Co-adaptive Human-Machine Interaction with Multimodal Inputs

Amber Hsiao-Yang Chou

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

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Samuel A. Burden, Chair

Amy L. Orsborn

Yiyue Luo

Program Authorized to Offer Degree:

Department of Electrical and Computer Engineering

University of Washington

Abstract

Co-adaptive Human-Machine Interaction with Multimodal Inputs

Amber Hsiao-Yang Chou

Chair of the Supervisory Committee:

Chair Samuel A. Burden

Department of Electrical and Computer Engineering

Human-machine interfaces record biosignals generated by humans – such as neural (brain), visual (gaze), or limb movement – and decode them into control commands for computational, robotic, and assistive systems. However, the widespread clinical and consumer adoption of biosignal-based human-machine interfaces remains limited. Conventional systems struggle to generalize across diverse user populations, leading to poor usability and high abandonment rates, particularly among individuals with motor impairments. There is a need to seamlessly tailor these interfaces to diverse users and contexts, providing an "out-of-the-box" solution that requires no expert setup. This problem has traditionally been challenging because biosignals have large variability, which leads to extensive interface calibration across people and within the same user over time. This dissertation builds on prior research to address this challenge, laying the groundwork for interfaces personalized to individual needs and abilities. Specifically, I combined theoretical frameworks from control theory with data-driven and neuroengineering methods to: **(1) model humans** as control systems interacting with machines to investigate user strategies and adaptations in multimodal interfaces, and use these insights to **(2) design intelligent interfaces that co-adapt in real time** — that is, adapt in response to users' ongoing adaptation.

This dissertation demonstrates how I integrated systems with multimodal inputs, including surface electromyography (EMG), eye tracking, and manual joystick, and used experimental data from human participants to model user behaviors. I then developed machine learning algorithms to engineer interfaces that automatically adapt to individual usage, allowing continuous calibration with multimodal inputs. This setup provided a new paradigm to investigate the emerging co-adaptation patterns between human and intelligent interfaces. In summary, this work provides new theoretical, computational, and experimental tools for personalizing adaptive, multimodal interfaces to diverse users and tasks. Moving forward, I

plan to implement the frameworks in rehabilitation programs, such as robot-guided therapy, to deliver personalized interventions for people with motor impairments. Together, this work will enhance the usability, generalizability, and clinical translation of biosignal-based human-machine interfaces.

TABLE OF CONTENTS

LIST OF FIGURES

Figure Number Page

ACKNOWLEDGMENTS

I am deeply grateful to all my mentors, especially my advisor Sam. Thank you, Sam, for being such a warm and supportive advisor, for understanding my struggles and worries, and for always making time whenever I needed guidance or feedback. Thank you to Amy, Momona, Maneeshika, and Sijia for your mentorship and collaboration. Your guidance extended well beyond our projects, and your career advice has meant a lot to me. Thank you to my labmates, especially Emmy and Jason, and everyone who has been part of the BioRobotics Lab for making the space warm and welcome. I also want to thank the undergraduate students, Liya, Lauren, and Andrew, for their help with our studies and data collection, and to all the participants — total of 106 people includes pilot subjects — who gave their time to our experiments. I want to thank my partner Eddie, for always being there for me and cheer me up during all the frustrating and stressful moments throughout my PhD. Doing a PhD was hard, pursuing our PhDs together was not easy either, but you were the reason why I could get through it and why I never gave up. Finally, my parents, my family, my friends, and my cats Mango and Nemo, thank you all for keeping me grounded and loved.

Chapter 1

INTRODUCTION

1 Biosignal-based human-machine interfaces: advantages and current challenges

Biosignals generated by humans – including neural, visual, audio, movement data – are increasingly integrated into consumer and medical devices. They drive a wide range of applications, such as everyday wearables, teleoperation (Farry et al., 1996; Morelli et al., 2016; Murphy, 2004), human-computer interfaces and VR/AR (Bilmes et al., 2005; Phelan et al., 2021), brain-computer interfaces (Collinger et al., 2013; Willett et al., 2023), upper-limb (Farina et al., 2023; Fougner et al., 2012; Kiguchi and Hayashi, 2012; Onose et al., 2012; Smith et al., 2016) and lower-limb prostheses (Song et al., 2024), shared driving (Gnatzig et al., 2012; Mars et al., 2014), neurorehabilitation (Hung et al., 2021; McDonald et al., 2020; Mugler et al., 2019; Rizzoglio et al., 2019). In these systems, an *interface* uses sensors to record biosignals and algorithms to decode them into control commands for an external machine. It also provides state feedback to the user, allowing them to adjust their actions to achieve the task goal. Here, I define *machines* broadly as external physical or virtual systems with dynamic properties, ranging from simple simulated cursors to multi-degree-of-freedom robotic devices.

Biosignals provide always-available (continuous), high-bandwidth (temporal richness) and high-dimensional (spatial richness) data (Drost et al., 2006; Saponas et al., 2009; Zheng et al., 2022), thus facilitating more degrees of freedom and more intuitive control than traditional manual interfaces like joysticks, mice, and keyboards (Lobo-Prat et al., 2014; Yamagami et al., 2020). While manual interfaces or kinematic interfaces (e.g. motion capture, inertial measurement units, IMU) remain the most familiar to users, they rely on physical mobility and are often inaccessible to individuals with neurological or acute injuries (Findlater et al., 2017). In contrast, modalities that do not rely on limb movement or joint kinematics, such as gaze, speech, or neural data, offer an accessible alternative for restoring motor functions across diverse populations (Bilmes et al., 2005; Huang et al., 2015b; Manero et al., 2022; Yamagami et al., 2020). For instance, movement intent can still be inferred from neural signals in cases of limb amputation, severe motor impairments, or limited dexterity (Borgestig et al., 2016; Elliott et al., 2019; Farina et al., 2023; Fleming et al., 2021). However, while assistive technologies utilizing gaze or speech are relatively mature, the widespread adoption of neural interfaces remains limited, particularly outside of laboratory settings.

1.1 Current challenges of neuromotor interfaces

Neuromotor interfaces record neuronal activity from the brain or periphery, such as intracortical recording, electromyography (EMG), or electroencephalography (EEG). They then use a computational model, or *decoder*, to translate these signals into control commands that drive robotic joint torque (Hochberg et al., 2012; Song et al., 2024), gesture-controlled systems (Saponas et al., 2008; Kaifosh et al., 2025), or brain-controlled cursor movements (Hochberg et al., 2006; Kim et al., 2008). Users receive sensory feedback, which may include visual, haptic, or proprioceptive cues. This feedback establishes a *closed-loop* interaction, placing the interface and the machine in-the-loop with the user.

The limited adoption of neuromotor interfaces stems not only from the inherent risks of invasive neural recordings; even non-invasive technologies, such as surface electromyography (sEMG), face significant translational challenges including low usability and high abandonment rates among individuals with motor impairments (Farina et al., 2014; Resnik et al., 2020). These challenges are primarily driven by the factors listed below. I will use sEMG interfaces as a primary example in discussing the challenges, but it is important to emphasize that these challenges are not unique to myoelectric interfaces; rather, they persist across a wide range of neuromotor modalities.

(1) Nonstationarity: Neural signals have high signal variability across sessions, and have high sensitivity to perturbations (Artemiadis, 2012; Farina et al., 2014). For instance, muscle signals measured from the skin change over time due to fatigue, sweating, and change in electrode positions (Artemiadis and Kyriakopoulos, 2008; Fougner et al., 2011; Jiang et al., 2012; Yousif et al., 2019). In addition, humans are not a stationary data source. Behavioral patterns dynamically adapt as people (or animals) receive sensory feedback while learning new interfaces and motor tasks, which could lead to day-to-day variability in performance (Albert and Shadmehr, 2016; Green and Kalaska, 2011; Jackson and Fetz, 2011; Shenoy and Carmena, 2014; Taylor et al., 2002). In healthcare and rehabilitation, patients' health and motor abilities also change during the recovery process (Levin and Demers, 2021; Navarro-Barrientos et al., 2011; Yao et al., 2017). Therefore, these changes often lead to a mismatch between the performance of conventional models trained in *open-loop* (user not in the loop; *offline*) and the performance during *closed-loop* (user in the loop; *online*) operation (Chase et al., 2009; Hahne et al., 2017; Jiang et al., 2012; Koyama et al., 2010; Shenoy and Carmena, 2014), requiring extensive and frequent re-calibration of neuromotor interfaces.

(2) Limited generalization: While training a generic interface on large datasets can yield high performance during closed-loop operation (Pan et al., 2018; Cote-Allard et al., 2019; Kaifosh et al., 2025), this paradigm relies on expensive offline training data spanning thousands of participants, making it difficult to translate to diverse tasks or clinical populations. Specifically, offline training frameworks often restrict users to a rigid, predefined set of tasks and movements established during calibration. For example, sEMG interfaces are commonly trained with specific gestures (kinematic labels (McIntosh et al., 2016; Muceli and Farina, 2012; Saponas et al., 2008)) or contractions (kinetic labels (Jiang et al., 2009)) for specific operations like reaching or selecting. Such constraints limit online utility, as real-world human interactions with machines and computers are highly variable (Mack et al., 2022; Pandarinath et al., 2017; Yamagami et al., 2023b). A model trained on a reaching task may have varied performance when transitioned to different contexts (Shenoy and Carmena, 2014), such as typing or grasping. This generalization challenge is further amplified in individuals with motor impairments. Because generic models are predominantly trained on healthy cohorts, they scale poorly to clinical use cases where patient data is scarce and underlying pathologies are highly varied.

(3) Subject variability in learning: Similar to how people learn sensorimotor skills, users naturally learn when interacting with neuromotor interfaces (Albert and Shadmehr, 2016; Carmena et al., 2003; He et al., 2015; Taylor et al., 2002). One could theoretically deploy a reasonable generic interface and rely entirely on user adaptation to bridge any performance deficits. However, neuromotor interfaces exhibit variable baseline performance across different users (Leeb et al., 2013; Pandarinath et al., 2017) and are often time-consuming or difficult to learn, frequently resulting in a high percentage of subject inefficiency (Guger et al., 2003; Zhang et al., 2020a). This highlights a need for personalized interfaces that are easily learnable and accommodate individual learning patterns, which requires a shift towards investigating how user and interface interact (Perdikis and Millan, 2020).

In rehabilitation contexts specifically, a primary objective is to promote motor learning, by providing *resistive* training with devices (Emken and Reinkensmeyer, 2005). Existing research on resistive robotic devices mostly relies on non-adaptive control strategies (McDonald et al., 2020), or adaptive systems trained on offline human data (Peña et al., 2019; Perez-Ibarra et al., 2019; Yao et al., 2018). However, functional recovery outcomes vary widely and often plateau after extensive training (Page et al., 2004; Schmidt et al., 2018). Introducing an intelligent interface in the loop — that can dynamically optimize for each patient's evolving recovery — may further increase patient engagement and facilitate learning (Zhang et al., 2023; Chemuturi et al., 2013; Furuya et al., 2025; Madduri et al., 2023).

1.2 The vision for neuromotor interfaces

The overarching vision for these neuromotor interfaces is twofold. First, from an **assistive perspective**, they aim to seamlessly tailor control parameters to diverse users and tasks, ultimately providing an "out-of-the-box" solution that eliminates the need for expert setup or frequent recalibration. Second, from a **rehabilitation or training perspective**, they seek to systematically shape user adaptation, promoting motor learning and guiding interaction outcomes. Achieving these long-term goals motivates a shift toward interfaces that *co-adapt*. Such systems continually monitor real-time deployment and adapt the interfaces in response to the user's ongoing adaptation.

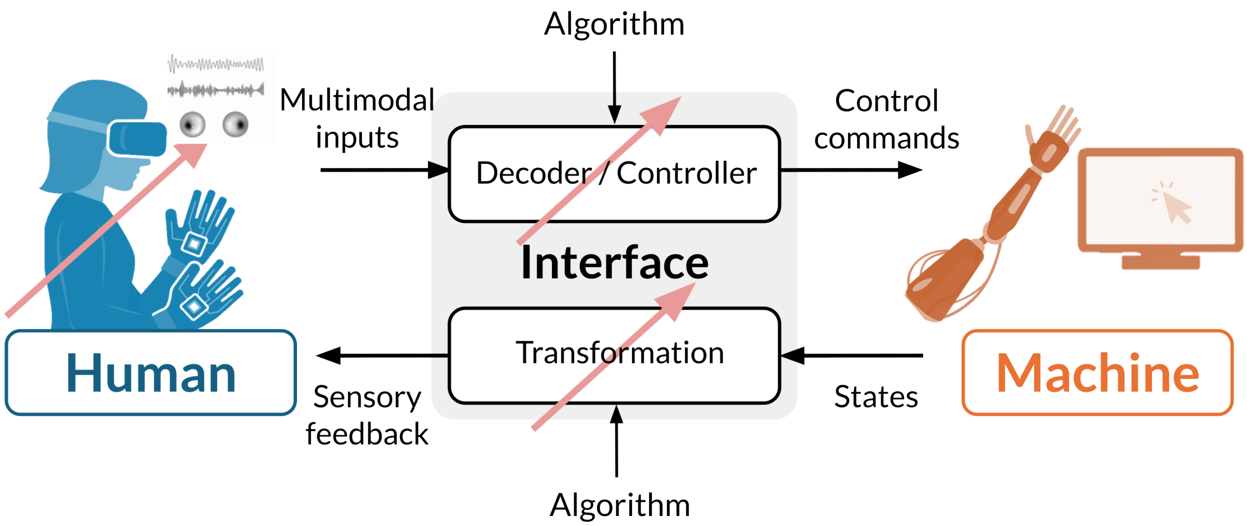

Figure 1.1: The closed-loop co-adaptation framework for human-machine interaction. The humans, adaptive interfaces, and machines interact in a closed-loop setting. The interface is composed of *motor* and *sensory* components, which are connected to the human and machine, respectively. (Top of the interface; *motor pathway*) Humans provide multimodal biosignal inputs that are continuous, such as neural recordings and eye tracking. These human inputs are mapped into control commands through a decoder or controller, to control machines such as a prosthetic limb or a computer cursor. (Bottom of the interface; *sensory pathway*) The machine states are transformed into sensory feedback or stimuli to humans, which can be visual, audio, or haptic cues. The interface is adaptive, meaning that the decoder, controller, or sensory feedback can be updated in real time, computed by algorithms. The algorithmic parameters can be adaptive as well. Humans' behaviors and signals change over time, and the interface is continuously updated to accommodate for the change in users. The adaptive nature is illustrated as a red arrow pointing to the top right.

2 Co-adaptation between humans and interfaces

Continuously adapting interfaces in a closed-loop — while the user operates the system and receives sensory feedback — can potentially bridge the performance gap between offline training and online testing, and guide user adaptation (Taylor et al., 2002; Gage et al., 2005; Ganguly and Carmena, 2009; Gilja et al., 2012; Gajos et al., 2012; Hahne et al., 2017; Madduri et al., 2023; Chasnov et al., 2025). In this context, machine learning algorithms are deployed in-the-loop with the user and the external device (Nothwang et al., 2016). The component that dynamically updates is the computational model, such as decoders (mappings or transformations), controller coefficients, and/or the algorithmic parameters, while the physical components like the sensors are usually not modified. Figure 1.1 illustrates a framework of co-adaptation between the human and the adaptive interface during human-machine interaction.

The concept of adapting algorithms in-the-loop of people has been studied in disparate fields, such as human-robot interaction (HRI) (Bengler et al., 2014; Li et al., 2019; Zemmar et al., 2020), brain-computer interfaces (BCI) (Madduri et al., 2026; Mahmood et al., 2019; Orsborn et al., 2012), body-machine interfaces (BoMI) (De Santis, 2021; Miehlbradt et al., 2018), human-computer interfaces (HCI) (Eisenstein and Puerta, 2000), personalized footwear in biomechanics (Tankink et al., 2024), assistive devices (Azimi et al., 2021; Felt et al., 2015; Slade et al., 2022; Zhang et al., 2017a; Zhuang et al., 2019b), and rehabilitations (Hardeman et al., 2019; Nahum-Shani et al., 2015).

The fundamental challenge of co-adaptation lies in determining *when* and *how much* the model or algorithm should adapt. A parallel idea is the *just-in-time adaptive interventions* (JITAIs) in the digital healthcare domain, which aims to provide interventions to address the dynamically changing needs of individuals "at the right time, in the right quantities" (Nahum-Shani et al., 2018). Studies in this domain measure patient's real-time psychological or physiological states to inform intervention decisions, such as mobile health notifications (Hardeman et al., 2019; Hsu et al., 2025), to optimize for fitness or recovery goals (Mistiri et al., 2025; Hekler et al., 2018).

Synthesizing optimal adaptive strategies at the right time and quantities would require an accurate human model to estimate the user's dynamic, time-varying behaviors. However, because optimal assistance patterns are highly varied across people (Zhang et al., 2017a) and user behaviors dynamically adapt in response to intelligent assistance (Green and Kalaska, 2011; Orsborn and Carmena, 2013), a fixed human model fitted on open-loop (offline) data (Mistiri et al., 2025) is often insufficient. Modeling these complex adaptations purely through offline data would demand extensive datasets spanning diverse, unconstrained human adaptive strategies. However, this is difficult to meet, particularly within clinical patient populations.

2.1 Human-in-the-loop optimization

Although a human model remains difficult to obtain, one prominent framework for achieving human-machine co-adaptation is *human-in-the-loop optimization*, wherein the algorithm updates the underlying model (i.e. the decoder, controller, or optimizer) to maximize systemic performance, without requiring an explicit human model (Slade et al., 2024). Instead, this approach treats the human, interface, and machine as a single unified system, focusing on optimizing toward a predefined performance metric, or measurable physiological objective, such as minimizing energy cost (Kantharaju et al., 2023; Slade et al., 2022; Zhang et al., 2017a), maximizing speed (Song and Collins, 2021), or increasing efficiency (De Santis, 2021). However, algorithmic optimization typically relies on an explicit human objective function to compute mathematical gradients, yet the true cost functions governing human co-adaptation during machine interaction remain largely unknown. Thus, studies have also investigated using more inclusive objectives that encompass several metrics, such as user preference (Ingraham et al., 2022; Tucker et al., 2020), to optimize for the user in-the-loop.

The human-in-the-loop optimization framework has been adopted to improve prostheses in gait research. For instance, in a lower-limb exoskeleton scenario aimed at reducing the energy consumption of walking (Slade et al., 2022), the optimizer continually tracked metabolic cost via respiratory data and periodically searched over optimal torque patterns that minimize energy expenditure. This computational loop is repeated continuously, allowing the system to update in real time as the user adapts. By directly optimizing control parameters with real-time measurements, this framework bypasses the need for an explicit human model or extensive offline pre-training datasets.

2.2 Closed-loop decoder adaptation

In applications like neuromotor interfaces, explicit knowledge of the user's intention is often used to deliver accurate control for actions such as reaching or typing. If we can identify user's task goal, we can use it as labeled data to train the decoder using machine learning. But this presents a new challenge: how to obtain labeled training data at the same time that the interface is being used? Current algorithms address this challenge by making use of information about calibration tasks that are prescribed to users, such as the target positions and velocities, and apply supervised learning algorithms to update the decoder over time (Dangi et al., 2013; Gilja et al., 2012; Madduri et al., 2022; Orsborn et al., 2012; Silversmith et al., 2020). These studies demonstrate that online decoder adaptation provides rapid initial calibration while continually retraining the decoder to account for signal nonstationarities (Hahne et al., 2015). The adaptive decoders presented in this framework also do not require a large dataset to be collected offline; instead, they are trained entirely online using the individual user's data.

Nevertheless, standard supervised approach assumes access to ground-truth intent labels, which are typically collected through prescribed, structured calibration tasks. While a structured calibration is acceptable for training an initial decoder, it becomes tedious if users must stop what they are doing to recalibrate every time the interface fails or when they switch to different tasks. Thus, BCI researchers have explored adaptation methods without the need for structured tasks via unsupervised approaches that eliminate labeled training data, including regression paradigms, principal component analysis (PCA), template matching, and autoencoders (De Santis, 2021; Jarosiewicz et al., 2015; Li et al., 2011; Rizzoglio et al., 2021; Sussillo et al., 2016). These unsupervised methods benefit the generalizability of neural interfaces during diverse, less structured, or self-paced tasks like free typing or writing (De Santis, 2021; Jarosiewicz et al., 2015).

2.3 Modeling co-adaptation between human and interface

While the above co-adaptive approaches demonstrate significant promise across various applications, they offer no theoretical guarantees regarding the equilibrium between the two intelligent agents, the human and the interface. Thus, the fundamental challenge of designing the optimal adaptive strategies remains largely unresolved. Predicting the outcomes of these complex interactions under different adaptive strategies is critical to guaranteeing safety and performance. To achieve this, it is helpful to establish a model of human-interface co-adaptation.

Because a complete empirical model is not yet available, we can begin with a simplified, parsimonious model as a foundation. Prior literature has investigated Bayesian models (Kording and Wolpert, 2004; Körding and Wolpert, 2006) and optimal feedback control laws (Diedrichsen et al., 2010; Todorov, 2004) to describe human sensorimotor control. Similarly, in the domain of human-machine interaction, classical frameworks such as the crossover model (McRuer and Jex, 1967) and inverse feedforward architectures (Drop et al., 2019; Wasicko et al., 1966; Yamagami et al., 2021b) have been deployed to capture human's control.

To apply an optimization framework under the assumption that human minimizes an internal cost, a fundamental question arises: which objective function best captures user adaptation during machine interaction? Although there is no answer to this question yet, prior studies have proposed cost functions with a composite structure consisting of task error and user effort (Berret et al., 2011; Emken et al., 2007), as the human brain tends to minimize metabolic or physical effort to optimize efficiency (Franklin et al., 2008). While these idealized human models and objective functions are inevitably approximations rather than exact representations, they still provide valuable conceptual contributions for understanding and predicting real-world human-machine interactions.

With a proposed human cost function, prior literature has modeled the interaction between the two intelligent agents, human and interface, as a two-learner problem. In this paradigm, agents employ optimal control (Hafs et al., 2026; Li et al., 2019; Pezeshki et al., 2023) or gradient descent (Felt et al., 2015; Chasnov et al., 2020; Madduri et al., 2021) to optimize either a joint objective function (Müller et al., 2017) or their respective individual cost functions (Madduri et al., 2026), thereby participating in a dynamic game (Başar and Olsder, 1998). Various game-theoretic equilibrium outcomes, such as convergence and stationary points, have been demonstrated in human-machine interaction (Chasnov et al., 2025; Madduri et al., 2026).

These game-theoretic frameworks are powerful tools for predicting co-adaptation outcomes and determining optimal adaptive strategies. Importantly, what is considered "optimal" depends on the application; for instance, in a cooperative setting, the priority may be for the interface to maximize assistance to the user, whereas in a rehabilitation setting, the objective is often to encourage larger user movement and effort. By leveraging this two-learner framework, researchers can theoretically predict whether and how the two agents will converge. This allows a systematic evaluation of different adaptive parameters, such as assistance levels (Hafs et al., 2026) or algorithmic adaptation rates (Chasnov et al., 2025; De Santis, 2021; Madduri et al., 2026), to induce targeted convergence patterns and ultimately shape the co-adaptive interaction.

2.4 Limitations in current co-adaptation frameworks

One limitation of the game-theoretic frameworks in prior work is that they formulate the human updates as deterministic optimal control problems in a dynamic game, with a standard assumption that both agents (human and interface) have noise-free information of the system (Hafs et al., 2026; Li et al., 2019; Merel et al., 2013). In reality, the human central nervous system cannot perfectly observe the machine state. We receive noisy sensory measurements and form a state estimation (Kording and Wolpert, 2004), while our motor outputs are corrupted by signal-dependent noise (Todorov, 2004). Simultaneously, the interface may not have perfect information of the human due to measurement noise, particularly in neural recordings. Although one can model human motor control as stochastic optimal control to handle disturbances (Berret et al., 2024; Todorov and Jordan, 2002), the standard linear-quadratic-Gaussian (LQG) solution does not generally extend to multi-player games when agents have asymmetric and delayed noisy observations (Başar and Olsder, 1998; Li et al., 2026; Ouyang et al., 2017; Vasal and Anastasopoulos, 2021). Human sensorimotor delays make the optimal control problem harder to solve analytically, especially when delays are further compounded in motor-impaired individuals (Song and Tong, 2013). Current human-machine frameworks do not account for these human delays and noise, which may lead to poor co-adaptation outcomes. Beyond assessing the robustness of current frameworks in simulation (Hafs et al., 2025), a new framework is needed to accommodate these challenges before being applied in real-world conditions.

Another limitation of the co-adaptation framework is that the objective function often only captures a narrow slice of true user objectives (Ingraham et al., 2022), which may not accurately reflect the user's actual behavioral goal or intention. However, modeling user intention remains difficult, particularly during unconstrained, free-form interaction (Jain and Argall, 2019; Jarosiewicz et al., 2015; Rigotti-Thompson et al., 2025; Willett et al., 2018). To address this challenge, prior research in robotics has proposed integrating additional modalities to infer user intentions during robotic interactions (Aronson and Admoni, 2022; Gardner et al., 2020). Adding additional modalities creates *multimodal* inputs. But how these multimodal inputs interact within a co-adaptive, human-in-the-loop system remains an open question.

3 Multimodal inputs to control machines

In human-machine interaction, humans interact and communicate with machines through various *sensory* and *motor* modalities (Figure 1.1). *Sensory modalities* enable humans to perceive stimuli or performance feedback from machines (e.g., a computer screen that provides visual cues to human eyes), while *motor modalities* enable humans to provide motor inputs to machines (e.g., a joystick that measures hand inputs). Here I refer to the traditional hand-held devices, such as joystick, mouse, steering wheel, as *manual* modalities.

Distinct modalities offer unique advantages and challenges when used as motor inputs for machine control. For example, neural modalities have higher dimensionality and a larger control bandwidth, but larger variability, compared to manual modalities (Artemiadis and Kyriakopoulos, 2008; Drost et al., 2006; Lobo-Prat et al., 2014; Yamagami et al., 2020). Combining distinct modalities together may leverage their complementary strengths while compensating their respective weaknesses (Blana et al., 2016; Novak et al., 2013; Pantic and Rothkrantz, 2003; Rizzoglio et al., 2020), and thus provides additional features or funtions in computer or robotic interaction (Müller-Putz et al., 2011; Zhang et al., 2019).

In addition, in human and animal sensorimotor control, redundant sensorimotor pathways commonly co-govern a single behavioral output (Gelfand and Latash, 1998; Latash et al., 2002; Scholz and Schöner, 1999). These redundant pathways provide an abundance of motor task solutions (Todorov, 2004), provide stability (Latash, 2012), assure robustness against uncertainties (Gelfand and Latash, 1998; Roth et al., 2016), and enhance movement efficiency (De Santis and Mussa-Ivaldi, 2020). In parallel to these biological observations, motor redundancy in human-machine interactions has shown to be beneficial, though prior studies have predominantly focused on multi-channel inputs derived from within the same modality such as multiple IMUs across different limbs (De Santis and Mussa-Ivaldi, 2020; Ranganathan et al., 2014).

Research is still limited regarding systems where multiple *distinct* modalities are combined with multiple motor pathways (e.g., a different modality on different limb). Here I call this setting *multimodal pathways*. Addressing this gap may provide new opportunities to promote neuromotor interface use in rehabilitation. For example, if a neuromotor interface is supplemented by an easily accessible manual modality, the user gains a two-“knob” control: when the neural signal degrades due to inherent biological variability or muscle fatigue, the user can seamlessly rely on the kinematic input to preserve control. However, it remains unknown how people integrate these multimodal pathways in machine interaction.

4 Research aims and dissertation chapters

Building upon the advantages and current challenges of co-adaptation paradigm and multimodal interfaces in human-machine interaction, a central research question emerges: how can **co-adaptation** and **multimodal inputs** shape interaction? This dissertation addresses the question by presenting new experimental methods and computational frameworks designed to bridge existing research gaps. This dissertation is structured around three research aims. Each aim is a different chapter, presenting theoretical, computational, and/or experimental results from multiple human-subject experiments.

4.1 Research methods

To model human behavior in machine interactions, this work leverages control theory alongside machine learning algorithms that adapt to individual users. Control theory provides an analytical tool for understanding human control strategies, which not only informs the underlying dynamics of behaviors, but also provides interpretable objective functions to apply in future optimization tools.

Evaluating theoretical results and computational algorithms requires initial testing within controlled environments, which was accomplished by conducting human-subject experiments with healthy participants. In these experiments, participants control a virtual *machine* (a computer cursor with dynamics) using different input modalities (manual, sEMG, and eye tracking), mimicking real-world machine interactions such as driving or robotic control.

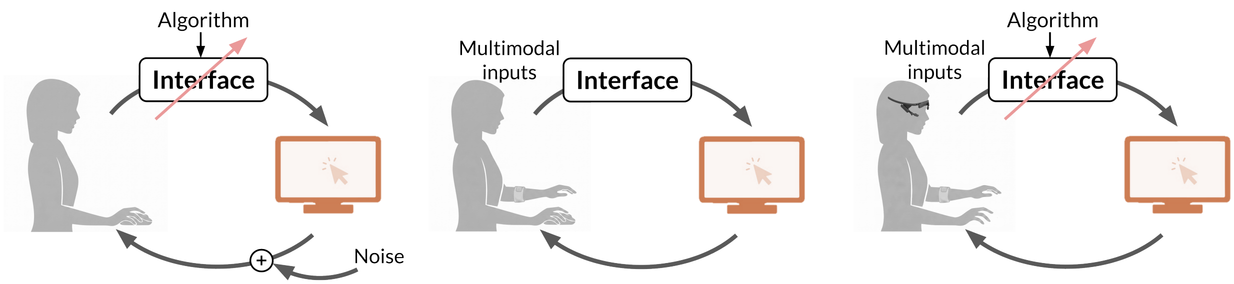

(a) AIM 1: Co-adaptive human-machine interaction with noisy measurements.

(b) AIM 2: Multimodal human-machine interaction with sEMG and kinematic interfaces.

(c) AIM 3: Co-adaptive, multi-modal human-machine interaction with sEMG and eye tracking.

Figure 1.2: Three research aims.

4.2 ***AIM 1. Co-adaptive human-machine interaction*** *(Chapter 2)*

Despite current advances in human-machine co-adaptation, there is still a gap between theoretical performance and real-world use of dynamic machines that co-adapt with people using optimization-based algorithms. One of the reasons that could cause this gap is because biosignal measurements are noisy, leading to poor translation of conventional theoretical equilibrium to applications outside of the lab.

This theoretical obstacle motivates us to conduct an experiment with human subjects ($N = 11$) controlling a dynamic machine under the influence of injected, known disturbance. We applied the human-in-the-loop optimization framework, where we tested an adaptive interface that continuously updates its dynamics to optimize for the objective of improving task performance while reducing "efforts", defined as input magnitude. The experimental results show co-adaptation between humans and interfaces, as well as an improvement in system performance.

This chapter discusses the diverse co-adaptive outcomes and tradeoffs that emerged when adaptive interfaces operate in the loop with humans and dynamic machine. Additionally, a model of this co-adaptation was developed using the empirical data and subsequently tested in simulation. This work provides a computational framework for synthesizing and evaluating co-adaptive systems with noise, and highlights the need to systematically engineer diverse co-adaptive outcomes in future applications.

4.3 ***AIM 2. Multimodal human-machine interaction*** *(Chapter 3)*

Motivated by the complementary advantages and distinct challenges of neural and kinematic inputs, this aim investigates how people integrate multiple input modalities during human-machine interaction. We developed a multimodal setup integrating sEMG and a manual joystick, and conducted an experiment where human subjects ($N = 15$) controlled a dynamic cursor using both modalities simultaneously, with each input provided by a different arm. This paradigm mirrors a clinical rehabilitation scenario where a patient operates a single-arm exoskeleton on their affected limb using sEMG, while supplementing control via a joystick operated by their unaffected limb.

This chapter explores how users integrate these motor pathways with distinct modalities, such as their control strategies and subjective preferences. The empirical results demonstrate that accessing abundant, redundant modalities yields enhanced performance and reduced sensorimotor noise when interacting with a dynamic machine. Furthermore, the results demonstrate that humans coordinated multiple modalities differently depending on their performance and preferences. These findings provide insights for designing personalized multimodal interfaces, such as algorithms that adapt the relative weights or magnitudes of each input modality.

In addition to multimodal *motor pathways*, in a preliminary study ($N = 3$), we evaluated human behaviors when there were multimodal *sensory pathways*. We added a kinesthetic haptic device to provide sensory feedback in addition to visual feedback (cursor on the computer screen). In this chapter, I show that the presence of haptic feedback did not significantly enhance the performance of our computer task, but it changed the user behaviors by encouraging larger user inputs.

4.4 ***AIM 3. Co-adaptive interfaces with multimodal inputs*** *(Chapter 4)*

In neuromotor interfaces, current decoder adaptation methods have shown to provide automatic re-calibration during user operation. However, many training methods require a prescribed calibration task to provide labeled data for supervised learning, limiting user autonomy in unconstrained, self-paced applications. The vision of this aim is to leverage an additional input modality to continuously infer user intention, which may remove the need for prescribed, repetitive calibrations whenever user switches tasks. Instead, the interface should seamlessly update to accommodate the user's changing needs in real time.

This chapter demonstrates a novel algorithm that leverages eye tracking to improve an adaptive sEMG interface that continuously adapts to users, without the need for task-specific labels (e.g., prescribed paths or targets). We benchmarked this new approach in a controlled task with human subjects ($N = 11$). Then in a preliminary study, we further demonstrated its feasibility in situations where the user's "ground truth" is not apparent, during a user-led reaching and writing.

I also demonstrate findings from another preliminary study ($N = 7$), where we investigated how eye movements may provide an easily-measured biomarker for neuromotor interface learning. We measured eye movements as users operated an sEMG interface to control cursor, and in post-hoc analysis, we investigated how gaze correlated with task performance and internal learning controllers. Results showed that eye movements provide useful information about changes in the user's learning strategy, which may inform rapid interface recalibration in future applications.

4.5 Interface design considerations (Chapter 5)

In Chapter 5, I describe design considerations for co-adaptive, multimodal interfaces derived from user experiences across five myoelectric experiments with healthy participants. I demonstrate that users' strategies for myoelectric control are diverse, and that user background — including prior experience with video games, musical instruments, and computers — influences interactions with co-adaptive interfaces.

4.6 Contributions, collaborations, and future work (Chapter 6)

Lastly, in Chapter 6, I summarize the contributions of this thesis, my past collaborations, and plans for future work. Moving forward, I plan to implement these frameworks in rehabilitation programs, including robot-assisted therapy, to deliver personalized interventions for people with motor impairments.

5 Significance

Biosignal-based human-machine interfaces hold promise for supporting people who rely on assistive and rehabilitative technologies, yet most current devices apply a fixed level of assistance or training regardless of how a patient's ability or needs change over time. The work presented in this dissertation is a step towards addressing that gap by providing experimental and computational frameworks that allow interfaces to adapt in response to user adaptation, informed by multimodal inputs. Specifically, this work models how humans interact with adaptive interfaces and how multimodal inputs can be leveraged to improve seamless device control.

Findings from this work inform future designs of adaptive programs for assistive devices and rehabilitative interventions. For instance, the control theory analysis may offer a new way to model patients' control strategies and motor recovery, while multimodal sensing may provide a new way to extract user intention or health patterns. Modeling a patient's dynamic behavior in this way tells the system when and how much to adapt the interface to deliver personalized interventions in real time, analogous to a clinician who adjusts therapy in person based on a patient's changing condition. This kind of automated system is especially valuable for remote, home-based rehabilitation, where it can provide tailored training that promotes recovery without adding burdens to patients or caregivers.

In summary, this work provides a foundation for designing future co-adaptive, multimodal interfaces, and adds to the scientific knowledge of human behavior in device interaction, ultimately enhancing the usability, generalizability, and clinical translation of biosignal-based interfaces.

Chapter 2

CO-ADAPTIVE HUMAN-MACHINE INTERACTION

Co-adaptation of a Dynamic Human-Machine Interface

Amber H.Y. Chou, Momona Yamagami, Maneeshika M. Madduri, Benjamin J. Chasnov, Finley Hutchison, Lauren N. Peterson, Samuel A. Burden

In preparation.

Abstract

Despite the growing prevalence of learning algorithms in daily life, methods for analysis and synthesis of how these systems interact with people are limited. We studied optimization-based algorithms that co-adapt with people in the presence of dynamic machines, finding limitations on current theory that motivated us to conduct an experiment where human subjects interact with a machine that has complex dynamics through an adaptive interface. Experimental results provided evidence of co-adaptation and a trade-off between performance and the "effort" of the human and interface, defined as the norm of their output signals. We developed a parsimonious model of the human adaptation strategy observed in our experiments and conducted a simulation study using this model. Our computational results matched the empirical results, suggesting our human subjects adapted to minimize a combination of error and effort. These results demonstrate how co-adaptation between humans and intelligent interfaces shapes behavior and performance, and introduces a modeling framework that can be used in future work to systematically design interaction outcomes.

1 Introduction

With practice and effort, people can learn to control dynamic systems like vehicles (McRuer, Duane T and Jex, Henry R, 1967; Zhang et al., 2020b; Yamagami et al., 2021a), tele- and co-robots (Nikolaidis et al., 2017; Warrier and Devasia, 2016; Zemmar et al., 2020), prostheses (De Santis, 2021; Zhuang et al., 2019a; Azimi et al., 2021), exoskeletons (Zhang et al., 2017a; Felt et al., 2015; Slade et al., 2022), or brain-computer interfaces (Orsborn et al., 2012; Willett et al., 2021; Mahmood et al., 2019). To facilitate and shape this learning, it is tempting to inject intelligence into the *interface* between the human and machine (Miehlbradt et al., 2018). Since people continually adapt to their sensorimotor context (Taylor et al., 2014; Heald et al., 2021), an intelligent interface in this closed-loop interaction creates a two learner problem (Müller et al., 2017; Perdikis and Millan, 2020), where the human and interface *co-adapt* (Madduri et al., 2023). Analyzing and synthesizing these systems is critically important in current and emerging applications, including driver assistance (Bengler et al., 2014), surgical robotics (Gomes, 2011), rehabilitative robotics (Marchal-Crespo and Reinkensmeyer, 2009), active prosthetics (Zhuang et al., 2019a), and neural interfaces (both invasive (Orsborn et al., 2012) and non-invasive (Yamagami et al., 2020; Chou et al., 2022)).

Motivated by the optimal feedback control hypothesis for human motor control (Diedrichsen et al., 2010; Todorov and Jordan, 2002; Selinger et al., 2019; Emken et al., 2007), we modeled the interaction between a human, adaptive interface, and dynamic machine using a robust control framework (Zhou et al., 1996). We assumed that both the human and the interface solve optimal control problems to determine their behaviors. In our theoretical analysis, we allowed arbitrarily large (but finite) system orders of the dynamic machine, and analyzed the equilibrium behavior of such systems. We found that a simple form of this problem has no solution due to a fundamental theoretical limitation neglected in prior work (Merel et al., 2013; Li et al., 2019). In prior work, it was assumed that the human and the interface had noise-free observations of the machine's state vector, which we regard as unrealistic in real-world applications. For example, in the context of neuroprosthetics, where an intelligent neural interface seeks to help a person control a dynamic device, intrinsic stochasticity of neural signals (Dayan and Abbott, 2005) injects noise into the system, precluding access to full information. This finding led us to conduct a human subjects experiment and simulation to understand and model how people behave in-the-loop with adaptive interfaces and dynamic machines under more realistic conditions.

In our experiment, we investigated co-adaptation between humans and adaptive interfaces in the presence of specific machine dynamics that are fundamentally challenging to control (Zhang et al., 2020b). We recruited participants ($N = 11$) to perform a continuous disturbance-rejection task using a one-dimensional manual input device (Yamagami et al., 2021a). We deployed adaptive interfaces with 0th-, 1st-, and 2nd-order dynamics, which correspond to the position-, velocity-, and acceleration-based systems that people routinely encounter in daily life (McRuer, Duane T and Jex, Henry R, 1967). The interface controller coefficients were updated as the participant completed a sequence of trials by estimating a model of the participant's feedback controller and optimizing a performance criterion of interest. We compared the system performance and human behavior with co-adaptation and without co-adaptation, finding that co-adaptation led to diverse outcomes exhibiting trade-offs between agent behavior and system performance.

Driven by this experimental finding, we conducted a computational analysis to test whether a model for co-adaptation can replicate the empirical outcomes we observed. Specifically, we modeled human behavior as arising from minimization of a cost function. We implemented a simulation where the human model and interface alternately optimize their costs and we conducted a parameter sweep over the space of penalty terms in the two agents' cost functions. This simulation yielded behaviors that aligned closely with empirical observations. Our experimental and computational findings will inform future frameworks and deployments of co-adaptive interfaces in human-machine interactions.

2 A model of co-adaptation

We consider systems where the machine has dynamics, that is, where the dimension of the machine's state vector – the system *order* (Aström and Murray, 2010, Sec. 2.2) – is greater than or equal to 1. The zero-dimensional case was studied in prior work (Madduri et al., 2021, 2022; Chasnov et al., 2023). We modeled dynamic human-machine interfaces (HMI) as the block diagrams in Fig. 2.1, where: H represents the *human* "in-the-loop", M is the *machine* that is being controlled, and I is the *interface* we seek to synthesize. When the blocks H, M, and I are linear time-invariant (LTI) transformations (Hespanha, 2009, Lec. 3), these diagrams are not solely conceptual – they provide mathematical specifications of the closed-loop transformation from input disturbance w to output error z. Although humans are generally nonlinear, they can behave remarkably linearly when interacting with finite-order LTI machine-and-interface dynamics (McRuer, Duane T and Jex, Henry R, 1967). These observations motivate us in what follows to focus on the finite-order LTI case where there exists a comprehensive toolkit for analysis and synthesis of feedback systems (Aström and Murray, 2010; Zhou et al., 1996; Hespanha, 2009; Benner et al., 1999).

