, & van der Kooij, 2020; Saracbasi et al., 2021) in which the location of the cursor/target is positionally mapped to the linear or planar movements of the apparatus. For instance, participants may be asked to complete a target tracing task in which both participants operate a robotic manipulandum in 2D space, and each person's cursor is mapped to the 2D location of their respective robotic end effector. These movements will then be subject to some force field or visuomotor rotation, eliciting adaptation learning as participants account for the presented perturbation. This reliance on low-dimensional adaptation tasks makes it difficult to generalize the benefits of joint dyadic practice to more complex, high-DoF tasks.

Not only that, but the feedback mechanisms used to guide these tasks inherently limits their dimensionality. The robotic manipulanda used to operate the task can allow for haptic coupling between participants (Takegi, Beckers, & Burdet, 2016; Avila Mireles et al., 2017; Kager et al., 2019; Ivanova et al., 2022; Kim et al., 2023), which serves to guide participants towards a proper solution. The positions of the two individual cursors are combined into a single "true" cursor that is observed by the participants, which provides haptic feedback via virtual springs that "pull" each person's manipulandum towards the location of the true cursor. While this feedback can be beneficial in instances in which participants are visually separated from their partner, its use is limited to the planarity of the robotic manipulandum, which causes similar problems in generalizing to higher-DoF tasks.

In cases in which dyadic practice does occur in higher dimensions, it is usually implemented via sequential dyadic practice instead of joint dyadic practice. In medical and surgical procedure simulator tasks, participants will engage in the “see one, do one” method of sequential practice, as they perform one trial of the given task, and then rest while observing their partner perform their trial (Shanks et al., 2013; Bjerrum et al., 2014). These tasks, unlike the aforementioned force field or visuomotor rotation tasks, require the learner to engage in de novo learning, as they develop the novel coordination pattern required to successfully complete the procedure. However, this method will inherently limit the magnitude of interaction participants have with one another, as they are no longer simultaneously completing the task or receiving feedback from one other, but instead alternating between observation and practice.

As a result, there is a lack of research for how to use dyadic practice for high-DOF de novo learning tasks like a body-machine interface. Specifically, there are few examples of dyadic practice that combine observational practice and shared control, in which individuals engage in simultaneous performance of a task with their partner (joint dyadic practice) while being able to observe their partner’s movements and implement them in future trials (sequential dyadic practice). Furthermore, the tasks used in typical experimental paradigms rarely involve de novo learning with coordination of several DOFs. This is likely due to (i) difficulties in studying de novo learning such as high within- and between-subject variability (Ranganathan and Scheidt, 2016), and (ii) providing feedback (such as haptic feedback) becomes less feasible due to the numerous dimensions about which this feedback could be provided.

In the present study, we seek to fill this gap by implementing a joint dyadic practice task that allows individuals to engage in de novo learning of a novel Body-Machine Interface task while being able to immediately observe the movements of their partner. This task will allow us to evaluate the effectiveness of a joint dyadic practice paradigm when participants can no longer rely on the haptic interaction between themselves and their partner, and how de novo learning may benefit from the combination of joint and sequential dyadic practice.

Chapter 3 – Solo Practice Provides Greater Benefits Compared to Joint Dyadic Practice when Learning a Cursor Control Task

Section 3.1. Abstract

Body-Machine Interfaces can help restore motion and independence to individuals with movement impairments, but learning the underlying movement coordination to control these interfaces can often take extensive time and effort. Joint dyadic practice is a practice strategy that could potentially mitigate some of this difficulty by allowing the learner to perform the task simultaneously with an expert. We conducted two experiments to observe the effectiveness of dyadic practice on the learning of a novel cursor control task using movements of the trunk. In Experiment 1, we compared a Solo group (n= 20) that practiced the task on their own to a group that learned the task within a dyad while being able to watch the movements of their trained expert partner ($\text{Dyad}_{\text{VISION}}$, n = 16). In Experiment 2, to examine the contribution of observational learning to dyadic practice, we collected an additional group that performed the same task as the $\text{Dyad}_{\text{VISION}}$ group, but without direct visual information of the partner ($\text{Dyad}_{\text{NOVISION}}$, n = 8). Learning was examined by analyzing movement time and angular error during testing conditions in which participants performed the task on their own.

Results showed that although the $\text{Dyad}_{\text{VISION}}$ group showed better task performance during training when partnered with the expert, the Solo group showed better task performance during the post-test conditions, as indicated by lower movement times (2.48 vs 4.16 sec, p = 0.001) and angular error (22.95 vs 35.76 deg, p = 0.004) compared to the $\text{Dyad}_{\text{VISION}}$ group. Results of Experiment 2 found no significant differences between the two dyadic groups, indicating that vision of the partner's movements did not have a direct impact on dyadic learning.

These findings suggest that dyadic practice with an expert might enhance performance but is not likely to be effective for learning if the expert can dominate task performance throughout practice.

Section 3.2. Introduction

Body-Machine Interfaces (BoMIs) can restore a sense of independence to individuals with movement impairments by allowing them to complete activities of daily living using an assistive device. These interfaces transform movement signals from the body, such as an individual's kinematics, kinetics, or muscular activity, into control commands for the assistive device (Casadio, Ranganathan, & Mussa-Ivaldi, 2012). BoMIs are generally more accessible than interfaces based on neural signals (i.e., Brain-Machine Interfaces) because they are non-invasive, affordable, and have high signal-to-noise ratios (Aspelund et al., 2020), which makes them useful for adapting to individuals based on the unique movement capabilities of an individual.

Despite these advantages, the issue of learning to control these interfaces remains a significant challenge. This is partly because individuals may have to learn a "non-intuitive" method for controlling the device. For example, a transhumeral amputee using an EMG-controlled prosthetic arm must find a way to activate different muscles to command the prosthetic arm to open and close (Fall et al., 2017). The intuitiveness of such tasks is based on the "map" used to transform the movement signals into commands of the device. Appropriate selection of this map will harness the movement capabilities of an individual and can be further customized based on the needs of the user; in the event that a user's movement capabilities change over time, the map can be modified to match these changes. In rehabilitation contexts, this map can also be used to encourage users to explore "less preferred" movements, which has been shown to improve movement capabilities for individuals recovering from spinal cord injury (Pierella et al., 2015; Seáñez-González et al., 2016). With these complexities in mind, it is important to understand how to best encourage learning of these devices to minimize risk of attrition, as the ease of use of such devices has been previously used as a predictor for device abandonment (Phillips & Zhao, 1993; Biddiss & Chau, 2007).

One method of doing so comes in the form of dyadic practice, in which two individuals work together to learn a task that is intended to be performed on an individual basis (Crook & Beier, 2010). Two distinct implementations of dyadic practice have been observed in the literature: (i) sequential dyadic practice, in which the learner observes the performance of their partner during their own rest periods, and (ii) joint dyadic practice, in which control of a task is "jointly shared" between the learner and their partner. Sequential dyadic practice is based on observation, and is typically most effective when employed in a "see one, do one" context (Shea, Wulf, & Whitacre, 1999; Shea et al., 2000; Wulf, Shea, & Lewthwaite, 2010; Bjerrum et al., 2014; Cordovani & Cordovani, 2016), as it allows for an optimal balance of observing potential movement options and exploring them on one's own. This approach is generally more effective than watching one's partner perform the entire task before being allowed to practice, as individuals do not get the chance to practice or implement the strategies they have observed (Shea, Wulf, & Whitacre, 1999).

On the other hand, joint dyadic practice tasks split control of the task in such a way that both individuals are in control of the given task; this is accomplished either by splitting the task into discrete subtasks that each participant has complete control over (referred to as the Active Interlocking Model [Shebilske et al., 1992]), or by allowing both individuals to control the entirety of the task and allocating a percentage of control based on the learner's skill (Nudehi, Mukherjee, & Ghodoussi, 2005). Participants in joint dyadic practice tasks are often placed out of view of one another, forcing them to instead communicate through the device itself via haptic feedback (Reed et al., 2006; Avila Mireles et al., 2017). While these two distinct implementations of dyadic practice have been shown to benefit learning of complex surgical (Bjerrum et al., 2014; Ritchie et al., 2024), robotics (Groten et al., 2013; Ganesh et al., 2014), and sport-specific (Shea, Wulf, & Whitacre, 1999; Scott et al., 2023) tasks, they are typically not implemented simultaneously in the learning of these tasks, especially for multi-DOF movements.

Here, we seek to combine these implementations into a single paradigm that allows the learner to observe the movements of their partner during joint control of a single task. Specifically, we collected three groups: a Solo group learning by trial-and-error, a dyad group that performs the task in view of their partner ($\text{Dyad}_{\text{VISION}}$), and a dyad group that was visually separated from their partner ($\text{Dyad}_{\text{NOVISION}}$). In our first experiment, we compared how individuals learning by trial-and-error compared to individuals in the $\text{Dyad}_{\text{VISION}}$ group and anticipated that the novice would be able to use their observation of the expert to understand the motions required to perform the task successfully. In our second experiment, we further examined the benefits of observation when learning the BoMI task by comparing the $\text{Dyad}_{\text{VISION}}$ group from experiment one to the $\text{Dyad}_{\text{NOVISION}}$ group and anticipated that visually separating the dyad would complicate the task by reducing partner visual feedback availability to the novice. To examine the novice's ability to learn the task, we evaluated task performance and movement synergy variables during testing conditions in which the novice is performing the task entirely on their own. We hypothesized that task performance and movement coordination variables would improve to a greater extent in the $\text{Dyad}_{\text{VISION}}$ group compared to the Solo group in experiment one, as well as compared to the $\text{Dyad}_{\text{NOVISION}}$ group in experiment two.

Section 3.3. Experiment 1

Methods

Participants

Thirty-six participants were recruited as a part of experiment one, (22 female, 14 male) aged 18-35 years (mean ± SD: 22.9 ± 4.1 years). Participants self-reported that they did not have any neuromuscular impairments that would affect their ability to move their upper body through a comfortable range of motion. All participants except for two reported as being right-handed (one reported as left-handed, and one did not report handedness). All participants completed informed consent prior to their participation, and all data collection procedures were approved by Michigan State University IRB.

Two groups were collected as a part of experiment one: a solo control group (n = 20) and a $\text{Dyad}_{\text{VISION}}$ group (n = 16). After completing a demographic questionnaire, participants were briefed on the structure of the data collection session. Individuals in the solo group were tasked with completing the task (described below) entirely on their own, while individuals in the dyad group were paired with a trained research assistant (hereby referred to as the "expert") to perform the task collaboratively. Before being permitted to serve as the expert within a dyad, three research assistants trained on the task on their own, and their performance was evaluated by the primary researcher to ensure the task was learned sufficiently. Recruitment for the solo group was completed first as the research assistants became proficient in the task.

Task

Four triaxial inertial measurement unit (IMU) sensors (YOST Labs, Portsmouth, Ohio) sampling at 50Hz were affixed to the anterior and posterior aspects of both individual's shoulders (eight IMUs within a dyad, four per person). Roll and pitch angles from each sensor were used to control the onscreen cursor (yaw angles were omitted due to their sensitivity to magnetic fields) via a custom Simulink program. Roll and pitch signals from each of the sensor were converted into a set of Cartesian coordinates for the screen cursor as follows (Farshchiansadegh et al., 2014):

$$X = A\,H + X_0$$

in which X is the Cartesian coordinates of the cursor (sized [2 x 1]), A is the mapping matrix (sized [2 x 8]) used to transform the set of IMU signals H (sized [8 x 1]), and X_0 (sized [2 x 1]) is an offset term so that the mean resting posture corresponded to the center of the screen. In the dyadic group, this calculation was done on both participants to obtain a virtual cursor position for each participant, and the average position of the two cursors with an equal weighting was calculated to obtain the cursor position:

$$X = \lambda\, X_{novice} + (1 - \lambda)\, X_{expert}$$

in which λ is a weighting factor used to assign a percentage of control to the expert that could vary between 0 and 1. Throughout the entirety of this study, λ was fixed at .5, which provided both expert and novice equal control of the task.

A critical difference from previous paradigms was that in prior work, the map A was customized to each participant based on exploratory movements of their own body (see the "body dance" from Lee et al., 2016). However, in the present paradigm, because there were two individuals involved in the dyad group, we used a constant pre-defined map for all participants that was developed by the experimenter. This allowed the task to still be novel, but also created a "fair comparison" for all groups.

To examine the effect on motor learning, two conditions were implemented in this study: testing blocks and training blocks. The training blocks differed based on the group – i.e., the solo group practiced these blocks on their own, and the dyad group had both the expert and novice jointly perform the task. In the testing blocks, all participants perform the task entirely on their own. The performance in the test blocks were our measure of “learning” since they only involved the learner. The full schedule of the data collection session can be found in Figure 1.

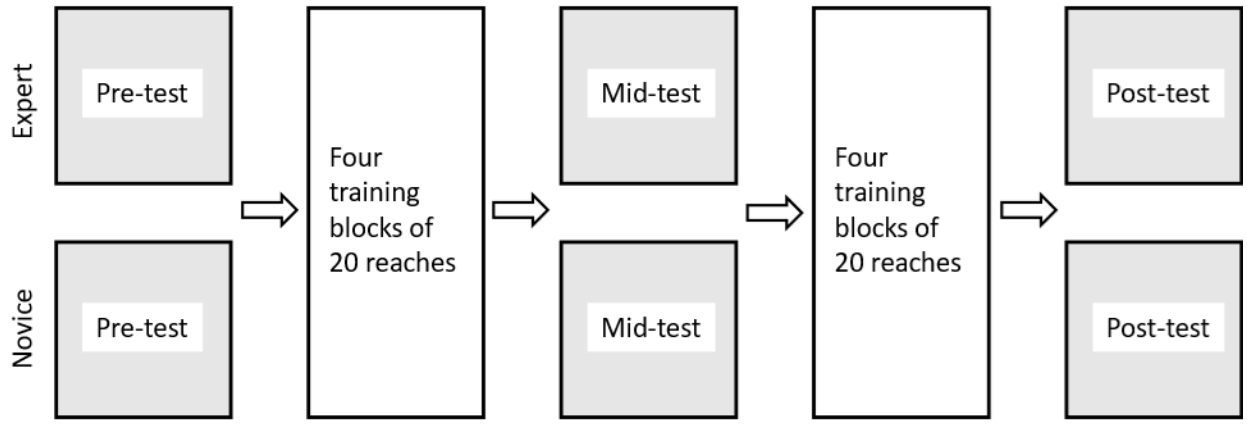

Figure 1 - Schedule of the data collection session. White blocks indicate training blocks where task was completed as a dyad and grey blocks indicate test blocks where participants performed the task on their own. For participants in the solo conditions, no expert was present and all blocks were performed solo.

Procedure

Participants were seated in front of a 23.8 inch computer monitor (Dell P2414H, resolution 1920x1080, refresh rate 59Hz) (Fig 2a), and individuals in the dyad groups were seated side-by-side with one another, with the expert always to the left of the novice when facing the screen (Fig 2b and 2c). Participants were tasked with navigating an onscreen cursor from a home position at the center of the screen to a set of presented targets as quickly and accurately as possible (Fig 3).

Upon affixing the sensors to participants, a familiarization test was performed to ensure the IMUs were functioning properly: participants were asked to move their upper body “any way they could think of while remaining seated” so that an onscreen cursor touched all four sides of the program window. Participants were informed that they were not timed or scored on the familiarization test, and that it only served to “make sure the task is doable.” No other information on how to complete the task was provided. In the dyadic group both the expert and novice shared equal control of the onscreen cursor, and the expert was instructed (prior to the novice’s arrival) to accommodate the novice’s movements to the best of their ability. The solo group underwent this same familiarization test, with the only difference being there was no expert present, and thus only four IMUs were required for control of the cursor.

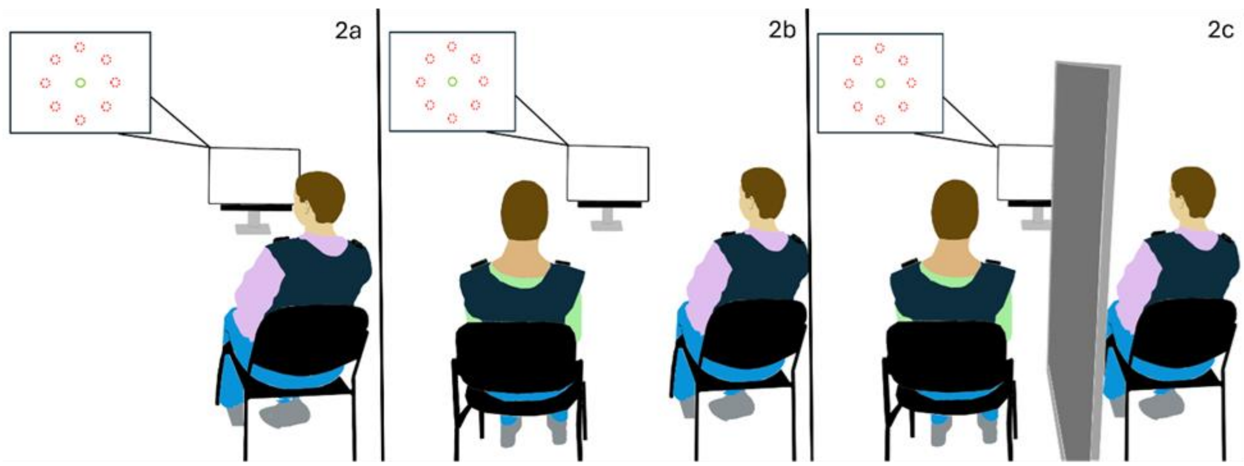

Figure 2 - Diagram of how participants were seated in the Solo (2a), DyadVision (2b) and DyadNoVision (2c) conditions. Participants performed a cursor control task using movements of the upper body. In the dyad conditions, the movement of the cursor corresponded to movements of both the novice and the expert.

