

©Copyright 2022
Momona Yamagami

Modeling and Enhancing Human-Machine Interaction for Accessibility and Health

Momona Yamagami

A dissertation submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2022

Reading Committee:

Samuel A. Burden, Chair

Katherine M. Steele, Chair

Amy L. Orsborn

Program Authorized to Offer Degree:
UW Electrical Engineering

University of Washington

Abstract

Modeling and Enhancing Human-Machine Interaction for Accessibility and Health

Momona Yamagami

Co-Chairs of the Supervisory Committee:

Associate Professor Samuel A. Burden

Electrical & Computer Engineering

Albert S. Kobayashi Endowed Professor Katherine M. Steele

Mechanical Engineering

Multi-channel closed-loop interfaces using biosignals measured from sensors in, on, or around our bodies have the potential to revolutionize the way we interact with devices, machines, and each other. Such signals can be used to more intuitively and accessibly control human-machine interfaces (HMIs) for machines like vehicles, computing devices, or virtual reality avatars. However, we do not yet have the tools to develop individualized adaptive interfaces that are stable in closed-loop with an adaptive human. This dissertation focuses on how individualized adaptive algorithms can be leveraged to improve accessibility and health toward a more inclusive and equitable society.

A significant challenge with multi-channel biosignals-based control is that the measured signals are noisy and require the user to correct for errors in addition to tracking a desired trajectory. For example, to control a powered wheelchair with biosignals like electromyography (EMG), users must correct for noisy signals that make the wheelchair deviate from the intended path using feedback control. In parallel, users can track a desired path using feedforward control. We first need to develop methods that can separate feedforward and feedback control. As an initial step, I recruited non-disabled participants to perform trajectory-tracking and disturbance-rejection tasks and used control theory methods to separately quantify the user's feedforward and feedback control. I used this model to examine whether handedness affects learned controllers with participants without disabilities. Par-

ticipants learned feedforward and feedback controllers equally well with both hands and participants' learned controllers transferred between hands. Reduction of motor noise (i.e., extraneous movements to the task at hand) was a large factor in improved trajectory-tracking, highlighting the need for an assistive algorithm to compensate for a person's intrinsic motor noise.

I next investigated how HMIs can be enhanced with EMG as a multi-channel control method. EMG is an attractive multi-channel, intuitive, always available control option because it is non-invasive, can be detected even if muscle contractions are small, and sensor placement can be individualized to the abilities and preferences of the user. I conducted a study with nondisabled participants comparing EMG and manual interfaces during a trajectory-tracking task when people are controlling the velocity and the acceleration of a cursor on a screen. People implemented better feedforward control (i.e., closer to the ideal controller to perfectly track the reference) when using an EMG interface than when using a manual interface for high-frequency acceleration-based trajectory-tracking tasks. I also had participants with upper-arm disability after stroke perform the tracking task with the EMG and manual interfaces, and found that they adapted their feedforward controllers similarly to participants without disabilities for both interfaces. This suggests that EMG interfaces could enable accessible device interactions for people with disabilities if an assistive algorithm helps minimize errors arising from motor disturbances.

I then used game theory to model and enhance HMIs during continuous disturbance-rejection tasks. There are significant challenges with applying traditional control theory methods to quantify how humans and interfaces co-adapt to reference-tracking and disturbance-rejection tasks. Game theory provides a framework to model such two-learner problems. I modeled the human and the interface as two separate agents who are both trying to minimize task error and effort. Based on the results of my previous two aims, I developed an adaptive interface that augments the person's feedback controller and found through simulation and experiment that co-adaptation improved performance and decreased human effort.

Lastly, I worked directly with people with disabilities to identify how technology can support people's health and accessibility. I qualitatively assessed through an interview study how wearable sensors could support physical therapy access for people with disabilities to improve adherence and function. People's access to physical therapy was hampered by both social (e.g., physically visiting a clinic) and physiological (e.g., chronic pain) barriers. I defined core design principles (flexibility, movement tracking, community building) and tensions (insurance) to consider when developing technology to support physical therapy access.

Individualized HMIs with multi-channel, closed-loop control provides exciting opportunities for improving the accessibility of current and future technologies and improving health outcomes for people with and without disabilities. To advance our field's capacity to design and optimize interfaces that can adapt to individual users, both user-centered and quantitative approaches to enhancing interfaces must be developed. My dissertation provides the foundation for developing biosignal-based interfaces that support accessibility and health for people with and without disabilities.

ACKNOWLEDGMENTS

Thank you so much everyone. I can confidently say that I would not have reached this milestone without the support and guidance of the many people around me who made my PhD a transformative experience.

Getting a PhD is a challenging journey, and one that I would have given up on many times without the support of my husband, Edgar. From taking a five hour flight every two weeks to moving halfway across the country, he has been my biggest support and I am so grateful for his steady presence in my life. Thank you for supporting me and my love for research, all the way back from when we would make a detour to lab after our dates during undergrad to split cells at midnight on Fridays.

I would like to express my gratitude towards all my mentors, especially to my advisors Sam and Kat. Sam, thank you for taking the time to listen to all my research ideas, and acknowledging that life exists outside of research and supporting a holistic PhD experience. Kat, thank you for inspiring me to be the best researcher that I can be and helping me shape my research identity. Thank you to my collaborators and mentors, Val, Amy, Martez, Jen, Eatai, and Marcie for introducing me to the research challenges and perspectives of the different fields they reside in. Their viewpoints and discussions have significantly shaped my perspective on interdisciplinary research. I also want to acknowledge Dr. Suh and Michelle for giving me the opportunity to explore research as an undergraduate and inspiring me to apply for PhD programs in the first place.

I also want to thank my labmates, friends and colleagues who helped me so much by celebrating with me, commiserating with me, and making my PhD experience an unforgettable one. There are too many of you to name, but please know that I would not have gotten to where I am now without you. To my undergraduate and graduate mentees, thank you for trusting me to mentor you.

Lastly, I want to thank my parents, Keiko and Kenji, and my sisters, Erika and Nozomi for giving me the opportunity to study in the US. Their sacrifices and support from across the Pacific Ocean means so much to me and this is a huge and shared accomplishment for the entire Yamagami family.

My dissertation was supported by the University of Washington (UW) College of Engineering Fellowship, the UW Institute for Neuroengineering Innovation Graduate Fellowship, the UW Amplifying Motion and Performance (AMP) Lab, the UW Electrical & Computer Engineering Irene Peden Fellowship, Meta Research, and the National Science Foundation Cyber-Physical Systems Program (Award #1836819).

TABLE OF CONTENTS

	Page
List of Figures	iii
List of Tables	viii
Chapter 1: Introduction	1
1.1 Dissertation Contributions and Organization	3
Chapter 2: Effect of Handedness on Learned Controllers and Sensorimotor Noise During Trajectory-Tracking	7
2.1 Introduction	9
2.2 Problem Formulation	12
2.3 Experimental Methods	15
2.4 Results	22
2.5 Discussion	27
2.6 Conclusion	34
Chapter 3: Decoding Intent with Control Theory: Comparing Muscle versus Man- ual Interface Performance	36
3.1 Introduction	38
3.2 Related Work	40
3.3 What is a Continuous Task?	43
3.4 Methods	45
3.5 Results	52
3.6 Discussion	56
3.7 Conclusion	60
Chapter 4: Co-Adaptation for Human-in-the-Loop Control Systems	62
4.1 Introduction	64
4.2 The co-adaptation game in human-machine interfaces	66
4.3 Experimental results	70

4.4	Simulation results	73
4.5	Discussion	77
4.6	Methods	80
Chapter 5:	“I’m Just Overwhelmed”: Investigating Physical Therapy Accessibility and Technology Interventions for People with Disabilities and/or Chronic Conditions	86
5.1	Introduction	88
5.2	Related Work	90
5.3	Methods	94
5.4	Results	97
5.5	Discussion	109
5.6	Conclusion	116
Chapter 6:	Conclusion and Future Work	118
6.1	Individualized HMI Model for Continuous Interactions	118
6.2	Electromyography as Machine Input for People With and Without Disabilities	120
6.3	Enhancing HMIs with Co-Adaptive Learning	121
6.4	Supporting Accessibility and Health with Technology	122
6.5	Human-machine Interfaces for Accessibility and Health	123
Appendix A:	Appendix	149
A.1	Aim 1 Appendix	149

LIST OF FIGURES

Figure Number	Page
<p>2.1 <i>Human-in-the-loop trajectory-tracking.</i> (a) Human response u is obtained with a one-dimensional manual slider and input to machine M to produce output y, which is overlaid on a display with 1 sec of a reference trajectory (0.5 sec preview). (b) The human H transforms reference r and output y to user response u; the machine M transforms the sum of control u and disturbance d to output y. We hypothesize that the human's transformation is the superposition of a <i>feedforward</i> F response to reference r and a <i>feedback</i> B response to tracking error $r - y$. Representative data from one trial of the linearity experiment are shown in (c) the time-domain and (d) the frequency-domain. The frequency content of r and d are confined to prime multiples of a base frequency (1/20 Hz). Magnitudes shown as percent of output or input space extent.</p>	11
<p>2.2 <i>Handles for linearity and handedness experiments.</i> (a) Participants pinched a rectangular handle with their fingers in the linearity experiments. (b) Participants grasped a cylindrical handle with their hand in the handedness experiments.</p>	16
<p>2.3 <i>Conditions for linearity experiment (cf. TABLE 2.2).</i> To assess whether the human's response to external reference r superimposes with the response to external disturbance d, we empirically estimated transfer functions using data from four experimental conditions: disturbance-only $((0, d)$, upper left); reference-only $((r, 0)$, upper right); reference and disturbance interleaved at different frequencies $((r, d)$, bottom left, right). The magnitude of \hat{r} is denoted with solid lines and filled circles, while dashed lines and open circles denote that of $\hat{M}d$; insets show corresponding time-domain signals r, $M(d)$. Magnitudes shown as percent of output or input space extent.</p>	19
<p>2.4 <i>Transfer function estimates in linearity experiment.</i> Distributions (median, interquartile) of transfer functions \hat{T}_{ud} (<i>left</i>), \hat{T}_{ur} (<i>right</i>) estimated from disturbance-only or reference-only trials, $(0, d)$ or $(r, 0)$, and reference-plus-disturbance trials (r, d), for the conditions in TABLE 2.2 and Fig. 2.3. Statistically significant differences (Wilcoxon signed-rank test: $p < 0.05$) in distribution magnitude or phase at each frequency indicated with †.</p>	23

2.5	<i>Predictive accuracy of models, linearity experiment.</i> Distribution (median, interquartile, range) of coefficient of determination (R^2) between human inputs u and predictions from feedback-only (B) and feedback-plus-feedforward ($B + F$) models. The $B + F$ model had significantly better prediction accuracy than the B model at all frequencies (Wilcoxon signed-rank test: $Z = 0.0, p = 0.016$; indicated with *).	24
2.6	<i>Tracking error from handedness experiment.</i> Distributions (median, interquartile) of time-domain tracking error $\ r - y\ ^2$ for 60 trials, with a switch between dominant (right ; red circles) and non-dominant (left ; blue squares) hands after trial 30, for two groups of 9 participants: (<i>top</i>) right then left (Group RL); (<i>bottom</i>) left then right (Group LR). Summary statistics in Fig. 2.7 use data from first five and last five trials with each hand, highlighted with light and dark gray boxes.	25
2.7	<i>Summary statistics from handedness experiment.</i> Distributions (median, interquartile, range) from first five (light gray box) and last five (dark gray box) of 30 trials with dominant (red solid background) and non-dominant (blue hatched background) hands: (a) tracking error $\ r - y\ ^2$; mean magnitude of (b) feedforward $ \widehat{F} $ and (c) feedback $ \widehat{B} $ controllers (shared y axis); mean magnitude of (d) disturbance rejection $ \widehat{T}_{yd} $ and (e) reference tracking $ \widehat{T}_{yr} - 1 $ errors (shared y axis). Statistically significant (Wilcoxon signed-rank test: $p < 0.05$) differences between adjacent distributions indicated with *. Group RL in top row, Group LR in bottom row, as in Fig. 2.6.	26
2.8	<i>Human feedback (B) and feedforward (F) controllers.</i> Distributions (median, interquartile) obtained by pooling data from the last five trials with each hand for both groups in the handedness experiment; we did not observe statistically significant differences between groups or hands (Wilcoxon signed-rank test: $p > 0.05$).	27
2.9	<i>Change in effect of sensorimotor noise.</i> (<i>top row:</i>) Distributions (median, interquartile) of magnitude of user response at non-stimulated frequencies from first and last five trials with first hand (trials #1–5 in light gray and #26–30 in green) in handedness experiment. (<i>bottom row:</i>) Ratio of user response magnitudes between first and last five trials with first hand decreases significantly below crossover (0.25 Hz). Group LR in left column, Group RL in right column.	28

3.1	Successfully completing continuous tasks with manual or muscle interfaces is crucial for many tasks including cursor navigation. While a user may intend to follow a desired reference path (dotted red line) with user intent (dotted red arrows), unexpected disturbances (sudden change in cursor position between the two blue circles) introduce errors that must be corrected with error correction (solid blue arrows). The user input (mouse position) combines user intent and error correction and maps to the cursor position on the screen (blue solid line).	39
3.2	(i) Block diagram representation of user interacting with device adapted from [105]. The user, contained within the purple dotted square, transforms external reference R and tracking error $R - Y$ through feedforward (user intent, F in red) and feedback (error correction, B in blue) controllers to produce user input U . The device transforms the sum of user input U and external disturbance D to device output Y via mapping M . (ii) While it is difficult to separate the two reference (dotted red line) and disturbance (solid blue line) signals in the time-domain (left), the task is much easier in the frequency-domain (right). We can easily go back and forth from time- and frequency-domain using the Fourier and inverse Fourier transform. . . .	45
3.3	Participants controlled a purple cursor on a computer screen using either a manual slider (top) or muscle EMG (bottom) interface.	48
3.4	Time-domain (left) and frequency-domain (right) measures of error for both simple and complex tasks. Lower values indicate better performance. Statistically significant differences are marked with their respective p values. . . .	53
3.5	Frequency-based performance across different frequencies for complex acceleration-based task. Participants performed significantly better with the muscle (yellow) than the manual (purple) interface.	55
3.6	Time-domain (left) and frequency-domain (right) measures of error for complex task for users with and without motor impairments. Lower values equals better performance. The time-domain error for users with motor impairments is much worse than users without motor impairments for both the muscle and manual interface. Users with motor impairments perform comparably to users without motor impairments in forming feedforward models in the frequency-domain.	57
3.7	Results from NASA TLI demonstrates similar subjective workload across all tasks for users with and without motor impairments.	58

4.1 *Block diagram models for HMIs.* These diagrams specify the flow of information, with signals illustrated by *arrows* and transformations of signals illustrated by *blocks*. We allow signals to be multidimensional, and blocks to be *multi-input multi-output* [195, Ch. 1]. For instance, M has three inputs (q, w, u) and outputs (p, z, y). We enumerate the inputs and outputs from top-to-bottom in the diagram, so q is the first input to M and y is the third output from M , and we use subscripts to denote transformations between (groups of) signals¹. When the blocks H, M , and I are linear time-invariant [195, Ch. 3], these diagrams are not solely conceptual – they provide precise mathematical specifications of the closed-loop transformation from input disturbance w to output error z . Indeed, the *left* diagram is equivalent to the two diagrams in the *middle* where the block M/I is determined by the feedback interconnection (formally, the *lower linear fractional transformation* [195, Ch. 10.1]) between M and I , and the block H/M is determined by the feedback interconnection (formally, the *upper linear fractional transformation* [195, Ch. 10.1]) between M and H . Finally, all of these are equivalent to the *right* diagram that is obtained by applying the same interconnection algebra to the feedback loops that remain in the *middle* diagrams to obtain the transformation $H/M/I$ from input disturbance w to output error z 67

4.2 *HMI model with series interaction and disturbance feedforward.* Left: When reference $r = 0$, the adaptive human H transforms output y to user response u_H ; the adaptive interface I transforms the user response u_H to interface response u_I ; the fixed machine M transforms the sum of interface response u_I and disturbance d to output y . Right: Human response u_H is obtained with one-dimensional manual slider and is input to the adaptive interface I and fixed machine M to produce output y 69

4.3 *Individualized interface controllers after co-adaptation.* Participants converged to a range of interfaces after co-adaptation. The black point represents the baseline interface (i.e., $I = 1$) and the blue points represent each participants’ final interface coefficients after co-adaptation. 71

4.4 *Baseline versus co-adaptation performance.* Distributions (median, interquartile) for baseline and co-adaptation for various performance metrics for 0th (left), 1st (middle), and 2nd (right) order interface dynamics. Top to bottom: time-domain cursor output error $\sum_t |y(t)|^2$; frequency-domain cursor output error for stimulated frequencies below crossover $\sum_\omega |\hat{y}(\omega)|^2$; potential function loss below crossover $\sum_\omega |\hat{p}(\omega)|^2$; disturbance rejection performance below crossover $\sum_\omega \hat{T}_{yd}(\omega)$. Statistical significance ($p < 0.05$) computed from Wilcoxon signed-rank test is indicated by the bracket. 72

4.5	<i>Human and interface effort.</i> Distributions (median, interquartile) for baseline and co-adaptation for various human and interface effort metrics for 0th (left), 1st (middle), and 2nd (right) order interface dynamics. Top to bottom: time-domain user response $\sum_t u_H(t) ^2$; frequency-domain user response for stimulated frequencies below crossover $\sum_\omega \hat{u}_H(\omega) ^2$; frequency-domain user response for non-stimulated frequencies below crossover $\sum_{\omega_{non-stim}} \hat{u}_H(\omega) ^2$; human controller magnitude below crossover $\sum_\omega \hat{H}(\omega) ^2$; interface controller magnitude below crossover $\sum_\omega \hat{I}(\omega) ^2$. Statistical significance ($p < 0.05$) computed from Wilcoxon signed-rank test is indicated by the bracket. . . .	73
4.6	<i>Final interface and human gain for 0th, 1st, and 2nd order interface dynamics.</i> Count for final interface (top) and human (bottom) gain for 0th (left), 1st (middle), and 2nd (right) order interface dynamics. Darker shade of grey represents lower gains; lighter shades represents higher gains. The simulation was bounded with a minimum gain of 0.2 and a maximum gain of 7.0. . . .	75
4.7	<i>HMI performance for 0th, 1st, and 2nd order interface dynamics.</i> HMI performance for 0th (left), 1st (middle), and 2nd (right) order human and interface. From top to bottom: potential function loss $\hat{p}(\omega)$; cursor output $\hat{y}(\omega)$; disturbance rejection performance $\hat{T}_{yd}(\omega)$. All performance metrics represent the difference between the performance with co-adaptation and baseline. Performance difference less than 0 represents better performance with co-adaptation compared to baseline for the given human penalty. . . .	76
4.8	<i>Human and interface effort for 0th, 1st, and 2nd order interface dynamics.</i> Human and interface effort for 0th (left), 1st (middle), and 2nd (right) order human and interface. From top to bottom: user response $\hat{u}_H(\omega)$; human effort $\hat{H}(\omega)$; interface effort $\hat{I}(\omega)$. All effort metrics represent the difference between the effort with co-adaptation and baseline. Effort difference less than 0 represents lower effort with co-adaptation compared to baseline for the given human penalty. . . .	77

LIST OF TABLES

Table Number	Page
2.1 <i>Table of symbols.</i>	12
2.2 <i>Conditions for linearity experiment (cf. Fig. 2.3).</i>	18
3.1 Participant Characteristics. Self-reported impairments adopted from Findlater et. al [49] and Mott et. al [109]	47
3.2 Two way ANOVA (<i>interface</i> \times <i>task</i>) results for time-domain (MSE_{time}) and frequency-domain (MSE_{freq}) measures of performance.	54
4.1 <i>Tested Conditions for Interface Controller Dynamics.</i>	82
5.1 The self-defined da/cc(s) of each participant.	96
5.2 Summary of participant suggested features to include when developing technology to support at-home PT and access barriers that are addressed with the features.	103

Chapter 1

INTRODUCTION

High-bandwidth closed-loop interfaces using biosignals measured from sensors in, on, or around our bodies have the potential to revolutionize the way we interact with devices, machines, and each other. Conventional interfaces like a computer mouse transfer information from our brains to machines slowly and inflexibly due to the limited number of communication channels available. Increasing the bandwidth of these interfaces through biosignals (e.g., multi-channel electromyography (EMG) signals, inertial measurement unit (IMU) signals), can take advantage of the multiple modes of communication that our bodies have to potentially enable faster, more intuitive, and always available transfer of information. **My research focuses on how biosignals can be used to more intuitively and accessibly control human-machine interfaces (HMIs) for machines like vehicles, computing devices, or virtual reality avatars, in a variety of situations and for many different goals.** HMIs are rapidly being developed for commercial use, but algorithms and sensor placements that make assumptions about the user’s abilities may be inaccessible for many groups, such as people with disabilities [95, 108, 183], and are not optimized for any particular user’s abilities, goals, or situation.

To support technology accessibility, the philosophy of ability-based design recommends developing “*a system that is aware of the abilities of the user and provides an interface better suited for those abilities*” [178]. Individualized interfaces can be developed to learn and be optimized to each user’s abilities [56, 177]. Such personalized technology design could be beneficial for healthcare applications as well [9], especially for encouraging actions that support functional recovery for people with short- or long-term disabilities. Individualized interface design could also be beneficial for improving the performance of biosignal-based interfaces for people without disabilities. Developing safe and high-performing interfaces that generalize to multiple users is challenging due to the variability in the efficacy of

biosignal-based interfaces across users [190] and variability in biosignals within a single user over time [185]. Therefore, for invasive and non-invasive biosignal control, individualized interfaces that are calibrated for each user are ideal [20, 35, 87, 124, 120, 140, 193]. However, many of these interfaces are machine learning-based and the stability and safety of such interfaces when in closed-loop with a human user are not guaranteed [3, 90].

Control theory [12] provides methods for assessing the stability and safety of closed-loop linear time-invariant (LTI) systems [12, Ch.3, pg. 4]. Many continuous HMIs, where the human is continuously providing input to the machine to produce a desired output, involve the human controlling a machine with LTI *dynamics*. Some common dynamical systems are zeroth- (e.g., position of mouse corresponding to cursor position on screen), first- (e.g., position of joystick corresponding to velocity of wheelchair), or second- (e.g., force on gas pedal corresponding to acceleration of car) order systems. One of the earliest instantiations of modeling the HMI as an LTI system arose from modeling humans manually controlling vehicles [105]. The McRuer gain-crossover model demonstrated that when a human user is in the loop with an LTI machine and performing a disturbance-rejection task, the human will behave like an LTI system. Additionally, near the crossover frequency (i.e., where the open-loop transfer function magnitude of the human and the machine is equal to 1), the open-loop behavior can be modeled as a first-order system with a delay [105]. Later studies demonstrated that during reference-tracking tasks, humans invert the machine dynamics to implement as their feedforward controller [82, 189, 191] and that feedforward model formulation improves with practice [192]. Although a control theory-based HMI model holds promise for enabling the development of individualized closed-loop HMIs, there has been limited work investigating the application of this model for biosignal-based HMIs, especially when humans are tasked with controlling machines with dynamics [31, 88]. Additionally, the gain-crossover model solely accounts for how the human user learns to control a static machine and does not consider how the stability of the closed-loop system will evolve when the interface adapts as well.

Co-adaptive HMIs, where the human and the interface adapt to jointly complete a desired task or goal, are ideal for developing individualized HMIs that adapt to the user’s abilities. However, co-adaptive HMIs present a “two-learner” problem that presents unique

challenges, due to the human and the interface jointly learning in closed-loop. Game theory has been proposed in recent years as a framework for studying these two-learner dynamics in sensorimotor control [23, 84, 85, 93, 112] and provides techniques for predicting convergence to and stability of equilibrium points in two-learner systems [19, 131]. Game theory provides a promising avenue for developing closed-loop individualized interfaces that are stable and adapt to changes in user needs.

The goal of this dissertation is to develop the tools necessary to build individualized, adaptable HMIs for people with and without disabilities. I approached this goal from both a quantitative and qualitative perspective. First, by integrating quantitative tools from control and game theory with biosignal-based interfaces, my work enables the development of stable, closed-loop, individualized HMI algorithms that adapt to users' abilities. These methods provide guarantees on the stability and safety of biosignal-based HMIs and enable the prediction of HMI convergence. Additionally, such methods potentially provide predictive methods to shape user learning as human users learn to control adaptive interfaces. Second, my qualitative work on identifying user needs for individuals with disabilities ensures that the algorithms and technologies that are developed will be directly beneficial to such users. Incorporating user perspectives during the development of novel interfaces is important to ensure that the developed systems are actually usable for the intended user groups. Both quantitative and qualitative perspectives are critical for ensuring that biosignal-based HMIs fully support the accessibility and health of people with and without disabilities.

1.1 Dissertation Contributions and Organization

The major insights of my work are twofold:

1. next-generation biosignal-based interfaces require the development of stable individualized interfaces where its effects are predictable in closed-loop with a human user,
2. ability-based interfaces present opportunities for improving the accessibility of interfaces, thus enabling people with disabilities to encourage actions that support functional recovery during activities of daily living.

Each chapter of my thesis aims to build tools that enable the development of individualized biosignal-based HMIs that support accessibility and health for people of all abilities. Below, I summarize the main findings and contributions of each aim.

1.1.1 Chapter 2: Model Continuous HMIs Using Control Theory

The first step towards individualized, co-adaptive HMIs is to **develop models of how users learn to control continuous HMIs**. We applied control theory methods to quantify how human feedback and feedforward controllers adapt when participants without disabilities are tasked with a combined disturbance-rejection and trajectory-tracking task [181]. We found that participants quickly improved tracking and disturbance-rejection performance and the performance improvement transferred when they switched hands [186]. We additionally found that feedback but not feedforward controllers adapted with practice, and sensorimotor noise decreased with practice. These findings suggest that learning transfers between hands and individuals without disabilities may benefit from an interface that can augment the person’s feedback controller. The main contribution of this aim was to establish a predictive model of how humans control static LTI machines during an unpredictable reference-tracking and disturbance-rejection task and leverage the model to quantify how learning is transferred between dominant and non-dominant hands.

1.1.2 Chapter 3: Leverage Electromyography Interfaces to Enhance HMI Control

Another step towards enabling closed-loop biosignal control is to **quantify how biosignal-based HMIs affect human feedback and feedforward controllers**. Understanding how biosignal control differs from manual control is beneficial for developing HMIs that take advantage of the benefits of biosignals. We quantified the effects of interfaces (EMG versus manual) and machine dynamics (first- versus second-order) on trajectory-tracking and disturbance-rejection performance for participants with and without disabilities [187]. Participants without disabilities formed better feedforward controllers that more closely approximated the inverse of the machine dynamics with the EMG interface than the manual interface when controlling a second-order machine. Participants with disabilities performed

much worse than participants without disabilities at the tracking task when comparing the time-domain tracking error. However, participants with disabilities had comparable feed-forward model formulation to participants without disabilities, suggesting that augmenting the user’s feedback controller with a co-adaptive interface may improve device accessibility. The main contribution of this aim was to demonstrate that control theory can be used to quantify differences between EMG and manual interfaces that can be leveraged to develop HMIs that take advantage of the benefits of both interfaces.

1.1.3 Chapter 4: Develop Co-Adaptive HMIs

One step towards developing co-adaptive HMIs is to **define challenges with co-adaptive HMIs for reference-tracking and disturbance-rejection tasks and quantify through simulation and experiment how humans and interfaces co-adapt to jointly complete a disturbance-rejection task**. We highlighted challenges with applying traditional control theory methods to quantify how humans and interfaces co-adapt to reference-tracking and disturbance-rejection tasks [184]. We then demonstrated in simulation and in experiment how co-adaptation improved performance and decreased human effort compared to baseline. This suggests that further development of theory that can be applied for continuous tasks is needed and that we can also begin to explore this application area through simulations and experiments. The main contribution of this aim was to define challenges with applying traditional methods to the problem of reference-tracking and disturbance-rejection in continuous HMIs and identify how co-adaptation affects final performance and human and interface effort.

1.1.4 Chapter 5: Define User Needs to Support Accessibility and Health with HMIs

Lastly, we cannot begin to develop biosignal-based HMIs that support people’s accessibility and health without direct input and co-development with people with disabilities. Through a qualitative interview study, we **characterized how people with disabilities envision biosignal-based HMIs supporting their health and accessibility**. Specifically, we investigated how biosignal-based HMIs could support access to physical therapy for people

with disabilities [183]. We identified recommendations and tensions to consider when leveraging technology to support physical therapy, and conclude with insights on how technology could potentially improve access to healthcare.

1.1.5 Chapter 6: Conclusions and Future Work

I conclude the dissertation by summarizing my main contributions and discussing potential future extensions in enabling accessible HMIs that support people's health. My work lays the foundation for developing biosignal-based HMIs that improve accessibility and health for people with and without disabilities.

Chapter 2

**EFFECT OF HANDEDNESS ON LEARNED CONTROLLERS AND
SENSORIMOTOR NOISE DURING TRAJECTORY-TRACKING**

IEEE Transactions on Cybernetics, September 2021

DOI 10.1109/TCYB.2021.3110187

Momona Yamagami

Lauren N. Peterson

Darrin Howell

Eatai Roth

Samuel A. Burden

Abstract

In human-in-the-loop control systems, operators can learn to manually control dynamic machines with either hand using a combination of reactive (feedback) and predictive (feed-forward) control. This article studies the effect of handedness on learned controllers and performance during a trajectory-tracking task. In an experiment with 18 participants, subjects perform an assay of unimanual trajectory-tracking and disturbance-rejection tasks through second-order machine dynamics, first with one hand then the other. To assess how hand preference (or dominance) affects learned controllers, we extend, validate, and apply a nonparametric modeling method to estimate the concurrent feedback and feedforward controllers. We find that performance improves because feedback adapts, regardless of the hand used. We do not detect statistically significant differences in performance or learned controllers between hands. Adaptation to reject disturbances arising exogenously (i.e., applied by the experimenter) and endogenously (i.e., generated by sensorimotor noise) explains observed performance improvements.

2.1 Introduction

Humans interact with dynamic machines and devices such as computers, quadrotors, and cars in daily life. These interactions give rise to a *human-in-the-loop* control system where the human and the machine jointly accomplish a task through one or more sensorimotor loops. For instance, in trajectory-tracking tasks, people can visually observe the machine and provide input through a manual interface like a mouse, joystick, or steering wheel [40, 105, 137, 164, 181, 189, 192, 187, 111]. In such cases, people learn to steer computer cursors, quadrotor drones, and personal vehicles using *visuomotor control*. Such manual interfaces often prescribe *how* we interact with the system: some tasks are performed with one hand, others require coordination between hands, and still others may use either or both left and right hands (e.g. the mouse, joystick, and steering wheel, respectively). Because performance in tasks involving fine motor control is affected by the hand used [51], we seek to understand how human visuomotor control differs between hands toward developing effective human-in-the-loop systems. For instance, modeling differences in control between hands could be used to improve bimanual interfaces or to assist unimanual interaction when someone’s preferred hand is unavailable due to injury, disease, or circumstance.

Colloquially understood as the “differences between the hands in terms of skill” [51], handedness can be quantitatively assessed with questionnaires (e.g. the Edinburgh Handedness Inventory [115] or Annett Handedness Questionnaire[76]) or observed from dexterity tasks [76] when questionnaires are difficult or unreliable to administer (such as for young children). These assessments suggest that about 63% prefer to use the right hand and about 7% prefer to use the left hand [76]. This means that about 70% of people have a *preferred* or *dominant* hand that is more dexterous than the *non-preferred* or *non-dominant* hand. Ongoing research indicates that the observed differences in dexterity between dominant and non-dominant hands may be due to each hemisphere of the brain specializing for different aspects of limb movements (termed *lateralization*) [43, 75, 139].

Studies in sensorimotor neuroscience suggest that participants learn different sensorimotor skills with their dominant versus non-dominant hand. For instance, when performing a reaching task under the influence of a force field applied by a robotic manip-

ulandum, participants learned to improve final position accuracy for both dominant and non-dominant hands [43]. However, initial movement direction improved only for the participants' dominant hand, which the researchers attribute to changes in predictive (i.e. feedforward) control, whereas the non-dominant hand primarily improved in final error correction, which the researchers attribute to changes in reactive (i.e. feedback) control. These findings suggest that participants rely more on feedforward than feedback control when using their dominant hand, and vice-versa when using their non-dominant hand, for reaching tasks [43, 14, 143, 139, 10].

For continuous trajectory-tracking and disturbance-rejection tasks through (smooth non)linear machines, prior research primarily focused on modeling participants using their dominant hand [40, 105, 137, 164, 162, 181, 187, 189, 192]. The results from these experiments support the hypothesis that humans learn to use a combination of feedback and feedforward control to reject disturbances and track references. However, little is known about the differences between controllers learned with different hands and whether learned controllers transfer between hands [111]. We seek to determine whether the differences in control mechanisms between left and right hands found in rapid reaching tasks [43, 14, 143, 139] extend to continuous trajectory-tracking tasks.

The goal of this paper is to determine whether participants learn different feedback or feedforward controllers when using their dominant versus non-dominant hand during a visuomotor trajectory-tracking task. We extend, validate, and apply a nonparametric system identification method to estimate feedback and feedforward controllers using unpredictable reference and disturbance signals and second-order machine dynamics. Then we experimentally assess differences in sensorimotor learning between the dominant and non-dominant hand and test whether controllers transfer between hands.

We previously reported preliminary results for first-order machine dynamics in a non-archival conference proceeding [181]; this paper extends those results to a second-order system and provides additional support for the underlying assumptions and hypotheses. More significantly, this paper presents new results comparing learned controllers and performance obtained with dominant and non-dominant hands.

Specifically, two groups learned to perform a unimanual trajectory-tracking and disturbance-

rejection task. One group started with their dominant right hand before switching to their non-dominant left hand, and vice-versa for the other group. To assess the effect of handedness on learning and transfer, we compared (i) feedback and feedforward controllers and (ii) performance obtained by the two groups with their dominant and non-dominant hands. We found that handedness did not affect the learned controller during a continuous trajectory-tracking and disturbance-rejection task. Additionally, we provide evidence that improvements in trajectory-tracking performance may be attributed to changes in feedback gain to reject disturbances applied (a) externally by the experimenter, leading to system-level performance improvements only for the group that learned the task with their non-dominant hand first, and (b) internally due to sensorimotor noise.