2.1 Interface synthesis

The field of *robust control theory* (Zhou et al., 1996) provides methods for synthesizing the interface I in Fig. 2.1 (conventionally termed the *controller*) to optimize a performance criterion with respect to fixed models of human H and machine M (and, optionally, uncertainty in the models). The performance criteria of interest in robust control are *induced norms* (Zhou et al., 1996, Ch. 4) that quantify how much signal power is transferred from input disturbance w to output error z. In this robust control paradigm, the interface I is synthesized by solving an optimization problem

$$\min_I ||H/M/I||_{\mathcal{I}} \tag{2.1}$$

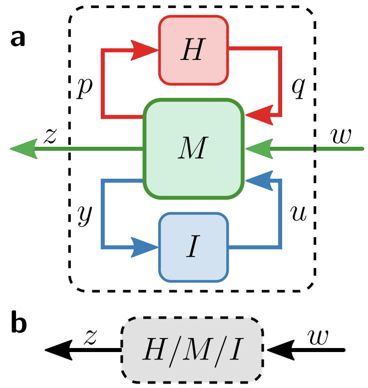

Figure 2.1: *Block diagram models for human-machine interfaces (HMI).* (**a**) This diagram specifies that the human H transforms signal p to signal q, i.e. the human observes output p from machine M and provides control input q to the machine. Similarly, the interface I transforms machine output y to control input u, and the machine M transforms both of the control inputs q, u and an external input w to produce outputs p, z, y. The external input w can contain a *disturbance* to reject (e.g. measurement or process noise) or a *reference* to track (e.g. a trajectory or stationary point). (**b**) This diagram shows a simplification of (a) that obscures details about the interconnection between H, M, and I, instead denoting the transformation $H/M/I$ resulting from this interconnection. The block $H/M/I$ is defined so that (a,b) both specify the same transformation from w to z.

where the induced norm $\|\cdot\|_{\mathcal{I}}$ encodes what components of z the interface seeks to make as small as possible. Under appropriate restrictions on the choice of norm $\|\cdot\|_{\mathcal{I}}$, models of human H and machine M, and statistics of input w and output z, a solution I^* to the optimization problem in (2.1) exists (Zhou et al., 1996) and can be computed using efficient numerical algorithms (Benner et al., 1999). As an example, the well-known *linear-quadratic Gaussian* (LQG) regulator (Hespanha, 2009, Lec. 24) is obtained by solving (2.1) when: the disturbance w is Gaussian; the error z consists of (linear transformations of) the state of the machine x and the control input u; the human-machine transformation H/M (defined in Fig. 2.2) is stabilizable and observable (Hespanha, 2009, Lec. 14, 15); and $\|\cdot\|_{\mathcal{I}}$ is the induced 2-norm.

If the statistics of the human H and machine M in Fig. 2.1 are stationary regardless of the implemented interface I (e.g. if H and M are given as fixed LTI transformations or distributions of LTI transformations), then the preceding paragraph describes a remarkably flexible framework for synthesizing an optimal interface I^* (Zhou et al., 1996; Anderson et al., 2019). However, although it may be reasonable to assume or ensure machine dynamics are stationary, ample evidence suggests that the human will naturally adapt to any perceived change in the interface, and moreover that this adaptation will not be random – rather, the human's transformation will be strongly influenced by the interface (McRuer, Duane T and

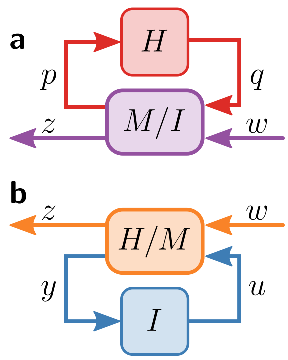

Figure 2.2: *Block diagrams from human and machine perspectives.* Given LTI transformations H, M, and I, the blocks M/I and H/M are defined so that the diagrams in (**a**, **b**) specify the same transformation from w to z as the diagrams in Fig. 2.1. Mathematically, the U/L operation is defined as the *linear fractional transformation* (Zhou et al., 1996, Ch. 3, 10) between blocks U and L. These simplified diagrams are conceptually useful when reasoning from the individual perspectives of the human H and interface I as they jointly interact with the machine M. Indeed, H interacts with the interconnection M/I between M and I as illustrated in (a), whereas I interacts with the interconnection H/M between H and M as in (b).

Jex, Henry R, 1967; Yamagami et al., 2020; Zhang et al., 2018, 2020b). Therefore if the interface I^* is synthesized by solving (2.1) with respect to an initial guess or estimate of the human's transformation H, it is reasonable to expect the human will adapt its transformation from H to $\widetilde{H}$ when the interface changes from I to I^*. Unfortunately, the synthesized interface I^* is not optimal with respect to the adapted human transformation $\widetilde{H}$. In fact, implementing I^* in-the-loop with $\widetilde{H}$ could yield arbitrarily bad performance (Doyle, 1978).

2.2 Co-adaptation

The preceding observations motivate the study of interfaces that *co-adapt* with the human and, hence, regard the human H and interface I as two *learners* (Müller et al., 2017; Perdikis and Millan, 2020) playing a *dynamic game* (Başar and Olsder, 1998) through their interaction with the machine M. As a starting point for modeling this interaction, prior work suggests the human may play this game by solving their own optimization problem (Diedrichsen et al., 2010; Todorov and Jordan, 2002; Selinger et al., 2019; Emken et al., 2007)

$$\min_{H} ||H/M/I||_{\mathcal{H}} \tag{2.2}$$

where the norm $\|\cdot\|_{\mathcal{H}}$ encodes what components of z the human seeks to make as small as possible. Other than in the special case where the goals of the human and interface are perfectly aligned so that $\|\cdot\|_{\mathcal{I}} = \|\cdot\|_{\mathcal{H}}$, the outcome of the game defined by simultaneously considering the optimization problems in (2.1) and (2.2) will generally represent a compromise between the player's conflicting goals. One such outcome considered in prior work on human-machine interfaces (Merel et al., 2013; Müller et al., 2017; Li et al., 2019; Madduri et al., 2021) is a *Nash equilibrium* (Nash, 1950; Başar and Olsder, 1998) defined by a pair of transformations H^*, I^* such that H^* minimizes $\|H/M/I^*\|_{\mathcal{H}}$ with respect to H and I^* minimizes $\|H^*/M/I\|_{\mathcal{I}}$ with respect to I.

2.3 Non-existence of Nash equilibria

This section provides theoretical analysis of the dynamic game defined in the preceding section that is played by a human H and interface I interacting with a machine M. We will assume in what follows that H, M, and I are finite-order linear time-invariant (LTI) transformations of continuous-time signals. The game of interest is specified by the coupled optimization problems

$$\min_{H} \|H/(M/I)\|_{\mathcal{H}} \tag{2.3a}$$

$$\min_{I} \|(H/M)/I\|_{\mathcal{I}} \tag{2.3b}$$

where $\|\cdot\|_{\mathcal{H}}$, $\|\cdot\|_{\mathcal{I}}$ denote the cost functions of the human and interface, respectively, and the closed-loop transformations $H/(M/I) = H/M/I = (H/M)/I$ are defined in Fig. 2.2. The simplified block diagrams in Fig. 2.2 are useful to illustrate (i) the combined machine-interface system (M/I) the human H interacts with and (ii) the combined human-machine system (H/M) the interface I interacts with.

Given LTI transformations M and I, the (central, strictly feasible) solution H^* of the optimization problem in (2.3a) with respect to the 2- or ∞-norm is, under standard stabilizability and detectability conditions, an LTI transformation (Zhou et al., 1996, Thm 14.7) that can be computed using efficient algorithms (Benner et al., 1999). With the exception of the special cases considered in prior work (Merel et al., 2013; Li et al., 2019; Hafs et al., 2026) (termed *full information* and *full control* in (Zhou et al., 1996; Doyle et al., 1989)), the number of state variables in the dynamics of H^* (the system's *order*) is equal to the sum of the number of state variables in M and I; if we let $\#(T)$ denote the number of state variables in the LTI transformation T, this property can be written $\#(H) = \#(M) + \#(I)$. Similarly, given LTI transformations H and M, the solution I^* of the optimization problem in (2.3b) is an LTI transformation, and the number of state variables in I^* equals the sum of the number of state variables in H and M: $\#(I) = \#(H) + \#(M)$. We will assume this generic condition holds in what follows.

We will show in the Theorem below that the game in (2.3) generally has no Nash equilibrium defined in terms of finite-order LTI transformations when M has dynamics, i.e., $\#(M) \geq 1$. To see why this may be the case, consider the co-adaptive interaction wherein H and I alternately solve their optimization problems ((2.3a) and (2.3b), respectively). Even starting from a static interface with $\#(I) = 0$, solving (2.3a) with respect to given M and I yields a solution H^* with $\#(H) = \#(M) \geq 1$, so subsequently solving (2.3b) with respect to given H^* and M yields solution I^* with $\#(I) = \#(H) + \#(M) = 2\,\#(M) \geq 2$. Iterating this process yields a sequence of H and I with ever-increasing numbers of state variables, preventing the existence of a stationary point for (2.3) in the sense of Nash defined with finite-order LTI transforms.

Theorem: *Suppose M is a finite-order linear time-invariant (LTI) transformation with dynamics so that $1 \leq \#(M) < \infty$, and suppose that $\|\cdot\|_{\mathcal{H}}$ and $\|\cdot\|_{\mathcal{I}}$ are system 2- or ∞-norms. If there exists a Nash equilibrium for* (2.3) *defined in terms of finite-order LTI transformations H^*, I^*, then at least one of H^*/M or M/I^* is full information or full control.*

Proof (by contradiction): Suppose there exists a Nash equilibrium H^*, I^* for (2.3) such that H^* and I^* are LTI and have finite order. If neither H^*/M nor M/I^* is full information or full control, then (Zhou et al., 1996, Thm. 14.7) implies $\#(H^*) = \#(M) + \#(I^*)$ and $\#(I^*) = \#(H^*) + \#(M)$. Substituting the second equation from the preceding sentence into the first and simplifying yields $0 = 2\,\#(M)$. But since $\#(M) \geq 1$, this equation is a contradiction.

Remark: *The co-adaptation games studied in prior work (Merel et al., 2013; Müller et al., 2017; Li et al., 2019; Madduri et al., 2021; Hafs et al., 2026) consider only the full information case, where it is assumed that the interface can observe the state of the machine M and human H with no measurement noise and similarly the human can observe the state of the machine and interface I with no noise. In this case, the game can admit a Nash equilibrium defined in terms of static state feedback transformations for H and I (Başar and Olsder, 1998, Sec. 6.2.2).*

Example: The neuroprosthetic example from the introduction illustrates why full information or full control assumptions are unrealistic. Indeed, full information would require both adaptive agents – human and interface – have noise-free measurements of all system states, including those internal to the other adaptive agent. Similarly, full control for either agent would require that agent to directly influence all system states, including those internal to the other; although it might be possible in principle to give the brain full control over the machine and interface, it seems implausible (and perhaps undesirable) for the interface to have full control over the human's neural state. Either assumption seems inconsistent with our understanding of neural dynamics (Dayan and Abbott, 2005).

Remark: *The preceding Theorem only applies to the case where H, M, and I are all finite-order LTI transformations as in (Zhou et al., 1996). Indeed, solutions of multi-agent stochastic optimal control problems need not be linear even if dynamics are linear and disturbances are Gaussian (Witsenhausen, 1968); we cannot say at present whether the human-machine game considered here has nonlinear stationary points. Furthermore, the techniques in (Zhou et al., 1996) are not applicable in the presence of delay or other common structural constraints, which necessitate the use of new tools (Anderson et al., 2019).*

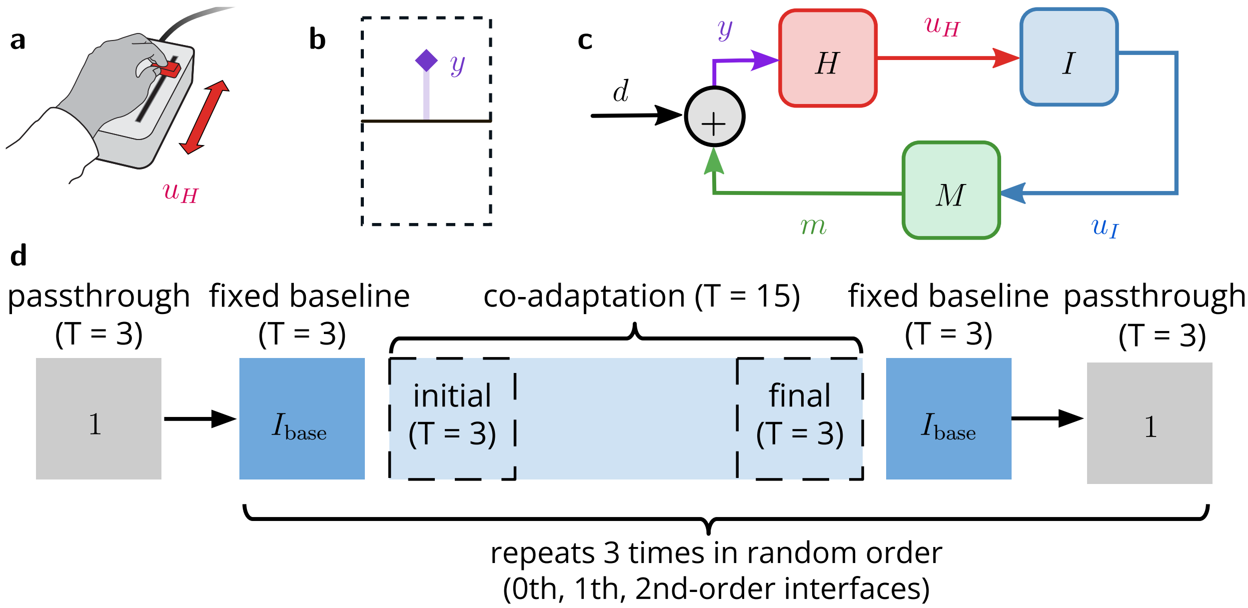

Figure 2.3: *Human subjects experiment design.* (**a**) Participants provide output u_H using a 1D manual device to (**b**) control the position of a cursor y on a computer display to minimize the distance to a reference line in the middle of the screen. (**c**) Human H, machine M, and interface I are connected in series, with the human viewing cursor y and producing response u_H that is input to I. The machine output is perturbed by external disturbance signals d. (**d**) Experiment conditions and trial structure. T indicates the number of trials (e.g., $T = 3$ means there are 3 trials in the block). Three conditions (0th-, 1st-, and 2nd-order interface) are presented in random order, followed by 3 pass-through trials after each block.

3 Methods

The experiment and computational analysis were designed to determine whether and how human-machine interfaces (HMI) co-adapt, and to characterize how co-adaptation affects human, interface, and system performance relative to a non-adaptive *baseline*.

3.1 Human subjects

Eleven participants were recruited (age: 19-36; gender: 6 women, 5 men; hand dominance: 10 right-handed, 1 left-handed). All were daily computer users. All participants provided informed consent according to the University of Washington, Seattle's Institutional Review Board (IRB #00000909).

3.2 Task

Participants controlled a one-dimensional (1D) cursor using a 1D slider (35×12×22 mm rectangular handle on a 10 cm slide potentiometer; Fig. 2.3a). The slider input u_H was transformed through interface I and a fixed machine M to produce cursor position y on the screen (Fig. 2.3c). Participants were instructed to keep the cursor as close to the center line, $y = 0$, as possible (Fig. 2.3b). We chose a nonminimum-phase second-order machine to increase task complexity (Zhang et al., 2020b):

$$\begin{aligned} \ddot{m} + 3.6\dot{m} + 4m &= -2(\dot{u}_I - 2.2u_I), \\ M(s) &= \frac{-2(s-2.2)}{s^2+3.6s+4}, \\ y &= m + d, \end{aligned} \tag{2.4}$$

where $\dot{m}$ represents the time-domain machine output m differentiated by time, $M(s)$ is the continuous-time representation of transformation M in the frequency domain, $s = j\omega$. The cursor position y was updated at 60 Hz, so the continuous-time machine was discretized (`scipy.signal.cont2discrete`).

We introduced an unpredictable disturbance d to the cursor (Fig. 2.3c), which mimics the measurement or process noise from external devices. Following prior work(Yamagami et al., 2018a, 2021a; Chou et al., 2022), disturbance signals were constructed as a sum of sinusoidal signals with the first eight prime multiples of a base frequency of 1/20 Hz (*stimulated frequencies* $\Omega = [0.10, 0.15, 0.25, 0.35, 0.55, 0.65, 0.85, 0.95]$ Hz). The magnitudes at each frequency component were normalized by its frequency to ensure constant signal power. The phase of each frequency component was randomized in each trial to produce pseudorandom time-domain signals.

3.3 Signal processing

Prior work demonstrated that when humans are tasked with tracking references and rejecting additive disturbances through a linear time-invariant (LTI) (Aström and Murray, 2010, Ch. 3, pg. 4) system M, humans behave approximately like LTI transformations for a range of reference and disturbance signals (Yamagami et al., 2018a, 2021a; McRuer, Duane T and Jex, Henry R, 1967). Thus, we can analyze our system in Fig. 2.3c using the frequency-domain representations (Phillips et al., 2003, Ch. 5) of signals and LTI systems. In this paper, time-domain signal x with a "hat" $\widehat{\cdot}$ denotes the Fourier transform, $\widehat{x}$, in the frequency-domain.

We recorded time-domain signals y and u_H at each sampling interval, and collected the updated interface dynamics I for each trial. In post-hoc analysis, we converted these signals into their frequency-domain forms, $\widehat{y}$ and $\widehat{u}_H$, for each trial, then derived the interface input

$\widehat{u}_I$ and the human transformation H (feedback controller) at stimulated frequencies:

$$\widehat{u}_I = I(j\omega)\widehat{u}_H(j\omega) \tag{2.5a}$$

$$H(j\omega) = \frac{\widehat{u}_H(j\omega)}{\widehat{y}(j\omega)}, \tag{2.5b}$$

where $\omega \in \Omega$. Given this H, the prescribed signals d, and transformations M, I, we can additionally derive the steady-state signals via block diagram algebra (Aström and Murray, 2010, Sec. 2.2):

$$\begin{aligned} \widehat{y}(j\omega) &= \frac{\widehat{d}(j\omega)}{1 + H(j\omega)M(j\omega)I(j\omega)}, \\ \widehat{u}_H(j\omega) &= H(j\omega)\widehat{y}(j\omega), \\ \widehat{u}_I(j\omega) &= I(j\omega)H(j\omega)\widehat{y}(j\omega), \end{aligned} \tag{2.6}$$

which were used in the interface updates (section 3.5).

3.4 Conditions

We tested interfaces that have 0th-, 1st- and 2nd-order dynamics (Table 2.1), corresponding to the position-, velocity-, and acceleration-based systems (McRuer, Duane T and Jex, Henry R, 1967) (Fig. 2.3d). Participants first completed 3 trials of a pass-through interface with an identity mapping ($I(z) = 1$, i.e., $u_I[t] = u_H[t]$), then performed the three interface conditions in random order. Each condition block consists of 3 baseline trials, followed by 15 trials of co-adaptation, and 3 additional baseline trials. At the end of each condition block, participants completed another 3 pass-through trials, with a total of 12 pass-through trials per participant. Each trial was 40 seconds after a 5-second ramp-up (45 seconds total). Participants were encouraged to rest between trials and required to take a three-minute break between condition blocks. Total experimental duration was approximately 2 hours per participant.

Table 2.1: *The discrete-time representations of 0th-, 1st-, 2nd-order, and pass-through interfaces tested in the experiment.* $z = e^{j\omega dt}$, *where* $dt = 1/60$ *second.*

condition	dynamics	baseline
0th	$I_0(z) = b$	$I_{0,\text{base}}(z) = 5$
1st	$I_1(z) = \frac{b_0 z + b_1}{z + a}$	$I_{1,\text{base}}(z) = \frac{1.55z - 1.55}{z - 0.98}$
2nd	$I_2(z) = \frac{b_0 z^2 + b_1 z + b_2}{z^2 + a_0 z + a_1}$	$I_{2,\text{base}}(z) = \frac{1.06z^2 - 2.12z + 1.06}{z^2 - 1.99z + 0.99}$
pass-through	$I_{\text{pass}}(z) = 1$	

3.5 Interface adaptation

Interface cost

We defined the interface cost as the induced 2-norm of $H/M/I$ (Fig. 2.1). The cost function was a linear combination of task error, interface effort, and human effort:

$$c_I(H, I) = \frac{\|(\widehat{y}, \widehat{u})\|}{\|\widehat{d}\|} = \frac{\|\widehat{y}\| + \lambda_I \|\widehat{u}_I\| + \lambda_H \|\widehat{u}_H\|}{\|\widehat{d}\|} \tag{2.7}$$

with penalty terms $\lambda_I = 0.5$ and $\lambda_H = 1.5$, which were estimated from pilot data to ensure the scaled interface and human efforts were within the same order of magnitude as the task error. We further scaled the interface penalty λ_I by half to avoid optimization from getting stuck at the smallest possible interface magnitudes. We used 2-norm to match the task instruction of minimizing the cumulative cursor displacement.

Interface parameters discretization

The parameter space for each interface order was discretized into three candidate lists I_{list} of transfer-function coefficients (a, b in Table 2.1).
1) 0th-order interfaces: gain $b \in [0.05, 5]$ to ensure interface remained responsive to user input without becoming unmanageably volatile, sampled at 100 equidistant points.
2) 1st-order interfaces: We generated 1st-order high-pass interfaces by sweeping gains (0.05 to 5) and cutoff frequencies (0.05 to 1 Hz), with 100 equidistant points per range. The cutoff range was selected to align with the low-frequency crossover observed in human manual control tasks (Yu et al., 2014a; Yamagami et al., 2021a).

Since the cursor updated at 60 Hz, we mapped continuous-time interface dynamics to their discrete-time representations. To ensure interfaces are stable or marginally stable, we excluded any interfaces where the eigenvalues of the denominator exceeded $|1|$ in discrete time. We also excluded configurations resulting in negative DC gain at $z = 1$, to prevent control inversion, which our pilot participants reported as confusing and counter-intuitive. This resulted in a global parameter set of 8,546 combinations of coefficients (b_0, b_1, a).
3) 2nd-order interface: We generated 2nd-order high-pass interfaces by sweeping gains (0.05 to 5), cutoff frequencies (0.05 to 1 Hz), and damping ratios (0.5 to 1), with 50 points per parameter for computational efficiency. Damping ratios below 1 were selected to ensure the interfaces remained stable and avoided aggressive noise amplification. We again computed the discrete-time representation of each interface, and then excluded configurations where interfaces were unstable or had negative DC gain, yielding a global set of 87,038 combinations of coefficients $(b_0, b_1, b_2, a_0, a_1)$.

For both 1st- and 2nd-order interfaces, we also explored low- and band-pass filters in pilot testing, but they consistently yielded higher task errors than high-pass filters. Thus, we restricted our adaptation framework to high-pass interfaces, as their phase-lead characteristic (Ogata, 2009, Ch. 7.11) has the potential to compensate for sensorimotor delays in the nervous system.

Interface baselines

We started and ended each condition block with a baseline interface: 0th-order $I_{0,\text{base}}$, 1st-order $I_{1,\text{base}}$, and 2nd-order $I_{2,\text{base}}$. Having different orders of baselines allowed for a controlled comparison within the same interface complexity, so we could verify whether improvements stem from co-adaptation or were simply an artifact of transitioning to a higher-order interface. We also use the baselines as the starting interfaces in co-adaptation.

We synthesized the baseline interfaces using pilot subject data, where 6 pilot subjects performed the same dynamic game with a pass-through interface ($I = 1$) and the second-order machine defined in (2.4). We estimated human transformations as (2.5b), then determined the optimal interface by searching through the parameter space I_{list} to minimize the interface cost $c_I(H, I)$ in (2.7), for each of the three conditions. We used these three optimal interfaces as the baselines, and all participants used the same baselines.

Interface updates

The interface update occurred at the end of every co-adaptation trial for a total of 15 trials for each condition. The interface updates occurred in three steps.
1) Human model: we estimated human transformation H in (2.5b) for each trial. This model was subsequently updated as a weighted combination between the historical estimate H^- and the current trial's estimate H:

$$H^+ = \alpha_H H^- + (1 - \alpha_H)H \tag{2.8}$$

where $\alpha_H = 0.75$ for gradual updates in the human model. For the first co-adaptation trial, the human model was initialized using data from the 3 preceding non-adaptive baseline trials.
2) Interface grid search: For each candidate interface in I_{list}, we computed the steady-state signals $\hat{y}, \hat{u}_H, \hat{u}_I$ using (2.6) and generated d with random phase, then performed a grid search to identify a parameter set that minimizes the interface cost:

$$I^* = \arg\min_I c_I(H^+, I). \tag{2.9}$$

3) Interface updates: *Smooth Batch* (Orsborn et al., 2012) was implemented to ensure gradual interface changes between trials. The interface for the subsequent trial I^+ was defined as a weighted combination of the previous trial's interface I^- and the optimal interface I^*:

$$I^+ = \alpha_I I^- + (1 - \alpha_I)I^*, \tag{2.10}$$

where α_I is the adaptation rate. We employed a conservative adaptation rate of $\alpha_I = 0.75$ in this study. We selected this adaptation rate because previous studies using the Smooth Batch algorithm suggested that aggressive adaptation rates may lead to poor convergence (Madduri et al., 2024). A pseudo-code summarizing the steps of the interface update is in **Algorithm 1**.

Algorithm 1 Interface update during co-adaptation

Human completes a trial with I^-.
Input: Human response $\widehat{u}_H$, disturbance $\widehat{d}$, and human transformation H^- of the previous trial.
Output: Next interface I^+.
Compute human transformation (2.5b) $H(j\omega) = \frac{\widehat{u}_H(j\omega)}{\widehat{y}(j\omega)}$.
Update human model as (2.8) $H^+ = \alpha_H H^- + (1 - \alpha_H)H$.
for all I in I_{list} **do**
    Compute $\widehat{y}, \widehat{u}_H, \widehat{u}_I$ (2.6).
    Compute interface cost (2.7): $c_I = \frac{\|\widehat{y}\| + \lambda_I \|\widehat{u}_I\| + \lambda_H \|\widehat{u}_H\|}{\|\widehat{d}\|}$.
end for
Find optimal interface $I^* = \arg\min_I c_I(H^+, I)$.
Update the next interface $I^+ = \alpha_I I^- + (1 - \alpha_I)I^*$.

3.6 Evaluation

We evaluated our experimental results against the following three hypotheses.

H1 (whether co-adaptation occurred): We tested whether humans and interfaces changed after 15 trials of adaptation compared to non-adaptive baselines. For each condition (0th-, 1st-, 2nd-order), we compared the distributions ($N = 11$) of transformations, human $\|H\|$ and interface $\|I\|$, between the co-adapted final interface (averaged of last 3) and the non-adaptive baseline (average of 6 baseline trials) using paired t-tests at each stimulated frequency ($\alpha = 0.05$). The null hypothesis $\mathbf{H1_0}$ is: there is no significant difference in the transformations of either the interface or the humans after adaptation compared to their respective baselines.

H2 (baselines vs. pass-through): We investigated whether the three baselines ($I_{0,\text{base}}$, $I_{1,\text{base}}$, $I_{2,\text{base}}$; Table 2.1) improved from having "no interface" (pass-through with $I = 1$, I_{pass}) in terms of the interface cost $c_I(H, I)$ defined in (2.7). We first confirmed normality of the average cost of each group using Shapiro-Wilk test ($p > 0.05$), then performed one-way ANOVA to compare the groups. We used paired t-tests with Bonferroni correction as post-hoc tests. The null hypothesis $\mathbf{H2_0}$ is: there is no significant difference in performance between baselines and pass-through.

H3 (adaptive vs. non-adaptive): Our primary outcomes of interest were the differences in performance with and without co-adaptation. Because our algorithm explicitly optimizes the

interface cost, we identified the reduction of this cost as the primary indicator of performance improvement. Additionally, we assessed changes in task error, human effort, and interface effort to reflect the practical goal of maximizing task performance while minimizing burdens from both agents. We defined our performance metrics as follows:
1) Task error: norm of the cursor displacement $\|\widehat{y}(j\omega)\|$ at the stimulated frequencies of a trial.
2) Human effort: norm of the manual device inputs $\|\widehat{u}_H(j\omega)\|$ at the stimulated frequencies of a trial.
3) Interface effort: norm of the interface inputs $\|\widehat{u}_I(j\omega)\|$ calculated in (2.5a) at the stimulated frequencies of a trial.
4) Interface cost: $c_I(H, I)$ defined in (2.7) at the stimulated frequencies of a trial.
5) Stability margin: the stability margin of the open loop transfer function $L(j\omega) = H(j\omega)M(j\omega)I(j\omega)$, by calculating the absolute distance between $L(j\omega)$ and the critical point -1 in Nyquist plot. Higher value indicates better stability margin.

We confirmed normality using Shapiro-Wilk test ($p > 0.05$), and used paired t-tests to compare the co-adapted final trials (average of last 3) and baseline trials (average of 6) in each performance metric, for each condition. The null hypothesis $\mathbf{H3_0}$ is: there is no significant difference in each performance metric between co-adapted finals and baselines.

We used `Python 3.11` for Shapiro-Wilks tests, ANOVA test, and paired t-tests. Data and analyses to reproduce the results are available in a publicly-accessible repository (Chou et al., 2026).

3.7 Human system identification

To characterize co-adaptation in simulation, we needed to first identify a parsimonious model of the human feedback controller in post-hoc. In a previous study, we estimated the human controller in 2nd-order machine interaction using the crossover model in (McRuer, Duane T and Jex, Henry R, 1967):

$$H(j\omega) \approx k(j\omega + 1)e^{-j\omega\tau}. \tag{2.11}$$

We estimated the human gain k and sensorimotor delay τ from prior experimental data in (Yamagami et al., 2021a), which had the same disturbance-rejection task but a simpler 2nd-order machine dynamics. We reported in (Peterson et al., 2022) with the average human gain and delay of $k = 2.75$ and $\tau = 0.321$ seconds, respectively.

In this study, since finite-order interfaces were introduced into the system, the crossover model with a 2nd-order machine may no longer apply. However, humans do not implement arbitrarily high-order transformations (Zhang et al., 2018; McRuer, Duane T and Jex, Henry R, 1967). Therefore, we started with models incorporating the lowest possible degrees of dynamics: 1st- and 2nd-order dynamics, each with a human sensorimotor delay of $\tau = 0.321$:

$$\begin{aligned} H_{\text{fit, 1st}}(j\omega) &= \frac{b_0 z + b_1}{z + a_0} e^{-j\omega\tau}, \\ H_{\text{fit, 2nd}}(j\omega) &= \frac{b_0 z^2 + b_1 z + b_2}{z^2 + a_0 z + a_1} e^{-j\omega\tau}. \end{aligned} \tag{2.12}$$

We fit the 1st- and 2nd-order models to the empirical human transformation by minimizing the difference between the modeled and empirical data. Specifically, we performed a grid search of parameters a, b over the range $[-10, 10]$, and performed optimization with the L-BFGS algorithm.

3.8 Co-adaptation simulation

Building upon our theoretical framework and experimental observations, we conducted a simulation to derive a computational model of the HMI co-adaptation. The objective was to identify the simulation parameters that best approximate our empirical results. Since our theory showed that there can be no Nash equilibrium when the order of the human and interface dynamics are unbounded, we constrained the simulation to low-order system dynamics for both human and interface. We evaluated the fit for both 1st- and 2nd-order human models with sensorimotor delay (Section 3.7), and selected the model order that most accurately captured the empirical human transformations. Following prior work (Madduri et al., 2021, 2024), we modeled human adaptation as the minimization of a cost function defined as a linear combination of task performance and human effort, mirroring the structure of the interface cost:

$$c_H(H, I) = \frac{\|(\widehat{y}, \widehat{u}_H)\|}{\|\widehat{d}\|} = \frac{\|\widehat{y}\| + \mu_H \|\widehat{u}_H\|}{\|\widehat{d}\|}. \tag{2.13}$$

We applied an alternating optimization procedure to synthesize the optimal human H^* and interface I^* in simulation. Starting with the baseline I_{base} as the initial interface, we synthesized H^* by minimizing human cost in (2.13). We imposed constraints on the human model to reflect physiological limitations: a phase constraint of $\angle H < \pi$ because humans cannot predict future task signals, and a magnitude constraint of $|H| \leq 2.5$ for frequencies ≥ 0.55 Hz to account for the limited human control bandwidth. Human parameters b_0, b_1, b_2, a_0, a_1 in (2.12) were randomly initialized over the range $[-20, 20]$ and the optimal human H^* was obtained by minimizing the non-convex cost using the Sequential Least Squares Programming (SLSQP).

Subsequently, we held H^* constant and synthesized I^* by minimizing the interface cost $c_I(H, I)$ using the experimental algorithm. This alternating optimization was iterated 15 times, mirroring the 15 trials performed in the empirical study.

We performed a parameter sweep of the penalty terms, λ_I, λ_H in (2.7) and μ_H in (2.13), across the range $[0.5, 10]$. The optimal set of penalties was identified by minimizing the difference between the simulated and the experimental data. Model performance was evaluated using the standardized distance (z-score): $z = \frac{\hat{x}-\bar{x}}{\sigma}$, where $\hat{x}$ denotes the simulated value, and $\bar{x}$ and σ represent the mean and standard deviation of the empirical subject data, respectively. A simulated result was considered consistent with experimental result if $|z| < 1$.

4 Experimental results

Eleven participants completed the disturbance-rejection task in the presence of a dynamic machine in (2.4), over a sequence of 21 trials (3 non-adaptive baselines; 15 co-adaptations; 3 non-adaptive baselines) in each of the three conditions presented in random order: 0th-, 1st, and 2nd-order interfaces (Table 2.1).

4.1 Humans and interfaces co-adapted

We tested our first hypothesis, whether humans and interfaces co-adapted in our experiment (**H1**). We found significant differences between the magnitudes of non-adaptive baseline interfaces (dashed lines; Fig. 2.4a) and the co-adapted final interfaces (blue lines; Fig. 2.4a), where $p < 0.05$ at all eight stimulated frequencies. This implies that in all three conditions, after 15 trials of co-adaptation, interfaces adapted to a different interface than the baseline.

We then examined human adaptation in response to these interface adaptations. We found significant differences between human transformations in baseline trials (dashed lines; Fig. 2.4b) and co-adapted final trials (red lines; Fig. 2.4b). These differences $p < 0.05$ were present across the majority of stimulated frequencies, particularly in the low-frequency range (< 0.55 Hz). Thus, we reject the null hypothesis $\mathbf{H1_0}$.

In addition, we observed a trade-off in magnitude shifts of the human and interface agents. In the 0th-order condition, interface magnitudes decreased relative to baseline while human magnitudes increased. Conversely, in the 1st- and 2nd-order conditions, we observed that interface magnitudes increased, whereas human magnitudes decreased. These results provide evidence of a trade-off between the two decision-making agents as they co-adapt in this dynamic game.

4.2 High pass interfaces improved performance

We tested our second hypothesis, whether the baseline interfaces would yield a different performance than the pass-through interface (**H2**). We defined our primary performance metric as the interface cost $c_I(H, I)$ in (2.7).

One-way ANOVA with post-hoc paired t-tests found no statistically significant difference in interface cost between the 0th-order baseline and the pass-through ($t(10) = 0.097$, corrected $p = 1$). Furthermore, participants subjectively reported in 0th-order baseline trials that the cursor was overly "sensitive", likely due to the large input scaling of 5 ($I_{0,\text{base}} = 5$).

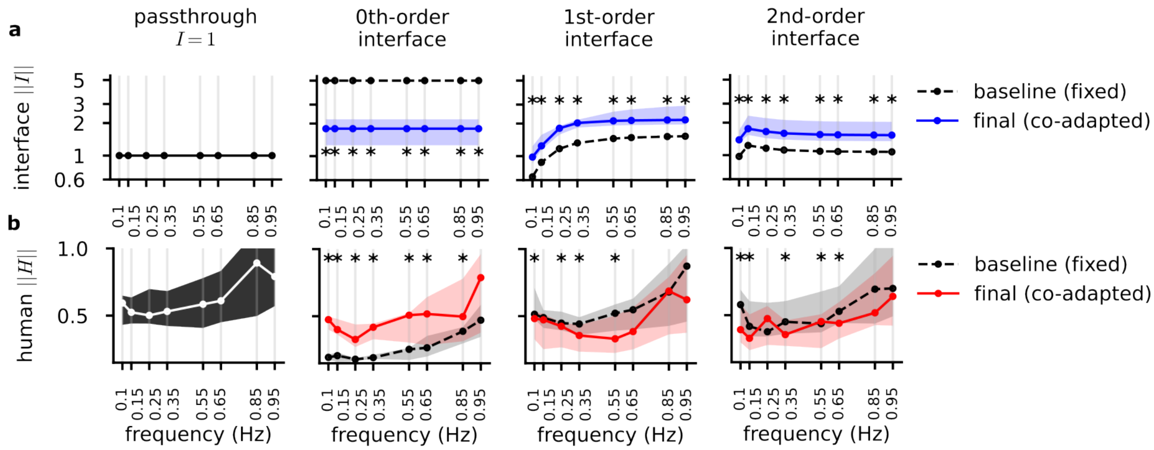

Figure 2.4: *Experiment results on human-interface co-adaptation.* Spectral density plots (median and interquartile, $N = 11$ participants) of **(a)** magnitudes of pass-through interface ($I = 1$), non-adaptive baseline interfaces (dashed lines), and co-adapted final interfaces (blue lines) for the 0th-, 1st-, and 2nd-order conditions at the stimulated frequencies. **(b)** Magnitudes of human transformations in pass-through trials, non-adaptive baseline trials (dashed lines), and the co-adapted final trials (red lines) for the 0th-, 1st-, and 2nd-order interfaces. Statistically significant differences between baselines and co-adapted final transformations are shown with an asterisk (* $p < 0.05$, paired t-test at each stimulated frequency).

In contrast, interface costs for the 1st- and 2nd-order baselines were significantly lower than those of the pass-through condition (1st-order: $t(10) = 8.76$, corrected $p < 0.001$; 2nd-order: $t(10) = 9.98$, corrected $p < 0.001$). These findings demonstrate that the order of the interface dynamics influenced system performance, with higher-order (1st- and 2nd-order) high-pass interfaces yielding better performance relative to the pass-through baseline. Thus, we fail to reject the null hypothesis $\mathbf{H2_0}$ for the 0th-order condition, and reject the null hypothesis for the 1st- and 2nd-order conditions.

4.3 Performance improved or maintained after co-adaptation

We evaluated the effect of co-adaptation using five performance metrics: task error $\|\widehat{y}\|$ (cursor displacements), human effort $\|\widehat{u}_H\|$ (human responses), interface effort $\|\widehat{u}_I\|$ in (2.5a), interface cost $c_I(H, I)$ in (2.7), and the stability margin of the open-loop transfer function $L = HMI$. We tested our third hypothesis, whether the co-adapted interfaces yielded a different performance than the non-adaptive baselines (**H3**).

In the 0th-order condition, task error (Fig. 2.5a) was significantly reduced in the co-adapted final interface compared to the 0th-order baseline ($p < 0.001$). For the 1st- and 2nd-order interfaces, while the median task errors showed a slight decrease, these differences did not show statistical significance (1st-order: $p = 0.14$; 2nd-order: $p = 0.65$). Given that 1st-and 2nd-order baselines significantly outperformed the pass-through condition in task error, these findings suggest that the 1st- and 2nd-order baselines already yielded relatively good task performance, and that task performance was maintained after co-adaptation.

We observed that human effort and interface effort exhibited divergent trends across conditions (Fig. 2.5b, c). Human effort was significantly higher in the co-adapted final interface compared to the baselines for the 0th-order condition ($p < 0.01$), whereas it was significantly reduced in both the 1st- and 2nd-order conditions (1st-order: $p < 0.01$; 2nd-order: $p < 0.01$). In contrast, interface effort was significantly lower in the co-adapted final interface compared to the baselines for the 0th-order condition ($p < 0.001$), and was significantly higher in both the 1st- and 2nd-order conditions (1st-order: $p = 0.044$; 2nd-order: $p = 0.048$).

Regarding the interface cost (Fig. 2.5d), which was the metric explicitly minimized by our algorithm, we observed a significant reduction for the 0th- and 1st-order co-adapted interfaces compared to their respective baselines (0th-order: $p = 0.013$; 1st-order: $p = 0.048$). No statistically significant difference was observed for the 2nd-order condition ($p = 0.1$).

We evaluated robustness by comparing the stability margin of the open-loop transfer function L between the co-adapted final trials and the non-adaptive baselines (Fig. 2.5e). We observed a significantly higher stability margin for the 0th-order co-adapted interface compared to the 0th-order baseline ($p < 0.01$). No significant difference was identified for the 1st-order condition ($p = 0.29$). For the 2nd-order condition, however, the co-adapted interface demonstrated a significant reduction in stability margin relative to the baseline ($p = 0.03$).

In summary, co-adaptation improved or maintained some system performance metrics, including task error and interface cost. The effort between human and interface agents shifted variably across conditions. However, co-adaptation did not necessarily enhance robustness. These findings led us to reject the null hypothesis $\mathbf{H3_0}$.

5 Computational results

Following the system identification procedure (Section 3.7), we modeled the human dynamics using both 1st- and 2nd-order system architectures with sensorimotor delay. We found that the 2nd-order model in (2.12) fit the empirical distribution well (Fig. 2.6a). The two distributions (magnitudes of H_{fit} and H_{exp}) exhibited Pearson correlation coefficients of at least $r > 0.88$ at every stimulated frequency ($p < 0.001$). The 1st-order model had a relatively poorer fit than the 2nd-order model, with Pearson correlation coefficients of at least $r > 0.77$ at every stimulated frequency ($p < 0.01$).