Following the familiarization test, participants performed a pre-test, which served as a baseline for the individual’s learning of the task. The testing conditions tasked each participant with navigating the onscreen cursor from the home position to one of eight randomly presented targets as quickly and accurately as possible, positioned equidistant from the home position (Fig 3). Each target was presented three times, for a total of 24 reaches in the testing conditions. This same test was performed halfway through the training (mid-test) as well as at the very end of all training blocks (post-test).

In the training blocks, dyad groups were asked to perform the cursor control task collaboratively, sharing equal control of a single cursor (similar to the familiarization test), while solo groups continued to control the cursor on their own (identical to their testing conditions). While the goal of the training blocks was identical to that of the testing condition, only four targets were presented instead of eight, with each target appearing a total of five times for a total of 20 trials per block. Eight blocks were completed in total, for a total of 160 training trials. To restrict the type of interaction between the participants primarily to observational learning, participants were instructed not to share direct instructions with their partner (e.g., commands such as “lean forward to make the cursor move to the current target” were discouraged), but other forms of communication (such as general encouragement and chatter) were allowed and encouraged to create cohesion within the dyad.

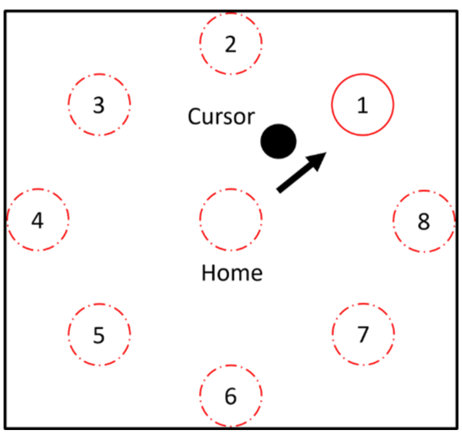

Figure 3 - Layout of the targets in the testing blocks. In the present example, Target 1 is presented, and all other targets are invisible to the participant. In the training blocks, only targets 2, 4, 6, and 8 can be presented.

Data Processing

Task Performance

Movement time of the cursor was calculated for each test condition and training block, calculated from the frame the cursor leaves the home position for the first time to the frame in which the target position is achieved and maintained for 0.5 seconds. Angular error was calculated as the angular displacement of the individual's cursor relative to the presented target 0.5 seconds after leaving the home position (an angular error of 0 indicates that the cursor was aligned with the line joining the home position and the presented target). These metrics were chosen as a method of quantifying performance – as all participants were instructed to move to the target position as quickly and accurately as possible, decreases in movement times and angular error are used to suggest learning of how to perform the task. In the dyadic conditions, movement time was calculated using the movement of the true cursor, while angular error was calculated using the novice's signals (e.g., the angular error of the novice's virtual cursor).

Movement Coordination

The ratio of the variance accounted for (VAF) by the first two principal components (PCs) was calculated using the roll and pitch signals during the testing conditions and each individual training block. This ratio (VAF-ratio) was calculated as (VAF2/VAF1) and used to provide a metric of exploration of the task (i.e., a smaller ratio indicates a greater reliance on PC1 [Lee & Ranganathan, 2019]). Task- (H_{task}) and null- (H_{null}) space variances were computed by first dividing the home position postures into their task- and null-space components, like so (Lokesh & Ranganathan, 2019):

$$H_{task} = A' * (A * A')^{-1} * A * H$$

$$H_{null} = (I_8 - A' * (A * A')^{-1} * A) * H$$

in which H is a matrix consisting of the eight IMU signals at the 24 instances in which the home target was achieved, A is the mapping matrix defined in the "Task" subsection, and I_8 is an identity matrix sized 8x8. These signals were only taken at the instance in which the home position was achieved, as each outer target was only presented three times during each testing condition. These components were then used to calculate the "per-DoF" null-space variance ratio (NullRatio; Scholz & Schöner, 1999), which was calculated as:

$$NullRatio = \frac{Var(H_{null})/6}{Var\ (H_{null})/6 + Var(H_{task})/2}$$

in which the variances about the task- and null-space are divided by the number of relevant signals in each dimension (i.e., two dimensions in the task-space, and six dimensions in the null-space). A ratio of 1 indicates equal distribution of variances along the task and null spaces (i.e., no preferential exploration). In the dyadic conditions, similar to the angular error in the previous section, VAF-ratio and task- and null-space variances were calculated using the novice's eight IMU signals.

Statistical Analyses

To ensure experts had learned the task well enough to help their novice partner, we began by performing analysis of variance on pre-test movement times between the Solo group, the trained expert, and the novice partner from the dyad, with Bonferroni post hoc tests performed as appropriate. Next, to examine if the expert's contribution benefitted novice performance during training, we performed independent samples t-tests with Bonferroni corrections on training block movement times between the Solo group and the $\text{Dyad}_{\text{VISION}}$ group. Alpha for these tests was set to .05.

One-way analysis of covariance was performed to assess between-subject (group membership) differences in movement time, angular error, VAF-ratio, and NullRatio. Pre-test values were used as the covariate to accommodate differences in initial performance. Partial eta-squared (η_p^2) values were calculated for effect size (in which values between 0.01 and 0.05 are considered small, values between .06 and .13 are considered medium, and values above .14 are considered large [Statistical Analysis in JASP: A Guide for Students, 2018]). JASP Software (University of Amsterdam, Netherlands) was used to calculate statistics. We report all the post-test comparisons first as primary comparisons, and the mid-test comparisons as exploratory comparisons.

Section 3.4. Results

Experts Helped their Novice Partner

Analysis of variance revealed statistically significant differences in pre-test movement times among the three groups ($F(2,49) = 18.787$, $p < .001$, Figure 4), with Bonferroni post hoc tests demonstrating that both novice and Solo pre-test movement times were significantly higher (novice: 14.76 sec, $p < .001$; Solo: 10.88 sec, $p < .001$) than expert pre-test movement times (expert: 1.58 sec), suggesting that experts understood how to properly perform the task. No statistically significant differences were observed between novice and Solo pre-test movement times ($p = 0.214$).

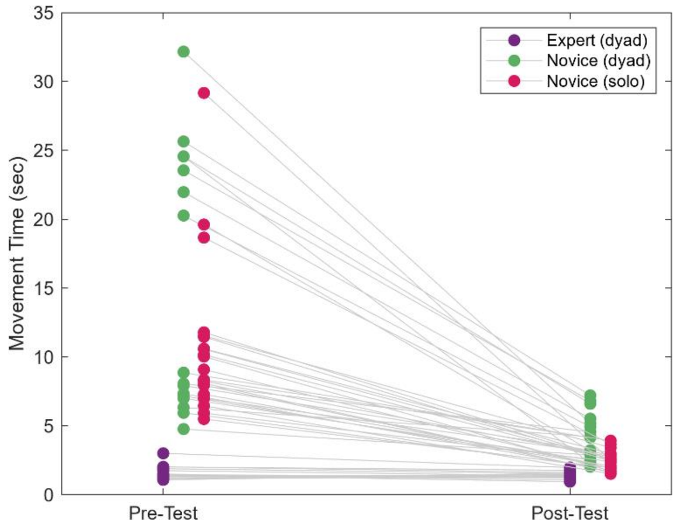

Figure 4 - Movement time comparisons between Experts and Novices in the DyadVision group, and Solo controls. Expert movement times were significantly lower than the novice.

Next, to examine if the expert's contribution benefitted novice performance during training, we compared movement times throughout each of the individual training blocks between the DyadVISION group and the Solo group. Movement times were found to be significantly lower in all training blocks for the DyadVISION group compared to the Solo group (t34 range: -5.911 to -3.309, all $p < 0.05$), suggesting that experts positively contributed to dyad performance and could account for potential errors from their novice partner. A sample of Solo and novice cursor trajectories for pre-, mid-, and post-test conditions are presented in Figure 5.

Solo group outperformed the dyad group in the post-test

Outcome Variables

Both movement times and angular errors decreased with practice. The ANCOVA found that novices in the Solo group demonstrated significantly lower movement times ($F[1,33] = 12.456$, $p = 0.001$, $\eta_p^2 = 0.274$; Figure 6) and angular error ($F[1,33] = 9.710$, $p = 0.004$, $\eta_p^2 = 0.227$; Figure 7) in the post-test compared to novices in the Dyad$_{\text{VISION}}$ group.

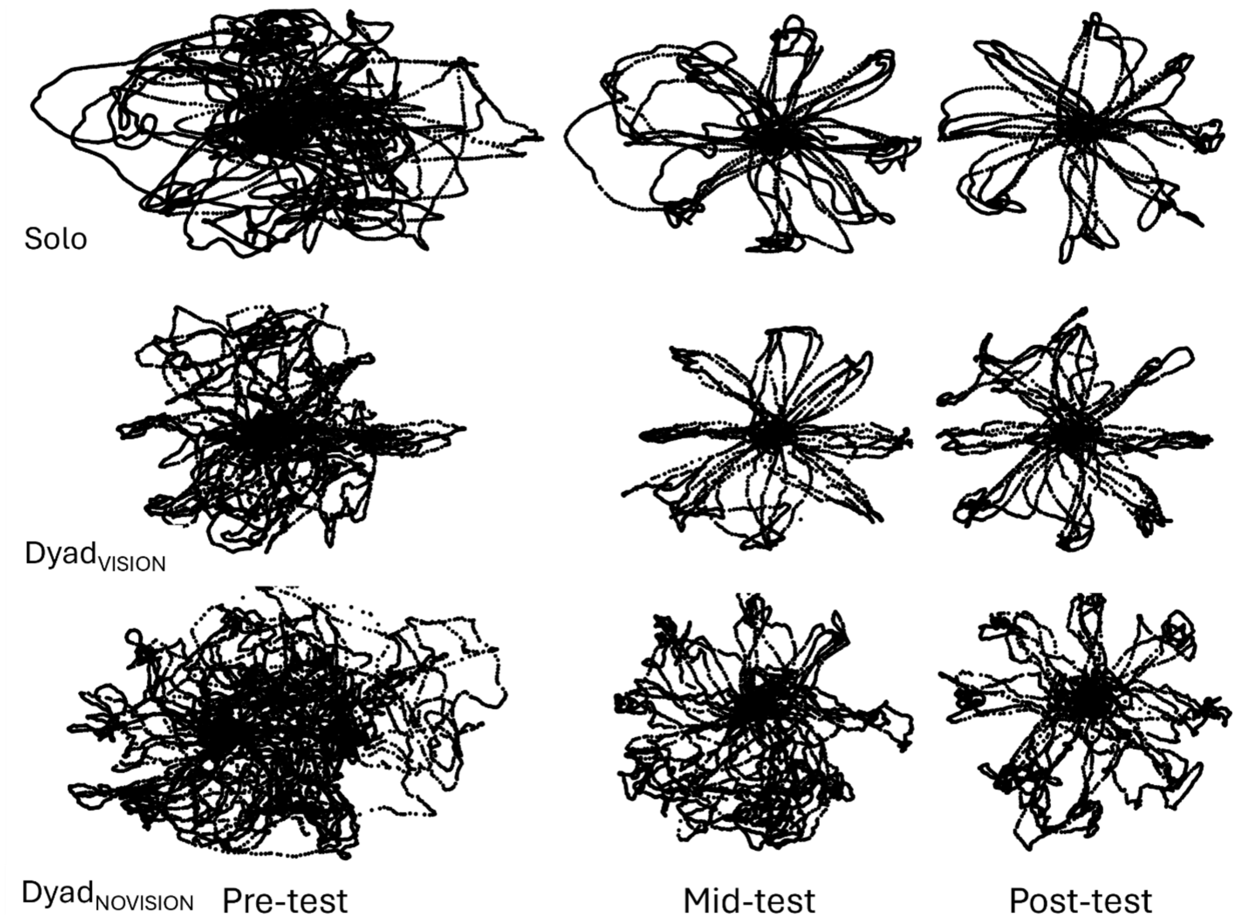

Figure 5 - Sample cursor trajectories for Solo, DyadVision, and DyadNoVision groups in the Pre-, Mid-, and Post-test conditions.

Coordination Variables

Movement coordination as indexed by VAF-ratio was found to be significantly lower for novices in the Solo group compared those in the $\text{Dyad}_{\text{VISION}}$ group in the post-test condition ($F[1,33] = 6.492$, $p = 0.016$, $\eta_p^2 = 0.164$; Figure 8). No statistically significant differences were found in null-space ratio ($F[1,33] = 0.485$, $p = 0.491$, $\eta_p^2 = 0.014$; Figure 9).

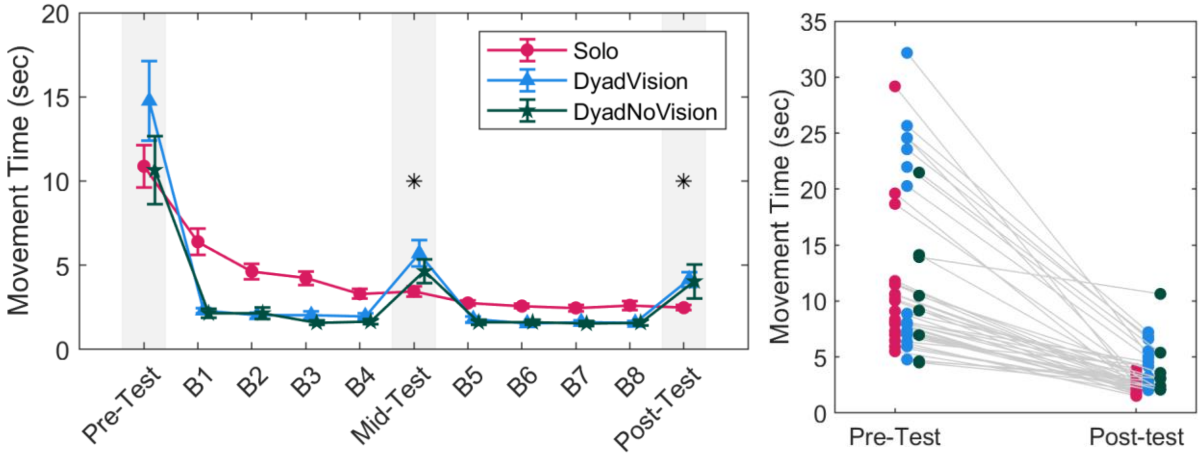

Figure 6 - Average movement time in all blocks between Solo, DyadVision, and DyadNoVision groups. Movement times for the Solo group were significantly lower than the DyadVision group at mid- and post-test time points. Error bars indicate standard error of the mean (between participant). Grey columns represent testing conditions in which participants were tested alone. B1-B8 indicates training blocks one through eight. Right panel shows individual data for the pre-test and post-test conditions.

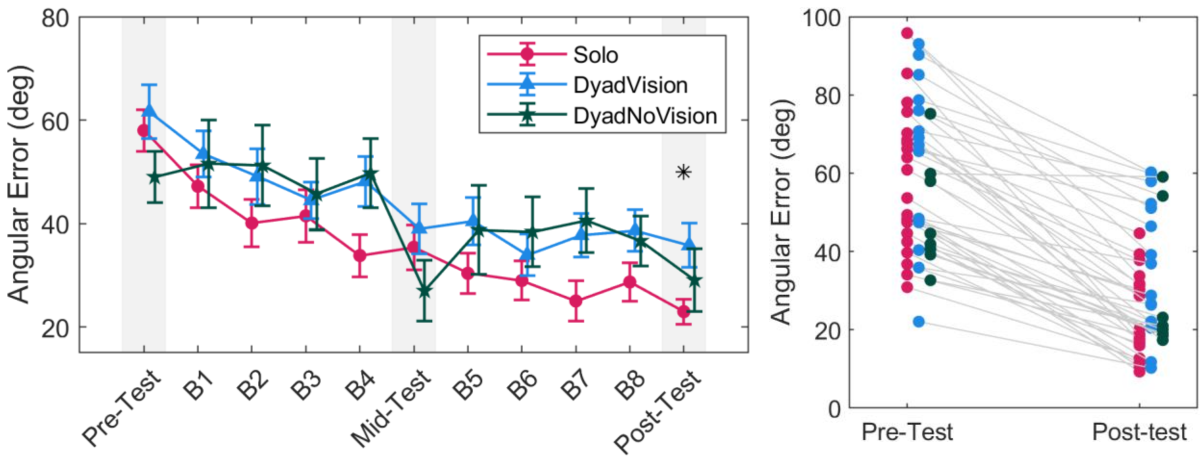

Figure 7 - Average angular error in all blocks between Solo, DyadVision, and DyadNoVision groups. Angular error was lower in the Solo group compared to the DyadVision group in the post-test condition. Error bars represent between-subject standard error of the mean. Grey columns represent testing conditions in which participants were tested alone. B1-B8 indicates training blocks one through eight. Right panel shows individual data for the pre-test and post-test conditions.