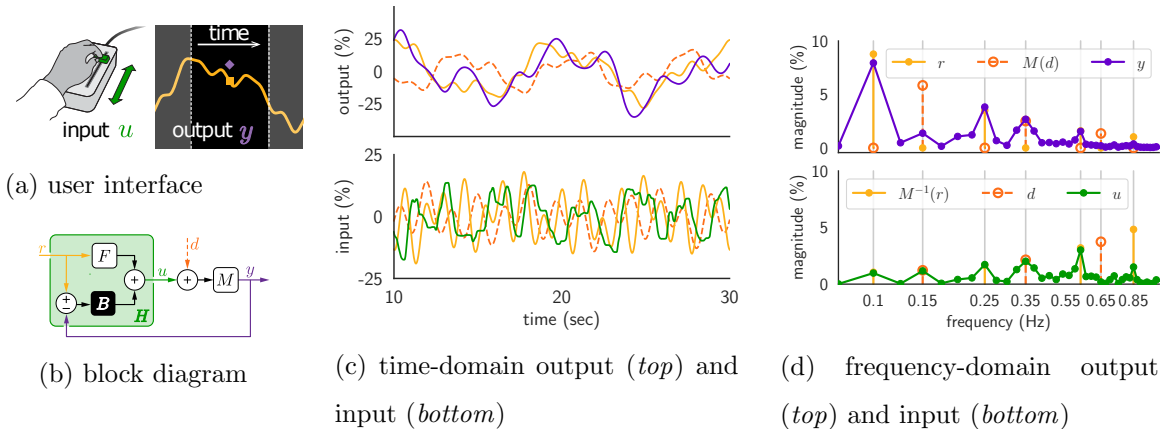


Figure 2.1: *Human-in-the-loop trajectory-tracking*. (a) Human response u is obtained with a one-dimensional manual slider and input to machine M to produce output y , which is overlaid on a display with 1 sec of a reference trajectory (0.5 sec preview). (b) The human H transforms reference r and output y to user response u ; the machine M transforms the sum of control u and disturbance d to output y . We hypothesize that the human’s transformation is the superposition of a *feedforward* F response to reference r and a *feedback* B response to tracking error $r - y$. Representative data from one trial of the **linearity** experiment are shown in (c) the time-domain and (d) the frequency-domain. The frequency content of r and d are confined to prime multiples of a base frequency (1/20 Hz). Magnitudes shown as percent of output or input space extent.

2.2 Problem Formulation

We adopt a tutorial expository style in this section for two reasons. First, to support validation of the assumptions underlying our modeling and analysis methodology, it is important that we explicitly state these assumptions. Second, to support the application of our methods outside the human-in-the-loop controls community, it is valuable to explicitly provide details and rationale that would ordinarily be taken as common knowledge in our niche community. The expert reader may wish to skim or skip this section after reviewing the following table of symbols and Fig. 2.1, returning only if questions arise in subsequent sections.

Table 2.1: *Table of symbols.*

symbol	reference	meaning
u	Fig. 2.1	human response signal
y	Fig. 2.1	machine output signal
r	Fig. 2.1	reference trajectory signal
d	Fig. 2.1	input disturbance signal
T_{zx}	Sec. 2.2.1	transformation from signal x to signal z
$T(x)$	Sec. 2.2.1	transformation of x by T
\hat{x}, \hat{T}	Sec. 2.2.1	Fourier transformation of x, T
M	Sec. 2.2.2	machine transformation: $y = M(u + d)$
H	Sec. 2.2.2	human transformation: $u = H(r, y)$
B	Sec. 2.2.2	human feedback controller
F	Sec. 2.2.2	human feedforward controller

2.2.1 Response to Reference and Disturbance Superimposes

In the laboratory, we instantiate the human-in-the-loop system as a one-degree-of-freedom reference-tracking and disturbance-rejection task (Fig. 2.1) [105]. The transformations that

must take place inside the human (i.e. to observe cursor position, generate a motor plan, and control muscles to move the hand) are known to be nonlinear. However, when tasked with tracking reference r and rejecting additive disturbance d through a linear time-invariant (LTI) [12, Ch. 3, pg. 4] system M , we assume that people behave approximately like LTI transformations for a range of reference and disturbance signals [105, 111, 116, 164, 181, 189]. When this assumption holds, the control signal u produced by the human in response to reference r and disturbance d satisfies the law of superposition,

$$u = T_{ur}(r) + T_{ud}(d), \quad (2.1)$$

where T_{ur} and T_{ud} are LTI transformations.

Hypothesis 2.2.1. *The user response for reference tracking with disturbance is consistent with a superposition of the user response to the reference and disturbance signals presented individually.*

Signals and LTI systems have time-domain and frequency-domain representations as in Fig. 2.1(c,d), related by the *Fourier transform* [126, Ch. 5]; we will adorn signal x and transformation T with a “hat” $\hat{\cdot}$ to denote the Fourier transform \hat{x} , \hat{T} . Importantly in what follows, the frequency-domain operation performed by an LTI system is particularly simple: each frequency component of the input is independently scaled and phase-shifted [12, Ch. 9]. Thus, frequency-domain LTI transformations (termed *transfer functions*) can be empirically estimated by dividing Fourier transforms of time-domain input and output signals at each frequency of interest ω and visualized using a *Bode plot* [126, Ch. 5] as in Fig. 2.4. Specifically, when disturbance $d = 0$ in (2.1) we have

$$\hat{T}_{ur}(\omega) = \frac{\hat{u}(\omega)}{\hat{r}(\omega)} \quad (2.2)$$

and when reference $r = 0$ in (2.1) we have

$$\hat{T}_{ud}(\omega) = \frac{\hat{u}(\omega)}{\hat{d}(\omega)}. \quad (2.3)$$

In contrast, an LTI system's time-domain operation (2.1) – *convolution* [126, Ch. 3] – is mathematically and computationally more complicated than frequency-domain multiplication. For this reason, we design and analyze experiments using frequency-domain representations of signals and systems.

2.2.2 Combined Feedback and Feedforward Improves Prediction

In the absence of reference, (i.e. $r = 0$), we assume that the human response is solely due to a *feedback* B transformation of *tracking error* $e = r - y = -y$ (i.e. $H(0, y) = B(-y)$). If a nonzero reference $r \neq 0$ is known to the human, we assume that it evokes an additive *feedforward* F transformation of r , so that the overall human response can be written as

$$u = H(r, y) = F(r) + B(r - y), \quad (2.4)$$

where $e = r - y$ is tracking error. Using a combination of feedback and feedforward control to model human trajectory-tracking has a long history in the field [40, 42, 41, 105, 111, 163, 137, 181, 189, 192], and is a well-known strategy to improve performance over error feedback alone [12, Ch 8]. We emphasize, however, that certain neurologic conditions like cerebellar ataxia could impair people's ability to perform feedforward control. In such cases, feedback alone may provide better predictions [155, 197].

Under Hypothesis 2.2.1, we can apply *block diagram algebra* [12, Sec. 2.2] to transcribe Fig. 2.1(b) into equations that can be manipulated to express the empirical and prescribed transfer functions \widehat{T}_{ur} (2.2), \widehat{T}_{ud} (2.3), \widehat{M} (2.7) in terms of the unknown transformations F and B ,

$$\widehat{u}(\omega) = \underbrace{\frac{\widehat{F}(\omega) + \widehat{B}(\omega)}{1 + \widehat{B}(\omega)\widehat{M}(\omega)}}_{\widehat{T}_{ur}(\omega)} \widehat{r}(\omega) + \underbrace{\frac{-\widehat{B}(\omega)\widehat{M}(\omega)}{1 + \widehat{B}(\omega)\widehat{M}(\omega)}}_{\widehat{T}_{ud}(\omega)} \widehat{d}(\omega), \quad (2.5)$$

and solve (2.5) to estimate the feedback B and feedforward F components of the human's controller at each stimulated frequency ω ,

$$\widehat{B}(\omega) = -\widehat{M}^{-1}(\omega) \frac{\widehat{T}_{ud}(\omega)}{1 + \widehat{T}_{ud}(\omega)}, \quad (2.6a)$$

$$\widehat{F}(\omega) = \left(1 + \widehat{B}(\omega)\widehat{M}(\omega)\right) \widehat{T}_{ur}(\omega) - \widehat{B}(\omega). \quad (2.6b)$$

If we instead assume that the human’s response to reference r is entirely due to feedback B , then the feedforward controller F estimated in (2.6b) will be approximately zero, so it can be neglected in (2.4) without affecting prediction accuracy.

Hypothesis 2.2.2. *The combined feedback and feedforward model predicts user responses better than a solely feedback model.*

2.2.3 Feedback and Feedforward Adapt with Experience

Previous studies on point-to-point reaching tasks suggest that improvements in end-point accuracy can be attributed to improvements in initial movement (feedforward control) for the dominant hand and improvements in error correction (feedback control) for the non-dominant hand [43, 14, 143], possibly due to specialization of each arm and the corresponding brain hemisphere that controls the arm [43, 75, 139]. These findings lead to the hypothesis that similar observations will hold in the trajectory-tracking task considered here.

Hypothesis 2.2.3. *Human feedback and feedforward controllers will adapt with practice. (a) Feedback will adapt when using the non-dominant hand. (b) Feedforward will adapt when using the dominant hand.*

Note that this hypothesis does not speculate about *how* controllers adapt in the trajectory-tracking task.

2.3 Experimental Methods

Two experiments approved by the University of Washington, Seattle’s Institutional Review Board (IRB #00000909) were conducted to:

(**linearity**) validate the proposed problem formulation and

(**handedness**) assess differences between dominant and non-dominant hands


during sensorimotor learning and control in a continuous trajectory-tracking task.

2.3.1 Manual Interface

Participants used a one-degree-of-freedom manual interface to control the position of a cursor on a screen to track a reference trajectory (Fig. 2.1a). The interface handle was attached to a linear potentiometer; the user response u was determined by measuring the potentiometer voltage using an Arduino Due (`Arduino.cc`). The linear potentiometer had a 10 cm extent, and trials were designed such that the input required to produce the reference trajectory was restricted to the middle third of this physical extent. The handle geometry changed between the **linearity** and **handedness** experiments to improve ergonomics (Fig. 2.2):

(**linearity**) participants used a $35 \times 12 \times 22$ mm (width \times height \times depth) rectangular handle;

(**handedness**) participants used a 35×150 mm (diameter \times height) cylindrical handle.



(a) **linearity** handle (b) **handedness** handle

Figure 2.2: *Handles for **linearity** and **handedness** experiments.* (a) Participants pinched a rectangular handle with their fingers in the **linearity** experiments. (b) Participants grasped a cylindrical handle with their hand in the **handedness** experiments.

2.3.2 Unpredictable Stimuli

Reference and disturbance signals were constructed as a sum of sinusoidal signals with distinct frequencies. Each frequency component's magnitude was normalized by the frequency squared to ensure constant signal power, and the phase of each frequency component was randomized in each trial to produce pseudorandom time-domain signals as in Fig. 2.1(c). A similar stimulus design procedure was employed in [189] to produce unpredictable reference and disturbance signals, and in [40] to produce unpredictable disturbance signals. However, to prevent harmonics from confounding user responses at different frequencies, we adopted the procedure from [136] that restricts stimuli frequency components to prime multiples of a base frequency (1/20 Hz in our experiments). Each trial consisted of two periods of the periodic stimuli (40 sec total) after a 5 sec ramp-up. The number of prime multiples changed between the **linearity** and **handedness** experiments to balance the experiment design:

(**linearity**) first seven prime multiples of base frequency;

(**handedness**) first eight prime multiples of base frequency.

2.3.3 Trajectory-Tracking Task

User response u was transformed through a second-order system with damping to produce output y :

$$M : \ddot{y} + \dot{y} = u + d, \quad \widehat{M} : \frac{1}{s^2 + s}. \quad (2.7)$$

In all experiments, 1 second of reference r was displayed with 0.5 second preview, participants were tasked with adjusting their control u to make a cursor positioned at y track the reference, and the user's response u was modified by an additive disturbance d to determine the machine output $y = M(u + d)$.

Conditions for **Linearity** Experiment

To test the superposition principle (2.1), the three different types of conditions illustrated in Fig. 2.3 were presented to the user in the order shown in TABLE 2.2. In disturbance-only

trials (condition $(0, d)$), the reference r was constant (zero) and the disturbance d was non-constant. In reference-only trials (condition $(r, 0)$), the reference r was non-constant and the disturbance d was zero. In reference-plus-disturbance trials (condition (r, d)), both signals were non-constant, but their frequency components were interleaved as in Fig. 2.3 (*bottom*) to distinguish the user’s response to both signals: specifically, reference or disturbance were active at even or odd multiples of the base frequency (indicated by an E or O subscript, respectively). The two types of (r, d) trials – (r_E, d_O) and (r_O, d_E) – were presented to the participants in alternating order.

Table 2.2: *Conditions for **linearity** experiment (cf. Fig. 2.3).*

Order	1	2	3	4	5
Condition	(r, d)	$(0, d)$	(r, d)	$(r, 0)$	(r, d)
# Trials	2	10	2	10	10

*Conditions for **Handedness** Experiment*

To assess the effects of handedness on feedback and feedforward control, participants were divided into two groups. All participants were right-handed, so we refer to the dominant hand as the “right” hand and the non-dominant hand as the “left” hand. The first group completed 30 (r, d) trials with their dominant right hand, then 30 (r, d) trials with their non-dominant left hand (Group RL). The second group completed the same number of trials, but with their non-dominant left hand first, followed by their dominant right hand (Group LR).

2.3.4 Data Analyses

Measured data and the code reproducing the analyses can be found at [1]. User response u , reference r , disturbance d , and output y were sampled at 60 Hz and converted to frequency-domain representations using the fast Fourier transform (FFT). Data were analyzed using

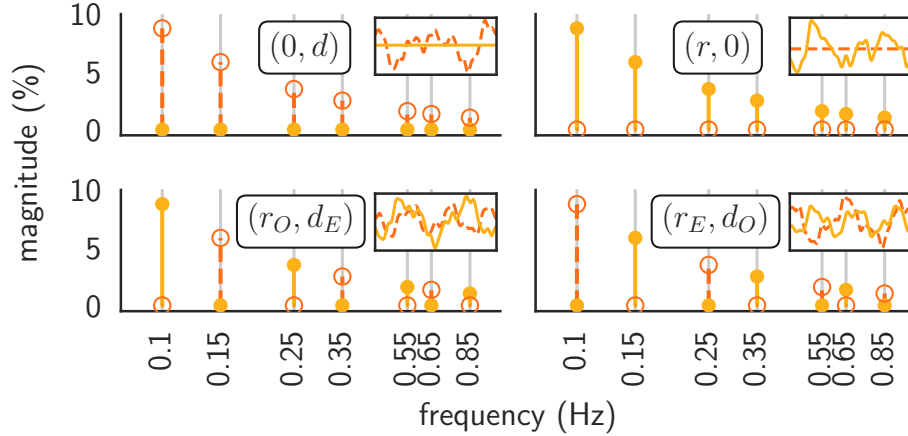


Figure 2.3: *Conditions for **linearity** experiment (cf. TABLE 2.2).* To assess whether the human’s response to external reference r superimposes with the response to external disturbance d , we empirically estimated transfer functions using data from four experimental conditions: disturbance-only $((0, d)$, upper left); reference-only $((r, 0)$, upper right); reference and disturbance interleaved at different frequencies $((r, d)$, bottom left, right). The magnitude of \hat{r} is denoted with solid lines and filled circles, while dashed lines and open circles denote that of $\widehat{M}\widehat{d}$; insets show corresponding time-domain signals r , $M(d)$. Magnitudes shown as percent of output or input space extent.

Python3.5. Transfer functions were estimated at stimulated frequencies from distributions obtained using equations (2.2), (2.3), and (2.6); this simple non-parametric modeling scheme is referred to as the *Fourier coefficients method* [113] or the *spectral measurement technique* [189].

Hypothesis 1

We computed frequency-domain representations of the transformation from disturbance d and reference r to response u (\widehat{T}_{ud} , \widehat{T}_{ur} , respectively) at each stimulated frequency using (2.2) and (2.3). We performed a Wilcoxon signed-rank test with significance threshold $\alpha = 0.05$ to assess whether the magnitudes and phases of \widehat{T}_{ud} and \widehat{T}_{ur} in the $(0, d)$ and $(r, 0)$ trials were

different from those in the (r, d) trials. The Wilcoxon signed-rank test is a non-parametric paired t -test for data that is not normally distributed [30, Sec. 5.7], selected for this study due to the small expected sample size of less than 10 participants (see Appendix A.1.3 for more details). If there are statistically significant differences between transformations estimated from different conditions, it suggests that the human is not well-modeled as an LTI system.

Hypothesis 2

Feedback B was estimated for each participant by applying (2.6a) to data from disturbance-only trials (condition $(0, d)$) and averaging across trials; similarly, feedforward F was estimated for each participant by applying (2.6b) to data from reference-only trials (condition $(r, 0)$), using B that was just estimated from the $(0, d)$ trials and averaging across trials. These controller estimates were used to predict user response \hat{u} by applying (2.5) to data from disturbance-plus-reference trials (condition (r, d)) for the last 10 trials. The coefficient of determination R^2 [60, Eqn. (3.9)] was used to assess prediction accuracy at each frequency (see Appendix A.1.1 for more details). We assessed differences between the R^2 value obtained from the feedback-only (B) model and the feedback-plus-feedforward ($B+F$) model with the Wilcoxon signed-rank test with significance threshold $\alpha = 0.05$. If there is a statistically significant improvement in the R^2 value for the $B + F$ model compared to the B -only model, it suggests that the human response is better modeled with a combination of feedback and feedforward control.

Hypothesis 3

We assessed the performance of each participant using time-domain tracking error computed as the mean-square error (MSE) between reference r and output y :

$$\|r - y\|^2 = \sum_{t \in [0, 40]} |r(t) - y(t)|^2. \quad (2.8)$$

Changes in performance over time were assessed by applying the Wilcoxon signed-rank test with $\alpha = 0.05$ to the average performance of each individual over the first and last five trials

with each hand. To assess differences between Group RL and Group LR, we performed the Mann-Whitney U test, a non-parametric unpaired t -test, with $\alpha = 0.05$.

To assess whether a transformation T changed with practice, we averaged the magnitude of the frequency-domain representation \widehat{T} at stimulated frequencies $\omega \in \{0.10 \text{ Hz}, 0.15 \text{ Hz}\}$,

$$|\widehat{T}| = \frac{1}{2} \left(|\widehat{T}(0.10 \text{ Hz})| + |\widehat{T}(0.15 \text{ Hz})| \right), \quad (2.9)$$

averaged this quantity over the first and last five trials with each hand for each participant, and applied the Wilcoxon signed-rank test with $\alpha = 0.05$. We only included the first two stimulated frequencies in (2.9) since the other stimulated frequencies exceeded the crossover frequency¹ observed in our population, and prior work indicates (and our results corroborate) that reference-tracking and disturbance-rejection performance degrades at frequencies higher than crossover.

This procedure was applied to the estimated human feedforward \widehat{F} and feedback \widehat{B} transformations, as well as the system-level transformations $|\widehat{T}_{yd}|$ and $|\widehat{T}_{yr} - 1|$. Our focus on the latter two transformations is motivated by the observations that the disturbance is rejected if $\widehat{T}_{yd} = 0$ and the reference is tracked if $\widehat{T}_{yr} = 1$. However, we note that $\widehat{T}_{yr} = \widehat{T}_{yd}(\widehat{F} + \widehat{B})$ (assuming F and B are LTI), so it is not possible for the user to simultaneously achieve $\widehat{T}_{yd} = 0$ and $\widehat{T}_{yr} = 1$ (assuming \widehat{F} and \widehat{B} have finite magnitude). We quantify *system-level performance* at each frequency using statistics for trajectory tracking and disturbance rejection, $|\widehat{T}_{yr} - 1|$ and $|\widehat{T}_{yd}|$. Smaller values correspond to better performance.

Non-Stimulated Frequencies

Our methods can only estimate transfer functions at stimulated frequencies; the denominators in (2.2) and (2.3) are undefined at frequencies ω where $\widehat{r}(\omega) = 0$ or $\widehat{d}(\omega) = 0$, respectively. Although we expect the power of the user response signal to be concentrated at these stimulated frequencies, we nevertheless measure user response at intermediate non-stimulated frequencies (see Fig. 2.1(d)). Since any user response at non-stimulated frequencies degrades task performance (there is no reference to track or disturbance to reject), we

¹Frequency where gain of loop transfer function $\widehat{L} = \widehat{B}\widehat{M}$ equals 1 [105].

use the deviation of $|\hat{u}|$ from 0 as another way to quantify task performance. Tracking error (3.1) is affected by user response across both stimulated and non-stimulated frequencies, which we separately quantify using system-level performance and $|\hat{u}|$.

2.4 Results

We recruited participants from the greater University of Washington community: 7 for the **linearity** experiment, and an additional 18 (9 male, 9 female; age 18-32; height 145-190 cm; weight 48-98 kg) for the **handedness** experiment.² The participants had no reported neurological or motor impairments and all were daily computer users.

2.4.1 Response to Reference and Disturbance (Approximately) Superimposed Across Conditions

We tested Hypothesis 2.2.1 with the **linearity** experiment to determine whether user response u in disturbance-only $(0, d)$ or reference-only $(r, 0)$ conditions was consistent with user response in disturbance-plus-reference conditions (r, d) (Fig. 2.4). The magnitude and phase of the transfer functions from d and r to u (\hat{T}_{ud} and \hat{T}_{ur} , respectively) estimated from these different conditions were indistinguishable at most stimulated frequencies ($p > 0.05$; exceptions denoted with † in Fig. 2.4), indicating that participants' response to reference and disturbance signals approximately satisfied the law of superposition across the qualitatively different conditions in Fig. 2.3.

2.4.2 Combined Feedback and Feedforward Improved Prediction

We tested Hypothesis 2.2.2 with the **linearity** experiment to determine whether a combined feedback-plus-feedforward $(B + F)$ model improves prediction compared to a feedback-only (B) model (Fig. 2.5). Predictions for both models were better (R^2 closer to 1) below crossover frequency (0.25 Hz, determined as the lowest stimulated frequency where the gain of the open-loop transfer function $\hat{L} = \hat{B}\hat{M}$ is less than 1 [105, 189]), and decreased in accuracy (R^2 closer to 0) at higher frequencies, suggesting the linear models were more

²Demographics were not recorded for the **linearity** experiment.

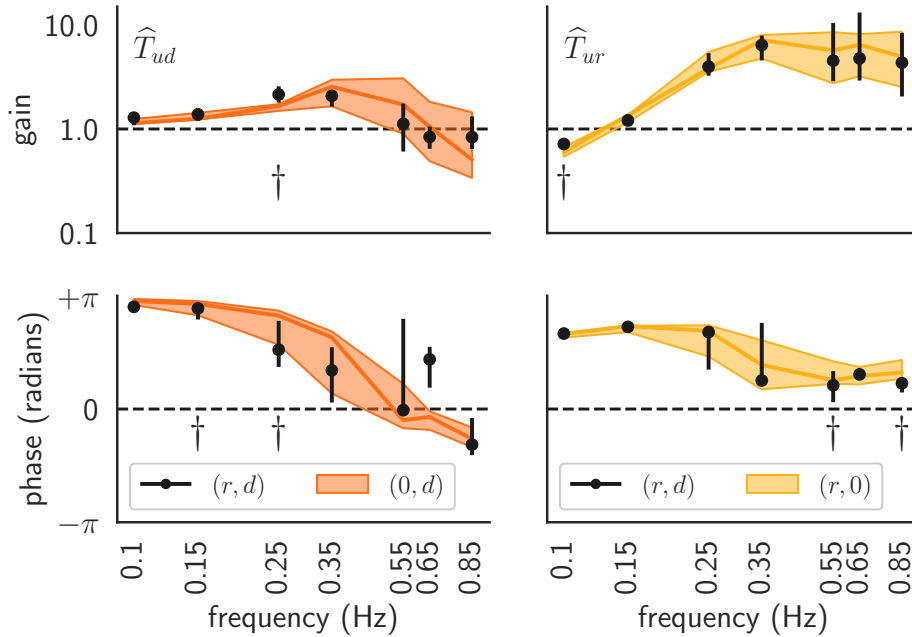


Figure 2.4: *Transfer function estimates in **linearity** experiment.* Distributions (median, interquartile) of transfer functions \hat{T}_{ud} (*left*), \hat{T}_{ur} (*right*) estimated from disturbance-only or reference-only trials, $(0, d)$ or $(r, 0)$, and reference-plus-disturbance trials (r, d) , for the conditions in TABLE 2.2 and Fig. 2.3. Statistically significant differences (Wilcoxon signed-rank test: $p < 0.05$) in distribution magnitude or phase at each frequency indicated with †.

accurate at lower frequencies. Prediction accuracy for the $B + F$ model was higher than the B model at all frequencies ($Z = 0.0, p = 0.016$), suggesting that user responses u to references r and disturbances d are better predicted with a combined feedback-plus-feedforward ($B + F$) model than a feedback-only (B) model.

2.4.3 Performance Improved and Feedback Adapted

We tested Hypothesis 2.2.3 with the **handedness** experiment (Fig. 2.6 and Fig. 2.7) to determine whether task performance changed with practice using time-domain reference

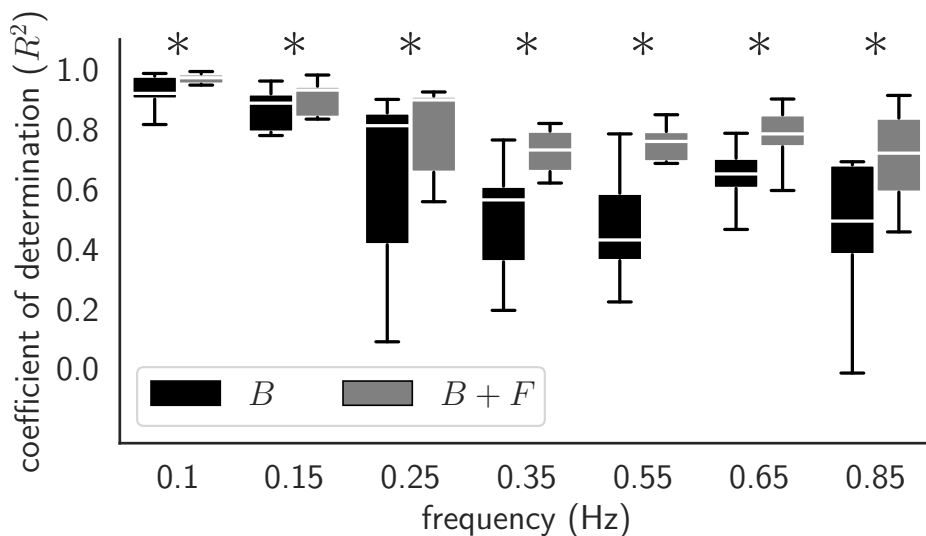


Figure 2.5: *Predictive accuracy of models, **linearity** experiment.* Distribution (median, interquartile, range) of coefficient of determination (R^2) between human inputs u and predictions from feedback-only (B) and feedback-plus-feedforward ($B + F$) models. The $B + F$ model had significantly better prediction accuracy than the B model at all frequencies (Wilcoxon signed-rank test: $Z = 0.0, p = 0.016$; indicated with *).

tracking error $\|r - y\|^2$ from (3.1). We found that performance improved rapidly within the first five trials and then did not change significantly, even after switching hands, regardless of which hand was used first (Fig. 2.7a). Performance improved significantly between the first and last five trials with the first hand (trials #1–5 and #26–30; Group RL: $Z = 0.00, p = 0.004$; Group LR: $Z = 0.00, p = 0.004$), and did not change significantly between the last five trials with the first hand and the first five trials of the second hand (trials #26–30 and #31–35; Group RL: $Z = 21.0, p = 0.86$; Group LR: $Z = 19.0, p = 0.68$). We did not find statistically significant differences between Group RL and Group LR in the first or last five trials with either hand (Mann-Whitney U : $p > 0.05$).

To determine whether improvements in $\|r - y\|^2$ could be attributed to changes in feedback or feedforward control, we assessed whether feedback B or feedforward F control changed with practice using the mean magnitude of the frequency-domain representation




Figure 2.6: *Tracking error from handedness experiment.* Distributions (median, interquartile) of time-domain tracking error $\|r - y\|^2$ for 60 trials, with a switch between dominant (right; red circles) and non-dominant (left; blue squares) hands after trial 30, for two groups of 9 participants: (*top*) right then left (Group RL); (*bottom*) left then right (Group LR). Summary statistics in Fig. 2.7 use data from first five and last five trials with each hand, highlighted with light and dark gray boxes.

$|\hat{B}|$ or $|\hat{F}|$ from (2.9). The mean magnitude of the feedback controller increased with practice for both groups ($Z = 3.0$, $p = 0.02$ in both) between the first and last five trials with the first hand, and did not change when switching to the second hand ($p > 0.05$). There was no statistically significant change in the mean magnitude of the feedforward controller across all conditions ($p > 0.05$) (Fig. 2.8). We did not find statistically significant changes in the phase of F and B at any stimulated frequency.

We observed system-level performance improvements at the first two stimulated frequencies (0.10, 0.15 Hz) solely for Group LR (Fig. 2.7). Group LR significantly decreased both $|\hat{T}_{yr} - 1|$ ($Z = 4.0$, $p = 0.028$) and $|\hat{T}_{yd}|$ ($Z = 0.0$, $p = 0.004$) through experience with their

first (left) hand, indicating significant improvements in reference tracking and disturbance rejection. This improved performance persisted even after switching from the left hand to the right hand, suggesting some transfer of knowledge between hands.

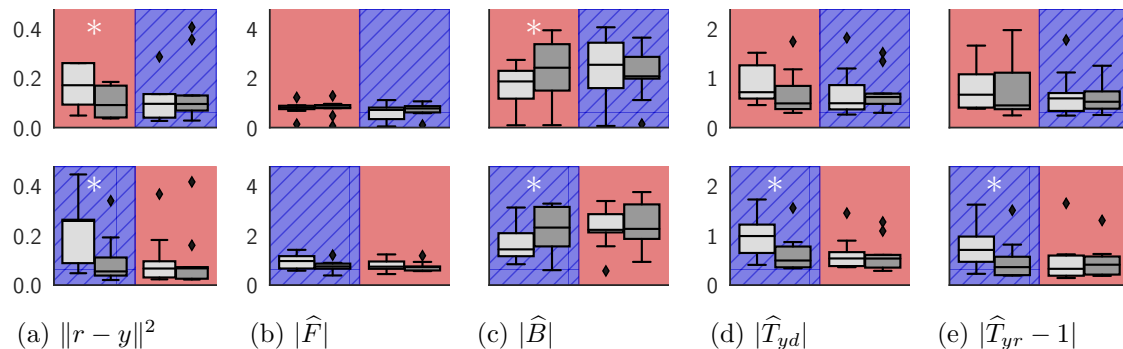


Figure 2.7: *Summary statistics from **handedness** experiment.* Distributions (median, interquartile, range) from first five (light gray box) and last five (dark gray box) of 30 trials with dominant (red solid background) and non-dominant (blue hatched background) hands: (a) tracking error $\|r - y\|^2$; mean magnitude of (b) feedforward $|\hat{F}|$ and (c) feedback $|\hat{B}|$ controllers (shared y axis); mean magnitude of (d) disturbance rejection $|\hat{T}_{yd}|$ and (e) reference tracking $|\hat{T}_{yr} - 1|$ errors (shared y axis). Statistically significant (Wilcoxon signed-rank test: $p < 0.05$) differences between adjacent distributions indicated with *. Group RL in top row, Group LR in bottom row, as in Fig. 2.6.

2.4.4 User Response Diminished at Non-Stimulated Frequencies

Although we saw significant improvements in tracking performance with practice, we only observed modest or no improvements in system-level performance at stimulated frequencies. These results led us to consider user response at non-stimulated frequencies, since attenuating such response improves tracking performance. For both groups, the magnitude of the response at non-stimulated frequencies below crossover (0.25 Hz) decreased significantly between the first and last five trials with the first hand (trials #1–5 and #25–30) (Fig. 2.9), and this diminished response transferred between hands.




Figure 2.8: *Human feedback (B) and feedforward (F) controllers.* Distributions (median, interquartile) obtained by pooling data from the last five trials with each hand for both groups in the **handedness** experiment; we did not observe statistically significant differences between groups or hands (Wilcoxon signed-rank test: $p > 0.05$).

2.5 Discussion

Prior work demonstrated that people adapt feedback and feedforward controllers differently with the dominant and non-dominant hands during reaching tasks [43, 14, 143, 139]. However, little is known about how handedness affects learned controllers in continuous trajectory-tracking tasks such as the one considered in this study. When subjects reach to targets, feedback and feedforward control are assumed to be episodic: the initial ballistic motion is attributed to solely feedforward control (since sensorimotor delays preclude feedback) whereas corrective motions in the latter stage of the reach are attributed to solely feedback control [43, 14, 143, 139]. In contrast, feedback and feedforward processes are engaged simultaneously when subjects track continuous trajectories as in our experiments.