We therefore used the 2nd-order human model with sensorimotor delay to simulate a human-like agent interacting with the 1st-order adaptive interface in the experiment. In this simulation, we modeled human behavior by adopting a cost function analogous to the interface agent's optimization strategy, defined in (2.13). We simulated the co-adaptation process exclusively for the 1st-order interface. This selection was based on experimental findings

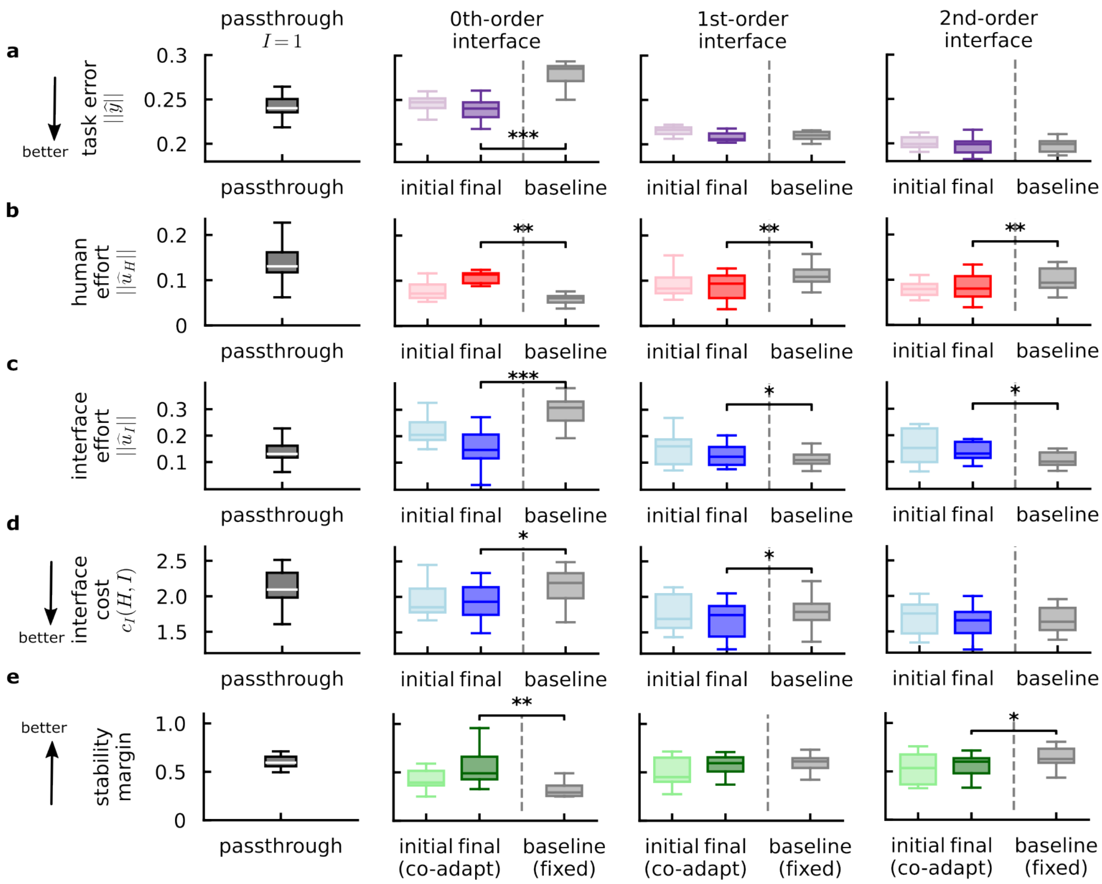

Figure 2.5: *Experiment results on performance.* Boxplots (0th, 25th, 50th, 75th, and 100th percentiles, $N = 11$ participants) of each performance metric for pass-through (average of 12 trials), baselines (average of 6 trials), and adaptive interfaces (average of initial 3 trials and average of final 3 trials). **(a)** task error $\|\widehat{y}\|$ (lower values are better), **(b)** human effort $\|\widehat{u_H}\|$, **(c)** interface effort $\|\widehat{u_I}\|$, **(d)** interface cost $c_I(H, I)$ defined in (2.7) (lower values are better), and **(e)** stability margin (higher values are better). Statistically significant differences are shown with horizontal lines and asterisks (* for $p < 0.05$, ** for $p < 0.01$, and *** for $p < 0.001$; paired t-test).

that only the 1st-order interface achieved significant improvements in interface cost relative to both the baseline ($p = 0.048$) and the pass-through interface ($p < 0.001$) (Fig. 2.5d).

We conducted a parameter sweep over the penalty terms (λ_I, λ_H in (2.7), and μ_H in (2.13)) in range $[0.5, 10]$ and identified the optimal configuration that minimized the deviation from empirical data ($\lambda_I = 1.32, \lambda_H = 5.5, \mu_H = 2.04$). The fit was validated by calculating the standardized distance (z-score) between the simulated and experimental magnitudes at each stimulated frequency. We found that $|z| < 1$ across all stimulated frequencies for the interface magnitudes ($\|I\|$, blue in Fig. 2.6b). Similarly, for the human magnitude ($\|H\|$, red in Fig. 2.6b), $|z| < 1$ across all frequencies except 0.85 Hz. These results indicate that the simulated result consistently remained within one standard deviation of the experimental mean, thus demonstrating good correspondence between the model and empirical outcomes.

We analyzed the evolution of the cost function components $\widehat{u}_I$, $\widehat{u}_H$, and $\widehat{y}$, normalized to each trial's difference to the initial trial (Fig. 2.6c). Simulation results indicate that as interface effort gradually increased and task error decreased, human effort increased and then declined slightly, with the system reaching equilibrium after 15 trials of co-adaptation. These simulation trends demonstrated the effort and error trade-off.

With the penalty term for human effort ($\mu_H = 2.04$), we computed the inferred human cost for the 1st-order interface condition in our experiment. It is important to note that these values represent model-based inferences rather than direct measurements from the experimental data. We observed that the inferred human cost significantly decreased after co-adaptation compared to the non-adaptive baseline (Fig. 2.6d). This finding demonstrates that our proposed framework with a specific, simplified human cost function structure replicated the observed experimental outcomes.

We swept through the penalty terms $\mu_H, \lambda_H, \lambda_I$ in range $[0.05, 8]$ and computed the moving average of interface effort, human effort, and task error. In this normalized plot, we observed clear trade-offs between effort and error (interface effort and task error; human effort and task error) (Fig. 2.7). However, the trade-offs between interface effort and human effort were less clear in this simulation, suggesting that the relationship between these two signals is more complex and may depend on the specific values of the penalty terms.

6 Discussion

Machine learning algorithms are being deployed in-the-loop with people (Nothwang et al., 2016). It is important to predict the outcome of these interactions to guarantee safety and performance in disparate applications including brain-machine interfaces (Orsborn et al., 2012; Madduri et al., 2024; Miehlbradt et al., 2018), human-robot interaction (Zemmar et al., 2020; Bengler et al., 2014; Gomes, 2011; Li et al., 2019) and assistive devices (Zhang et al., 2017a; Zhuang et al., 2019a; Felt et al., 2015; Slade et al., 2022; Azimi et al., 2021; Mahmood et al., 2019). These interactions are increasingly viewed through the lens of game theory (Nikolaidis et al., 2017; De Santis, 2021; Müller et al., 2017; Perdikis and Millan, 2020; Li et al., 2019, 2016; Madduri et al., 2021, 2024; Chasnov et al., 2023; Merel et al., 2013; Isa et al., 2025), where it is hypothesized that the human's behavior arises from the solution of an optimization problem (Diedrichsen et al., 2010; Todorov and Jordan, 2002; Selinger et al., 2019; Emken et al., 2007; Tversky and Kahneman, 1974) and the adaptive agent similarly optimizes its own cost. However, prior work makes assumptions on the ability of people to solve these optimization problems that we regard as unreasonable.

In this study, we sought to systematically test the outcome of HMI co-adaptation and com-

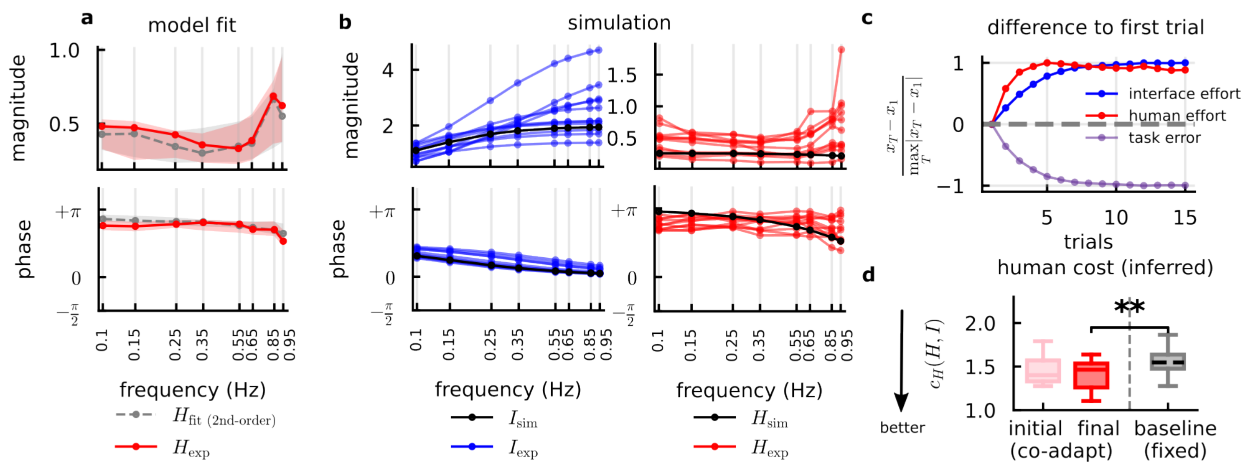

Figure 2.6: *Computational results with a 1st-order interface.* **(a)** Spectral density plots with distributions (median and interquartile, $N = 11$ participants) of (top) magnitudes and (bottom) phases of the 2nd-order model fit H_{fit} and the empirical human transformation H_{exp}. **(b)** Spectral density plots of the (left) simulated interface I_{sim} and (right) simulated human H_{sim} overlaid on empirical data $I_{\text{exp}}, H_{\text{exp}}$ of every subject ($N = 11$ participants). Each line is the average of 3 co-adapted final trials from one of the subjects. **(c)** Difference between each trial T and the first trial $T = 1$ in simulation, for each element in the interface cost defined in (2.7): interface effort, human effort, and task error. Trial-wise metrics were plotted as deviations from the first trial, normalized by the maximum deviation, so the largest absolute change from baseline corresponds to ± 1. **(d)** Boxplot (0th, 25th, 50th, 75th, and 100th percentiles, $N = 11$ participants) of the inferred human cost, with the μ_H value found in simulation, of 1st-order baseline and the 1st-order adaptive trials (initial and final). Statistically significant difference is shown with a horizontal line and asterisks (** for $p < 0.01$; paired t-test).

pare experimental results with predictions from theory and simulations. We adopted a human-in-the-loop control paradigm wherein participants stabilized a complex (nonminimum-phase) control system through a dynamic interface as in Fig. 2.3. This paradigm was intended to instantiate the fundamental challenges inherent in more embodied applications involving neural measurements or physical interaction with a robot or assistive device. Our testbed enables complete control over the system dynamics and external disturbances that facilitate precise comparison of experimental results with theoretical and computational predictions.

6.1 Limitations of theoretical predictions

A leading hypothesis in the human motor control literature is that people solve an $\mathcal{H}_2$ optimal control problem and implement a *linear-quadratic Gaussian* (LQG) controller to move their bodies (Diedrichsen et al., 2010; Todorov and Jordan, 2002). Prior work on co-adaptation (Merel et al., 2013; Li et al., 2019) specialized this model further by considering

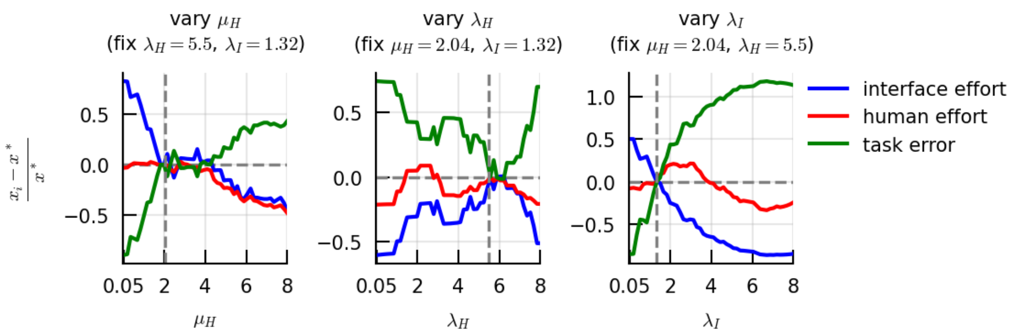

Figure 2.7: *Tradeoffs between signals in simulation.* Moving average of interface effort, human effort, and task error when sweeping penalty terms: (*left*) μ_H, (*middle*) λ_H, and (*right*) λ_I, normalized to each signal in the simulation result in Fig. 7. x_i is the norm of the each signal given the varying penalty terms, x^* is the norm of each signal given the penalty terms set ($\lambda_I = 1.32, \lambda_H = 5.5, \mu_H = 2.04$; gray vertical dashed lines) that yielded simulation results that best matched the empirical data. The penalty parameters were swept in range $[0.05, 8]$, with 50 equidistant points. Here we display the moving average of each signal with window of 10 data points.

the *full information* case where the human is assumed to have noise-free observations of the system state. The full information case has the appeal that it admits a Nash equilibrium defined in terms of static state feedback matrices (Başar and Olsder, 1998, Sec. 6.2.2), but we argue that it is unrealistic to assume each agent has perfect knowledge of both the environment and the other agent.

Considering the more general model without full information (or full control), we find a different theoretical objection: if both the human and adaptive system solve $\mathcal{H}_2$ optimization problems, there can be no Nash equilibrium for the co-adaptive system defined in terms of finite-order LQG controllers. The reason for this is a simple counting argument: since an LQG controller generally has the same order as the system it controls, and since the system one agent controls includes the dynamics of the other agent as in Fig. 2.2, it is not possible for these system orders to “add up” to consistent quantities. Of course the Nash equilibrium is not the only candidate outcome (Chasnov et al., 2023). But this thought experiment reveals an even more fundamental objection to this model for human behavior: it is simply not plausible that people are able to implement LQG controllers of arbitrary system order – they must have bounded rationality (Tversky and Kahneman, 1974; Gershman et al., 2015). These observations motivated us to conduct an experiment with adaptive systems that had fixed system orders toward the goal of understanding how people actually play this class of human-in-the-loop game.

6.2 Co-adaptation shaped behavior and performance

We observed co-adaptation in our experiments: the baseline and final co-adapted human and interface transformations were different for all interface orders in Fig. 2.4. Of course, the interface adapts by design – it repeatedly solves the optimization problem in (2.9). But there is no guarantee that (i) the human will adapt in response and (ii) the outcome of this interaction will differ from baseline. The fact that we observed both (i) and (ii) is particularly significant since we started with a baseline that was synthesized using prior knowledge of human behavior (Peterson et al., 2022), i.e., our baseline was not naïve.

Our experimental results demonstrate potential advantages – and drawbacks – of deploying adaptive interfaces in-the-loop with humans and dynamic machines. Since our interface sought to minimize the cost in (2.7), it seems natural to expect this cost to decrease in the experiment. However, this decrease is not guaranteed due to the presence of human adaptation, as co-adaptive interactions may converge to a constellation of equilibrium outcomes (Chasnov et al., 2023). We observed significant decreases in this cost for the 0th- and 1st-order interfaces but not for the 2nd-order interface (Fig. 2.5d), confirming performance gains are not guaranteed.

A similar observation holds for the stability margin of the closed-loop system: solving the interface's $\mathcal{H}_2$ optimization problem can yield arbitrarily small margins in theory (Doyle, 1978). Experimentally, we found a significant increase, decrease, or no change in the stability margin across the different interface orders (Fig. 2.5e). Future work may seek to prioritize robustness using $\mathcal{H}_\infty$ controller synthesis (Doyle et al., 1989; Zhou et al., 1996; Benner et al., 1999; Nakahira et al., 2021).

A consistent finding across the three interface conditions was that the human and interface efforts traded off, as an increase in one was accompanied by a decrease in the other (Fig. 2.5b,c). This outcome comports with other theoretical and experimental findings (Madduri et al., 2021, 2024) and supports the hypothesis that the human and interface play a *game* (Chasnov et al., 2023; Madduri et al., 2023) whose outcome represents a tradeoff between competing objectives.

6.3 A computational model for human adaptation

We found that the transformations our human subjects implemented in-the-loop with a 1st-order interface could be well-approximated as a second-order system with sensorimotor delay (Fig. 2.6a). Importantly, the dimension of the state for the human's transformation H in this case (2) is less than the combined dimension of the machine M and interface I $(2 + 1 = 3)$, undermining the hypothesis that the human implements the LQG controller whose state dimension would match that of the combined machine-interface system M/I. This finding is consistent with prior work and supports the hypothesis that humans do not implement transformations of arbitrarily-high order (Zhang et al., 2018; McRuer, Duane T and Jex, Henry R, 1967).

We found that modeling the human as optimizing a cost function consisting of a linear combination of error and effort – inspired by Emken *et al.* (Emken et al., 2007) – can produce similar outcomes to what we observed in our experiment (Fig. 2.6b) and yield an error-effort tradeoff (Fig. 2.6c). We regard this result as provocative but far from definitive. In particular, we do not believe our human subjects literally optimized a cost function with the precise parameterization in (2.13), as human motor control can be influenced by cognitive load, fatigue, injury, or distractions (Hemond et al., 2010; Cole and Shields, 2019; Branscheidt et al., 2019). All models are wrong, but some are useful (Box and Draper, 1987); interrogating the utility of our proposed model is an important goal for future work.

6.4 Principles for interface design

Motivated by the observation that human sensorimotor delay constrains performance in trajectory-tracking tasks (Nakahira et al., 2021), we employed high-pass interfaces in our experiments. Indeed, the human brain attempts to compensate for delay internally through feedback projections, but this compensation is imperfect and is subject to speed-accuracy trade-offs (Li et al., 2023). Adding a high-pass interface can supplement this phase-lead mechanism in the human's internal delay compensation, potentially improving performance of the human-in-the-loop system. Fig. 2.5a,d demonstrates that the high-pass baselines (1st- and 2nd-orders) and co-adapted final interfaces decreased task error and interface cost compared to the pass-through.

If human behavior in co-adaptive interactions arises from optimization as corroborated by our simulation results, it opens the possibility of systematically shaping interaction outcomes by modifying the interface cost function. A proof-of-concept of this idea was demonstrated in (Madduri et al., 2024), where it was shown that increasing a penalty term in the interface cost decreased interface effort and increased human effort. This observation has broad implications for the design of both assistive devices, which seek to minimize user effort, and rehabilitation devices, which aim to promote active patient engagement (Cha et al., 2025; Bonanno et al., 2025). Game-theoretic methods that explicitly consider the strategic interests of both decision-making agents may enable designers to systematically select from a variety of outcomes (Chasnov et al., 2023).

7 Conclusion

Our work highlights limitations in prior theory, experiment, and simulation work on co-adaptation between humans and intelligent interfaces interacting in closed loop with dynamic machines in real-world conditions. We demonstrate that a Nash equilibrium does not exist under standard assumptions of optimal adaptation, and explore in experiment and simulation how a human and interface with *bounded rationality* (Tversky and Kahneman, 1974) adapt to improve performance in a disturbance-rejection task. Understanding how to analyze and synthesize co-adaptive systems, where both human and algorithmic agents adapt to control dynamic systems like vehicles, robots, and prostheses is critically important in current and emerging applications including driver assistance, rehabilitation robotics, and neuroprosthetics. We contribute new theory, experiment, and simulation results that will inform creation of these systems for real-world deployment.

Chapter 3

MULTIMODAL HUMAN-MACHINE INTERACTION

Redundant Sensorimotor Pathways in a Multimodal Neuromotor Interface

Amber H.Y. Chou, Momona Yamagami, Samuel A. Burden

In preparation.

Evaluating a Human/Machine Interface with Redundant Motor Modalities for Trajectory-Tracking

Amber H.Y. Chou, Momona Yamagami, Samuel A. Burden

4th IFAC Conference on Cyber-Physical Human Systems. IFAC-PapersOnLine, 55(41), pp.125–130. 2022. `https://doi.org/10.1016/j.ifacol.2023.01.114`

Abstract

Motor redundancy enables the human control system to find abundant solutions to motor tasks. However, how multiple distinct motor modalities are integrated while humans learn to interface with dynamic machines and computing devices like vehicles or touchscreens remains unanswered. Our motivation to consider disparate modalities comes from the observation that neurologic disorders can limit use of conventional manual interfaces that rely on mobility (e.g. joysticks, touchscreens), whereas central or peripheral neural interfaces (e.g. invasive intracortical recording, non-invasive electromyography) have the potential to enable interaction in the presence of injury or disease. Prior work focused on humans coordinating redundant channels of a single modality, whereas we seek to understand how people integrate multiple distinct modalities. In this work, we integrated kinematic (manual joystick) and neural (muscle EMG) modalities to collectively control a 1-degree-of-freedom (DOF) trajectory-tracking task. Our central hypothesis is that redundant modalities enable the motor system to overcome limitations of single-mode interfaces. We introduced novel control-theoretic methods based on linearity and the Uncontrolled Manifold (UCM) hypothesis to analyze sensorimotor noise and demonstrated that humans suppress the effect of noise by coordinating redundant pathways.

1 Introduction

In human-machine interfaces (HMI), humans interact and communicate with dynamic machines through one or more *sensory* and *motor* modality. *Sensory* modalities enable humans to perceive stimuli from machines (e.g., a computer display that provides visual stimuli to human eyes), while *motor* modalities enable humans to provide motor inputs to machines (e.g., a joystick that measures hand movement on one or more axes). We refer to traditional hand-held devices such as joysticks, mice, and steering wheels as *manual* modalities.

We are interested in understanding whether integrating both motor modalities could combine benefits and compensate for weaknesses from both sides, and thus improve the performance of HMI. Recently, researchers have begun implementing biological signals like muscle activity as an alternative motor input. For instance, electromyographic (EMG) signals are widely used in developing body-machine interfaces (BoMI) (Casadio et al., 2012; Antuvan et al., 2014), classifying gestures and motions (Saponas et al., 2009) (always-available inputs), and controlling assistive prosthetic devices (Kiguchi and Hayashi, 2012; La Scaleia et al., 2014; Fall et al., 2017) and unmanned aerial vehicles (Miehlbradt et al., 2018). We refer to the recordings of muscle inputs in HMI as *muscle* modalities in this paper. Both muscle and manual modalities have unique advantages as human interfaces. Interfaces with muscle modality (i.e., muscle interfaces) provide high-density measurements (Drost et al., 2006; Yang et al., 2025), reduce delays in motion (Artemiadis, 2012), and have large control bandwidth (Yamagami et al., 2020; Lobo-Prat et al., 2014; Corbett et al., 2011). However, there are still many challenges in current upper-limb muscle interfaces such as high signal variability between and within individuals, and high sensitivity to perturbations (Artemiadis, 2012; Farina et al., 2014). However, there are still many challenges in current upper-limb muscle interfaces such as high signal variability between and within individuals, and high sensitivity to perturbations (Farina et al., 2014; Artemiadis, 2012). On the other hand, manual interfaces have lower control bandwidth but provide less variable signals and are currently much more familiar to users. Fusing the inputs from both motor modalities has the potential to combine benefits and compensate for weaknesses from both sides (Pantic and Rothkrantz, 2003; Rizzoglio et al., 2020).

To develop human-machine interfaces with redundant motor modalities, it is instructive to first understand how humans and other animals integrate motor pathways in their sensorimotor systems. Previously, researchers have assessed the integration of parallel *sensory* pathways (Roth et al., 2016; Peterka, 2018) and *motor* pathways (Gelfand and Latash, 1998; Scholz and Schöner, 1999; Latash et al., 2002) to explain how redundant inputs collectively govern a single behavior. Researchers have also investigated motor redundancy using body-machine interfaces (Ranganathan et al., 2014; De Santis and Mussa-Ivaldi, 2020), in which body inputs were collected from multiple channels with the same modality to control lower-dimensional systems. Studies have found that sensorimotor redundancy could provide abundant solutions to achieve tasks (Todorov, 2004), provide stability (Latash, 2012), assure robustness against uncertainties (Gelfand and Latash, 1998; Roth et al., 2016), and enhance movement efficiency (De Santis and Mussa-Ivaldi, 2020). Stability here meant "the capacity of the system to return to a given state after a (phasic) perturbation has driven the system away from that state" (Scholz and Schöner, 1999). In this work, we extend these results to investigate how the inherent redundancies in human sensorimotor control could be integrated in designing multimodal interfaces to improve the performance of human-machine interaction.

The advantages of multimodal interfaces have been suggested in previous studies in the fields of human- and brain-computer interfaces. Multimodal interfaces can incorporate diverse features from different modalities, and thus enable human-computer interfaces to better analyze complex human behaviors (Pantic and Rothkrantz, 2003; Jaimes and Sebe, 2007), recognize hand gestures (Wu et al., 2016), and classify intent of movement (Zhang et al., 2019). Similarly, fusing multiple data sources with different frequency bands could enhance the accuracy of brain decoders (Fazli et al., 2015) and improve usability of brain-computer interfaces (Müller-Putz et al., 2011). Most relevant for the present study is the muscle-and-motion interface studied in (Rizzoglio et al., 2020), where it was found that the multimodal (*hybrid*) interface performed similarly to motion-only (implemented with inertial measurement units, IMU) and outperformed muscle-only (implemented with surface electromyography, sEMG) in a reaching task.

The objective of this study is to evaluate a human-machine interface that combines redundant muscle and manual motor modalities to perform a second-order reference-tracking and disturbance-rejection task. Our experimental results with 15 human subjects demonstrate that the multimodal interface outperforms the traditional manual interface and performs comparably to the muscle interface. We also found that users can suppress the effect of sensorimotor noise in task-relevant dimensions by coordinating noise between muscle and manual modalities. Our results demonstrate potential advantages for human-machine interfaces that incorporate motor redundancy.

2 Background

In *human-in-the-loop* systems, humans provide inputs to control dynamic machines and receive feedback to learn and adapt their actions. This creates a closed-loop interaction. Modeling human behaviors, such as intent, strategies, and preference, in this closed-loop can provide insights for designing personalized human-in-the-loop systems. In this study, we model how humans behave in-the-loop when controlling dynamic machines with multimodal inputs.

2.1 Decoding Human Controllers

Using theoretical frameworks in control theory, we can model humans as controllers in a closed-loop system.

Previous studies (Yamagami et al., 2021b) assumed humans H behave like LTI controllers if given a LTI machine M and stimuli in a specific range of frequency. Assuming this assumption still holds for HMI, the output y can be written as the linear combination in response to reference $\widehat{r}$ and disturbance $\widehat{d}$:

$$\widehat{y} = \widehat{T}_{yr}\ \widehat{r} + \widehat{T}_{y(Md)}\ (\widehat{M}\widehat{d}) \tag{3.1}$$

where $\widehat{T}_{yr}$ and $\widehat{T}_{y(Md)}$ are LTI transfer functions for signals in the output space. We can then estimate these transfer functions empirically. Since reference and disturbance stimuli are interleaved at the stimulated frequencies ω, when $\widehat{d}(\omega) = 0$,

$$\widehat{T}_{yr}(\omega) = \frac{\widehat{y}(\omega)}{\widehat{r}(\omega)}, \tag{3.2}$$

and when $\widehat{r}(\omega) = 0$,

$$\widehat{T}_{y(Md)}(\omega) = \frac{\widehat{y}(\omega)}{\widehat{M}\widehat{d}(\omega)} \tag{3.3}$$

respectively. The feedback controller B can be estimated using $\widehat{T}_{ud}(\omega)$, where

$$\widehat{B_0}(\omega) = -\widehat{M}^{-1}(\omega)\frac{\widehat{T_{u_0d}}(\omega)}{1+\alpha\widehat{T_{u_0d}}(\omega)+(1-\alpha)\widehat{T_{u_1d}}(\omega)} \tag{3.4a}$$

$$\widehat{B_1}(\omega) = -\widehat{M}^{-1}(\omega)\frac{\widehat{T_{u_1d}}(\omega)}{1+\alpha\widehat{T_{u_0d}}(\omega)+(1-\alpha)\widehat{T_{u_1d}}(\omega)} \tag{3.4b}$$

$$\begin{aligned}\widehat{F_0}(\omega) = -\widehat{B_0}(\omega) + \widehat{T_{u_0r}}(\omega) + \alpha\widehat{B_0}(\omega)\widehat{M}(\omega)\widehat{T_{u_0r}}(\omega) +\\ (1-\alpha)\widehat{B_0}(\omega)\widehat{M}(\omega)\widehat{T_{u_1r}}(\omega).\end{aligned} \tag{3.5a}$$

$$\begin{aligned}\widehat{F_1}(\omega) = -\widehat{B_1}(\omega) + \widehat{T_{u_1r}}(\omega) + \alpha\widehat{B_1}(\omega)\widehat{M}(\omega)\widehat{T_{u_0r}}(\omega) +\\ (1-\alpha)\widehat{B_1}(\omega)\widehat{M}(\omega)\widehat{T_{u_1r}}(\omega).\end{aligned} \tag{3.5b}$$

The open loop transfer function of this system is defined as

$$\widehat{L}(\omega) = \widehat{B}(\omega)\widehat{M}(\omega). \tag{3.6}$$

We used the McRuer's model to select the crossover frequency of our feedback system (McRuer and Jex, 1967; Yu et al., 2014a). The crossover is selected to be 0.25 Hz, which is the first stimulated frequency that has the gain of open-loop transfer function below 1, i.e., $|\widehat{L}(\omega)| < 1$ (Yamagami et al., 2021b).

2.2 Crossover Frequency

According to the McRuer's crossover model, the ability of human operators to track and reject stimuli degrades above the crossover frequency due to humans' neuro-muscular limits (McRuer and Jex, 1967; Yu et al., 2014a). We used the McRuer's model to select the crossover frequency of our feedback system. The crossover is selected to be 0.25Hz, which is the first stimulated frequency that has the gain of open loop transfer function below 1, $|\widehat{L}(\omega)| < 1$ (Yamagami et al., 2021b). However, prior work discovered an improvement of performance above crossover using a muscle interface (Yamagami et al., 2020). Thus, we investigated the task performance at all stimuli (below and above crossover) in this study, but used the notion of crossover to investigate the effect of inputs from the individual modalities (u_{EMG} and u_{Manual}).

2.3 Sensorimotor Noises

We followed the analyses in Yamagami et al. (2021b) and defined sensorimotor noise as the imputed disturbance $\hat{\delta}$, which was the input disturbance if removing the effect of reference and disturbance stimuli in the feedback system:

$$\hat{\delta} = (1 + \widehat{L})\hat{u}_{\mathrm{x}}, \tag{3.7}$$

where $\hat{u}_{\mathrm{x}}$ is the user input at the non-stimulated frequencies below 1 Hz. We linearly interpolated $\hat{L}(\omega)$ at the non-stimulated frequencies to get $\hat{L}$ since the transformation was only derived at the stimulated frequencies (eq.3.6). The sensorimotor noises at both the *EMG* and *Manual* pathways can be calculated using each modality's inputs, $\hat{u}_{\mathrm{EMG}}$ and $\hat{u}_{\mathrm{Manual}}$, and the feedback transformations, B_{EMG} and B_{Manual}.

3 Methods

We developed a multimodal interface with muscle (EMG) and manual (joystick) modalities. We then conducted human-subject experiments and data analysis to empirically evaluate performance of the multimodal interface.

3.1 Multimodal Interface

In this study, the EMG electrodes were placed on participants' dominant arms and the manual interface was placed in front of their non-dominant hands (Fig. 3.1A). In conditions that involve muscle modality, participants laid their non-dominant arm on a hard surface and controlled the cursor by activating their biceps (pulling up) or triceps (pushing down). It was previously discovered that handedness does not affect learned controllers using a manual interface (Yamagami et al., 2021b). We picked the non-dominant arm for the muscle modality because prior work has shown that users performed significantly better with a muscle interface versus a manual interface at higher frequencies using their non-dominant arms (Yamagami et al., 2020).

The manual signals were measured using a 1-DOF 10 $k\Omega$ slide potentiometer. The slider's handle has a travel distance of 11.4 cm. The muscle signals were obtained via surface EMG electrodes (Muscle SpikerShields, Backyard Brains, Inc.). The EMG electrodes were placed on participants' biceps and triceps according to the guidelines of Surface Electromyography for the Non-Invasive Assessment of Muscle (SENIAM) (Hermens et al., 2000). The potentiometer values and the EMG data were recorded with an Arduino Uno (`Arduino.cc`) at 1000 Hz.

An EMG calibration was conducted at the beginning of the experiment. We asked each participant to do maximum voluntary contraction (MVC) of their biceps and triceps separately for three times. Each MVC was held for two seconds. We took the 95th percentile of the two-second data and then averaged the three trials to get an averaged MVC. We converted the raw EMG data to the Average Rectified Value (ARV) by taking the absolute value of raw EMG signals and applying a moving average filter with a 200 ms window. We then normalized the ARV with the average MVC of each participant. To prevent muscle fatigue, only a maximum of 33% of the MVC was required to track the largest stimuli of the task. The normalized ARV of biceps and triceps were used to generate the EMG input $u_{EMG} \in [-1, 1]$. The u_{EMG} was returned positive if biceps' scaled value was larger than triceps', and negative if the other way around. The u_{EMG} was zero if both biceps' and triceps' scaled values are below 2.5% of the average maximum contraction.

3.2 Participants

We recruited 15 participants (6 female, 9 male; 14 right-handed, 1 left-handed; aged between 22 to 31 years; height: 177±15.25 cm; weight: 71.5±19 kg). Participants gave their written consent prior to the study. This study was approved by the University of Washington's Institutional Review Board (IRB #STUDY00000909). All participants use computers and smartphones daily. Two participants reported themselves as non-professional experts in video games and eleven reported themselves as causal gamers. Six participants have seen or used EMG devices, and two participants work with EMG regularly. None of them had prior experience with the sEMG device we used in this study and they were not aware of the study's purpose or conditions.

3.3 Continuous Tracking Task

In the following sections, we denote a signal and a transfer function as q and T in time-domain; and $\widehat{q}$ and $\widehat{T}$ in frequency-domain. The frequency-domain representations of signals are calculated using the fast Fourier Transform (FFT). We adopted a 1-degree-of-freedom (DOF) continuous human-machine task previously conducted by (Yamagami et al., 2021b). Human users H are tasked with controlling a second-order linear-time invariant (LTI) ma-

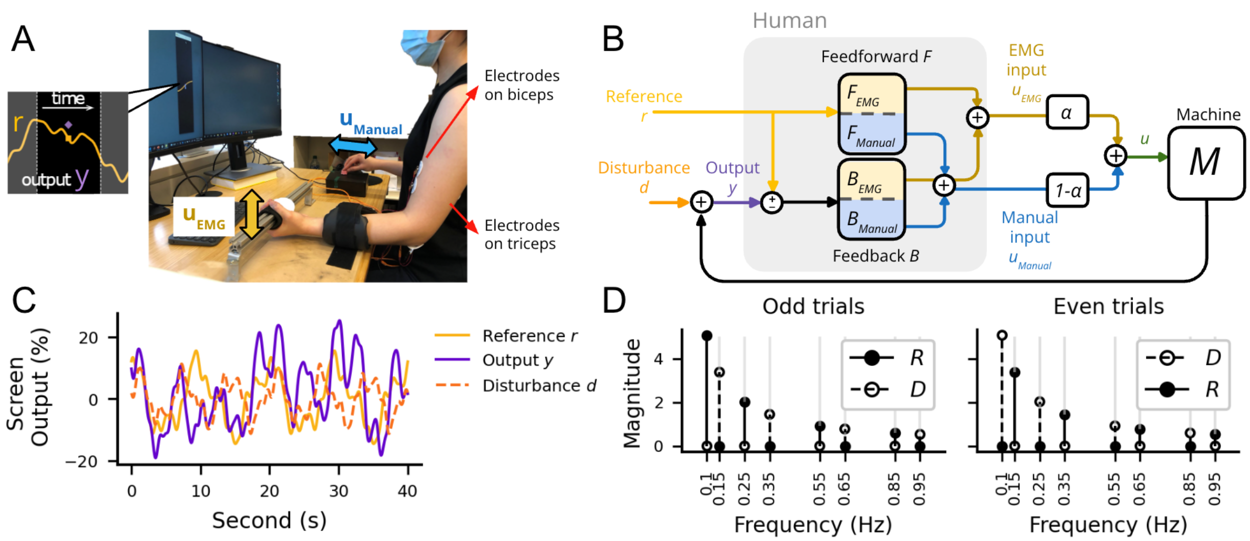

Figure 3.1: (A) Multimodal human-machine interface: reference trajectory r and cursor output y are displayed on a computer screen; participants used their non-dominant arm to control the muscle modality (left arm in this image) and the dominant hand to control the manual modality (right hand in this image). Electrodes were placed on participants' biceps and triceps. (B) Block diagram of the multimodal human-machine interface. The human user transforms reference r and output y to produce *muscle* and *manual* inputs, u_{EMG} and u_{Manual}, respectively. The inputs are scaled by α and summed to yield input u to the second-order machine M. The feedforward controller F is the transformation between reference r and input u, and the feedback controller B is the transformation between tracking error $r - y$ and input u. (C) Example reference r, cursor output y, and disturbance signals d in a trial. (D) Reference and disturbance stimuli were interleaved at the stimulated frequencies. (Top row) magnitudes of signals in the output space at the stimulated frequencies; (Bottom row) magnitudes of signals in the input space at the stimulated frequencies. (Left column) In odd trials (1, 3, 5, ...), reference and disturbance stimuli happened at odd and even frequencies, respectively. (Right column) In even trials (2, 4, 6, ...), reference and disturbance stimuli happened at even and odd frequencies, respectively.

chine dynamics M:

$$M : \ddot{y} + \dot{y} = u, \quad \widehat{M} : \frac{1}{s^2 + s}, \tag{3.8}$$

where $s = j\omega$. We chose a second-order M in this study since previous studies suggest that muscle and manual modalities perform differently during acceleration-based tasks (Yamagami et al., 2020).

Second-order task dynamics were chosen because previous studies suggest that muscle and manual modalities perform differently during acceleration-based tasks (Yamagami et al., 2020).

The objective of human users is to control their cursor output y to track a reference trajectory r on a computer screen and reject the invisible input disturbance d. Users control the system using both muscle and manual modalities through a custom interface that records motor inputs. We define the muscle and manual input signals as u_μ and u_ν, respectively. The two inputs were linearly combined with a weight constant, α:

$$u = \alpha \ u_{EMG} + (1 - \alpha) \ u_{Manual} \tag{3.9}$$

To investigate the optimal weighting of the two modalities, we selected α to be evenly spaced numbers in [0, 1], where $\alpha \in \{0, 0.25, 0.5, 0.75, 1\}$. We refer to the weighted sum u as the overall user input, and the transformation between tracking error $r - y$ and u as the feedback controller B. A block diagram of the human-machine system is shown in Figure 3.1B. The time-domain reference r, disturbance d, cursor output y, and user inputs (u, u_μ and u_ν) are collected at sampling rate of 60 Hz. The task was programmed using `pygame 1.9.2` package.

3.4 Task Stimuli

We constructed unpredictable reference r and disturbance d as sums of sinusoidal signals interleaved at the stimulated frequencies (Yamagami et al., 2021b; Yu et al., 2014a). Stimulated frequencies Ω were designed to be eight prime multiples of the base frequency (0.05 Hz) below 1 Hz, where $\Omega = \{0.1, 0.15, 0.25, 0.35, 0.55, 0.65, 0.85, 0.95 \text{ Hz}\}$ to avoid harmonics (Roth et al., 2011). The magnitude of each stimulus was scaled by the inverse of its frequency, and the phase was randomized. The magnitudes of stimuli in the output space, $|\widehat{r}|$ and $|\widehat{M}\widehat{d}|$, were scaled by its frequency, and random phases were introduced for each stimuli. An illustration of magnitudes of reference and disturbance stimuli are shown in figure 3.1. Reference trajectories were provided to users with a look-ahead time of two seconds in order to allow feedforward control of human users (Sheffler et al., 2019; van der El et al., 2016).

3.5 Experiment Conditions

Participants were asked to conduct the continuous tracking task by minimizing the distance between cursor output y and reference trajectory r and rejecting the invisible disturbance *d*. A score (time-domain mean-squared-error) indicating participants' performance was displayed on screen after they completed each trial. Participants were asked to perform the task with both of their upper limbs: manipulating the slider with their dominant hand and activating the biceps or triceps of their non-dominant arm.

Each participant performed three sessions in the following order (Table 3.1):
1^{st} **session**: 10 trials for each single-mode condition (Manual-only $\alpha = 0$, followed by EMG-only $\alpha = 1$);
2^{nd} **session**: 14 trials for each multimodal condition ($\alpha = 0.25,\ 0.5,\ 0.75$ in random order);
3^{rd} **session**: 4 trials for each single-mode condition ($\alpha = 0,\ 1$ in random order) as the control groups.

Table 3.1: Condition orders

Sessions	# of Trials	α	%EMG	random order
1^{st}	10	0, 1	0,100	no
2^{nd}	14	0.5, 0.25, 0.75	25, 50, 75	yes
3^{rd}	4	0, 1	0,100	yes

We started each experiment with the manual interface ($\alpha = 0$) to allow participants to familiarize themselves with the task since most computer users are more familiar with a manual device. Each trial was 45 seconds starting with a 5 seconds ramp-up period (ramp-up was not included in the data analysis), thus the total time period $T = 40\ s$. Participants were asked to take mandatory breaks in-between conditions to avoid muscle and eye fatigue.

3.6 Evaluation

We evaluated the performance of multimodal interfaces (25%, 50%, 75% EMG) and single-modal interfaces (EMG-only, Manual-only) with the following metrics:
(1) Overall task performance. We quantified the overall task performance using the time-domain mean-square error (MSE) between reference r and cursor output y:

$$||r - y||^2 = \sum_{t \in [0,T]} |r(t) - y(t)|^2 \tag{3.10}$$

where $T = 40$ seconds. A lower MSE corresponds to better performance.