Minimal differences observed in mid-test conditions

Outcome Variables

ANCOVA results revealed movement times for novices in the Solo group were significantly lower than those in the Dyad$_{\text{VISION}}$ group in the mid-test conditions (F[1,33] =

5.472, p = 0.026, η_p^2 = 0.142; Figure 6), similar to post-test conditions. No statistically significant differences were observed in angular error in mid-test conditions (F[1,33] = 0.037, p = 0.849, η_p^2 = 0.001; Figure 7).

Coordination Variables

No statistically significant results were found as a result of the ANCOVA for VAF-ratio (F[1,33] = 0.402, p = 0.531, η_p^2 = 0.012; Figure 8), or null-space ratio (F[1,33] = 0.362, p = 0.551, η_p^2 = 0.011; Figure 9) in the mid-test.

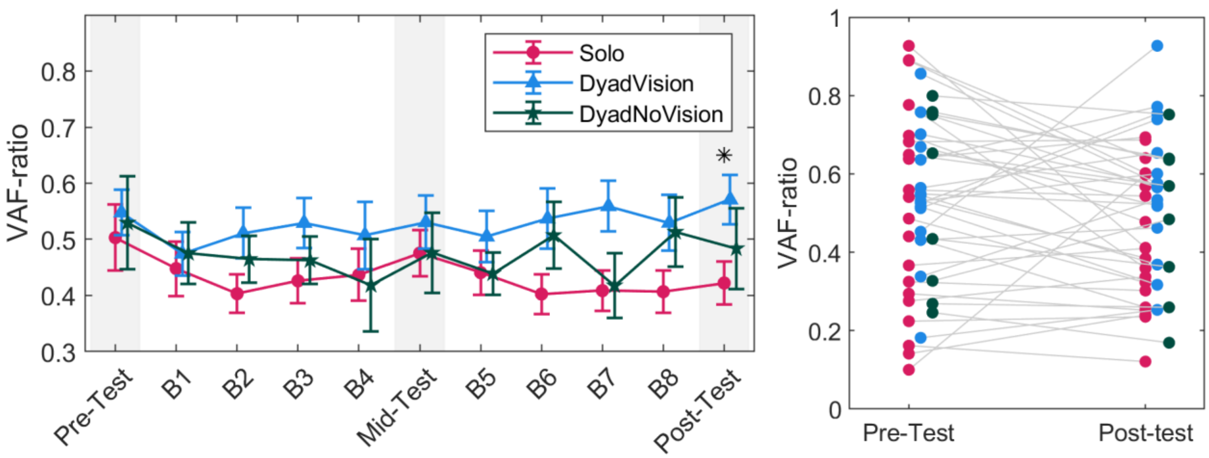

Figure 8 - VAF-ratio in all blocks between the Solo, DyadVision, and DyadNoVision groups. VAF-ratio was significantly lower in the Solo group compared to the DyadVision group in the post-test condition. Error bars represent between subject standard error of the mean. Grey columns represent testing conditions in which participants were tested alone. B1-B8 indicates training blocks one through eight. Right panel shows individual data for the pre-test and post-test conditions.

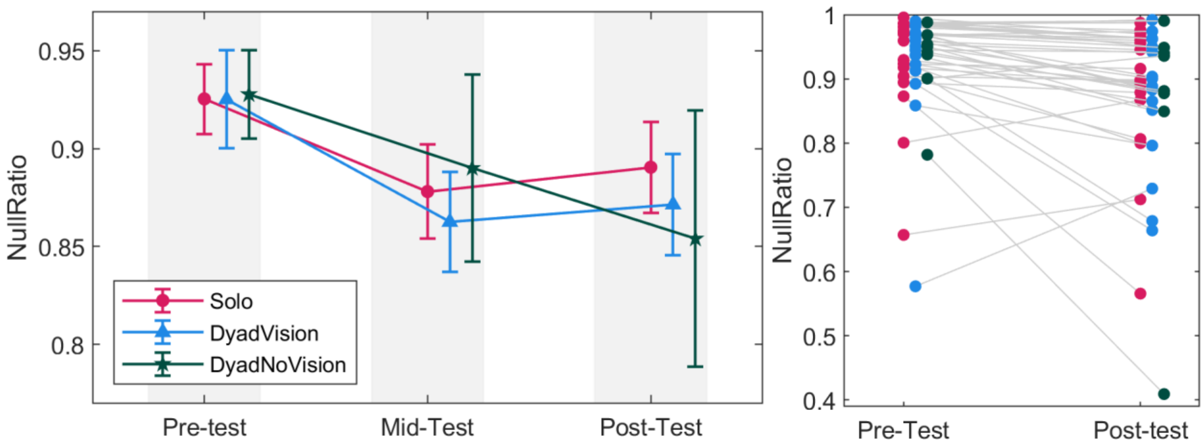

Figure 9 - Per signal null-space ratio during testing blocks for the Solo, DyadVision, and DyadNoVision groups. No significant differences were observed between any groups. Error bars represent between subject standard error of the mean. Grey columns represent testing conditions in which participants were tested alone. Right panel shows individual data for the pre-test and post-test conditions.

Section 3.5. Experiment 2

Methods

We sought to further explore the effects of observational practice in this paradigm by collecting an additional dyadic group ($Dyad_{NOVISION}$, n = 8; 6 female, mean ± SD: 20.5 ± 1.4 years) which was identical to our $Dyad_{VISION}$ group, except that participants in the $Dyad_{NOVISION}$ group were not able to view their partner throughout learning of the task (Fig 1c). All other components of the task were identical to the task performed by the $Dyad_{VISION}$ group.

Experts helped their partners even when out of sight

Similar to Experiment 1, analysis of variance revealed statistically significant differences in pre-test movement times among the three groups ($F(2,33) = 10.330$, $p < 0.001$; Figure 10), with Bonferroni post hoc tests revealing that movement times for the Solo group (10.88 sec, $p < 0.001$) and the novice group (10.65 sec, $p = 0.003$) were significantly higher than expert movement times (1.63 sec), suggesting that experts in the $Dyad_{NOVISION}$ group were sufficiently trained on the task. No statistically significant differences were observed between Solo and novice pre-test movement times (p = 1.000).

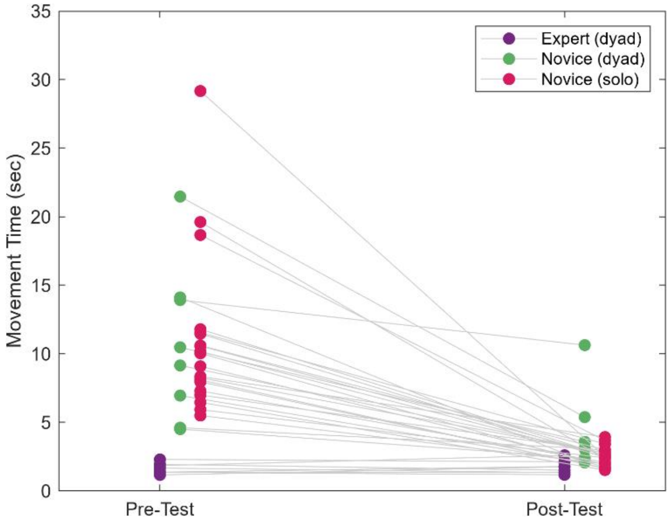

Figure 10 - Movement time comparison between the Expert and Novice in the DyadNoVision group, as well as solo controls. Expert movement times were lower compared to the novice.

We then compared training block movement times between the $Dyad_{NOVISION}$ group and the Solo group to determine if the expert's knowledge of the task positively benefitted training block performance. Movement times for blocks one through six (B1-B6) were found to be significantly lower in the $Dyad_{NOVISION}$ group compared to the Solo group (t_{26} range: -4.199 - -3.246, all $p < 0.05$), while there were no statistically significant differences between the two groups in training blocks seven ($t_{26} = -2.879$, p = 0.064) and eight ($t_{26} = -2.353$, p = 0.208).

No differences between groups in post-test conditions

No statistically significant differences were observed as a result of the ANCOVA in movement times (F[1,21] = 0.341, p = 0.565, $\eta_p^2 = 0.016$; Figure 6), angular error (F[1,21] = 0.112, p = 0.741, $\eta_p^2 = 0.005$; Figure 7), VAF-ratio (F[1,21] = 1.284, p = 0.270, $\eta_p^2 = 0.058$; Figure 8) or null-space ratio (F[1,21] = 0.189, p = 0.668, $\eta_p^2 = 0.009$; Figure 9) in the post-test.

No differences between groups in mid-test conditions

No statistically significant differences were found in movement times ($F[1,21] = 0.093$, $p = 0.763$, $\eta_p^2 = 0.004$, Figure 6), angular error ($F[1,21] = 0.173$, $p = 0.682$, $\eta_p^2 = 0.008$, Figure 7), VAF-ratio ($F[1,21] = 0.435$, $p = 0.517$, $\eta_p^2 = 0.020$, Figure 8) or null-space ratio ($F[1,21] = 0.464$, $p = 0.503$, $\eta_p^2 = 0.022$, Figure 9) during the mid-test.

Section 3.6. Discussion

The purpose of this study was to evaluate the effectiveness of dyadic practice on the learning of a novel Body-Machine Interface (BoMI) task. We hypothesized that dyadic practice would lead to improved task performance compared to learning the task solo. However, contrary to our hypotheses, movement times were lower in the solo control group in the mid- and post-test conditions, and angular error was lower in the solo group in the post-test condition, suggesting that dyadic practice was not as effective at improving performance compared to learning the task solo. Experiment one analyses also demonstrated that post-test VAF-ratios were lower in the solo group compared to the $\text{Dyad}_{\text{VISION}}$ group, suggesting that dyadic practice was less effective in developing movement coordination compared to the solo group.

One possible explanation for the observed learning differences in the $\text{Dyad}_{\text{VISION}}$ group is that experts tended to dominate performance of the task when coupled with their novice partner, providing the novice with insufficient time and/or opportunity to receive feedback on their own movements. Given that experts had a thorough understanding of how to perform the task properly, it is possible that experts were simply able to adapt to the malperformance of their novice partner. If this is the case, this would have caused novices to gain an incorrect assumption of how the task is to be performed: if the expert is able to correct for any novice errors in the moment, the novice has now become reliant on the presence of the expert to complete the task (i.e., they have not developed a movement pattern that would allow them to complete the task on their own). This explanation is further supported by evaluating movement times between testing and training conditions; novice movement times in pre-, mid-, and post-test conditions were higher than movement times observed in training blocks, in which experts had some magnitude of control. This improved performance is likely attributable to the expert's ability to lead the novice through the task despite their lack of experience (effectively dragging their novice partner along as they go through their own motions).

Previous research has found similar results: the improved task performance observed in expert-novice pairs is due mainly to the expert's prior experience on the task, and less to the novice's acquisition of proper task performance (Avila Mireles et al., 2017; Nishimura et al., 2021). Due to the expert's prior experience, this could result in a reduction in effort put forth by the novice during the learning process (Crook & Beier, 2010; Mace et al., 2017). In fact, it has been suggested that the only time this expert-novice pairing is beneficial for learning is when the novice gets prior exposure to the task on their own before being paired with their expert partner (Avila Mireles et al., 2017). While novices in the $\text{Dyad}_{\text{VISION}}$ group did demonstrate improvements in task performance at the mid- and post-test time points (as demonstrated by decreases in movement time), they were unable to reach the same level of performance as the solo controls.

It is also possible that our method of sharing control between participants was not the most effective way of doing so. Our rationale for implementing a 50/50 split between both participants came from the idea that in the early stages of learning, novice participants would directly mimic the actions of their partner in order to complete the task. By allocating control of the task in this 50/50 fashion, a given movement performed by the expert (e.g., forward lean of the trunk 10 degrees) would result in the same magnitude of movement of the cursor if performed by the novice, reinforcing the movements required to properly navigate the cursor. However, this 50/50 split meant it was difficult to quickly attribute movement errors during training, which could have reinforced incorrect movement patterns for the novice.

Additionally, the guidance hypothesis (Salmoni, Schmidt, & Walter, 1984) suggests that while performance feedback can be beneficial to a novice if provided in the right amounts, an excess amount of this feedback can interfere with important between-trial aspects of the learning process such as encoding, storage, or problem solving (Winstein, Pohl, & Lewthwaite, 1994). In the present study, we observed similar results: while the presence of the expert assisted in the performance of the task as a whole (e.g., by keeping movement times during training blocks low), once the expert's contribution was removed in the testing blocks, novice performance worsened as they could no longer rely on the expert to carry them through the task. The presence of the expert in the training blocks likely caused the same issues as excessive feedback: enhancing immediate performance but worsening the learner's ability to learn the task well enough to perform on their own.

It is possible that introducing a "fading" method of allocating expert control could provide greater benefits (Winstein, Pohl, & Lewthwaite, 1994; Aoyagi et al., 2019), in that it would allow for increased expert contribution in the early stages of learning, and then reduce that contribution as learning continues. In similar studies implementing joint dyadic practice, this fading concept provides control to the learner based on their prior performance – if performance is high during early stages of learning, expert contribution in future blocks decreases (Nudehi, Mukherjee, & Ghodoussi, 2005). Future research should consider methods of dynamically modifying the amount of control allocated to the novice during training.

Given our original motivation of using dyadic practice to combine beneficial effects of observational learning and shared control, in our second experiment, we sought to further explore the contribution of observational learning by comparing task performance in the $\text{Dyad}_{\text{VISION}}$ group with the $\text{Dyad}_{\text{NOVISION}}$ group that visually separated the novice learners from their expert partner with a partition. Surprisingly, we found no significant differences as a result of group membership among any of our variables, suggesting that individuals visually separated from their partner were able to learn the task just as well as their counterparts who were able to see one another. It is possible that, in the present task, the main focus of the novice was on how their movements affected the cursor, meaning less attention was given to how their expert partner was moving, or that the movements of their partner were negatively affecting the novice's ability to move on their own. Previous research has suggested that focusing on the effect of the performed bodily movement (known as an external focus, in this case the movement of the cursor) instead of the bodily movement itself (known as an internal focus, in this case the movement of the body) provides greater advantages in motor learning (Wulf, Höß, & Prinz, 1998). It has also been suggested that the movements of an observer are "contagious" to the movements of another individual (referred to as "motor contagion"; Karlinsky, Welsh, & Hodges, 2019), as the observed movements will appear in the behavior of the observer (Roberts et al., 2017), even when remaining still (Tia et al., 2012). Thus, it is possible that the movements of the expert partner become represented internally to the novice via motor contagion, and are thus considered an "internal focus", rendering them less effective compared to the external focus of the cursor. This would explain the similarities in performance between our two dyad groups, as even when placed in view of a trained expert, novices instead chose to focus on the movement of the cursor (which is what novices in the $\text{Dyad}_{\text{NOVISION}}$ group did by default).

This research offers insight on the effectiveness of dyadic practice in cursor-based Body-Machine Interface tasks. Although solo controls demonstrated greater improvements to performance compared to our dyad groups, it is important to note that all of our groups showed improvement relative to their starting point, meaning working on this task in a dyad still resulted in learning benefits. The rationale for beginning with a two-dimensional cursor control task arose from the steep difficulty increases that have been observed in three-dimensional tasks; even tasks that exist in 3D space but are locked to two dimensions of control have shown high rates of attrition among nine- and twelve-year olds during learning (Ranganathan et al., 2019), and tasks that allow for operation in three dimensions often require control of more than three degrees of freedom (e.g., the JACO arm used in [Aspelund et al., 2020] allows for control of seven degrees of freedom: two at the shoulder, one at the elbow, three at the wrist, and an endpoint gripper), which introduces difficulties in how all of these degrees of freedom are simultaneously controlled. By providing evidence that a two-dimensional cursor control task can benefit from dyadic practice, future research can begin to explore more complex tasks to gain a better understand of how dyadic practice can benefit learning of BoMI tasks.