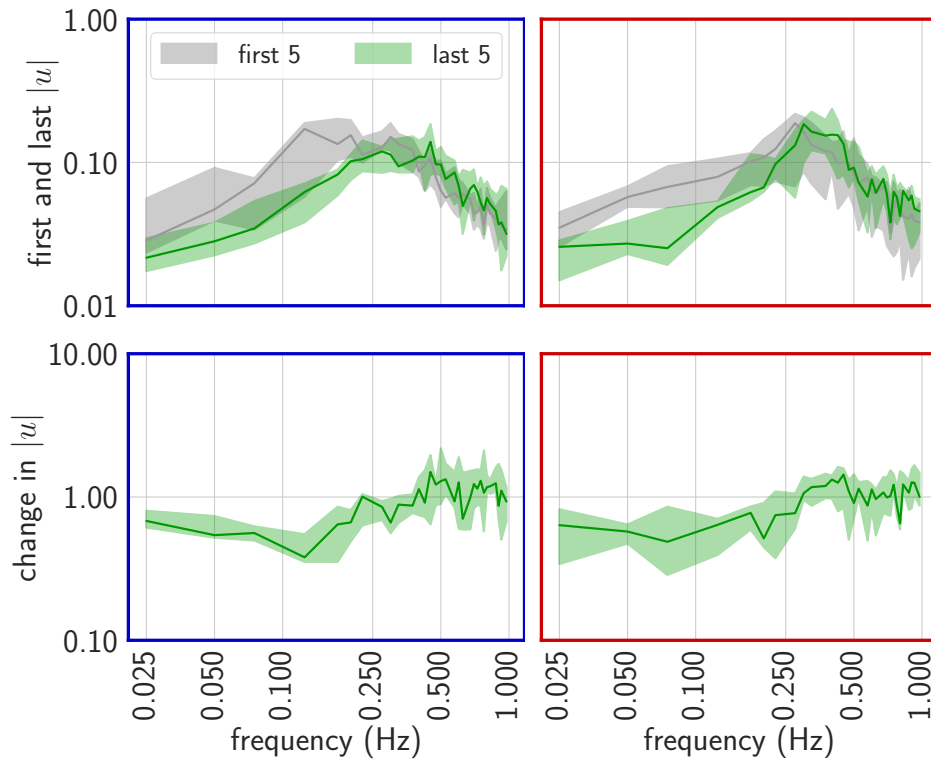


Figure 2.9: *Change in effect of sensorimotor noise.* (top row:) Distributions (median, interquartile) of magnitude of user response at non-stimulated frequencies from first and last five trials with first hand (trials #1–5 in light gray and #26–30 in green) in **handedness** experiment. (bottom row:) Ratio of user response magnitudes between first and last five trials with first hand decreases significantly below crossover (0.25 Hz). Group LR in left column, Group RL in right column.

To assess how feedback and feedforward controllers are learned through experience and transferred between hands in a trajectory-tracking task, we extended, validated, and applied a non-parametric system identification method (adapted from [105, 181, 189, 187]). We found that feedback and feedforward controllers estimated for different hands were not distinguishable and that learned controllers transferred between hands. Trajectory-tracking performance improved significantly with practice, but system-level performance improvements were significant only for the group that learned the trajectory-tracking task with

their non-dominant hand first. Surprisingly, we did not find significant adaptation of the feedforward controller across the sample population. Instead, performance improvements can be attributed to a significant increase in feedback gain below crossover frequency; this accounts for significant changes in the effect of disturbances applied both externally by the experimenter and internally by sensorimotor noise.

2.5.1 Response to Reference and Disturbance (Approximately) Superimposed Across Conditions

We found small but statistically significant differences between the transformations T_{ud}, T_{ur} estimated using data from disturbance-only $(0, d)$ and reference-only $(r, 0)$ trials and the combined reference-and-disturbance (r, d) trials. Thus, the controllers implemented by our participants to control a second-order system do not satisfy the superposition principle (2.1) as well as in our previous findings for first-order systems [181]. We attribute this difference to the increased difficulty of the trajectory-tracking task for a second-order system. However, considering how different each of the experimental conditions in Fig. 2.3 are from the user’s perspective – namely, that $(r, 0)$ trials only have reference, $(0, d)$ trials only have disturbance, and (r, d) trials have both reference and disturbance – we regard the empirical transformations in Fig. 2.4 as remarkably consistent across the qualitatively different conditions in Fig. 2.3.

Similarly to our previous findings for first-order systems [181], we found higher variability in estimates of transformation magnitude at higher frequencies compared to lower frequencies. Thus, although we found evidence that our human-in-the-loop control system is mildly nonlinear, neglecting this nonlinearity nevertheless yields good predictions for the human’s learned controllers, so our results support Hypothesis 2.2.1 with caveats.

Although human behavior is richly varied and nonlinear in general, our results support the assumption that people can behave remarkably linearly after sufficient experience interacting in closed-loop with a linear time-invariant system [40, 170, 171, 189, 181, 187, 192]. Previous studies have ensured that human-in-the-loop-systems are approximately linear by using experts such as pilots [105] or only collecting data after participants undergo prac-

tice [40, 171]. Because our experiments commenced immediately without providing time for participants to explore the interface or machine dynamics (let alone become experts), this lack of practice may have contributed to the mild nonlinearities we observed. Future studies may benefit from estimation of nonlinearity [170], especially during learning.

2.5.2 *Combined Feedback and Feedforward Improved Prediction*

We observed significant improvements in prediction accuracy with the feedback-plus-feedforward model compared to the feedback-only model at all frequencies in Fig. 2.5. This improvement in prediction accuracy implies that a model selection procedure based on an information criterion [98] would favor the combined feedback-plus-feedforward model over the feedback-only model if prediction accuracy was prioritized over model simplicity. Thus, our results lend further support for Hypothesis 2.2.2, consistent with previous results for first-order [40, 181, 189, 187] and fourth-order [171, 192] systems.

Our system identification method assumes the human controller consists of parallel feedback and feedforward controllers. However, the method does not assume or require either controller to be non-zero; in particular, if participants did not employ feedforward control, our method would yield a feedforward estimate with negligible magnitude. We emphasize that including *both* reference-tracking *and* disturbance-rejection in the task is necessary to ensure we can solve two independent equations in two unknowns (2.6) to uniquely determine feedback and feedforward controllers using our non-parametric modeling method.

2.5.3 *Performance Improved Because Feedback Adapted*

Regardless of which hand was used first, participants significantly improved tracking performance through experience with their first hand. This improvement in time-domain performance persisted when participants switched hands, suggesting that learned controllers transferred between hands. Since we observed corresponding significant increases in feedback gain and observed no significant change in feedforward, we attribute this performance improvement to changes in feedback. These findings lead us to *reject* Hypothesis 2.2.3.

Our Hypothesis 2.2.3 was motivated by previous studies of human sensorimotor learning during reaching tasks that suggest improvements in end-point precision were due to improvements in initial movement (feedforward control) for the dominant (right) hand and improvements in error correction (feedback control) for the non-dominant (left) hand [43, 14, 143, 139]. However, there are differences between target-reaching tasks and the trajectory-tracking task used in this current experiment. For instance, the target-reaching tasks in [43, 14, 143, 139] are brief (approximately 1 sec in duration), so feedforward control is thought to dominate user response for a significant fraction of each trial since visual feedback is delayed by approximately 250 msec, and the target’s location changes discontinuously when the trial begins. In contrast, feedback and feedforward are engaged simultaneously for the entire 40 sec duration of each of our trajectory-tracking trials, and the reference changes continuously throughout the trial. The differences in experiment design could account for the differences we observed in how feedback and feedforward adapt. Since increasing the difficulty of a target-reaching task affects adaptation of feedback and feedforward [147, 176], it is possible that changing the machine dynamics or user interface may affect adaptation of feedback and feedforward in trajectory-tracking tasks.

Our inability to detect adaptation in feedforward control over a 1-hour period is inconsistent with previously published research that demonstrated adaptation of feedforward control over a 2-week period [192]. However, there are significant differences between our study methodology and [192] that may explain why we did not observe feedforward adaptation. First, the participants in [192] were tasked with learning to track a fourth-order system, which is significantly more complex than the second-order system used here, and the differing location of machine poles and zeros may affect learned controllers and tracking performance [191]. Second, since many of our participants reported prior experience controlling second-order systems (e.g. driving cars, playing video games), they may have employed a previously-learned feedforward controller in our experiment. Third, there was a significant difference in practice time between the two studies. In [192], participants learned the system dynamics over two weeks, whereas in our study, participants learned the system dynamics over 1 hour. Although we observed performance plateau during the 1-hour study, a longer practice time over the course of days or weeks may result in significant adapta-

tion of feedforward control. Finally, and most significantly, while participants in [192] were tasked with following a predictable chirp trajectory, we tasked our participants to track unpredictable sum-of-sine trajectories. Stimuli predictability is known to affect tracking performance for human-in-the-loop systems [189, Fig. 5] [110, Fig. 5], possibly due to the use of internal *signal generators* [59, 32] (as opposed to the internal *controllers* posited here).

2.5.4 Adaptation of Feedback Improved System-Level Performance For Group LR

To determine whether adaptations in feedback controller gain lead to system-level improvements in performance, we looked for differences in $\hat{T}_{yd}, \hat{T}_{yr}$ at the first two stimulated frequencies (0.10, 0.15 Hz) by comparing the first five and last five trials with each hand. For Group LR, we saw improvements in both \hat{T}_{yd} and \hat{T}_{yr} with their first hand, suggesting that reference tracking and disturbance rejection both improved. Despite clear improvements in time-domain performance for both groups, we did not observe statistically significant improvements in system-level performance at the stimulated frequencies for Group RL.

One possible explanation for these findings is that participants may initially find it more challenging to perform the trajectory-tracking task with their non-dominant hand, producing a larger effect that was easier to detect statistically. Consideration of sample size provides an alternative explanation for observed system-level differences in group performance that points to interesting directions for future study. Group LR and Group RL were relatively small populations (9 participants in each group), so there may have been unmeasured group-level differences. For instance, participants reported subjective differences in the strategy they employed to improve tracking performance. Some participants acknowledged that they were controlling the cursor acceleration and consciously altered their response accordingly, while others mainly focused on reactively minimizing tracking error. Future experiments with a larger number of participants are needed to determine whether different subpopulations employ different strategies when learning controllers.

2.5.5 *Adaptation of Feedback Affected the Effect (but not the Source) of Sensorimotor Noise*

Since time-domain tracking performance improved significantly for both groups of participants but rejection of disturbance stimuli and tracking of reference stimuli only improved for one group, we are led to consider user response at frequencies we measured but did not stimulate. Any user response at non-stimulated frequencies degrades time-domain tracking performance, so it is in the users' best interest to suppress this response [166]. We observed nonzero user response at non-stimulated frequencies, and this response decreased significantly with practice for frequencies below crossover for the first hand in both groups (Fig. 2.9). This suggests that, instead of (or in addition to) improving performance of disturbance rejection and trajectory-tracking at stimulated frequencies, the participants suppressed their response at low non-stimulated frequencies, leading to improved time-domain performance.

Because the machine dynamics and feedback in Fig. 2.1(b) are linear time-invariant, the user response at non-stimulated frequencies arises due to (i) nonlinearity in the human's transformation and/or (ii) sensorimotor noise. Although we found evidence for (i) mild nonlinearities (see Fig. 2.4 and Sec. 2.5.1), we tested for but did not find significant coherent responses in the user response at harmonics of the stimulated frequencies (i.e. non-stimulated frequencies), so nonlinearity alone does not appear to explain our observations. Assuming instead that user response at non-stimulated frequencies arises solely due to (ii) additive sensorimotor noise, we did not find statistically significant changes in this noise with experience. Indeed, despite the fact that we observed significant changes in feedback B and user response u at non-stimulated frequencies, we observed no significant changes in the power spectrum of the imputed disturbance $\delta = (1 + MB)u$. Instead, the *effect* of the noise was attenuated by the increase in feedback gain below crossover. This result is consistent with prior studies from sensorimotor control that found the presence of significant noise whose statistics did not change with the limited amount of practice (less than 1 hour) considered here [73].

2.5.6 *Does Stimulus or Noise Drive Learning?*

When learning to perform novel tasks like controlling a cursor on a screen or reaching under a force field, sensorimotor noise and movement variability are crucial for driving learning [38, 68, 180]. As people explore the action space for a particular task, certain movements (e.g. tracking a trajectory with specific frequency components) result in greater reward (e.g. improved tracking) [38]. With significant practice, noise and variability decreases, leading to improved performance in ballistic throwing [66, 73] and reaching [180] tasks. Similarly, we argue here that our observations that 1) there was time-domain improvement, 2) there was no corresponding system-level performance improvement at stimulated frequencies, and 3) user response decreased at non-stimulated frequencies below crossover, suggests that reducing the effect of sensorimotor noise may be a crucial aspect of performance improvement in continuous trajectory-tracking tasks. Although out of scope for our study, our results indicate that changes in sensorimotor noise at non-stimulated frequencies should be considered in addition to feedback and feedforward control at stimulated frequencies in studies of human-in-the-loop control systems.

2.6 *Conclusion*

Understanding how humans learn to track continuous trajectories with their dominant and non-dominant hands is crucial for enabling bimanual device control when teleoperating a surgical robot or manipulating objects in augmented or virtual reality. To this end, we first validated a non-parametric modeling method to simultaneously estimate feedback and feedforward control during a second-order continuous trajectory-tracking and disturbance-rejection task with seven participants. We then investigated adaptation of feedback and feedforward control and corresponding system-level changes in performance when nine participants learned to track with their right hand before their left hand, and when nine other participants learned to track with their left hand before their right hand.

Our study demonstrated that: (1) feedback control adapted with practice and transferred between hands in both groups; (2) feedback adaptation improved system-level performance in tracking prescribed references and rejecting externally-applied disturbances for the group

that first learned the task with their non-dominant (left) hand; and (3) feedback adaptation improved tracking performance by attenuating the effect of a user's sensorimotor noise in both groups. These findings suggest that handedness may not affect learned controllers, demonstrate that learned controllers may be transferred between hands, and highlight the importance of attenuating sensorimotor noise for human-in-the-loop control systems.

Chapter 3

**DECODING INTENT WITH CONTROL THEORY: COMPARING
MUSCLE VERSUS MANUAL INTERFACE PERFORMANCE**

In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems
(CHI 2020), April 2020

DOI 10.1145/3313831.3376224

Momona Yamagami

Katherine M. Steele

Samuel A. Burden

Abstract

Manual device interaction requires precise coordination which may be difficult for users with motor impairments. Muscle interfaces provide alternative interaction methods that may enhance performance, but have not yet been evaluated for simple (eg. mouse tracking) and complex (eg. driving) continuous tasks. Control theory enables us to probe continuous task performance by separating user input into intent and error correction to quantify how motor impairments impact device interaction. We compared the effectiveness of a manual versus a muscle interface for eleven users without and three users with motor impairments performing continuous tasks. Both user groups preferred and performed better with the muscle versus the manual interface for the complex continuous task. These results suggest muscle interfaces and algorithms that can detect and augment user intent may be especially useful for future design of interfaces for continuous tasks.

3.1 Introduction

Users predominantly interact with devices using manual interfaces such as mice, touchscreens, steering wheels, and joysticks. However, many of these interfaces may be difficult or impossible to use for individuals with upper-extremity motor impairments after neurologic injury. Such users may have difficulty precisely coordinating arm and hand function to control manual interfaces due to weakness of the arm muscles, spasticity provoking unintended movement, and muscle tightness limiting mobility [63]. The lack of accessibility of manual interfaces for users with motor impairments is well-documented [50, 62, 94, 109]. People with neurologic injuries that impact one side of the body like stroke or cerebral palsy tend to solely use their unaffected side for device interaction [165]. This leads to slower and more error-prone technology use and increases fatigue [158]. Alternatives that can be personalized and require less strength and coordination could encourage greater use and utility of the affected side. Muscle interfaces are one potential alternative to manual interfaces that may enable users with and without motor impairments to interact effectively and unobtrusively with their device [140]. The placement of the muscle sensors can be personalized so that users can adapt the interface to their own ability level [178]. Such interfaces may decrease errors, increase use of the affected side, and enhance long-term function.

In this paper, we investigate the performance of a muscle versus a manual interface for continuous trajectory tracking tasks in users with and without motor impairments using modeling techniques from control theory. While other performance metrics for modeling continuous task performance exist [5, 92], we demonstrate that control theory techniques provide powerful insights not available with other techniques. To the best of our knowledge, there are no methods in human-computer interaction (HCI) that separate and quantify *user intent (feedforward control)* from *error correction (feedback control)*. This could be particularly useful for users with motor impairments. Users with motor impairments after neurologic injury often retain the ability to determine the input needed to control a device to follow a desired trajectory in the absence of errors. However, they may have difficulty correcting for errors that arise from unexpected disturbances like arm tremor [63] (Fig. 3.1). We apply techniques from control theory to decode user intent, providing a foundation for




Figure 3.1: Successfully completing continuous tasks with manual or muscle interfaces is crucial for many tasks including cursor navigation. While a user may intend to follow a desired reference path (dotted red line) with user intent (dotted red arrows), unexpected disturbances (sudden change in cursor position between the two blue circles) introduce errors that must be corrected with error correction (solid blue arrows). The user input (mouse position) combines user intent and error correction and maps to the cursor position on the screen (blue solid line).

future development of HCI algorithms that assist users as they perform continuous tasks like mouse tracking and driving.

We used frequency-domain analysis to separate and quantify feedforward and feedback control for simple (velocity-based) and complex (acceleration-based) continuous tasks using a muscle and manual interface in eleven users without motor impairments. We also studied muscle and manual interface performance for the complex task for three participants with motor impairments. We computed two performance metrics: (i) time-domain error between a desired trajectory and actual cursor position and (ii) frequency-domain error between the user’s feedforward controller and the controller required to perfectly follow a reference in the absence of errors.

The contributions of this paper are threefold:

- C1 extend control theory-based quantitative modeling techniques that separate user intent (feedforward control) and error correction (feedback control) to muscle interfaces;
- C2 experimentally compare muscle versus manual interface performance for simple (velocity-based) and complex (acceleration-based) continuous trajectory tracking tasks;
- C3 conduct preliminary evaluations of muscle versus manual interface performance for a complex continuous task for users with motor impairments.

We report two key experimental findings:

- F1 users without motor impairments were 49% better at tracking continuous trajectories using the muscle than the manual interface for the complex continuous task;
- F2 users without motor impairments were 61% better at tracking high-frequencies above 0.35 Hz with the muscle versus the manual interface.

Our paper proposes and extends an experimental and analytical method to guide future development of accessible interfaces like muscle interfaces using control theory. The results demonstrate the feasibility of using methods from control theory to inform future interface design and develop assistive algorithms to aid users with motor impairments in achieving desired tasks.

3.2 Related Work

3.2.1 Accessible Interfaces for Users With Motor Impairments

Despite technological advancements in personal computing, device accessibility for users with motor impairments and alternate abilities remains a challenge [178]. Researchers have demonstrated how ability-based assumptions underlying traditional interfaces such as touchscreens [50, 62, 159] and mice [50] are inappropriate for users with motor impairments, and how these assumptions can be modified to encompass users of all abilities. Other researchers have worked on using artificial intelligence to adapt current interfaces such as

touchscreens [109, 194] and screen layouts for use with a mouse [56] such that they take into account each user’s ability level. Researchers have also worked on modeling stroke gestures on touchscreens for users with motor impairments [161]

We are interested in whether alternative interfaces could provide performance advantages for continuous tasks. Although it is crucial to understand how traditional interfaces can be adapted for users with motor impairments, novel interfaces like smart watches [94] and headsets [96] are quickly being developed for commercial use. Understanding whether muscle interfaces provide performance benefits for users of all abilities is important for encouraging development of muscle interfaces.

3.2.2 Electromyography as Non-Invasive Muscle Sensors

Although muscle interfaces have gained popularity in research as a hands-free interaction method, muscle electrical signals are most often used in clinical research to quantitatively assess impairments level, track progress, and evaluate clinical interventions for various clinical populations [16, 34, 153]. Electromyography (EMG) sensors are commonly used in these settings to noninvasively measure muscle electrical activity from the skin surface. Dry or wet electrodes passively measure these electrical signals, which can then be relayed to a computing unit for analysis. EMG technology is still limited to short-term use due to low battery life, bulky form factor, high cost, and lack of comfort [13, 48, 123]. Researchers are currently addressing these limitations by developing novel electrodes and hardware for long-term EMG use [129, 185, 188].

3.2.3 Electromyography in Human-Computer Interaction

Muscle interfaces for HCI have mainly focused on gesture classification tasks for hands-free device use for users without motor impairments. Such interfaces have been demonstrated to have high gesture classification accuracy even when: hands are occupied with other objects [140, 141], EMG signals are weak [71], consumer-level EMG sensors are used [64], and a large number of gestures are attempted [8]. These studies demonstrate that users without motor impairments can successfully use muscle interfaces to reliably perform discrete tasks

like tapping and swiping.

Work on enabling discrete interactions with EMG data is crucial for the adoption of muscle interfaces into everyday life, but little work has studied the use of muscle interfaces for continuous tasks or for users with upper-extremity motor impairments. In addition, prior research on gesture classification with EMG data demonstrates the strength of muscle interfaces in scenarios where manual interfaces would be difficult to use, but have not directly compared performance of muscle and manual interfaces.

3.2.4 Continuous Control Using Muscle Interfaces

Prior work on continuous muscle interfaces focused on measuring EMG signals from residual muscles of amputees for prosthetic control. EMG control is desirable for prosthesis users because it requires minimal effort, allows for intuitive device manipulation, and is noninvasive [47, 145]. In this application, EMG signals are measured from the user and fed into a proportionality controller to manipulate position, speed, or acceleration of the prosthesis [54].

Preliminary research on continuous muscle interfaces for prosthetic control is limited and compares manual and muscle interfaces for simple tasks that map the user input to the position or velocity of a cursor on a screen. Researchers [31, 88] performed investigations where they compared force-based, EMG-based, and position-based interfaces for controlling position and velocity of a cursor. They demonstrated that users tracked a desired reference more accurately with force-based and position-based interfaces. However, they also found that users could track higher frequency signals with the muscle interface than with the manual interfaces.

Our study builds on this work by using metrics from control theory to compare simple and complex task performance for users with and without motor impairments. Understanding tradeoffs between muscle and manual interface performance for simple and complex tasks may lead to greater incorporation of muscle interfaces that are more accurate, easier to use, and encourage muscle use in users with motor impairments.

3.2.5 Feedforward Controller Formulation for Manual Interfaces

An emerging technique for modeling continuous human and device interactions is using control theory to separate user input into a feedforward component that expresses the intended output and a feedback component that corrects for errors [40, 105, 137, 182, 189, 192]. The feedforward component can be considered a performance metric to determine whether the user has learned how their user input maps to the device output in the absence of errors (Fig. 3.1). Control theory provides established frequency-domain techniques for separating and quantifying feedforward and feedback controllers for users without motor impairments. In the 1960's, McRuer et al. [105] used trajectory tracking data collected from pilots to demonstrate feasibility of estimating a user's feedforward and feedback controllers from data. More recent work focuses on quantifying user performance using the estimated feedforward controller. Researchers demonstrated that users without motor impairments using a manual interface develop good feedforward controllers for predictable [40, 137, 192] and unpredictable [182, 189] trajectories. In addition, researchers also demonstrated that users' feedforward controllers improve as users gain experience performing a trajectory tracking task [192].

Our study extends the control theory-based experimental methods and analyses previously used to study how users without motor impairments use manual interfaces. Our paper focuses on how users with and without motor impairments use muscle interfaces. Understanding how feedforward and feedback controllers are affected by alternative interfaces and motor impairments is crucial for improving device interaction for all users.

3.3 What is a Continuous Task?

Continuous tasks can range from simple to more complex. We define simple tasks as being position- (eg. mouse tracking, where the position of the mouse determines the cursor position) or velocity-based (eg. wheelchair navigation, where the joystick position determines the velocity of the wheelchair). We define complex tasks as being acceleration-based (eg. automobile or robot control, where the user input determines the acceleration of the mechanical system). Mathematically, the increase in task complexity arises from the increased

number of derivatives that relate the user input to the device output. These complex tasks require more abstraction (derivatives) for the user to determine the input they should apply to produce the desired device output. In continuous tasks, the user input is theorized to be a combination of i) *user intent* (feedforward control; the input that yields the desired device output in the absence of any errors) and (ii) *error correction* (feedback control; the input that corrects for errors that can arise from unexpected perturbations, inappropriate inputs (eg. due to motor impairments), or unexpected changes in the task) [57]. In this paper, *controller* or *control* refer to the process by which the user determines their input in response to device output. Mathematically, a controller is a function that transforms time- and/or frequency-domain signals. Taking the example of mouse tracking as a continuous task, user intent could express the user's desire to move the cursor along a specific path, while error correction could compensate for deviations from the intended path caused by unintentional tremors of the user's arm (Fig. 3.1).

Neurologic injuries like stroke or cerebral palsy that result in motor impairments usually do not affect the cerebellum, where user intent (feedforward control) is formed [18]. Instead, the injury usually occurs in the motor cortex that coordinates and transmits signals to arm muscles [63]. The injury to the motor cortex can cause errors between the user's intended motion and the implementation of the desired motion (eg. causing unintentional arm tremor), making it difficult to perform error correction (feedback control). Thus, we hypothesize that neurologic injury may impair feedback but not feedforward elements of user input. To test this hypothesis, we separately quantify user intent and error correction in continuous tasks for users with and without motor impairments using frequency-domain techniques from control theory. These techniques have previously been applied to participants without motor impairments using manual interfaces [105, 182, 189, 192]. This study extends the applicability of these tools to include simple and complex tasks using muscle interfaces and users with motor impairments after neurologic injury (post-stroke).

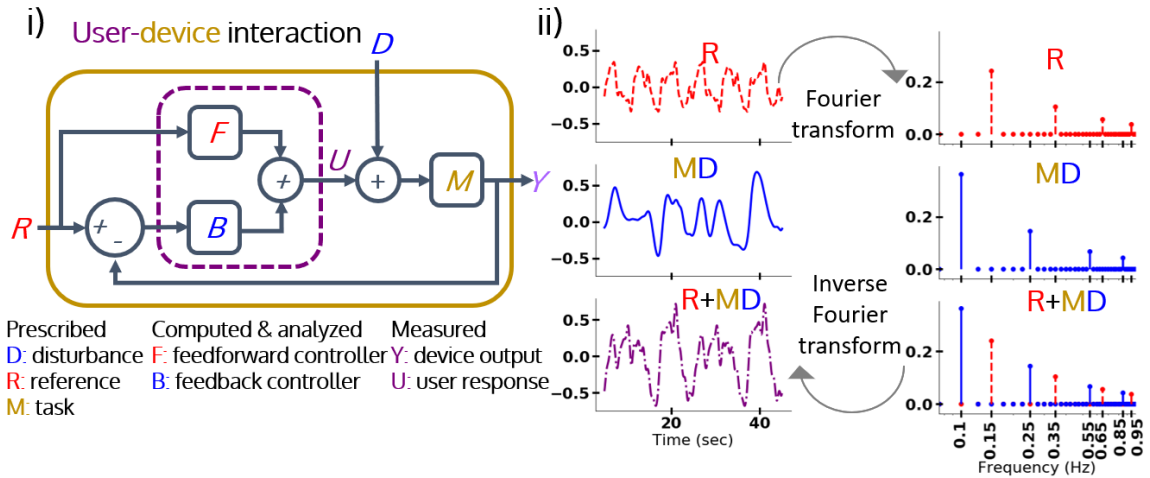


Figure 3.2: (i) Block diagram representation of user interacting with device adapted from [105]. The user, contained within the purple dotted square, transforms external reference R and tracking error $R - Y$ through feedforward (user intent, F in red) and feedback (error correction, B in blue) controllers to produce user input U . The device transforms the sum of user input U and external disturbance D to device output Y via mapping M . (ii) While it is difficult to separate the two reference (dotted red line) and disturbance (solid blue line) signals in the time-domain (left), the task is much easier in the frequency-domain (right). We can easily go back and forth from time- and frequency-domain using the Fourier and inverse Fourier transform.

3.4 Methods

3.4.1 Experimental Design

We manipulated the following conditions: *interface* (muscle versus manual), *task* (simple versus complex), and *population* (with versus without motor impairments). We conducted two types of experiments. First, we ran a 2×2 factorial design study with 11 participants without motor impairments (*interface*: muscle versus manual, *task*: simple versus complex). Second, we conducted a case series study with 3 participants with motor impairments after stroke, and compared muscle versus manual interface performance for the complex task.

We compared the results from this experiment against the participants without motor impairments. To shorten the study and avoid fatigue among participants who had a stroke, data was only collected for the complex task for participants with motor impairments. The order of presentation for the conditions was randomized for each participant.

3.4.2 Participants

We recruited 11 participants without motor impairments for this study from the broader community (4 female, 7 male; 1 left-handed, 10 right handed; age: 25 ± 3.7 years, height: 171 ± 11.3 cm; weight: 68 ± 10.5 kg). All were daily computer users and played video games monthly or yearly. Six participants were familiar with the concept of EMG signals, and one participant regularly worked with EMG signals.

We also recruited 3 participants who had a stroke that affected one side of their body from clinics and local stroke survivor support groups (Table 5.1). P1 and P2 predominantly used their unaffected arm for activities of daily living, including using a computer or phone. As shown by the self-reported impairments, P3 had fairly good control over her affected arm, and used her affected side for mouse navigation and writing. However, she only uses her affected side to use the mouse, and types solely with her non-affected side. Potential participants were asked if they could touch their shoulder and move their arm back as a measure of bicep and tricep control.

3.4.3 Task

Participants used their muscles or a slider to control a cursor on a screen to track a yellow trajectory (Fig. 3.3). Since one goal of this work is to encourage bilateral device interaction, participants without motor impairments used their non-dominant arm and participants with motor impairments used their affected arm to complete the tasks. When using their muscles, participants were strapped into a padded rigid device with their palms facing up (Fig. 3.3 bottom). Participants moved the cursor up by pulling up against the rigid device to activate the biceps, and moved the cursor down by pushing down into the rigid device to activate the triceps. We previously found during a pilot study on participants without

	Age yrs	Sex	Yrs since str- oke	Aff- ec- ted side	Self-reported impairments										
					Mo	Sp	St	Tr	Co	Fa	Gr	Ho	Se	Dr	Ds
P1	48	M	2	L	✓	✓	✓	✓	✓	✓		✓		✓	✓
P2	47	M	11	R	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓
P3	51	F	6	R	✓	✓		✓					✓		

Mo = slow movements, Sp = spasm, St = low strength, Tr = tremor, Co = poor coordination, Fa = rapid fatigue, Gr = difficulty gripping, Ho = difficulty holding, Se = lack of sensation, Dr = difficulty controlling direction, Ds = difficulty controlling distance.

Table 3.1: Participant Characteristics. Self-reported impairments adopted from Findlater et. al [49] and Mott et. al [109]

motor impairments that participants moved the slider in many different ways from flicking the slider to using their whole arm to move the slider. To standardize how participants moved the slider, participants were asked to lay their elbow on a hard surface and move the slider with their biceps and triceps (Fig. 3.3 top).

The user input was mapped to the output of the device as either the velocity (simple task) or the acceleration (complex task) of the cursor. Participants without motor impairments performed 30 trials per condition, each 45 seconds long. At the end of each trial the error between the reference trajectory and the cursor position was displayed as a scaled number between 0 and 100%. Participants were asked to make this number as small as possible. Participants with and without motor impairments were highly encouraged to take breaks between trials and between conditions and were reminded that they were free to stop the experiment at any time. Reset screens where participants could take breaks were shown between each 45 second trial. Participants with motor impairments performed at least 20




Figure 3.3: Participants controlled a purple cursor on a computer screen using either a manual slider (top) or muscle EMG (bottom) interface.

trials per condition, depending on fatigue. Continued clonus or spasticity (more than once per 45-second trial) was also used to indicate muscle fatigue as a break or stop point during the experiment.

After each condition, participants filled out the NASA Task Load Index (TLI) [65] to subjectively quantify the difficulty of completing the trajectory tracking task across six different categories – mental demand, physical demand, temporal demand, performance, effort, and frustration. The NASA TLI rates the workload of a task from 0 (low workload) to 100 (high workload). At the end of the experiment, we asked participants whether they preferred the muscle or slider interface.

3.4.4 *Game Development*

The experiment was described to participants as a trajectory tracking game (Fig. 3.3). Participants were asked to control a purple diamond cursor on the screen using a slider

or their muscles. The cursor was restricted to motion in one-dimension (up or down). Participants controlled the cursor by either manipulating a manual interface (slider) towards or away from the body, or activating the muscle interface by pulling up or pushing down against a rigid device.

The trajectory tracking task was visualized using pygame 1.9.4 in Python3.5. We used a randomly phase-shifted sum-of-sines at eight fixed frequencies between 0.1-0.95 Hz and amplitude to generate pseudorandom references and disturbances (Fig. 3.2(ii)). Researchers previously found that frequencies much higher than 1 Hz are difficult to track in the context of this experiment [105, 182, 189]. The position of the cursor on the screen was updated by the user input at 60 Hz, the same update frequency as a standard computer screen. This game was adapted from work by [105] and more recently by [182, 189].

3.4.5 Muscle Interface Development

We used the Delsys Trigno EMG System (Delsys Inc. Massachusetts, USA) to collect EMG activity from the biceps and triceps of our participants. The Delsys sensor is a wireless dry electrode commonly used in clinical settings, and collects EMG data at 1926 Hz. The electrodes were placed on the biceps and triceps according to Surface Electromyography for the Non-Invasive Assessment of Muscle (SENIAM) [67] guidelines. The Delsys software development kit was used to import raw EMG signals from the Delsys unit to Python for further processing.

EMG values were normalized by calibrating the EMG activity against participants' maximum expected contraction. At the beginning of the trial, we asked each participant to flex their biceps or triceps as hard as they could three times, each for two seconds, while secured by the rigid device or by a researcher. The 95th percentile of the EMG data collected was saved for each 2-second trial, and the average of the three trials was saved as the maximum contraction.

Raw EMG signals were processed similarly to [31, 88]. EMG signals were filtered by processing 100 ms of EMG data at a time. Each 100 ms window was further split up into two, 50 ms windows and delinearized before taking the average of the two windows. We

then scaled the filtered EMG activity by the value of the maximum contraction for each muscle. If scaled user input for both the biceps and triceps were below a specified threshold (defined as 2.5% of the maximum contraction for participants without motor impairments), the user input was set to zero. This ensured that participants could reach zero despite minor fluctuations in EMG signal from measurement noise. Otherwise, the muscle with the larger scaled value was returned as the user input. If the biceps had a larger scaled value than the triceps, the cursor would move up, and if the triceps had a larger scaled value than the biceps, the cursor would move down.

All of the participants who had a stroke could not sufficiently relax their muscles to obtain a zero user input with the 2.5% threshold due to weaker maximum contractions. The threshold for zero user input for the muscle interface was adapted to a maximum of 12% of the maximum contraction, depending on the level of EMG activity we observed during rest.

3.4.6 *Slider Interface Development*

Participants manipulated a custom slider connected to a 10 $k\Omega$ potentiometer. An Arduino Due (Arduino.cc) was used to measure and import the potentiometer values into Python for further processing. The slider was 35 mm wide \times 12 mm tall \times 22 mm deep and printed with a 3D printer using ABS filament. Pushing the slider required very little strength similar to pushing a pen across a table.