(2) Feedforward and feedback transformations at stimulated frequencies. We assumed that a perfect feedforward operator inverts the machine dynamics M^{-1}, especially at frequencies below crossover (Yamagami et al., 2020; Drop et al., 2013). We computed the MSE between M^{-1} and the estimated feedforward transformation F over the stimulated frequencies:

$$||F - M^{-1}||^2 = \sum_{\omega \in [0.1, 0.15, 0.25]} |F(\omega) - M^{-1}(\omega)|^2. \tag{3.11}$$

Although there is no theoretical ideal of a feedback transformation, we could still investigate the steady state error of the feedback controller at the stimulated frequencies:

$$\widehat{y}(\omega) = \widehat{r}(\omega) + \frac{\widehat{d}(\omega)}{1 + M(\omega)B(\omega)}. \tag{3.12}$$

(3) Stability margin of the closed-loop system. We computed the stability margins using Nyquist plots to assess robustness, where the locus of the open-loop transfer function $L(\omega) = M(\omega)B(\omega)$ did not cross -1. The stability margin was calculated as the smallest distance between $L(\omega)$ and -1 in the complex plane.

We hypothesized that (**H1**) **the performance of multimodal interfaces (EMG + Manual) is not different from both single-modal interfaces (EMG or Manual)**. To test this hypothesis, we compared each multimodal condition (%EMG = 25, 50, and 75) with each single-modal condition (%EMG = 0 and 100).

In addition to interface performance, we also evaluated the user inputs of each modality. Prior research has shown that EMG interfaces benefit high-frequency responses and have shorter delays than manual interfaces in human-machine interaction (Peterson et al., 2022). Thus, we hypothesized that (**H2**) **EMG input dominates at high frequencies and leads in phase compared to manual input in multimodal interfaces.**

To further understand coordination between redundant modalities, we introduced a novel frequency-domain interpretation to the established uncontrolled manifold (UCM) hypothesis (Latash, 2012). Given redundant DOF, the UCM method enabled us to project the variability of 2D signals (e.g., inputs from the two modalities) onto 1D task-relevant and task-irrelevant directions (Fig. 3.4B). The variance along the task-irrelevant directions, or the uncontrolled manifold (UCM), does not affect performance but allows abundant solutions; while the variance orthogonal to that space should be suppressed (Latash, 2012; Todorov and Jordan, 2002). Prior research found that if the variance in the UCM direction is greater than the variance in the direction orthogonal to it (ORT), the overall responses are stabilized by the coordination of redundant inputs. We investigated the variances in uncontrolled manifold (UCM, task-irrelevant) and orthogonal (ORT, task-relevant) directions for inputs signals at the non-stimulated frequencies (Fig. 3.4B). Non-stimulated frequencies were chosen for this analysis because the two inputs had the same task goal at every non-stimulated frequency, which is $\alpha u_{EMG} + (1 + \alpha) u_{Manual} = 0$. We calculated the inputs at the non-stimulated frequency in the frequency domain as complex numbers, and then used inverse fast Fourier Transform (IFFT) to compute the inputs in the time domain as real values. We hypothesize that users coordinate multiple inputs aligned with the uncontrolled manifold, and thus the (**H3**) **variance is larger in uncontrolled manifold than the orthogonal space.**

We then examined the multimodal pathways of users by calculating the feedforward and feedback controllers (eq. 3.4 and 3.5) and sensorimotor noises (eq. 3.7) from empirical data. We hypothesized that users combined multiple pathways linearly (Roth et al., 2016); however, due to the interaction between pathways in the closed-loop, the multimodal pathways may not be estimated via the linear combination of pathways in single modalities. We hypothesized that (**H4**) **the feedforward and feedback transformations in EMG and manual pathways in multimodal interaction are not different from the feedforward and feedback transformations in single modalities.** We also hypothesize that (**H5**) **the sensorimotor noises are no different between the two pathways.** The sensorimotor noise was calculated by first linear interpolating the feedback transformations B at the non-stimulated frequencies, then followed eq. 3.6 and 3.7 to find the imputed disturbances $\hat{\delta}$ (section 2.3).

3.7 Statistical Tests

To verify the five hypotheses above, we considered the average of the last four trials of each condition in the 2^{nd} and the 3^{rd} sessions to ensure adequate user learning time.

To test hypothesis **H1**, we first ran a Shapiro-Wilk Test on the residuals to test if the data was normally distributed. We found that our data was not normally distributed so we used a non-parametric Friedman Test to detect an overall difference between the 5 conditions. If the result showed a significant difference between the condition, we would use Wilcoxon signed-rank tests (non-parametric paired t-test) with a confidence level of 5% with Holm-Sidak correction on each multimodal condition to each single-modal condition (i.e., six pairwise comparisons). For hypotheses **H2-H5**, we applied Wilcoxon signed-rank test with a confidence level of 5%. We used `Scipy 1.11.1` for Shapiro-Wilks tests, Friedman tests, Wilcoxon signed-rank tests, and Holm-Sidak correction.

4 Results

4.1 Multimodal interface had better performance than the manual interface

We tested our hypothesis **H1** that the performance of multimodal interfaces ($EMG + Manual$) was not different from single-modal interfaces (EMG or $Manual$). We quantified overall task performance using the time-domain MSE between reference r and cursor output y, $||r - y||^2$. Results showed no significant difference between EMG and Manual-only interfaces ($p > 0.05$), and the equal-weighted condition ($\alpha = 0.5$) achieved a significantly lower error than the Manual-only condition ($\alpha = 0$) (Fig. 3.2A; $p < 0.01$). While the multimodal interface did not significantly outperform the EMG-only condition ($\alpha = 1$), it exhibited significantly lower output at non-stimulated frequencies than both single-modal

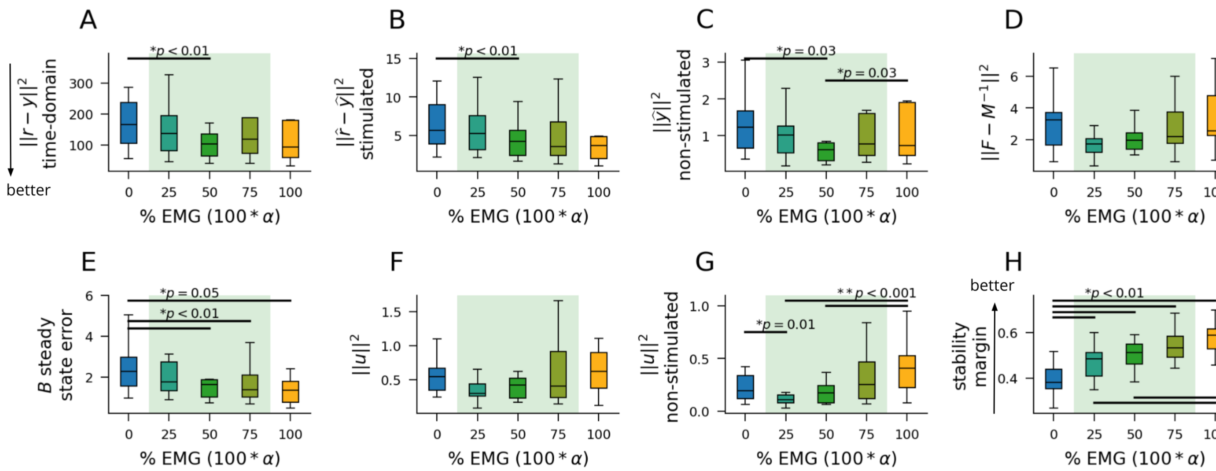

Figure 3.2: Boxplots (N = 15) of the performance and stability of the five experimental conditions. (A) Tracking error between reference r and cursor y at the stimulated frequencies. (B) Error y at the non-stimulated frequencies below 1Hz. (C) Magnitudes of the difference between the feedforward controller F and the inverse of the machine M^{-1} at the stimulated frequencies. (D) Magnitudes of the steady state error with the feedback controller B at the stimulated frequencies. (E) Stability margins, calculated as the minimum distance between open loop transfer function L and -1 on the Nyquist plot. The multimodal conditions ($\alpha = 0.25, 0.5, 0.75$) are labeled with a green background. Horizontal lines with $*$ indicate significant differences between pairs of distributions (Wilcoxon signed-rank tests for pair-wise comparisons with Holm-Sidak corrections, $p < 0.05$).

baselines (Fig. 3.2B; $p < 0.05$). The magnitude of output at the non-stimulated frequencies $|y|$ was approximately 16% of the magnitude of $|r - y|$ at all frequencies, averaged across subjects. This suggests that although stimulated frequencies dominate the response, the multimodal integration specifically suppresses the unwanted noise responses at non-stimulated frequencies, and those noise responses were not trivial. These findings lead us to reject **H1**.

4.2 EMG inputs dominated at high frequencies and manual inputs at low frequencies

To understand how users partitioned control across modalities, we examined the magnitudes of EMG inputs u_{EMG} and Manual inputs u_{Manual} in the equal-weighted condition. We tested our hypothesis **H2** that u_{EMG} dominates at higher frequencies and leads in phase compared to u_{Manual}. We found that in the equal-weighted condition (% EMG = 50), u_{Manual} magnitudes were significantly larger at frequencies below the crossover frequency[1] (median 0.25 Hz across subjects), while u_{EMG} magnitudes dominated above the crossover frequency

[1]frequency at which the open-loop transfer function magnitude is below 1, $|\widehat{L}| = |\widehat{M}(\omega)\widehat{B}(\omega)| < 1$ (McRuer and Jex, 1967).

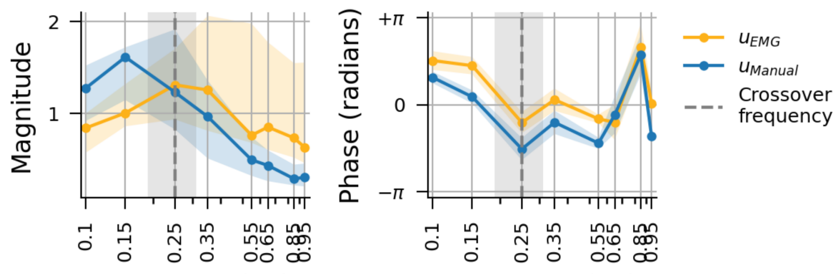

Figure 3.3: Distributions (median, interquartile, N = 15) of the magnitudes and phases of the EMG and manual inputs (u_{EMG} and u_{Manual}) at the stimulated frequencies in the equal-weighted multimodal condition ($\alpha = 0.5$). The gray region indicates the distribution (median, interquartile, N = 15) of crossover frequencies. The median (dashed vertical line) was rounded to the nearest stimulated frequency.

(left plot of Fig. 3.3). Furthermore, u_{EMG} led in phase compared to u_{Manual} (right plot of Fig. 3.3), consistent with prior findings that muscle signals have a shorter time delay and can better predict movements compared to manual signals (Artemiadis, 2012). This supports our **H2** and aligns with prior observations that multimodal interfaces combine distinct features of different signals (Rizzoglio et al., 2020; Zhang et al., 2019).

4.3 Multimodal inputs were anti-correlated below crossover

As we found a significant difference between interface performance at the non-stimulated frequencies, we investigated how users coordinated redundant degrees of freedom (DOF) at those frequencies. We tested our **H3** that the variance of multimodal inputs followed the uncontrolled manifold (UCM) hypothesis. In the frequency domain, we separated inputs u_{EMG} and u_{Manual} at stimulated frequencies Ω and at non-stimulated frequencies $\widetilde{\Omega}$; then we used inverse Fourier Transform to calculate the time-domain signals of the non-stimulated frequencies portion (Fig. 3.4A). Since the task goal at the non-stimulated frequencies is 0 (i.e., noise suppression), the uncontrolled manifold (UCM) spans the task-irrelevant direction $(1 - \alpha, -\alpha)$. If the variance is greater in UCM than the variance of the direction orthogonal to it (ORT), the overall input is stabilized by the coordination of redundant DOF. We found that in all multimodal conditions (25%, 50%, and 75% EMG), the multimodal inputs roughly spanned the UCM directions (Fig. 3.4B). In addition, variances in the uncontrolled manifold (V_{UCM}) were significantly greater than that in the orthogonal direction (V_{ORT}), supporting our **H3** (Fig. 3.4C; $p < 0.01$ for all three multimodal conditions).

We further investigated the magnitudes and phases of inputs at non-stimulated frequencies. We found that users increased the *magnitude* of an input modality if there was a “shortage” in that modality. For instance, in the Manual-dominated condition (25% EMG), users increased

u_{EMG} below crossover to compensate the lower weighting in EMG (yellow line in Fig. 3.5). Interestingly, phase analysis of inputs at non-stimulated frequencies revealed an antagonistic relationship between u_{EMG} and u_{Manual}. The phase difference ($\angle\widehat{u}_{EMG} - \angle\widehat{u}_{Manual}$) remained approximately π for frequencies below crossover (median 0.25 Hz) (grey lines in Fig. 3.5). This non-random anti-phase coordination is consistent across multimodal conditions, which effectively suppresses input magnitudes at these non-stimulated frequencies (green lines in Fig. 3.5). However, we did not observe a similar trend in inputs at the stimulated frequencies. These findings imply that people combined both modalities to work jointly toward tracking/rejecting the stimuli, but used one modality to compensate for the other to suppress their inputs at the non-stimuli, effectively reducing the effect of sensorimotor noise. This finding also aligns with some users' self-reported strategies, which we collected at the end of the experiment via a survey, that they used one modality to "fight" for the other modality to gain a more stable control.

4.4 Multimodal feedforward pathways differ from single-modal feedforward pathways

We investigated whether multimodal pathways could be predicted by simple scaling of single-modal behavior. We tested our hypothesis **H4** to assess whether the feedforward and feedback transformations in single-modal conditions can inform the transformations in multimodal conditions.

We observed that the feedforward transformation of EMG pathway F_{EMG} in multimodal conditions (yellow lines) were different from the F_{EMG} in the EMG-only condition (black lines), and they are significantly different in phase below the 0.25 Hz crossover (Fig. 3.6A, $p < 0.05$ in phase). Similarly, the feedforward of Manual pathway F_{Manual} in multimodal conditions (blue lines) were different from that in the Manual-only condition (black lines) (Fig. 3.6B, $p < 0.05$ in phase). Both F_{EMG} and F_{Manual} in the multimodal conditions (yellow and blue lines respectively) deviated further from the theoretical ideal (dashed lines in Fig. 3.6, where $F = M^{-1}$ in open-loop transformation) than their single-modal counterparts (black lines). These results indicate that the presence of an additional modality reshapes the underlying feedforward control strategy, meaning multimodal behavior is not a linear superposition of isolated modal behaviors.

However, we observed that the feedback transformations of EMG and Manual pathways (B_{EMG} and B_{Manual}, yellow and blue lines respectively) in multimodal conditions were closer to, and were not significantly different from, their single-modal counterparts (black lines) (Fig. 3.7, $p > 0.05$). This suggests that users maintain similar feedback control strategy across single-modal and multimodal conditions.

4.5 Sensorimotor noises were not different between the two pathways

We followed methods in Yamagami et al. (2021b) to compute the imputed disturbances to estimate the sensorimotor noise in each pathway (eq. 3.7). We tested hypothesis **H5** that

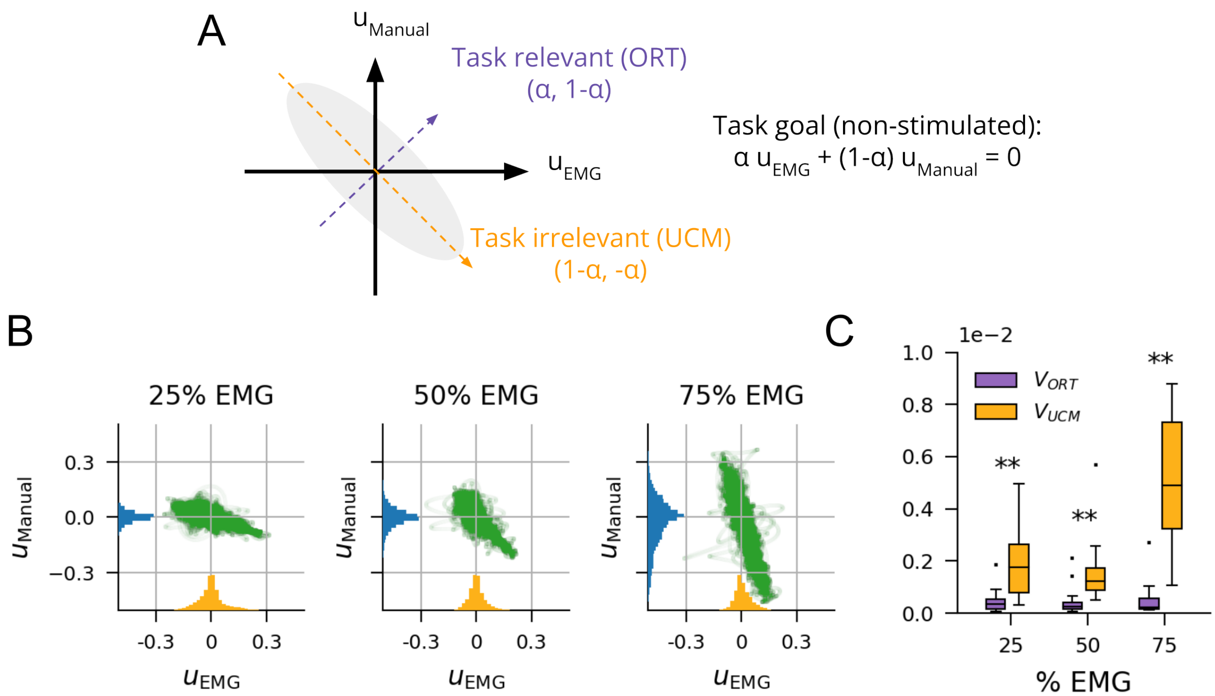

Figure 3.4: (A) Boxplots (N = 15) of magnitudes of (*left*) the combined user inputs in every condition, and (*right*) the combined user inputs at the non-stimulated frequencies below 1 Hz. Horizontal lines with $*$ indicate significant differences between pairs of distributions (Wilcoxon signed-rank tests, $p < 0.05$). (B) Distributions (median, interquartile, N = 15) of the magnitudes and phases of the EMG and manual inputs (u_{EMG} and u_{Manual}) at the stimulated frequencies in the equal-weighted multimodal condition ($\alpha = 0.5$). (C) Example of task-relevant (ORT) and task-irrelevant (UCM) directions of a multimodal interface. The UCM hypothesis predicts greater variance in the UCM than the ORT direction. (D) Scatter and histograms of non-stimulus inputs below crossover for the multimodal conditions (see methods). The x-axis and y-axis are the EMG modality and the Manual modality, respectively. (E) Boxplots (N = 15) of the variance in UCM (V_{UCM}) and the variance in ORT (V_{ORT}) of inputs at the non-stimulated frequencies for the multimodal conditions. Significant differences between V_{UCM} and V_{ORT} are indicated with $*p < 0.05$ and $**p < 0.01$ (Wilcoxon signed-rank tests).

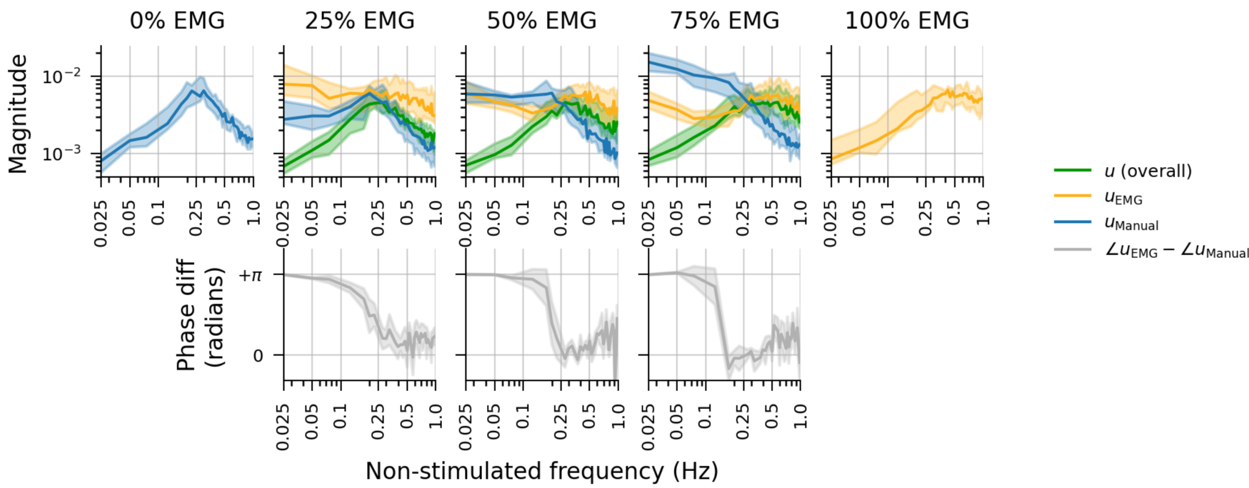

Figure 3.5: Distributions (median, interquartile, N = 15) of user inputs at non-stimulated frequencies in multimodal conditions. (*Top*) magnitudes of combined input (green: $\alpha u_{\text{EMG}}+(1+\alpha)u_{\text{Manual}}$) and inputs from individual two modalities (yellow: u_{EMG}; blue: u_{Manual}). (*Bottom*) phase differences between u_{EMG} and u_{Manual} at the non-stimulated frequencies.

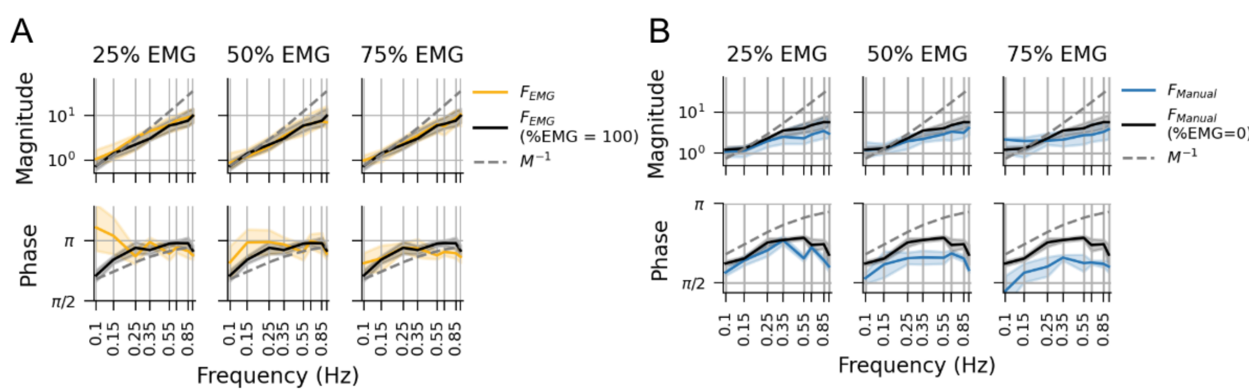

Figure 3.6: (A) Distributions (median, interquartile, N = 15) of the EMG feedforward controller F_{EMG} in the multimodal conditions. The black line indicates F_{EMG} in the EMG-only single-modal condition. (B) Distributions (median, interquartile, N = 15) of the Manual feedforward controller F_{Manual} in the multimodal conditions. The black line indicates F_{Manual} in the Manual-only single-modal condition. The dashed line is the inverse of the machine.

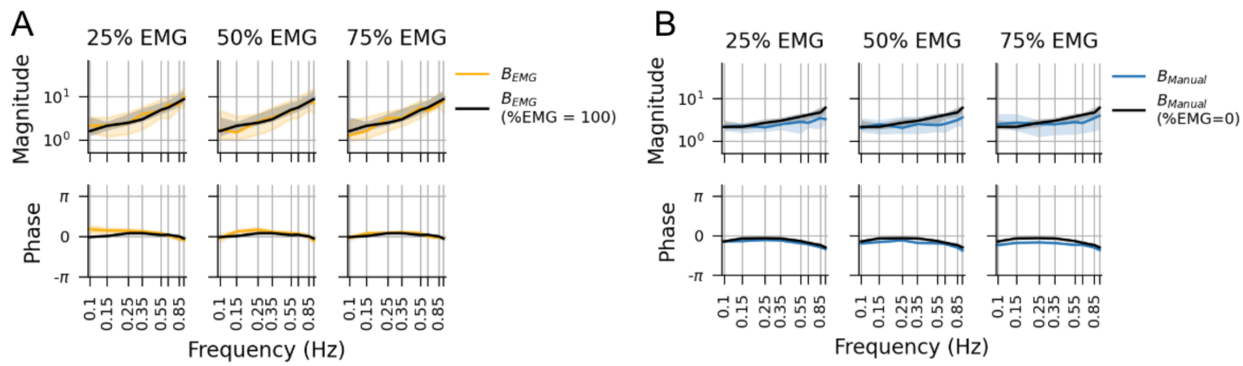

Figure 3.7: (A) Distributions (median, interquartile, N = 15) of the EMG feedback controller B_{EMG} in the multimodal conditions. The black line indicates B_{EMG} in the EMG-only single-modal condition. (B) Distributions (median, interquartile, N = 15) of the Manual feedback controller B_{Manual} in the multimodal conditions. The black line indicates B_{Manual} in the Manual-only single-modal condition. The dashed line is the inverse of the machine.

the sensorimotor noises were no different between the two pathways in the equal-weighted condition. We found that the resulting imputed disturbances δ_{EMG} and δ_{Manual} were not significantly different in the power spectrum ($p > 0.05$). However, the sensorimotor noises were affected by the weightings. The δ of non-dominated modality was significantly higher than the dominated modality ($p < 0.01$ in 25% and 75% EMG in Fig. 3.8). This is unsurprising as our subjects reported in 25% and 75% EMG conditions that they felt that the non-dominated modality was not sensitive enough, so that they would "try harder" in using that modality. This is consistent with prior findings that sensorimotor noise is signal-dependant, which means noise could be affected by the magnitude of user inputs (Harris and Wolpert, 1998).

4.6 Users adopted diverse strategies in combining modalities

Subjects exhibited three distinct strategies based on their feedforward/feedback magnitudes: *EMG-dominant* ($|F_{EMG}| - |F_{Manual}| > 0$), *Manual-dominant* ($|F_{EMG}| - |F_{Manual}| < 0$), and *Mixed strategy* ($|F_{EMG}|$ was larger in high frequencies and $|F_{Manual}|$ was larger in low frequencies); users' strategies in feedback fall in the same strategy groups as that in feedforward (Fig. 3.9A). *EMG-dominant* subjects self-reported in a post-experiment survey that they used EMG to control large movements and Manual for fine adjustments. On the contrary, *Manual-dominant* subjects reported that they used small amount of EMG to compensate for larger Manual overshoot, or had more steady Manual movements compared to EMG.

Additionally, we looked at the averaged task errors of every individual in each strategy group, and found that all three strategies were capable of achieving high performance (Fig. 3.9B),

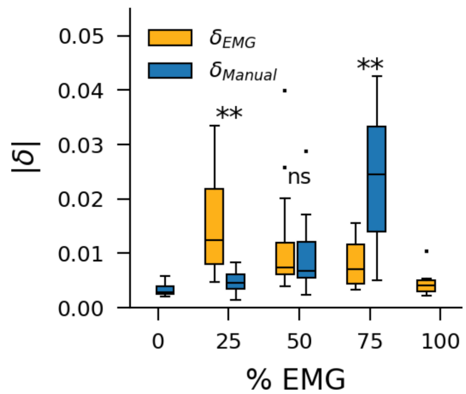

Figure 3.8: Boxplots of magnitudes of sensorimotor noises $|\delta|$ of each pathways. Significance differences are indicated with $**p < 0.01$ (Wilcoxon signed-rank tests) and no significant difference is indicated with *ns*.

highlighting the flexibility afforded by motor redundancy. We also found that people in the *EMG-dominant* group had better performance using EMG-only interface than the Manual-only interface; while people in the *Manual-dominant* and *Mixed* groups have varied performance in EMG-only and Manual-only interfaces. These results show that individuals used their dominant modality in multimodal interfaces, based on performance and personal preferences.

4.7 Subjective user experience

At the end of the experiment, participants provided subjective inputs, including the difficulty of the system (on a scale of 1−5, with 5 being the most difficult), the accuracy of the system (1−5, with 5 being the most accurate), and how much effort did it take to use the system (1−5, with 5 being a lot of effort). The median values of the three indices showed that participants rated the multimodal system to have relatively medium (3) difficulty, high (4) accuracy, and low (2) effort. Regarding subjective preferences, most participants favored using both channels, with seven subjects selecting the multimodal interface, six preferring EMG-only, and two choosing manual-only control.

5 Discussions

Previous research in human-device interaction has shown the advantages of redundancy in both motor pathways (De Santis and Mussa-Ivaldi, 2020) and modalities. Specifically, distinct modalities can provide complementary, overlapping signals for the same movement (Rizzoglio et al., 2020) or offer alternative control modes for robotic systems (Zhang et al., 2019).

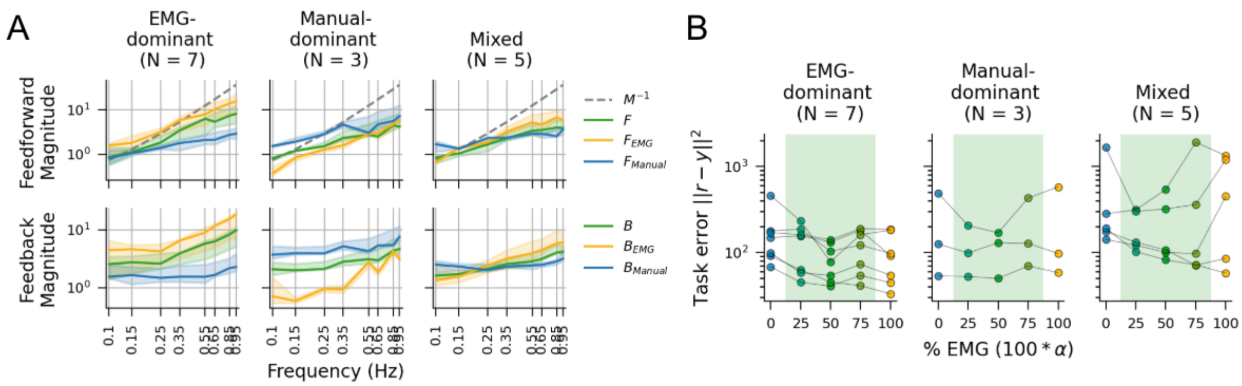

Figure 3.9: (A) Distributions (median, interquartile, N = 15) of the magnitudes of (*top row*) feedforward controllers, F_{EMG}, F_{Manual}, and F, and (*bottom row*) feedback controllers, B_{EMG}, B_{Manual}, and B. Each column is a subject group that had similar strategies: *EMG-dominant*, *Manual-dominant*, and *Mixed*. (B) The task error of each subject in each condition. Each dot is an average across trials. Each subject's data is connected by a line.

However, few studies have investigated how users integrate distinct modalities across different motor pathways to control a single device.

This integration is particularly critical for neural interfaces used in assistive and rehabilitative contexts. Consider a myoelectric-controlled single-arm exoskeleton as an example (McDonald et al., 2020). In this scenario, a user might primarily employ EMG on their affected arm to perform rehabilitation tasks. However, when EMG control degrades due to signal nonstationarity or user fatigue, a redundant manual channel on the less affected limb could maintain system-level accuracy, ensuring continuous device functionality and reducing user frustration. Providing distinct modalities across separate motor pathways gives users two independent control mechanisms ("knobs") that may overcome the inherent limitations of single-mode interfaces. Furthermore, the synergy between these two pathways may optimize the advantages between modalities and thus enhance overall control performance and robustness.

In this study, we picked EMG and Manual modalities in our experiment due to their distinctive trade-off between control bandwidth and physical effort (Lobo-Prat et al., 2014; Yamagami et al., 2020). In our experiment with 15 participants, we evaluated whether multimodal inputs from distinct pathways, or *multimodal pathways*, of combined EMG and Manual signals improve device control performance and robustness. Additionally, we investigated how users coordinate these separate channels, specifically examining the synergy and control strategies between neural (EMG) and manual (joystick) modalities across two motor pathways during the control of complex machines.

5.1 Synergy between multimodal pathways reduced the effect of sensorimotor noise

Our experimental results demonstrated that multimodal inputs generated by distinct motor pathways allowed users to control the virtual machine with significantly better tracking performance than manual (joystick) control alone. Furthermore, the EMG modality demonstrated better performance and robustness compared to manual control (Fig. 3.2A,E). This latter finding is consistent with prior studies showing that EMG signals surpass kinematic signals in continuous tasks (Lobo-Prat et al., 2014; Yamagami et al., 2020). While previous multimodal literature reported benefits of hybrid signals relative to muscle but not motion (IMU) modalities alone, our results specifically highlight performance improvements over manual (joystick) modalities in continuous tracking tasks (Rizzoglio et al., 2020). Additionally, we observed a reduction in the effect of sensorimotor noise in the multimodal condition compared to both the manual-only and EMG-only control conditions (Fig. 3.2B).

To investigate the mechanisms underlying this multimodal advantage, we examined the two input signals (u_{EMG} and u_{Manual}) at both the stimulated and non-stimulated frequencies. Although overall performance is determined by a combination of both responses, isolating them allowed us to identify whether the improvements stemmed from the user's tracking and disturbance-rejection abilities (at stimulated frequencies) or their ability to suppress internal sensorimotor noise (at non-stimulated frequencies). At stimulated frequencies, we observed a higher-frequency phase lead in the EMG input alongside a lower-frequency dominance in the manual input (Fig. 3.3), which means the multimodal interface provided two distinct control signal "knobs" from each modality. This aligns with prior research demonstrating that neural signals from muscle have a shorter time delay and offer better movement prediction than manual signals (Artemiadis, 2012).

Regarding inputs at the non-stimulated frequencies, we found that participants coordinated the two channels to suppress task-relevant noise while allowing task-irrelevant variability (Fig. 3.4B,C), consistent with the Uncontrolled Manifold (UCM) hypothesis (Scholz and Schöner, 1999; Latash, 2012). Despite the disparate nature of these modalities, that one being a peripheral neural signal and the other is kinematic signal, the motor system treats them as a unified, redundant system. The modalities exhibited a synergy analogous to how the central nervous system internally coordinates motor pathways: if one effector deviates from its intended trajectory, the other compensates to maintain stable system-level performance. Interestingly, we observed a larger control effort, defined by the input magnitudes of each modality, in multimodal conditions compared to the inputs in the EMG-only and manual-only conditions (yellow and blue lines, top row of Fig. 3.5). This finding has implications for rehabilitation settings; when muscle engagement needs to be encouraged, a multimodal control paradigm may promote larger EMG magnitudes.

Phase analysis of the inputs at non-stimulated frequencies further revealed an antagonistic relationship between u_{EMG} and u_{Manual}. Specifically, the phase difference ($\angle\widehat{u}_{\mathrm{EMG}} - \angle\widehat{u}_{\mathrm{Manual}}$) maintained a non-random, anti-phase value of approximately π for frequencies below the 0.25

Hz crossover frequency (bottom row of Fig. 3.5). This implies that anti-phase coordination effectively cancels out and suppresses noise magnitude at these non-stimulated frequencies. These results suggest that users actively prioritize the reduction of noisy inputs over minimizing physical effort, achieving this by utilizing opposing forces between redundant degrees of freedom (DOF). We propose that this anti-phase coordination of redundant modalities is functionally similar to the mutual opposing forces observed in animal locomotion, where animals generate counteracting forces to simultaneously enhance maneuverability and locomotor stability (Sefati et al., 2013).

5.2 Multimodal pathways reshape behaviors

Our results suggest that the feedforward pathways for EMG and manual inputs (F_{EMG}, F_{Manual}) in multimodal conditions differ significantly from those in single-modal conditions (yellow and blue lines vs. black lines in Fig. 3.6). This suggests that multimodal behavior cannot be predicted simply from single-modal performance. Furthermore, these feedforward pathways varied depending on the weighting parameter α assigned to each modality. Specifically, when a modality was weighted less, its feedforward magnitudes were larger and its phase deviated further from the theoretical ideal (M^{-1}, represented by the dashed lines in Fig. 3.6). This finding aligns with subjective reports from several participants, who noted that they used the less sensitive modality to fine-tune cursor movements and complement the more sensitive modality. The fact that control strategies adapt in the presence of a redundant modality suggests that the user's internal model shifts based on the available control degrees of freedom (DOF) and their relative weightings.

Beyond behavioral shifts between conditions, we also observed a diversity of individual strategies across participants. These could be broadly categorized into three distinct control behaviors: some subjects relied on the EMG pathway as the dominant modality, others favored the manual pathway, and some adopted a frequency-dependent mixed strategy (Fig. 3.9A), and these also matched the participants' self-reported experiences. Interestingly, more subjects used EMG as their dominant modality ($N = 7$) compared to those who relied primarily on manual control ($N = 3$). These individuals reported preferring EMG for large, coarse adjustments and the joystick for fine-tuning. Additionally, more people reported a higher subjective preference for the EMG modality than the manual. This preference is intuitive because physical force is directly tied to acceleration, which matches the second-order acceleration dynamics of the virtual machine controlled by the participants. Nevertheless, some participants struggled with muscle-based control; three subjects reported that the EMG interface was "somewhat difficult" to use. This observation evokes the concept of neural interface inefficiency, where prior literature indicates that inefficient users can make up a considerable fraction of an experimental population (Guger et al., 2003; Zhang et al., 2020a). This variation further highlights that incorporating a redundant manual modality allows these individuals to achieve high-performance cursor control while simultaneously utilizing and practicing their muscle pathways.

Furthermore, there are subjects who can still achieve relatively good task performance within each strategy group (lower task error in Fig. 3.9B). This indicates that no single control strategy is inherently flawed; rather, it is possible to yield good task performance regardless of the user's strategic preference, though overall performance variance across subjects remains high. These findings emphasize that individual preferences and distinct motor abilities must be evaluated when designing future human-machine interfaces.

5.3 Future work

In this study, we proposed a framework of evaluating multimodal interfaces and demonstrated the benefits of redundant motor modalities. We would like to further investigate whether multimodal interfaces can be predicted using optimal combination of multiple modalities. Furthermore, we are interested in extending the current interface to an adaptive interface (Chasnov et al., 2019; Madduri et al., 2021), in which we will iteratively update the machine according to users' performance. Through these approaches, we wish to provide a generalizable framework to optimally fuse multiple sensory-and-motor modalities using an adaptive interface.

6 Preliminary experiment with multimodal sensory feedback

In addition to multimodal motor pathways, we are also interested in investigating the benefits of multimodal sensory feedback. Prior research has shown that multimodal sensory feedback can improve performance and synergy of motor tasks (Koh et al., 2016). Here we present the preliminary findings of engineering the sensory-and-motor transformations in a human-computer interface, as an extension to the paper "Redundant Sensorimotor Pathways in a Multimodal Neuromotor Interface" in this Chapter.

6.1 Motivation

While our results suggest that users benefit from multiple motor output pathways, it is unclear if and how users would benefit from multiple sensory input pathways. We further seek to extend methods and findings in this Chapter to incorporate multiple sensory pathways. Haptic feedback naturally complements the visual feedback we already use since it can provide lower-latency task information to the user. We conducted a pilot experiment where we systematically vary the "weight" assigned to the different sensory pathways (haptics and visual) and estimate the sensorimotor transforms that result.

6.2 Methods

Subjects use their dominant hands (left hand in Fig. 3.10) to control a slider joystick for a disturbance rejection task. The subject sees the position of a cursor on a computer screen, and feels a force from a haptic device with their non-dominant hand (right hand in Fig. 3.10).

This haptic feedback is achieved by applying a horizontal opposing force to the user using an actuated mechanism. The larger the cursor is away from the central line, the greater the force is applied. The task goal of the subject is to achieve zero displacement of the cursor to the center line and to obtain zero force with the haptic device.

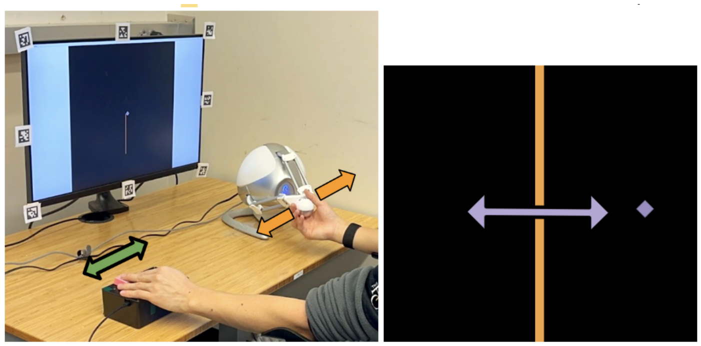

Figure 3.10: Experimental setup. (left) A participant uses a slider with their dominant hand to control a purple cursor to track a yellow center line while feeling horizontal force feedback on their non-dominant hand. (right) The cursor moves horizontally in the task screen; the center line stays stationary.

The experiment includes two distinct sensory feedbacks: visual feedback via a computer screen and the force feedback via a haptic device. Subjects conduct a disturbance rejection task using a slider. Subjects first conduct the task given single feedback (visual-only or haptic-only), and then given multimodal feedback (visual + haptic) with different weightings for the gains ($\omega = 0, 0.25, 0.5, 0.75, 1$). A block diagram of the multi-sensory system is shown in Fig. 3.11. Subjects conduct the task again with single feedback at the end of the experiment. The experiment was approximately 1.5 hours for each subject. The schematic of the trial structure is shown in Fig. 3.12.

6.2.1 Experiment A ($N = 1$ subject): conflicting multisensory feedback

We conducted another exploratory experiment with 1 subject to test the multi-sensory system with conflicting feedback signals. We hypothesized that users rely predominantly on the visual sensory information when given conflicting visual and haptic feedback. The stimulated frequencies of the haptic disturbance are prime multiples (2,5,11,17) of the base frequency of 0.05 Hz, while the stimulated frequencies of the visual disturbance are prime multiples (3,7,13,19) of the base frequency. This design leads to conflicting feedback signals from the two channels. Five different conditions with different weightings $\omega \in [0, 0.25, 0.5, 0.75, 1]$ are tested. Each condition has 14 trials, one minute each.

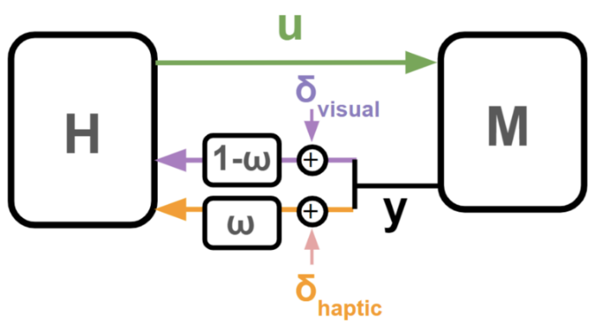

Figure 3.11: Block diagram of the multi-sensory system. Human H uses slider input u to control a first-order (velocity-based) machine M. The output of the machine y is perturbed by disturbances δ and is sent back to the user via distinct channels: visual and haptic with different gains ω. The larger the ω, the larger the haptic input.