In conclusion, we found that solo practice was a more effective method of learning a novel Body-Machine Interface task compared to dyadic practice. Additionally, we found that visually isolating individuals from their partner during learning did not affect learning outcomes to a significant degree. Given the additional complexities observed when completing a task in greater than two dimensions, future research should implement a similar dyadic practice schedule when operating more complex high degree of freedom interfaces and consider more dynamic methods of allocating control to the learner.

Chapter 4 – Adaptive Strategies During Collaborative Learning of a Body-Machine Interface

Section 4.1. Abstract

The challenge of learning to operate an unintuitive Body-Machine Interface can be mitigated through dyadic practice with an expert. However, previous research implementing dyadic practice to expedite the learning process was not effective when compared to solo practice. A potential cause for this was because the expert had too much control of the task during training, causing novices to become dependent on them. To address this issue, in the present study, we sought to explore strategies to reduce the control given to the expert throughout practice by comparing open- and closed-loop control allocation strategies during a two-dimensional cursor control task. In the open-loop strategy, the control allocated to the expert decreased with practice regardless of the performance of the novice, whereas in the closed-loop strategy, the control allocated to the expert decreased depending on the performance of the novice. These two strategies were compared to a control group where the control allocated to the expert was always constant. Results showed that decreasing the control of the expert in both strategies showed effects on performance, but we found no statistically significant differences between the constant control and open-loop control strategies on learning. There were also no statistically significant differences between the open-loop and closed-loop control strategies on learning. These findings demonstrate the need for novice learners to receive error feedback in the early learning stages and provide insight to the development of effective closed-loop control allocation strategies.

Section 4.2. Introduction

Body-Machine Interfaces (BoMIs) are technologies that capture and transform movements from an individual into control commands for a machine. While these interfaces are often used for common tasks such as general computer usage, they are also employed for control of complex, high degree-of-freedom machines, such as surgical simulators and assistive devices. These devices tend to have low margins for error in their operation, or require learning of unintuitive controls, making them prone to extended learning times or even device withdrawal due to their inherent complexity. Given that ease of use and operability of these kinds of devices has been previously used as a predictor for withdrawal in prosthetic users (Phillips & Zhao, 1993; Biddiss & Chau, 2007), finding methods to simplify the learning of these devices without compromising how effectively they are learned is important.

Previous research has evaluated the effectiveness of dyadic practice on learning these interfaces (see Chapter 3 above), in which two individuals are paired together to learn a task that is ultimately to be performed by a single individual (Crook & Beier, 2010). While dyadic practice has been demonstrated to be an effective method of teaching complex skills in industrial, military, and sport contexts (Shea, Wulf, & Whitacre, 1999), it has typically used one of two distinct paradigms: sequential dyadic practice, in which partners alternate between hands-on practice and observing their partner, or joint dyadic practice, in which both partners contribute to completion of the task simultaneously. Our previous study attempted to combine these two paradigms into a single approach, to mixed results: while this combination of two paradigms did allow novices to learn the task well enough to be performed on their own, the benefits observed were less than if novices were to train entirely on their own.

One explanation for the suboptimal results of dyadic practice lies in how control of the task was allocated to the users: in the previous study, control was divided evenly between participants (e.g., a given movement as performed by the expert would result in the same magnitude of cursor movement if performed by the novice) and stayed constant throughout practice. Therefore, novices in the dyad groups could have become “dependent” on the expert when performing the task (guidance hypothesis, Salmoni, Schmidt, & Walter, 1984), or reduced the amount of effort put forth during training (social loafing, Latané, Williams, Harkins, 1979; Crook & Beier, 2010), making the task more difficult when it came time to perform the task on their own. This effect has been observed in earlier work evaluating expert-novice dyads performing a task simultaneously (Avila Mireles et al., 2017; Mace et al., 2017; Nishimura et al., 2021), suggesting novices are not acquiring sufficient knowledge of the task during training.

To reduce this reliance on the expert, we employed a new strategy of allocating control that instead focuses on fading expert guidance as training progresses (Winstein, Pohl, & Lewthwaite, 1994). In the feedback literature, the “faded practice” approach suggests providing increased frequency of feedback during the early stages of learning as the novice gains familiarity with the task, and gradually reducing the frequency as learning goes on and the novice becomes more comfortable with task performance. A similar concept has been suggested in joint dyadic practice tasks, in which the amount of control given to the expert is high during the early stages of learning, and then reduced (i.e., faded) as the novice learns the task (Nudehi, Mukherjee, & Ghodoussi, 2005; Khademian & Hashtrudi-Zaad, 2011; Li, Tavakoli, & Huang, 2014; Shamaei, Kim, & Okamura, 2015). In order to test this application of faded practice, two methods of control allocation were developed: an open-loop method, which linearly increased the amount of control given to the novice in each training block (regardless of the performance of the novice), and a closed-loop method, that instead provided an amount of control based on the performance of the novice on the previous block. Here, we hypothesize that this faded practice concept, when applied to the magnitude of control given to the novice, will demonstrate greater benefits to learning when compared to a constant control allocation.

Section 4.3. Methods

The task performed here is identical to that which was performed in Chapter 3, with the exception being the groups examined. The task description is repeated here for completeness.

Participants

An additional twenty-five participants were recruited for this study (15 female, 10 male) aged 18-35 years (mean +/- SD: 24.80 +/- 4.33 years). Participants reported no injuries that would affect their ability to move their upper body through a comfortable range of motion. All but four participants self-reported as right-hand dominant. Informed consent was obtained from all participants prior to participation, and all data collection procedures were approved by Michigan State University's IRB.

Two additional groups were collected in Chapter 4: a $\text{Dyad}_{\text{GRADUAL}}$ group (n = 12), and a $\text{Dyad}_{\text{ADAPTIVE}}$ group (n = 13). As the present study was a part of a larger data collection project, these two groups were compared to the $\text{Dyad}_{\text{VISION}}$ group from Chapter 3 (n = 16), which will be referred to as "$\text{Dyad}_{\text{CONSTANT}}$" throughout the remainder of Chapter 4. Following completion of a demographic questionnaire, individuals were paired with a trained research assistant to complete the given task (described in the following section) collaboratively. The three research assistants (hereby referred to as the "expert") trained on the task on their own, and their performance was evaluated by the primary researcher to ensure they learned the task well enough to assist their partner.

Task

Four triaxial inertial measurement unit (IMU) sensors (YOST Labs, Portsmouth, Ohio) were affixed to the anterior and posterior aspects of both individual's shoulders (eight IMUs within a dyad, four per person). Roll and pitch angles from each sensor were used to control the onscreen cursor (yaw angles were omitted due to their sensitivity to magnetic fields) via a custom Simulink program. Roll and pitch signals from each of the sensor were converted into a set of Cartesian coordinates for the screen cursor as follows (Farshchiansadegh et al., 2014):

$$X = A\,H + X_0$$

in which X is the Cartesian coordinates of the cursor (sized [2 x 1]), A is the mapping matrix (sized [2 x 8]) used to transform the set of IMU signals H (sized [8 x 1]), and X_0 (sized [2 x 1]) is an offset term so that the mean resting posture corresponded to the center of the screen. In the dyadic group, this calculation was done on both participants to obtain a virtual cursor position for each participant, and a weighted average position of the two cursors was calculated to obtain the cursor position:

$$X = \lambda\, X_{expert} + (1 - \lambda)\, X_{novice}$$

in which λ is a weighting factor used to assign a percentage of control to the expert that could vary between 0 and 1. Both groups began the first training block by assigning the expert with complete control, and the final training block by assigning the novice with complete control. For the novices in the $\text{Dyad}_{\text{GRADUAL}}$ group, this provided a near-linear increase in their percentage of control throughout the training duration. In the $\text{Dyad}_{\text{ADAPTIVE}}$ group, lambda values were instead calculated based on the magnitude of the novice's average angular error on the previous training block, calculated as:

$$\lambda_i = \frac{Novice\ Error_{i-1}}{Solo\ Error_{i-1}} * \lambda\ Gradual_i$$

in which Novice Error_{i-1} is the Novice's angular error from the previous training block, Solo Error_{i-1} is the angular error of our solo controls from Chapter 3 on the same block, and $\lambda\text{Gradual}_i$ is the amount of control allocated to our $\text{Dyad}_{\text{GRADUAL}}$ group in the upcoming block. In the event that the novice performed worse than a solo performer on a given block, the novice was provided the same linear increase provided to the $\text{Dyad}_{\text{GRADUAL}}$ group on the same block. This was to ensure that novices did not regress beyond those in the $\text{Dyad}_{\text{GRADUAL}}$ group in the amount of control they were allocated during training.

We pilot tested this algorithm in a controlled setting using two of our experts paired in a dyad. One of the experts was asked to serve as a "novice-agent," in which they would purposefully perform the task well or poorly at the request of the researcher to test the algorithm. When asked to perform the task "poorly," the novice-agent would move in the opposite direction of the presented target (i.e., increasing their angular error), and when asked to perform the task well, the novice-agent would move as they normally would (i.e. reducing their angular error). The algorithm performed as expected in this controlled test: novice control adhered to the baseline (floor) values after performing the task poorly, and increased dramatically after performing the task well.

A critical difference from previous paradigms was that in prior work, the map A was customized to each participant based on exploratory movements of their own body (see the "body dance" from Lee et al., 2016). However, in the present paradigm, because there were two individuals involved in the dyad groups, we used a constant pre-defined map for all participants that was developed by the experimenter. This allowed the task to still be novel, but also created a "fair comparison" for all groups.

To examine the effect on motor learning, two conditions were implemented in this study: testing blocks and training blocks. The training blocks differed based on the group – i.e., the solo group practiced these blocks on their own, and the dyad group had both the expert and novice jointly perform the task. In the testing blocks, all participants perform the task entirely on their own. The performance in the test blocks were our measure of "learning" since they only involved the learner. The full schedule of the data collection session can be found in Figure 1 in Chapter 3.

Procedure

Participants were seated in front of a 23.8 inch computer monitor (Dell P2414H, resolution 1920x1080, refresh rate 59Hz) (Fig 2a in Chapter 3), and individuals in the dyad groups were seated side-by-side with one another, with the expert always to the left of the novice when facing the screen (Fig 2b in Chapter 3). Participants were tasked with navigating an onscreen cursor from a home position at the center of the screen to a set of presented targets as quickly and accurately as possible (Fig 3 in Chapter 3).

Upon affixing the sensors to participants, a familiarization test was performed to ensure the IMUs were functioning properly: participants were asked to move their upper body "any way they could think of while remaining seated" so that an onscreen cursor touched all four sides of the program window. Participants were informed that they were not timed or scored on the familiarization test, and that it only served to "make sure the task is doable." No other information on how to complete the task was provided. In the dyadic groups, both the expert and novice shared equal control of the onscreen cursor, and the expert was instructed (prior to the novice's arrival) to accommodate the novice's movements to the best of their ability.

Following the familiarization test, participants performed a pre-test, which served as a baseline for the individual's learning of the task. The testing conditions tasked each participant with navigating the onscreen cursor from the home position to one of eight randomly presented targets as quickly and accurately as possible, positioned equidistant from the home position (Fig 3). Each target was presented three times, for a total of 24 reaches in the testing conditions. This same test was performed halfway through the training (mid-test) as well as at the very end of all training blocks (post-test).

In the training blocks, dyad groups were asked to perform the cursor control task collaboratively, sharing control of a single cursor (similar to the familiarization test) according to the lambda values described in the Task subsection above. While the goal of the training blocks was identical to that of the testing condition, only four targets were presented instead of eight, with each target appearing a total of five times for a total of 20 trials per block. Eight blocks were completed in total, for a total of 160 training trials. To restrict the type of interaction between the participants primarily to observational learning, participants were instructed not to share direct instructions with their partner (e.g., commands such as "lean forward to make the cursor move to the current target" were discouraged), but other forms of communication (such as general encouragement and chatter) were allowed and encouraged to create cohesion within the dyad.

Data Processing

Task Performance

Movement time of the cursor was calculated for each test condition and training block, calculated from the frame the cursor leaves the home position for the first time to the frame in which the target position is achieved and maintained for 0.5 seconds. Angular error was calculated as the angular displacement of the individual's cursor relative to the presented target 0.5 seconds after leaving the home position (an angular error of 0 indicates that the cursor was aligned with the line joining the home position and the presented target). These metrics were chosen as a method of quantifying performance – as all participants were instructed to move to the target position as quickly and accurately as possible, decreases in movement times and angular error are used to suggest learning of how to perform the task. Movement time was calculated using the movement of the true cursor, while angular error was calculated using the novice's signals (e.g., the angular error of the novice's virtual cursor).

Movement Coordination

The ratio of the variance accounted for (VAF) by the first two principal components (PCs) was calculated using the roll and pitch signals during the testing conditions and each individual training block. This ratio (VAF-ratio) was calculated as (VAF2/VAF1) and used to provide a metric of exploration of the task (i.e., a smaller ratio indicates a greater reliance on PC1 [Lee & Ranganathan, 2019]). Task- (H_{task}) and null- (H_{null}) space variances were computed by first dividing the home position postures into their task- and null-space components, like so (Lokesh & Ranganathan, 2019):

$$H_{task} = A' * (A * A')^{-1} * A * H$$

$$H_{null} = (I_8 - A' * (A * A')^{-1} * A) * H$$

in which H is a matrix consisting of the eight IMU signals at the 24 instances in which the home target was achieved, A is the mapping matrix defined in the "Task" subsection, and I8 is an identity matrix sized 8x8. These signals were only taken at the instance in which the home position was achieved, as each outer target was only presented three times during each testing condition. These components were then used to calculate the "per-DoF" null-space ratio (NullRatio; Scholz & Schöner, 1999), which was calculated as:

$$NullRatio = \frac{Var(H_{null})/6}{Var\ (H_{null})/6 + Var(H_{task})/2}$$

in which the variances about the task- and null-space are divided by the number of relevant signals in each dimension (i.e., two dimensions in the task-space, and six dimensions in the null-space). A ratio of 1 indicates equal distribution of variances along the task and null spaces (i.e., no preferential exploration). In the dyadic conditions, similar to the angular error in the previous section, VAF-ratio and task- and null-space variances were calculated using the novice's eight IMU signals.

Statistical Analysis

To ensure experts had learned the task well enough to help their novice partner, analysis of variance was performed on pre-test movement times between the Solo group, the trained expert, and the novice partner within each dyad, with Bonferroni post hoc tests performed as appropriate. Next, to examine if the expert's contribution benefitted novice performance during training, we performed independent samples t-tests with Bonferroni corrections on training block movement times between the Solo group and the $\text{Dyad}_{\text{GRADUAL}}$ group, as well as between the Solo group and $\text{Dyad}_{\text{ADAPTIVE}}$ group. Alpha for these tests was set to .05.

One-way analysis of covariance was performed to analyze differences as a result of how control was allocated to the learner (group membership), using the pre-test values for a given variable as the covariate to account for differences in initial performance. Significance was set to 0.05, and partial eta-squared values were calculated for effect size (in which values of .01 to .05 are considered small, values of .06 to .13 are considered medium, and values of .14 or higher are considered large [Statistical Analysis in JASP: A Guide for Students, 2018]). JASP Software (University of Amsterdam, Netherlands) was used to calculate statistics.

Section 4.4. Results

While examining the data, one participant in the $Dyad_{ADAPTIVE}$ group demonstrated significantly worse performance compared to the rest of the group, and as such was excluded from statistical analyses, leaving us with ($n = 12$) participants in the $Dyad_{ADAPTIVE}$ group.

Experts Assisted Novice Performance

Analysis of variance revealed statistically significant differences in pre-test movement times between the Solo group, novices, and experts within the $Dyad_{GRADUAL}$ group ($F(2,41) = 20.503$, $p < 0.001$; Figure 11), with Bonferroni post hoc tests demonstrating significantly higher pre-test movement times in the Solo group (10.88 sec, $p < 0.001$) and novices (13.92 sec, $p < 0.001$) when compared to the experts (1.40 sec). This suggests that experts were sufficiently trained on the task. No statistically significant differences were observed between novice and Solo pre-test movement times ($p = 0.324$).