3.4.7 *Data Analysis*

User input from either the muscle or manual interface, reference and disturbance trajectories, and position of the cursor on the screen was collected at 60 Hz. Collected data was analyzed in Python3.5. To quantify user performance taking into account both user intent and error correction, we compute the mean-square error (MSE) between the prescribed reference R and the measured position of the cursor Y over time t :

$$MSE_{time} = \sum_t (|R - Y|)^2. \quad (3.1)$$

To quantify user performance solely taking intent into account and ignoring difficulties with error correction arising from motor impairments, we compute the MSE between the inverse of the device dynamics M^{-1} and the estimated feedforward controller F over frequencies w :

$$MSE_{freq} = \sum_w (|M^{-1} - F|)^2. \quad (3.2)$$

The performance for the last five trials was averaged as a measure of error after learning for both performance metrics.

For the 2×2 factorial design with participants without motor impairments, we looked for potential differences between conditions (*interface, task*) for the two performance metrics defined in eq. (3.1, 3.2) using the two-way analysis of variance (ANOVA) test. We tested the normality distribution assumption of our data using the Shapiro-Wilks test and allowed for minor violations in normality because of the robustness of the ANOVA. We hypothesized that participants will perform worse when performing the complex task compared to the simple task due to the added abstraction (derivative). Additionally, we hypothesized that muscle interfaces will perform worse than manual interfaces because participants will be more acquainted with manual interfaces than the muscle interface. Paired t-tests with $\alpha = 0.05$ were used as a post-hoc test.

Similarly to previous studies [31, 88], we also hypothesized that we will see performance differences between muscle and manual interfaces at higher frequencies. Although researchers previously only compared muscle and manual interfaces for the simple task, we hypothesized that their finding will extend to the complex task as well. We tested for differences in frequency-domain performance at each frequency between the muscle and manual interface with the paired t-test for the complex task with $\alpha = 0.05$.

As we only had three participants with motor impairments, comparisons between users with and without motor impairments are descriptive. This experiment was mainly to assess the viability of a muscle interface for users without motor impairments. We hypothesized that users with motor impairments will perform worse than users without motor impairments with the time-domain performance metric, but perform similarly for the frequency-domain performance metric. This is because motor impairments after neurologic injury usually affect the error correction (feedback), not user intent (feedforward) contributions to

user input [63].

3.5 Results

3.5.1 Study 1: Muscle versus Manual Interfaces for Simple and Complex Tasks

To determine when muscle interfaces provide performance advantages over manual interfaces, this study compared performance for users without motor impairments using muscle and manual interfaces for simple and complex continuous tasks.

Muscle Interface Improves Performance for Complex Task

The time-domain performance metric (MSE_{time}) quantifies the performance of continuous trajectory tracking tasks when taking into account both user intent and error correction. The two-way ANOVA only found a main effect for the task difficulty (Table 3.2). Contrary to our hypothesis that participants will perform better with the manual interface, participants without motor impairments performed equally well with either interface for both the simple and complex tasks (simple: $t=1.88$, $p = 0.09$; complex: $t=-0.15$, $p = 0.88$) (Fig. 3.4). However, participants performed significantly worse for the complex task compared to the simple task using the manual interface ($t=-5.76$, $p < 0.001$) (Fig. 3.4). This suggests that participants did find the increased abstraction (derivative) more difficult but only when using the manual interface.

The frequency-domain performance metric (MSE_{freq}) quantifies the performance of continuous trajectory tracking tasks when only taking into account the user intent. This metric ignores how well or poorly participants perform error correction. We found a significant main (*interface, task*) and interaction (*interface* \times *task*) effect for the frequency-domain performance (Table 3.2). As expected, users developed a better feedforward controller for the simple task compared to the complex task when using the manual interface ($t=-7.81$, $p < 0.001$) (Fig. 3.4). This suggests that participants found it easier to determine the input required to track the desired trajectory for the simple task compared to the complex task when using the manual interface. Participants performed equally well for simple and complex tasks when using the muscle interface ($t=-2.09$, $p = 0.063$). Surprisingly, users had 49%

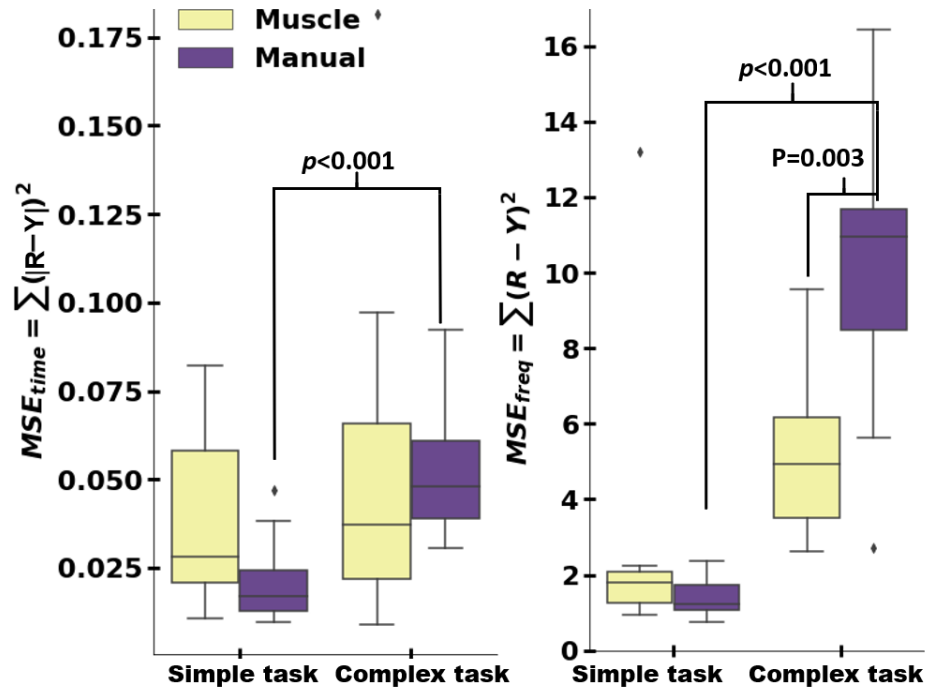


Figure 3.4: Time-domain (left) and frequency-domain (right) measures of error for both simple and complex tasks. Lower values indicate better performance. Statistically significant differences are marked with their respective p values.

more accurate feedforward controllers when performing the complex task with the muscle interface than the manual interface ($t = -4.66$, $p < 0.001$). This means that in the absence of errors users can track reference trajectories more accurately with the muscle interface than the manual interface, but only for the complex acceleration-based task. This suggests that interface performance is task-dependent and feedforward controller accuracy is dependent on the type of interface used.

Comparing the performance of muscle and manual interfaces by solely quantifying user intent (MSE_{freq}) enabled us to detect differences between the two interfaces that were not readily apparent when also taking into account error correction (MSE_{time}). While the two interfaces performed similarly in the time-domain performance metric, the muscle interface performed significantly better than the manual interface in the frequency-domain metric.

Factor	$F_{1,40}$	p	Partial η^2
<i>MSE_{time}</i>			
Interface	1.19	0.28	0.025
Task	6.55	0.014	0.14
Interface \times task	0.74	0.39	0.15
<i>MSE_{freq}</i>			
Interface	4.83	0.034	0.047
Task	43.5	<0.001	0.43
Interface \times task	13.5	<0.001	0.13

Table 3.2: Two way ANOVA (*interface \times task*) results for time-domain (MSE_{time}) and frequency-domain (MSE_{freq}) measures of performance.

Having an accurate prediction of what the user intends to do is critical for developing algorithms that assist the user in performing tasks.

Muscle Interface Accurately Tracks High-Frequency Signals

To more deeply understand why the muscle interface performed better than the manual interface for the complex task in the frequency-domain, we compared the frequency-domain performance at each stimulus frequency. For the complex task, participants performed significantly better at frequencies above 0.35 Hz with the muscle interface than the manual interface (Fig. 3.5). Overall, participants performed 61% better at high frequencies above 0.35 Hz with the muscle than the manual interface. Without accounting for error corrections, users' inputs more accurately tracked faster moving components of the reference trajectory with the muscle than the manual interface. This suggests that if a task requires users to track rapidly changing trajectories like navigating a drone in a forest at a high speed, they may find it easier to do so with the muscle than the manual interface.




Figure 3.5: Frequency-based performance across different frequencies for complex acceleration-based task. Participants performed significantly better with the muscle (yellow) than the manual (purple) interface.

3.5.2 Study 2: Comparing Interface Performance for Users With Motor Impairments

We learned from the first study that users without motor impairments perform better using muscle compared to manual interfaces while conducting complex tasks in the absence of errors. In this study, we compared differences and similarities in performance for users with and without motor impairments as they navigate a complex task with a muscle or manual interface. This study was a proof-of-concept case study with 3 participants (P1, P2, P3) who had a stroke to verify whether muscle interfaces were a viable alternative to manual interfaces for users with motor impairments.

Muscle Interface Improves Performance

Participants with motor impairments successfully completed the complex trajectory tracking task with both muscle and manual interfaces. P1 learned to use the muscle interface

quickly and preferred it over the manual interface. P2 and P3 had more difficulties learning to use the muscle interface and isolating bicep and tricep activation. They expressed that they would have performed better if they had more time to practice. Despite the limited practice time, users with motor impairments performed 24% and 44% better using the muscle than the manual interface with the time- and frequency-domain performance metrics respectively (Fig. 3.6). As we only had three participants, the performance of each participant with motor impairments is shown as a dot. As we hypothesized, users with motor impairments had much worse time-domain performance (MSE_{time}) than users without motor impairments (Fig. 3.6). However, frequency-domain performance (MSE_{freq}) for users with motor impairments were within the range observed for users without motor impairments. Since frequency-domain performance excludes contributions from error correction while time-domain performance accounts for both user intent and error correction, we can conclude that users with motor impairments were worse than users without motor impairments with error correction, but not with forming user intent.

NASA Task Load Index (TLI)

Users with and without motor impairments perceived no differences in task load across tasks and interfaces (Fig. 3.7). We found no significant main (*interface*: muscle versus manual; *task*: simple versus complex) or interaction (*interface* \times *task*) effects from the results of the NASA TLI for participants without motor impairments. NASA TLI for users with motor impairments ranged from 45 to 80, well within the range of users without motor impairments. This suggests that users subjectively found all interfaces equally easy to manipulate, despite the muscle interface being a novel interface for many participants.

3.6 Discussion

We demonstrate for the first time that users with and without motor impairments perform better when using a muscle interface compared to a manual interface for a complex (acceleration-based) continuous task. Muscle interfaces provide an attractive alternative interaction method to manual interfaces that encourage bilateral interaction for users with motor impairments after neurologic injury. The fine coordination of multiple arm and finger




Figure 3.6: Time-domain (left) and frequency-domain (right) measures of error for complex task for users with and without motor impairments. Lower values equals better performance. The time-domain error for users with motor impairments is much worse than users without motor impairments for both the muscle and manual interface. Users with motor impairments perform comparably to users without motor impairments in forming feedforward models in the frequency-domain.

muscles required to use manual interfaces are simplified to activating one or two user-chosen muscles with muscle interfaces. For users with difficulties performing error correction like users with motor impairments, quantifying the intent from the user input while ignoring error correction is an important metric for quantifying interface performance. To the best of our knowledge, we proposed the first performance metric for users with motor impairments that quantifies user intent, that is, the user input needed to control a device to follow a desired trajectory in the absence of errors.

We found that both users with and without motor impairments preferred and performed 44% and 49% better respectively with the muscle than the manual interface for the complex

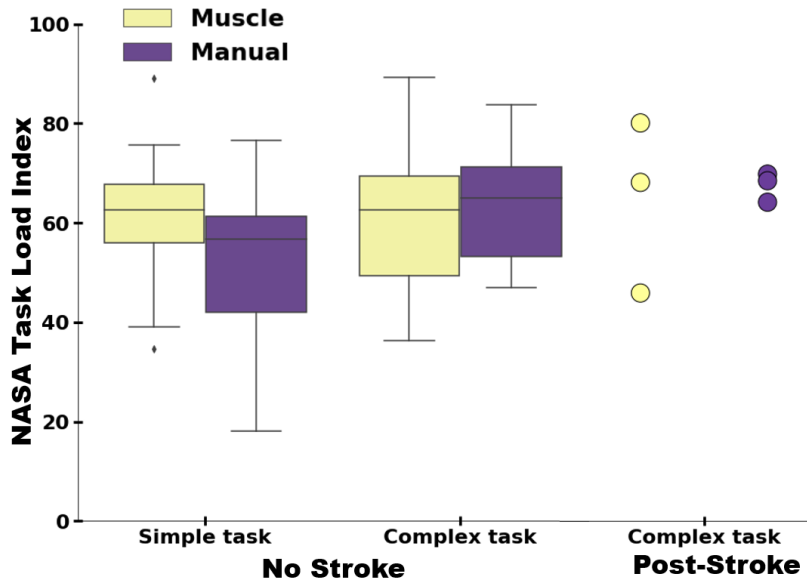


Figure 3.7: Results from NASA TLI demonstrates similar subjective workload across all tasks for users with and without motor impairments.

task but not the simple task. We additionally found that users without motor impairments performed 61% at frequencies above 0.35 Hz. It may be that muscle interfaces are particularly intuitive for acceleration-based tasks. The electrical activity that we measure as the user input for the muscle interface is a result of electrical signals sent from the brain to the muscle fibers, which then produce force that generates movement [34]. Since force (F) is correlated to acceleration (a) and the mass of the system (m) by $F = ma$, EMG activity can be directly mapped to acceleration without the need for abstraction. This is one possible explanation for why muscle interfaces were preferred and performed better for the complex acceleration-based task. Something to consider for future iterations of this study is comparing the muscle interface against a force-based manual interface instead of a position-based manual interface that we used for this study. If direct mapping between the user input and device output is important for performance, then muscle interfaces and force-based interfaces should perform similarly for the complex task.

Muscle interfaces performing better for complex acceleration-based high-frequency tasks

have implications for interface design for users with and without motor impairments. Muscle interfaces may be beneficial for tasks where the user controls the acceleration of the device that require quick maneuvers like flying a drone through a dense forest or remotely controlling a legged robot through rocky terrain. Continuing research on comparing how various interfaces perform with frequency-domain analysis from control theory is useful for informing intuitive interface design for device control.

A key contribution of this study is the use of techniques from control theory to decode user intent for users with motor impairments for the first time. Previous studies solely compared performance between manual and muscle interfaces from users without motor impairments [31, 88]. Our study demonstrated that while users with motor impairments have higher time-domain error (MSE_{time}) than users without motor impairments, they performed better with the muscle interface compared to the manual interface. Users with motor impairments had to coordinate and activate their whole arm and upper body to hold and move the manual interface. With the muscle interface, however, they solely activated their biceps and triceps, requiring less coordination and movement than the manual interface. In the frequency-domain, we successfully derived user intent for users with motor impairments, and showed that the quantified feedforward controllers were within range of users without motor impairments. This is consistent with what is known about motor impairments after stroke, as stroke that results in the impairments often affects the motor and sensory cortices of the brain affecting muscle recruitment and sensory feedback, but not the planning of the movement in cerebellum [63]. These promising results suggest that muscle interfaces may be a viable alternative to manual interfaces to enable bilateral device interaction for users with motor impairments. It is especially exciting that users with motor impairments appear to perform better using the muscle than the manual interface even when the performance metric takes into account error correction. However, we expect that users with motor impairments will perform even better when their intent is used by artificial intelligence to assist them with error correction. In future work the derived user intent could also be used to design algorithms that adapt to user capabilities to assist with error correction.

3.6.1 *Limitations*

With such a small population of participants with motor impairments (3), it is not possible to draw statistically significant conclusions. In addition, motor control ability in users who have had a stroke is diverse, and even in our case study we saw a large heterogeneity in user capabilities, making it challenging to draw general conclusions about users with motor impairments from this study. However, our preliminary results suggest that muscle interfaces are a viable alternative to manual interfaces and modeling methods from control theory can be used to quantify user intent separate from error correction for users with motor impairments.

We only compared the muscle interface against a custom-built slider, one type of manual interface that is not as commonly used in daily life and was not designed for user comfort or performance. Additionally, users were not able to customize how the physical displacement of the slider mapped to the movement on the screen. In the future, comparing the muscle interface against commercially-available manual interfaces like touchscreens, joysticks, and mice and allowing for customization of interface sensitivity would inform when muscle interfaces are a desirable alternative to manual interfaces for complex high-frequency continuous tasks.

Lastly, there were a number of restrictions placed on the participants during this study that would not be in place during everyday use that may have affected the results of the study. To standardize how participants interacted with the manual interface, we asked participants to place their elbow on a hard surface and solely use their biceps and triceps, rather than using their wrists or fingers to manipulate the slider. We also do not know the effect of handedness on muscle or manual interface performance. Participants may have performed better with the manual than the muscle interface if they had used their dominant hand since users generally have better coordination with their dominant hand.

3.7 *Conclusion*

This is the first paper to report on the performance of muscle versus manual interfaces for simple (velocity-based) and complex (acceleration-based) continuous tasks for users with

and without motor impairments. We introduced techniques from control theory to quantify the performance of user intent in the absence of errors (like unintended tremor from motor impairments). Users without motor impairments performed 49% better with the muscle than the manual interface for tasks that required rapid changes to user inputs. Users with motor impairments performed 44% better with the muscle than the manual interface and had similar intent to users without motor impairments, suggesting that the modeling method successfully separated the users' intent from errors arising from motor impairments.

Muscle interfaces provide performance advantages for users with motor impairments and for complex tasks that require users with and without motor impairments. Such alternate interfaces should continue being developed to support users of all abilities.

Chapter 4

**CO-ADAPTATION FOR HUMAN-IN-THE-LOOP CONTROL
SYSTEMS**

To be Submitted

Momona Yamagami

Maneeshika M. Madduri

Benjamin Chasnov

Amber H.Y. Chou

Lauren N. Peterson

Samuel A. Burden

Abstract

In continuous human-machine interfaces such as robot teleoperation, both the human and interface must work together to achieve a common goal. To optimally assist the human user, co-adaptation schemes that enable the prediction of how the human and interface will adapt to stimuli and parameters in closed-loop with fixed dynamics is ideal. However, co-adaptive HMIs present a “two-learner” problem with unique challenges, due to the human and interface jointly learning in closed-loop. In this paper, we demonstrate that the co-adaptation game variation where the human and interface are series interactions with the goal of reference tracking or disturbance feedforward are theoretically ill-posed. We find in simulation and in experiment that minimizing for task error and human and interface effort in a co-adaptive HMI results in improved performance and decreased human effort compared to baseline. Further investigation is needed to identify the optimal parameterization of the human controller and how adaptive interfaces can be best synthesized in closed-loop with an adaptive human to optimally control a fixed machine during reference tracking or disturbance feedforward tasks.

4.1 Introduction

When controlling human-machine interfaces (HMIs) such as prostheses [35, 196], brain-computer interfaces [120, 174], and robots [84, 114, 172], it is often desirable for the human and the interface to jointly control a fixed machine by co-adapting to achieve a common goal or task. In many of these scenarios, the human and interface must control the output of a machine with *dynamics*. To optimally assist the human user, co-adaptation schemes that enable prediction of how the human and interface will adapt to certain stimuli and parameters when in closed-loop with fixed dynamics are ideal.

However, co-adaptive HMIs present a “two-learner” problem with unique challenges, due to the human and the interface jointly learning in closed-loop. Game theory has been proposed in recent years as a framework for studying these two-learner dynamics in sensorimotor control [23, 84, 85, 93, 112] and provides techniques for predicting convergence to and stability of equilibrium points in two-learner systems [19, 131]. Previous studies have demonstrated that for HMIs where the human and interface have the same inputs and outputs (i.e., the human and the interface are *parallel interactions*), closed-form solutions to the two-learner problem can be found with classical techniques [84, 114, 195]. However, the case where the human and the interface have different inputs and outputs (i.e., the human and the interface are *series interactions*) and are both interacting with a fixed machine with dynamics has not been as well studied.

Examples such as invasive [120] or non-invasive [187] neural interfaces, rehabilitation [100], driving [2], prosthetics [196, 35], motor learning [79], and surgical robotics [61] highlight the ubiquity of existing and emerging HMIs that fall within the category where the human and interface are series interactions. To more clearly define this problem category, we can take the example of a neural interface [93]. A common framework for a neural interface is to measure neural signals, such as firing rates, (i.e., the *human*) and use a decoder (i.e., the *interface*) to translate those signals to control an external device, such as a computer cursor to follow a target on a screen [93, 120]. In this scenario, the input to the human is the target position and the output is the neural activity. The input to the interface is the neural activity and the output is the cursor position. Similarly, there are many other application

areas in which we can define the human and the interface as being series interactions, with the input to the interface being the output of the human and the output from the interface potentially interacting with a fixed machine with dynamics.

In the application areas listed above, we aim for the machine output to track references and reject disturbances that occur in the form of additive noise. For example, when controlling a computer cursor via a neural interface, there are references that the user may want to track (e.g., moving a cursor from the center of the screen to a desired location) as well as additive noise that the user may want to reject (e.g., unintended cursor movement due to instability in the neural interface) [187]. In this paper, we demonstrate that the co-adaptation game variation of the human and the interface as series interactions in an HMI with the goal of reference tracking or disturbance feedforward is theoretically ill-posed [195]. Yet, the ubiquity and utility of such two-learner systems suggest that this is an important problem to investigate. Therefore, we focus on the experimental results and computational analysis of co-adaptation games in HMIs with series interactions and disturbance feedforward.

We demonstrate that by minimizing task error and human and interface effort in a co-adaptive HMI, we can synthesize usable, adaptive interface controllers in both simulation and during real-time experimentation. We test the co-adaptation strategies of the human and interface in a disturbance-feedforward experiment with a fixed 2nd order machine and adaptive 0th, 1st, and 2nd order interfaces. We find improved performance and decreased human effort in only the 1st order adaptive interface compared to baseline. We do not find significant differences in performance or human effort between baseline and co-adaptation for the 0th or 2nd order adaptive interfaces. Our simulation results demonstrate that performance improves with co-adaptation compared to baseline at low human penalty parameters, but human effort improves with co-adaptation compared to baseline only for the 2nd order human and interface controllers. Combined, our experimental and computational results suggest that co-adaptation does improve performance compared to baseline, and the 2nd order parameterization of the human controller with a low human penalty produces similar computational results as our experiments. Further investigation is needed to identify the optimal parameterization of the human controller and how adaptive interfaces can be best synthesized in closed-loop with an adaptive human to optimally control a fixed machine

during reference tracking or disturbance feedforward tasks.

4.2 The co-adaptation game in human-machine interfaces

For a generalized co-adaptation game, we can model the two-learner HMI using the block diagrams in Fig. 4.1, where: H represents the *human*, M the (fixed) *machine* that is being controlled, and I the (adaptive) *interface* we seek to synthesize. The analysis and synthesis of feedback systems like these is the *raison d'être* for the field of *control theory* [12, 103]. In particular, a sub-field termed *robust control* [195] is concerned with the synthesis of the interface I (conventionally termed the *controller*) to optimize a performance criterion with respect to a model of the machine M and human H (and, optionally, uncertainty in the given models). The performance criteria of interest in robust control are (*induced*) *norms* [195, Ch. 4] that quantify how much signal power is transferred from disturbance w to error z . From the perspective of synthesizing the interface I , the optimal control problem takes the form¹

$$\min_I \|H/M/I\|_I \quad (4.1)$$

where the norm $\|\cdot\|_I$ is determined from components of z that involve y or u .

In a classical robust control framework, the controller to be synthesized is the only unknown or nonstationary transformation – once the optimal controller is found (e.g. using efficient numerical algorithms [21]), it can be implemented with confidence that it will extremize the chosen performance criterion. The problem in (4.1) is therefore non-classical due to the presence of the human H , who will undoubtedly adapt [105] in response to perceived changes in the interface I , for instance by solving its own optimal control problem [39, 146, 157],

$$\min_H \|H/M/I\|_H, \quad (4.2)$$

where the norm $\|\cdot\|_H$ is determined from components of z that involve p or q . If the interface adapts to observed changes in the human, for instance by iteratively (re-)solving its optimal control problem (4.1), it is natural to regard this co-adaptive interaction as a *game* [169]

¹As an example, when the disturbance w is Gaussian and the error z consists of the state of the machine x and the control input u , solving (4.1) using the frequency-domain 2-norm as the performance criterion yields the well-known *linear-quadratic Gaussian* (LQG) regulator [195, Ch. 14].




Figure 4.1: *Block diagram models for HMIs.* These diagrams specify the flow of information, with signals illustrated by *arrows* and transformations of signals illustrated by *blocks*. We allow signals to be multidimensional, and blocks to be *multi-input multi-output* [195, Ch. 1]. For instance, M has three inputs (q , w , u) and outputs (p , z , y). We enumerate the inputs and outputs from top-to-bottom in the diagram, so q is the first input to M and y is the third output from M , and we use subscripts to denote transformations between (groups of) signals¹. When the blocks H , M , and I are linear time-invariant [195, Ch. 3], these diagrams are not solely conceptual – they provide precise mathematical specifications of the closed-loop transformation from input disturbance w to output error z . Indeed, the *left* diagram is equivalent to the two diagrams in the *middle* where the block M/I is determined by the feedback interconnection (formally, the *lower linear fractional transformation* [195, Ch. 10.1]) between M and I , and the block H/M is determined by the feedback interconnection (formally, the *upper linear fractional transformation* [195, Ch. 10.1]) between M and H . Finally, all of these are equivalent to the *right* diagram that is obtained by applying the same interconnection algebra to the feedback loops that remain in the *middle* diagrams to obtain the transformation $H/M/I$ from input disturbance w to output error z .

¹e.g. $M_{y,q} = M_{3,1}$ denotes the transformation that M defines from q to y and $M_{(z,y),q} = M_{(2,3),1}$ denotes the transformation that M defines from q to (z, y) ; the order of subscripts is chosen to match the position of the inputs and outputs in the equation $y = M_{y,q} \cdot q$.

– specifically a *dynamic game* [19], owing to the presence of dynamics in the machine M (and, presumably [105], the human H and interface I).

We now discuss several variants of this co-adaptation game. The dynamic non-cooperative game defined by the optimal control problems in (4.1) and (4.2) has been studied extensively in the case that the human H and interface I share input and output spaces [19, Ch. 6] – this *parallel* interaction has applications in collaborative and rehabilitation robotics [84, 114]. Motivated by applications in teleoperated robots, neuroprosthetics, and other assistive devices where the human’s input or output are processed through the interface, we consider the less well-studied case of *series* interactions. It is natural in our applications of interest that the input disturbance w contains a *reference* r we want the machine’s output y to track, or additive *noise* d we want the machine’s output to reject, so that the output error z contains tracking error $e = r - (y + d)$. However, including these types of *disturbance feedforward* in the output errors renders the optimal control problems in (4.1) and (4.2) theoretically ill-posed [195, Ch. 14.7]. Thus we focus in the remainder of this paper on experimental and computational studies of co-adaptation games in HMIs with series interactions and disturbance feedforward.

4.2.1 Modeling the HMI with series interaction and disturbance feedforward

We focus the remainder of this paper on the variant of the co-adaptation game with series interaction and disturbance feedforward. We model the adaptive human and interface as two players that interact in a closed-loop *dynamic game* [19] where the human provides a response with a one-degree-of-freedom slider $u_H(t) \in \mathbb{R}$ that the interface then transforms into the interface response $u_I(t) \in \mathbb{R}$ to reject disturbances $d(t) \in \mathbb{R}$ (Fig. 4.2).

We previously demonstrated that when humans are tasked with tracking references r and rejecting additive disturbances d through a linear time-invariant (LTI) [12, Ch. 3, pg. 4] system M , humans behave approximately like LTI transformations for a range of reference and disturbance signals [181, 186]. As such, we can analyze our system in Figure 4.2 using the frequency-domain representations [126, Ch. 5] of signals and LTI systems; we will adorn signal x and transformation T with a “hat” $\hat{\cdot}$ to denote the Fourier transform \hat{x} , \hat{T} .




Figure 4.2: *HMI model with series interaction and disturbance feedforward.* Left: When reference $r = 0$, the adaptive human H transforms output y to user response u_H ; the adaptive interface I transforms the user response u_H to interface response u_I ; the fixed machine M transforms the sum of interface response u_I and disturbance d to output y . Right: Human response u_H is obtained with one-dimensional manual slider and is input to the adaptive interface I and fixed machine M to produce output y .

Therefore, for a given prescribed and measured signals and transformations $\hat{r}, \hat{d}, \hat{y}, \hat{I}, \hat{M}$, we can apply block diagram algebra [12, Sec. 2.2] to transcribe Fig. 4.2 into equations that can be manipulated to express the empirical and prescribed transfer functions $\hat{T}_{u_H d} = \frac{\hat{u}_H}{\hat{d}}$ as a function of the unknown human transfer function $\hat{H}(\omega)$:

$$\hat{u}_H(\omega) = \frac{-\hat{H}(\omega)\hat{M}(\omega)}{\underbrace{1 + \hat{H}(\omega)\hat{M}(\omega)\hat{I}(\omega)}_{\hat{T}_{u_H d}(\omega)}} \hat{d}(\omega). \quad (4.3)$$

We can then estimate the human's controller $\hat{H}(\omega)$ at each stimulated frequency ω as:

$$\hat{H}(\omega) = -\hat{M}^{-1}(\omega) \frac{\hat{T}_{u_H d}(\omega)}{1 + \hat{I}(\omega)\hat{T}_{u_H d}(\omega)} \quad (4.4)$$

We can additionally apply block diagram algebra to obtain the user controlled cursor position \hat{y} as a function of prescribed and measured signals and transfer functions:

$$\hat{y}(\omega) = \frac{\hat{M}(\omega)}{1 + \hat{H}(\omega)\hat{M}(\omega)\hat{I}(\omega)} \hat{d}(\omega) \quad (4.5)$$

For a given sinusoid of frequency ω , the adaptive interface can be synthesized as minimizing the task error $\|\hat{y}(\omega)\|_2^2$ and effort $\|\hat{I}(\omega)\|_2^2$. Assuming that the human is also minimizing the task error and effort $\|\hat{H}(\omega)\|_2^2$, we can assume that the human and the interface are

playing a *potential game* [69, Ch. 12] and define a corresponding potential function:

$$\hat{p}(\hat{H}(\omega), \hat{I}(\omega)) = \|\hat{y}(\omega)\|_2^2 + \lambda_H \|\hat{H}(\omega)\|_2^2 + \lambda_I \|\hat{I}(\omega)\|_2^2, \quad (4.6)$$

where $\lambda_H, \lambda_I > 0$ corresponds to penalty parameters of the human and interface effort, respectively. Therefore, for a given human controller \hat{H} , we can synthesize the interface controller \hat{I} by minimizing the potential function (4.6), where the user-controlled cursor position $\hat{y}(\omega)$ is defined by (4.5).

4.3 Experimental results

The goal of this experiment was to determine the effects of human-interface co-adaptation on final interface dynamics, performance, and human and interface effort as participants completed a disturbance-rejection task with a one-degree-of-freedom slider. We recruited eleven participants (age: 28 ± 7 years (mean \pm standard deviation); gender: 8 women, 3 men, 1 non-binary (some identified with multiple genders); hand dominance: 11 right-hand dominant). All participants provided consent according to the University of Washington, Seattle’s Institutional Review Board (IRB #00000909). All were daily computer users.

Participants completed three conditions of disturbance-rejection tasks in random order: 1) 0th order interface controller; 2) 1st order interface controller; 3) 2nd order interface controller. Each condition started with a randomized interface initialization, and the interface was updated every three trials by minimizing for the human effort, interface effort, and the joint performance of the human and the interface to make the cursor position magnitude as close to 0 as possible (4.6). Participants completed 21 trials per condition. Between each condition, participants completed three baseline trials where the interface was set to 1 (passthrough). Participants were asked to keep the randomly disturbed cursor as close to the center of the screen as possible.

4.3.1 Participants converge to different interface controllers

For all tested interface controllers, participants converged to a range of interface controller parameters after co-adaptation that were not the baseline (Fig. 4.3).




Figure 4.3: *Individualized interface controllers after co-adaptation.* Participants converged to a range of interfaces after co-adaptation. The black point represents the baseline interface (i.e., $I = 1$) and the blue points represent each participants' final interface coefficients after co-adaptation.

4.3.2 Co-adaptation improved performance for 1st order interface

Compared to baseline, participants improved or maintained performance with co-adaptation (Fig. 4.4). For cursor position output in the time-domain, regardless of interface order, there was no difference between baseline and co-adaptation ($p > 0.05$). For the frequency-domain performance metrics, we solely computed performance for stimulated frequencies below crossover frequency of 0.25 Hz^2 (i.e., we solely took into consideration performance at $\omega = 0.10, 0.15 \text{ Hz}$). For the 1st order interface, participants performed better for all frequency-domain performance metrics with interface co-adaptation compared to baseline (cursor output error $\sum_{\omega} |\hat{y}(\omega)|^2$: $p = 0.03$; potential function loss $\sum_{\omega} |\hat{p}(\omega)|^2$: $p = 0.04$; disturbance rejection performance $\sum_{\omega} \hat{T}_{yd}(\omega)$: $p = 0.04$). However, for the 0th order interface, there was no difference in any of the frequency-domain performance metrics between co-adaptation and baseline ($p > 0.05$). For the 2nd order interface, only the potential function loss $\sum_{\omega} |\hat{p}(\omega)|^2$ was significantly lower with interface co-adaptation compared to baseline ($p = 0.02$); there was no difference with other performance metrics.