6.2.2 Experiment B ($N = 2$ subjects): non-conflicting multisensory feedback

We conducted an exploratory experiment with two sensory modalities that have identical disturbance signals. We hypothesized that augmenting haptic feedback improves task performance compared to having visual feedback alone. The disturbances are sum-of-sines signals at the stimulated frequencies. The stimulated frequencies are prime multiples (2,5,11,17) of the base frequency (0.05 Hz). Five different conditions with different weightings (0, 0.25, 0.5, 0.75, 1) are tested. Each condition has three trials. Each trial is 4 minutes.

6.2.3 Evaluation

In the multi-sensory system, user simultaneously perceives two feedback signals: (1) visual feedback signal: $(1-\omega) * (y + \delta_{visual})$, and (2) haptic feedback signal: $\omega * (y + \delta_{haptic})$.

In time domain, the visual feedback error is the normalized distance between the computer cursor and the stationary center line.

$$\|y + \delta_{visual}\|^2 = \sum_{t \in [0,T]} |y(t) + \delta_{visual}(t)|^2. \tag{3.13}$$

The haptic feedback error is the normalized force that the end effector applied to the user. These errors are normalized by the gains ω.

$$\|y + \delta_{haptic}\|^2 = \sum_{t \in [0,T]} |y(t) + \delta_{haptic}(t)|^2, \tag{3.14}$$

where T = 1 minute (experiment A) or 4 minutes (experiment B). Note that δ_{visual} and δ_{haptic} are identical in experiment B.

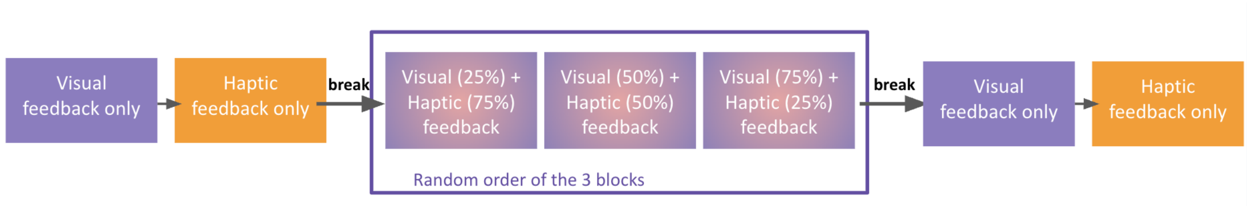

Figure 3.12: Schematic of trial structure.

6.3 Preliminary findings and discussions

6.3.1 Experiment A: sensory conflict yields tradeoff in modality performance.

In experiment A, different error signals are provided along two sensory pathways, and the weighting for each pathway varies between conditions: $\omega = 0$ corresponds to visual-only feedback and $\omega = 1$ corresponds to haptic-only feedback.

We found that task performance was better in single-modality conditions than when multiple modalities were active across all three multi-modal weighting conditions (Fig. 3.13). Taken together, these results indicate that the users exert more effort (quantified by the magnitude of the user input signal) to negotiate the tradeoff between conflicting task goals (defined by the distinct error signals displayed along the sensory pathways).

6.3.2 Experiment B: in the absence of sensory conflict, introducing haptic feedback affected behavior but not performance.

In experiment B, we removed the sensory conflict, instead displaying the same error signal with different weightings along the two sensory pathways.

We found that task performance in the haptic-only condition was significantly worse than all other conditions, which all had similar performance (Fig. 3.14a). At the same time, input signal magnitude in the visual-only condition was significantly smaller than all other conditions, which had similar magnitudes (Fig. 3.14b). This indicates that the presence of haptic feedback in the multi-modal conditions changed user behavior – in particular, by yielding more hand movement – without significantly changing performance. This finding may be valuable in future work that seeks to elicit user exploration while learning a novel interface.

7 Conclusion

In this study, we propose a new experimental method to evaluate multimodal human-device interaction across multimodal motor (manual + EMG) or sensory (visual + haptic) path-

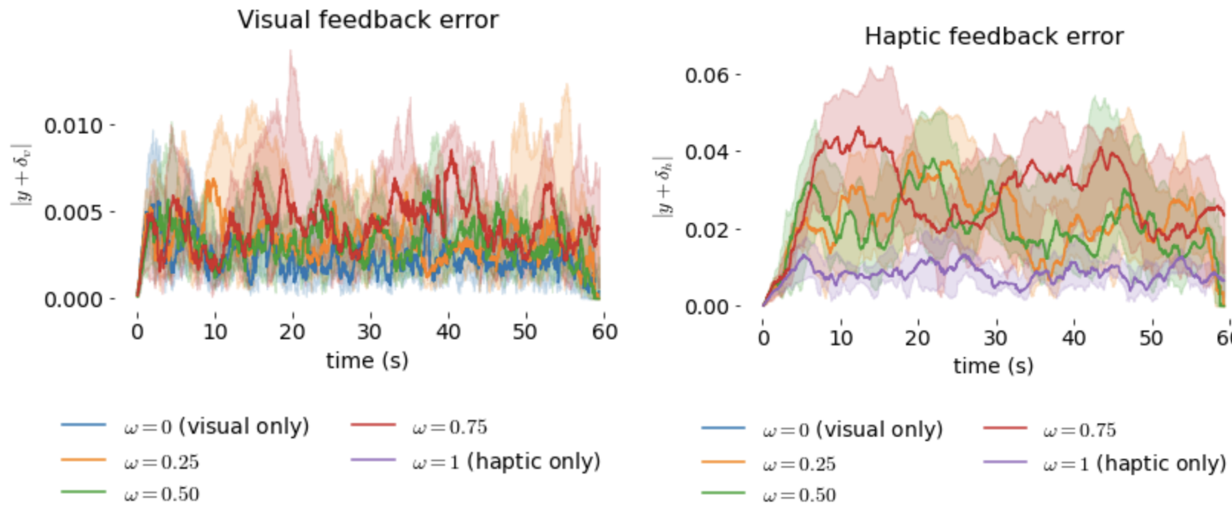

Figure 3.13: Time-domain task error for each condition, with sensory conflict. Distribution (median, interquartile) of normalized (left) visual feedback errors and (right) haptic feedback error , N = 3 (1 subjects × 3 trials for each condition). The weight ω is the gain of haptic feedback; $(1-\omega)$ is the gain of visual feedback.

ways. Through human-subject experiments involving a continuous trajectory-tracking and disturbance-rejection task, we demonstrate that multimodal motor pathways improve performance and reduce the effects of sensorimotor noise compared to single-modal interfaces. We also find that users adapt their control strategies to effectively coordinate the two modalities, and that synergy between the two motor pathways allows users to suppress their sensorimotor noise. This suggests that future human-machine systems requiring minimal sensorimotor noise may benefit from redundant motor modalities, especially combined manual and muscle inputs. Furthermore, we observed diverse strategies across users, which suggests that future systems should consider individual preferences and motor abilities when designing multimodal interfaces. Additionally, our preliminary findings in multimodal sensory feedback show that user behaviors change when multiple sensory pathways are available. These findings provide a foundation for future work in designing multimodal human-machine systems that incorporate both motor and sensory pathways to improve performance and robustness. Because human behaviors are shaped by these additional modalities, future studies are needed to investigate how to optimally combine multiple sensory and motor pathways in human-machine systems.

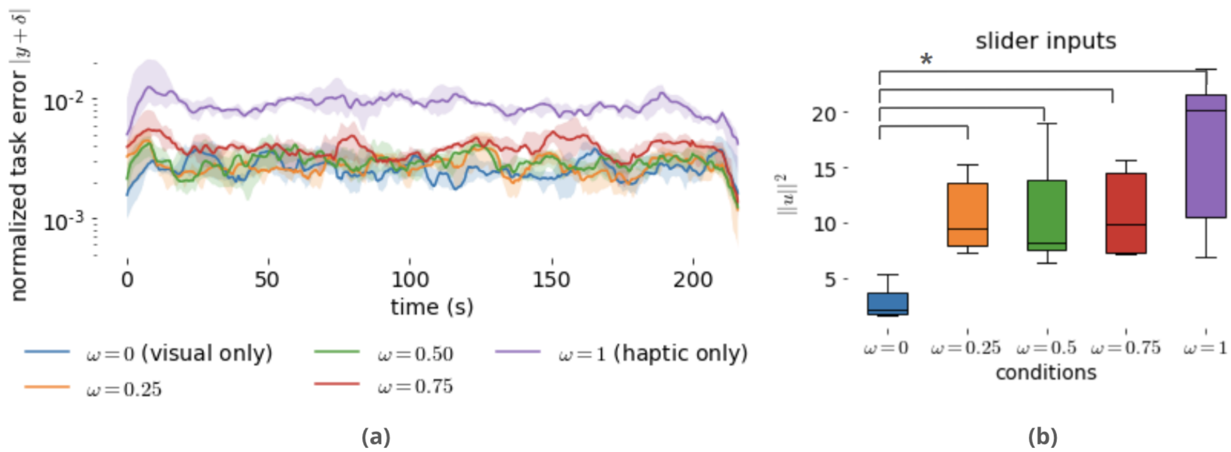

Figure 3.14: (a) **Time-domain task error for each condition, without sensory conflict.** Distribution (median, interquartile) of normalized task errors in the visual and haptic pathways, N = 6 (2 subjects × 3 trials for each condition). Moving average with a 10-second window. The weight ω is the gain of haptic feedback; $(1 - \omega)$ is the gain of visual feedback. (b) **Time-domain slider input for each condition.** The user input is the norm squared of the slider input. Error bars were from N = 6 (2 subjects × 3 trials). Significant differences are shown with a $*$ (Wilcoxon signed-rank test, $p < 0.05$).

Chapter 4

CO-ADAPTIVE NEUROMOTOR INTERFACES WITH MULTIMODAL INPUTS

Using Gaze to Train a Closed-loop Adaptive Decoder for Neuromotor Interfaces

Amber H.Y. Chou, Maneeshika M. Madduri, Si Jia Li, Jason Isa, Andrew Christensen, Finley Hutchison, Samuel A. Burden, Amy L. Orsborn

Under review. bioRxiv. `https://doi.org/10.1101/2024.04.08.588608`

Abstract

Neuromotor interfaces enable control of devices like robots and computers, thereby providing accessible interactions and restoring motor function. Current interfaces still face challenges in decoding high-dimensional neural recordings to maintain consistent control across users and tasks. Algorithms that adapt decoders based on closed-loop performance can improve accuracy and provide avenues to continuously account for changes in their users. Current approaches for closed-loop decoder adaptation commonly use task-specific labels (e.g., prescribed paths or target positions) for supervised training, which restricts the algorithms to be used only in instructed calibrations. There is a need for a new approach to extend decoder adaptation beyond instructed settings to support real-world, self-paced applications without explicit targets. We propose a new paradigm that approximates task goals from natural gaze, and uses gaze to replace task-specific labels when training the decoder in real time. This method offers a less supervised alternative for closed-loop decoder adaptation while still supporting continuous, rapid calibration achieved via supervised training. We evaluated our gaze-trained method using a myoelectric interface in a controlled task, where 11 participants used high-density surface electromyographic (sEMG) signals measured from forearm muscles to control a 2D continuous cursor. Experimental results demonstrate the feasibility of our gaze-trained method, with the decoder and users co-adapting to achieve performance nearly comparable to the current task-dependent training method.

1 Introduction

Neuromotor interfaces that measure neural activity, muscle activity, movement, and other physiological recordings have become increasingly popular in *human-device interaction* for healthcare applications, including myoelectric prostheses (Farina et al., 2023), brain-computer interfaces (BCI) (Willett et al., 2023), gaze-based assistive devices (Eid et al., 2016), and neurorehabilitation (Mugler et al., 2019; McDonald et al., 2020). These interfaces translate rich, high-dimensional neural recordings into low-dimensional control commands for external devices, such as computer cursors or bionic limbs. Neuromotor interfaces provide always-available inputs (Saponas et al., 2008) and a higher control bandwidth than conventional manual inputs such as mice, keyboards, and touch-based screens (Yamagami et al., 2020; Lobo-Prat et al., 2014; Wolpaw, 2007), making them a promising approach to let individuals with motor impairments operate different devices.

Transforming high-dimensional neural inputs into effective commands is key to providing intuitive device control. For example, a non-invasive *myoelectric interface* records and processes surface electromyographic (sEMG) activities from multiple muscles, and uses a transformation, known as a *decoder*, to convert multi-channel recordings into control signals such as discrete gestures or continuous cursor kinematics. This decoder is commonly programmed without the user interacting with the device or decoder (user not in the loop; *open-loop*), through methods such as dimensionality reduction, classification, or neural networks (Mugler et al., 2019; Jiang et al., 2014; Saponas et al., 2008; Zia Ur Rehman et al., 2018; Kaifosh et al., 2025). Open-loop decoding often results in a mismatch between open- and closed-loop interface performance (Hahne et al., 2017; Woodward and Hargrove, 2019; Chase et al., 2009; Koyama et al., 2010; Eddy et al., 2023). Differences in open and closed loop performance can occur due to signal quality changes over time due to variability in electrode displacement or sensor properties (Yamagami et al., 2018b; Young et al., 2011), as well as user changes. Users are not stationary data sources – they adapt to the interface and task, and may feel fatigued over time. Decoding algorithms need to account for the dynamic nature of human users in the loop to ensure closed-loop performance (Hahne et al., 2017; De Santis, 2021; Slade et al., 2024).

Closed-loop decoder adaptation (CLDA) – continuously updating the model while the user is operating the interface and receiving feedback – holds promise to close the train/test performance gap (Gajos et al., 2012; Orsborn et al., 2012, 2014; Yamagami et al., 2023a; Madduri et al., 2026). Current adaptation approaches include supervised methods that require labeled training data (Taylor et al., 2002; Orsborn et al., 2012; Gilja et al., 2012; Dangi et al., 2013; Shanechi et al., 2016), unsupervised methods to identify latent patterns without labels (e.g., regression, feature matching, etc.) (Li et al., 2011; Gürel and Mehring, 2012; Jarosiewicz et al., 2015; Sussillo et al., 2016; Zhang et al., 2022; Liu et al., 2024; Li et al., 2025), and reinforcement learning that updates decoder parameters based on rewards (Mahmoudi and Sanchez, 2011; Pilarski et al., 2011). Supervised methods have been shown to improve interface performance and provide rapid calibration – both in generating an initial decoder and retraining the decoder (Taylor et al., 2002; Orsborn et al., 2012; Gilja et al., 2012; Jarosiewicz et al., 2013). Additionally, they are useful in personalizing the interface to diverse users and their usage, without the need for extensive offline data collection (Madduri et al., 2023). They are promising in enhancing long-term usability, but present the challenge: how to obtain labeled training data at the same time that the interface is being used? Current algorithms in this category address this challenge by making use of information about the tasks prescribed to users, such as the target positions and velocities (Orsborn et al., 2012; Gilja et al., 2012; Madduri et al., 2026; Hahne et al., 2015). However, requiring this information limits the generalizability of interfaces to tasks without explicit goals, especially for less structured, self-paced tasks like writing or exploring a webpage. Unsupervised or reinforcement learning methods, on the contrary, do not require task goals to be explicitly defined. Prior studies have shown that these methods are robust to neural signal nonstationarities (Li et al., 2011; Sussillo et al., 2016). Existing unsupervised algorithms primarily focused on maintaining the performance of already high-performing decoders and have not been systematically evaluated for generating an initial decoder, which imposes different computational goals and algorithmic needs (Dangi et al., 2013). Moreover, methods that implement neural networks often require a large amount of training data, and are difficult to retrain online, thus constraining the ability to adapt to a new operation or to switch between operations, such as switching between computer tasks or changes in user strategies. Our vision is to leverage the advantages of supervised methods to provide rapid initial calibration for each user and retrain the interface on the fly, while also leveraging unsupervised methods to allow seamless calibration in diverse operations, including tasks without explicit goals, for continuous interactions in neuromotor interfaces. Toward this long-term vision, we investigated the use of *gaze* as a new task-agnostic source of training data labels to adapt decoders in closed-loop (Fig. 4.1A).

Eye-tracking technology has found widespread applications in virtual reality headsets to enrich user understanding and augment interface functionalities (Jacob and Karn, 2003; Zhao et al., 2021). Conventional gaze-controlled interfaces generally require users to focus their (*intentional*) gaze to produce input, which limits eye movements during daily tasks and can cause fatigue (Jacob and Karn, 2003). *Natural* eye behavior, in contrast, captures user-specific information, such as user's movement goals (Matthis et al., 2018; Johansson et al., 2001; Aronson et al., 2018), fatigue levels (Horng et al., 2004), and other physiological states (e.g., emotional state, attention) (Skaramagkas et al., 2023) that neural recordings like sEMG alone may not provide. Thus, natural gaze holds potential to augment neuromotor interfaces with additional relevant information about users. Several studies in robotic or prosthetic control have integrated people's natural gaze, which revealed advantages in anticipating user selections or intentions to improve system efficiency and accuracy (Huang and Mutlu, 2016; Stolzenwald and Mayol-Cuevas, 2019; David-John et al., 2021; Aronson and Admoni, 2022; Admoni and Srinivasa, 2016; Karrenbach et al., 2022). For instance, Aronson *et al.* (Aronson and Admoni, 2022) used natural gaze to predict users' task goals and then used the predictions to improve a learning algorithm for manual (joystick)-based robot manipulation.

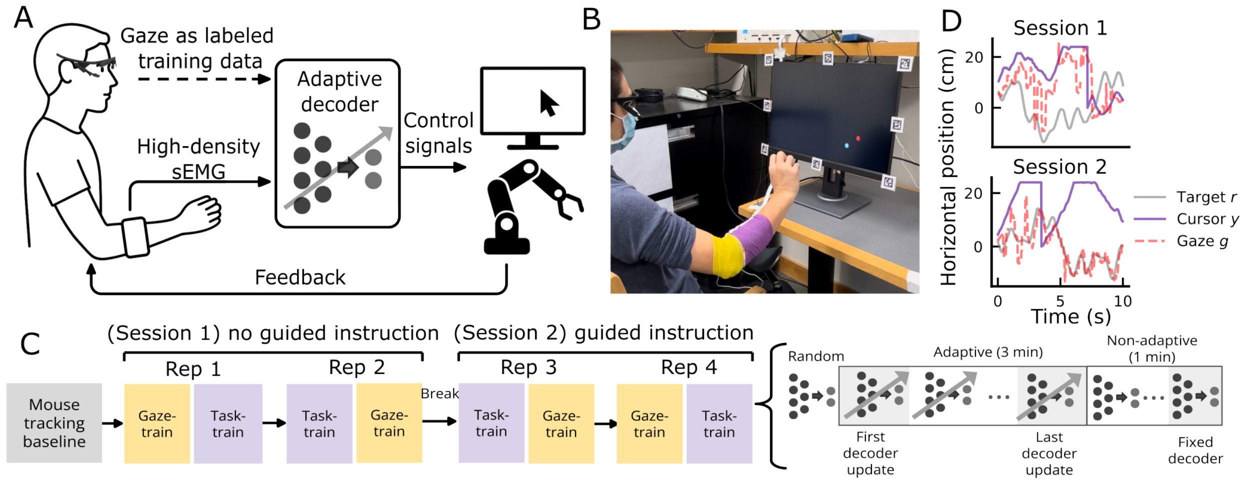

Figure 4.1: (A) Overview of a closed-loop myoelectric interface that records high-density sEMG and gaze from the user, and uses gaze to train an adaptive decoder to control an external device, such as a computer cursor or a robotic arm. The angled arrow indicates decoder adaptation. The user receives visual feedback while operating the interface. (B) The user controls a cursor (white) to track a continuous target (red) on a computer screen using sEMG measured from the forearm. Gaze is measured with an eye-tracking headset. (C) Trial structure of the experiment. After calibration of the eye tracker, participants conducted two baseline trials using a mouse, and then conducted four repetitions of each condition (task- and gaze-trained) in a random order. Participants were asked to fixate on the continuous target (red) in the second Session. Participants had a mandatory five-minute break between the two sessions. During each 4-minute trial, the decoder was initialized randomly, adapted every 20 seconds for three minutes, and then stopped adapting (fixed) for another one minute. (D) Example gaze positions g over time in (*top*) *Session 1* and (*bottom*) *Session 2*.

Karrenbach *et al.* (Karrenbach et al., 2022) also leveraged gaze to predict wrist postures in prosthetic control for grasping.

In this study, we hypothesized that eye movements can replace the prescribed task data in closed-loop decoder adaptation by assuming that natural gaze serves as a proxy for user intent or task goals. We recorded real-time gaze and used gaze as ground truth to train an adaptive decoder. We call this new method the *gaze-trained* paradigm in contrast to the previous supervised methods that depend on a prescribed task (*task-trained* approach). In this study, we used a high-density myoelectric interface to test the feasibility of our method because of its non-invasiveness and ease of setup. We benchmarked our gaze-trained paradigm to the task-trained paradigm to demonstrate its practical use in a controlled task. We recruited 11 naive subjects to control a computer cursor using a 64-channel sEMG electrode on their forearm to follow a 2-dimensional (2D) moving target using any hand movements they preferred (Fig. 4.1B). Previous literature defined this cursor-tracking as a *continuous interaction* (Eddy et al., 2023; Yamagami et al., 2020), which has the potential to be translated to other continuous computer or human-device tasks, such as handwriting (Huang et al., 2010; Kaifosh et al., 2025), prostheses control (Fleming et al., 2021; Zhuang et al., 2019b), or wheelchair navigation (Moon et al., 2003). To our knowledge, this work is the first myoelectric interface that uses gaze to adapt its decoder in closed-loop continuous operation, contrasting with methods in previous myoelectric studies, such as open-loop decoding or discrete operations (e.g., gesture classifications).

We analyzed the interface performance and leveraged tools from control theory to understand the *co-adaptation* between users and decoders. The contributions of our work are:

1. We develop a new training paradigm that uses gaze as labeled training data in closed-loop decoder adaptation, and highlight its task-agnostic potential for diverse tasks.

2. We benchmark the performance of the gaze-based training algorithm with that of task-based algorithm, both when users are provided instructions about gaze and when they are not.

3. We show that users responded differently when the decoder was adapted based on gaze compared to when it was adapted based on task information.

This study will enable future neuromotor interfaces that seamlessly self-calibrate to diverse users and tasks during operation. Our method might address the challenges associated with generalizability in existing adaptive decoders by removing the need for prescribed tasks, gestures, or guided eye movements.

1.1 Relation to Prior Work in HCI

Noninvasive myoelectric sensing has been explored in human-device interaction studies, including user understanding (Mandryk et al., 2006), VR/AR training (Phelan et al., 2021), limb prostheses (Farina et al., 2014; Fleming et al., 2021), and rehabilitation (McDonald et al., 2020; Mugler et al., 2019; Hung et al., 2021; Rizzoglio et al., 2019). From an HCI perspective, these studies demonstrated an accessible alternative to traditional manual devices like mice and keyboards for users with upper-limb motor impairments (Yamagami et al., 2020). Additionally, myoelectric interfaces provide always-available inputs that enable everyday gesture detection to replace touch-based mobile interfaces (Huang et al., 2015b; Saponas et al., 2009; Kaifosh et al., 2025). This is particularly desirable when a user's hands are occupied or when a touch screen is not available.

Particularly, prior HCI studies have shown the promise of using upper body (e.g., forearm) muscle signals as an input modality to a computer device. Saponas et al. (Saponas et al., 2008, 2009) and Huang et al. (Huang et al., 2015b) used sEMG armbands to classify hand gestures such as thumb tapping and finger lifting for mobile or keyboard interactions. Others have also studied sEMG in gesture recognition for smartwatch (Yang et al., 2015) and smart garment (Benatti et al., 2014) interaction. In addition to discrete hand gestures, many studies have demonstrated the use of sEMG for continuous cursor navigation through hand movements or muscle contractions (Yamagami et al., 2020; Rosenberg, 1998; Lobo-Prat et al., 2014; Kim et al., 2004). CTRL-labs at Reality Labs (Kaifosh et al., 2025) has recently reported an sEMG decoding model trained with thousands of participants' data for multiple computer tasks like continuous navigation, discrete gestures, and handwriting. They have found that their deep-learning model can be accurately generalized across populations and sessions.

While the interfaces in these studies worked effectively for their specific purposes, they were trained using offline data, which hindered their adaptability across users and tasks when a dynamic user is present in the loop. While findings with a generic interface were exceptional, it is challenging to replicate their model built with such large subject sizes. This work also uses sEMG as an input modality for computer cursor control; however, the main difference between this work and the prior work is the continuous adaptation of the myoelectric interface while users operate a task. In addition, our adaptive myoelectric algorithm does not require a large dataset collected offline; instead, the mapping is trained online with data from a single participant in a trial. This can potentially avoid issues such as replicability or biosignal data privacy.

Eye-tracking technology has found widespread applications in virtual reality headsets to enrich user understanding and create more immersive experiences. In addition, eye-tracking serves as a readily measurable input modality to control devices without hands, and is faster in cursor control compared to conventional mice (Jacob and Karn, 2003; Yu et al., 2014b). Prior studies have shown that eye inputs can augment interface functionalities (Jacob and Karn, 2003; Zhao et al., 2021) and enhance interface accessibility (Hemmingsson and Borgestig, 2020; Zhang et al., 2017b).

Several studies in HCI have shown the advantages of integrating interfaces with eye-tracking and upper-body movements. For instance, the combination of gaze with gesture inputs enables rapid, touch-free computer interactions (Chatterjee et al., 2015) and can be used to enhance user intention estimation (Schwarz et al., 2014). Other studies developed multimodal interfaces that used gaze for cursor navigation and sEMG (with facial muscles, forehead, forearm, etc.) for discrete operations such as selection, switching, or clicking, demonstrating performance surpassing that of traditional mice (Mateo et al., 2008; Pai et al., 2019). While combining multiple sources of biosignals enhances interactions, the potential benefits of multiple biosignals in adaptive interfaces have not yet been fully explored.

In this work, we extend previous research on myoelectric interfaces by incorporating eye gaze data. What sets our approach apart from other gaze-controlled devices is that we did not directly use intentional gaze as a control input. Rather, we used natural gaze as training data for our adaptive algorithm. Several studies have also integrated people's natural gaze in human-robot interaction and have proved its advantages in anticipating user selections or intentions, thus improving system efficiency and task accuracy (Huang and Mutlu, 2016; Stolzenwald and Mayol-Cuevas, 2019; David-John et al., 2021; Aronson and Admoni, 2022; Admoni and Srinivasa, 2016). Aronson et al. (Aronson and Admoni, 2022) used natural gaze to predict users' task goals and then used the predictions to improve a learning algorithm for manual (joystick)-based robot manipulation. Similarly, this study also leverages gaze to approximate the user's desired actions. The main distinction lies in the interface adaptation using high-dimensional biosignal inputs, contrasting with the fixed mapping of low-dimensional manual devices used in previous studies.

2 Methods

2.1 Gaze facilitated adaptive decoder

We built upon the closed-loop adaptive algorithm explored in prior literature (Orsborn et al., 2012; Madduri et al., 2026) to design a noninvasive myoelectric interface with an adaptive decoder that continuously adapts to its user during operation. The adaptive decoder, matrix $D \in \mathbb{R}^{2\times N}$, maps the high-dimensional sEMG signals $s \in \mathbb{R}^N$ to the velocity of a 2D cursor $v_y \in \mathbb{R}^2$:

$$v_y = D \cdot s, \tag{4.1}$$

where N is the number of sEMG input features (64 in this study). The cursor position y was calculated in discrete time:

$$y_t = y_{t-1} + v_y \Delta t \tag{4.2}$$

where Δt is the time interval.

The decoder D was adapted using a gradient-based learning algorithm designed to minimize the cost with a regularization term. In prior work that used predefined task information for training, it was assumed that the user would intend to move the cursor toward the task target (Orsborn et al., 2012; Gilja et al., 2012). Thus, the intended goal was the target position, and the cost function was defined as the norm difference between the intended velocity – the instantaneous velocity of a vector between the target r and cursor y at a given time – and the actual velocity of the cursor:

$$c_{\text{task}}(D) = ||(r-y)/\Delta t - D \cdot s||_2^2 + \lambda ||D||_F^2. \tag{4.3}$$

The $||D||_F$ denotes the Frobenius norm of the decoder and λ is the weight of the regularization. Both $r \in \mathbb{R}^2$ and $y \in \mathbb{R}^2$.

The innovation of this work is to infer the user's intended goal from their gaze and use gaze to train the decoder. We replaced the target position in the cost function to be the gaze position $g \in \mathbb{R}^2$ in the plane defined by the computer display:

$$c_{\text{gaze}}(D) = ||(g - y)/\Delta t - D \cdot s||_2^2 + \lambda ||D||_F^2. \tag{4.4}$$

Thus, the adaptive decoder was not explicitly given the target information in the gaze-trained decoder.

Following the closed-loop adaptation framework in Madduri *et al.* (Madduri et al., 2022), the decoder D was updated iteratively to minimize the cost (c_{task} or c_{gaze}) using 20-second batches of data. The 20-second batch length was selected to train data with intended velocities spanning various directions on the computer screen. The optimal decoder D^* was computed using the Broyden–Fletcher–Goldfarb–Shanno (BFGS) method:

$$D^* = \arg\min_D c(D), \tag{4.5}$$

and the updated matrix D^+ is the convex combination of the previous D and the optimal D^* for a smooth adaptation:

$$D^+ = \alpha D^* + (1 - \alpha)D, \tag{4.6}$$

where α is the rate of adaptation. We picked a rate of $\alpha = 0.25$ based on a prior study that evaluated performance as a function of adaptation rate (Madduri et al., 2022).

2.2 Participants

We recruited 11 participants (5 males, 5 females, 1 non-binary; 2 left-handed, 9 right-handed; aged 19 to 32). They were included in the study if they had normal or corrected-to-normal eyesight and no history of upper-limb disability. Participants with corrected-to-normal vision were required to wear contact lenses to wear the headset. All participants reported using computers or touchscreen devices daily. Three participants had used an sEMG device before, and two participants regularly worked with sEMG devices. No participants had prior experience with the sEMG device we used in this study and they were not aware of the study's purpose or conditions. Participants gave their written consent prior to the study. All procedures were approved by the University of Washington's Institutional Review Board (IRB #STUDY00014060).

2.3 Data acquisition and signal pre-processing

User myoelectric activity was recorded from a high-density sEMG 64-channel electrode (5x13 rectangular electrode layout, 4mm inter-electrode spacing from Quattrocento, Bioelectronica, Italy). Participants' arms were first shaved and then cleaned with alcohol pads before applying the electrodes. The electrode array was covered with conductive gel, and then was placed on the user's dominant forearm, targeting the Extensor Carpi Radialis, and was secured with self-adherent tape. EMG signals were recorded using OT Bio-Light Software (Bioelettronica, Italy) at 2048 Hz on Differential Mode with a built-in high-pass filter at 10

Hz, and a low-pass filter at 150 Hz. Then, raw EMG signals were digitally rectified and low-pass filtered by averaging the delinearized signals with 100 ms time bins (Yamagami et al., 2020) to obtain an EMG envelope that was input to the decoder.

Eye data were collected using PupilLabs Core headset in real time (PupilLabs, Germany) (Kassner et al., 2014). Gaze data were streamed from two infrared cameras at 250 Hz using PupilCapture Software (PupilLabs, Germany). Eight AprilTags markers were placed around a 24-inch computer monitor to define an area of interest (AOI) (Fig. 4.1B). Eye data was calibrated using a PupilCapture routine that asks users to gaze at static circles on the screen, which appeared one by one at the four corners and the center (two trials per target for each subject). PupilCapture recorded the gaze positions on AOI, gaze confidence level (the quality assessment of pupil detection), and the pupil diameters of each eye. The monitor was 27 inches in front of the participant. The heights of the monitor, table, and chair were adjustable to maximize participants' comfort. In the condition where gaze measurements were used online to train the adaptive decoder, the gaze data in every 20-second batch was pre-processed (details in the next section).

2.4 Eye data pre-processing

To accurately estimate the task goals or the task targets from gaze, we pre-processed the raw gaze data using the following steps:

1. In a pilot data collection, we found that raw gaze delayed the target reference by approximately 250-350 ms and had a phase lag of 0.4 radians below 1 Hz when a participant intentionally gazed at the target. This roughly aligned with the human visuomotor delay in a first-order tracking task (Peterson et al., 2022). To estimate the target position from gaze position, we removed the phase lag using the fast Fourier Transform (FFT) in the frequency-domain.

2. After an initial calibration in PupilCapture, we asked participants to fixate on nine static dots on the screen, eight forming a circle and one at its center, to examine any bias in gaze measurements. These biases were due to imperfect eye-tracker calibration or inaccurate transformation from camera coordinates to screen coordinates. We then removed biases in gaze positions in both horizontal and vertical axes.

3. We masked low-confidence eye measurements by removing gaze data below a confidence threshold of 0.5, in a range from 0 to 1 given by the PupilCapture Software. Low confidence gaze could be due to the user blinking or gazing off of the computer screen.

The pre-processed gaze data had a significantly better correlation with the target trajectory r compared to the raw gaze data when a user was intentionally gazing at the target (Fig. 4.2).

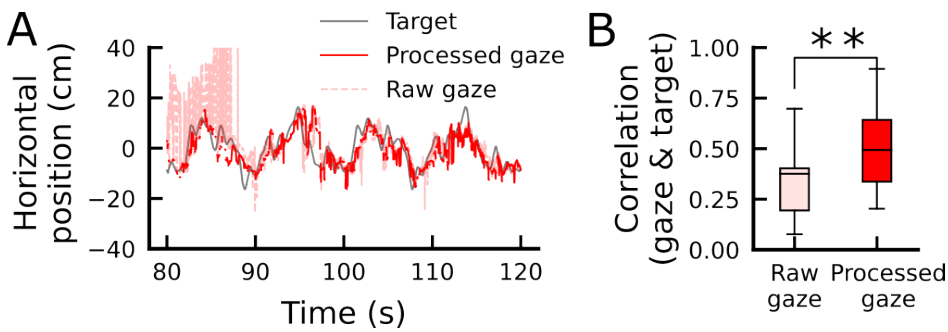

Figure 4.2: (a) An example of raw and pre-processed gaze data when the user is gazing at the reference target r. Pre-processed gaze data has fewer unwanted spikes compared to raw gaze data. (b) Boxplots (N = 11) of the Pearson correlation coefficient (range from 0 to 1) between raw gaze position and target r versus correlation coefficient between pre-processed gaze position and target r, averaged across all trials. Significant difference was labeled with $**$ (paired t-test, $p < 0.01$).

2.5 User task

The experiment was described to the participants as a continuous tracking task. We chose a continuous task for this experiment because it allows the use of control theory tools to mathematically model the user behavior, and has the potential to be extended to other continuous device-interaction scenarios, such as website browsing, robot navigation, or driving. Participants were instructed to control a white cursor with the goal of staying as close to a 2D moving red target as possible at all times. Both the target and the cursor had a radius of 1 cm. We constructed the target trajectory as sums of sinusoidal signals with random phases to create pseudorandom movements. To prevent harmonics from user responses at different frequencies, the sinusoidal signals were designed to be at frequencies of prime multiples of the base frequency (0.05 Hz) (Roth et al., 2011; Yamagami et al., 2021b; Chou et al., 2022). We adopted procedures from (Yamagami et al., 2021b) to choose stimulated frequencies below 1 Hz, as prior studies found frequencies higher than 1 Hz were difficult to track given pseudorandom stimuli and a cursor with 1st-order dynamics (Yu et al., 2014b; McRuer and Jex, 1967). We separated the stimulated signals to be at 0.1, 0.25, 0.55, 0.85 Hz in the horizontal direction, and 0.15, 0.35, 0.65, 0.95 Hz in the vertical direction. The magnitudes of stimuli were the inverse of the stimulated frequencies to distribute signal power equally. The target and cursor positions on the computer screen were updated at 60 Hz. The task was programmed using *pygame 1.9.2* and *LabGraph 2.0* packages.

2.6 Conditions and procedure

At the beginning of the experiment, participants were instructed to use a conventional computer mouse to perform two baseline trials of the tracking task. In these baseline trials, the cursor position was directly mapped from the physical movements of the mouse, and this mapping remained constant. After the baseline trials, participants performed the same tracking task but with the myoelectric interface. They controlled the myoelectric interface under two conditions: **(1) Task-trained:** using target position as decoder training data (Eq. 4.3), and **(2) Gaze-trained:** using gaze position as decoder training data (Eq. 4.4). Each condition was repeated four times. The condition orders in each repetition were randomized to control for order effects. Each trial was four minutes, a three-minute adaptive period followed by a one-minute non-adaptive (fixed) period (Fig. 4.1C). Participants were not aware of the decoder training mechanism in this experiment, nor which training condition was used in each phase of the experiment. The decoder D was initialized randomly, as a matrix sampled uniformly within the range $[0, 10^{-2}]$, at the start of each trial and then updated every 20 seconds during the adaptation period.

During the first two repetitions, which we defined as *Session 1*, we did not provide any explicit instructions to the users about gaze in order to test the use of natural gaze as a training data source. To probe whether the user's gaze intent would impact our gaze-training algorithm, we included an additional session (two repetitions; *Session 2*) where we instructed participants to gaze at the target as closely as possible (Fig. 4.1D). There was a five-minute break between the two sessions.

At the end of the experiment, participants filled out a brief questionnaire to subjectively quantify this system across three categories: task difficulty, interface accuracy, and effort.

2.7 Evaluation

We measured the performance of the adaptive decoder by quantifying the continuous task performance as the time-domain tracking ability (Eddy et al., 2023; Yamagami et al., 2020). We quantified task performance during the time interval $[t_0, t_1]$ using the mean square error between the 2D target r and cursor y positions:

$$|r - y|^2 = \frac{1}{n} \sum_{t \in [t_0, t_1]} ||r(t) - y(t)||^2. \tag{4.7}$$

where n is the number of samples in the interval $[t_0, t_1]$. Lower $|r - y|^2$ values correspond to better task performance.

To assess whether performance improved within trials, we compared performance from the *first decoder update* to the *last decoder update* and the *fixed decoder period* of a trial (as illustrated in Fig. 4.1C). Since the decoder was initialized randomly and updated after 20 seconds, we removed the initial 20 seconds of data and used the subsequent 20 seconds, i.e., the *first decoder update*, for comparisons. We verified normality of our data using Shapiro-Wilks test (task-trained: $p = 0.67$; gaze-trained: $p = 0.89$) and applied analysis of variance (ANOVA) test with a significance level of $\alpha = 0.05$. We used paired t-tests with Bonferroni correction as post-hoc tests. We had the following hypotheses for each condition:

H1: Performance improved within trials.
H2: Performance did not change after the decoder is fixed.

In addition to within-trial performance, we compared performance between trials and under different conditions. User performance in *Session 2* may be expected to differ from that of *Session 1* for at least two reasons. First, users had more experience with the interface during the second session than in the first, so although the decoders were initialized randomly, the user may have changed their behavior due to learning. Second, users were explicitly instructed to fixate their gaze on the target in the second session. We used a two-way repeated measure ANOVA to test the effect of session orders (*Session 1* and *2*) and conditions (task- and gaze-trained decoder) on performance in the *last decoder update*. We used paired t-tests with Bonferroni correction for post-hoc tests. We hypothesized performance across sessions and conditions as:
H3: Performance improved from session 1 to session 2.
H4: Performance of the gaze-trained decoder was not different from the task-trained decoder.

To evaluate the feasibility of our myoelectric interface, we compared its performance with that of the conventional mouse interface. We used ANOVA tests to compare the performance of task- and gaze-trained decoders in the *last decoder update* with the baseline. We hypothesized that:
H5: Performance of task- and gaze-trained decoders is not different from the mouse baseline.

To assess how users' gaze matched the target or cursor, we examined the relationship between gaze behaviors and task information in the frequency domain. In the frequency domain, we can decompose complex gaze signals into sinusoidal components using fast Fourier Transform, and then calculate the transformations between gaze and target or cursor:

$$T_{gr} = \widehat{g}/\widehat{r} \tag{4.8}$$

$$T_{gy} = \widehat{g}/\widehat{y} \tag{4.9}$$

where $\widehat{g}$, $\widehat{r}$, and $\widehat{y}$ are gaze, target, and cursor after Fourier transform, respectively. We hypothesized that the pre-processed natural gaze was a proxy of the task target:
H6: Magnitude and phase of the transformation T_{gr} are approximately I (identity) and 0, respectively.

We used *R 4.3* for ANOVA tests and *Python 3.9* for Shapiro-Wilks tests and paired t-tests.

2.8 User encoder models

In this co-adaptive system, the user and decoder influence each other via closed-loop interactions. To characterize how users behaved and changed, we used a model to mathematically separate the contributions of the user from those of the decoder as developed in Madduri *et al.* (Madduri et al., 2026). We modeled each user as a closed-loop *encoder* that mapped visual feedback of the target and cursor (inputs) into sEMG activity (outputs, s) (Fig. 4.4A):

$$s = E \cdot \begin{bmatrix} r \\ \dot{r} \\ r - y \\ \dot{r} - \dot{y} \end{bmatrix} + \beta, \tag{4.10}$$

where E is the encoder and $\mathrm{E} \in \mathbb{R}^{64\times 8}$, r and $\dot{r}$ are 2D target position and velocity, y and $\dot{y}$ are 2D cursor position and velocity, and β is the offset. We defined the encoder as a combination of feedforward (F) and feedback (B) controllers (Yamagami et al., 2021b):

$$E = [F_0 \ F_1 \ B_0 \ B_1], \tag{4.11}$$

where F_0 and F_1 are associated with tracking the target and target velocity, B_0 and B_1 are associated with rejecting error and error velocity, respectively, and each is a 64×2 matrix. The subscripts 0 and 1 represent the 0th- and 1st-order dynamics, respectively. In a perfectly tracking scenario, the user's feedforward controller is the pseudo-inverse of decoder D and a first-order machine dynamics M. Thus, the theoretical ideal for feedforward is: $D \cdot F_0 = 0$ and $D \cdot F_1 = I$, where I is a 2×2 identity matrix. The theoretical ideal for feedback $D \cdot B_0$ and $D \cdot B_0$ are negative-definite matrices, also 2×2, to ensure closed-loop stability (Åströn and Murray, 2021). We fit the encoder model using linear least squares with L2 regularization on 20-second batches of data.