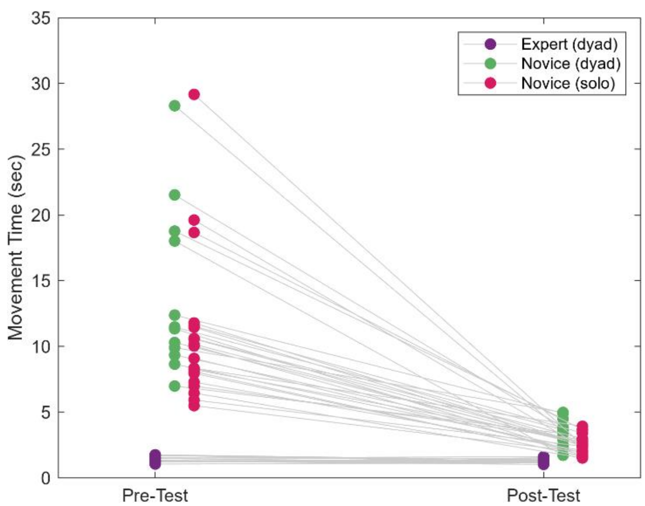

Figure 11 - Movement times in pre- and post-test conditions for Experts and Novices in the DyadGradual group, and Solo controls. Movement times were significantly shorter for experts compared to novices and solo controls.

An analysis of variance was performed between the Solo group and individuals within the $Dyad_{ADAPTIVE}$ group, and showed statistically significant differences in pre-test movement times between Solo controls, experts, and novices ($F(2,41) = 18.615$, $p < 0.001$; Figure 12), with Bonferroni post hoc analyses finding significantly higher pre-test movement times for the Solo (10.88 sec, $p < 0.001$) and novice groups (12.51 sec, $p < 0.001$) when compared to the experts (1.17 sec). No statistically significant differences were observed between novice and Solo pre-test movement times ($p = 1.000$).

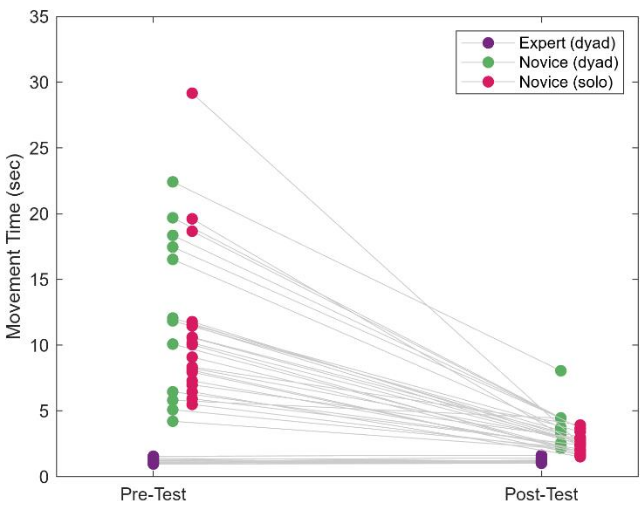

Figure 12 - Movement times in pre- and post-test conditions for Experts and Novices in the DyadAdaptive group, and Solo controls. Movement times were significantly shorter for experts compared to novices and solo controls.

Next, we compared movement times of the dyads during the training blocks to those of the solo controls to ensure that the expert was able to complete the task even when the novice was contributing to the task. Movement times during the training blocks were found to be significantly lower in the $Dyad_{GRADUAL}$ group compared to the solo control during blocks 1 through 6 (t_{30} range: -3.979 to -5.488, all $p < 0.05$), while no statistically significant differences were found in block 7 ($t_{30} = 1.080$, $p = 1.000$) or block 8 ($t_{30} = 1.605$, $p = 0.952$).

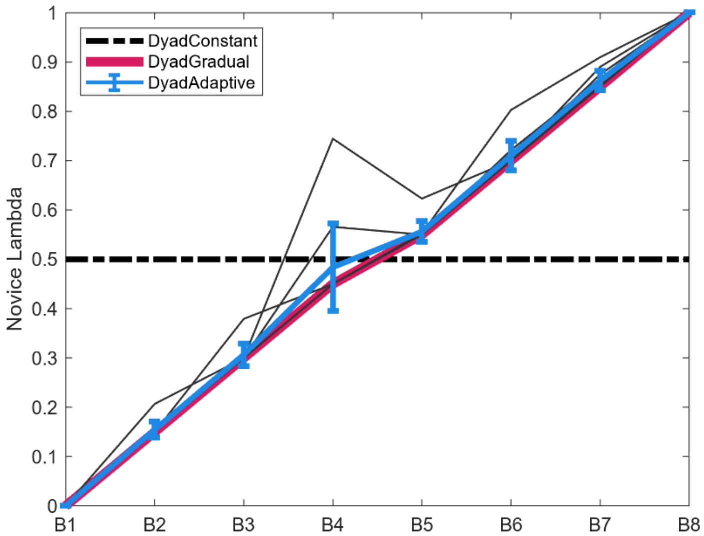

Figure 13 - Novice lambda values for the DyadConstant, DyadGradual, and DyadAdaptive groups.

Similar results were found in our Dyad$_{\text{ADAPTIVE}}$ group: significant differences in movement time were observed in blocks 1 through 5 (t_{30} range: -3.859 to -5.724, all $p < 0.05$), with dyads demonstrating reduced movement times in these blocks. No significant differences were observed in block 6 ($t_{30} = -1.898$, $p = 0.536$), block 7 ($t_{30} = 1.000$, $p = 0.468$), or block 8 ($t_{30} = 2.147$, $p = 0.32$). The lack of significant findings in the later blocks is likely due to the expert's control of the task decreasing, meaning their ability to compensate for novice errors is reduced. Lambda values for the Dyad$_{\text{GRADUAL}}$ group and average values for the Dyad$_{\text{ADAPTIVE}}$ group (as well as the values for each individual in the Dyad$_{\text{ADAPTIVE}}$ group) are displayed in Figure 13. A sample of novice cursor trajectories for the Dyad$_{\text{GRADUAL}}$ and Dyad$_{\text{ADAPTIVE}}$ groups are presented in Figure 14.

No observed differences between groups in post-test conditions

Outcome Variables

No statistically significant differences were found between the Dyad$_{\text{CONSTANT}}$ and Dyad$_{\text{GRADUAL}}$ groups as a result of the ANCOVA for movement times ($F[1,25] = 1.734$, $p =$

0.200, η_p^2 = 0.065; Figure 15) or angular error (F[1,25] = 3.203, p = 0.086, η_p^2 = 0.114; Figure 16). There were also no statistically significant differences found between the Dyad$_{GRADUAL}$ and Dyad$_{ADAPTIVE}$ groups for movement times (F[1,21] = 1.083, p = 0.310, η_p^2 = 0.049; Figure 15) or angular error (F[1,21] = 1.091, p = 0.308, η_p^2 = 0.049; Figure 16).

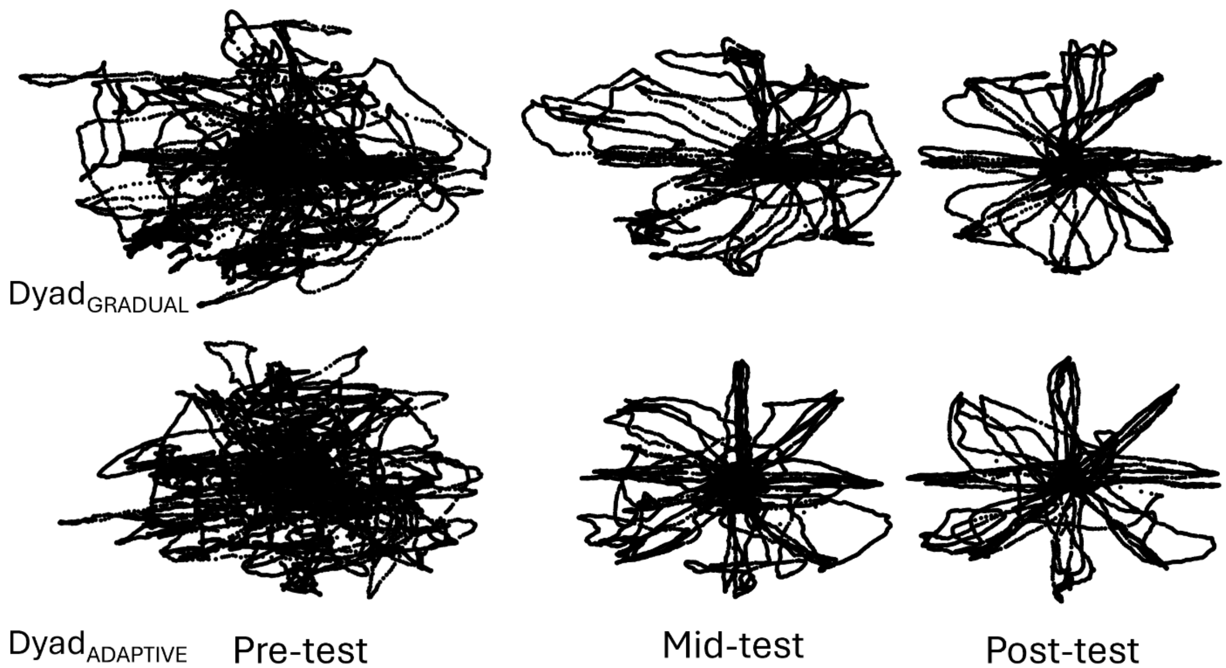

Figure 14 - Sample cursor trajectories for a novice in the DyadGradual and DyadAdaptive groups for the Pre-, Mid-, and Post-test conditions.

Coordination Variables

No statistically significant differences were observed between the Dyad$_{CONSTANT}$ and Dyad$_{GRADUAL}$ groups as a result of the ANCOVA for VAF-ratio (F[1,25] = 1.698, p = 0.204, η_p^2 = 0.064; Figure 17) or null-space ratio (F[1,25] = 1.506, p = 0.231, η_p^2 = 0.057; Figure 18). There were also no statistically significant differences between our Dyad$_{GRADUAL}$ and Dyad$_{ADAPTIVE}$ groups for VAF-ratio (F[1,21] = 0.157, p = 0.696, η_p^2 = 0.007; Figure 17) or null-space ratio (F[1,21] = 0.349, p = 0.561, η_p^2 = 0.016; Figure 18).

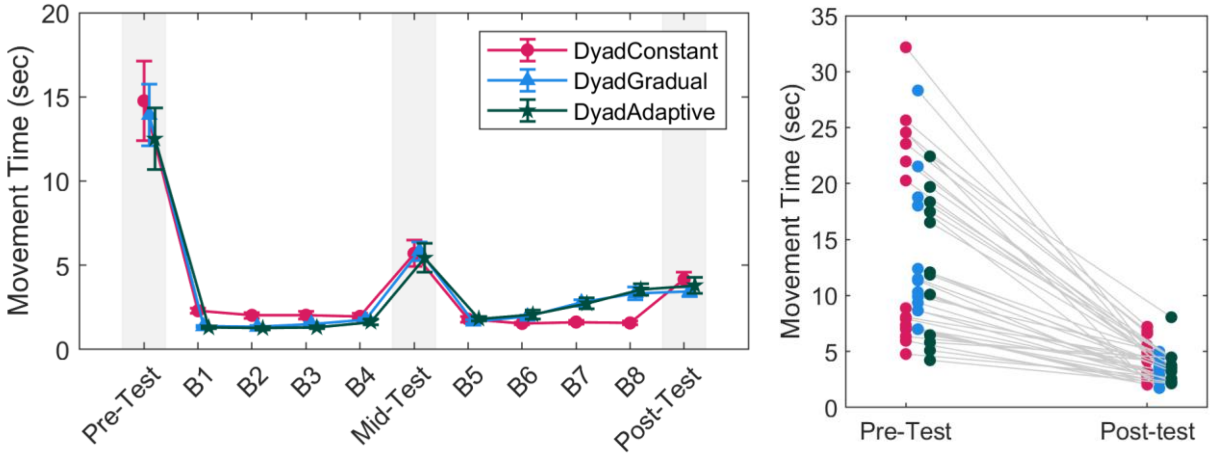

Figure 15 - Average movement times in all blocks between the DyadConstant, DyadGradual, and DyadAdaptive groups. No significant differences were observed between any groups. Grey columns represent testing conditions in which participants were tested alone. B1-B8 represents training blocks one through eight. Right panel shows individual data for the pre-test and post-test conditions.

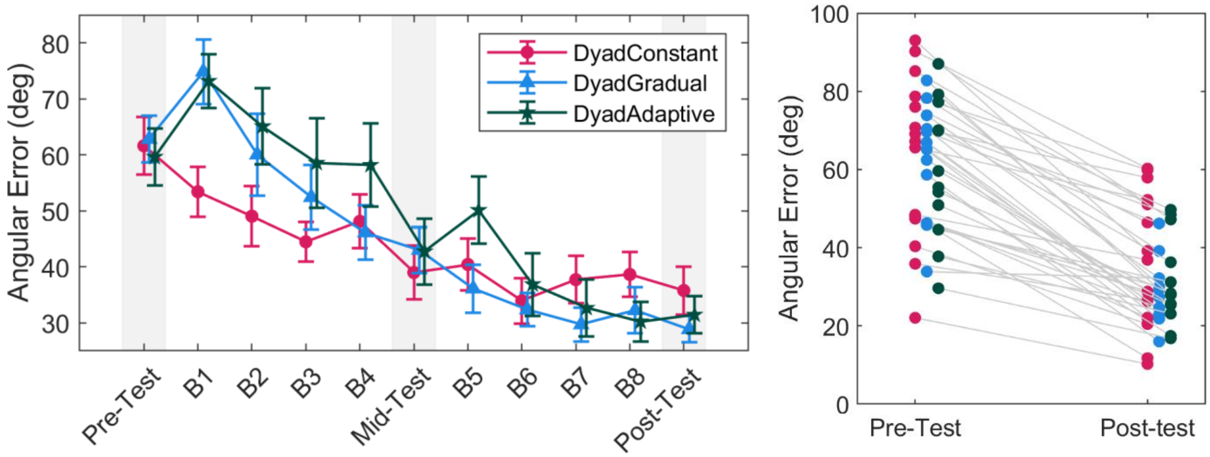

Figure 16 - Average angular error in all blocks between the DyadConstant, DyadGradual, and DyadAdaptive groups. No significant differences were observed between any groups. Grey columns represent testing conditions in which participants were tested alone. B1-B8 represents training blocks one through eight. Right panel shows individual data for the pre-test and post-test conditions.

No differences between groups in mid-test conditions

Outcome Variables

No statistically significant differences were observed between the Dyad$_{\text{CONSTANT}}$ and Dyad$_{\text{GRADUAL}}$ groups during the mid-test for movement times ($F[1,25] = 0.072$, $p = 0.790$, $\eta_p^2 = 0.003$, Figure 15) or angular error ($F[1,25] = 0.429$, $p = 0.519$, $\eta_p^2 = 0.017$, Figure 16). There were also no statistically significant differences found between the Dyad$_{\text{GRADUAL}}$ and Dyad$_{\text{ADAPTIVE}}$ groups for movement times ($F[1,21] = 0.012$, $p = 0.915$, $\eta_p^2 < 0.001$, Figure 15) or angular error ($F[1,21] = 0.090$, $p = 0.768$, $\eta_p^2 = 0.004$, Figure 16).

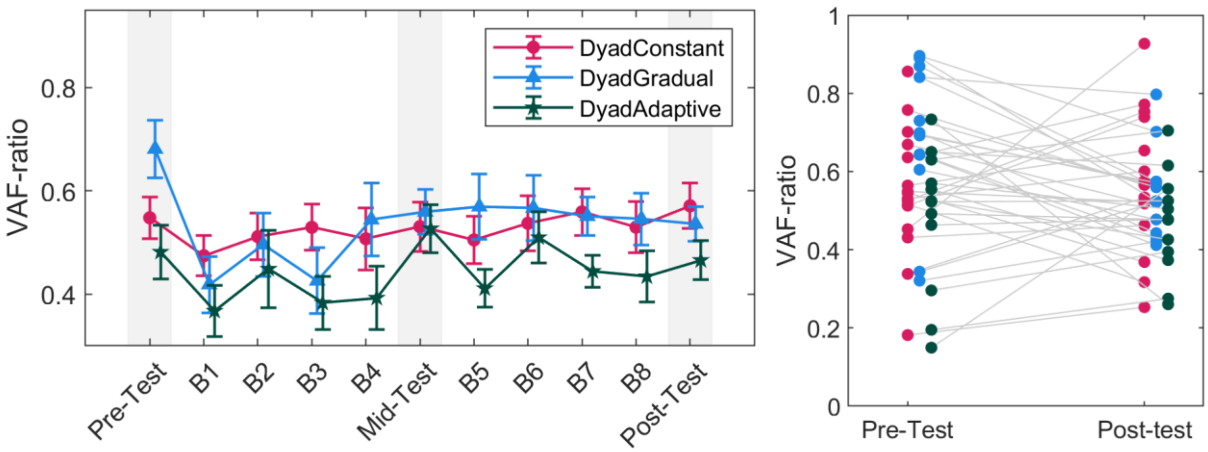

Figure 17 - VAF-ratio in all blocks between the DyadConstant, DyadGradual, and DyadAdaptive groups. No significant differences were observed between any groups. Grey columns represent testing conditions in which participants were tested alone. B1-B8 represents training blocks one through eight. Right panel shows individual data for the pre-test and post-test conditions.