²frequency at which the open-loop transfer function magnitude is below 1, $|\hat{L}(\omega)| = |\hat{H}(\omega)\hat{M}(\omega)\hat{I}(\omega)|$ [105]

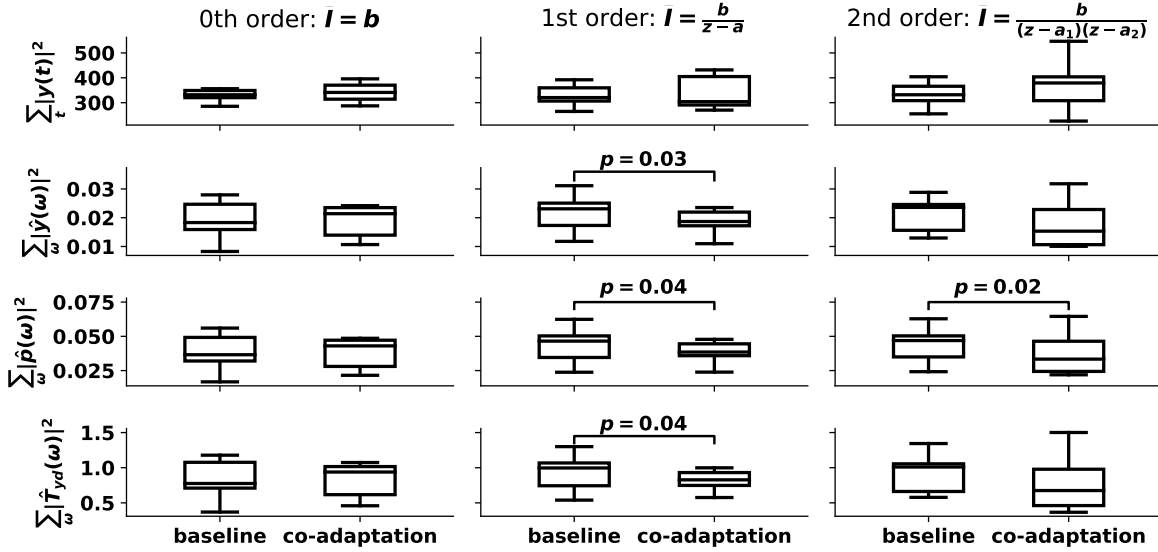


Figure 4.4: *Baseline versus co-adaptation performance.* Distributions (median, interquartile) for baseline and co-adaptation for various performance metrics for 0th (left), 1st (middle), and 2nd (right) order interface dynamics. Top to bottom: time-domain cursor output error $\sum_t |y(t)|^2$; frequency-domain cursor output error for stimulated frequencies below crossover $\sum_{\omega} |\hat{y}(\omega)|^2$; potential function loss below crossover $\sum_{\omega} |\hat{p}(\omega)|^2$; disturbance rejection performance below crossover $\sum_{\omega} \hat{T}_{yd}(\omega)$. Statistical significance ($p < 0.05$) computed from Wilcoxon signed-rank test is indicated by the bracket.

4.3.3 Co-adaptation decreased human effort for 1st order interface

We computed various time- and frequency-domain measures of human and interface effort, including time-domain user response magnitude, frequency-domain user response magnitude at stimulated and non-stimulated frequencies, human controller magnitude, and interface controller magnitude. All frequency-domain measures were computed solely below crossover. For the 1st order interface, all measures of human effort was lower with interface co-adaptation compared to the baseline ($p < 0.05$; Fig. 4.5). However, for the 0th and 2nd order interface, there was no significant difference in human effort between co-adaptation and baseline ($p > 0.05$). All three interface controllers resulted in the interface effort being

higher with interface co-adaptation compared to baseline ($p < 0.05$).




Figure 4.5: *Human and interface effort*. Distributions (median, interquartile) for baseline and co-adaptation for various human and interface effort metrics for 0th (left), 1st (middle), and 2nd (right) order interface dynamics. Top to bottom: time-domain user response $\sum_t |u_H(t)|^2$; frequency-domain user response for stimulated frequencies below crossover $\sum_{\omega} |\hat{u}_H(\omega)|^2$; frequency-domain user response for non-stimulated frequencies below crossover $\sum_{\omega_{non-stim}} |\hat{u}_H(\omega)|^2$; human controller magnitude below crossover $\sum_{\omega} |\hat{H}(\omega)|^2$; interface controller magnitude below crossover $\sum_{\omega} |\hat{I}(\omega)|^2$. Statistical significance ($p < 0.05$) computed from Wilcoxon signed-rank test is indicated by the bracket.

4.4 Simulation results

The goal of the simulation was to establish simulation parameters that approximate the experimental results, towards developing a predictive model of co-adaptive HMIs. We wanted to determine whether it was possible to obtain similar results as the experiment with a

simple co-adaptation scheme.

To simulate our experimental results, we assumed that the human model was parameterized similarly to the interface (i.e., 0th order human controller when testing 0th order interface controller, 1st order human controller when testing 1st order interface controller, and 2nd order human controller when testing 2nd order interface controller). We randomly initialized the human and interface controllers within the parameter bounds tested experimentally, and found the human controller with the lowest potential function value for the randomized interface controller. We then held the human controller constant and found the interface with the lowest potential function value, and repeated the alternating optimization until convergence. We repeated the random initialization and convergence 100 times for various human penalties to determine whether initialization affected convergence. We held the interface penalty at the same penalty parameter used in the experiment ($\lambda_I = 10^{-4}$). We additionally performed the same optimization but solely for the human controller, holding the interface constant at $\hat{T} = 1$ to obtain how the human controller adapts for a baseline interface.

4.4.1 *Human and interface co-adaptation convergence is order-dependent*

Regardless of the human controller order, as the human penalty term decreased, the human gain increased, ranging from the bounded minimum of $b = 0.2$ at the highest simulated human penalty $\lambda_H = 1e - 01$ to the bounded maximum of $b = 7.0$ at the lowest simulated human penalty $\lambda_H = 1e - 08$ (Fig. 4.6). The 0th and 1st order human controllers converged to the same gain regardless of initialization, but the 2nd order human controller gain was dependent on initialization location. The interface convergence was order-dependent. For 0th and 1st order interfaces, higher human penalties resulted in lower interface gain, and lower human penalties resulted in higher interface gain. However, for the 2nd order interface, the final interface gain was dependent on initialization location, and there was not a clear dependence of gain value on the human penalty. We found that this may be due to the increased degrees-of-freedom provided by the poles that enables the 2nd order interface to have multiple stable points for a given human penalty.

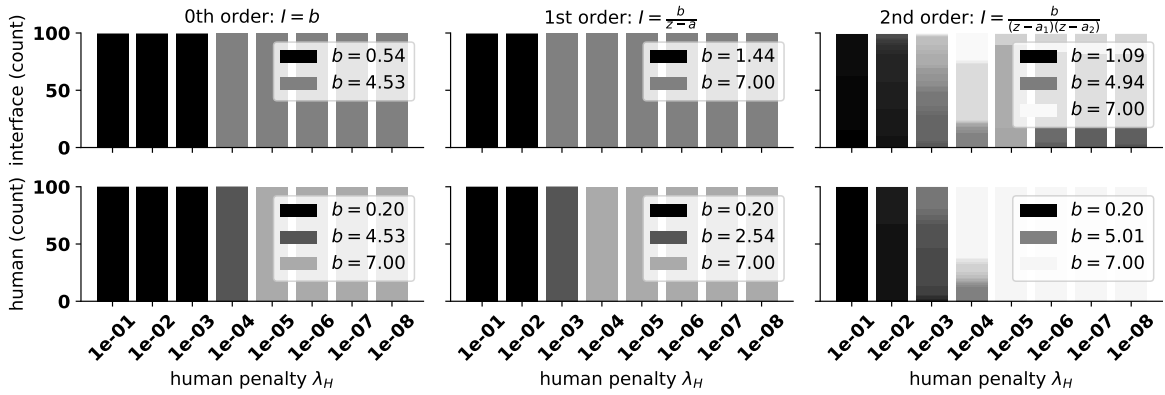


Figure 4.6: *Final interface and human gain for 0th, 1st, and 2nd order interface dynamics.* Count for final interface (top) and human (bottom) gain for 0th (left), 1st (middle), and 2nd (right) order interface dynamics. Darker shade of grey represents lower gains; lighter shades represents higher gains. The simulation was bounded with a minimum gain of 0.2 and a maximum gain of 7.0.

4.4.2 HMI co-adaptation improves performance at low human penalties

We next simulated the effects of human and interface order and human penalty on performance metrics. For all performance metrics, we took the difference between the computed performance metric with co-adaptation and the baseline, so values less than 0 indicate that co-adaptation resulted in better performance than the baseline. Low human penalties resulted in better performance with co-adaptation compared to baseline across all tested human and interface orders (Fig. 4.7). At high penalty parameters, co-adaptation had similar performance to baseline.

4.4.3 Changes in human and interface effort with co-adaptation is order-dependent

We computed human effort in two ways: 1) with the user response magnitude $|\hat{u}_H(\omega)|$ and 2) with the human controller magnitude $|\hat{H}(\omega)|$. We additionally computed the interface effort with the interface controller magnitude $|\hat{I}(\omega)|$. For the user response magnitude, co-adaptation resulted in similar or increased user response magnitude at high hu-

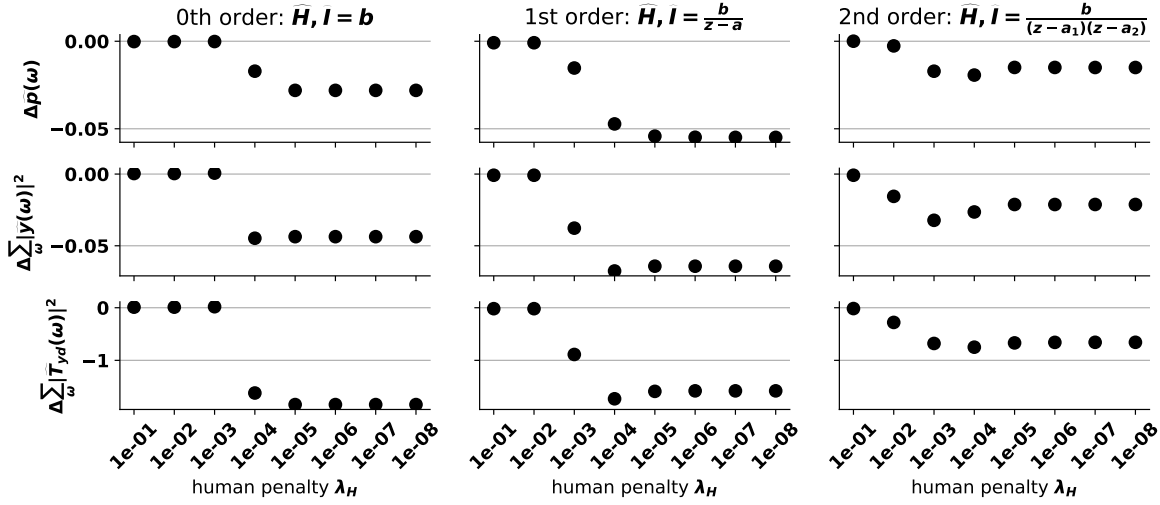


Figure 4.7: HMI performance for 0th, 1st, and 2nd order interface dynamics. HMI performance for 0th (left), 1st (middle), and 2nd (right) order human and interface. From top to bottom: potential function loss $\hat{p}(\omega)$; cursor output $\hat{y}(\omega)$; disturbance rejection performance $\hat{T}_{yd}(\omega)$. All performance metrics represent the difference between the performance with co-adaptation and baseline. Performance difference less than 0 represents better performance with co-adaptation compared to baseline for the given human penalty.

man penalties and decreased user response magnitude at low human penalties compared to baseline (Fig. 4.8). The effects of co-adaptation on human controller magnitude was order-dependent. For 0th and 1st order human and interface controllers, co-adaptation did not affect or increased human controller magnitudes compared to baseline. However, for the 2nd order human and interface controller, co-adaptation increased human controller magnitude at low and high human penalties and decreased human controller magnitudes at medium human penalties compared to baseline. We found a similar order-dependence for the interface controller magnitudes as well. For 0th and 1st order human and interface controllers, co-adaptation decreased interface controller magnitudes at high human penalties and increased interface controller magnitudes at low human penalties compared to baseline. However, for the 2nd order human and interface controllers, co-adaptation increased

interface controller magnitudes regardless of the human penalty.

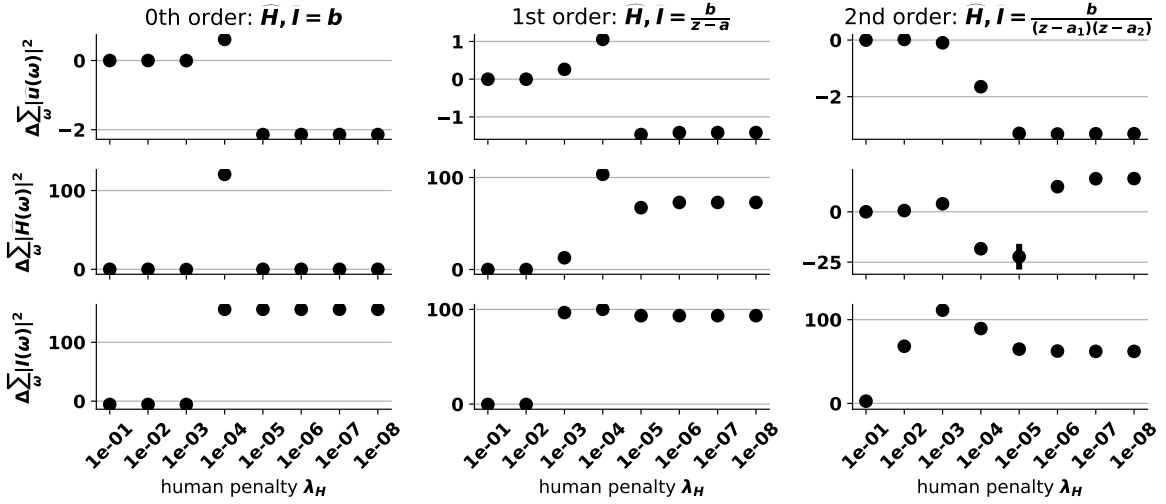


Figure 4.8: *Human and interface effort for 0th, 1st, and 2nd order interface dynamics.* Human and interface effort for 0th (left), 1st (middle), and 2nd (right) order human and interface. From top to bottom: user response $\hat{u}_H(\omega)$; human effort $\hat{H}(\omega)$; interface effort $\hat{I}(\omega)$. All effort metrics represent the difference between the effort with co-adaptation and baseline. Effort difference less than 0 represents lower effort with co-adaptation compared to baseline for the given human penalty.

4.5 Discussion

There are many application areas including prosthetics [35, 196], brain-computer interfaces [174, 120], and robotics [84, 172, 114] where the adaptive human and interface are series interactions and both interacting with a fixed machine with dynamics. However, we demonstrated in this paper that in the scenario of the HMI tracking references and rejecting disturbances, the co-adaptation game is theoretically ill-posed. Previous studies have investigated the use of classical optimal control solutions when the adaptive human and interface are parallel connections [84, 114]. In these scenarios, however, they do not take into consideration tracking references or disturbance feedforward, as the goal is to find the optimal

action that the robot can take to support human-robot teams [114] or support reaching movements using contact robots [84]. Therefore, we performed experimental and computational studies of co-adaptation games in HMIs with series interactions and disturbance feedforward to examine how the human-interface co-adaptation game evolves.

To begin to identify how a HMI where the human and interface are series interactions will adapt, we investigated a simple case where the adaptive interface is updated by minimizing a potential function that accounts for the performance to minimize the cursor output and human and interface effort. The co-adaptation game where the human and interface are series interactions with disturbance feedforward in the output error has not been addressed previously [84, 114] and it was not clear whether or what equilibria the HMI will converge to and the effects of co-adaptation on performance and effort compared to baseline. We found in experiment that co-adaptation improved performance and decreased human effort only for the 1st order interface compared to baseline. We did not observe significant improvements with co-adaptation compared to baseline for the 0th and 2nd order interface compared to baseline. We found in simulation that co-adaptation improved performance at low human penalties compared to baseline for 0th, 1st, and 2nd order human and interface controllers. Co-adaptation decreased human response for all tested controller orders at low human penalties, but decreased human effort only for the 2nd order human and interface controllers compared to baseline. Interface controller effort increased with co-adaptation compared to baseline at low human penalties for all tested controller orders.

Our mixed experimental results as to whether co-adaptation will improve performance is consistent with prior experimental results investigating co-adaptation for a myoelectric interface [196]. The researchers found that shared control may be especially beneficial to users who have limited experience with myoelectric interfaces [196]. In our study, the use of a manual slider interface that can commonly be found in daily life as well as the fixed machine dynamics being second-order non-minimum phase [12, 191] may have contributed to the lack of performance improvement with co-adaptation compared to baseline for the 0th and 2nd order interface dynamics. Non-minimum phase systems, where the inverse of the system dynamics \widehat{M} is unstable are somewhat commonly encountered in everyday life, such as making a turn on a bicycle or driving a car in reverse [12, Ch 14] so our participants may

have had too much familiarity with the task to observe significant performance improvement with co-adaptation compared to baseline. Additionally, previous studies that looked at how humans optimize gait demonstrated that people have difficulty finding optimal gait parameters to minimize energetic cost [179] – participants need guided exploration and experience to find optimal gait parameters [148]. It is possible that our experimental protocol did not provide participants with enough variability in which to explore different human and interface parameters to optimize performance, and a broad exploration of parameters prior to human adaptation and optimization may have improved performance.

It is not immediately evident why we experimentally found improved performance and lower human effort with co-adaptation compared to baseline solely for the 1st order interface controller. One reason could be that the 0th order interface does not provide enough degrees of freedom to significantly affect the performance and human effort, whereas the 2nd order interface was too complex for the HMI to fully co-adapt with the number of trials tested. Considering that the human is tasked with controlling a 4th order system since the interface and fixed machine dynamics are both 2nd order, it is reasonable for participants to require a longer adaptation time than what was provided in this experiment. The trend shown with the potential function loss being lower with co-adaptation compared to baseline for the 2nd order interface, but not the other performance metrics supports this theory as well. In general, while most of the performance and effort metrics for the 0th and 2nd order interface controllers were not statistically significant, we can observe a trend of co-adaptation improving performance and decreasing human effort. This also highlights that the small number of participant was a limitation of our pilot study and it is unclear whether the lack of statistical significance was due to the small number of participants or some other factor. Testing a larger range of parameters (e.g., different fixed machine dynamics, longer co-adaptation time, different interface penalty parameters, different learning rates) with more participants could further elucidate how the different parameters affect performance and effort.

Our simulation results also supports the theory that we may not have provided sufficient learning time with our experimental protocol for the 2nd order interface. The goal of the simulation was to find human parameters that reflect our experimental results, towards the

goal of developing a HMI model that could be used to test parameters in simulation prior to experiment. Our performance metrics results suggest that across the three human and interface controllers tested, having a lower human penalty results on the order of 10^{-4} and below resulted in improved performance and decreased user response with co-adaptation compared to baseline, in line with our experimental results from the 1st order interface controller. However, for the human controller magnitude, only the 2nd order human and interface controllers with human penalty between 10^{-5} to 10^{-4} resulted in decreased human effort with co-adaptation compared to baseline. This suggests that we need at least a 2nd order controller to model the human during HMI co-adaptation. Together, these simulation results suggests that for the experiment that we ran, humans can be modeled as a 2nd order adaptive system with low human penalty on the order of $\lambda_H = 10^{-4}$. Such HMI models are valuable for making qualitative predictions of human-interface co-adaptation [84, 112], especially to rapidly test the effects of different parameters on performance and effort prior to running a user study. Convergence of HMI co-adaptation is not always guaranteed and can be dependent on parameter choices [33, 58], so it is important to ensure that the chosen experimental parameters will lead to convergence of HMI co-adaptation [84, 120].

Lastly, theoretical developments for the case of co-adaptation games in HMIs with series interactions and reference tracking or disturbance feedforward is needed to design optimal interfaces for the HMI applications described in this paper. Such applications are increasingly ubiquitous in our daily lives, and it is important to ensure that the adaptive interfaces lead to stable HMIs that are optimal for the user.

4.6 Methods

4.6.1 Experimental methods

All participants provided consent according to the University of Washington, Seattle’s Institutional Review Board (IRB #00000909). The goal of this experiment was to determine the effects of human and interface co-adaptation on final interface dynamics, performance, and human and interface effort. This was a pilot study to determine whether humans co-adapt to individualized interface controllers or if they co-adapt to the baseline or to a single

equilibrium.

Task

Participants were tasked with controlling a cursor on the screen with a one-dimensional slider. A $35 \times 12 \times 22$ mm (width \times height \times depth) rectangular handle was attached to a linear potentiometer with a 10 cm extent. Similarly to [181, 186], unpredictable disturbance signals were constructed as a sum of sinusoidal signals with the first eight prime multiples of a base frequency of 1/20 Hz ($\Omega = [0.10, 0.15, 0.25, 0.35, 0.55, 0.65, 0.85, 0.95]$ Hz). Each frequency component's magnitude was normalized by the frequency squared to ensure constant signal power, and the phase of each frequency component was randomized in each trial to produce pseudorandom time-domain signals. The disturbances moved the cursor in an unpredictable fashion, and participants were asked to keep the cursor as close to the center of the screen as possible. Each trial was two periods of the periodic stimuli (40 seconds total) following a 5 second ramp-up.

User response u_H were transformed through a fixed non-minimum phase second-order machine M to produce output y to increase the complexity of the task [191]:

$$M : \ddot{y} + 3.6\dot{y} + 4 = 2(\dot{u} + 2.2u) + d, \quad \widehat{M}(s) = \frac{2(s + 2.2)}{s^2 + 3.6s + 4} \quad (4.7)$$

As the updates to the cursor position y occurred at 60 Hz, the machine was discretized prior to implementation in the code.

Experimental conditions

Three conditions (0th, 1st, and 2nd order) of interface controllers were tested (Table 4.1). As the updates to the interface controller response u_I occurred at 60 Hz, discrete dynamics were used to update the interface controller output from one time point to the next.

Each condition started with three trials of a baseline where the user response u_H was unaffected by the interface controller dynamics ($u_I[t] = u_H[t-1]$, $\widehat{I}_{baseline}(\widehat{z}) = 1$), followed by 21 trials of co-adaptation for each condition (~ 30 min per condition). The condition order was randomized for each participant. After all three conditions were completed, the participants performed three more baseline trials. All participants were encouraged to take

Table 4.1: *Tested Conditions for Interface Controller Dynamics.*

	time-domain update	frequency-domain transfer function
0th	$u_I[t] = bu_H[t - 1]$	$\hat{I}_0(\hat{z}) = b$
1st	$u_I[t] = au_I[t - 1] + bu[t - 1]$	$\hat{I}_1(\hat{z}) = \frac{b}{z-a}$
2nd	$u_I[t] = (a_1 + a_2)u_I[t - 1]$ $-a_1a_2u_I[t - 2] + bu[t - 1]$	$\hat{I}_2(\hat{z}) = \frac{b}{(z-a_1)(z-a_2)}$

breaks between the 45-second trials, and participants were asked to take at least a one-minute break after each condition.

Interface controller adaptation

Interface controller updates occurred every three trials. This was to ensure that participants could co-adapt to the new interface controller and was based on results obtained from pilot studies (not shown). Interface controller parameters were restricted such that $0.2 < b < 7$, $-0.95 < a, a_1, a_2 < 0.7$ and initially randomly assigned within those ranges. We initially restricted the poles to have any magnitude less than 1 to ensure a stable interface controller, but found during pilot studies that poles smaller than -0.95 or larger than 0.7 resulted in an uncontrollable interface. We additionally initially did not restrict the 2nd order interface controller poles to be solely real numbers, but found during pilot studies that participants were only converging to real values and so chose to only search over real poles.

The interface controller updates occurred in four steps. First, three trials of the disturbance rejection task was completed by the participant. Next, the human controller \hat{H} was estimated with (4.4) solely using the data from the last two trials. Subsequently, a grid search was conducted between the coefficient ranges noted above for the minimum potential function value (4.6) with penalty terms $\lambda_H = 10^{-6}$, $\lambda_I = 10^{-4}$, and the resulting minimizing interface controller $\hat{I}_{optimal}$ was noted. The grid search was initialized with 100 equidistant points between the ranges noted above. Lastly, to ensure gradual changes between inter-

face controllers from trial to trial and slower changes with increasing trial numbers, *Smooth Batch* [120] was implemented. The next interface controller \widehat{I}^+ was defined as a weighted combination of the previous interface controller \widehat{I}^- and the computed optimal interface controller $\widehat{I}_{optimal}$:

$$\widehat{I}^+ = \alpha \widehat{I}^- + (1 - \alpha) \widehat{I}_{optimal}. \quad (4.8)$$

The parameter α was used to adjust the weighting of the previous interface controller and computed optimal interface controller, and linearly increased from 0 (no previous interface controller) to 1 (solely previous interface controller) as the number of trials increased. This ensured that the interface controller would update rapidly initially, and then update slower as the number of trials increased.

A pseudo-code summarizing the steps of the interface controller update is as follows:

Algorithm 1 Interface update algorithm performed every three trials

Human completes three tasks with $\widehat{I}(\omega) = \widehat{I}^-(\omega)$.

Input: User response $\widehat{u}_H(\omega)$, disturbance $\widehat{d}(\omega)$, previous interface controller $\widehat{I}^-(\omega)$ from the last two disturbance-rejection trials.

Output: Next interface controller $\widehat{I}^+(\omega)$.

for all ω in Ω **do**

- Compute human transfer function (4.4) $\widehat{H}(\omega) \leftarrow -\widehat{M}^{-1}(\omega) \frac{\widehat{T}_{u_H d}(\omega)}{1 + \widehat{I}^-(\omega) \widehat{T}_{u_H d}(\omega)}$.

end for

- Grid search to find optimal interface controller (4.6) $\widehat{I}_{optimal}(\omega) \leftarrow \min_{\widehat{I}(\omega)} \widehat{p}(\widehat{H}(\omega), \widehat{I}(\omega))$.

- Use *Smooth Batch* to update next machine (4.8) $\widehat{I}^+(\omega) = \alpha \widehat{I}^-(\omega) + (1 - \alpha) \widehat{I}_{optimal}(\omega)$.

Human performs three more disturbance rejection tasks.

Data analysis

Time-domain and frequency-domain metrics were computed to analyze performance and effort. For time-domain metrics, the squared sum of the time-series $x(t)$ was computed as: $\sum_t |x(t)|^2$. For frequency-domain metrics, the squared sum of the signal and transfer

function magnitudes $\hat{x}(\omega), \hat{T}(\omega)$ was computed as: $\sum_{\omega} |\hat{x}(\omega)|^2, \sum_{\omega} |\hat{T}(\omega)|^2$. All frequency-domain metrics were solely computed for frequencies below crossover frequency (i.e., at $\omega = [0.10, 0.15]$ Hz). Additionally, some frequency-domain metrics were computed at non-stimulated frequencies below crossover frequency (i.e., at $\omega < 0.25$ Hz). As the participants were shown disturbances at eight *stimulated* frequencies, metrics such as cursor position $\hat{y}(\omega)$ and user response $\hat{u}_H(\omega)$ can be computed at *non-stimulated* frequencies outside of the stimulated frequencies as well.

All metrics were computed by taking the average metric across the last two co-adaptation and baseline trials for each participant such that each participant was an individual data point. Statistical significance was assessed using the Wilcoxon signed-rank test [30, Sec. 5.7].

4.6.2 Simulation methods

We simulated our experimental results by taking the same parameter ranges used in the experiment to develop synthetic human and interface controllers. Because we solely simulated human and interface controllers at stimulated frequencies $\omega \in \Omega$, we did not take into account human response at non-stimulated frequencies. We assumed that the human controller was parameterized similarly to the interface (i.e., 0th order human controller for 0th order interface controller; 1st order human controller for 1st order interface controller; 2nd order human controller for 2nd order interface controller; Table 4.1). To mimic our experiment, interface and human controller parameters were restricted such that $0.2 < b < 7, -0.95 < a, a_1, a_2 < 0.7$ and initially randomly assigned within those ranges. The interface penalty parameter was set at 10^{-4} , same as in the experiment, but as we cannot control the human penalty parameter experimentally, we tested a range of human penalty parameters $10^{-1} > \lambda_H > 10^{-8}$.

Co-adaptation and baseline simulations

For simulating the co-adaptation, we first randomized both the human and interface controller parameters. We then held the interface controller parameters constant, and found the human controller with the lowest potential function value (4.6). We next held the human

controller parameters constant, and found the interface controller with the lowest potential function value (4.6). This was repeated until parameter convergence. We repeated the random initialization and convergence 100 times for the tested human penalty parameters to determine whether initialization affected convergence.

For simulating the baseline, we solely randomized the human controller parameters, and held the interface controller constant at $\hat{I}(\omega) = 1$. We then found the human controller with the lowest potential function value (4.6) for each tested human penalty parameter.

Simulation analysis

Once we obtained the final human and interface controllers for each human penalty parameter, we computed various performance and effort metrics such as potential function loss $\hat{p}(\omega)$ (4.6), cursor output $\hat{y}(\omega)$ (4.5), user response $\hat{u}_H(\omega)$ (4.3), and human controller magnitude $\hat{H}(\omega)$ (4.4). While there was a distribution of final human and interface controller parameters with the co-adaptation simulations, all baseline simulations converged to the same human controller parameters for all human penalties and human and interface order. Therefore, for each performance and effort metric $\sum |x|^2$, we plotted the difference Δ between the mean of the co-adaptation and baseline simulations:

$$\Delta \sum |x|^2 = \sum |x_{co-adaptation}|^2 - \sum |x_{baseline}|^2. \quad (4.9)$$

As lower values indicated improved performance and lower effort for all metrics, $\Delta \sum |x|^2 < 0$ indicated better performance and lower effort with co-adaptation compared to baseline, whereas $\Delta \sum |x|^2 > 0$ indicated worse performance and higher effort with co-adaptation compared to baseline. We indicated the distribution of $\Delta \sum |x|^2$ with error bars indicating one standard deviation from the mean.

Chapter 5

**“I’M JUST OVERWHELMED”: INVESTIGATING PHYSICAL
THERAPY ACCESSIBILITY AND TECHNOLOGY
INTERVENTIONS FOR PEOPLE WITH DISABILITIES AND/OR
CHRONIC CONDITIONS**

Under Review

Momona Yamagami

Kelly Mack

Jennifer Mankoff

Katherine M. Steele

Abstract

Many individuals with disabilities and/or chronic conditions (da/cc) experience symptoms that may require intermittent or ongoing medical care. However, healthcare is an often-overlooked domain for accessibility work, where access needs associated with temporary and long-term disability must be addressed to increase the utility of physical and digital interactions with healthcare workers and spaces. Our work focuses on a specific domain of healthcare often used by individuals with da/ccs: Physical Therapy (PT). Through a twelve-person interview study, we examined how people's access to PT for their da/cc is hampered by social (e.g., physically visiting a PT clinic) and physiological (e.g., chronic pain) barriers, and how technology could improve PT access. In-person PT is often inaccessible to our participants due to lack of transportation and insufficient insurance coverage. As such, many of our participants relied on at-home PT to manage their da/cc symptoms and worked towards PT goals. Participants felt that PT barriers, such as having particularly bad symptoms or feeling short on time, could be addressed with well-designed technology that flexibly adapts to the person's dynamically changing needs while supporting their PT goals. We introduce core design principles (flexibility, movement tracking, community building) and tensions (insurance) to consider when developing technology to support PT access. Rethinking da/cc access to PT from a lens that includes social and physiological barriers presents opportunities to integrate accessibility and flexibility into PT technology.

5.1 Introduction

Chronic conditions are highly prevalent, with six in ten adults in the US having a chronic condition [78]. Many individuals with chronic conditions experience “*uncertain, unpredictable, [and potentially] progressively deteriorating illness*” [15] that may require medical care [128]. Yet, prior work around the inaccessibility of medical care [74, 44, 80, 125] and physician bias towards patients with disabilities [74] raises significant concerns about people with disabilities’ equitable access to healthcare. Physical therapy (PT) is a particularly important domain in which to investigate healthcare accessibility concerns, given its frequent use by individuals with a variety of chronic conditions that may impact access to PT. A better understanding of social and physiological¹ barriers to PT accessibility can inform the design of technology supports, as well as improve our understanding of the PT service ecosystem. In this paper, we examine the inter-relationship of managing embodied illness experiences (eg., physical pain) and fighting socially created, exclusionary accessibility barriers [29] in the context of PT access. We discuss the potential for technology to alleviate these barriers [99] for people with disabilities and/or chronic conditions (da/cc)² and introduce design principles to guide future innovation.

PT is an important health service for acute care such as after an injury (e.g., ankle sprain, back pain), and as part of long-term care strategies for chronic conditions (e.g., Ehlers-Danlos syndrome, fibromyalgia) [11]. During PT, an individual and their physical therapist(s) work together to identify and achieve goals by assessing and implementing an (often movement-based) intervention [52]. These interventions can be prescribed as in-person sessions, where the individual meets with the physical therapist to do exercises together, or as at-home sessions, where the individual is prescribed exercises to do at-home. Adherence to the PT exercise routines is important for achieving the individual’s goals [81, 89, 138], but is often low [53, 106, 167], due to challenges such as not having time for the exercises, daily stress, and lack of social support and guidance [46]. Engagement

¹We refer to the set of factors within a person’s body/mind including physical body parts but also psychological/emotional processes.

²We specify da/cc because many people with chronic conditions do not necessarily identify as having a disability, and vice versa.

in PT exercises is one challenge where design recommendations for technology interventions [101, 7] have been developed, such as collecting at-home compliance and performance data [72] or gamifying PT exercises with commercially available gaming devices (e.g., Nintendo Wii, virtual reality headsets) [6, 156, 55, 83, 102]. However, many of these technology interventions do not take into account the unpredictable, fluctuating symptoms that people may experience from their da/cc and how those symptoms interact with social barriers to influence PT access in complex ways. Moreover, while prior work focuses mostly on either physiological [72] or social [101, 7] barriers to PT, considering the nuanced interaction between physiological and social barriers is important to support PT access for people with da/ccs and inform technology design. Rethinking da/cc access to PT from this more holistic lens presents opportunities to develop technology to improve PT accessibility.

We present an interview study with twelve people with da/ccs in the United States who do PT exercises at-home or with a physical therapist in-person. Our study examines the interaction of disability with their PT to address three questions:

1. What motivates or demotivates people with da/ccs to do PT?
2. What are the social and physiological barriers to PT access for people with a variety of da/ccs?
3. How can technology address these barriers and support people's PT goals?