3 Results

3.1 Performance of both task- and gaze-trained decoders improved during co-adaptation

We tested our hypotheses that task performance improves within trials and remains stable after the decoder is fixed (**H1** and **H2**) using one-way ANOVA tests comparing task errors between the *first decoder update*, the *last decoder update*, and the *fixed decoder* period. We found significant within-trial improvement in the task-trained condition (task-trained: $F_{2,30} = 5.94$, $p < 0.01$; gaze-trained: $F_{2,30} = 1.89$, $p = 0.17$). Further comparison with post-hoc paired t-tests, with corrected significance level of $\alpha/3 = 0.017$, showed that the performance of both task- and gaze-trained conditions improved significantly from the *first* to the *last decoder update* (task-trained: $t(10) = 4.93, p < 0.001$; gaze-trained: $t(10) = 3.4, p < 0.01$) (Fig. 4.3AB). The averaged performance from the *last decoder update* to the *fixed decoder* decreased 15.34% and 7.67%, in task- and gaze-trained conditions, respectively, though not statistically significant (Fig. 4.3B). Changes in task performance co-occurred with changes in the decoder. Specifically, the magnitudes of both task- and gaze-trained decoders changed significantly from the *first* to the *last decoder update* (Fig. 4.3C; $p < 0.001$).

We then tested our hypotheses about whether performance varied across sessions and between decoder training conditions (**H3** and **H4**, respectively). In *Session 2*, participants were instructed to gaze at the target, yielding a significantly higher Pearson correlation between gaze than in *Session 1* (Fig. 4.3D). Performance also improved from *Session 1* to *2*, with task errors during the *last decoder update* reduced 13.5% and 11.2% from *Session 1* to *Session 2* in task- and gaze-trained conditions, respectively (Fig. 4.3F). However, these improvements were not found to be significant in the two-way repeated measure ANOVA test comparing both sessions and conditions (Table 4.1; $p = 0.24$).

Eight out of the eleven subjects in this study had worse performance learning to use the gaze-trained decoder compared to using the task-trained decoder (Fig. 4.5). Looking at the sEMG activities of each subject (Fig. 3E in the paper), larger sEMG magnitudes did not necessarily lead to better or worse performance in a condition.

Table 4.1: Two-way repeated measure ANOVA (2 sessions × 2 conditions) results for time-domain error in the *late* period.

Factor	$F_{1,40}$	p-value	Partial η^2
Sessions	1.40	0.24	0.03
Conditions	5.17	0.028	0.11
Interaction	< 0.001	0.99	< 0.001

3.2 Task-trained decoder had better performance than gaze-trained decoder

The two-way ANOVA test for hypotheses on performance between sessions and conditions (**H3** and **H4**) did not show a significant difference between sessions; however, it showed a main effect on conditions (Table 4.1; $p = 0.028$). Thus, we performed post-hoc paired t-tests to compare the performance of the two conditions during the *last decoder update* and the *fixed decoder* phase, averaged across all trials. Results of the post-hoc tests showed that performance of the task-trained condition was significantly better than the gaze-trained condition during the *last decoder update* ($t(10) = -3.88, p < 0.01$), but not significant during the *fixed decoder* phase ($t(10) = -2.12, p = 0.06$) (Fig. 4.3B).

We also compared the two conditions separately in *Session 1* and *Session 2* with post-hoc paired t-tests. Looking at performance during the *last decoder update*, we found no significant differences between the two conditions in either session (task- vs. gaze-trained in *Session 1*: $t(10) = -2.11, p = 0.061$; task- vs. gaze-trained in *Session 2*: $t(10) = -2.51, p = 0.031$, with corrected significance of $\alpha/2 = 0.025$) (Fig. 4.3F). These results indicate that the task-trained decoder performed better than the gaze-trained decoder, but only marginally so.

We then compared the performance of the task- and the gaze-trained decoders with that of the mouse interface (**H5**). We compared the performance of the three groups (*task-trained last decoder*, *gaze-trained last decoder*, and *mouse baseline*) using one-way ANOVA test. The test result showed significant differences between the three groups ($F_{2,30} = 20.64$; $p < 0.001$).

Post-hoc tests further showed that the performance of the baseline was significantly better than both conditions (task-trained vs. baseline: $t(10) = 8.34, p < 0.001$; gaze-trained vs. baseline: $t(10) = 4.38, p < 0.01$). However, the difference between task-trained and gaze-trained conditions (effect size = 0.85) was relatively small compared to their differences from the baseline (task-trained vs. mouse, effect size = 29.91; gaze-trained vs. mouse, effect size = 47.07). Moreover, the magnitudes of decoders were not significantly different between the two conditions (Fig. 4.3C). These results further emphasize the relatively small differences in performance between the gaze- and task-trained decoders, while also highlighting clear performance gaps in both sEMG interfaces.

3.3 User controllers were less stable in gaze-trained condition

We then examined how users controlled the interfaces by fitting the encoder model with the empirical data and teasing apart the user's feedforward and feedback control contributions (Madduri et al., 2026). We calculated the combinations of decoder and encoder, or the *decoder-encoder pairs*, and compared them to theoretical ideals. The theoretical ideals of the decoder-encoder pairs are $DF_0 = 0$, $DF_1 = I$, and negative definite matrices for DB_0 and DB_1 for a perfectly tracking scenario. Since the decoder-encoder pairs are 2 by 2 matrices, we displayed each of their four elements in Fig. 4.4B.

We found that the magnitudes of decoder-encoder pairs were roughly consistent with the theoretical ideals in both task- and gaze-trained conditions. Moreover, the near-zero magnitudes in DB_0 indicated that users exerted little feedback control over positional error, but instead relied more on feedback control of velocity. We also found that in the gaze-trained condition, both the feedforward and feedback elements deviated more from the theoretical ideals than in the task-trained condition, but we only observed significant differences in velocity tracking and in the horizontal direction, i.e., $DF_1(0,0)$ and $DB_1(0,0)$ (Fig. 4.4B, $DF_1(0,0)$: $p = 0.015$; $DB_1(0,0)$: $p < 0.01$).

In addition to the decoder-encoder pairs, we investigated decoder and encoder adaptation separately by assessing each component's within-trial changes. The within-trial changes were calculated as the magnitudes of differences between consecutive decoders (Fig. 4.4C) or encoders (Fig. 4.4D). We observed that in the trajectories of both decoder and encoder, the task-trained condition showed a steadier decrease, whereas the gaze-trained condition exhibited slight instability. Both the decoder and encoder differences within trials were larger in the gaze-trained condition than in the task-trained condition, although the differences were only significant in the encoder during the last minute of decoder adaptation (Fig. 4.4C, $p = 0.1$; Fig. 4.4D, $p = 0.025$).

We also assessed the magnitudes of EMG signals, $||\text{sEMG}||$, to investigate user efforts under the two conditions. Some subjects had significantly different effort between the two conditions (subjects 4, 5, 9 in Fig. 4.4E), but the EMG signals were not always larger or smaller under a single condition. Moreover, larger magnitudes of EMG signals did not consistently

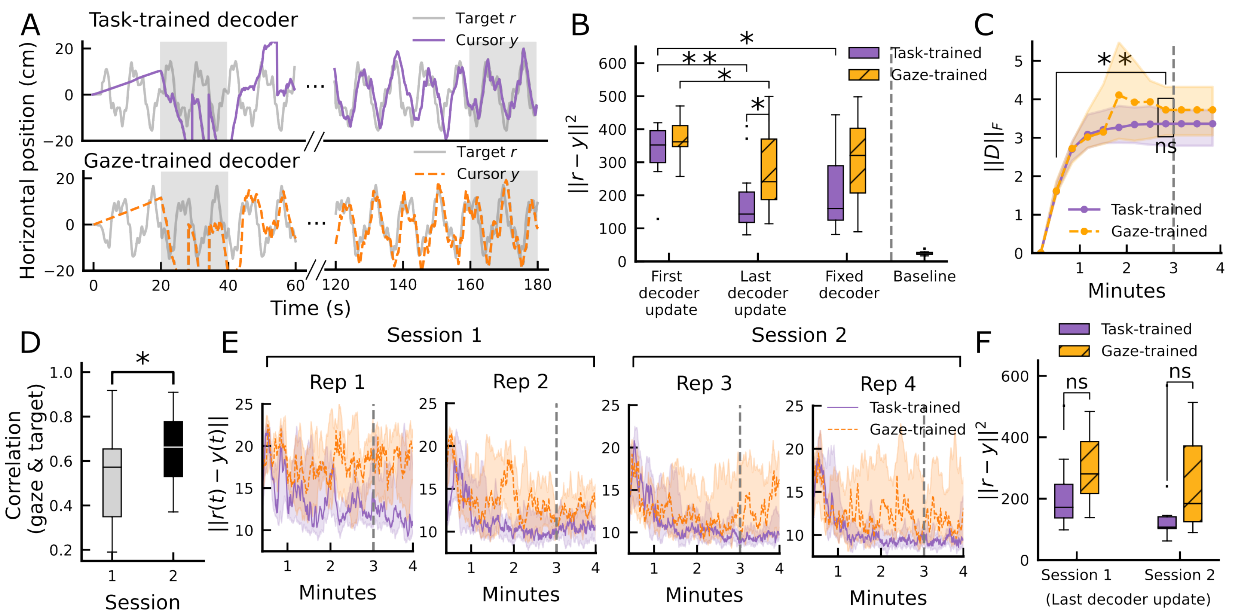

Figure 4.3: (A) Two example trials of the continuous tracking task. Both the (*top*) task-trained and the (*bottom*) gaze-trained decoder had better tracking performance during the last decoder update (gray region on the right) compared to the first decoder update (gray region on the left) of a trial. (B) Boxplots (N = 11 subjects) of task error during the *first decoder update*, the *last decoder update*, and the *fixed decoder* period, averaged across all repetitions. The average task error of the mouse baseline trials is shown for comparison. Lower task error corresponds to better task performance. Significant differences with post-hoc paired t-tests are labeled with $*$ ($p < 0.05$) and $**$ ($p < 0.01$). (C) Distributions (median, interquartile, N = 11) of magnitudes of task- and gaze-trained decoders in 4-minute trials, averaged across all repetitions. Left of the vertical dashed line is the adaptive period, and right of the dashed line is the non-adaptive (fixed) period. (D) Boxplots (N = 11) of Pearson correlation between pre-processed gaze positions and target positions in *Session 1* and *2*. (E) Distributions (median, interquartile, N = 11) of task error at each timestamp, in each repetition. The initial 20-second period with a random decoder was removed. Lower task error corresponds to better task performance. The dashed lines separate adaptive and non-adaptive periods. (F) Boxplots (N = 11) of the *last decoder update* data in (B) separated into *Session 1* and *2*. No significant difference is labeled with *ns*.

correspond to better or worse performance in either condition (Fig. 4.5).

3.4 Saccade rates and gaze confidence levels

Saccades are high-velocity periods in eye movements, and saccade rates are the number of saccades per second (Sailer et al., 2005). We defined the saccades as the baseline $\pm$ 3 standard deviation of the second-order derivative of positional gaze data. The baseline was defined as the gaze when users fixate on the center of a screen. Multiple saccades with an inter-saccadic interval less than 50 ms were considered as one saccade.

In addition to gaze positions, we recorded other gaze information, such as pupil diameters and confidence level from the PupilCapture Software. Pupil diameters are pupil sizes of the left and right eyes in pixels from 2D eye images. Gaze confidence level indicates the real-time quality of pupil detection given the eye images. Confidence level ranges from 0 to 1, with 1 being the highest confidence. Low confidence levels are due to poor pupil detections, and the reasons include blinking, blocked images (by hair, eyelashes, etc.), slippage, or lighting/exposure.

We found no significant differences in saccade rates nor the averaged pupil diameters of both eyes between the two conditions (Fig. 4.6A). In addition, we did not observe a strong correlation between saccade rates and task performance, nor gaze confidence level and task performance (Fig. 4.6B). However, the gaze confidence level is correlated with saccade rates, in both conditions (Fig. 4.6C).

In summary, we did not observe significant differences in gaze between the two conditions, in terms of saccade rates, gaze confidence levels, or pupil diameters.

3.5 Gaze behaviors influenced performance of gaze-trained decoder

We hypothesized that natural gaze (from *Session 1*) is a proxy of the task target (**H6**). We tested this hypothesis by investigating the transformations between gaze and task positions in the frequency domain (Fig. 4.7A). We found that the transformations between raw gaze and target, gray line in T_{gr}, had magnitudes close to 1 at lower frequencies (below 0.65 Hz) and negative phase angles (i.e., $|T_{gr}| \sim 1$, $\angle T_{gr} < 0$). This shows that the user's natural gaze could reach similar magnitudes as the target, but was delayed in time. Although our gaze pre-processing pipeline had already accounted for the 250-350 ms of sensorimotor delays by removing phase lag in the gaze training data, we found that the delays were larger than anticipated. The raw gaze had a larger phase delay to the target (gray line with approximately 400-700 ms delays), and after pre-processing, the gaze signals had a reduced phase delay to the target (purple and orange lines with approximately 150-350 ms delays). This means that the natural gaze, either pre-processed or raw, could approximate the task goal, but was still not a perfect representation of it.

In addition, the transformations between gaze and cursor, T_{gy}, showed that the gaze positions

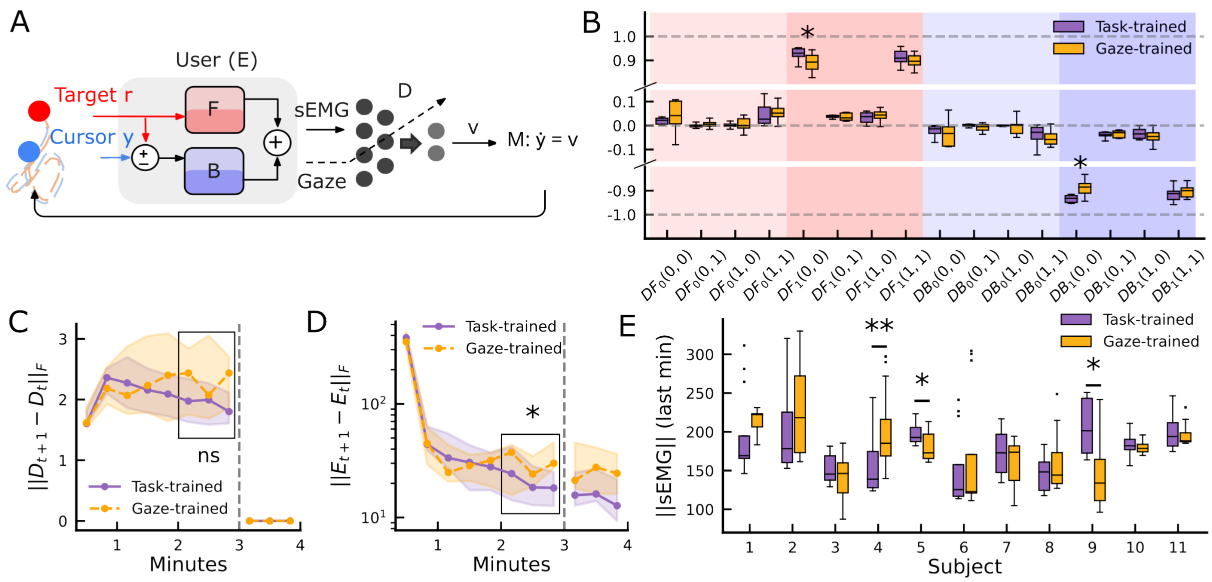

Figure 4.4: (A) Closed-loop block diagram of the adaptive myoelectric interface. User controls a 2D cursor y to follow the target using 64-channel sEMG signals. We define the encoder model E of the user as a combination of feedforward (F_0 and F_1) and feedback (B_0 and B_1) controllers. The adaptive decoder D outputs cursor velocity v_y. The cursor positions are calculated via a first-order machine dynamics M. The dashed arrow shows decoder adaptation trained with gaze data. (B) Boxplots (N = 11 subjects) of the decoder-encoder pairs DF_0 (light pink), DF_1 (dark pink), DB_0 (light blue), and DB_1 (dark blue), averaged across decoder updates and repetitions. Each matrix is 2×2, labeled as (i, j), where i is the row index and j is the column index. Significant differences are labeled with $*$ ($p < 0.05$). (C and D) Distributions (median, interquartile, N = 11) of the magnitudes (Frobenius norms) of (C) decoder changes and (D) encoder changes within trials, averaged across repetitions. The dashed vertical lines separate adaptive and non-adaptive periods. Significant difference is labeled with $*$ and no significant difference is labeled with *ns*. (E) Boxplots (N = 4 repetitions × 12 decoder updates) of the magnitudes (2-norm) of sEMG signals for each subject.

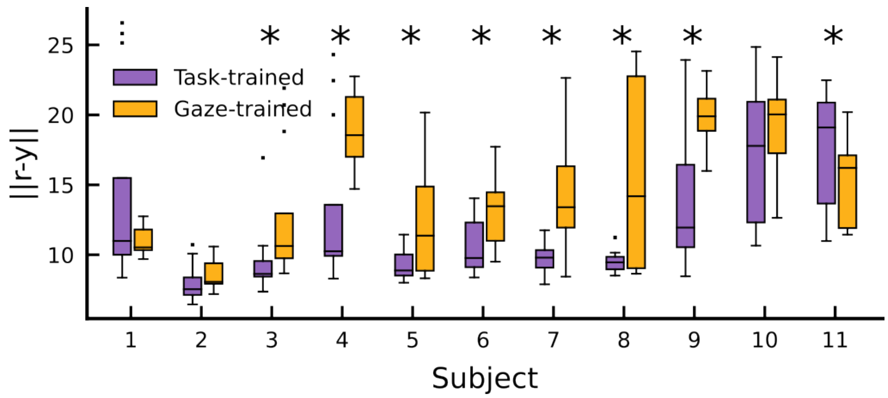

Figure 4.5: Boxplots (N = 4 repetitions × 12 decoder updates) of task error during the last minute of adaptation for each subject. Significant differences were labeled with ∗ (paired t-test, $p < 0.05$).

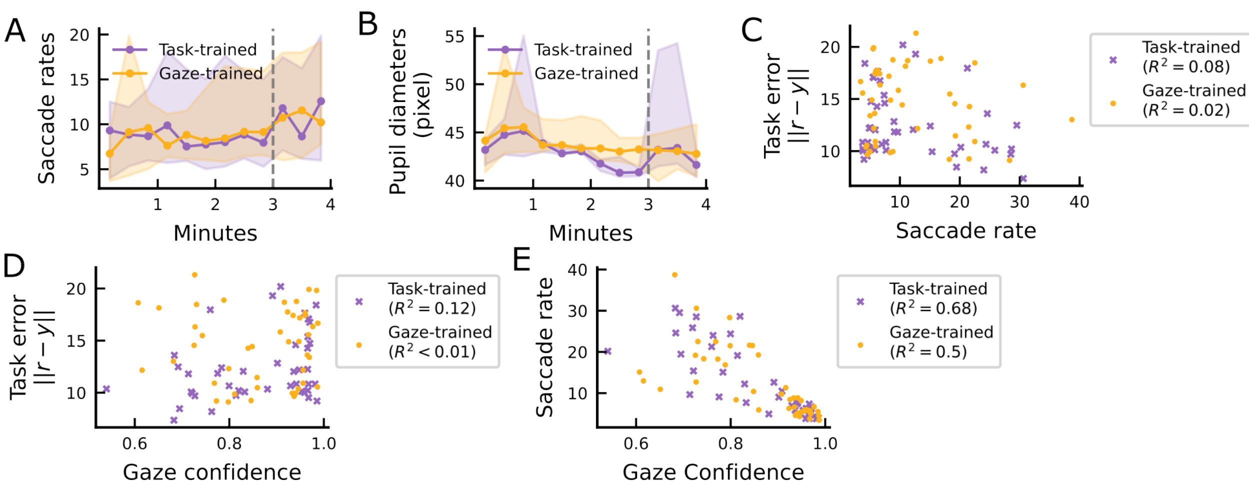

Figure 4.6: (A & B) Distributions (N = 11 subjects) of (A) saccade rates and (B) pupil diameters, averaged across every 20-second update window. The pupil diameters are averaged between both eyes. The vertical dashed line separates adaptive and non-adaptive (fixed) periods. (C – E) Relationships between (C) saccade rates and task error, (D) gaze confidence level and task error, and (E) gaze confidence level and saccade rates. Each dot is an average value of a trial (N = 11 subjects × 4 repetitions). Data with low gaze confidence below 0.5 are removed from these plots. The R^2 values of each condition are reported.

moved further than the cursor and led the cursor in phase (i.e., $|T_{gy}| > 1$, $\angle T_{gy} > 0$). Moreover, the transformation between target and cursor, T_{yr}, was only close to 1 and had an acceptable phase delay at the lowest frequencies (below 0.15 Hz). These results indicated that the cursor tried to follow but could not reach the target or gaze, and the phase delay worsened at high frequencies.

We further investigated how gaze behaviors were associated with task performance (Fig. 4.7B). When task error was low, the Pearson correlation between gaze and target, or the *gaze-target correlation*, was higher (closer to 1), and task error was more related to that correlation in the gaze-trained condition ($R^2 = 0.28$) than in the task-trained condition ($R^2 = 0.1$). However, there was insufficient evidence to determine the causal direction between gaze-target correlation and gaze-trained performance (i.e., whether gaze influenced performance or performance influenced gaze). We investigated how gaze behaviors may have contributed to a diverse performance in the gaze-trained condition by comparing performance at the beginning (hollow dots) versus the end of a trial (solid dots) for each subject (Fig. 4.7C). We selected eight subjects who were able to achieve good performance in either condition, that is, had performance better than the averaged performance during the last adaptation minute, because they had successfully learned our myoelectric interface. At the beginning of a trial, all of these subjects started with high task error but had varied gaze-target correlations. During the final minute of trials, no matter what their gaze-target correlation was during the first minute, all of them achieved above-average performance with the task-trained decoder (purple solid dots). However, those who had low gaze-target correlations (i.e., in the gray region) during the first minute had below-average performance using the gaze-trained decoder (orange solid dots). These results show that people had diverse gaze behaviors, and low gaze-target correlations (bad training data) corresponded to bad gaze-trained decoder performance.

3.6 Subjective user experience

At the end of the experiment, participants provided subjective inputs, including the difficulty of the system (on a scale of 1−5, with 5 being the most difficult), the accuracy of the system (1−5, with 5 being the most accurate), and how much effort did it take to use the system (1−5, with 5 being a lot of effort). The median value of the three indices showed that participants rated the system to have relatively low (2) difficulty, slightly high (4) effort, and medium (3) accuracy (Fig. 4.8).

4 Discussion

This work used gaze to train a closed-loop adaptive decoder for myoelectric cursor control. The decoder was continuously adapted based on real-time sEMG and gaze while users performed a continuous tracking task. We evaluated the performance of this gaze-trained approach compared to the task-dependent approach and investigated the user behaviors and gaze behaviors during co-adaptation.

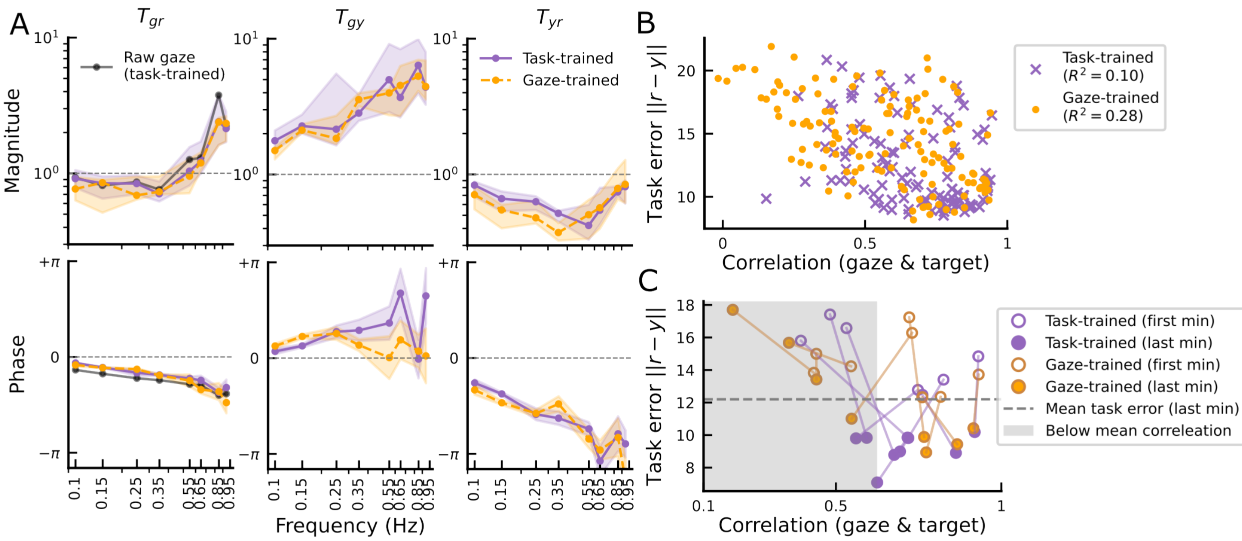

Figure 4.7: (A) Distributions (N = 11 subjects) of transformations between gaze and target T_{gr}, gaze and cursor T_{gy}, and cursor and target T_{yr} in the frequency domain below 1 Hz. For simplicity, we combined the interleaved data into a single plot, displaying data from the odd stimulated frequencies (0.1, 0.25, 0.55, 0.85 Hz) in the horizontal direction, and from the even stimulated frequencies (0.15, 0.35, 0.65, 0.95 Hz) in the vertical direction. The transformation of raw gaze and target was overlaid in the T_{gr} plot, showing the phase shift from raw gaze (gray) to pre-processed gaze (purple/orange) data. (B) The relationship between task error and gaze-target Pearson correlation. Each dot (or cross) is the averaged value in a 20-second decoder update window, and averaged across trials (N = 11 subjects × 12 decoder updates). The R^2 values of task error and gaze-target correlation for each condition are reported. Data with low gaze confidence levels (i.e., averaged gaze confidence below 0.8) is removed. (C) The relationship between task error and gaze-target Pearson correlation of each subject during the first (hollow dots) and the last minute (solid dots) of adaptation (N = 8, two conditions each). The horizontal dashed line indicates the mean task error during the last minute of adaptation, averaged across subjects. The gray region marks below the mean gaze-target correlation across subjects.

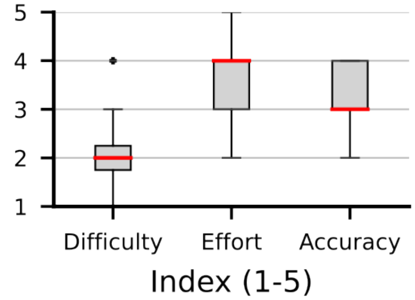

Figure 4.8: Boxplots (N = 11 subjects) of the difficulty, effort, and accuracy indices reported by the participants. The indices range from 1 to 5, with 5 being the most difficult, largest effort, and highest accuracy. The red lines indicate the median values.

4.1 The gaze-trained decoder worked but its performance fell behind task-trained decoder and baseline

Our results support hypothesis **H1** that both gaze- and task-trained methods worked effectively for training the adaptive decoder (Fig. 4.3AB). This is consistent with the findings in Madduri *et al.* (Madduri et al., 2022) and Aronson *et al.* (Aronson and Admoni, 2022), where the former showed that their (task-trained) adaptive decoder improved tracking performance within minutes of training and the latter showed that natural gaze assisted in goal predicting for device control. Our results also showed a slight, though not statistically significant, decrease in performance after transitioning from continuous adaptation to the fixed decoder period, especially for the gaze-trained decoder (Fig. 4.3B). This finding fails to reject our **H2**. It not only highlights the importance of continuous decoder adaptation, but also indicates that the gaze-trained decoder might not yet have converged toward a solution that would work well when the decoder stopped adapting. Prior studies have shown that in decoder-encoder co-adaptation, the decoder adaptation rate could influence the convergence to stationary points (equilibria) that correspond to stable interaction (Madduri et al., 2026). Thus, the gaze-trained adaptive algorithm might require a longer training time or a different adaptation rate than that of task-based training methods. Our results, indeed, suggest that user behavior differed between the two decoder training algorithms (Fig. 4.4).

We also found that the task-trained decoder still performed better and had smaller between-subject variability than the gaze-trained decoder, which leads us to reject **H4** (Fig. 4.3BE). It is not surprising that access to true task information yielded higher performance; nevertheless, post-hoc analyses showed that even without knowledge of the task, the gaze-trained decoder still demonstrated its feasibility under both natural or intentional gaze (Fig. 4.3F). Although we showed feasibility in both task- and gaze-trained decoders, the conventional mouse (baseline) still had significantly better tracking performance than both decoders, leading us to reject **H5**. This underscores the challenges in designing personalized neuromotor interfaces with high-dimensional inputs. Although our findings reject both hypotheses **H4** and **H5**, the performance of gaze-trained decoder was relatively close to the performance of task-trained decoder compared to their differences from the baseline. This suggests its potential to achieve performance comparable to the current task-based method in real-world deployment.

Several factors may also contribute to the overall lower performance we observed in the myoelectric interface. The continuous decoding that we were solving was more challenging compared to the discrete decoding used in many other myoelectric studies. For instance, prior studies restrict users to a fixed set of gestures for discrete tasks, such as gesture recognition (Saponas et al., 2008; Li et al., 2025). In contrast, we allowed users to customize their continuous hand movements at their own pace, which can support diverse strategies from different individuals (Yamagami et al., 2023b). Users could also modify their strategies using different hand movements during decoder adaptation.

Another factor that made our interface less comparable with the baseline was that our algorithm performed initial calibration from scratch (i.e., a random decoder initialization) in every trial; however, we still achieved reasonable cursor control within three minutes of training. In the future, initializing with a generic decoder trained from a variety of users would likely increase the speed of personalization with adaptive decoding approaches. Such fine-tuning approaches may also reduce the need for high-quality training data, increasing the feasibility of gaze-based training.

4.2 Both natural and intentional gaze worked for the gaze-trained decoder

Unlike interfaces that use someone's gaze as direct control inputs (Mateo et al., 2008; Pai et al., 2019; Tall et al., 2009), we intended to use natural gaze as training data for our adaptive algorithm to minimize fatigue associated with intentional eye-tracking (Jacob and Karn, 2003). However, we anticipated that intentional gaze provides a more accurate representation of the task goal than natural gaze, so we tested how instructing gaze affected our training paradigm by giving participants additional instruction on gaze fixation in *Session 2* of our experiment. We observed that the performance of both decoders improved in *Session 2*, with a lower error shown in Fig. 4.3EF, but was not statistically significant, which partially supported our **H3**. However, our study design could not determine whether this improvement was related to our gaze instructions or the user's additional experience with the interface in the latter session. In neuromotor interfaces, practice has been shown to lead users to develop more consistent control strategies and better performance (Orsborn et al., 2014; Ganguly and Carmena, 2009; Madduri et al., 2026).

We did observe that gaze was closer to the task target in *Session 2* than in *Session 1* (Fig. 4.3D), but the improvement in the gaze-trained decoder between sessions was not as apparent as we expected (11.2%) compared to the improvement in the task-trained decoder (13.5%) (Fig. 4.3F). This could be due to several factors, including a shift in user strategy prompted by the instruction, or data quality issues such as eye-tracking headset slippage. We only calibrated the headset at the beginning of the experiment, which might lead to accumulated errors over time (Kassner et al., 2014), and thus led to a reduced performance of the gaze-trained decoder.

Nevertheless, our results demonstrated that the gaze-trained decoder operated effectively without guided instruction on where to look (*Session 1*), and the gaze-trained decoder did not substantially benefit from intentional gaze (*Session 2*) (Fig. 4.3F). These findings highlight the robustness of our training paradigm and suggest that fixation instruction might not be necessary to train for a gaze-trained decoder; instead, general practice across sessions or guidance in strategies might be sufficient to account for the improvement across trials. However, additional studies are still needed to know whether and how guided instruction affects performance.

4.3 User behaviors were less consistent when using gaze-trained decoders than task-trained decoders

We also investigated how decoders and users' encoders evolved during co-adaptation, which provides additional insights into our decoder algorithm's performance. We observed a similar progression in co-adaptation between the two conditions, where we found no significant difference in magnitudes of the last decoders (Fig. 4.3C), and we observed generally similar trends in their decoder-encoder pairs (Fig. 4.4B). However, the changes in the decoder within trials show that the gaze-trained decoder had a less stable adaptation trajectory compared to the task-trained decoder (Fig. 4.4C). This is not a surprising finding, as we observed that the gaze-trained condition not only had a worse overall performance, but also had a larger decrease in performance when the decoder stopped adapting compared to the task-trained condition. These observations again indicate potentially a poorer convergence between the decoder and the encoder in the gaze-trained than the task-trained condition.

Moreover, the changes in the encoder within trials revealed that users responded to the decoders differently when the algorithm was adapting based on different signals (task or gaze) (Fig. 4.4D). The gaze-trained condition had larger within-trial changes in the encoder than the task-trained condition, suggesting that the user strategies were less consistent. Although the user strategies changed more, we did not find that the gaze-trained condition necessarily required a larger user effort, i.e., larger magnitudes of muscle contractions, than the task-trained condition (Fig. 4.4E). In addition, some subjects had significantly different magnitudes of sEMG activities when operating the two different decoders (Fig. 4.4E), but a larger sEMG magnitude was not associated with better performance of a condition (Fig. 4.5). These suggest that different training paradigms might lead to different user strategies or effort, but the differences were subject-dependent, which further highlights the need for personalized adaptation.

4.4 Diverse gaze behaviors led to diverse performance of the gaze-trained decoder across subjects

Our results demonstrated that using natural gaze to train an adaptive decoder was feasible because natural gaze served as a proxy of the task goal, which supports our **H6** (T_{gr} in Fig. 4.7A). In this study, we pre-processed the raw gaze data to account for its phase delay to the target. Although our data pre-processing slightly reduced the phase delay, the remaining delay to the target was still substantial (Fig. 4.7A). This led to less accurate training labels for the gaze-trained decoder than the task-trained decoder, which had direct access to the target without delay. Moreover, our pre-processing pipeline did not capture moments when the gaze prediction might fail, such as when a user was distracted or had eye fatigued. These suggest that how we pre-processed the gaze data was just a start; we need a better model to more accurately predict the user's intention and attention from gaze measurements, such as data-driven models to classify real-time goals from gaze features in various scenarios (Huang et al., 2015a).

Although transformations of gaze and target/cursor showed no significant difference between the task- and gaze-trained conditions in the population level (Fig. 4.7A), the gaze behaviors were different across subjects. We observed that where people looked, the *gaze-target correlation*, was associated with performance, especially in the gaze-trained condition (Fig. 4.7B). This is reasonable because the gaze-trained method depended on the gaze data. To explain how an individual's performance was affected by their gaze, we further showed that where people looked, especially during the first minute of training, was associated with gaze-trained decoder performance in the last minute (Fig. 4.7C). These findings suggest that the subject-varied gaze led to a more diverse performance of the gaze-trained than the task-trained decoder.

4.5 Potential biomedical and HCI applications

Imagine a wearable device that a user can inattentively put on their arm that automatically calibrates as they perform various tasks or switch between tasks in a virtual environment. Our work is a step toward realizing this vision, laying the groundwork for future investigations into continuously adaptive interfaces across a variety of applications. We conclude by listing consumer and health applications that can potentially benefit from our interface.

4.5.1 Wearable health devices in AR/VR

Our gaze-trained myoelectric interface has the potential to enable computer control when eye-tracking is available, such as using hand movements to navigate a virtual environment for therapy. For example, consider a patient controlling tasks in VR to regain muscle control using a sEMG-based device and a headset. Over time, the user may begin to sweat, feel fatigued and opt for smaller movements, or switch to another task. Each of these scenarios leads to changes in sEMG signals. If using a conventional myoelectric interface, the user would need to perform additional offline calibrations to accommodate such signal changes. In contrast, our interface provides continuous adaptation on the fly, eliminating the need for offline re-calibration or prescribed tasks. This holds promise for seamlessly transitioning between activities using user-defined hand movements.

4.5.2 Assistive devices

While we tested our interface with healthy participants, the insights gained from this study have implications for assistive devices to restore motor abilities in healthcare settings. For example, our approach can potentially benefit sEMG- or electroencephalogram (EEG)-controlled wheelchair (Al-Qaysi et al., 2018). As gaze reveals a user's intended moving direction, it can be used to train an adaptive interface to navigate a wheelchair, enhancing the adaptability of current powered wheelchairs. Unlike previous gaze-controlled wheelchairs (Matsumotot et al., 2001; Tall et al., 2009), this approach integrates high-density myoelectric signals to control the interface and does not require intentional gaze.

Another potential application lies in powered prostheses or robotic limbs controlled with high-density measurements like sEMG, EEG, or intracortical activities. Myoelectric prostheses still face high rates of abandonment, partly due to poor ease of use (Farina et al., 2014). If eye information is accessible, our gaze-trained paradigm could be used to train an adaptive interface to control those devices, thus enhancing the usability of myoelectric prostheses across a diverse set of users. It is important to note that our study focused only on 2D cursor navigation. To extend its utility to the above applications, we will need to modify and validate the paradigm for 3D and multi-degree-of-freedom robotic limb control. Additionally, this will require depth information into eye-tracking data to estimate user intention in the 3D space.

4.6 Future work

This work serves as a proof-of-concept and showcases the feasibility of training an adaptive decoder using a task-agnostic approach. A preliminary exploration suggests that the gaze-trained decoder could extend beyond the tracking task presented in this paper. It has potential for broader applications, such as self-paced handwriting using forearm muscle contractions (Fig. 4.9). Our immediate next step is to expand the preliminary exploration and validate the applicability of our gaze-trained paradigm to other less structured tasks in-the-wild. Specifically, we plan to train the adaptive decoder while novice users engage in activities such as browsing web pages or writing with a cursor.

Another future step is to incorporate the dynamic nature of eye movements in our gaze-trained paradigm. Eye movements have been shown to evolve as people learn novel sensory-motor mappings (Sailer et al., 2005; de Brouwer et al., 2021) and are different between novice and skilled users (Foerster et al., 2011). This suggests that eye-tracking data will have additional benefits for interfaces beyond the training method presented in this paper. For instance, eye-tracking can be particularly helpful in detecting gradual user learning or abrupt shifts in user strategy, thereby facilitating prompt adaptation of neuromotor interfaces.

In addition, we have only tested our interface with participants without disabilities. As myoelectric and eye-tracking devices hold promise for improving device accessibility for people with motor disabilities (Yamagami et al., 2020), our plan moving forward is to test the gaze-trained interface with individuals who have upper-limb motor disabilities, such as stroke survivors. Understanding how patients co-adapt with neuromotor interfaces can help improve devices designed to aid mobility within a patient-in-the-loop framework (Cashaback et al., 2024).

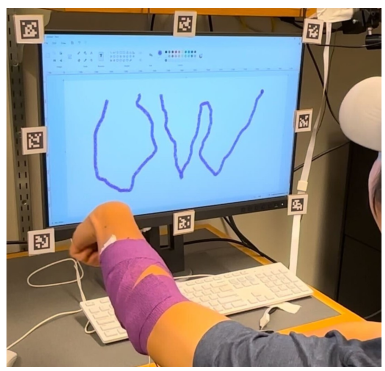

Figure 4.9: User controls a computer cursor using forearm muscle contractions to write letters ("UW") on a screen. No prescribed target or reference was given.

5 Two preliminary experiments with eye tracking

5.1 (Experiment 1) Gaze behaviors during myoelectric interface learning

While control-theoretic methods enable modeling of users controlling familiar interfaces like mice and joysticks, questions remain about how users learn to control novel interfaces, especially high-dimensional neural interfaces. Similarly, eye movement dynamics have been shown to evolve as people learn novel sensory-motor mappings (Sailer et al., 2005), and are quite different for novice versus skilled athletes (Foerster et al., 2011); eye movements provide an easily-measured read-out but have not been explored in neural interface learning. Quantifying users' internal controllers and identifying biomarkers of learning will allow neural interfaces to incorporate principles of motor control to yield intuitive and effective devices. This preliminary study aims to quantify the evolution of users' internal controllers and eye behavior as they learn to control a novel peripheral nervous system interface. Users controlled a multichannel myoelectric interface in a visuomotor tracking task that allowed us to estimate user controllers while also monitoring their eye behavior with eye-tracking. This study allowed us to identify relationships between user learned controllers and gaze, which may lead to new ways to monitor sensorimotor strategies.

5.1.1 Methods

We recruited seven healthy participants (5 females and 2 males; aged 20 to 28) to control a computer cursor using muscle activity from their dominant forearm. Muscle activity was measured with a 64-channel surface EMG system (Quattrocento, Bioelettronica). This myoelectric interface provided continuous user input for a two-dimensional (2D) trajectory tracking task (Madduri et al., 2022). To expedite learning this novel interface, we integrated an adaptive mapping between muscle activity and cursor velocity, represented by a decoder matrix D. Each participant completed a total of 12 trials, each lasting 4 minutes. The decoder was initialized randomly in each trial and adapted every 20 seconds. This adaptation was achieved through a supervised learning algorithm, incorporating a cost function that considered both task error and decoder effort, quantified as the Frobenius matrix norm. Previous research has demonstrated that users engaged in continuous tracking tasks can be effectively modeled by a combination of computed feedforward (user intent, F) and feedback (error correction, B) controllers (Yamagami et al., 2021a; Madduri et al., 2026). We set out to explore the relationship between eye movements and users' proficiency in managing both the feedforward and feedback aspects of this task during co-adaptation.