Coordination Variables

No statistically significant differences were observed between the Dyad$_{\text{CONSTANT}}$ and Dyad$_{\text{GRADUAL}}$ groups during the mid-test for VAF-ratio ($F[1,25] = 0.634$, $p = 0.433$, $\eta_p^2 = 0.025$, Figure 17) or null-space ratio ($F[1,25] = 0.104$, $p = 0.749$, $\eta_p^2 = 0.004$, Figure 18). There were also no statistically significant differences found between the Dyad$_{\text{GRADUAL}}$ and Dyad$_{\text{ADAPTIVE}}$ groups for VAF-ratio ($F[1,21] = 0.657$, $p = 0.427$, $\eta_p^2 = 0.030$, Figure 17) or null-space ratio ($F[1,21] = 0.481$, $p = 0.496$, $\eta_p^2 = 0.022$, Figure 18).

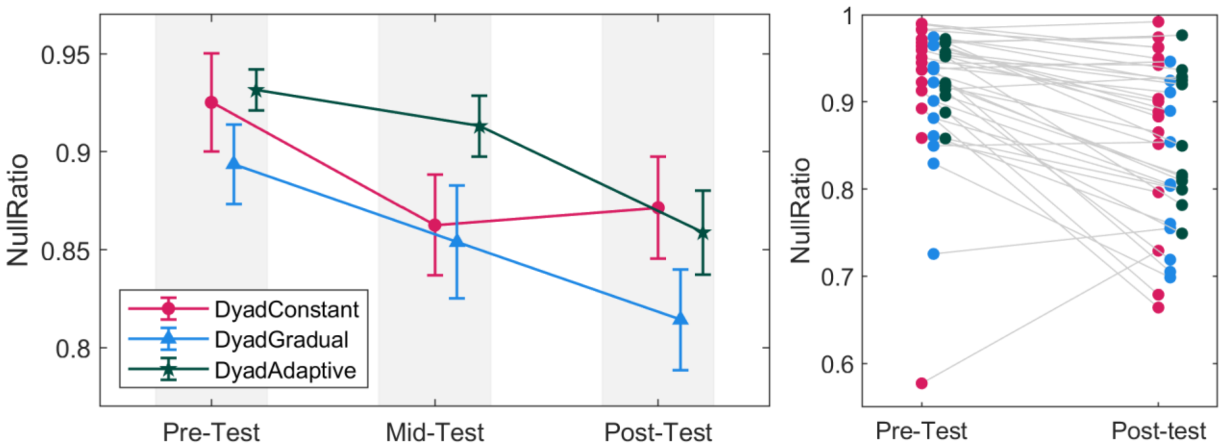

Figure 18 - Per signal null-space ratio during testing blocks for the DyadConstant, DyadGradual, and DyadAdaptive groups. No statistically significant differences were observed between any groups. Grey columns represent testing conditions in which participants were tested alone.

Section 4.5. Discussion

In this study, we evaluated the effectiveness of two methods of allocating control to a novice participant during a shared control task: an open-loop method that increased control linearly to the learner as training progressed, and a closed-loop method that provided control based on performance in a previous block. We hypothesized that these methods of decreasing the expert control during training would result in greater learning benefits than a constant control scheme. Our findings suggest no statistically significant differences between the Dyad$_{\text{CONSTANT}}$ and Dyad$_{\text{GRADUAL}}$ groups, nor any differences between the Dyad$_{\text{GRADUAL}}$ and Dyad$_{\text{ADAPTIVE}}$ groups. There are a few possible explanations for these findings.

Regarding our Dyad$_{\text{GRADUAL}}$ group, it is possible that no differences were observed relative to the Dyad$_{\text{CONSTANT}}$ group due to how we chose to allocate control in the early stages of the learning process. To emulate the "see one, do one" concept that has been used in observational practice, the first training block that our novices were exposed to provided them with no control whatsoever, instead providing the experts with total control of the task. Novice participants were not explicitly told this, but were able to tell rather quickly that their movements were not contributing to the task. Previous research in dyadic practice has suggested that working on a task within a dyad can cause participants to under-perform due to an expectation for their partner to "pick up the slack" (known as "social loafing," Latané, Willimas, and Harkins, 1979) which, in the present task, could have set a negative precedent for the novice participant early (as demonstrated by one participant moving completely erratically as to "prove a point" on the first block; this participant would later go on to say "it seems like you're giving me more control each time"). This social loafing idea has been shown to persist even when novices are informed that they will be tested on their ability to perform the task entirely on their own (Crook & Beier, 2010), and given that similar results manifested in our study (i.e., participants were informed early and often that pre-, mid-, and post-test conditions would be performed without the assistance of a partner), it is possible that this risk of future solo performance was not sufficient to reduce loafing during training blocks.

Similarly, it is possible that the reduced responsibility during the early stages of learning contributed to the development of incorrect movement patterns that later needed fixing. When presented with insufficient feedback on how their movements affect task performance, while simultaneously observing low movement times during training, it is possible that whatever movement patterns adopted by the novice early were the patterns they associated with proper performance. The only time these movements were "checked" for accuracy was during the testing conditions, during which novices would quickly be exposed to the errors of their movement strategy. It is possible that their magnitude of control was not enough for them to notice a difference until the second half of training, in which novices had a majority of control of the task – one participant commented in training block six that "it feels like my movements do more in the later training, [like] my cursor is more sensitive." This could have caused novices to develop an early reliance on the expertise of their expert partner, similar to our findings in Chapter 3. This is further explained by our initial findings when comparing differences in movement time during the training blocks: experts were able to compensate for novice errors (and thus keep movement times low) only through approximately block 5. At the onset of Block 6, novice participants would have (at minimum) 70% control of the task, and it was during this block that movement times during the dyad were no longer significantly different from those of the solo controls, suggesting it was around this time that novices truly felt “in control” of the task.

In an attempt to combat this, we developed an algorithm that instead allocated control to the novice based on their performance on the previous block – if participants’ error on a given block was low, they were given more control of the task on the next block. This error metric was based on the angular error of the novice’s cursor during the initial half-second of the reach, and was compared to the angular error of the solo participants from Chapter 3 on the same block (which acted as our “gold standard” for performance). Despite our efforts, it is possible that this method of closed-loop allocation instead rewarded those who loafed during training: by reducing effort during training, participants were presented with reduced responsibility on the next block compared to those who put forth increased effort and thus received more responsibility. As shown in Figure 13, most participants in the Dyad$_{\text{ADAPTIVE}}$ group coasted on the baseline control magnitude, and even those that “beat” the algorithm to get additional responsibility did not necessarily maintain that increased level of responsibility throughout the entire training.

It has been suggested that tasks in which both participants have full control of all components are more susceptible to social loafing, as theoretically one participant could perform the task entirely on their own while their partner remains static (Karau & Williams, 1993). This tendency to loaf seems to stem from a lack of uniqueness or individuality in the tasks allocated to the participants; essentially, if an individual's performance is not directly measured/identifiable in the task performance, the individuals will have an increased tendency to loaf (Williams et al., 1989; Swain, 1996). One method of addressing this lack of individuality in collaborative tasks is to instead employ a control scheme in which each participant has control of a discrete aspect of the task that renders the full task impossible to complete without contribution of the partner (referred to as the Active Interlocking Model [AIM], Shebilske et al, 1992). This allows for a similar dyadic practice paradigm to take place in which participants are both performing the task themselves and observing the movements of their partner, but the discrete aspects that make the entirety of the task reduce the likelihood of loafing as the learner's partner may only be able to "pick up the slack" to a certain degree (Mace et al., 2017). Interestingly, one novice participant in our study asked if they could communicate with their partner in such a way to allocate control akin to the AIM, asking "can I talk to my partner to split up the directions? Like I take up and down, and they take left and right?". This communication was discouraged by the researcher, on the grounds that it would be considered "providing instructions" to the partner, but provides interesting insight into how participants would simplify the task if given the option.

As the tendency to loaf is based (at least partially) on the ability to hide behind total task performance, another method of reducing social loafing is to identify the learner's performance/contribution throughout performance of the task. In the present study, participants were not shown any feedback on their own performance, only the performance of the collaborative task, meaning if movement times were low, they could just as easily (in the moment) be attributed to the novice learning the task as they could to the expert dragging the novice along behind them. By strictly identifying how the learner contributed to the task, they lose the ability to hide behind the overall task performance, making them less inclined to loaf. This would, in turn, address our previous issue in which novice participants lacked sufficient feedback on their own movements: by identifying individual performance as it contributed to total task completion, novices will be less inclined to loaf (as their contribution to the task will be revealed) and will receive sufficient feedback on their own performance to make proper corrections in future attempts.

While the algorithm we developed performed well in controlled testing, it relied on the assumption that a straight-line cursor trajectory was a direct correlate for improved performance, which may have negatively affected how control was allocated. The instructions given to the participants were to move the cursor to the target "as quickly and accurately as possible," but offered nothing about the trajectory taken by the cursor itself. Previous research has suggested that when subjected to perturbations in a reaching task (such as a force-field or visuomotor rotation), trajectories taken are less likely to align with rectilinear axes and instead follow a curved trajectory, and that the accuracy of a reaching movement is not necessarily predicated on rectilinearity of hand paths (Scheidt & Ghez, 2007). It is possible that similar effects were observed in the present task, in which participants settled on some trajectory that allowed them to move quickly and accurately to the presented target while simultaneously moving in a curvilinear fashion, causing our algorithm to falsely assume malperformance. With this in mind, future attempts at quantifying knowledge of proper task performance should focus less on the straightness of the path and instead consider focusing on the endpoint of the novice's virtual cursor upon achieving the presented target: as the magnitude of control given to each participant is calculated after calculating the location of each participant's virtual cursor, the trajectory and location of each participant's virtual cursor exists "behind the scenes," regardless of the amount of control given to each person. If, instead of looking at how closely the novice's trajectory resembles a straight line, the endpoint accuracy of the novice's virtual cursor is collected once the true cursor reaches the presented target, we could observe how the novice would have controlled the task on their own during that same trial.

This research offers additional insights into how to structure dyadic practice in order to encourage learning of Body-Machine Interfaces. In the present study, we employed a "fading" method of reducing expert contribution in an attempt to reduce novice reliance on their expert partner. While no statistically significant differences were observed between the $\text{Dyad}_{\text{CONSTANT}}$, $\text{Dyad}_{\text{GRADUAL}}$, or $\text{Dyad}_{\text{ADAPTIVE}}$ groups, all groups demonstrated improvements to performance relative to their pre-test performance, suggesting that each of these methods of allocating control were effective methods of learning the task. It is possible that our method of implementing fading actually put the novice at a disadvantage, as we began both of our faded groups ($\text{Dyad}_{\text{GRADUAL}}$ and $\text{Dyad}_{\text{ADAPTIVE}}$) with complete expert control in Block 1, which appeared to discourage some novices early, or at the very least make their movements in the first block appear inconsequential. Given that training block movement times in both these faded groups stayed consistently low until around Block 6 or so (Figure 18), it could be that the proper starting point for this faded practice sits around a 70:30 ratio between novice and expert control: enough control for the novice to feel like they are actively contributing to the task, and enough control for the expert to be able to "wrangle" the novice in if needed.

In summary, we found that all three groups were effective methods of learning the Body-Machine Interface task. Our findings suggest that novices should be provided more control of a joint dyadic practice task in the early stages in order to receive sufficient feedback on their movements, reduce reliance on the expert, and minimize the risk of loafing. Future research should continue to explore methods to reduce the risk of loafing in these overlapping tasks.

Chapter 5 – Discussion

Section 5.1. Summary of Findings

The purpose of this dissertation was to assess the effectiveness of a joint dyadic practice paradigm on learning a novel Body-Machine Interface (BoMI) task. This paradigm sought to combine two unique implementations of dyadic practice – sequential dyadic practice, in which the learner observes the movements of their partner during rest periods, and joint dyadic practice, in which both participants are given control of a given task simultaneously. This was accomplished by developing a paradigm that allowed us to modify both the degree of observational learning the novice could engage in (by placing their expert partner either in- or out-of-sight) as well as how control of the task was allocated to the novice (giving both participants equal control, providing a linear control increase to the novice, or by allocating control based on prior performance).

Our first study looked at the effectiveness of sequential dyadic practice on a novice's ability to learn a novel Body-Machine Interface task. Three groups were collected: a solo control group, a $\text{Dyad}_{\text{VISION}}$ group, and a $\text{Dyad}_{\text{NOVISION}}$ group. We anticipated that the $\text{Dyad}_{\text{VISION}}$ group would show greater benefits to learning relative to the solo controls, as novices in the $\text{Dyad}_{\text{VISION}}$ group could observe the movements of their expert partner to quickly learn the appropriate movement pattern. However, this was not the case: solo controls ended up learning the task better than the $\text{Dyad}_{\text{VISION}}$ group, as demonstrated by decreased movement time and angular error on testing conditions, as well as increased movement coordination, as demonstrated by decreased VAF-ratios compared to the $\text{Dyad}_{\text{VISION}}$ group.

Next, we sought to evaluate just how effective sequential dyadic practice was within the dyads by collecting an additional $\text{Dyad}_{\text{NOVISION}}$ group that was visually separated from their partner so as to restrict the novice's ability to view their partner's movements, causing them instead to rely of the visual feedback of the cursor's position onscreen. We expected the $\text{Dyad}_{\text{NOVISION}}$ group to demonstrate decreased performance as a result, but surprisingly found no significant differences between the $\text{Dyad}_{\text{VISION}}$ and $\text{Dyad}_{\text{NOVISION}}$ groups, suggesting that the ability to directly observe the partner's movements did not impact learning of the task. One possible explanation (that will be expanded upon later in this chapter) is that when given equal control of the task (i.e., a movement of a given magnitude as performed by the expert will result in the same cursor deviation if performed by the novice), experts may end up dominating task performance and dragging their novice partner along, giving the novice insufficient opportunity to learn the task.

To address this reliance on the expert, we employed a fading method of practice in which the amount of control allocated to the expert is gradually reduced over time. Two additional groups were collected as a part of Experiment 2: a $\text{Dyad}_{\text{GRADUAL}}$ group, in which the novices began with no control of the task and were gradually provided with control as the task progressed, and a $\text{Dyad}_{\text{ADAPTIVE}}$ group, in which novices were instead given control based on their performance on the previous block. The $\text{Dyad}_{\text{VISION}}$ group from Experiment 1 was used as a comparison for these two groups in Experiment 2, and is referred to as "$\text{Dyad}_{\text{CONSTANT}}$" for any comparisons made within Experiment 2 (to represent the equal control provided to both participants throughout the entirety of the task). We hypothesized that faded practice (i.e., the $\text{Dyad}_{\text{GRADUAL}}$ and $\text{Dyad}_{\text{ADAPTIVE}}$ groups) would be more beneficial to learning compared to a fixed control allocation. However, we found no significant differences among the three dyad groups, suggesting that this method of implementing faded practice offered similar learning benefits to the constant control allocation.

It is important to note that novices in all dyad groups (in Experiments 1 and 2) were still able to learn the task well enough to perform it on their own, as demonstrated by decreases in movement times and angular error on their post-test relative to their pre-test. This suggests that all implementations of this dyadic practice paradigm are still effective methods of learning a novel BoMI task. However, none of the dyadic practice implementations were able to reach the level of learning achieved by those learning entirely on their own. In the following sections, we discuss potential reasons why this may be the case.

Section 5.2. Main Themes

Theme 1 – Experts tend to "drag"' the novice along, which makes novices reliant on them

Experts in our task were instructed (prior to the participant's arrival) that when paired with the novice, they should still aim to complete the task as quickly and accurately as possible (i.e., their focus should be on completing the task, not accommodating their partner). As no direct instruction was allowed between participants, there was not any opportunity for the dyad to verbally collaborate on a shared strategy. Rather, it is possible that the expert saw the contributions from the novice as a "perturbation" or "noise," something they would have to accommodate for in order to complete the task. What we saw as a result was that movement times from the dyad as a whole during training blocks (i.e. when both participants were contributing) were consistently low, but then spiked when the novice engaged in their mid- and post-test conditions. This result aligns with previous research suggesting that while dyadic performance of a task tends to demonstrate better performance than solo performance, this improved performance only exists due to the expert's existing knowledge of the task "dragging" the novice along, and disappears once the novice is asked to perform the task on their own (Avila

Mireles et al., 2017; Nishimura et al., 2021). This implies that the novice is not given ample opportunity to learn the task, as their performance now worsens when unpaired from their expert partner. Previous research suggests that the only time this expert-novice pairing is beneficial for the novice is when the novice has an opportunity to explore the task on their own first before being paired with the expert (Avila Mireles et al., 2017), and that novice-novice pairings may end up being more beneficial as each novice can benefit from the trial-and-error processes undergone by their partner (Ganesh et al., 2014; Nishimura et al., 2021; Saracbasi et al., 2021).