Our interviews revealed the complexities of having one or multiple da/ccs and participating in PT. We found that da/ccs often served as a motivator to engage in PT, for example, to maintain current physical abilities, relieve chronic symptoms, or achieve a physical feat (e.g., engaging in a sport). At the same time, people with da/ccs sometimes found that PT exercises required modification to be accessible, such as when having a “bad symptoms day”. In addition to physiological barriers, our participants discussed accessibility challenges that reflected socially constructed barriers, including trouble finding transportation to get to and from in-person PT and challenges with insurance. Our findings also revealed the complex interactions between physiological and social barriers to PT, such as feeling too sick on

the day of an appointment to drive to and perform PT. Finally, we present insights from participants about how technology could support at-home PT. Participants were excited for such systems because of the potential to customize daily exercises or converse with a physical therapist outside of the clinic; both are examples of important features when one's abilities fluctuate daily due to a da/cc or if in-person PT is not available to them.

Our contributions include:

1. an assessment of lived experiences of people with da/ccs who are doing PT,
2. identification of opportunities for technology to alleviate social and physiological barriers to PT, and
3. an introduction of design principles from participant technology design suggestions that could improve at-home PT access.

While technology development is only one aspect of improving accessibility to PT for people with da/ccs, our results suggest that there are some PT access barriers that technology is uniquely situated to address, such as dynamically updating exercises or providing opportunities to connect with other people with da/ccs doing PT. Considering the complex interplay between social and physiological barriers to PT presents opportunities to develop technology that holistically supports PT accessibility for people with da/ccs.

5.2 Related Work

There has recently been a shift towards a more nuanced interpretation of disability and impairment, that encourages dismantling social access barriers while acknowledging that for some, their bodies are at the center of their experience with disability [99, 175, 149, 150]. PT sits squarely in this intersection. We identify the post-modern model of disability as a useful framework for understanding the complex interplay between social and physiological barriers. Applying this model to PT and healthcare accessibility highlights still underexamined access barriers and opportunities for technology interventions.

5.2.1 Defining Disability

In examining the PT experience for people with da/ccs, it is first important to examine how we define disability and impairments. Two of the most prevalent models of disability – the social and the medical model – have limitations that do not fully address the experiences of people with disabilities. Although less prevalent, other models of disability highlight different factors core to the disabled experience, such as interdependence [22] or social effects [118]. Notably, the post-modern model takes into account both structural/social factors and medical/physiological factors can be helpful to address these limitations [99, 175, 149, 150]. Key to this approach is to consider *impairment* (i.e., physical or mental differences or limitations) and *disability* (i.e., social factors that limit a person’s participation due to differences in ability) as interrelated in nuanced ways [29, 150].

The medical and social models of disability differ in describing what “causes” the disabling experiences. The *medical model* has often been used in the medical or assistive technology fields because it “*focuses on the physical and functional limitations a person may demonstrate*” [99]. In contrast, a traditional interpretation of the *social model* frames disability as arising from mismatches between a person’s ability and the world. When people with disabilities face an access challenge, the blame for inaccessibility falls not on the individual, but on society’s laws, architecture, and ableist enforcers of exclusionary, discriminatory practices in life, work, and education [29, 175, 149, 150]. However, a pure social model interpretation does not recognize the tight relation between disability and the body [28]. Wendell argues: “*some unhealthy disabled people, as well as some healthy people with disabilities, experience physical or psychological burdens that no amount of social justice can eliminate*” [173].

Alternative models of disability aim to address the tension between the social and medical models (e.g., [22, 118, 99]). Among these, the interdependence model [22] can be a useful lens to view medical care, and physical therapists in particular, as part of an interdependent network that supports individuals across the lifespan. In medicine, the International Classification of Functioning, Disability and Health (ICF) model [118] has been used more heavily in recent years to understand the environmental and social barriers that limit function in ad-

dition to physiological barriers. In HCI, the post-modern model of disability “*privileges each individual’s unique lived experience...disability, illness, impairment, functional limitation, and bodily anomaly are separate but complementary issues, and successful assistive technology must account for all of these perspectives*” [99]. While an interdependence or ICF model highlight key factors in disability, a post-modern model emphasizes the interplay between social and physiological barriers.

Thus, the post-modern model of disability is a useful perspective to adopt in analyzing PT, a system which is socially constructed and maintained through processes like insurance, but one that is focused deeply on quality of physical movement. Indeed, the disabling situations, both created by an ableist society and physical differences in bodies, have important roles to play in understanding someone’s experience with disability [127, 15]. Particularly for people who may be involved in PT, not all daily pain and discomfort can be assuaged solely by a social interpretation of accessible practices. On the other hand, adopting the pervasive, and harmful assumption that physical and functional limitations are problematic and must be overcome by technology also does people a disservice. Instead of making solely social or physiologically based assumptions about why PT is inaccessible, the post-modern model of disability provides a lens to examine each individual’s experience of the complex interplay between social and physiological access barriers.

5.2.2 *Defining Physical Therapy*

The primary goals of doing PT are defined by the person doing PT and often include improving or maintaining movement and function. Physical therapists and the person doing PT strive to achieve this through a feedback loop of evaluation, plan development, implementation, and assessment [77]. The development of the plan of care generally involves setting measurable short- and long-term outcome goals and considering the resources available to the person. During the course of treatment, people generally meet with their physical therapists anywhere from weekly to monthly. People are encouraged to follow a (usually daily) PT exercise program that is tailored towards them and their goals to supplement the in-person PT sessions. When doing PT to improve symptoms of da/cc, high intensity or

dosage can be important for recovery and to maintain quality of life [89]. For example, the Levine protocol [152] is a rigorous eight-month cardiovascular and strength training routine for people with postural orthostatic tachycardia syndrome (POTS) where people work up to exercising 5-6 hours per week. The exercises, frequency of the exercises, duration, and number of repetitions are usually provided to people on a sheet of paper, and adapted during the in-person PT sessions as people progress.

However, many accessibility barriers exist throughout the PT process that makes completing a treatment plan less accessible, or even hinder starting PT in the first place. Technology has been developed to address a limited number of these barriers, but few studies provide guidelines or develop technology to address the intersectional needs of someone with da/ccs whose symptoms may fluctuate. We detail both of these areas below.

Barriers to PT

Access and adherence to in-person and at-home PT can be challenging for numerous reasons. For example, it can be challenging to obtain the in-PT dosage required for a da/cc due to the prohibitive cost and limited number of PT appointments a person can make with insurance [102, 151, 37]. In rural and/or low-and-middle income areas, lack of experienced physical therapists, insufficient community-based programs, long travel times, and high cost to travel are also significant barriers to attending in-person PT [27, 142, 119]. Unfortunately, adherence to at-home PT programs also comes with its own challenges such as limited time, daily stress, and lack of social support and guidance [17, 46]. There is room for innovation in the space of PT by adapting a new perspective around adherence failure. Lack of adherence is often viewed as a failure of the person doing PT to comply with medical advice [46]. A better understanding of PT success from an accessibility perspective is needed. For example, people with chronic conditions may have strong, uncomfortable symptoms that may worsen when doing exercises or symptoms that vary day-to-day [70, 127, 130, 151]. This variability within (e.g., daily fluctuation) and between people with chronic conditions is understudied, or viewed simply as an adherence issue, rather than an opportunity for improved technology design to help manage this variability [91].

5.3 Methods

We performed semi-structured interviews with twelve U.S.-based, da/cc individuals who are currently doing PT either with a physical therapist and/or on their own at home. Participants filled out a screening survey with demographic information and we selected participants to include multiple disabilities, genders, and races. The interviews were one hour-long on a video calling platform, due to COVID-19 restrictions. All participants were compensated with a \$15 Amazon gift card, and the interviews were recorded and transcribed by hand. The protocol was approved by our Institutional Review Board.

5.3.1 Interview Protocol

Our interviews were composed of three main parts. First, we discussed the participant's background with technology and disability. We discussed the technologies that the participants owned to understand potential sensors to use for PT tracking. We then asked about the participant's disability, including when they became disabled and how their disability affects their daily life. We also recorded what assistive technologies they used and accessibility challenges they faced with their devices.

In the second part of the interview, we asked the participant about their experience with PT, including the exercises they perform, their motivation for doing PT, and their goals that they are working towards. We asked them to differentiate their PT routines done with professional physical therapists (if they saw one) and routines they performed at home. We also discussed what factors demotivated them to do PT and the effect that their da/cc had on their exercises.

In the third part of the interview, we asked participants about potential sensor-based systems to aid in at-home PT. We provided examples like using accelerometers to track a PT exercise that then enables technology rewards like playing an episode of Netflix or opening social media. We talked with participants about when they would want the system to prompt them to do exercises, what exercises they would be willing to do, and how they could dismiss the notifications to do exercises. Given our imagined form factor involving smart devices, we asked if there were unique ways that our system could motivate them to

complete exercises (e.g., the user must do 10 squats before they can open Instagram), and if they would want the exercises to be logged manually or tracked automatically via sensor data.

We concluded the interview with demographic questions and information about compensation.

5.3.2 Analysis

Interview transcripts were analyzed using reflexive thematic coding [24, 25]. One author conducted all of the interviews and reviewed all of the interview transcripts to take notes and develop codes. Another author performed the same process on interview segments to check for gaps or biases in code coverage. The full list of 233 codes were discussed and revised by the first two authors. After the code book was complete, one author applied the final codes to all of the transcripts. After all the data was coded, the authors met and through discussion arrived at the broader themes presented in the results.

5.3.3 Positionality Statement

We recognize that, in performing reflexive thematic analysis, our backgrounds and positionality shaped our findings. This work was conducted by a mix of disabled and non-disabled white and Asian scholars who work in US institutions. Of the four authors, two are engineers who predominantly do research in improving rehabilitation outcomes. The other two authors are computer scientists engaged in mostly accessibility research. Three of the authors have experience receiving PT for a chronic condition and one author has supported multiple family members through receiving PT for chronic conditions, both via telehealth and in-person, as well as fighting for insurance supported access to PT for a chronic condition.

5.3.4 Participants

In total, we recruited and interviewed twelve participants from December 2020 - March 2021. Six participants identified as men, five as women, and one as non-binary. The participants' races included Caucasian/white (9), white Hispanic (1), African American (1), and Asian

Table 5.1: The self-defined da/cc(s) of each participant.

ID	Disability	ID	Disability
P01	Spinal injury affecting lower body	P07	Vestibular neuritis, depression
P02	Lower-body left-side hemiparesis, learning disability	P08	Gastroparesis, anxiety, depression, chronic back pain, torn miniscus
P03	Epidermolysis bullosa, acid reflux	P09	Herniated disc, radiculopathy
P04	Carpal tunnel, herniated disc, hydrocephalus, epilepsy, inverted scoliosis, fibromyalgia	P10	Autism, Ehlers-Danlos, chronic pain, complex PTSD, anxiety, depression, dyslexia
P05	Pulmonary lymphoma, irritable bowel syndrome, bronchospasm	P11	Hypermobility Ehlers-Danlos, partially deaf in one ear
P06	Chronic ankle pain after surgery	P12	Ehlers-Danlos, autoimmune disease, POTS, orthostatic hypotension, celiac disease, chronic gastritis

Disabilities included autism, cognitive disabilities, learning disabilities, mental health disabilities, health-related disabilities, and motor disabilities. Ten out of twelve participants had multiple da/ccs

Abbreviations: PTSD = post-traumatic stress disorder; POTS = postural tachycardia syndrome

(1). The mean age was 32.4 (range = 20-58). The highest degrees earned were high school diploma (3), Bachelor's degree (3), Master's degree (3), Associate degree (1), and other graduate degree (2). The da/ccs of each individual are listed in Table 5.1, and we briefly summarize below the participants' self-reported background on how their da/cc affects technology accessibility and daily life.

The participants' da/ccs affected the accessibility of activities such as work and employment (5), socialization with family and friends (4), and keeping up with non-disabled peers (2). Access barriers stemmed from da/cc symptoms such as chronic pain (6), fatigue (3), and nausea (1), their da/cc making it hard to sit, stand, or walk (6), and the effect of

their da/cc on their mental health (3). Technology adaptations like speech to text (5), ergonomic keyboards and mice (5), and mobility devices (5) were helpful in improving access to activities.

Further, the da/ccs that our participants had affected their technology use. Seven participants discussed challenges physically interacting with devices, such as hand fatigue (5), interacting with a standard mouse or keyboard (4), using a touchscreen, trackpad, or fingerprint detector (2), and pressing multiple buttons at once (4). Consuming digital content caused similar challenges. Six participants discussed how technology use exacerbates da/cc symptoms, including pain, strain, headaches, and motion sickness. Consequently, participants used a variety of low- and high-tech assistive technologies ranging from screen readers, to ergonomic devices, to placing adhesive tape on devices to allow for better grips. Accessibility barriers meant that five participants had to abstain from using technology to alleviate da/cc symptoms; these cases demonstrate the importance of designing PT-focused technologies with accessibility in mind.

5.4 Results

Our participants described many benefits and motivations for doing PT; however, in doing so, they encountered access barriers that were based in the embodiment of participant impairments, social constraints like financial and physical access to PT, and interactions between barriers. We then discuss how this information around motivators and challenges influenced participant's visions for future technologies supporting PT.

5.4.1 PT Improving Quality of Life is Motivating

Many people were motivated to do PT because of its potential to make areas of life more accessible or comfortable given their da/cc, such as working, socializing, and exercising.

For some participants, aiming to close the performance gap between their current abilities and those of the peers or themselves pre-da/cc was motivational: *"I know what it's like to be healthy and strong and I want that again"* (P10). For others, their motivation stemmed from maintaining quality of life. P02 discussed: *"my function is deteriorating, so I want to maintain my function and my current level of mobility as long as possible"* and

P05 was motivated *“to slow down the progression of disease”*. Improving and maintaining performance are different goals that affect what “progress” looks like and can impact how PT may be presented with technology.

Ten participants discussed improving mobility and strength as a tangible goal for maintaining or regaining quality of life (P01-P04, P06, P08-P12). *“Increasing mobility and the amount that I can walk”* (P02) was a goal for P01-P03, P06, P10, and P12. P01 describes how PT could increase their physical abilities, allowing them more freedom from social barriers: *“so I would be able to walk around a store or go outside without having to bring my wheelchair...it’s kind of a pain because I have to make sure the buildings that I go into are ramp accessible before I go there”* (P01), as many places are not accessible for people who use mobility devices. Gaining strength and stamina was also an important goal for P02, P09, P10 and P12 so they can participate in different activities like *“playing sled hockey”* (P02). These smaller, more actionable steps that support higher level goals provide opportunities for technology to support tracking and celebrating progress, which may increase participants’ motivation.

Participants also faced chronic symptoms due to their da/cc, which lowered their quality of life and made activities less comfortable. PT served as a way to combat these negative symptoms, such as pain (P02, P04-P06, P08, P09, P11) or dizziness (P07). Similarly, P03, P04 and P12 felt that, by improving their strength and stamina, they could avoid symptoms of their da/cc such as pain and scar tissue buildup. Two participants (P02, P11) performed PT intermittently and used their pain levels as a gauge for when to restart PT exercises.

Overall, we saw that participants had clear motivations for engaging in PT, which usually revolved around maintaining or increasing physical capability, or increasing quality of life through symptom mitigation. A purely medical goal for PT exercise compliance is completing the number of exercises set by the physical therapist; framing this goal in the context of their higher level goal (e.g., maintaining physical ability) could be motivating and improve adherence to PT routines [72].

5.4.2 Context-Dependent Barriers to PT

In-person (i.e., in clinics) and at-home are two common contexts for performing PT, each of which has their own benefits and challenges. While in person PT can be more engaging and allows access to therapist expertise, it can be physically and monetarily inaccessible. At-home PT, while less expensive, makes it more challenging to stay engaged and safely adapt and correctly perform exercises based on their current da/cc symptoms.

In-person PT was inaccessible to many of our participants due to insufficient insurance coverage and lack of transportation. Insurance issues, such as a gap in coverage (P10, P12), insurance only covering a limited number of sessions per year (P03, P09, P12), high out-of-pocket cost (P11), and being too far out from the initial diagnosis to get PT covered by insurance (P02) were all noted as barriers to in-person PT. The current insurance system in the US only allows for a certain number of PT visits with a physical therapist per year, and consequently, people often end up doing more at-home PT between or in lieu of in-person visits. Such financial barriers can be viewed as accessibility barriers both in the literal sense (since they remove access to PT) and also because people with disabilities are more likely to be living in poverty [26]. As a result of these insurance challenges, participants were forced to choose between going to unaffordable PT or exacerbating the symptoms of their da/cc. P11 discusses, *“as a way to try to save me money [P11’s physical therapist] said let’s try to meet less often.”* Transportation was another commonly noted challenge of in-person PT. Not all of our participants owned their own form of transportation. Public transportation is not always available, or not an option during a pandemic. Relying on relatives or friends can also cause additional tensions (P01). The pandemic further demotivated participants from attending in-person PT due to potential exposures (P04, P09, P11). For example, P11 discussed how meeting with their PT in-person is *“a calculated risk”* due to their da/cc making them high-risk for COVID-19.

Because of the many social barriers to in-person PT, many participants relied on at-home PT exercises to alleviate da/cc symptoms and improve mobility and strength. However, at-home PT also came with significant barriers such as requiring time and effort, causing discomfort and fatigue, dealing with fluctuating symptoms, managing safety, and feeling a

lack of engagement and variety. Nine participants discussed how doing the PT exercises took effort, time, and caused discomfort (P02-P04, P06-P08, P10-P12). For example, P07 discussed how their PT exercises exacerbates their da/cc symptoms, discouraging them and making the rest of their day less comfortable. Symptoms of da/cc can also limit the overall number of activities that a person can do in a day, making it even more challenging for people to incorporate PT into their everyday routines. For example, P11 discussed how they were *“just piled on with more and more exercises, it is getting harder to stay on track...it’s a challenge being motivated, it’s a challenge dealing with the fatigue [one of their da/cc symptoms], I’ve been feeling overwhelmed by it all. I just feel like it’s an endless list of things to do”*. In P11’s broader narrative, we see that PT was a time-burden that was compounded by the time to seek other medical care, which is an experience specific to people with illness-based da/ccs [173]. P11 additionally had multiple disabilities that they went to multiple physical therapists for, who all gave PT exercises that were solely specific to the problem at hand, without considering the larger picture of P11’s multiple conditions. P11’s narrative highlights how physiological barriers (e.g., fatigue due to da/cc) can interact with social barriers (e.g., navigating the complex healthcare system to find and receive care for multiple conditions), which made it challenging to fit PT exercises into their life.

Physiological barriers to PT access also affected participants’ confidence and motivation to do PT. For example, several participants reported physically not being able to do the PT exercises or not progressing in their PT (P01-P03, P08). P02 discussed how *“it’s very discouraging trying to use the theraband [a resistance band used to do their ankle PT exercise] sometimes I can’t even do it without the theraband resistance...that is one reason why I do those ones less.”* Other health symptoms not related to the da/cc that the individual is doing PT for could also affect people’s motivation. For example, P10 discusses how: *“if you’re really depressed, getting into the whole, ‘what’s the point of [doing PT], why bother”*’ (P10). During in-person PT, many of these physiological PT barriers can be mitigated with the help of the physical therapist because they can dynamically adapt the exercises. Adapting the PT exercises when doing at-home PT is challenging without access to this expertise. Many of our participants were given sheets of paper describing the PT exercises, the number of repetitions, and how often to do the PT exercises, which can make it difficult

to track and adapt the exercises to suit the person's needs and symptoms. The challenges adapting PT exercises to account for fluctuating symptoms can be considered as both a physiological barrier stemming from their da/cc symptoms as well as a social barrier of not having an accessible PT routine that can be easily adapted when people are having a bad symptoms day.

Another challenge of at-home PT was getting feedback on whether participants were doing the PT exercises safely and correctly to prevent injury. Seven participants (P01, P04, P06, P07, P09-P11) discussed the importance of checking in with their physical therapists when at-home PT exercises didn't alleviate da/cc symptoms or if they were not making progress. If people were not currently doing in-person PT or they needed feedback at that moment, they resorted to *"comparing myself as best as I can to the visuals on the screen"* (P12) using YouTube videos that they had searched. However, this can be a challenge because *"you have to be careful and make sure you're not looking at something that's not a good source of information"* (P11). Additionally, it is hard to ensure safety when doing at-home PT. For example, P07 discussed how falls were a safety concern for some of their PT exercises so they could only do them *"if there's someone home."* To alleviate these challenges, P10 wished they had *"an option to communicate [virtually and asynchronously] with my physical therapist...having that check-in monitoring and what to do next"*. Such a system could mitigate both safety and exercise progression issues.

Lack of engagement and variety in exercises (P02, P05, P09, P12) was another barrier to at-home PT. P02 wished the PT exercises *"were more engaging...[the physical therapists] used to try when I was a kid to make them more engaging, like stand on one leg and throw a ball back and forth. But when you get older, [making the PT exercises engaging] is not a thing anymore."* Lack of variety in exercises over time exacerbated the disinterest in completing at-home PT exercises. P12 also commented that performing exercises in a PT clinic was more engaging than at home because *"at least physically going into PT, there were people you could talk to, you had different tools that we're using, now it's just kind of the same old, same old"*. P12 alludes to the fact that the PT clinic was a source of socialization, or perhaps even community, which was key to their motivation for PT. Therefore, in addition to considering how to make people with da/ccs feel safe, technology that supports at-home

PT must address the issues of engagement and community.

Although in-person PT supplemented with at-home PT is considered the norm for exercise-based care for da/ccs, our participants discussed how in-person PT is often physically and monetarily inaccessible for them, forcing them to solely rely on at-home PT without physical therapist supervision to manage the symptoms of their da/cc. While at-home PT removes some barriers like transportation and risk of exposure to illness, it introduces new access barriers that are specific to people with da/ccs like concerns with safety when doing the exercises. Given the lack of communication channels with physical therapists outside of clinic visits, people with da/ccs often need to choose which set of barriers (those associated with at-home or in-person) are more appropriate, given their unique situation.

5.4.3 Participant Technology & Design Suggestions

All twelve participants were interested in incorporating technology into their at-home PT routines to alleviate at-home PT barriers such as time and effort, fluctuating symptoms, lack of motivation, and safety and injury (Table 5.2). Additionally participants cited insurance as a potential concern for several features. The participants suggested seven technology features that would improve at-home PT access: exercise presentation, tracking, technology rewards, notifications, hardware preferences, data security, and sharing progress.

Table 5.2: Summary of participant suggested features to include when developing technology to support at-home PT and access barriers that are addressed with the features.

Feature	Access Barriers	Participant Suggestions	Sample Quote
Exercise Presentation	TE FS	Choose between different exercises; customize depending on day; customize number of repetitions; support exercise progression; minimal setup (e.g., sit-to-stand)	<i>“[being able to choose exercises is helpful] because certain exercises are targeted to help with certain areas...so only I would know what I need for the day” (P04)</i>
Tracking	FS MO SI	Tracking PT progress is motivating; feedback useful to prevent injury or worsening da/cc symptoms	<i>“[I want to] make sure that I’m doing [the PT exercises] right because if I don’t, it’s gonna mess [up the ankle] even more” (P06)</i>
Technology Rewards	MO	Non-essential apps (e.g., not email); time sinks (e.g., social media); unlocking devices; commercial breaks (e.g., Netflix); congratulatory or fun elements	<i>“If I’m really in the thick of it with Grey’s Anatomy, that would absolutely motivate me [to do PT]” (P12)</i>
Notifications Timing	TE FS SI	Personalized notification timing and frequency; preemptive notifications for da/cc symptoms; remind to do later; offer easier exercises before dismissing	<i>“I want it to detect my pain somehow and then prompt me [to do the PT exercise]” (P09)</i>

Hardware Preferences	IN SI	Cameras with depth perception; wearables; cost and sensor failure concerns; some preferred software only	<i>“[software] would make it easier to get updates... hardware would be a never-ending investment” (P08)</i>
Data Security	IN	No data security concerns; want data to be anonymized; HIPAA compliant; data security against insurance is concerning	<i>“[The data tracked by technology should not be] abused by insurance companies” (P04)</i>
Sharing Progress	FS MO SI	Helpful for modifying exercises or proving that PT isn't working; sharing can improve accountability	<i>“[sharing progress with PT] would help if a modification needed to be introduced if something was severely going wrong” (P12)</i>

TE = time and effort; FS = fluctuating symptoms; IN = insurance; MO = lack of motivation; SI = safety and injury

Exercise Presentation

Customization of how and what types of exercises were presented was important to participants so that they could account for fluctuating symptoms and decrease the time and effort required to complete their PT exercises. Four participants wanted to customize their daily PT routine by choosing between multiple exercises in an application (P02, P04, P09, P11) so they could pick and choose according to their symptoms that day or equipment restrictions. As P04 explains, this option would be helpful *“because certain exercises are targeted to help with certain areas...so only I would know what I need for the day.”* Three participants wanted different groups of exercises to be presented on different days (P10, P11, P12). Other participants had a subset of exercises that they would want to get prompted to do throughout the day such as exercises with minimal setup. For example, P01 and P10 were interested in doing sit-to-stands because *“that’s something you could do anywhere...even*

if you just had your wheelchair” (P01). Spreading out exercises that could be done with minimal setup throughout the day was appealing to participants because it alleviated the need to carve out time specifically to do the PT exercises.

The other element that participants were most interested in customizing was the number of repetitions for each exercise (P05, P10). Customizing the number of repetitions could support participants in making an exercise goal more attainable. For example, on a bad symptoms day, having a smaller goal may lead to more long-term success in adherence and goal completion than trying and failing to reach the same number of repetitions each day.

Manually accounting for fluctuating symptoms took extra time and effort and was a barrier to PT exercise completion. An application that can decrease time and effort needed to make those daily adjustments around which exercises, what types of exercises, and the number of repetitions to complete can help alleviate such PT barriers.

Tracking

Participants were excited to potentially use camera or sensor data to track at-home PT. They thought about tracking both individual movements (for quality and safety purposes) and overall progress (e.g., how many days did they do PT). If combined with da/cc symptom tracking, participants thought this data could support self experimentation and data-informed conversations with physical therapists or other medical care providers.

Tracking movement quality is key to ensure exercise quality and prevent injury. All but two (P05, P08) participants had one or more exercises they did at-home for which they wanted technology to provide feedback on movement quality. P04, P07, and P09 wanted feedback on movement quality to prevent injuries and pain, while others wanted to ensure they were doing the exercise correctly (P01, P02, P03, P06, P10-P12) and using the right muscle groups (P11). P10 was interested in leveraging exercise accuracy tracking to self-experiment to see how their movement quality changed over the course of the day as they got more tired.

Participants were also excited to track PT progress so that they could keep track of and account for fluctuating symptoms. Six participants (P01, P04, P06, P09, P11-P12)

discussed tracking PT progress to get *“an actual physical visualization of what I’m able to do”* (P12). P06 suggested using the technology for self experimentation to better understand the efficacy of their PT; they commented that, if they could track their pain levels and the amount of PT exercises they are doing, they can see *“if I stretch every day consistently for a month, does that decrease my pain levels or not.”* In addition to supporting PT tracking, six participants (P01, P02, P06, P07-P09) were also interested in tracking other aspects of their life, such as tracking PT-like movements like walking during daily life or supporting tracking of da/cc symptoms such as pain, fatigue, or mental health, especially since symptoms and PT success and outcomes are so tightly intertwined.

Participants appreciated that better tracking encapsulated a more accurate representations of their daily practices into data that they could then share with their physical therapists. For example, four of our participants (P06, P08, P09, P11) were excited about being able to show more accurate pain levels and adherence to the PT routine over time. People felt that when their physical therapist asked them to rate pain levels in a session *“it’s just my answer at that time...if the number changes [later], I can’t call them and tell them differently”* (P09). By tracking their PT progress and symptoms of da/cc, the participants felt that *“the data can really help [the physical therapists] see daily [changes in da/cc symptoms]”* (P09) and it would be helpful to prove to physical therapists that *“I am indeed doing ... what you’ve told me to”* (P08).

Two participants (P03, P12) thought that tracking exercise progression could be especially useful when in-person PT is not available to them: *“I could actually keep using it even if I wasn’t actively in [in-person] PT”* (P12). In particular, adjusting the difficulty of the PT exercises between potentially infrequent PT visits can be challenging. Tracking exercises could be helpful to support the unlocking of *“new exercises that are similar to the ones that you’ve already been doing or if there’s a way to build up the exercises that you’re already doing...just make [the PT exercise] a little bit more challenging and a little bit more complex”* (P12). Similarly, two participants wanted the technology to encourage progression by increasing the difficulty of the exercises or by increasing the number of repetitions required automatically (P03, P12).

Technology Rewards

Technology-based rewards could help keep PT engaging and motivating but participants differed in preferences as to the rewards looked like. Some participants (P02, P03, P04, P06, P08) were interested in unlocking a device (e.g., phone, laptop), while others (P01, P02, P04, P11) preferred to only gate access to less critical "time sink" activities (e.g., social media). Using PT as a "commercial break" either before or as an interruption in a content consumption activity like watching TV, podcasts, or online videos was another popular suggestion (P02, P04, P07, P09, P10, P12). P12 mentioned how if they were *"really in the thick of it with Grey's Anatomy, that would absolutely motivate me"*. Eight participants (P01, P03, P04, P06-P08, P10, P11) requested congratulatory or fun elements to be embedded into the technology to increase motivation. P08 suggested a *"sound notification....where it does a clap...or says 'nice job on your walk'....[that] would make the user feel good about what they were doing"* and P04 mentioned *"it should help me celebrate and reward when I'm keeping a good track...we like apps that make us feel good when we're doing something right"*. P10 suggested a *"wheel of PT"*, where a randomized feature can pick different exercises for them to do. The final popular suggestion was having some type of game-like reward system, where they could accumulate points that add up to some sort of reward (P08, P09, P10, P11, P12). The rewards that were suggested varied, from free subscription services, to bloopers and special features for a show they are watching, and personalized quips and positive visual feedback for hitting certain goals.

Notifications

Participants saw the potential for personalized notifications to decrease the time and effort of remembering to do the exercises. Some participants preferred to spread out their exercises and notifications throughout the day (P03-P04, P08, P09, P11-P12), while others preferred to receive one notification and do all their exercises (P06-P07). Additionally, P10 discussed the importance of the system customizations extending to include notifications. For example, in addition to adjusting what and how many exercises to show on a *"bad pain day"*, the system could time notifications to be at low pain-points in the day. Other

participants saw the potential for more algorithmically driven notification timings based on their symptoms. For example, P09 was interested in receiving a PT exercise notification as a preemptive measure for or as a consequence of back pain: *“I want it to detect my pain somehow and then prompt me”* (P09).

All participants defined the need to dismiss notifications when they were busy or having a bad symptoms day. The most popular suggested dismissal method was a simple “remind to do later” option (P02, P04, P08, P11, P12). P03 suggested that, before fully dismissing, offering an easier exercise may encourage some PT rather than none in the moment. P04 and P12 additionally suggested disabling dismissals after a certain point to strongly encourage people to do the PT exercises or letting the person know that they have broken the record for the number of dismissals they have done. Finally, P12 suggested having the technology check-in with the person if they dismissed the notification too many times, which might indicate that a PT exercise is too challenging or that a person’s symptoms are worsening.

Hardware Preferences

As movement tracking and feedback was of interest, many participants were open to purchasing hardware such as cameras with depth perception (e.g., Kinect; P3, P07), wearables (e.g., Apple Watch; P04, P06, P09, P11), and implantable chips (e.g., Neuralink; P09). However, there were several concerns around the cost, upkeep, and comfort of hardware devices. For example, if the hardware device was a wearable, P02 mentioned *“I’m not sure I would put it on everyday”* as a potential challenge, and P03 mentioned that *“I’d be much, much less likely to use [wearables] ... because they would cause too much chafing,”* which is a concern for their da/cc. Others were concerned about both the base cost of hardware (P03, P05, P06-P07, P10-P11) as well as updates: *“[hardware] would be a never-ending investment”*. Participants noted that insurance coverage could help ease the cost burden, but P03 was not optimistic: *“I don’t feel like insurance is ready for the video game technology yet”*. Without insurance support, many participants seemed reluctant to acquire extra hardware sensors for PT tracking.

Data Security

Insurance was a significant fiscal barrier to in-person PT access, and data security against insurance companies was a significant concern for tracking at-home PT. For example, P05 commented that *“while movements on my arms might seem innocuous right now...”*, that information in the future could be used to prevent insurance coverage. P02-P04 were also concerned about the data being *“abused by insurance companies”* (P03) because *“the insurance company doesn’t need to know that I missed a week of PT because I was hospitalized...the physical therapist understands that, and she isn’t going to hold that against you, but the insurance company can and will”* (P04) because *“...they already deny you for everything”* (P02). While tracking provides powerful affordances to support people in their PT experiences, technology designers must always take precautions to safeguard against this insurance abuse of their data.

Sharing Progress

All but two participants (P02, P05) were interested in sharing their at-home PT progress with their physical therapist to help with motivation, adjust exercises for fluctuating symptoms, and prevent injury. Seven participants (P01, P03, P04, P06, P09, P10, P12) thought that sharing progress with their physical therapist would be helpful *“if a modification needed to be introduced”* (P12) and for the physical therapist *“to have some control over how [the exercises] are programmed”* (P04). P06 felt that sharing the data would be helpful to work with their physical therapist to understand *“where this PT regimen is working or it’s not working”*.

5.5 Discussion

Clare writes about the tensions between *“the wisdom that tells us the causes of the injustice we [disabled people] face lie outside our bodies, and also to the profound relationships our bodies have to that injustice”* [28]; the relationship between the body and built environment is evident in the experiences of people with da/ccs in performing PT. Our work highlights the tensions and intersections between social- and physiological-based barriers in PT access.

For example, in-person PT can be inaccessible for people due to lack of public transportation (social barrier) because they can't drive due to their da/cc symptoms (physiological barrier). While solutions like performing more at-home PT or sparsely attending in-person PT mitigate some of these issues, they introduce new access barriers. At-home PT can be difficult to complete if the person's da/cc symptoms are fluctuating or if there are concerns with safety and injury. Upon first glance, barriers like lack of transportation and fluctuating symptoms could be construed as individual barriers or compliance failures. However, when viewed from an accessibility perspective, the locus of responsibility for addressing such barriers may shift (such as reconsidering the accessibility of the transportation system as a social responsibility) as can the strategy for addressing challenges (such as expanding technology from monitoring to addressing fluctuating symptoms). In the words of P11, *"I'm just overwhelmed, and if technology can help make sort of a systematic way to address chronic illness challenges and hopefully the medical system will follow suit, I think that would be a huge benefit to people like me"*.