5.1.2 Preliminary results and discussions

We examined task performance in the time-domain and users' internal controllers in the frequency-domain. Using the time-domain task error (Euclidean distance between the target and cursor), we observed that, on average, task performance across all trials improved in the final minute compared to the first minute of each trial ($p = 0.037$, N = 7 subjects, 3 decoder updates, Wilcoxon signed-rank test) (Fig. 4.11B, left). We then investigated users' feedforward F and feedback B controllers in the frequency domain, using a control-theoretic nonparametric estimation method (Yamagami et al., 2021a). If users have accurately estimated the sensory-motor mapping, the pairing of the decoder D with the feedforward controller F equals the inverse of the interface dynamics M^{-1} (velocity-based in our experiments). We quantified the accuracy of users' feedforward controller by computing the norm difference between DF and M^{-1} at the lowest two reference stimulus frequencies (Fig. 4.10B), and found improvements in median performance at the beginning and the end of trials ($p = 0.13$, N = 7 subjects, 3 decoder updates), although not significant (Fig. 4.11B, middle). Similar analyses on users' feedback controllers (norm of the decoder D and the feedback controller B) revealed significant enhancement in error correction ($p < 0.01$, N = 7, 3 decoder updates) (Fig. 4.11B, right).

Having established that users update their internal controllers as they learn a novel neural interface, we then examined the relationship between users' gaze and their learned controllers.

We measured gaze proximity to the target relative to the cursor (Fig. 4.12A), and found that users who exhibited significant task performance improvement tended to have their gaze positions progressively shift closer to the target than the cursor. We found significant changes with respect to this metric comparing the initial and final minute ($p = 0.016$, N = 4 subjects, 3 decoder updates) (Fig. 4.12B, right). These changes in gaze behavior were positively correlated with user learning, measured either by tracking error or the user's feedforward controller, with averaged R-squared values of 0.26 and 0.12 across all seven subjects, respectively (Fig. 4.12C).

In summary, our results with seven subjects showed that as users learned to control a myoelectric interface with an adaptive mapping, system performance improved significantly and users' proficiency in controlling both the intent (feedforward) and error correction (feedback) aspects of the interface also improved. We found that eye gaze increasingly aligned with the target compared to the cursor position as users learned the interface, and that eye movement patterns were correlated with both system performance and user intent (feedforward controller accuracy).

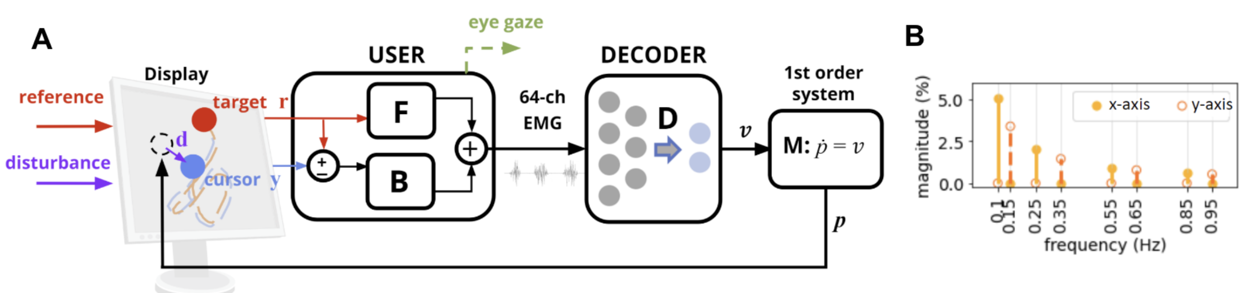

Figure 4.10: Overview of co-adaptive myoelectric interface and experimental design. (A) Users conduct a novel tracking task by controlling a computer cursor using a 64-channel surface EMG array on their dominant forearm. Reference stimuli r and cursor y are displayed on a computer screen. Users are modeled as a combination of Feedforward F and Feedback B controllers. Closed-loop decoder D adaptation is performed to update the mapping between EMG and cursor velocity. Decoder updates occur every 20 seconds. A first order system M (integrator) outputs cursor position, which was then disturbed by an external disturbance stimulus d. Users' eye gaze on the display plane are measured throughout the experiment for offline analysis. (B) Reference r and disturbance d are designed to be pseudo-random sums-of-sinusoids at stimulus frequencies. Stimulus frequencies are eight prime multiples of a base frequency (0.05 Hz). Sets of stimulus frequencies swap between the x and y axes on each trial. Frequency-domain analysis leverages established techniques from control theory to isolate user responses to independent task stimuli (reference tracking and disturbance rejection).

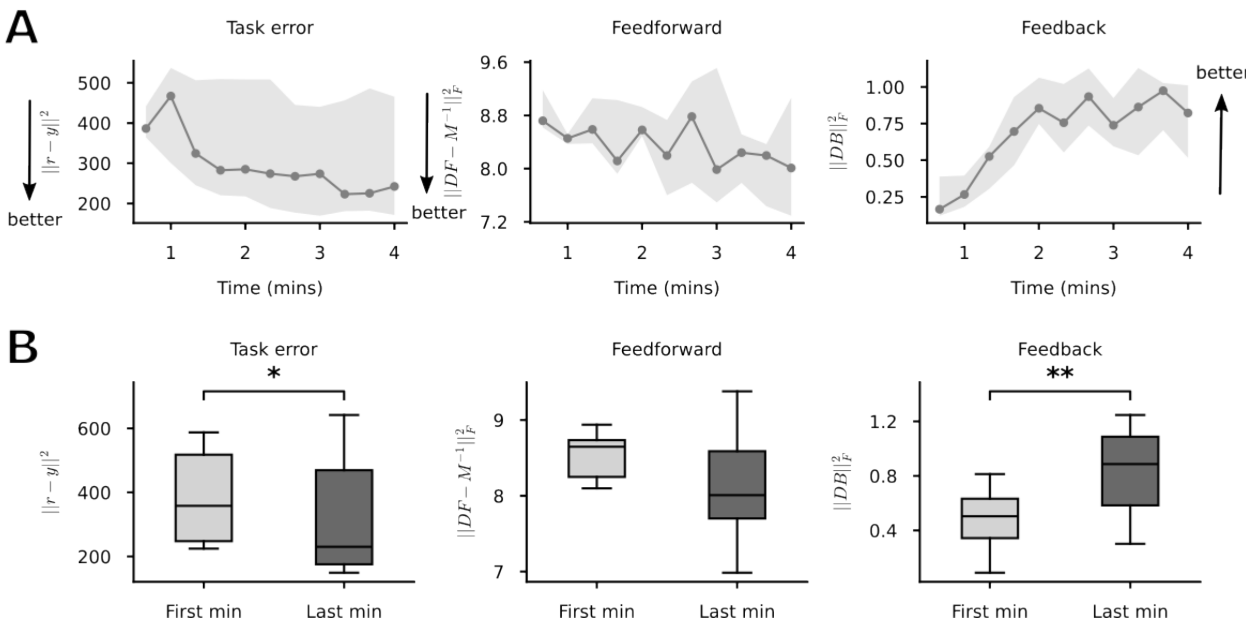

Figure 4.11: Task performance, feedforward, and feedback improved. (A) (Left) Distribution (median, interquartile, N = 7 subjects) of task error, which is defined as euclidean distance between reference and cursor in time-domain. (Middle) Distribution (median, interquartile, N = 7) of feedforward performance, which is defined as the pairing of the decoder D with the feedforward controller F and its difference between the inverse of system dynamics M^{-1} in the frequency-domain. (Right) Distribution (median, interquartile, N = 7) of feedback performance, which is defined as the pairing of the decoder D with the feedback controller B. (B) Boxplots of performance of first minute vs. last minute for each performance metric. Significant differences are marked with * (N = 7 subjects $\times$ 3 decoder updates in a minute, $p < 0.05$, Wilcoxon signed-rank test).

5.2 (Experiment 2) Using gaze to train adaptive neuromotor interfaces for diverse tasks

Here we present the preliminary findings of training the adaptive interface in diverse tasks, as an extension to the controlled continuous tracking task shown in Section 2 in this Chapter.

5.2.1 Methods

To demonstrate that the trained mapping could be generalized to a different task than the tracking task (Fig. 4.13C), a subset of our 11 participants (N = 4; 1 male, 3 females) performed a one-minute point-to-point task after each four-minute trial in the fourth block. The mapping trained in the tracking task was held fixed and used to control the cursor in the point-to-point task. We compared the performance of conditions by applying a single-factor ANOVA test on the average deviations. We hypothesized that the three conditions would have no difference in point-to-point task performance.

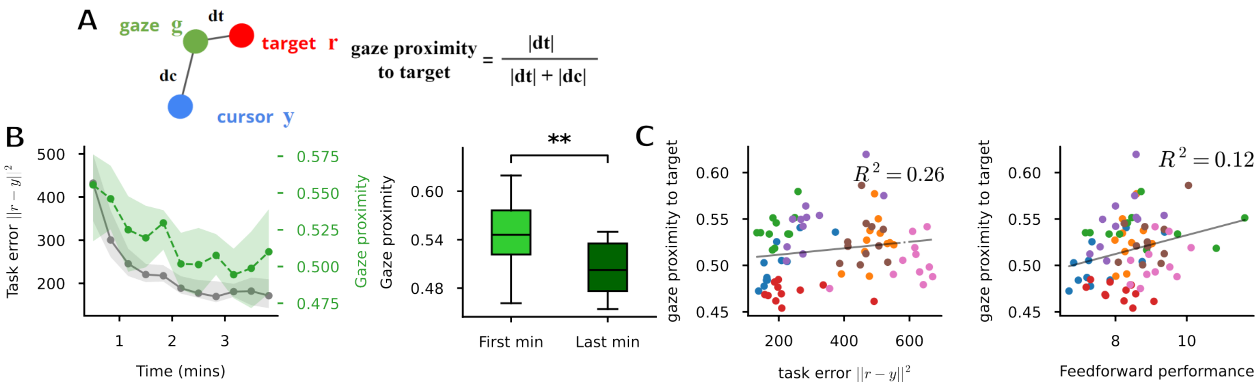

Figure 4.12: Gaze shifts closer to the target compared to the cursor as user feedforward control improves. (A) Gaze proximity to target is defined as the normalized gaze distance to target reference. Lower value means participant gaze position got relatively closer to the target compared to the cursor. (B) We compared the first and last minute of the time-domain task error for each user and pulled four users who had significant improvement in task performance. (Left) Distribution (median, interquartile, N = 4) of gaze proximity to target (green dash line) overlays on averaged task error (grey solid line), averaged across trials. Values are averaged in 20-second windows. First 20-second ramp is removed. (Right) Boxplot of gaze proximity to target of the first minute vs. the last minute after the ramp. Significant differences are marked with $*$ (N = 4 subjects $\times$ 3 decoder updates in a minute, $p < 0.05$, Wilcoxon signed-rank test). (C) (Left) Correlation between gaze proximity to target and time-domain task error, each dot is an averaged value of a 20-second window. R2 values of the 7 subjects are averaged. (Right) Correlation between gaze proximity to target and feedforward performance $\|DF - M^{-1}\|^2$, 7-subjects.

To further demonstrate that this novel gaze-based training approach is a step closer to our overarching goal of adapting while users perform diverse tasks, we conducted a preliminary trial to train the adaptive myoelectric interface in a point-to-point task. During this trial, one participant performed gaze-based training in a five-minute point-to-point task, starting with a randomly initialized mapping.

5.2.2 Preliminary results and discussions

The results presented above are derived from trials conducted in the tracking task. To assess the adaptability of interfaces to diverse contexts, we tested these trained interfaces in a one-minute discrete (point-to-point) task. Preliminary findings with a subset of subjects indicate that users could effectively complete a point-to-point task using the mapping trained in a tracking task (Fig. 4.14A). We found that the average deviations were comparable across conditions (ANOVA test, $p = 0.41$, $\eta^2 = 0.18$) (Fig. 4.14B). This supported our hypothesis that there was no difference in performance transfer between tasks in both gaze-trained and

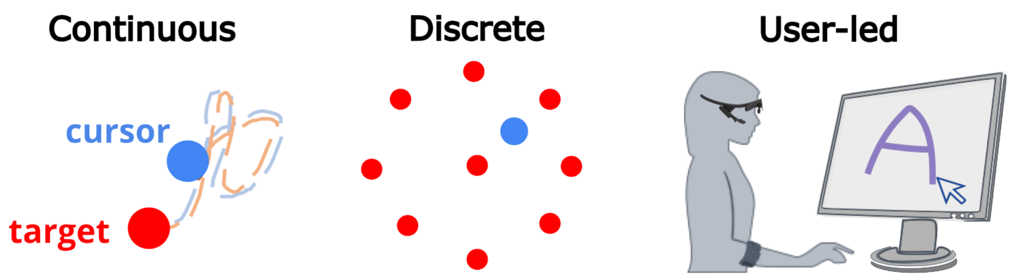

Figure 4.13: Illustrations of diverse tasks including continuous task (tracking a moving target), discrete task (pointing to a static goal), and user-led task (drawing on a screen with a computer cursor).

target-trained conditions.

In addition to testing the trained interfaces, we further demonstrated that we could train the gaze-based interface in a second task (point-to-point movements). Starting with a random initialization, the participant initially had limited control over the cursor. Over the course of five minutes of training using gaze, the participant's myoelectric control gradually improved (Fig. 4.14C). We quantified this by computing the cursor deviation from the idealized trajectory, which decreased throughout training (Fig. 4.14D). While gaze-based training generalized to the point-to-point task, we noticed a difference in the timescale of performance improvements between tasks. Preliminary evidence suggests performance improved slower in the point-to-point task, possibly due to a sparser coverage of cursor positions on the screen imposed by the task structure.

6 Conclusion

We proposed a novel method for training closed-loop adaptive decoders using natural eye movements, leveraging gaze as a proxy for task goals. Our gaze-trained approach eliminated the need for explicit task information to train adaptive decoders, offering a less supervised alternative compared to the task-dependent method. We demonstrated its feasibility in a continuous myoelectric interaction, where users controlled a dynamic cursor using forearm muscle contractions. Additionally, we showed that our method did not require users to focus their gaze on the target, which avoided drawbacks associated with intentional eye-tracking. This gaze-trained approach has the potential to extend to a wide range of tasks, including less structured computer operations with undefined task goals.

Our first preliminary results with seven subjects show that gaze increasingly aligns with the target rather than the cursor position as users gain proficiency in controlling the interface. Understanding how eye movement patterns evolve as users adapt to novel sensorimotor map-

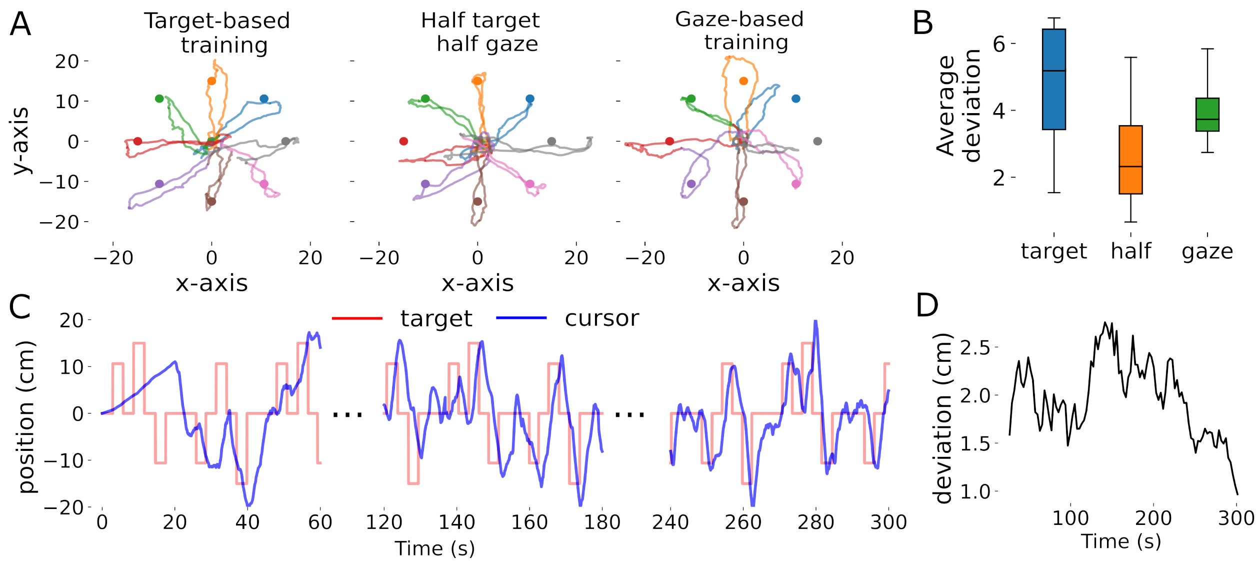

Figure 4.14: (A) **Interfaces trained in the tracking task can be generalized to a point-to-point task.** Example of a point-to-point task using a fixed mapping that was trained in a three-minute tracking task. The dots represent the targets and the different colors represent different target angles. The solid lines are the cursor trajectories reaching outward and coming back to center. (B) Boxplot (N = 4 subjects) of the average deviation of the cursor trajectory from the idealized trajectory, averaged across all targets in a trial. (C) **The gaze-trained interface could be trained in a point-to-point task.** Example of interface training in a point-to-point movement task. Performance improved throughout the five minutes of training. (D) The deviation between the cursor trajectories and idealized trajectories in the five-minute trial, moving average with a 30-second window.

pings holds great potential to improve the performance and accessibility of peripheral neural interfaces. Our second preliminary results with four subjects demonstrate the potential of our gaze-trained approach for rapid calibration across various tasks, such as discrete point-to-point movements, to generate a feasible initial decoder. This study will inform the ongoing work to develop out-of-the-box neuromotor interfaces that dynamically adapt during diverse operations, and thus contribute to a more user-friendly and accessible class of interfaces.

Chapter 5

DESIGN CONSIDERATIONS FOR CO-ADAPTIVE MULTIMODAL INTERACTIONS

Design Principles for Co-adaptive, Multimodal Interfaces

Amber H.Y. Chou, Maneeshika M. Madduri, Si Jia Li, Sasha Burckhardt, Amy L. Orsborn, Samuel A. Burden

ACM SIGCHI '24: Human-Factors in Computer Systems, PhysioCHI Workshop. May 2024.

Abstract

Biosignal-based human-machine interfaces are increasing in popularity but need to account for diversity across user physiology, strategies, and environments. Ideally, we want interfaces that can seamlessly adapt to the user during interface operation, providing customized interfaces for each user. Our prior work has generated these co-adaptive, multimodal interfaces that can be personalized to each user, ultimately expanding the accessibility and generalizability of human-machine interfaces. In this paper, we describe the advantages of co-adaptive, multimodal interfaces and report user experiences from 43 participants across five myoelectric experiments. We find that users' strategies for myoelectric control are diverse, and user backgrounds, specifically prior experience with video games, musical instruments, and computers, influence user interactions with the myoelectric interface. Additionally, we find that user onboarding with biosignals needs to be well-designed to convey myoelectric control intuitively to users. Our observations highlight the need for co-adaptive, multimodal methods to personalize interfaces for a diverse set of users and provide suggestions for designing future interfaces to facilitate user adaptation.

1 Introduction

Biosignal-based human-machine interfaces show promise across various applications, yet their usability remains limited due to variable performance among users (Farina et al., 2014; Artemiadis, 2012). This is not surprising given the diverse physiology of individuals, resulting in vastly different biosignals. Moreover, each user interacts with the interface differently, and biosignals can vary over time and across environments. Interface algorithms, which translate user biosignals to control external devices like computer cursors (Chapin et al., 1999; Taylor et al., 2002) or robotic arms (Flesher et al., 2021; Collinger et al., 2013), are typically trained offline by extracting features from biosignals and mapping them to computer commands (e.g., (Mugler et al., 2019; De Santis and Mussa-Ivaldi, 2020)). We instead consider algorithms that adapt online during user operation to personalize interfaces to users and accommodate signal variability (Orsborn et al., 2014).

In a prior study, we implemented a novel, adaptive algorithm to personalize a myoelectric interface during interface operation (Madduri et al., 2022). Then, we extended the work by training the algorithm with data from multiple modalities – surface EMG and eye tracking – to enhance interface adaptability in diverse tasks (Chou et al., 2024). In this paper, we discuss the benefits of our adaptive, multimodal approach for personalizing biosignal interfaces. We also discuss user strategies and experiences observed during our experiments to understand how users interact with co-adaptive interfaces.

Specifically, we analyze qualitative and quantitative data regarding task performance, user experiences, and diverse user strategies in learning a novel myoelectric interface. We investigate the effect of prior experiences, including video game playing, musical instrument proficiency, and computer usage, on user strategies. Avid video game players demonstrate the ability to grasp the nuances of myoelectric interface control, while individuals with experience in musical instruments and computer devices adjust their hand gestures accordingly. We also investigated electromyography (EMG) onboarding to aid users in learning myoelectric control. Surprisingly, our onboarding task confused users rather than helped them, emphasizing the need for more intuitive EMG training to facilitate learning. Lastly, we conclude with suggestions for future research in co-adaptive, multimodal interfaces that account for user experiences.

1.1 Personalizing human-machine interfaces

Many current biosignal interfaces are trained offline, often by applying machine learning or deep learning algorithms on data collected from several subjects and training the algorithms to work well on the distribution of subjects in their datasets (Saponas et al., 2008; Zhuang et al., 2019b; De Santis and Mussa-Ivaldi, 2020; Mugler et al., 2019; Sussillo et al., 2016). These approaches require lengthy setup procedures, rely heavily on experimenter-guided data collection, and need large numbers of user data to train a generalizable interface (Ctrl-Labs et al., 2024). Moreover, despite training with large datasets, the performance observed online falls short of that obtained offline (Hahne et al., 2017; Chase et al., 2009; Koyama et al., 2010; Eddy et al., 2023).

Adapting the interface online can potentially eliminate the offline-to-online performance mismatch. During online adaptation, both the interface and the user are updated synergistically in the loop, leading to co-adaptation (Orsborn et al., 2014; Hahne et al., 2017; Danziger et al., 2009; Ding et al., 2019; Madduri et al., 2023; Cashaback et al., 2024). Co-adaptation improves closed-loop performance and personalizes interfaces to diverse users during operation. It also accounts for changes within a user such as fatigue and learning. As human-machine interfaces grow more complex, interface adaptation alone might not be sufficient for proficient interface control. Users also need to adapt and develop strategies to learn the complex interfaces. With proper designs of the adaptive algorithm, interface adaptation has the potential to assist or guide user learning in human-machine interaction (Silversmith et al., 2020; De Santis, 2021; Oby et al., 2019). Leveraging co-adaptation will be key for engineering complex biosignal interfaces that require personalization and user adaptation.

1.2 A multimodal approach for human-machine interfaces

Co-adaptation has many advantages for personalizing human-machine interfaces, but future research is still needed to create seamless user experiences in real-world applications. Our overarching vision is that biosignal interfaces should be usable out-of-the-box for any user and context. To achieve this goal, the interface must provide users autonomy in choosing and changing their operations. Current approaches for co-adaptive interfaces use task structure for training the algorithms by updating the algorithm based on the error between user intention and task goals (Orsborn et al., 2014; Dangi et al., 2013; Madduri et al., 2022). While this has significantly improved user time-to-operation, this is still a task-based approach and requires users to operate a prescribed task. Using task-based training poses limitations for translation to real-world devices as users might find these training tasks laborious or constraining. As biosignal interfaces integrate into our everyday lives, we want seamless adaptation across various operations.

In prior work, we proposed a task-agnostic approach that augments eye-tracking inputs to adapt interfaces (Chou et al., 2024). We leveraged users' eye gaze to train an adaptive myoelectric interface, treating gaze as a proxy of the user intent. Our preliminary findings show the potential of generalizing the gaze-trained interface in diverse tasks, such as handwriting with a computer cursor. Adding modalities, such as eye gaze, has the potential to provide flexibility in training and re-training interfaces while operating different tasks. Moreover, multimodal interfaces incorporate diverse features from distinct modalities and thus enable more thorough measurements of complex human behaviors (Pantic and Rothkrantz, 2003).

2 Summary of experimental methods

Previously, we conducted five human-subject experiments to investigate co-adaptation in myoelectric interfaces, with a total of 43 healthy subjects (17 males, 24 females, 2 non-binary, 1 reported multiple, 1 preferred not to say; 4 left-handed, 39 right-handed; aged 19 to 34). Four participants conducted research with EMG. Our exclusion criteria included anyone who had completed this task before or had upper-limb motor impairment. All participants gave their written consent before the study and all procedures were approved by the University of Washington's Institutional Review Board (IRB #STUDY00014060).

In these experiments, users completed a continuous tracking task with a myoelectric-controlled computer cursor (Madduri et al., 2022). Users were instructed to keep their cursor close to a 2-dimensional (2D) moving target at all times. User myoelectric activity was input to an adaptive algorithm, which generated a velocity that was then integrated to determine the cursor position on the computer screen (Fig. 5.1). The algorithm was adapted every 20 seconds during user operation.

2.1 EMG and eye tracking set-up

User myoelectric activity was recorded from a 64-channel surface EMG electrode (5x13 rectangular electrode layout, 8mm inter-electrode spacing) that was placed on the user's dominant forearm, targeting the Extensor Carpi Radialis (Quattrocento, Bioelectronica, Italy). EMG signals were recorded at 2048 Hz on Differential Mode with a built-in high-pass filter at 10 Hz, and a low-pass filter at 150 Hz. Then, the EMG signals were digitally rectified and filtered with high and low-pass filters to obtain an EMG envelope, which was the input to the interface algorithm.

User eye behavior (gaze, pupil sizes) was recorded using the PupilLabs Core (PupilLabs, Germany) (Kassner et al., 2014). Eye tracking data was calibrated using the PupilCapture software (PupilLabs, Germany). Eight AprilTags were placed around a 24-inch computer monitor to define the area of interest for the PupilCapture software. Calibration was conducted by a process where users had to gaze at static targets on the screen. Gaze data was streamed at 250 Hz through the PupilCapture Software, combining the streaming data from two infrared cameras in real time. One of the experiments had conditions where gaze measurements were used online to train the adaptive interface, in which gaze data was filtered to remove low-confidence estimations and de-biased for any calibration errors (Chou et al., 2024).

2.2 Summary of past experiments

Five human-subject myoelectric interface experiments were conducted. While the experiments tested different hypotheses, the following similarities across the five experiments exist:

- All experiments tested myoelectric interface control of a computer cursor in a 2D continuous trajectory-tracking task.

- All experiments used the same EMG methods and targeted muscle activation of the Extensor Carpi Radialis for myoelectric input to the interface.

- All participants were instructed to control their cursor by activating their forearm muscles but were given no explicit instructions on which movements they should use. All participants were allowed to freely move their upper limbs.

- In all experiments, the cursor would reset to the center of the screen if participants lost control and the cursor got stuck to a border for too long (greater than 1.67 seconds).

- The task performance of the continuous tracking task was defined as the error between target and cursor positions in the time domain.

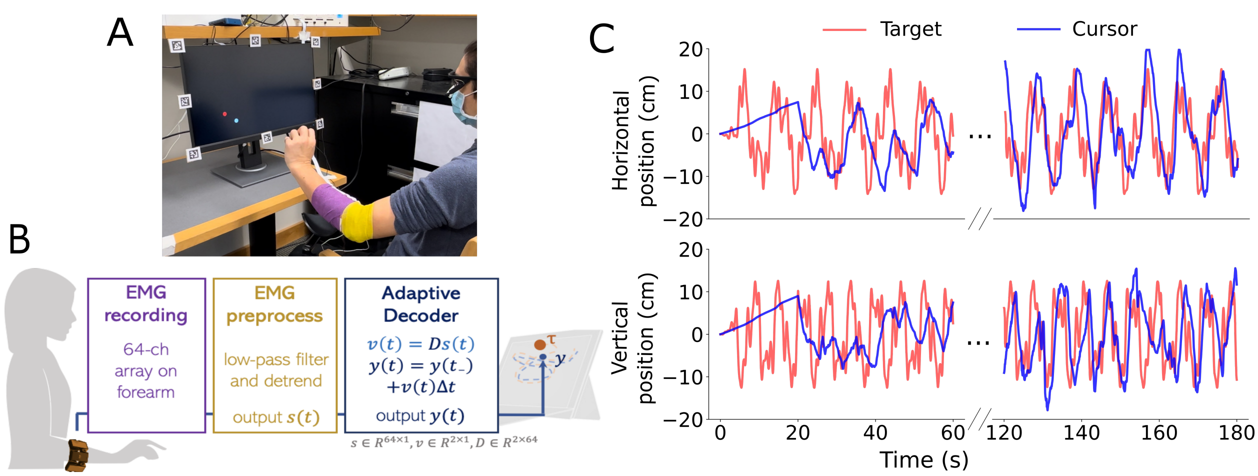

Figure 5.1: (A) The participant controls a cursor (white dot) to track a continuous target (red dot) using a 64-channel surface EMG on the dominant forearm (left arm in this figure). An eye-tracking headset measures the user's gaze on the computer screen. (B) Block diagram illustration of EMG measurements to the computer display (modified from Fig. 1 in (Madduri et al., 2022)). Image credit to Maneeshika Madduri. (C) One example trial of the continuous tracking task performance. The mapping was initialized randomly and updated every 20 seconds. The tracking performance improved throughout the three-minute trial.

3 User experiences in co-adaptive, multimodal myoelectric interfaces

In the following sections, we will present experiences reported by the participants in our co-adaptive, multimodal experiments. Specifically, we will discuss the hand movements, task strategies, the effect of prior experience, and the effect of an onboarding task.

3.1 User strategy and experience during the tracking task

Participants in our co-adaptive myoelectric interface experiments chose diverse strategies for activating the forearm muscles and controlling the cursor to track the moving target. Twelve participants described that they used their wrist as a joystick to control the direction of the cursor; eight participants reported opening palms and closing fists for control; five participants used their pointer finger or thumb to point to their intended direction; one participant reported waving their hands to activate the forearm muscles. Fourteen participants reported that they explored multiple movements during the experiment. One of them discussed that they switched from large strength-based movements to smaller, more relaxed movements. Another said they tried inverted movements (i.e. hand movements that are opposite to their original direction) when they lost control of the cursor.

We also asked the participants about their impressions of the myoelectric interface after each trial and at the end of the experiment. Eight participants said that they felt one direction (either up/down or left/right) was easier or better than the other direction. Eleven participants reported that they experienced gradual improvements in control throughout the experiment. One participant discussed that "I felt like it (the machine) was learning with me", while another said, "I relied more heavily on the algorithm". These results suggest that the users can distinguish their efforts from the algorithm's and that the user experience also includes the algorithm adaptation.

3.2 Effect of user's prior experience on strategies

In all the experiments above, we collected participants' self-reported video gaming experience, including the gaming frequency (number of hours per week) and their confidence level (novice, casual, and non-professional expert). We observed that individuals who play more than 20 hours a week on average or rated themselves as non-professional experts generally performed well in the continuous tracking task, with overall lower task error (mean Euclidean distance between target and cursor) than other participants (Fig. 5.2). Some of these high-performance subjects also comprehended the task structure quickly. For instance, one of them realized the task dynamics just after the first few trials and asked "Am I controlling the velocity/acceleration of the cursor?" In addition, one participant who identified as a non-professional expert in video gaming discussed that they "imagined a vector between target and cursor and imitated it with hand direction", which was also how we calculated the intended direction to train the adaptive algorithm. However, we did not observe exceptional performance from this participant.

We only had eight gamers in our subject pool, which may be insufficient to demonstrate a significant difference in task performance between gamers and non-gamers. Although the correlation between average gaming hours and task performance was low ($R^2 = 0.087$), we observed that all subjects who usually play video games for over 20 hours a week had low task error (Fig. 5.2). This roughly aligned with studies that found extensive gamers tend to have better eye-hand coordination and spatial attention required to adapt to complex visually guided movement using novel interfaces (Granek et al., 2010).

We also observed that users' prior experience with musical instruments and computer devices could potentially influence their strategies. For example, one participant had prior experience with piano playing, and the gesture that they picked was single-finger tapping at a time, similar to playing piano keys. Another participant mimicked the mouse control when using the EMG to control the cursor, for instance, navigating and clicking the buttons of an imaginary mouse on the table. If a participant used those movements, the experimenter would remind them that the EMG control is different from mouse control, and users could choose to move their hand by their wrists to achieve most forearm muscle control. Moreover, some participants also moved their entire arms in the air to control the cursor. Many of these movements were not recorded by the EMG electrodes and did not translate to any cursor control. These scenarios made us realize the importance of EMG-based onboarding that could help users understand how to activate task-relevant EMG muscles.

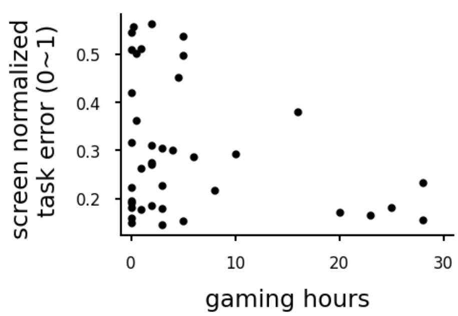

Figure 5.2: Scatter plot of normalized task error vs. gaming hours (hours/week). Each black dot is the data from one subject. The task error is the Euclidean distance between the target and the cursor, averaged across all trials and normalized to the maximum computer screen size.

3.3 Effect of an onboarding task on performance

To reduce user confusion and help users understand the task-relevant muscle activation in the tracking task, we conducted a four-subject preliminary experiment with an onboarding task before the 2D continuous trajectory-tracking task. We designed a simple "radar plot" that visualizes the absolute amplitudes of 64-channel EMG signals with 64 spikes radiating from the center of a circle (Fig. 5.3). Each spike represents the amplitude of each EMG channel. The radar plot dynamically updated on the computer screen as the participant explored which movements contributed to forearm muscle activation. Participants completed the same 2D trajectory-tracking task after the onboarding. We then compared their tracking task error with prior data from participants who did not receive a radar plot (the controls). The hypothesis was that the radar plot (visualization) training would discourage movements that did not change EMG measurements and improve user performance compared to the controls. Contrary to our expectations, we found that the task performance of participants who received the radar plot was significantly worse than that of the four worst-performing control subjects. Interestingly, the radar plot seemed to confuse the users rather than help, as users reflected that "The first trial felt confusing." and "I wasn't sure what was the point of the radar plot." We observed that the confusion might be from the directions of spikes radiating from the center, which did not correspond to the hand-moving directions nor the cursor directions in the later tracking task. While this was an exploratory investigation, these findings highlight the importance of onboarding design for biosignal interfaces.

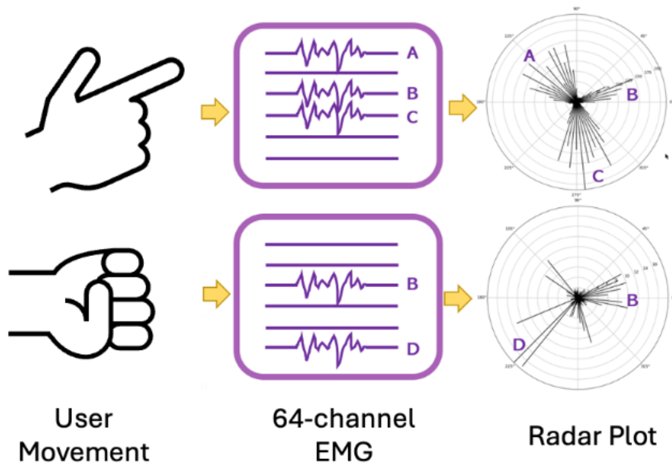

Figure 5.3: Example of a radar plot visualization from hand gestures and EMG signals. Image credit to Sasha Burckhardt.

4 Suggestions for future research

4.1 Lessons learned from co-adaptive experiment on user experience

Our co-adaptive myoelectric interface enabled participants to choose their strategies to perform the task. We observed diverse muscle activation strategies even though participants were conducting the same tracking task with electrodes on roughly the same muscle groups. We also observed that many participants selected distinct movements based on their prior experiences, such as video gaming, computer usage, or musical instruments. An open question is how much, if any, agency participants felt operating a co-adaptive interface. Autonomy and agency can positively impact users' experiences and feelings of independence when interacting with interfaces (Bennett et al., 2023). While users could perform diverse and self-defined sets of movements and gestures, users also had to account for the adaptation of the interface. Not all user explorations led to improvements in co-adaptive interface performance, which might have affected users' sense of agency and, potentially, their interface experiences. User-defined hand movements have been characterized in prior work to design intuitive gestures for human-computer interactions (Wobbrock et al., 2009). They have advantages such as feeling more natural, easier to remember, and personalized for each user's upper body. From our analysis of user prior experiences, some of the user-defined movements, like finger tapping or shoulder rotation, were difficult to measure with the EMG electrodes in our study. Thus, it is important to increase sensor coverage or add modalities to detect diverse movements, especially for individuals with limited computer experience or those with motor impairments (Yamagami et al., 2023b).

Another advantage of our co-adaptive myoelectric interface was the ability to re-train the algorithm as users operate the task. Many participants in our study discussed that they explored different gestures or changed their task strategies during the experiment. Some reported that they had muscle or eye fatigue over the two hours of the experiment. Co-adaptive interfaces are powerful as they accommodate user adaptation, such as strategies, cognitive load, and fatigue, and can dynamically update interfaces during user operation.

4.2 Challenges and future work

Our findings from our novel EMG onboarding experiment (section 3.3) emphasized the need for designing biosignal onboarding more intentionally, as some visualizations could confuse users. In addition to facilitating user learning with an onboarding task, we could combine co-adaptation with a *curriculum* to shape user learning. A curriculum could define the schedule of interface and/or algorithm adaptation with the goal of effective training in novel, complex tasks (Karpathy and van de Panne, 2012). Curriculum learning is well-studied in machine learning and has been shown to improve model performance and generalizability (Bengio et al., 2009). For example, we may design a curriculum that modulates interface algorithm parameters (i.e. learning rate) to further incentivize learning or exploration (Madduri et al., 2022; Pfister et al., 2023). Moreover, with multimodal interfaces, we can potentially guide users in learning complex tasks by adjusting the weightings between modalities from different channels (Chou et al., 2022). The challenge of curriculum design for interface adaptation is the increased number of parameters and decision variables. It becomes important to determine when and how to schedule the adaptations (Wang et al., 2022). Pfister et al. (Pfister et al., 2023) explored the effect of modulating the learning rate based on task error in a co-adaptive system on user strategies and performance. Although the modulation did not improve the overall performance, it seemed to affect how users conducted the task. A more nuanced update rule for the hyperparameters could potentially enhance users' ability to learn the novel interface more effectively. Therefore, future co-adaptive interfaces must be capable of detecting user learning stages by measuring changes in strategy, physiological behavior, the time required to achieve stability, and the capacity to integrate multiple modalities. This will allow for further modulation of the adaptive algorithms to facilitate seamless and effective learning.

Chapter 6

CONTRIBUTIONS, COLLABORATIONS, AND FUTURE RESEARCH PLANS

In this section, I summarize the main contributions of each chapter, list work led by my collaborators, and identify avenues for future work.

1 Contributions of this work

Chapter 2 demonstrates that realistic co-adaptive systems must account for delays and measurement noise — constraints that render the idealized full-information game studied in prior theoretical work impractical for real-world applications such as those involving neural recordings or arm tremors. In a human-subject experiment, we found that adaptive high-pass interfaces changed user input magnitudes and improved system performance in a dynamic human-in-the-loop game. This potentially suggests that adaptively tuning interface controller coefficients could serve as a mechanism for modulating task difficulty or engagement intensity in rehabilitation settings. We also observed distinct effort tradeoffs between the human and intelligent interface agents: as interface effort increased, human effort decreased, and vice versa. Future work is needed to explicitly model the strategic interests of both agents and use them to systematically shape co-adaptive outcomes. Additionally, we showed in simulation that a simplified model — in which the human minimizes a combination of task error and effort — reproduces the empirical co-adaptive outcomes, providing a starting point for future predictive models of human-in-the-loop behavior.

Chapter 3 demonstrates that incorporating multiple motor modalities reduces the effect of internal sensorimotor noise during machine interaction. By coordinating sEMG and manual joystick inputs in an anti-phase pattern, participants suppressed the effect of sensorimotor noise while jointly meeting the task objective of tracking the continuous target trajectory and rejecting external disturbances. We further found that the weighting assigned to each modality shaped how people coordinated their inputs: when a modality was down-weighted, users increased their effort through that channel to compensate. These findings suggest that adding a redundant motor modality may yield more precise and stable device control, and in rehabilitation contexts, may be leveraged to actively encourage muscle engagement from an impaired limb. In addition, we observed a diverse control strategies in coordinating the two input modalities across subjects, again highlighting the need to personalize multimodal interfaces for each user.

Chapter 4 introduces natural gaze as a task-agnostic source of training labels for closed-loop adaptive myoelectric decoders, which removes the need for prescribed target information and thus extend the supervised decoder adaptation to more unstructured, real-world tasks. Our experimental results with a controlled task showed that the sEMG decoder trained with gaze label data achieves practical performance compared to the decoder trained with the prescribed target signals, which demonstrates that eye tracking can meaningfully augment neuromotor interfaces. We also demonstrated the potential of gaze-trained decoder across diverse operating conditions without constraining users to specific tasks. In addition to decoder performance, we found that the choice of decoder training labels — whether derived from gaze or task targets — shaped how users adapted their strategies during co-adaptation, revealing that the training paradigm itself influences user behavior. How this influence can be systematically leveraged to steer co-adaptive outcomes in multimodal systems remains an important open question for future work.

Chapter 5 presents design considerations for co-adaptive, multimodal interfaces derived from observations and user feedback across five myoelectric experiments. We found diverse user control strategies and demonstrated that prior experience with video games, musical instruments, and computers could affect interface interactions. Furthermore, several lessons emerge from these studies, highlighting the importance of onboarding guidance, accounting for changing user strategies, and incorporating user-defined movements into interface design.

Together, this thesis addresses key research gaps in understanding how co-adaptation and multimodal inputs shape human-machine interaction, and introduces new experimental and computational frameworks for personalizing biosignal-based interfaces across diverse users and tasks. A unifying theme across chapters is that interfaces must adapt online together with the user, and that modeling this co-adaptation is essential for engineering interactions with predictable and desirable outcomes. Multimodal inputs further shape user behavior and can augment interface generalization to tasks that lack explicit structure or prescribed goals. This work is a step toward realizing continuous, automatic, and seamless calibration of biosignal-based interfaces for real-world assistive device control in diverse operations.