It is possible, then, that the expert's ability to accommodate for the errors of the novice could have negatively affected the novice's ability to learn the task, instead forcing the novice to become dependent on the knowledge of the expert. This is referred to as the "guidance hypothesis" (Salmoni, Schmidt, & Walter, 1984), and suggests that there is an optimal frequency to which feedback on task performance during training is provided so that novice learners can benefit from the feedback while not becoming reliant on it. In the present study, we applied this concept to the magnitude of expert control, with the idea that there is an optimal amount of control that can be allocated to the expert in order for the novice to benefit from their control but not become reliant on it. Our findings suggest that the initial method of allocating control, in which both participants have equal control of the task, actually favored the expert too heavily. In this equal "50/50" split, the expert was still able to compensate for movement errors of the novice and keep movement times low, which may have caused novices to ill-perform once asked to complete the task themselves.

Relatedly, given how evident the expert's knowledge was (indicated by several novice participants, during the expert's pre-test, commenting about how "you [the expert] have done this before"), novices could have "hidden behind" the knowledge of the expert in a phenomenon known as “loafing.” In group tasks, loafing occurs due to an individual’s expectation for other members of the group to pick up the slack, as their own contribution may not be directly identifiable (Karau & Williams, 1993). It is possible that this phenomenon manifested in the present study, as the indicators of improved performance were based on group metrics (e.g. movement times of the true cursor to the target). As the expert’s knowledge was evident early, the novice may have put forth reduced effort with the hopes that their expert partner could pick up the slack. If, instead, both partners were presented with feedback of their own performance following each block (such as their own cursor trajectories as they contributed to the trajectory of the true cursor), the novice may have been less inclined to hide behind their expert partner.

One method of reducing this reliance on the expert is to introduce a “fading” method of expert contribution (Winstein, Pohl, & Lewthwaite, 1994; Aoyagi et al., 2019), in which (as mentioned above) expert contribution to task performance is more salient in the early stages to account for the initial intricacies of task performance, and then is gradually reduced over time to allow the novice to perform the task on their own. In Experiment 2, we attempted to replicate this fading concept by increasing the expert’s contribution in the early stages, and then gradually tapering it off throughout the training blocks so that the novice finishes the training with complete control of the task. However, the method in which we implemented this fading idea likely still favored the expert too much in the early stages of learning. Instead of providing the novice with an opportunity to observe the proper movements of their partner to be implemented in future trials, novices instead seemed to focus on the outcome of their own movements, which led many of them to the correct conclusion that their movements did not contribute to the performance of the first block at all. As the purpose of implementing this faded practice is to provide experts with enough control to correct for errors, and not take complete control of the task, starting the practice by giving the expert complete control ended up being counterproductive to our desired outcome.

Just as the guidance hypothesis suggests an optimal frequency of feedback that is beneficial, we also suggest an optimal control allocation that is beneficial for joint dyadic practice. While experts were able to keep movement times low throughout most of the training blocks (relative to solo controls), they were unable to compensate to the same degree once novices had approximately 70% control of the task. It could be that this “70% novice control” could approach the “sweet spot” of control between an expert and a novice within a dyad that would allow novices enough control of the task to actually feel as though their movements affected the overall outcome of the task while allowing experts to compensate, to some extent, for the errors of the novice.

Theme 2 – Novice attention was on the cursor, not their partner

As discussed in the previous section, our hope for placing participants in view of their expert partner was that the novice would observe the movements of their partner and implement them on future trials, potentially guiding them towards correct coordination patterns. However, as suggested by the lack of differences observed between the $\text{Dyad}_{\text{VISION}}$ and $\text{Dyad}_{\text{NOVISION}}$ groups, being able to observe the movements of the expert partner did not directly benefit learning of the task. There are several possible explanations for this.

One explanation for this is that due to the unintuitive nature of the BoMI’s control scheme, novice participants weren’t able to focus on the movements of their partner, whether due to choice (i.e. consciously prioritizing the movement of the cursor over those of their partner) or necessity (i.e. as the outcome of the task was the movement of the cursor, the participant may have had no choice BUT to concentrate on movements of the cursor). This aligns with other observations made by the researcher throughout the duration of the study: a few participants in the early stages of data collection reported what they referred to as "eye strain" that appeared to result from concentrating heavily on movements of the cursor and "forgetting to blink," while other participants throughout data collection were noticeably unresponsive to attempts at conversation while trials were ongoing, mentioning that "their brain" could not handle conversation and doing the task simultaneously. This could imply that participants were indeed focused on the cursor to such an extent that their partner's movements just became another stimulus that the novice had to ignore in order to provide sufficient attention to the task.

It is also possible that participants sought to focus on the outcomes that their movements had on the cursor as opposed to the movements themselves. Previous research has suggested that when learning a novel motor task, having participants concentrate on an external focus (in our case, the movement of the cursor) provides greater benefits than if they concentrate on an internal focus (such as their bodily movements). This is potentially due to the "constrained action hypothesis," which suggests that when presented with some reference to an internal focus, individuals will constrain their motor system as they become more aware of how they are moving, causing "micro-choking" episodes, which result in decreased performance (Wulf & Lewthwaite, 2010; Wulf, 2013). It is possible that, when presented with this novel task, participants defaulted to focusing on the movements of the cursor to such an extent that bodily movements (whether it be their own or their partner's) were less of a concern, meaning participants were not able to benefit from being able to observe the movements of their partner. Some participants in the Dyad$_{\text{GRADUAL}}$ and Dyad$_{\text{ADAPTIVE}}$ groups engaged in wide sweeping motions of their body during the first training block, potentially in an attempt to elicit some reaction from the cursor (rather, they were less concerned with what their body was doing, and more concerned with observing some response from the cursor).

It is also possible that participants did not understand that the movements of their partner were intended to guide them to a proper solution. When providing instructions to the novice participants, they were informed that “any movements [they] can think of doing with [their] upper body while remaining seated are fair game,” and were only told that they would be moving collaboratively with their partner to control a single cursor. While some participants seemed to understand that the movement map was identical between the expert and themselves (as demonstrated by some participants asking, “Should I be moving how my partner moves?” or others mirroring the movements of their expert partner during expert testing conditions), others were under the impression that the expert had their own “set of controls” and that their task was to “figure out” their own control scheme, while others still seemed to think that the control scheme changed between testing and training conditions. It is possible that by communicating how the two participants were connected, participants would have “used” their partner more as a resource to guide them towards a proper solution.

Section 5.3. Retrospective

In the present study, we evaluated the effectiveness of a dyadic practice paradigm on the learning of a novel Body-Machine Interface task. Our findings suggest that while our method of implementing dyadic practice (via a combination of sequential and joint dyadic practice paradigms) was an effective method of learning the task, novices were unable to reach the same level of performance as those who learned the task entirely on their own. We found that the ability to observe the movements of one’s partner does not appear to affect the learning of this task, and that our method of fading expert contribution throughout the task likely favored the expert too much. From these findings, we found that the point in which experts begin to lose their ability to completely accommodate novice errors sits around a 70:30 ratio of novice:expert control, which may be able to serve as a starting point for future research that seeks to implement a "fading" method of expert contribution.

Dyadic Interaction

Regarding our experimental structure, there are some slight adjustments that could be made to the paradigm that could allow participants to benefit from our dyadic paradigm. One option would be to lift the restriction on otherwise often-seen dyadic practice paradigms. For instance, previous research suggests that part of the effectiveness of dyadic practice stems from the ability for participants to engage in strategy sharing between trials. While communication between participants was permitted throughout the study, direct instruction on how to move was restricted, in part due to our expert-novice pairing; it was our impression that if permitted to share strategies in this way, experts would simply share the solution with the novice, which would turn the task into "how well can novices follow direct instruction" (it was this same concern that caused us to collect the $\text{Dyad}_{\text{NOVISION}}$ group, as we were concerned that novices were simply mimicking the movements of their partner throughout). If participants were instead allowed to engage in discussion of movement strategies between trials, it may have cleared up previously discussed confusion as to how the movements of each person affected the cursor's position.

In a similar vein, the greatest benefits to dyadic practice are typically observed in novice-novice pairings, as individuals are less likely to take over control of the task throughout. By instead pairing novices with other novices of similar skill and allowing for between-trial discussion of movement strategies, it is possible that we would observe greater benefits to learning compared to expert-novice dyads. While we attempted to collect these novice-novice pairings early in data collection, it proved to be difficult to get two participants to sign up for the same time slot, due either to schedule conflicts on the part of the novice or a desire to sign up for a unique time (i.e., if presented with a time slot that was completely empty and a slot with one existing signup, potential participants appeared to prefer the former).

Task Requirements

It was mentioned earlier that tasks in which both participants can control the entirety of the task are more susceptible to loafing compared to tasks with discrete components. Based on this, the present task could be manipulated in such a way that it visually reflects an interlocked model (i.e., requiring discrete tasks from each participant) but maintains the same task constraints as how it was presented in the present study. As the present study relied on the (weighted) average position of two virtual cursors to determine the position of the true cursor, we could instead present the virtual cursors to the participants while hiding the position of the true cursor. If combined with a set of virtual targets for each participant, a task could be constructed that appeared to represent an interlocked task (in that both participants must navigate their virtual cursor to their own virtual target) while simultaneously causing the true cursor to reach the true target "behind the scenes."

All things considered, it is possible that the cursor control task (as well as the constant map used to transform individual movements into cursor control) was simply too easy to observe benefits from dyadic practice greater than those observed when learning the task solo. To start, outcome metrics such as movement time appeared to "floor out" around the mid-test, implying that much of the difficulty found in the task was reconciled within about 4 blocks of 20 reaches each. It could be that the "simplicity" of the cursor control task could have been the cause of the observed loafing: participants did not feel the need to perform well in training blocks as they believed it would not take very long to "figure it out" once they reached the testing blocks, echoing the sentiment previously mentioned by Crook & Beier (2010) which suggested that the fear of solo performance was not enough of a motivator to reduce loafing in training conditions.

The map that was used by participants was developed by the researcher during the development of the paradigm. Originally, this paradigm first tasked participants with performing a "body dance," which allowed them to move through a comfortable range of motion, and it was these movements performed in the body dance that would be the basis of the map used throughout the remainder of the study. In research involving individuals with movement impairments, this body dance was used so that the movements assigned to control of the device were movements of which the user was capable (demonstrating the customizability of the interface). However, several contributing factors led us to modify how the map was customized to the participant.

First, as the purpose of this research was to evaluate the effectiveness of joint dyadic practice on the BoMI task, we wanted to ensure that all participants were learning an "identical task" so that any differences observed between groups could more easily be attributed to the type of training received (as opposed to differences in task difficulty). If we had continued in allowing participants to create their own map, this would mean that solo participants would create the map that they would use throughout the entire study, while individuals in the dyad would still be at the mercy of the map that experts were trained on, potentially giving the solo controls a leg up before even beginning the pre-test.

Second, in the event that the researcher's instructions were unclear or misunderstood during this body dance, it could leave participants with a map that is frustrating or confusing to understand. Prior to the change enforcing a single map, the researcher performed some pilot testing that allowed participants to set their own map via the body dance, and found that maps were inconsistent between groups – one set of participants had a very intuitive map that enabled them to use their body as a joystick, in that leaning forward moved the cursor forward, etc., and subsequently observed minimal changes to performance between pre-, mid-, and post-test conditions due to how quickly they were able to understand the controls, while another (albeit extreme) case provided participants with a map so difficult that by the time the two-hour session maximum was reached, participants had only just completed the mid-test.

5.4. Final Summary

Given the lack of differences observed between all of our dyad groups, it is possible that the learning of the task did not directly benefit from the presence of an expert. However, two points have to be noted – (i) dyadic practice groups were able to learn the task without making errors during the learning stage, which might make it attractive for some applications where trial-and-error learning is not feasible (e.g., for safety or cost issues), and (ii) dyadic practice could potentially provide greater learning benefits in more complex high degree-of-freedom tasks, as the expert could theoretically control the more complex degrees of freedom in the early stages of learning, and then concede control as the novice becomes familiar with the unique components of total task performance.

In summary, we were able to demonstrate that our method of dyadic practice did allow for sufficient learning of the Body-Machine Interface task, but individuals within a dyad were unable to reach the level of performance of those that earned the task on their own. With all of the previous discussion in mind, future research can begin to implement dyadic practice into more complex high degree-of-freedom Body-Machine Interface tasks. In doing so, we can continue to restore motion and independence to individuals with movement impairments.

BIBLIOGRAPHY

Abdollahi, F., Farshchiansadegh, A., Pierella, C., Seáñez-González, I., Thorp, E., Lee, M.-H., Ranganathan, R., Pedersen, J., Chen, D., Roth, E., Casadio, M., & Mussa-Ivaldi, F. (2017). Body-machine interface enables people with cervical spinal cord injury to control devices with available body movements: Proof of concept. *Neurorehabilitation and Neural Repair*, *31*(5), 487–493. https://doi.org/10.1177/1545968317693111

Aoyagi, Y., Ohnishi, E., Yamamoto, Y., Kado, N., Suzuki, T., Ohnishi, H., Hokimoto, N., & Fukaya, N. (2019). Feedback protocol of 'fading knowledge of results' is effective for prolonging motor learning retention. *Journal of Physical Therapy Science*, *31*(8), 687–691. https://doi.org/10.1589/jpts.31.687

Aspelund, S., Patel, P., Lee, M.-H., Kagerer, F. A., Ranganathan, R., & Mukherjee, R. (2020). Controlling a robotic arm for functional tasks using a wireless head-joystick: A case study of a child with congenital absence of upper and lower limbs. *PLOS ONE*, *15*(8), e0226052. https://doi.org/10.1371/journal.pone.0226052

Avila Mireles, E. J., Zenzeri, J., Squeri, V., Morasso, P., & De Santis, D. (2017). Skill learning and skill transfer mediated by cooperative haptic interaction. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *25*(7), 832–843. https://doi.org/10.1109/TNSRE.2017.2700839

Bastian, A. J. (2008). Understanding sensorimotor adaptation and learning for rehabilitation. *Current Opinion in Neurology*, *21*(6), 628. https://doi.org/10.1097/WCO.0b013e328315a293

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. https://doi.org/10.1109/EMBC44109.2020.9176261

Beckers, N., van Asseldonk, E. H. F., & van der Kooij, H. (2020). Haptic human–human interaction does not improve individual visuomotor adaptation. *Scientific Reports*, *10*(1), 19902. https://doi.org/10.1038/s41598-020-76706-x

Bernstein, N. (1967). The co-ordination and regulation of movements. *The Co-Ordination and Regulation of Movements*. https://cir.nii.ac.jp/crid/1571698599600323072

Biddiss, E., & Chau, T. (2007). Upper-limb prosthetics: Critical factors in device abandonment. *American Journal of Physical Medicine & Rehabilitation*, *86*(12), 977. https://doi.org/10.1097/PHM.0b013e3181587f6c

Bjerrum, A. S., Eika, B., Charles, P., & Hilberg, O. (2014). Dyad practice is efficient practice: A randomised bronchoscopy simulation study. *Medical Education*, *48*(7), 705–712. https://doi.org/10.1111/medu.12398

Casadio, M., Ranganathan, R., & Mussa-Ivaldi, F. A. (2012). The Body-Machine Interface: A new perspective on an old theme. *Journal of Motor Behavior*, *44*(6), 419–433. https://doi.org/10.1080/00222895.2012.700968

Centers for Disease Control and Prevention, National Center on Birth Defects and Developmental Disabilities, Division of Human Development and Disability. Disability and Health Data System (DHDS) Data [online]. [accessed May 6, 2024]. URL: https://dhds.cdc.gov

Cordovani, L., & Cordovani, D. (2016). A literature review on observational learning for medical motor skills and anesthesia teaching. *Advances in Health Sciences Education*, *21*(5), 1113–1121. https://doi.org/10.1007/s10459-015-9646-5

Crook, A. E., & Beier, M. E. (2010). When training with a partner is inferior to training alone: The importance of dyad type and interaction quality. *Journal of Experimental Psychology: Applied*, *16*, 335–348. https://doi.org/10.1037/a0021913

Fall, C. L., Gagnon-Turcotte, G., Dube, J.-F., Gagne, J. S., Delisle, Y., Campeau-Lecours, A., Gosselin, C., & Gosselin, B. (2017). Wireless sEMG-based body-machine interface for assistive technology devices. *IEEE Journal of Biomedical and Health Informatics*, *21*(4), 967–977. https://doi.org/10.1109/JBHI.2016.2642837

Farshchiansadegh, A., Abdollahi, F., Chen, D., Lee, M.-H., Pedersen, J., Pierella, C., Roth, E. J., Gonzalez, I. S., Thorp, E. B., & Mussa-Ivaldi, F. A. (2014). A body machine interface based on inertial sensors. *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 6120–6124. https://doi.org/10.1109/EMBC.2014.6945026

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), Article 1. https://doi.org/10.1038/srep03824

Goss-Sampson, M. A. (n.d.). *Statistical Analysis in JASP: A Guide for Students*.