5.5.1 Design Recommendations to Support PT With Technology

People with da/ccs who are doing PT inherently face physiologically-based barriers to PT access: symptoms of the chronic condition that they are doing PT for. Although a traditional medical interpretation of such physiological barriers to PT access suggests a physiologically based solution (e.g., improve motivation and engagement through tracking and gamification of the PT exercises), considering both physiological and social access barriers, and how such barriers can be reinterpreted to focus on the individual's lived experience is important for developing technology that holistically supports the person's PT goals. We present three design recommendations (flexibility, movement tracking, community building) and one tension (insurance) that highlights the nuanced interplay between social and physiological access barriers and how technology solutions to PT access barriers can be reconsidered to take into account the lived experience of people with da/ccs. Although we do not specifically call out accessibility as a design recommendation, we emphasize that first and foremost, any PT technology must be built with digital accessibility (e.g., screen reader accessibility) in

mind from the start to prevent further inaccess to PT [134]. Technologies that continuously adapt to the person's abilities [6] and enable accessible inputs and calibration [7, 72] are ideal.

Recommendation: Design with Flexibility in Mind

Perhaps the most complex and pervasive issues mentioned in our interviews focused on the physiological barriers that arose between PT and a person's body. These barriers included feeling too symptomatic to perform exercises, exercises triggering symptoms that persist throughout the day, fluctuating symptoms, and complex interactions with multi-faceted and/or multiple conditions. Such barriers are traditionally considered "adherence issues", both externalized by physical therapists who encourage adherence by encouraging patients to remember their larger goals for doing PT [72] and internalized by the patients themselves, who can have feelings of guilt when they do not adhere to the exercise routines.

However, if we adopt an accessibility lens, we can view the lack of adherence as not an physiological issue, but one that arises from a mismatch between the prescribed PT exercises and how the individual is feeling on that day. This mismatch can be especially potent for people who have multiple conditions. For example, an exercise that alleviates one impairment or symptom may exacerbate another, or an individual is tasked with doing an insurmountable number of PT exercises for each of their multiple conditions. These complex, contradictory health needs could be alleviated with attentive care and guidance from a physical therapist that can adapt exercises to match people's abilities for that day or holistically evaluate a person's multiple conditions. However, this barrier can often be compounded by social barriers like insurance preventing people from accessing in-person PT or the structure of the medical system making it challenging for a person to get a holistic evaluation and exercise recommendations for their multiple conditions.

In such scenarios, the individual doing the PT is the expert in their own condition. Technology that supports people with da/ccs doing PT should ensure that people can flexibly adapt the exercises and the number of repetitions to match what they need for their body on that day and provide people with the data they need to make those decisions.

For example, applications to support PT should be built in a way that records the abilities (e.g., pain level less than 4/10, dizziness less than 3/10) and time needed to complete the exercise. Then, when a user chooses to perform PT, the exercises and the number of repetitions can be tailored based on the individual's symptoms and abilities to increase the odds of success. To fully and holistically support people in managing their da/cc, flexibility must be built-in to technology supports.

Recommendation: Allow for Tracking Movement Quality and Progress

Due to recent advances in wearable sensors and decreased computing cost, technology is opportunistically situated to enable tracking for movement quality and exercise progress. Tracking movement quality was a highly requested feature that our participants felt would help ensure safety and prevent injury. Tracking exercises to prevent injury using wearables or camera-based technology is currently being explored in literature [72], but not necessarily within a PT context [133]. With the ability to track the quality of PT exercises, at-home PT becomes more accessible to people with da/ccs who are concerned about the potential for muscle or joint injury without physical therapist supervision.

Additionally, technology can reduce the burden of tracking exercise progress and completion to highlight people's successes and improve motivation, aid in self-experimentation and adaptation of PT exercises, and open the door to having conversations with physical therapists about adapting exercises. Self-tracking can be burdensome and another way that managing a medical condition can shorten the day for people with da/ccs [45, 168]. To alleviate this time-burden, wearable sensors or camera-based technologies can be used to automatically track the type and number of exercises completed. Algorithmic methods that enable safe and automated tracking of PT exercises is an unsolved issue that requires further research. During algorithm development, it is important to consider our participants' concerns about purchasing additional hardware. Therefore, supporting tracking with wearable technologies that people may already own such as smartwatches or enabling tracking with solely software-based solutions using a smartphone may make it easier for people to incorporate automated tracking into their PT exercise routines.

Recommendation: Create Technology That Encourages Community-Building

Our participants felt that being accountable to a community, whether it be with other people doing PT or with a physical therapist, is helpful for staying motivated to complete their exercises. In one study of a community-building PT application, participants found that the community was helpful for improving motivation and for comparing their PT exercises to other people who had similar conditions so they could experiment with new PT exercises [97]. Although there were concerns with misleading information [97], information sharing could be a useful work-around for when people are unable to see a physical therapist to get updated exercises. Additionally, virtual communities could facilitate encouragement and engagement between users, especially those who have rare or multiple conditions [97].

Facilitating supportive interactions between the person and their physical therapist through sharing of PT progress between in-person sessions may also increase familiarity and a sense of accomplishment between the patient and the physical therapist. It can be challenging for physical therapists to keep track of how patients are doing in-between sessions [72]. Consistent data-sharing could help catch issues with the PT routines in a timely manner and prevent communication issues. In situations where people are doing less than the prescribed number of exercises, it is important to frame the incomplete adherence as a sign that the PT routine is too time consuming, hard, or challenging because of other symptoms, rather than simply an adherence issue, to prioritize the lived experiences of the person doing PT. This reframing highlights where the PT routine might not be the best suited for the person, not that the person is primarily or solely to blame for non-adherence. These discussions, supported by data tracking, will hopefully lead to more appropriate exercise selection given a more holistic view of the person and lead to increased PT success, and therefore, adherence.

Tension: Designing with Insurance in Mind

Finally, to support tracking of PT with technology, we cannot ignore the constant presence of insurance concerns. Insurance was a major barrier to PT access for many participants, all of whom were located in the US. While technology could help participants work around

the in-person limits insurance imposes, participants were concerned about the ability for data collected with technology to be used against them. Therefore, technology that is built to support PT must be built with data privacy and use in mind.

Difficulties with and mistrust of insurance were near-ubiquitous in our interviews. The first major barrier to PT that many people faced was lack of coverage. Generally, one needs to be employed to receive benefits that include PT coverage, which already excludes a considerable number of disabled people from receiving PT. Even if someone with a da/cc wanted more flexible work hours or less working hours to better manage their symptoms, that is often not feasible because benefits are directly tied to the job. While health insurance is available in the US for unemployed individuals, navigating that infrastructure with a da/cc can be challenging, and the coverage can be insufficient [135, 122]. One major improvement that could be made in this area is to deploy (accessible) technology to help navigate finding and obtaining insurance that provides sufficient coverage for people's needs.

Further, these problems around achieving and maintaining PT coverage were so pervasive in our participants' lives that it affected their willingness to use technology to support PT. While a few participants thought a PT system that tracks progress and movements could serve as evidence for skeptical doctors or physical therapists, many other participants were uncomfortable with the PT app tracking progress because of the potential insurance ramifications; they were concerned that their progress, or lack thereof, could be acquired by insurance companies and used to deny coverage. Therefore, though technology can be part of the solution for increasing the accessibility of PT, insurance politics indubitably need to change before the access challenges are fully removed.

With these insights in mind, technology that supports PT must consider data privacy. Most work on data privacy in healthcare focuses on securing medical data collected by the healthcare industry against security breaches [121, 4] Our results bring up intriguing considerations for data privacy policies and systems to protect private patient data from being misused by the healthcare system, including insurance companies. Further, our work uncovers complex ethical considerations when creating such technology. If technology that could have negative consequences for users is built and adopted by insurance companies as "required for coverage," the technology we use and the data it collects could actively harm

participants with respect to PT and healthcare access.

5.5.2 Limitations and Future Work

One area of PT that we did not discuss in this paper is the rising popularity of telemedicine or virtual PT, where the physical therapist and the individual receiving care meet on video call for diagnosis and to receive care [36, 144]. Similar to reflections on virtual workplaces [86, 107, 104], a new, virtual setting for PT can improve some accessibility issues while introducing new challenges. Telemedicine can alleviate some aspects of the physical inaccessibility of in-person PT, such as removing the need to find transportation and the risks of exposure for people with weakened immune systems. However, similar barriers to at-home PT such as the potential for injury, lack of equipment, and lack of safety without the physical presence of the physical therapists exist in the virtual setting [160]. Additionally, virtual PT can be difficult because video calling can only provide so much information to properly diagnose and provide recommendations to the person doing the PT exercises [160]. Here, new and emerging sensors such as wearables and 3D cameras could be leveraged to provide the physical therapist with the information they need to care for their patients.

Additionally, PT access is not only challenging for people with chronic health conditions. People who are blind or low vision, d/Deaf or hard-of-hearing, and/or neurodiverse could also need PT access for acute (e.g., breaking a leg) or chronic (e.g., having a stroke) conditions. Consequently, the accessibility of non-visual PT, for example, is a key part of injury recovery for people who are blind or low vision, but it is not the focus of current accessibility research or standards. Although technology can never replace the knowledge and expertise of physical therapists, our study begins to identify potential ways in which technology can improve PT access, particularly if technologies are designed with access in mind from the start.

In the broader context of our conversations with our participants, many of the social and physiological PT access barriers that participants mentioned were barriers to healthcare access in general. Yet, accessibility research in healthcare remains limited. Our work highlights the many areas (reminders, tracking, motivational rewards, data security) where

technology could improve healthcare access for people with da/cc.

Lastly, beyond PT access, our study revealed several HCI areas in which needs arose at the intersection of social and physiological access barriers for people with da/ccs. For example, mental health was a PT barrier for many participants and is generally an understudied chronic condition in HCI accessibility literature [91] that limits participation in daily life due to both social and physiological barriers [117]. Further, almost all of our participants discussed overlapping social and physiological access barriers to performing their jobs, such as having to limit work hours due to chronic pain. We also suggest taking a post-modern lens of disability to other important situations for people with da/ccs, such as education, leisure activities, or care work. Given the pervasiveness of access barrier stemming from the body and society in many areas of life for people with chronic conditions, we see the intersection of physiological and social access barriers as a rich area of future HCI accessibility research.

5.6 Conclusion

Through our interview study with twelve individuals who self-identify as having a disability and/or chronic condition (da/cc), we gained insights into the complex issues that this group of individuals face at the intersection of social and physiological barriers when accessing physical therapy (PT). Our findings characterize both in-person and at-home PT as imperfect solutions due to transportation, insurance, safety, and equipment issues and identify barriers that technology could help alleviate. Notably, PT routines that do not center the daily experiences of people with da/ccs are inherently challenging and less effective. For example, daily symptom fluctuations can make it difficult or impossible to complete prescribed PT. Through our study, we characterized factors that influence a person's PT experience including motivators and access challenges, which outline a clear set of needs that technology can fill to support PT. Building technology that is motivating and adaptable between and within users and that includes key factors (e.g., tracking progress, data privacy) may support PT accessibility. Finally, our work highlights the tightly interwoven threads of social and physiological barriers to PT access. Our analysis would be incomplete looking at one or the other, and they cannot be disentangled. We call for other researchers

to thoughtfully consider the nuanced inter-relation between physiological and social access needs in accessibility work moving forward, particularly when studying people with chronic conditions.

Chapter 6

CONCLUSION AND FUTURE WORK

This dissertation provides an important contribution toward applying control theory methods to improve accessibility and health for people with and without disabilities. To achieve this, I developed individualized HMI models (Ch 2), quantified how EMG interfaces could enhance HMI control (Ch. 3), leveraged co-adaptation to design interfaces that adapt to individual users' needs (Ch. 4), and identified how technology could improve accessibility and health for individuals with disabilities (Ch. 5). Below, I summarize the main contributions of each chapter and identify avenues for future work.

6.1 Individualized HMI Model for Continuous Interactions

The objective of Chapter 2 was to develop a predictive model for continuous HMIs and leverage the model to quantify how handedness affects learned feedforward and feedback controllers. We found that a combined feedforward and feedback model predicted user's interactions with a dynamic machine better than a solely feedback model [181]. We additionally found that handedness does not affect learned controllers, and feedback controllers adapt, but feedforward controllers do not [186]. Future research based upon and informed by this work could include:

6.1.1 *How is learning of feedback and feedforward controllers affected by neurologic injury or neurologic disorders?*

In motor control literature, researchers have identified that neurologic disorders like Huntington disease [154] and cerebellar ataxia [155] affect the learning of feedback and feedforward mechanisms during a center-out reaching task using a robotic manipulandum. The experimental method developed in Chapter 2 demonstrates potential to similarly separately quantify how feedback and feedforward controllers are affected by neurologic disor-

ders. While the center-out reaching tasks enable researchers to identify disorders that affect learning of feedback and feedforward control, there are no clear methods to predict and augment the person's affected controller with an intelligent agent. The methods investigated in this dissertation demonstrate how control theory and game theory could be used to develop predictive models of HMI and use those models to enhance HMIs. Chapter 4 in particular suggests how we could begin to augment a user's feedback controller using co-adaptation assuming that we can first model the person's feedback controller. Future work on extending co-adaptation for people with disabilities will first require identifying how learning of feedback and feedforward controllers are affected by neurologic injury or neurologic disorders to determine whether continuous tracking tasks produce similar results to the well-studied center-out reaching tasks.

6.1.2 How does time-scale affect learned feedback and feedforward controllers?

Although we observed that learned controllers transferred between the left and right hands for the prescribed task, there are many real-world examples where learning does not transfer between hands. One example of this is writing. For many individuals who are right- or left-hand dominant, writing with the opposite hand is a challenging and frustrating experience. One area of future work would be to have participants interact with our experimental setup over a much longer time scale than what we have investigated so far (~30 minutes) to identify how learning of feedback and feedforward controllers and transfer of learning is affected by the length of time that participants have to learn a task. For example, although we did not observe changes in feedforward controllers over the 30-minute task, previous studies have shown feedforward controller adaptation when participants learn the tracking task over a two-week period [192]. Identifying the time-scale in which learning of novel dynamical systems no longer transfers across hands would be important to develop HMIs that are easy to use regardless of the hand used to interact with the machine.

6.2 *Electromyography as Machine Input for People With and Without Disabilities*

The objective of Chapter 3 was to determine whether EMG interfaces improve continuous cursor control and whether disability affects EMG versus manual control. We found that EMG interfaces improve feedforward model formulation for a second-order task but not a first-order task [187]. This seemed to be due to participants being able to track much higher frequencies with the EMG interface compared to the manual interface. We additionally found that while participants with disabilities performed much worse than participants without disabilities, participants with disabilities had similar feedforward models as participants without disabilities [187]. This suggests that EMG interfaces could be useful for acceleration-based tasks like flying a drone or driving a car and that EMG interfaces could improve interface accessibility if we augment the person's feedback controller. Several future directions could build on this work:

6.2.1 How can EMG and manual interfaces be smartly combined?

We identified that EMG interfaces may be useful for high-frequency acceleration-based tasks. An exciting avenue for future work is to consider how EMG and manual interfaces could be smartly combined depending on the task at hand to enhance performance. For example, if the user is controlling a drone flying through a forest (a high-frequency acceleration-based task), the user could rely more heavily on the EMG interface, whereas if the user was controlling a powered wheelchair (a low-frequency velocity-based task), the user could rely more heavily on the manual interface. Having a smart interface that could identify situations where one interface may perform better than another and seamlessly combine the two interfaces optimally could improve user performance and improve user interactions.

6.2.2 How can co-adaptation improve EMG interfaces?

Because biosignal-based interfaces like EMG interfaces suffer from variability within and between individuals, co-adaptation could be a useful strategy for ensuring that the EMG interfaces are continuously calibrated. The time-scale of how often the EMG interface would

need to co-adapt to the individual and how quickly the co-adaptation should occur are open questions that must be answered before such interfaces could be deployed in the real world. Additionally, understanding how often users would be willing to calibrate an EMG interface and how they would like to be prompted to do that are also open questions that need to be answered.

6.3 *Enhancing HMIs with Co-Adaptive Learning*

The objective of Chapter 4 was to determine whether game theory could be used to augment the person's feedback controller and to understand what differences exist between a co-adaptive interface and a baseline. We found in experiment that participants converged to different interface controllers, and performance and human effort improved for the 1st order interface controller with co-adaptation compared to baseline. We additionally found in simulation that a 2nd order human controller at medium human penalties approximates our experimental results of improved performance and human effort. This suggests that co-adaptive interfaces could improve user interactions by improving performance and decreasing user effort with appropriate parameterization. Future research that can build on this work includes:

6.3.1 How can we develop a more accurate model of human interface co-adaptation?

The human interface co-adaptation model simulated in Ch. 4 simplifies the HMI to solely the stimulated frequencies. However, our experimental investigations demonstrate that adaptation of the user response at non-stimulated frequencies also plays an important role in improving performance. Future simulation and modeling work should investigate how time- and frequency-domain models can be developed that take into account how the human adapts to an adaptive interface and fixed machine at both stimulated and non-stimulated frequencies. Additionally, we did not investigate the effects of modeling the human as a higher-order controller. Modeling the human as a more complex dynamical system than a second-order system may be more effective at simulating our experimental results.

6.3.2 Can we shape user learning by changing penalty terms and learning rates?

In my dissertation work, I solely tested one penalty term and learning rate experimentally. Future studies should consider how penalty terms and learning rates affect performance and human and interface effort and whether we can shape user learning by adapting these parameters. Being able to shape user learning has implications for healthcare applications in particular, as being able to predictably adapt user effort could be beneficial for rehabilitation applications.

6.4 Supporting Accessibility and Health with Technology

The objective of Chapter 5 was to identify how biosignal-based interfaces could best be leveraged to improve accessibility and health for people with disabilities. We found that biosignal-based HMIs that provide continuous feedback from the human and the machine could be useful in supporting people's health [183] but the way people wanted to interact with the technology differed from individual to individual. The study also highlighted the importance of co-developing accessible and health-focused interfaces with people with disabilities. Several future directions that could build on this work:

6.4.1 How can individualized and adaptive HMIs support people's accessibility and health?

This dissertation laid the groundwork for how continuous interactions can be modeled and enhanced when adults without disabilities use manual interfaces like joysticks, sliders, and cursors. Potential areas for future work is to identify how we can adapt these modeling techniques and algorithms to be used with people with disabilities to improve device accessibility. For example, we can develop adaptive HMIs where the human and the interface can continuously co-adapt to a task at hand despite fluctuations in the user's abilities due to a disability or situational impairment. Key to this approach is to holistically consider not just optimal algorithm parameters for the task, but also what movements users find comfortable and intuitive to perform during human-machine interaction. Additionally, it is important to consider the placement of biosignal sensors such that it enables accessible input for the individual based on their abilities.

6.4.2 How can we ensure that multi-channel biosignals-based HMIs are accessible to people with disabilities?

To maximize the potential of individualized and co-adaptive HMIs, they must be translated to be used in daily life and broadly disseminated. This could be achieved by developing toolkits that enable application developers to include accessible biosignals-based HMIs as a plugin to their applications. Such toolkits on the developer side could enable plug-and-play HMI control that matches each user's abilities and needs. This includes both continuous HMIs, like controlling the movements of a virtual avatar in VR, and discrete HMIs, like tapping on an icon or gripping an object in VR.

6.4.3 How can we leverage HMIs to improve healthcare?

Always-available multi-channel biosignals present exciting opportunities for continuous monitoring and analysis of the person's abilities and health as they fluctuate over time. Currently, the standard of care is to diagnose and treat people's health based on the person's symptoms when they go to a doctor's office. Technology that closes the loop by continually monitoring, adapting, adjusting, and reporting on the person's health could help individualize treatments, potentially improving treatment outcomes and decreasing negative side effects. One domain of healthcare that I have begun to assess user needs in this dissertation is physical therapy [183]. With continuous monitoring of biosignals with wearable sensors, rehabilitation movements can be tracked, monitored, and adjusted based on the person's progress and goals. Additionally, rehabilitation movements can be integrated into people's daily interactions with technology and could help improve adherence and function.

6.5 Human-machine Interfaces for Accessibility and Health

The goal of this dissertation was to develop the necessary tools to build individualized, adaptive HMIs for people with and without disabilities. The key contributions of developing a predictive model for continuous HMIs, leveraging the model to identify how EMG interfaces could improve HMIs, and building a game-theoretic framework for co-adaptive interfaces provide a strong foundation for applying closed-loop biosignals-based interfaces

towards real-world accessibility and health applications. In addition, my work on identifying user needs highlight the need for co-adaptive and individualized HMIs for accessibility and health for people with disabilities. The potential future work outlined here provides ways in which my dissertation work can be extended to build better predictive HMI models and begin to translate adaptive interfaces into the real world.

BIBLIOGRAPHY

- [1] Dataset on trajectory tracking, 2021.
- [2] David A Abbink, Erwin R Boer, and Mark Mulder. Motivation for continuous haptic gas pedal feedback to support car following. In *2008 IEEE Intelligent Vehicles Symposium*, pages 283–290. IEEE, 2008.
- [3] Reza Abiri, Soheil Borhani, Eric W Sellers, Yang Jiang, and Xiaopeng Zhao. A comprehensive review of EEG-based brain–computer interface paradigms. *Journal of Neural Engineering*, 16(1):011001, 2019.
- [4] Karim Abouelmehdi, Abderrahim Beni-Hssane, Hayat Khaloufi, and Mostafa Saadi. Big data security and privacy in healthcare: A review. *Procedia Computer Science*, 113:73–80, 2017.
- [5] Johnny Accot, Shumin Zhai, et al. Beyond fitts’ law: Models for trajectory-based HCI tasks. In *CHI*, volume 97, pages 295–302. Citeseer, 1997.
- [6] Gazihan Alankus and Caitlin Kelleher. Reducing compensatory motions in motion-based video games for stroke rehabilitation. *Human–Computer Interaction*, 30(3-4):232–262, 2015.
- [7] Gazihan Alankus, Amanda Lazar, Matt May, and Caitlin Kelleher. Towards customizable games for stroke rehabilitation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2113–2122, 2010.
- [8] Christoph Amma, Thomas Krings, Jonas Böer, and Tanja Schultz. Advancing muscle-computer interfaces with high-density electromyography. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 929–938. ACM, 2015.

- [9] Javier Andreu-Perez, Daniel R Leff, Henry MD Ip, and Guang-Zhong Yang. From wearable sensors to smart implants—toward pervasive and personalized healthcare. *IEEE Transactions on Biomedical Engineering*, 62(12):2750–2762, 2015.
- [10] Joaquin A Anguera, Colleen A Russell, Douglas C Noll, and Rachael D Seidler. Neural correlates associated with intermanual transfer of sensorimotor adaptation. *Brain Research*, 1185:136–151, 2007.
- [11] American Physical Therapy Association. *Guide to Physical Therapist Practice*. 2nd ed. Phys Ther., 2001.
- [12] Karl Johan Aström and Richard M Murray. *Feedback systems: An introduction for scientists and engineers*. Princeton University Press, 2010.
- [13] Ju-Yeoul Baek, Jin-Hee An, Jong-Min Choi, Kwang-Suk Park, and Sang-Hoon Lee. Flexible polymeric dry electrodes for the long-term monitoring of ECG. *Sensors and Actuators A: Physical*, 143(2):423–429, 2008.
- [14] Leia B Bagesteiro and Robert L Sainburg. Handedness: Dominant arm advantages in control of limb dynamics. *Journal of Neurophysiology*, 88(5):2408–2421, 2002.
- [15] MG Baker and MR Pinder. Relations between the disabled and chronic sick and society: Towards a better understanding. *International Journal for the Advancement of Counselling*, 12(2):137–142, 1989.
- [16] John V Basmajian. Muscles alive. their functions revealed by electromyography. *Academic Medicine*, 37(8):802, 1962.
- [17] S Frances Bassett. The assessment of patient adherence to physiotherapy rehabilitation. *New Zealand Journal of Physiotherapy*, 31(2):60–66, 2003.
- [18] Amy J Bastian. Learning to predict the future: The cerebellum adapts feedforward movement control. *Current opinion in neurobiology*, 16(6):645–649, 2006.
- [19] Tamer Başar and Geert Jan Olsder. *Dynamic noncooperative game theory*. Number 23 in Classics in applied mathematics. SIAM, Philadelphia, 2nd ed edition, 1999.

- [20] Hrvoje Benko, T Scott Saponas, Dan Morris, and Desney Tan. Enhancing input on and above the interactive surface with muscle sensing. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*, pages 93–100, 2009.
- [21] Peter Benner, Volker Mehrmann, Vasile Sima, Sabine Van Huffel, and Andras Varga. SLICOT—A subroutine library in systems and control theory. In Biswa Nath Datta, editor, *Applied and Computational Control, Signals, and Circuits*, pages 499–539. Birkhäuser Boston, 1999.
- [22] Cynthia L Bennett, Erin Brady, and Stacy M Branham. Interdependence as a frame for assistive technology research and design. In *Proceedings of the 20th International ACM SIGACCESS Conference on Computers and Accessibility*, pages 161–173, 2018.
- [23] Daniel A Braun, Pedro A Ortega, and Daniel M Wolpert. Nash equilibria in multi-agent motor interactions. *PLoS computational biology*, 5(8):e1000468, 2009.
- [24] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101, 2006.
- [25] Virginia Braun and Victoria Clarke. Reflecting on reflexive thematic analysis. *Qualitative Research in Sport, Exercise and Health*, 11(4):589–597, 2019.
- [26] Debra L Brucker, Sophie Mitra, Navena Chaitoo, and Joseph Mauro. More likely to be poor whatever the measure: Working-age persons with disabilities in the United States. *Social Science Quarterly*, 96(1):273–296, 2015.
- [27] Neale R Chumbler, Xinli Li, Patricia Quigley, Miriam C Morey, Dorian Rose, Patricia Griffiths, Jon Sanford, and Helen Hoenig. A randomized controlled trial on stroke telerehabilitation: The effects on falls self-efficacy and satisfaction with care. *Journal of Telemedicine and Telecare*, 21(3):139–143, 2015.
- [28] Eli Clare. Stolen bodies, reclaimed bodies: Disability and queerness. *Public Culture*, 13(3):359–365, 2001.

- [29] David Copley. *Disability and International Development: a Guide for Students and Practitioners*. Routledge, 2018.
- [30] W J Conover. *Practical nonparametric statistics*. John Wiley & Sons, 1999.
- [31] Elaine A Corbett, Eric J Perreault, and Todd A Kuiken. Comparison of electromyography and force as interfaces for prosthetic control. *Journal of Rehabilitation Research and Development*, 48(6):629, 2011.
- [32] S Cutlip, J Freudenberg, N Cowan, and R B Gillespie. Haptic feedback and the internal model principle. In *IEEE World Haptics Conference (WHC)*, pages 568–573, July 2019.
- [33] Siddharth Dangi, Amy L Orsborn, Helene G Moorman, and Jose M Carmena. Design and analysis of closed-loop decoder adaptation algorithms for brain-machine interfaces. *Neural Computation*, 25(7):1693–1731, 2013.
- [34] Carlo J De Luca. The use of surface electromyography in biomechanics. *Journal of Applied Biomechanics*, 13(2):135–163, 1997.
- [35] Dalia De Santis. A framework for optimizing co-adaptation in body-machine interfaces. *Frontiers in Neurobotics*, 15:40, 2021.
- [36] Benedetta Demartini, Federica Bombieri, Diana Goeta, Orsola Gambini, Lucia Ricciardi, and Michele Tinazzi. A physical therapy programme for functional motor symptoms: A telemedicine pilot study. *Parkinsonism & Related Disorders*, 76:108–111, 2020.
- [37] Gail D Deyle, Stephen C Allison, Robert L Matekel, Michael G Ryder, John M Stang, David D Gohdes, Jeremy P Hutton, Nancy E Henderson, and Matthew B Garber. Physical therapy treatment effectiveness for osteoarthritis of the knee: A randomized comparison of supervised clinical exercise and manual therapy procedures versus a home exercise program. *Physical Therapy*, 85(12):1301–1317, 2005.

- [38] Ashesh K Dhawale, Maurice A Smith, and Bence P Ölveczky. The role of variability in motor learning. *Annual Review of Neuroscience*, 40:479–498, 2017.
- [39] Jörn Diedrichsen, Reza Shadmehr, and Richard B Ivry. The coordination of movement: optimal feedback control and beyond. *Trends in Cognitive Sciences*, 14(1):31–39, 2010.
- [40] Frank M Drop, Daan M Pool, Herman J Damveld, Marinus M van Paassen, and Max Mulder. Identification of the feedforward component in manual control with predictable target signals. *IEEE Transactions on Cybernetics*, 43(6):1936–1949, 2013.
- [41] Frank M Drop, Daan M Pool, Marinus M van Paassen, Max Mulder, and Heinrich H Bülthoff. Effects of target signal shape and system dynamics on feedforward in manual control. *IEEE Transactions on Cybernetics*, 49(3):768–780, 2018.
- [42] Frank M Drop, Daan M Pool, Marinus René M van Paassen, Max Mulder, and Heinrich H Bülthoff. Objective model selection for identifying the human feedforward response in manual control. *IEEE Transactions on Cybernetics*, 48(1):2–15, 2016.
- [43] Susan V Duff and Robert L Sainburg. Lateralization of motor adaptation reveals independence in control of trajectory and steady-state position. *Experimental Brain Research*, 179(4):551–561, 2007.
- [44] Deborah J Edwards, Dikaios Sakellariou, and Sally Anstey. Barriers to, and facilitators of, access to cancer services and experiences of cancer care for adults with a physical disability: A mixed methods systematic review. *Disability and Health Journal*, 13(1):100844, 2020.
- [45] Daniel A Epstein, An Ping, James Fogarty, and Sean A Munson. A lived informatics model of personal informatics. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 731–742, 2015.
- [46] Rosie Essery, Adam WA Geraghty, Sarah Kirby, and Lucy Yardley. Predictors of adherence to home-based physical therapies: A systematic review. *Disability and Rehabilitation*, 39(6):519–534, 2017.

- [47] Dario Farina, Ning Jiang, Hubertus Rehbaum, Aleš Holobar, Bernhard Graimann, Hans Dietl, and Oskar C Aszmann. The extraction of neural information from the surface EMG for the control of upper-limb prostheses: emerging avenues and challenges. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 22(4):797–809, 2014.
- [48] Heather A Feldner, Darrin Howell, Valerie E Kelly, Sarah Westcott McCoy, and Katherine M Steele. “Look, your muscles are firing!”: A qualitative study of clinician perspectives on the use of surface electromyography in neurorehabilitation. *Archives of Physical Medicine and Rehabilitation*, 100(4):663–675, 2019.
- [49] Leah Findlater, Alex Jansen, Kristen Shinohara, Morgan Dixon, Peter Kamb, Joshua Rakita, and Jacob O Wobbrock. Enhanced area cursors: Reducing fine pointing demands for people with motor impairments. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, pages 153–162. ACM, 2010.
- [50] Leah Findlater, Karyn Moffatt, Jon E Froehlich, Meethu Malu, and Joan Zhang. Comparing touchscreen and mouse input performance by people with and without upper body motor impairments. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 6056–6061. ACM, 2017.
- [51] Kenneth Flowers. Handedness and controlled movement. *British Journal of Psychology*, 66(1):39–52, 1975.
- [52] World Confederation for Physical Therapy. Policy statement: Description of physical therapy, May 2019.
- [53] Rebecca Forkan, Breeanna Pumper, Nicole Smyth, Hilary Wirkkala, Marcia A Ciol, and Anne Shumway-Cook. Exercise adherence following physical therapy intervention in older adults with impaired balance. *Physical Therapy*, 86(3):401–410, 2006.
- [54] Anders Fougner, Øyvind Stavdahl, Peter J Kyberd, Yves G Losier, and Philip A Parker. Control of upper limb prostheses: Terminology and proportional myoelectric

- control – a review. *IEEE Transactions on neural systems and rehabilitation engineering*, 20(5):663–677, 2012.
- [55] Nizan Friedman, Vicky Chan, Andrea N Reinkensmeyer, Ariel Beroukhim, Gregory J Zambrano, Mark Bachman, and David J Reinkensmeyer. Retraining and assessing hand movement after stroke using the musicglove: Comparison with conventional hand therapy and isometric grip training. *Journal of Neuroengineering and Rehabilitation*, 11(1):1–14, 2014.
- [56] Krzysztof Z Gajos, Daniel S Weld, and Jacob O Wobbrock. Automatically generating personalized user interfaces with supple. *Artificial Intelligence*, 174(12-13):910–950, 2010.
- [57] Francesca Gandolfo, FA Mussa-Ivaldi, and Emilio Bizzi. Motor learning by field approximation. *Proceedings of the National Academy of Sciences*, 93(9):3843–3846, 1996.
- [58] Karunesh Ganguly and Jose M Carmena. Emergence of a stable cortical map for neuroprosthetic control. *PLoS biology*, 7(7):e1000153, 2009.
- [59] R B Gillespie, A H Ghasemi, and J Freudenberg. Human motor control and the internal model principle. In *IFAC Symposium on Analysis, Design, and Evaluation of Human-Machine Systems (HMS)*, volume 49, pages 114–119, 2016.
- [60] Stanton A Glantz, Bryan K Slinker, and Torsten B Neilands. *Primer of Applied Regression & Analysis of Variance, Third Edition*. McGraw-Hill Education, April 2016.
- [61] Paula Gomes. Surgical robotics: Reviewing the past, analysing the present, imagining the future. *Robotics and Computer-Integrated Manufacturing*, 27(2):261–266, 2011.
- [62] Tiago Guerreiro, Hugo Nicolau, Joaquim Jorge, and Daniel Gonçalves. Towards accessible touch interfaces. In *Proceedings of the 12th international ACM SIGACCESS conference on Computers and accessibility*, pages 19–26. ACM, 2010.