Griffin, W. B., Provancher, W. R., & Cutkosky, M. R. (2003). Feedback strategies for shared control in dexterous telemanipulation. *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003) (Cat. No.03CH37453)*, *3*, 2791–2796 vol.3. https://doi.org/10.1109/IROS.2003.1249293

Groten, R., Feth, D., Klatzky, R. L., & Peer, A. (2013). The role of haptic feedback for the integration of intentions in shared task execution. *IEEE Transactions on Haptics*, *6*(1), 94–105. IEEE Transactions on Haptics. https://doi.org/10.1109/TOH.2012.2

Haith, A. M., Yang, C. S., Pakpoor, J., & Kita, K. (2022). De novo motor learning of a bimanual control task over multiple days of practice. *Journal of Neurophysiology*, *128*(4), 982–993. https://doi.org/10.1152/jn.00474.2021

Ivanova, E., Eden, J., Carboni, G., Krüger, J., & Burdet, E. (2022). Interaction with a reactive partner improves learning in contrast to passive guidance. *Scientific Reports*, *12*(1), Article 1. https://doi.org/10.1038/s41598-022-18617-7

Kager, S., Hussain, A., Cherpin, A., Melendez-Calderon, A., Takagi, A., Endo, S., Burdet, E., Hirche, S., Ang, M. H., & Campolo, D. (2019). The effect of skill level matching in dyadic interaction on learning of a tracing task. *2019 IEEE 16th International Conference on Rehabilitation Robotics (ICORR)*, 824–829. https://doi.org/10.1109/ICORR.2019.8779485

Karau, S. J., & Williams, K. D. (1993). Social loafing: A meta-analytic review and theoretical integration. *Journal of Personality and Social Psychology*, *65*(4), 681–706. https://doi.org/10.1037/0022-3514.65.4.681

Karlinsky, A., Welsh, T. N., & Hodges, N. J. (2019). Learning together: Observation and other mechanisms which mediate shared practice contexts. In *Skill Acquisition in Sport: Research, Theory, and Practice* (3rd).

Khademian, B., & Hashtrudi-Zaad, K. (2011). Shared control architectures for haptic training: Performance and coupled stability analysis. *The International Journal of Robotics Research*, *30*(13), 1627–1642. https://doi.org/10.1177/0278364910397559

Kim, S. J., Wen, Y., Ludvig, D., Küçüktabak, E. B., Short, M. R., Lynch, K., Hargrove, L., Perreault, E. J., & Pons, J. L. (2023). Effect of dyadic haptic collaboration on ankle motor learning and task performance. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *31*, 416–425. IEEE Transactions on Neural Systems and Rehabilitation Engineering. https://doi.org/10.1109/TNSRE.2022.3225759

Kitago, T., Ryan, S. L., Mazzoni, P., Krakauer, J. W., & Haith, A. M. (2013). Unlearning versus savings in visuomotor adaptation: Comparing effects of washout, passage of time, and removal of errors on motor memory. *Frontiers in Human Neuroscience*, *7*. https://doi.org/10.3389/fnhum.2013.00307

Kluzik, J., Diedrichsen, J., Shadmehr, R., & Bastian, A. J. (2008). Reach adaptation: What determines whether we learn an internal model of the tool or adapt the model of our arm? *Journal of Neurophysiology*, *100*(3), 1455–1464. https://doi.org/10.1152/jn.90334.2008

Krakauer, J. W., Pine, Z. M., Ghilardi, M.-F., & Ghez, C. (2000). Learning of visuomotor transformations for vectorial planning of reaching trajectories. *The Journal of Neuroscience*, *20*(23), 8916–8924. https://doi.org/10.1523/JNEUROSCI.20-23-08916.2000

Krakauer, J. W., Hadjiosif, A. M., Xu, J., Wong, A. L., & Haith, A. M. (2019). Motor learning. In R. Terjung (Ed.), *Comprehensive Physiology* (1st ed., pp. 613–663). Wiley. https://doi.org/10.1002/cphy.c170043

Latané, B., Williams, K., & Harkins, S. (1979). Many hands make light the work: The causes and consequences of social loafing. *Journal of Personality and Social Psychology*, *37*(6), 822–832. https://doi.org/10.1037/0022-3514.37.6.822

Latash, M. L., & Anson, J. G. (1996). What are "normal movements" in atypical populations? *Behavioral and Brain Sciences*, *19*(1), 55–68. https://doi.org/10.1017/S0140525X00041467

Lee, M.-H., & Ranganathan, R. (2019). Age-related deficits in motor learning are associated with altered motor exploration strategies. *Neuroscience*, *412*, 40–47. https://doi.org/10.1016/j.neuroscience.2019.05.047

Lee, M.-H., Ranganathan, R., Kagerer, F. A., & Mukherjee, R. (2016). Body-machine interface for control of a screen cursor for a child with congenital absence of upper and lower limbs: A case report. *Journal of NeuroEngineering and Rehabilitation*, *13*(1), 34. https://doi.org/10.1186/s12984-016-0139-4

Li, J., Tavakoli, M., & Huang, Q. (2014). Absolute stability of a class of trilateral haptic systems. *IEEE Transactions on Haptics*, *7*(3), 301–310. IEEE Transactions on Haptics. https://doi.org/10.1109/TOH.2014.2321616

Liu, X., Mosier, K. M., Mussa-Ivaldi, F. A., Casadio, M., & Scheidt, R. A. (2011). Reorganization of finger coordination patterns during adaptation to rotation and scaling of a newly learned sensorimotor transformation. *Journal of Neurophysiology*, *105*(1), 454–473. https://doi.org/10.1152/jn.00247.2010

Lokesh, R., & Ranganathan, R. (2019). Differential control of task and null space variability in response to changes in task difficulty when learning a bimanual steering task. *Experimental Brain Research*, *237*(4), 1045–1055. https://doi.org/10.1007/s00221-019-05486-2

Mace, M., Kinany, N., Rinne, P., Rayner, A., Bentley, P., & Burdet, E. (2017). Balancing the playing field: Collaborative gaming for physical training. *Journal of NeuroEngineering and Rehabilitation*, *14*(1), 116. https://doi.org/10.1186/s12984-017-0319-x

Nishimura, K., Saracbasi, O. O., Hayashi, Y., & Kondo, T. (2021). Cooperative visuomotor learning experience with peer enhances adaptability to others. *Advanced Robotics*, *35*(13–14), 835–841. https://doi.org/10.1080/01691864.2021.1913445

Nudehi, S. S., Mukherjee, R., & Ghodoussi, M. (2005). A shared-control approach to haptic interface design for minimally invasive telesurgical training. *IEEE Transactions on Control Systems Technology*, *13*(4), 588–592. IEEE Transactions on Control Systems Technology. https://doi.org/10.1109/TCST.2004.843131

Orthotics and Prosthetics. (n.d.). Atrium Health Wake Forest Baptist. Retrieved November 9, 2023, from https://www.wakehealth.edu/treatment/o/orthotics-and-prosthetics

Phillips, B., & Zhao, H. (1993). Predictors of assistive technology abandonment. *Assistive Technology*, *5*(1), 36–45. https://doi.org/10.1080/10400435.1993.10132205

Pierella, C., Abdollahi, F., Farshchiansadegh, A., Pedersen, J., Thorp, E., Mussa-Ivaldi, F. A., & Casadio, M. (2015). Remapping residual coordination for controlling assistive devices and recovering motor functions. *Neuropsychologia*, *79*(0 0), 364–376. https://doi.org/10.1016/j.neuropsychologia.2015.08.024

Pierella, C., Abdollahi, F., Thorp, E., Farshchiansadegh, A., Pedersen, J., Seáñez-González, I., Mussa-Ivaldi, F. A., & Casadio, M. (2017). Learning new movements after paralysis: Results from a home-based study. *Scientific Reports*, *7*(1), 4779. https://doi.org/10.1038/s41598-017-04930-z

Pierella, C., Galofaro, E., De Luca, A., Losio, L., Gamba, S., Massone, A., Mussa-Ivaldi, F. A., & Casadio, M. (2021). Recovery of distal arm movements in spinal cord injured patients with a body-machine interface: A proof-of-concept study. *Sensors (Basel, Switzerland)*, *21*(6), 2243. https://doi.org/10.3390/s21062243

Ranganathan, R., Lee, M.-H., Padmanabhan, M. R., Aspelund, S., Kagerer, F. A., & Mukherjee, R. (2019). Age-dependent differences in learning to control a robot arm using a body-machine interface. *Scientific Reports*, *9*(1), Article 1. https://doi.org/10.1038/s41598-018-38092-3

Ranganathan, R., & Scheidt, R. A. (2016). Organizing and reorganizing coordination patterns. In J. Laczko & M. L. Latash (Eds.), *Progress in Motor Control: Theories and Translations* (pp. 327–349). Springer International Publishing. https://doi.org/10.1007/978-3-319-47313-0_18

Reed, K., Peshkin, M., Hartmann, M., Patton, J., Vishton, P., & Grabowecky, M. (2006). Haptic cooperation between people, and between people and machines. *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2109–2114. https://doi.org/10.1109/IROS.2006.282489

Ritchie, A. M. A., Chaudhry, N., Robinson, S. J. A., Pacilli, M., & Nataraja, R. M. (2024). Dyadic practice for the acquisition of laparoscopic skills (DALS)—A randomized controlled trial. *World Journal of Surgery*, *48*(2), 278–287. https://doi.org/10.1002/wjs.12074

Rizzoglio, F., Pierella, C., De Santis, D., Mussa-Ivaldi, F., & Casadio, M. (2020). A hybrid Body-Machine Interface integrating signals from muscles and motions. *Journal of Neural Engineering*, *17*(4), 046004. https://doi.org/10.1088/1741-2552/ab9b6c

Rizzoglio, F., Giordano, M., Mussa-Ivaldi, F. A., & Casadio, M. (2023). A non-linear body machine interface for controlling assistive robotic arms. *IEEE Transactions on Biomedical Engineering*, *70*(7), 2149–2159. IEEE Transactions on Biomedical Engineering. https://doi.org/10.1109/TBME.2023.3237081

Roberts, J. W., Bennett, S. J., Welsh, T. N., Elliott, D., Lyons, J. L., & Hayes, S. J. (2017). The influence of environmental context in interpersonal observation–execution. *Quarterly Journal of Experimental Psychology*, *70*(1), 154–162. https://doi.org/10.1080/17470218.2015.1127982

Salmoni, A. W., Schmidt, R. A., & Walter, C. B. (1984). Knowledge of results and motor learning: A review and critical reappraisal. *Psychological Bulletin*, *95*(3), 355–386. https://doi.org/10.1037/0033-2909.95.3.355

Saracbasi, O. O., Harwin, W., Kondo, T., & Hayashi, Y. (2021). Mutual skill learning and adaptability to others via haptic interaction. *Frontiers in Neurorobotics*, *15*. https://www.frontiersin.org/articles/10.3389/fnbot.2021.760132

Scheidt, R. A., & Ghez, C. (2007). Separate adaptive mechanisms for controlling trajectory and final position in reaching. *Journal of Neurophysiology*, *98*(6), 3600–3613. https://doi.org/10.1152/jn.00121.2007

Scholz, J. P., & Schöner, G. (1999). The uncontrolled manifold concept: Identifying control variables for a functional task. *Experimental Brain Research*, *126*(3), 289–306. https://doi.org/10.1007/s002210050738

Scott, M. W., Howard, J., Karlinsky, A., Mehta, A., Welsh, T. N., & Hodges, N. J. (2023). Short-term evidence of partner-induced performance biases in simultaneous and alternating dyad practice in golf. *Scientific Reports*, *13*(1), 21099. https://doi.org/10.1038/s41598-023-48133-1

Seáñez-González, I., Pierella, C., Farshchiansadegh, A., Thorp, E. B., Wang, X., Parrish, T., & Mussa-Ivaldi, F. A. (2016). Body-machine interfaces after spinal cord injury: Rehabilitation and brain plasticity. *Brain Sciences*, *6*(4), 61. https://doi.org/10.3390/brainsci6040061

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

Shamaei, K., Kim, L. H., & Okamura, A. M. (2015). Design and evaluation of a trilateral shared-control architecture for teleoperated training robots. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual International Conference*, *2015*, 4887–4893. https://doi.org/10.1109/EMBC.2015.7319488

Shanks, D., Brydges, R., den Brok, W., Nair, P., & Hatala, R. (2013). Are two heads better than one? Comparing dyad and self-regulated learning in simulation training. *Medical Education*, *47*(12), 1215–1222. https://doi.org/10.1111/medu.12284

Shea, C. H., Wright, D. L., Wulf, G., & Whitacre, C. (2000). Physical and observational practice afford unique learning opportunities. *Journal of Motor Behavior*, *32*(1), 27–36. https://doi.org/10.1080/00222890009601357

Shea, C. H., Wulf, G., & Whitacre, C. (1999). Enhancing training efficiency and effectiveness through the use of dyad training. *Journal of Motor Behavior*, *31*(2), 119–125. https://doi.org/10.1080/00222899909600983

Shebilske, W. L., Regian, J. W., Arthur, W., & Jordan, J. A. (1992). A dyadic protocol for training complex skills. *Human Factors*, *34*(3), 369–374. https://doi.org/10.1177/001872089203400309

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

Swain, A. (1996). Social loafing and identifiability: The mediating role of achievement goal orientations. *Research Quarterly for Exercise and Sport*, *67*(3), 337–344. https://doi.org/10.1080/02701367.1996.10607961

Takagi, A., Beckers, N., & Burdet, E. (2016). Motion plan changes predictably in dyadic reaching. *PLOS ONE*, *11*(12), e0167314. https://doi.org/10.1371/journal.pone.0167314

Taylor, J. A., & Ivry, R. B. (2011). Flexible cognitive strategies during motor learning. *PLOS Computational Biology*, *7*(3), e1001096. https://doi.org/10.1371/journal.pcbi.1001096

Telgen, S., Parvin, D., & Diedrichsen, J. (2014). Mirror reversal and visual rotation are learned and consolidated via separate mechanisms: Recalibrating or learning de novo? *Journal of Neuroscience*, *34*(41), 13768–13779. https://doi.org/10.1523/JNEUROSCI.5306-13.2014

Tia, B., Paizis, C., Mourey, F., & Pozzo, T. (2012). Do equilibrium constraints modulate postural reaction when viewing imbalance? *Brain and Cognition*, *79*(2), 89–95. https://doi.org/10.1016/j.bandc.2012.02.008

Tolsgaard, M. G., Madsen, M. E., Ringsted, C., Oxlund, B. S., Oldenburg, A., Sorensen, J. L., Ottesen, B., & Tabor, A. (2015). The effect of dyad versus individual simulation-based ultrasound training on skills transfer. *Medical Education*, *49*(3), 286–295. https://doi.org/10.1111/medu.12624

Williams, K. D., Nida, S. A., Baca, L. D., & Latane, B. (1989). Social loafing and swimming: Effects of identifiability on individual and relay performance of intercollegiate swimmers. *Basic and Applied Social Psychology*, *10*(1), 73–81. https://doi.org/10.1207/s15324834basp1001_7

Winstein, C. J., Pohl, P. S., & Lewthwaite, R. (1994). Effects of physical guidance and knowledge of results on motor learning: support for the guidance hypothesis. *Research Quarterly for Exercise and Sport*, *65*(4), 316–323. https://doi.org/10.1080/02701367.1994.10607635

Wulf, G. (2013). Attentional focus and motor learning: A review of 15 years. *International Review of Sport and Exercise Psychology*, *6*(1), 77–104. https://doi.org/10.1080/1750984X.2012.723728

Wulf, G., Höß, M., & Prinz, W. (1998). Instructions for motor learning: Differential effects of internal versus external focus of attention. *Journal of Motor Behavior*, *30*(2), 169–179. https://doi.org/10.1080/00222899809601334

Wulf, G., & Lewthwaite, R. (2010). Effortless motor learning?: An external focus of attention enhances movement effectiveness and efficiency. In *Effortless attention: A new perspective in the cognitive science of attention and action* (pp. 75–101). Boston Review. https://doi.org/10.7551/mitpress/9780262013840.003.0004

Wulf, G., Shea, C., & Lewthwaite, R. (2010). Motor skill learning and performance: A review of influential factors. *Medical Education*, *44*(1), 75–84. https://doi.org/10.1111/j.1365-2923.2009.03421.x

Yang, C. S., Cowan, N. J., & Haith, A. M. (2021). De novo learning versus adaptation of continuous control in a manual tracking task. *eLife*, *10*, e62578. https://doi.org/10.7554/eLife.62578