- [63] Mark Hallett. Plasticity of the human motor cortex and recovery from stroke. *Brain Research Reviews*, 36(2-3):169–174, 2001.
- [64] Faizan Haque, Mathieu Nancel, and Daniel Vogel. Myopoint: Pointing and clicking using forearm mounted electromyography and inertial motion sensors. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 3653–3656. ACM, 2015.
- [65] Sandra G Hart and Lowell E Staveland. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in Psychology*, volume 52, pages 139–183. Elsevier, 1988.
- [66] Christopher J Hasson, Zhaoran Zhang, Masaki O Abe, and Dagmar Sternad. Neuro-motor noise is malleable by amplifying perceived errors. *PLoS Computational Biology*, 12(8):e1005044, 2016.
- [67] Hermie J Hermens, Bart Freriks, Roberto Merletti, Dick Stegeman, Joleen Blok, Günter Rau, Cathy Disselhorst-Klug, and Göran Hägg. European recommendations for surface electromyography. *Roessingh Research and Development*, 8(2):13–54, 1999.
- [68] David J Herzfeld and Reza Shadmehr. Motor variability is not noise, but grist for the learning mill. *Nature Neuroscience*, 17(2):149–150, 2014.
- [69] João P Hespanha. *Noncooperative Game Theory: An Introduction for Engineers and Computer Scientists*. Princeton University Press, June 2017.
- [70] Paula Holland and Alison M Collins. “Whenever I can I push myself to go to work”: A qualitative study of experiences of sickness presenteeism among workers with rheumatoid arthritis. *Disability and Rehabilitation*, 40(4):404–413, 2018.
- [71] Donny Huang, Xiaoyi Zhang, T Scott Saponas, James Fogarty, and Shyamnath Gollakota. Leveraging dual-observable input for fine-grained thumb interaction using forearm EMG. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, pages 523–528. ACM, 2015.

- [72] Kevin Huang, Patrick J Sparto, Sara Kiesler, Asim Smailagic, Jennifer Mankoff, and Dan Siewiorek. A technology probe of wearable in-home computer-assisted physical therapy. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2541–2550, 2014.
- [73] Meghan E Huber, Nikita Kuznetsov, and Dagmar Sternad. Persistence of reduced neuromotor noise in long-term motor skill learning. *Journal of Neurophysiology*, 116(6):2922–2935, 2016.
- [74] Lisa I Iezzoni, Sowmya R Rao, Julie Ressler, Dragana Bolcic-Jankovic, Nicole D Agaronnik, Karen Donelan, Tara Lagu, and Eric G Campbell. Physicians’ perceptions of people with disability and their health care: Study reports the results of a survey of physicians’ perceptions of people with disability. *Health Affairs*, 40(2):297–306, 2021.
- [75] Roland S Johansson, Anna Theorin, Göran Westling, Mikael Andersson, Yukari Ohki, and Lars Nyberg. How a lateralized brain supports symmetrical bimanual tasks. *PLoS Biology*, 4(6):e158, 2006.
- [76] Diana Kilshaw and Marian Annett. Right-and left-hand skill I: Effects of age, sex and hand preference showing superior skill in left-handers. *British Journal of Psychology*, 74(2):253–268, 1983.
- [77] Carolyn Kisner, Lynn Allen Colby, and John Borstad. *Therapeutic Exercise: Foundations and Techniques*. Fa Davis, 2017.
- [78] Donald Knuth. About chronic diseases, 2021.
- [79] John W Krakauer, Alkis M Hadjiosif, Jing Xu, Aaron L Wong, and Adrian M Haith. Motor learning. *Compr Physiol*, 9(2):613–663, 2019.
- [80] Tara Lagu, Nicholas S Hannon, Michael B Rothberg, Annalee S Wells, K Laurie Green, McAllister O Windom, Katherine R Dempsey, Penelope S Pekow, Jill S Avrunin, Aaron Chen, et al. Access to subspecialty care for patients with mobility impairment: a survey. *Annals of Internal Medicine*, 158(6):441–446, 2013.

- [81] Peter Langhorne, Fiona Coupar, and Alex Pollock. Motor recovery after stroke: A systematic review. *The Lancet Neurology*, 8(8):741–754, 2009.
- [82] Vincent A Laurence, Daan M Pool, Herman J Damveld, Marinus René M van Paassen, and Max Mulder. Effects of controlled element dynamics on human feedforward behavior in ramp-tracking tasks. *IEEE Transactions on Cybernetics*, 45(2):253–265, 2014.
- [83] Kate E Laver, Belinda Lange, Stacey George, Judith E Deutsch, Gustavo Saposnik, and Maria Crotty. Virtual reality for stroke rehabilitation. *Cochrane Database of Systematic Reviews*, (11), 2017.
- [84] Yanan Li, Gerolamo Carboni, Franck Gonzalez, Domenico Campolo, and Etienne Burdet. Differential game theory for versatile physical human–robot interaction. *Nature Machine Intelligence*, 1(1):36–43, 2019.
- [85] Yanan Li, Keng Peng Tee, Rui Yan, Wei Liang Chan, and Yan Wu. A framework of human–robot coordination based on game theory and policy iteration. *IEEE Transactions on Robotics*, 32(6):1408–1418, 2016.
- [86] Maureen A. Linden and Karen Milchus. Teleworkers with disabilities: Characteristics and accommodation use. *Work*, 47:473–483, 2014.
- [87] Jiayang Liu, Lin Zhong, Jehan Wickramasuriya, and Venu Vasudevan. uWave: Accelerometer-based personalized gesture recognition and its applications. *Pervasive and Mobile Computing*, 5(6):657–675, 2009.
- [88] Joan Lobo-Prat, Arvid QL Keemink, Arno HA Stienen, Alfred C Schouten, Peter H Veltink, and Bart FJM Koopman. Evaluation of EMG, force and joystick as control interfaces for active arm supports. *Journal of Neuroengineering and Rehabilitation*, 11(1):68, 2014.
- [89] Keith R Lohse, Catherine E Lang, and Lara A Boyd. Is more better? Using metadata to explore dose–response relationships in stroke rehabilitation. *Stroke*, 45(7):2053–2058, 2014.

- [90] Fabien Lotte, Laurent Bougrain, Andrzej Cichocki, Maureen Clerc, Marco Congedo, Alain Rakotomamonjy, and Florian Yger. A review of classification algorithms for EEG-based brain-computer interfaces: A 10 year update. *Journal of Neural Engineering*, 15(3):031005, 2018.
- [91] Kelly Mack, Emma McDonnell, Dhruv Jain, Lucy Lu Wang, Jon Froehlich, and Leah Findlater. What do we mean by “accessibility research”? A literature survey of accessibility papers in CHI and ASSETS from 1994 to 2019. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021.
- [92] I Scott MacKenzie. Fitts’ law as a research and design tool in human-computer interaction. *Human-computer interaction*, 7(1):91–139, 1992.
- [93] Maneeshika M Madduri, Samuel A Burden, and Amy L Orsborn. A game-theoretic model for co-adaptive brain-machine interfaces. In *2021 10th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 327–330. IEEE, 2021.
- [94] Meethu Malu, Pramod Chundury, and Leah Findlater. Exploring accessible smart-watch interactions for people with upper body motor impairments. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, page 488. ACM, 2018.
- [95] Meethu Malu, Pramod Chundury, and Leah Findlater. Motor accessibility of smart-watch touch and bezel input. In *The 21st International ACM SIGACCESS Conference on Computers and Accessibility*, pages 563–565, 2019.
- [96] Meethu Malu and Leah Findlater. Personalized, wearable control of a head-mounted display for users with upper body motor impairments. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 221–230. ACM, 2015.
- [97] Meethu Malu and Leah Findlater. Sharing automatically tracked activity data: implications for therapists and people with mobility impairments. In *Proceedings of the 11th*

- EAI International Conference on Pervasive Computing Technologies for Healthcare*, pages 136–145, 2017.
- [98] N M Mangan, J N Kutz, S L Brunton, and J L Proctor. Model selection for dynamical systems via sparse regression and information criteria. *Proceedings of the Royal Society A: Mathematical, Physical, and Engineering Sciences*, 473(2204):20170009, August 2017.
- [99] Jennifer Mankoff, Gillian R Hayes, and Devva Kasnitz. Disability studies as a source of critical inquiry for the field of assistive technology. In *Proceedings of the 12th international ACM SIGACCESS conference on Computers and accessibility*, pages 3–10, 2010.
- [100] Laura Marchal-Crespo and David J Reinkensmeyer. Review of control strategies for robotic movement training after neurologic injury. *Journal of Neuroengineering and Rehabilitation*, 6(1):1–15, 2009.
- [101] Liam Mason, Kathrin Gerling, Patrick Dickinson, and Antonella De Angeli. Design goals for playful technology to support physical activity among wheelchair users. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2019.
- [102] Marta Matamala-Gomez, Marta Maisto, Jessica Isbely Montana, Petar Aleksandrov Mavrodiev, Francesca Baglio, Federica Rossetto, Fabrizia Mantovani, Giuseppe Riva, and Olivia Realdon. The role of engagement in teleneurorehabilitation: A systematic review. *Frontiers in Neurology*, 11:354, 2020.
- [103] O Mayr. The origins of feedback control, 1970.
- [104] David McNaughton, Tracy Rackensperger, Dana Dorn, and Natasha Wilson. “Home is at work and work is at home”: Telework and individuals who use augmentative and alternative communication. *Work*, 48:117–126, 2014.
- [105] Duane T McRuer and Henry R Jex. A review of quasi-linear pilot models. *IEEE transactions on Human Factors in Electronics*, (3):231–249, 1967.

- [106] Karen Devereaux Melillo, May Futrell, Eileen Williamson, Claire Chamberlain, Anne Marie Bourque, Marilyn MacDonnell, and Julie P Phaneuf. Perceptions of physical fitness and exercise activity among older adults. *Journal of Advanced Nursing*, 23(3):542–547, 1996.
- [107] Nathan W. Moon, Maureen A. Linden, John C. Bricout, and Paul M.A. Baker. Telework rationale and implementation for people with disabilities: Considerations for employer policymaking. *Work*, 48:105–115, 2014.
- [108] Martez Mott, John Tang, Shaun Kane, Edward Cutrell, and Meredith Ringel Morris. “I just went into it assuming that I wouldn’t be able to have the full experience” Understanding the accessibility of virtual reality for people with limited mobility. In *The 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, pages 1–13, 2020.
- [109] Martez E Mott, Radu-Daniel Vatavu, Shaun K Kane, and Jacob O Wobbrock. Smart touch: Improving touch accuracy for people with motor impairments with template matching. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 1934–1946. ACM, 2016.
- [110] S Alireza Seyyed Mousavi, Faina Matveeva, Xingye Zhang, T Michael Seigler, and Jesse B Hoagg. The impact of command-following task on human-in-the-loop control behavior. *IEEE Transactions on Cybernetics*, 2020.
- [111] Max Mulder, Daan M Pool, David A Abbink, Erwin R Boer, Peter MT Zaal, Frank M Drop, Kasper van der El, and Marinus M van Paassen. Manual control cybernetics: State-of-the-art and current trends. *IEEE Transactions on Human-Machine Systems*, 48(5):468–485, 2017.
- [112] Jan Saputra Müller, Carmen Vidaurre, Martijn Schreuder, Frank C Meinecke, Paul Von Büнау, and Klaus-Robert Müller. A mathematical model for the two-learners problem. *Journal of Neural Engineering*, 14(3):036005, 2017.

- [113] F M Nieuwenhuizen, P M T Zaal, M Mulder, M M Van Paassen, and J A Mulder. Modeling human multichannel perception and control using linear Time-Invariant models. *Journal of Guidance, Control, and Dynamics*, 31(4):999–1013, July 2008.
- [114] Stefanos Nikolaidis, Swaprava Nath, Ariel D Procaccia, and Siddhartha Srinivasa. Game-Theoretic modeling of human adaptation in Human-Robot collaboration. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, pages 323–331, 2017.
- [115] Richard C Oldfield. The assessment and analysis of handedness: The Edinburgh inventory. *Neuropsychologia*, 9(1):97–113, 1971.
- [116] Mario Olivari, Frank M Nieuwenhuizen, Joost Venrooij, Heinrich H Bühlhoff, and Lorenzo Pollini. Methods for multiloop identification of visual and neuromuscular pilot responses. *IEEE Transactions on Cybernetics*, 45(12):2780–2791, 2015.
- [117] World Health Organization. The world health report 2001: Mental health: New understanding, new hope. 2001.
- [118] World Health Organization. *International Classification of Functioning, Disability, and Health: Children & Youth Version: ICF-CY*. World Health Organization, 2007.
- [119] World Health Organization et al. *Telemedicine: Opportunities and Developments in Member States. Report on the Second Global Survey on eHealth*. World Health Organization, 2010.
- [120] Amy L Orsborn, Siddharth Dangi, Helene G Moorman, and Jose M Carmena. Closed-loop decoder adaptation on intermediate time-scales facilitates rapid BMI performance improvements independent of decoder initialization conditions. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 20(4):468–477, 2012.
- [121] Harsh Kupwade Patil and Ravi Seshadri. Big data security and privacy issues in healthcare. In *2014 IEEE International Congress on Big Data*, pages 762–765. IEEE, 2014.

- [122] Brendan M Patterson, Jeffrey T Spang, Reid W Draeger, Erik C Olsson, Robert A Creighton, and Ganesh V Kamath. Access to outpatient care for adult rotator cuff patients with private insurance versus medicaid in North Carolina. *Journal of Shoulder and Elbow Surgery*, 22(12):1623–1627, 2013.
- [123] Thamizhisai Periyaswamy and Mahendran Balasubramanian. Ambulatory cardiac bio-signals: From mirage to clinical reality through a decade of progress. *International Journal of Medical Informatics*, 130:103928–103928, 2019.
- [124] Huy Phan. A recurrent neural network for hand gesture recognition based on accelerometer data. In *Engineering in Medicine and Biology Society (EMBC), Annual International Conference of the IEEE*. IEEE, 2019.
- [125] Jennifer R Pharr, Tamara James, and Yeu-Li Yeung. Accessibility and accommodations for patients with mobility disabilities in a large healthcare system: How are we doing? *Disability and Health Journal*, 12(4):679–684, 2019.
- [126] Charles L Phillips, John M Parr, and Eve A Riskin. *Signals, Systems, and Transforms*. Prentice Hall, 2003.
- [127] Ruth Pinder. Sick-but-fit or fit-but-sick? Ambiguity and identity at the workplace. *Exploring the Divide*, pages 135–156, 1996.
- [128] Ruth Pinder. Zones of danger, zones of safety: Disabled people’s negotiations around sickness and the sick record. In *Health and Work*, pages 161–179. Springer, 1999.
- [129] Hugo F Posada-Quintero, Ryan T Rood, Ken Burnham, John Pennace, and Ki H Chon. Assessment of carbon/salt/adhesive electrodes for surface electromyography measurements. *IEEE journal of Translational Engineering in Health and Medicine*, 4:1–9, 2016.
- [130] Mark Priestley, Mairian Corker, and Nick Watson. Unfinished business: Disabled children and disability identity. 1999.

- [131] L. J. Ratliff, S. A. Burden, and S. S. Sastry. Characterization and computation of local Nash equilibria in continuous games. In *2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, October 2013.
- [132] Nornadiah Mohd Razali, Yap Bee Wah, et al. Power comparisons of Shapiro-Wilk, Kolmogorov-Smirnov, Lilliefors and Anderson-Darling tests. *Journal of Statistical Modeling and Analytics*, 2(1):21–33, 2011.
- [133] Kyle Rector, Cynthia L Bennett, and Julie A Kientz. Eyes-free yoga: An exergame using depth cameras for blind & low vision exercise. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility*, pages 1–8, 2013.
- [134] Paul Rinne, Michael Mace, Tagore Nakornchai, Karl Zimmerman, Susannah Fayer, Pankaj Sharma, Jean-Luc Liardon, Etienne Burdet, and Paul Bentley. Democratizing neurorehabilitation: How accessible are low-cost mobile-gaming technologies for self-rehabilitation of arm disability in stroke? *PloS One*, 11(10):e0163413, 2016.
- [135] Miranda J Rogers, Ian Penrose, Emily J Curry, Anthony DeGiacomo, and Xinning Li. Medicaid health insurance status limits patient accessibility to rehabilitation services following ACL reconstruction surgery. *Orthopaedic Journal of Sports Medicine*, 6(4):2325967118763353, 2018.
- [136] E Roth, K Zhuang, S A Stamper, E S Fortune, and N J Cowan. Stimulus predictability mediates a switch in locomotor smooth pursuit performance for *Eigenmania virescens*. *Journal of Experimental Biology*, 214(7):1170–1180, 2011.
- [137] Eatai Roth, Darrin Howell, Cydney Beckwith, and Samuel A Burden. Toward experimental validation of a model for human sensorimotor learning and control in teleoperation. In *Micro-and Nanotechnology Sensors, Systems, and Applications IX*, volume 10194, page 101941X. International Society for Optics and Photonics, 2017.
- [138] Geert M Rutten, Saskia Degen, Erik J Hendriks, Jozé C Braspenning, Janneke Harting, and Rob A Oostendorp. Adherence to clinical practice guidelines for low back

- pain in physical therapy: Do patients benefit? *Physical Therapy*, 90(8):1111–1122, 2010.
- [139] Robert L Sainburg and D Kalakanis. Differences in control of limb dynamics during dominant and nondominant arm reaching. *Journal of Neurophysiology*, 83(5):2661–2675, 2000.
- [140] T Scott Saponas, Desney S Tan, Dan Morris, and Ravin Balakrishnan. Demonstrating the feasibility of using forearm electromyography for muscle-computer interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 515–524. ACM, 2008.
- [141] T Scott Saponas, Desney S Tan, Dan Morris, Ravin Balakrishnan, Jim Turner, and James A Landay. Enabling always-available input with muscle-computer interfaces. In *Proceedings of the 22nd annual ACM symposium on User interface software and technology*, pages 167–176. ACM, 2009.
- [142] Fred S Sarfo, Sheila Adamu, Dominic Awuah, Osei Sarfo-Kantanka, and Bruce Ovbi-agele. Potential role of tele-rehabilitation to address barriers to implementation of physical therapy among West African stroke survivors: A cross-sectional survey. *Journal of the Neurological Sciences*, 381:203–208, 2017.
- [143] Christopher N Schabowsky, Joseph M Hidler, and Peter S Lum. Greater reliance on impedance control in the nondominant arm compared with the dominant arm when adapting to a novel dynamic environment. *Experimental brain research*, 182(4):567–577, 2007.
- [144] Ruth B Schneider and Kevin M Biglan. The promise of telemedicine for chronic neurological disorders: The example of parkinson’s disease. *The Lancet Neurology*, 16(7):541–551, 2017.
- [145] RN Scott and PA Parker. Myoelectric prostheses: State of the art. *Journal of Medical Engineering & Technology*, 12(4):143–151, 1988.

- [146] Stephen H Scott. Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews. Neuroscience*, 5(7):532–546, 2004.
- [147] RD Seidler, DC Noll, and G Thiers. Feedforward and feedback processes in motor control. *Neuroimage*, 22(4):1775–1783, 2004.
- [148] Jessica C Selinger, Jeremy D Wong, Surabhi N Simha, and J Maxwell Donelan. How humans initiate energy optimization and converge on their optimal gaits. *Journal of Experimental Biology*, 222(19):jeb198234, 2019.
- [149] Tom Shakespeare. *Disability Rights and Wrongs Revisited*. Routledge, 2013.
- [150] Tom Shakespeare and Nick Watson. Beyond models: Understanding the complexity of disabled people’s lives. In *New Directions in the Sociology of Chronic and Disabling Conditions*, pages 57–76. Springer, 2010.
- [151] Steven R Shaw and Paul C McCabe. Hospital-to-school transition for children with chronic illness: Meeting the new challenges of an evolving health care system. *Psychology in the Schools*, 45(1):74–87, 2008.
- [152] Shigeki Shibata, Qi Fu, Tiffany B Bivens, Jeffrey L Hastings, Wade Wang, and Benjamin D Levine. Short-term exercise training improves the cardiovascular response to exercise in the postural orthostatic tachycardia syndrome. *The Journal of Physiology*, 590(15):3495–3505, 2012.
- [153] Benjamin R Shuman, Marije Goudriaan, Kaat Desloovere, Michael H Schwartz, and Katherine M Steele. Muscle synergies demonstrate only minimal changes after treatment in cerebral palsy. *Journal of Neuroengineering and Rehabilitation*, 16(1):46, 2019.
- [154] Maurice A Smith, Jason Brandt, and Reza Shadmehr. Motor disorder in huntington’s disease begins as a dysfunction in error feedback control. *Nature*, 403(6769):544–549, 2000.

- [155] Maurice A Smith and Reza Shadmehr. Intact ability to learn internal models of arm dynamics in huntington's disease but not cerebellar degeneration. *Journal of Neurophysiology*, 93(5):2809–2821, 2005.
- [156] Katie Thomson, Alex Pollock, Carol Bugge, and Marian Brady. Commercial gaming devices for stroke upper limb rehabilitation: A systematic review. *International Journal of Stroke*, 9(4):479–488, 2014.
- [157] E Todorov and M I Jordan. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11):1226–1235, 2002.
- [158] Shari Trewin and Helen Pain. Keyboard and mouse errors due to motor disabilities. *International Journal of Human-Computer Studies*, 50(2):109–144, 1999.
- [159] Shari Trewin, Cal Swart, and Donna Pettick. Physical accessibility of touchscreen smartphones. In *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility*, page 19. ACM, 2013.
- [160] Andrea Turolla, Giacomo Rossettini, Antonello Viceconti, Alvisa Palese, and Tommaso Geri. Musculoskeletal physical therapy during the COVID-19 pandemic: Is telerehabilitation the answer? *Physical Therapy*, 100(8):1260–1264, 2020.
- [161] Ovidiu-Ciprian Ungurean, Radu-Daniel Vatavu, Luis A Leiva, and Réjean Plamondon. Gesture input for users with motor impairments on touchscreens: Empirical results based on the kinematic theory. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems*, page LBW537. ACM, 2018.
- [162] Kasper van der El, Sharon Padmos, Daan M Pool, Marinus M van Paassen, and Max Mulder. Effects of preview time in manual tracking tasks. *IEEE Transactions on Human-Machine Systems*, 48(5):486–495, 2018.
- [163] Kasper van der El, Daan M Pool, Herman J Damveld, Marinus René M van Paassen, and Max Mulder. An empirical human controller model for preview tracking tasks. *IEEE Transactions on Cybernetics*, 46(11):2609–2621, 2015.

- [164] Kasper van der El, Daan M Pool, Marinus René M van Paassen, and Max Mulder. Effects of preview on human control behavior in tracking tasks with various controlled elements. *IEEE Transactions on Cybernetics*, 48(4):1242–1252, 2017.
- [165] Arturo Vega-González and Malcolm H Granat. Continuous monitoring of upper-limb activity in a free-living environment. *Archives of physical medicine and rehabilitation*, 86(3):541–548, 2005.
- [166] Joost Venrooij, Mark Mulder, David A Abbink, Marinus M Van Paassen, Frans CT Van Der Helm, Heinrich H Bühlhoff, and Max Mulder. A new view on biodynamic feedthrough analysis: Unifying the effects on forces and positions. *IEEE Transactions on Cybernetics*, 43(1):129–142, 2012.
- [167] Etienne Vermeire, Hilary Hearnshaw, Paul Van Royen, and Joke Denekens. Patient adherence to treatment: Three decades of research. a comprehensive review. *Journal of Clinical Pharmacy and Therapeutics*, 26(5):331–342, 2001.
- [168] Lisa M Vizer, Jordan Eschler, Bon Mi Koo, James Ralston, Wanda Pratt, and Sean Munson. “It’s not just technology, it’s people”: Constructing a conceptual model of shared health informatics for tracking in chronic illness management. *Journal of Medical Internet Research*, 21(4):e10830, 2019.
- [169] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.
- [170] Saskia E Wagenaar, Daan M Pool, Herman J Damveld, Marinus M van Paassen, and Max Mulder. Estimation of nonlinear contributions in human controller frequency response functions. In *IEEE Conference on Systems, Man, and Cybernetics (SMC)*, pages 3434–3439, 2018.
- [171] Rahul B Warriar and Santosh Devasia. Inferring intent for novice human-in-the-loop iterative learning control. *IEEE Transactions on Control Systems Technology*, 25(5):1698–1710, 2016.

- [172] Rahul B Warriier and Santosh Devasia. Iterative learning from novice human demonstrations for output tracking. *IEEE Transactions on Human-Machine Systems*, 46(4):510–521, 2016.
- [173] Susan Wendell. Unhealthy disabled: Treating chronic illnesses as disabilities. *Hypatia*, 16(4):17–33, 2001.
- [174] Francis R Willett, Donald T Avansino, Leigh R Hochberg, Jaimie M Henderson, and Krishna V Shenoy. High-performance brain-to-text communication via handwriting. *Nature*, 593(7858):249–254, 2021.
- [175] Simon J Williams. Is anybody there? Critical realism, chronic illness and the disability debate. *Sociology of Health & Illness*, 21(6):797–819, 1999.
- [176] Carolee J Winstein, Scott T Grafton, and Patricia S Pohl. Motor task difficulty and brain activity: Investigation of goal-directed reciprocal aiming using positron emission tomography. *Journal of Neurophysiology*, 77(3):1581–1594, 1997.
- [177] Jacob O Wobbrock, James Fogarty, Shih-Yen Sean Liu, Shunichi Kimuro, and Susumu Harada. The angle mouse: Target-agnostic dynamic gain adjustment based on angular deviation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1401–1410. ACM, 2009.
- [178] Jacob O Wobbrock, Shaun K Kane, Krzysztof Z Gajos, Susumu Harada, and Jon Froehlich. Ability-based design: Concept, principles and examples. *ACM Transactions on Accessible Computing (TACCESS)*, 3(3):9, 2011.
- [179] Jeremy D Wong, Jessica C Selinger, and J Maxwell Donelan. Is natural variability in gait sufficient to initiate spontaneous energy optimization in human walking? *Journal of Neurophysiology*, 121(5):1848–1855, 2019.
- [180] Howard G Wu, Yohsuke R Miyamoto, Luis Nicolas Gonzalez Castro, Bence P Ölveczky, and Maurice A Smith. Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nature Neuroscience*, 17(2):312–321, 2014.

- [181] Momona Yamagami, Darrin Howell, Eatai Roth, and Samuel A Burden. Contributions of feedforward and feedback control in a manual trajectory-tracking task. In *IFAC Conference on Cyber-Physical-Human Systems (CPHS)*, volume 51, pages 61–66. Elsevier, 2018.
- [182] Momona Yamagami, Darrin Howell, Eatai Roth, and Samuel A Burden. Contributions of feedforward and feedback control in a manual trajectory-tracking task. In *Proceedings of the 2018 Conference on Cyber-Physical Human Systems*, volume 51, pages 61–66. IFAC, 2019.
- [183] Momona Yamagami, Kelly Mack, Jennifer Mankoff, and Katherine M Steele. “I’m just overwhelmed”: Investigating physical therapy accessibility and technology interventions for people with disabilities and/or chronic conditions. *arXiv preprint arXiv:2202.02281*, 2022.
- [184] Momona Yamagami, Maneeshika Madduri, Benjamin Chasnov, Hsiao-Yang Chou, Lauren N Peterson, and Samuel A Burden. Co-adaptation for human-in-the-loop control systems. *In Preparation*.
- [185] Momona Yamagami, Keshia Peters, Ivana Milovanovic, Irene Kuang, Zeyu Yang, Nanshu Lu, and Katherine Steele. Assessment of dry epidermal electrodes for long-term electromyography measurements. *Sensors*, 18(4):1269, 2018.
- [186] Momona Yamagami, Lauren N Peterson, Darrin Howell, Eatai Roth, and Samuel A Burden. Effect of handedness on learned controllers and sensorimotor noise during trajectory-tracking. *IEEE Transactions on Cybernetics*, 2021.
- [187] Momona Yamagami, Katherine M Steele, and Samuel A Burden. Decoding intent with control theory: Comparing muscle versus manual interface performance. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2020.
- [188] Kaiwen Yang, Luke Nicolini, Irene Kuang, Nanshu Lu, and Dragan Djurdjanovic.

Long-term modeling and monitoring of neuromusculoskeletal system performance using tattoo-like EMG sensors.

- [189] Bo Yu, R Brent Gillespie, James S Freudenberg, and Jeffrey A Cook. Human control strategies in pursuit tracking with a disturbance input. In *53rd IEEE Conference on Decision and Control*, pages 3795–3800. IEEE, 2014.
- [190] Rui Zhang, Fali Li, Tao Zhang, Dezhong Yao, and Peng Xu. Subject inefficiency phenomenon of motor imagery brain-computer interface: Influence factors and potential solutions. *Brain Science Advances*, 6, September 2020. Publisher: SAGE Publications Ltd.
- [191] Xingye Zhang, T Michael Seigler, and Jesse B Hoagg. The impact of nonminimum-phase zeros on human-in-the-loop control systems. *IEEE Transactions on Cybernetics*, 2020.
- [192] Xingye Zhang, Shaoqian Wang, Jesse B Hoagg, and T Michael Seigler. The roles of feedback and feedforward as humans learn to control unknown dynamic systems. *IEEE Transactions on Cybernetics*, 48(2):543–555, 2017.
- [193] Xu Zhang, Xiang Chen, Yun Li, Vuokko Lantz, Kongqiao Wang, and Jihai Yang. A framework for hand gesture recognition based on accelerometer and EMG sensors. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 41(6):1064–1076, 2011.
- [194] Yu Zhong, Astrid Weber, Casey Burkhardt, Phil Weaver, and Jeffrey P Bigham. Enhancing android accessibility for users with hand tremor by reducing fine pointing and steady tapping. In *Proceedings of the 12th Web for All Conference*, page 29. ACM, 2015.
- [195] Kemin Zhou, John Comstock Doyle, and Keith Glover. *Robust and Optimal Control*, volume 40. Prentice Hall, 1996.
- [196] Katie Z Zhuang, Nicolas Sommer, Vincent Mendez, Saurav Aryan, Emanuele Formento, Edoardo D’Anna, Fiorenzo Artoni, Francesco Petrini, Giuseppe Granata, Gio-

vanni Cannaviello, et al. Shared human–robot proportional control of a dexterous myoelectric prosthesis. *Nature Machine Intelligence*, 1(9):400–411, 2019.

- [197] Amanda M Zimmet, Di Cao, Amy J Bastian, and Noah J Cowan. Cerebellar patients have intact feedback control that can be leveraged to improve reaching. *eLife*, 9, October 2020.

Appendix A

APPENDIX

A.1 Aim 1 Appendix

A.1.1 Computing User Input Prediction

To evaluate Hypothesis 2, where we use data from the **linearity** experiment to compare the predictive accuracy of B -only and $B + F$ models, we partitioned data from the conditions in TABLE 2.2 into disjoint *train* and *test* subsets as follows. First, we compute B at each stimulated frequency for each participant by averaging (2.6a) across the ten $(0, d)$ trials. Subsequently, for the $B + F$ model, we use each participant's estimated B to compute their F at each stimulated frequency by averaging (2.6b) across the ten $(r, 0)$ trials. Thus, the *train* dataset consisted of ten $(0, d)$ trials and ten $(r, 0)$ trials. We applied (2.5) to predict the (frequency-domain) user response for the $B + F$ model in the last ten (r, d) trials using the user's B and F estimates.

For the B -only model, we set $F = 0$ in (2.5) so that

$$\hat{u}(\omega) = \underbrace{\frac{\hat{B}(\omega)}{1 + \hat{B}(\omega)\widehat{M}(\omega)}}_{\widehat{T}_{ur}(\omega)} \hat{r}(\omega) + \underbrace{\frac{-\hat{B}(\omega)\widehat{M}(\omega)}{1 + \hat{B}(\omega)\widehat{M}(\omega)}}_{\widehat{T}_{ud}(\omega)} \hat{d}(\omega). \quad (\text{A.1})$$

We can again compute B at each stimulated frequency for the ten $(0, d)$ trials by averaging (2.6a). Then, we used (A.2) to compute B during the ten $(r, 0)$ trials, and took the average to obtain an estimate of B .

$$\hat{B}(\omega) = \frac{\widehat{T}_{ur}(\omega)}{1 + \widehat{T}_{ur}(\omega)\widehat{M}(\omega)} \quad (\text{A.2})$$

We applied (A.1) to predict user response for the B -only model in the last ten (r, d) trials using the user's B estimates.

By performing this analysis at each stimulated frequency for each of the ten (r, d) trials for each participant, we obtained a predicted user response \hat{u}_{pred} which we compared to the measured user input \hat{u}_{meas} . Thus, the *test* dataset consisted of ten (r, d) trials.

A.1.2 Computing R^2 for User Input Prediction

We used the coefficient of determination R^2 [60, Eqn. (3.9)] to assess prediction accuracy of the user input \hat{u} at each frequency. For each set of \hat{u}_{pred} and \hat{u}_{meas} obtained from the (r, d) trials (see Appendix A.1.1), we computed an R^2 value at each stimulated frequency for each trial i with the equation,

$$R^2(\omega) = 1 - \frac{\sum_i |\hat{u}_{meas,i}(\omega) - \hat{u}_{pred,i}(\omega)|^2}{\sum_i |\hat{u}_{meas,i}(\omega) - \tilde{u}_{meas}(\omega)|^2}. \quad (\text{A.3})$$

We defined $\tilde{u}_{meas}(\omega)$ as the average user input of the measured input in the ten (r, d) trials that were used to make a prediction. We computed each of these quantities in the frequency domain using complex numbers, so we computed an R^2 value at each stimulated frequency that represents deviation in both magnitude and phase. If the predictions perfectly match the measured user inputs, then $R^2 = 1$. If the predictions do no better than predicting the mean of the observed user inputs, then $R^2 = 0$. If the predictions are worse than predicting the mean of the observed user inputs, then $R^2 < 0$.

A.1.3 Wilcoxon Signed-Rank Test

The Wilcoxon signed-rank test is a non-parametric paired t -test for data that is not normally distributed [30, Sec. 5.7]. Parametric statistical tests like the paired t -tests come with several assumptions that must be verified, such as that the data must be normally distributed [30]. We determined that the assumption of normality does not hold for our dataset using the Shapiro-Wilk test ($\alpha = 0.05$) [132]. Therefore, we chose to use the Wilcoxon signed-rank test.

The Wilcoxon signed-rank test compares whether the differences between two conditions for a single group of N individuals have statistically different medians or not. The test does

this by ranking the absolute difference between the two conditions for each participant, with 1 being assigned to the individual with the smallest difference, and N being assigned to the individual with the largest difference. Then, the rank for each individual is multiplied by 1 or -1 depending on whether the difference between the two conditions were positive or negative. The test statistic Z is computed as the sum of the signed ranks, and the p -value can be defined from the computed Z value [30, Sec. 5.7].

Because the test compares differences between all samples from two datasets, it is generally not possible to determine whether there is a statistically significant difference between the two datasets using only the median and interquartile statistics represented in box plots.