2 Other collaborations

The work presented in this thesis proposal has resulted in some exciting collaborations, including (in chronological order):

2.1 Modeling human delays in disturbance-rejection task

In research led by mentee Lauren Peterson (former undergraduate, UW ECE), we conducted a secondary data analysis to model human parameters (sensorimotor gain and delay) in EMG interfaces when participants are tasked with controlling a complex (first- and second-order) dynamical system. This work was published (Peterson et al., 2022).

2.2 Gaming experiences and sEMG interface control

Video-game experience enhances eye-hand coordination and spatial attention required to interact with machines (Granek et al., 2010). We investigated whether gaming experiences generalized to our EMG-controlled visuomotor tasks. In research led by Finley Hutchison (former undergraduate, UW Informatics), we recruited 10 participants and examined the relationship between their video game usage and dynamic machine interaction. Preliminary results show positive correlations between gaming hours and task performance.

2.3 Adapting user learning rates as a metric of error

In research led by Annika Pfister (former undergraduate, Wellesley College), Annika tested an adaptation of learning rate, adding a layer of adaptation to adaptive interfaces. This work was presented at the 2023 IEEE Conference on Neural Engineering and Rehabilitation.

2.4 A novel hybrid myoelectric interfaces with decoder adaptation

In research led by Si Jia Li (PhD, UW Bioengineering) in the Orsborn Lab, we extended this game-theoretic framework to a task with both continuous and discrete dynamics and have seen interesting results about how this framework can help users learn orthogonal tasks.

2.5 Patient-in-the-loop framework for neurorehabilitation

In collaboration with other researchers in modeling neurorehabilitation, a perspective piece was published (Cashaback et al., 2024).

2.6 User preference in multimodal interaction

In research led by Finley Hutchison, we extended the analysis in Chapter 3 to understand the user preferences when they were given multiple modalities to control machines. This work was presented at the 2025 Conference of the IEEE Engineering in Medicine and Biology Society (EMBC).

2.7 Synthesize generic objective functions of humans (ongoing)

In research led by Emmy Chow (PhD student, UW ECE), we aim to generate a cost function of humans interacting with dynamic machines, through a large amount of participants' data collected from an online platform.

3 Future research plans

The frameworks developed in this thesis lay the groundwork for clinical translation, particularly in device-assisted neurorehabilitation, where personalized and adaptive interfaces could improve access to and quality of care for people with motor impairments.

My plan moving forward is to implement the theoretical and experimental frameworks in this dissertation in device-assisted rehabilitation programs, such as robot-guided therapy, to deliver **patient-in-the-loop** treatments and interventions (Cashaback et al., 2024). Figure 6.1 illustrates this framework.

Neurorehabilitation improves recovery and development, but patients often receive limited in-clinic physical therapy delivered by a therapist. On the contrary, devices such as rehabilitation robots deliver high-intensity and reproducible treatments that can be delivered remotely at home. However, unlike in-clinic rehabilitation, where therapists can monitor and adjust the FITT (frequency, intensity, time, and type) of treatments to match each patient's recovery, remote rehabilitation often lacks adaptive interventions in real time. Home rehabilitation needs to be personalized and adaptive to each patient to optimize recovery outcomes. To achieve this long-term vision, I aim to develop (1) adaptive curricula and (2) adaptive device-assisted treatments.

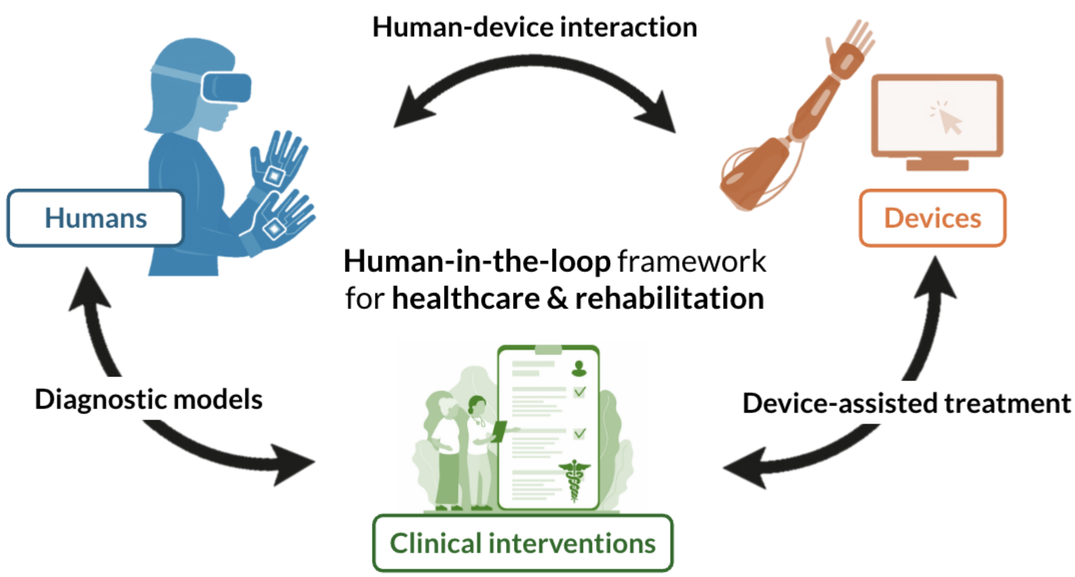

Figure 6.1: The **patient-in-the-loop** framework in designing personalized devices and clinical interventions. Humans (patients), devices (robots or computers), and clinical interventions (recovery plans or curricula) interact in a closed-loop setting. Humans' behaviors and conditions change over time, and devices need to update their algorithmic parameters to accommodate for the change in users. Clinical interventions should also change accordingly based on diagnostic models.

3.1 Adaptive curricula for home rehabilitation

A key challenge to implementing adaptive intervention is to establish the decision rules – whether, when, and based on what criteria to modify the treatment. There remains a need to identify declines (states of vulnerability) and improvements (states of opportunity) in patients, and know how they could influence outcomes, so that we can determine the appropriate modifications at the right time. I aim to characterize the dynamically changing patient needs via wearable sensor biomarkers, and then determine the frequency, timing, and rate of the just-in-time adaptive interventions (Nahum-Shani et al., 2015; Hsu et al., 2025; Pehlivan et al., 2016).

Additionally, modeling time-varying human sensorimotor control adds significant scientific value by uncovering the underlying biological mechanisms inherent to the co-adaptation paradigm, while offering valuable opportunities within the healthcare domain to predict patient progression and rehabilitation recovery.

While optimizing the control of moving devices and understanding or monitoring human health represent fundamentally distinct objectives, this dissertation focuses primarily on the former. Nevertheless, many of the computational ideas and modeling frameworks developed herein can inform the latter, particularly regarding the integration of multimodal physiological signals (e.g., respiration, heart rate, metabolic metrics, etc.).

3.2 Adaptive device-assisted treatments

To achieve adaptive interventions, once the curricula are identified, the next step is to determine the appropriate modification – such as the intensity or type of treatment. These adjustments may involve program parameters, like exercise amount or duration, or device settings, like resistance or joint torque in an exoskeleton, with constraints to ensure the exercises remain sufficiently challenging. I aim to leverage control-theoretic methods and optimization tools to provide interpretable interventions. This will benefit the automatic and real-time interventions in home settings, without adding burdens to patients or caregivers.

BIBLIOGRAPHY

Henny Admoni and Siddhartha Srinivasa. Predicting user intent through eye gaze for shared autonomy. In *2016 AAAI Fall Symposium Series*, 2016.

Z T Al-Qaysi, B B Zaidan, A A Zaidan, and M S Suzani. A review of disability EEG based wheelchair control system: Coherent taxonomy, open challenges and recommendations. *Comput. Methods Programs Biomed.*, 164:221–237, October 2018.

Scott T Albert and R Shadmehr. The neural feedback response to error as a teaching signal for the motor learning system. *J. Neurosci.*, 36:4832–4845, April 2016.

James Anderson, John C Doyle, Steven H Low, and Nikolai Matni. System level synthesis. *Annu. Rev. Control*, 47:364–393, January 2019.

Chris Wilson Antuvan, Mark Ison, and Panagiotis Artemiadis. Embedded human control of robots using myoelectric interfaces. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 22(4): 820–827, July 2014.

Reuben M Aronson and Henny Admoni. Gaze complements control input for goal prediction during assisted teleoperation. In *Robotics: Science and Systems XVIII*. Robotics: Science and Systems Foundation, June 2022.

Reuben M Aronson, Thiago Santini, Thomas C Kübler, Enkelejda Kasneci, Siddhartha Srinivasa, and Henny Admoni. Eye-hand behavior in human-robot shared manipulation. In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, New York, NY, USA, February 2018. ACM.

Panagiotis Artemiadis. EMG-based robot control interfaces: Past, present and future. *Advances in Robotics & Automation*, 01(02), 2012.

Panagiotis K Artemiadis and Kostas J Kyriakopoulos. Assessment of muscle fatigue using a probabilistic framework for an EMG-based robot control scenario, 2008.

Karl Johan Aström and Richard M Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton university press, Princeton, NJ, 2010.

Vahid Azimi, Tony Shu, Huihua Zhao, Rachel Gehlhar, Dan Simon, and Aaron D Ames. Model-based adaptive control of transfemoral prostheses: Theory, simulation, and experiments. *IEEE Transactions on Systems, Man, and Cybernetics*, 51(2):1174–1191, February 2021.

Tamer Başar and Geert Jan Olsder. *Dynamic noncooperative game theory.* SIAM, 1998.

Simone Benatti, Elisabetta Farella, and Luca Benini. Towards EMG control interface for smart garments. In *Proceedings of the 2014 ACM International Symposium on Wearable Computers: Adjunct Program*, ISWC '14 Adjunct, pages 163–170, New York, NY, USA, September 2014. Association for Computing Machinery.

Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09, pages 41–48, New York, NY, USA, June 2009. Association for Computing Machinery.

Klaus Bengler, Klaus Dietmayer, Berthold Farber, Markus Maurer, Christoph Stiller, and Hermann Winner. Three decades of driver assistance systems: Review and future perspectives. *IEEE Intell. Transp. Syst. Mag.*, 6(4):6–22, 2014.

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.

Dan Bennett, Oussama Metatla, Anne Roudaut, and Elisa D Mekler. How does HCI understand human agency and autonomy? In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, number Article 375 in CHI '23, pages 1–18, New York, NY, USA, April 2023. Association for Computing Machinery.

Bastien Berret, Enrico Chiovetto, Francesco Nori, and Thierry Pozzo. Evidence for composite cost functions in arm movement planning: an inverse optimal control approach. *PLoS Comput. Biol.*, 7(10):e1002183, October 2011.

Bastien Berret, Dorian Verdel, Etienne Burdet, and Frédéric Jean. Co-contraction embodies uncertainty: An optimal feedforward strategy for robust motor control. *PLoS Comput. Biol.*, 20(11):e1012598, November 2024.

Jeff A Bilmes, Xiao Li, Jonathan Malkin, Kelley Kilanski, Richard Wright, Katrin Kirchhoff, Amar Subramanya, Susumu Harada, James Landay, Patricia Dowden, and Howard Chizeck. The vocal joystick: A voice-based human-computer interface for individuals with motor impairments. In Raymond Mooney, Chris Brew, Lee-Feng Chien, and Katrin Kirchhoff, editors, *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 995–1002, Vancouver, British Columbia, Canada, October 2005. Association for Computational Linguistics.

D Blana, T Kyriacou, J M Lambrecht, and E K Chadwick. Feasibility of using combined EMG and kinematic signals for prosthesis control: A simulation study using a virtual reality environment. *Journal of Electromyography and Kinesiology*, 29:21–27, 2016.

Mirjam Bonanno, Beatrice Saracino, Irene Ciancarelli, Giuseppe Panza, Alfredo Manuli, Giovanni Morone, and Rocco Salvatore Calabrò. Assistive technologies for individuals with a disability from a neurological condition: A narrative review on the multimodal integration. *Healthcare (Basel)*, 13(13):1580, July 2025.

Maria Borgestig, Jan Sandqvist, Richard Parsons, Torbjörn Falkmer, and Helena Hemmingsson. Eye gaze performance for children with severe physical impairments using gaze-based assistive technology-a longitudinal study. *Assist. Technol.*, 28(2):93–102, 2016.

G E Box and N R Draper. *Empirical model-building and response surfaces*. John Wiley & Sons, New York, NY, 1987.

Meret Branscheidt, Panagiotis Kassavetis, Manuel Anaya, Davis Rogers, Han Debra Huang, Martin A Lindquist, and Pablo Celnik. Fatigue induces long-lasting detrimental changes in motor-skill learning. *Elife*, 8(e40578), March 2019.

Jose M Carmena, Mikhail A Lebedev, Roy E Crist, Joseph E O'Doherty, David M Santucci, Dragan F Dimitrov, Parag G Patil, Craig S Henriquez, and Miguel A L Nicolelis. Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol.*, 1 (2):E42, November 2003.

Maura Casadio, Rajiv Ranganathan, and Ferdinando A Mussa-Ivaldi. The body-machine interface: a new perspective on an old theme. *Journal of Motor behavior*, 44(6):419–433, 2012.

Joshua G A Cashaback, Jessica L Allen, Amber Hsiao-Yang Chou, David J Lin, Mark A Price, Natalija K Secerovic, Seungmoon Song, Haohan Zhang, and Haylie L Miller. NSF DARE-transforming modeling in neurorehabilitation: a patient-in-the-loop framework. *J. Neuroeng. Rehabil.*, 21(1):23, February 2024.

Jun Min Cha, Juntaek Hong, Jehyun Yoo, and Dong-Wook Rha. Wearable robots for rehabilitation and assistance of gait: A narrative review. *Ann. Rehabil. Med.*, 49(4):187–195, August 2025.

J K Chapin, K A Moxon, R S Markowitz, and M A Nicolelis. Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nat. Neurosci.*, 2(7): 664–670, July 1999.

Steven M Chase, Andrew B Schwartz, and Robert E Kass. Bias, optimal linear estimation, and the differences between open-loop simulation and closed-loop performance of spiking-based brain-computer interface algorithms. *Neural Netw.*, 22(9):1203–1213, November 2009.

B Chasnov, M Yamagami, B Parsa, and others. Experiments with sensorimotor games in dynamic human/machine interaction. *Micro-and Nanotechnology Sensors, Systems, and Applications XI*, 10982, 2019.

Benjamin Chasnov, Lillian Ratliff, Eric Mazumdar, and Samuel Burden. Convergence analysis of Gradient-Based learning in continuous games. In *Conference on Uncertainty in Artificial Intelligence (UAI)*, volume 115 of *Proceedings of Machine Learning Research*, pages 935–944. PMLR, 2020.

Benjamin J Chasnov, Lillian J Ratliff, and Samuel A Burden. Human adaptation to adaptive machines converges to game-theoretic equilibria. *arXiv [cs.AI]*, May 2023.

Benjamin J Chasnov, Lillian J Ratliff, and Samuel A Burden. Human adaptation to adaptive machines converges to game-theoretic equilibria. *Sci. Rep.*, 15(1):29364, August 2025.

Ishan Chatterjee, Robert Xiao, and Chris Harrison. Gaze+Gesture: Expressive, precise and targeted free-space interactions. In *Proceedings of the 2015 ACM on International Conference on Multimodal Interaction*, ICMI '15, pages 131–138, New York, NY, USA, November 2015. ACM.

Radhika Chemuturi, Farshid Amirabdollahian, and Kerstin Dautenhahn. Adaptive training algorithm for robot-assisted upper-arm rehabilitation, applicable to individualised and therapeutic human-robot interaction. *J. Neuroeng. Rehabil.*, 10(1):102, September 2013.

Amber H Y Chou, Momona Yamagami, and Samuel A Burden. Evaluating a human/machine interface with redundant motor modalities for trajectory-tracking. *IFAC-PapersOnLine*, 55(41):125–130, January 2022.

Amber H Y Chou, Maneeshika Madduri, Si Jia Li, Jason Isa, Andrew Christensen, Finley (liya) Hutchison, Samuel A Burden, and Amy L Orsborn. Using eye gaze to train an adaptive myoelectric interface. *bioRxiv*, page 2024.04.08.588608, April 2024.

Amber H Y Chou, Momona Yamagami, and Samuel A Burden. Co-adaptation of dynamic human-machine interfaces [source code], April 2026.

Keith R Cole and Richard K Shields. Age and cognitive stress influences motor skill acquisition, consolidation, and dual-task effect in humans. *J. Mot. Behav.*, 51(6):622–639, January 2019.

Jennifer L Collinger, Brian Wodlinger, John E Downey, Wei Wang, Elizabeth C Tyler-Kabara, Douglas J Weber, Angus J C McMorland, Meel Velliste, Michael L Boninger, and Andrew B Schwartz. High-performance neuroprosthetic control by an individual with tetraplegia. *Lancet*, 381(9866):557–564, February 2013.

Elaine A Corbett, Eric J Perreault, and Todd A Kuiken. Comparison of electromyography and force as interfaces for prosthetic control. *J. Rehabil. Res. Dev.*, 48(6):629–641, 2011.

Ulysse Cote-Allard, Cheikh Latyr Fall, Alexandre Drouin, Alexandre Campeau-Lecours, Clement Gosselin, Kyrre Glette, Francois Laviolette, and Benoit Gosselin. Deep learning for electromyographic hand gesture signal classification using transfer learning. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 27(4):760–771, April 2019.

Ctrl-Labs, David Sussillo, Patrick Kaifosh, and Thomas Reardon. A generic noninvasive neuromotor interface for human-computer interaction. *bioRxiv*, February 2024.

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 Comput.*, 25(7):1693–1731, July 2013.

Zachary Danziger, Alon Fishbach, and Ferdinando A Mussa-Ivaldi. Learning algorithms for human-machine interfaces. *IEEE Trans. Biomed. Eng.*, 56(5):1502–1511, May 2009.

Brendan David-John, Candace Peacock, Ting Zhang, T Scott Murdison, Hrvoje Benko, and Tanya R Jonker. Towards gaze-based prediction of the intent to interact in virtual reality. In *ACM Symposium on Eye Tracking Research and Applications*, ETRA '21 Short Papers, page 1–7, New York, NY, USA, May 2021. ACM.

Peter Dayan and Laurence F Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. MIT press, London, England, 2005.

Anouk J de Brouwer, J Randall Flanagan, and Miriam Spering. Functional use of eye movements for an acting system. *Trends Cogn. Sci.*, 25(3):252–263, March 2021.

Dalia De Santis. A framework for optimizing co-adaptation in body-machine interfaces. *Front. Neurorobot.*, 15:662181, April 2021.

Dalia De Santis and Ferdinando A Mussa-Ivaldi. Guiding functional reorganization of motor redundancy using a body-machine interface. *Journal of neuroengineering and rehabilitation*, 17(1):61, May 2020.

Jörn Diedrichsen, Reza Shadmehr, and Richard B Ivry. The coordination of movement: Optimal feedback control and beyond. *Trends Cogn. Sci.*, 14(1):31–39, 2010.

Qichuan Ding, Xingang Zhao, Jianda Han, Chunguang Bu, and Chengdong Wu. Adaptive hybrid classifier for myoelectric pattern recognition against the interferences of outlier motion, muscle fatigue, and electrode doffing. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 27 (5):1071–1080, May 2019.

J Doyle. Guaranteed margins for LQG regulators. *IEEE Trans. Automat. Contr.*, 23(4): 756–757, August 1978.

John C Doyle, Keith Glover, Pramod P Khargonekar, and Bruce A Francis. State-space solutions to standard H_2 and H_∞ control problems. *IEEE Trans. Automat. Contr.*, 34(8): 831–847, 1989.

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 Trans. Cybern.*, 43(6):1936–1949, December 2013.

Frank M Drop, Daan M Pool, Marinus M van Paassen, Max Mulder, and Heinrich H Bulthoff. Effects of target signal shape and system dynamics on feedforward in manual control. *IEEE Trans. Cybern.*, 49(3):768–780, March 2019.

Gea Drost, Dick F Stegeman, Baziel G M van Engelen, and Machiel J Zwarts. Clinical applications of high-density surface EMG: a systematic review. *Journal of Electromyography and Kinesiology*, 16(6):586–602, December 2006.

Ethan Eddy, Erik J Scheme, and Scott Bateman. A framework and call to action for the future development of EMG-based input in HCI. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, number Article 145 in CHI '23, pages 1–23, New York, NY, USA, April 2023. Association for Computing Machinery.

Mohamad A Eid, Nikolas Giakoumidis, and Abdulmotaleb El Saddik. A novel eye-gaze-controlled wheelchair system for navigating unknown environments: Case study with a person with ALS. *IEEE Access*, 4:558–573, 2016.

Jacob Eisenstein and Angel Puerta. Adaptation in automated user-interface design. In *Proceedings of the 5th international conference on Intelligent user interfaces*, IUI '00, pages 74–81, New York, NY, USA, January 2000. Association for Computing Machinery.

Michael A Elliott, Henrique Malvar, Lindsey L Maassel, Jon Campbell, Harish Kulkarni, Irina Spiridonova, Noelle Sophy, Jay Beavers, Ann Paradiso, Chuck Needham, Jamie Rifley, Maggie Duffield, Jeremy Crawford, Becky Wood, Emily J Cox, and James M Scanlan. Eye-controlled, power wheelchair performs well for ALS patients. *Muscle Nerve*, 60(5): 513–519, November 2019.

Jeremy L Emken and David J Reinkensmeyer. Robot-enhanced motor learning: accelerating internal model formation during locomotion by transient dynamic amplification. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 13(1):33–39, March 2005.

Jeremy L Emken, Raul Benitez, Athanasios Sideris, James E Bobrow, and David J Reinkensmeyer. Motor adaptation as a greedy optimization of error and effort. *J. Neurophysiol.*, 97(6):3997–4006, June 2007.

Cheikh Latyr Fall, Gabriel Gagnon-Turcotte, Jean-Francois Dube, Jean Simon Gagne, Yannick Delisle, Alexandre Campeau-Lecours, Clement Gosselin, and Benoit Gosselin. Wireless sEMG-based body-machine interface for assistive technology devices. *IEEE journal of biomedical and health informatics*, 21(4):967–977, July 2017.

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 Trans. Neural Syst. Rehabil. Eng.*, 22(4):797–809, July 2014.

Dario Farina, Ivan Vujaklija, Rickard Brånemark, Anthony M J Bull, Hans Dietl, Bernhard

Graimann, Levi J Hargrove, Klaus-Peter Hoffmann, He Helen Huang, Thorvaldur Ingvarsson, Hilmar Bragi Janusson, Kristleifur Kristjánsson, Todd Kuiken, Silvestro Micera, Thomas Stieglitz, Agnes Sturma, Dustin Tyler, Richard F Ff Weir, and Oskar C Aszmann. Toward higher-performance bionic limbs for wider clinical use. *Nat Biomed Eng*, 7 (4):473–485, April 2023.

Kristin A Farry, Ian D Walker, and Richard G Baraniuk. Myoelectric teleoperation of a complex robotic hand. *IEEE Trans. Rob. Autom.*, 12(5):775–788, October 1996.

Siamac Fazli, Sven Dähne, Wojciech Samek, Felix Bießmann, and Klaus-Robert Müller. Learning from more than one data source: Data fusion techniques for sensorimotor rhythm-based Brain–Computer interfaces. *Proceedings of the IEEE*, 103(6):891–906, June 2015.

Wyatt Felt, Jessica C Selinger, J Maxwell Donelan, and C David Remy. "body-in-the-loop": Optimizing device parameters using measures of instantaneous energetic cost. *PLoS One*, 10(8):e0135342, August 2015.

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*, CHI '17, pages 6056–6061, New York, NY, USA, May 2017. Association for Computing Machinery.

Aaron Fleming, Nicole Stafford, Stephanie Huang, Xiaogang Hu, Daniel P Ferris, and He Helen Huang. Myoelectric control of robotic lower limb prostheses: a review of electromyography interfaces, control paradigms, challenges and future directions. *J. Neural Eng.*, 18 (4), July 2021.

Sharlene N Flesher, John E Downey, Jeffrey M Weiss, Christopher L Hughes, Angelica J Herrera, Elizabeth C Tyler-Kabara, Michael L Boninger, Jennifer L Collinger, and Robert A Gaunt. A brain-computer interface that evokes tactile sensations improves robotic arm control. *Science*, 372(6544):831–836, May 2021.

Rebecca M Foerster, Elena Carbone, Hendrik Koesling, and Werner X Schneider. Saccadic eye movements in a high-speed bimanual stacking task: changes of attentional control during learning and automatization. *J. Vis.*, 11(7):9, June 2011.

Anders Fougner, Erik Scheme, Adrian D C Chan, Kevin Englehart, and Øyvind Stavdahl. Resolving the limb position effect in myoelectric pattern recognition. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 19(6):644–651, December 2011.

Anders Fougner, Oyvind Stavdahl, Peter J Kyberd, Yves G Losier, and Philip A Parker. Control of upper limb prostheses: terminology and proportional myoelectric control-a review. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 20(5):663–677, September 2012.

David W Franklin, Etienne Burdet, Keng Peng Tee, Rieko Osu, Chee-Meng Chew, Theodore E Milner, and Mitsuo Kawato. CNS learns stable, accurate, and efficient movements using a simple algorithm. *J. Neurosci.*, 28(44):11165–11173, October 2008.

Shinichi Furuya, Takanori Oku, Hayato Nishioka, and Masato Hirano. Surmounting the ceiling effect of motor expertise by novel sensory experience with a hand exoskeleton. *Sci. Robot.*, 10(98):eadn3802, January 2025.

Gregory J Gage, Kip A Ludwig, Kevin J Otto, Edward L Ionides, and Daryl R Kipke. Naïve coadaptive cortical control. *J. Neural Eng.*, 2(2):52–63, June 2005.

Krzysztof Z Gajos, Amy Hurst, and Leah Findlater. Personalized dynamic accessibility. *Interactions*, 19(2):69–73, March 2012.

Karunesh Ganguly and Jose M Carmena. Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol.*, 7(7):e1000153, July 2009.

Marcus Gardner, C Sebastian Mancero Castillo, Samuel Wilson, Dario Farina, Etienne Burdet, Boo Cheong Khoo, S Farokh Atashzar, and Ravi Vaidyanathan. A multimodal intention detection sensor suite for shared autonomy of upper-limb robotic prostheses. *Sensors (Basel)*, 20(21):E6097, October 2020.

I M Gelfand and M L Latash. On the problem of adequate language in motor control. *Motor Control*, 2(4):306–313, October 1998.

Samuel J Gershman, Eric J Horvitz, and Joshua B Tenenbaum. Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245): 273–278, July 2015.

Vikash Gilja, Paul Nuyujukian, Cindy A Chestek, John P Cunningham, Byron M Yu, Joline M Fan, Mark M Churchland, Matthew T Kaufman, Jonathan C Kao, Stephen I Ryu, and Krishna V Shenoy. A high-performance neural prosthesis enabled by control algorithm design. *Nat. Neurosci.*, 15(12):1752–1757, December 2012.

S Gnatzig, F Schuller, and M Lienkamp. Human-machine interaction as key technology for driverless driving - a trajectory-based shared autonomy control approach. In *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*, pages 913–918. ieeexplore.ieee.org, September 2012.

Paula Gomes. Surgical robotics: Reviewing the past, analysing the present, imagining the future. *Robot. Comput. Integr. Manuf.*, 27(2):261–266, 2011.

Joshua A Granek, Diana J Gorbet, and Lauren E Sergio. Extensive video-game experience alters cortical networks for complex visuomotor transformations. *Cortex*, 46(9):1165–1177, October 2010.

Andrea M Green and John F Kalaska. Learning to move machines with the mind. *Trends Neurosci.*, 34(2):61–75, February 2011.

C Guger, G Edlinger, W Harkam, I Niedermayer, and G Pfurtscheller. How many people are able to operate an EEG-based brain-computer interface (BCI)? *IEEE Trans. Neural Syst. Rehabil. Eng.*, 11(2):145–147, June 2003.

Tayfun Gürel and Carsten Mehring. Unsupervised adaptation of brain-machine interface decoders. *Front. Neurosci.*, 6:164, November 2012.

Abdelwaheb Hafs, Dorian Verdel, Olivier Bruneau, and Bastien Berret. Game-theoretic interaction control for assistive exoskeletons: A 2-DOF simulation study. *IFAC-PapersOnLine*, 59(18):265–270, 2025.

Abdelwaheb Hafs, Anaïs Farr, Dorian Verdel, Olivier Bruneau, Etienne Burdet, and Bastien Berret. Model predictive game control for personalized and targeted interactive assistance. *Commun Eng*, 5(1):57, February 2026.

Janne M Hahne, Sven Dahne, Han-Jeong Hwang, Klaus-Robert Muller, and Lucas C Parra. Concurrent adaptation of human and machine improves simultaneous and proportional myoelectric control. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 23(4):618–627, July 2015.

Janne M Hahne, Marko Markovic, and Dario Farina. User adaptation in myoelectric man-machine interfaces. *Sci. Rep.*, 7, June 2017.

Wendy Hardeman, Julie Houghton, Kathleen Lane, Andy Jones, and Felix Naughton. A systematic review of just-in-time adaptive interventions (JITAIs) to promote physical activity. *Int. J. Behav. Nutr. Phys. Act.*, 16(1):31, April 2019.

Christopher M Harris and Daniel M Wolpert. Signal-dependent noise determines motor planning. *Nature*, 394(6695):780–784, August 1998.

Jiayuan He, Dingguo Zhang, Ning Jiang, Xinjun Sheng, Dario Farina, and Xiangyang Zhu. User adaptation in long-term, open-loop myoelectric training: implications for EMG pattern recognition in prosthesis control. *J. Neural Eng.*, 12(4):046005, August 2015.

James B Heald, Máté Lengyel, and Daniel M Wolpert. Contextual inference underlies the learning of sensorimotor repertoires. *Nature*, 600(7889):489–493, December 2021.

Eric B Hekler, Daniel E Rivera, Cesar A Martin, Sayali S Phatak, Mohammad T Freigoun, Elizabeth Korinek, Predrag Klasnja, Marc A Adams, and Matthew P Buman. Tutorial for using control systems engineering to optimize adaptive mobile health interventions. *J. Med. Internet Res.*, 20(6):e214, June 2018.

Helena Hemmingsson and Maria Borgestig. Usability of eye-gaze controlled computers in sweden: A total population survey. *Int. J. Environ. Res. Public Health*, 17(5), March 2020.

Christopher Hemond, Rachel M Brown, and Edwin M Robertson. A distraction can impair or enhance motor performance. *The Journal of Neuroscience*, 30(2):650–654, January 2010.

H J Hermens, B Freriks, C Disselhorst-Klug, and G Rau. Development of recommendations for SEMG sensors and sensor placement procedures. *Journal of electromyography and Kinesiology*, 10(5):361–374, October 2000.

João P Hespanha. *Linear Systems Theory.* Princeton University Press, Princeton, NJ, 2009.

Leigh R Hochberg, Mijail D Serruya, Gerhard M Friehs, Jon A Mukand, Maryam Saleh, Abraham H Caplan, Almut Branner, David Chen, Richard D Penn, and John P Donoghue. Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature*, 442 (7099):164–171, July 2006.

Leigh R Hochberg, Daniel Bacher, Beata Jarosiewicz, Nicolas Y Masse, John D Simeral, Joern Vogel, Sami Haddadin, Jie Liu, Sydney S Cash, Patrick van der Smagt, and John P Donoghue. Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature*, 485(7398):372–375, May 2012.

Wen-Bing Horng, Chih-Yuan Chen, Yi Chang, and Chun-Hai Fan. Driver fatigue detection based on eye tracking and dynamic template matching. In *IEEE International Conference on Networking, Sensing and Control, 2004*, volume 1, pages 7–12. IEEE, 2004.

Ting-Chen Chloe Hsu, Pauline Whelan, Julie Gandrup, Christopher J Armitage, Lis Cordingley, and John McBeth. Personalized interventions for behaviour change: A scoping review of just-in-time adaptive interventions. *Br. J. Health Psychol.*, 30(1):e12766, February 2025.

Chien-Ming Huang and Bilge Mutlu. Anticipatory robot control for efficient human-robot collaboration. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 83–90, March 2016.

Chien-Ming Huang, Sean Andrist, Allison Sauppé, and Bilge Mutlu. Using gaze patterns to predict task intent in collaboration. *Front. Psychol.*, 6:1049, July 2015a.

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*, UIST '15, pages 523–528, New York, NY, USA, November 2015b. Association for Computing Machinery.

Gan Huang, Dingguo Zhang, Xidian Zheng, and Xiangyang Zhu. An EMG-based handwriting recognition through dynamic time warping. In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, pages 4902–4905. IEEE, August 2010.

Na-Teng Hung, Vivek Paul, Prashanth Prakash, Torin Kovach, Gene Tacy, Goran Tomic, Sangsoo Park, Tyler Jacobson, Alix Jampol, Pooja Patel, Anya Chappel, Erin King, and Marc W Slutzky. Wearable myoelectric interface enables high-dose, home-based training in severely impaired chronic stroke survivors. *Ann Clin Transl Neurol*, 8(9):1895–1905, September 2021.

K A Ingraham, C D Remy, and E J Rouse. The role of user preference in the customized control of robotic exoskeletons. *Sci. Robot.*, 7(64):eabj3487, March 2022.

Jason T Isa, Lillian J Ratliff, and Samuel A Burden. A learning algorithm that attains the human optimum in a repeated human-machine interaction game. *arXiv [cs.GT]*, January 2025.

Andrew Jackson and Eberhard E Fetz. Interfacing with the computational brain. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 19(5):534–541, October 2011.

Robert J K Jacob and Keith S Karn. Commentary on section 4 - eye tracking in human-computer interaction and usability research: Ready to deliver the promises. In J Hyönä, R Radach, and H Deubel, editors, *The Mind's Eye*, pages 573–605. North-Holland, Amsterdam, January 2003.

Alejandro Jaimes and Nicu Sebe. Multimodal human–computer interaction: A survey. *Computer vision and image understanding*, 108(1):116–134, October 2007.

Siddarth Jain and Brenna Argall. Probabilistic human intent recognition for shared autonomy in assistive robotics. *ACM Trans. Hum. Robot Interact.*, 9(1):2, December 2019.

Beata Jarosiewicz, Nicolas Y Masse, Daniel Bacher, Sydney S Cash, Emad Eskandar, Gerhard Friehs, John P Donoghue, and Leigh R Hochberg. Advantages of closed-loop calibration in intracortical brain-computer interfaces for people with tetraplegia. *J. Neural Eng.*, 10(4):046012, August 2013.

Beata Jarosiewicz, Anish A Sarma, Daniel Bacher, Nicolas Y Masse, John D Simeral, Brittany Sorice, Erin M Oakley, Christine Blabe, Chethan Pandarinath, Vikash Gilja, Sydney S Cash, Emad N Eskandar, Gerhard Friehs, Jaimie M Henderson, Krishna V Shenoy, John P Donoghue, and Leigh R Hochberg. Virtual typing by people with tetraplegia using a self-calibrating intracortical brain-computer interface. *Sci. Transl. Med.*, 7(313): 313ra179, November 2015.

Ning Jiang, Kevin B Englehart, and Philip A Parker. Extracting simultaneous and proportional neural control information for multiple-DOF prostheses from the surface electromyographic signal. *IEEE Trans. Biomed. Eng.*, 56(4):1070–1080, April 2009.

Ning Jiang, S Dosen, K-R Muller, and D Farina. Myoelectric control of artificial limbs—is there a need to change focus? [in the spotlight]. *IEEE Signal Process. Mag.*, 29(5):152–150, September 2012.

Ning Jiang, Hubertus Rehbaum, Ivan Vujaklija, Bernhard Graimann, and Dario Farina. Intuitive, online, simultaneous, and proportional myoelectric control over two degrees-of-freedom in upper limb amputees. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 22(3):501–510, May 2014.

R S Johansson, G Westling, A Bäckström, and J R Flanagan. Eye-hand coordination in object manipulation. *J. Neurosci.*, 21(17):6917–6932, September 2001.

Patrick Kaifosh, Thomas R Reardon, and CTRL-labs at Reality Labs. A generic non-invasive neuromotor interface for human-computer interaction. *Nature*, 645(8081):702–711, September 2025.

Prakyath Kantharaju, Sai Siddarth Vakacherla, Michael Jacobson, Hyeongkeun Jeong, Meet Nikunj Mevada, Xingyuan Zhou, Matthew J Major, and Myunghee Kim. Framework for personalizing wearable devices using real-time physiological measures. *IEEE Access*, 11:81389–81400, 2023.

Andrej Karpathy and Michiel van de Panne. Curriculum learning for motor skills. In *Advances in Artificial Intelligence*, pages 325–330. Springer Berlin Heidelberg, 2012.

Maxim Karrenbach, David Boe, Astrini Sie, Rob Bennett, and Eric Rombokas. Improving automatic control of upper-limb prosthesis wrists using gaze-centered eye tracking and deep learning. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 30:340–349, February 2022.

Moritz Kassner, William Patera, and Andreas Bulling. Pupil: an open source platform for pervasive eye tracking and mobile gaze-based interaction. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*, UbiComp '14 Adjunct, pages 1151–1160, New York, NY, USA, September 2014. Association for Computing Machinery.

K Kiguchi and Y Hayashi. An EMG-based control for an upper-limb power-assist exoskeleton robot. *IEEE Trans. Syst. Man Cybern. B Cybern.*, 42(4):1064–1071, August 2012.

Jong-Sung Kim, Huyk Jeong, and Wookho Son. A new means of HCI: EMG-MOUSE. In *2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat. No.04CH37583)*, volume 1, pages 100–104 vol.1. IEEE, 2004.

Sung-Phil Kim, John D Simeral, Leigh R Hochberg, John P Donoghue, and Michael J Black. Neural control of computer cursor velocity by decoding motor cortical spiking activity in humans with tetraplegia. *J. Neural Eng.*, 5(4):455–476, December 2008.

Kyung Koh, Hyun Joon Kwon, Yang Sun Park, Tim Kiemel, Ross H Miller, Yoon Hyuk Kim, Joon-Ho Shin, and Jae Kun Shim. Intra-auditory integration improves motor performance and synergy in an accurate multi-finger pressing task. *Front. Hum. Neurosci.*, 10:260, June 2016.

Konrad P Kording and Daniel M Wolpert. Bayesian integration in sensorimotor learning. *Nature*, 427(6971):244–247, January 2004.

Shinsuke Koyama, Steven M Chase, Andrew S Whitford, Meel Velliste, Andrew B Schwartz, and Robert E Kass. Comparison of brain–computer interface decoding algorithms in open-loop and closed-loop control. *J. Comput. Neurosci.*, 29(1):73–87, August 2010.

Konrad P Körding and Daniel M Wolpert. Bayesian decision theory in sensorimotor control. *Trends Cogn. Sci.*, 10(7):319–326, July 2006.

Valentina La Scaleia, Francesca Sylos-Labini, Thomas Hoellinger, Letian Wang, Guy Cheron, Francesco Lacquaniti, and Yuri P Ivanenko. Control of leg movements driven by EMG activity of shoulder muscles. *Front. Hum. Neurosci.*, 8:838, October 2014.

Mark L Latash. The bliss (not the problem) of motor abundance (not redundancy). *Exp. Brain Res.*, 217(1):1–5, 2012.

Mark L Latash, John P Scholz, and Gregor Schöner. Motor control strategies revealed in the structure of motor variability. *Exercise and sport sciences reviews*, 30(1), 2002.

R Leeb, S Perdikis, L Tonin, A Biasiucci, M Tavella, M Creatura, and dR Millán. Transferring brain–computer interfaces beyond the laboratory: successful application control for motor-disabled users. *Artificial intelligence in medicine*, 59(2):121–132, 2013.

Mindy F Levin and Marika Demers. Motor learning in neurological rehabilitation. *Disabil. Rehabil.*, 43(24):3445–3453, December 2021.

Jing Shuang Li, Anish A Sarma, Terrence J Sejnowski, and John C Doyle. Internal feedback in the cortical perception-action loop enables fast and accurate behavior. *Proc. Natl. Acad. Sci. U. S. A.*, 120(39):e2300445120, September 2023.

Wei Li, Ping Shi, Sujiao Li, and Hongliu Yu. Enhancing and optimizing user-machine closed-loop co-adaptation in dynamic myoelectric interface. *IEEE Trans. Neural Syst. Rehabil. Eng.*, April 2025.

Xin Li, Qingyuan Qi, and Kemi Ding. Noncooperative game in multi-controller system under delayed and asymmetric information. *arXiv [math.OC]*, March 2026.

Y Li, G Carboni, F Gonzalez, D Campolo, and E Burdet. Differential game theory for versatile physical human–robot interaction. *Nat. Mach. Intell.*, 1(1):36–43, January 2019.

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 Trans. Rob.*, 32(6): 1408–1418, 2016.

Zheng Li, Joseph E O'Doherty, Mikhail A Lebedev, and Miguel A L Nicolelis. Adaptive decoding for brain-machine interfaces through bayesian parameter updates. *Neural Comput.*, 23(12):3162–3204, December 2011.

Yan Liu, Xinhao Peng, Yingxiao Tan, Tolulope Tofunmi Oyemakinde, Mengtao Wang, Guanglin Li, and Xiangxin Li. A novel unsupervised dynamic feature domain adaptation strategy for cross-individual myoelectric gesture recognition. *J. Neural Eng.*, 20(6), January 2024.

Joan Lobo-Prat, Arvid Q L Keemink, Arno H A Stienen, Alfred C Schouten, Peter H Veltink, and Bart F J M Koopman. Evaluation of EMG, force and joystick as control interfaces for active arm supports. *Journal of neuroengineering and rehabilitation*, 11:68, April 2014.

Kelly Mack, Emma J McDonnell, Leah Findlater, and Heather D Evans. Chronically under-addressed: Considerations for HCI accessibility practice with chronically ill people. In *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility*, New York, NY, USA, October 2022. ACM.

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, May 2021.

Maneeshika M Madduri, Momona Yamagami, Augusto X T Millevolte, Si Jia Li, Sasha N Burckhardt, Samuel A Burden, and Amy L Orsborn. Co-adaptive myoelectric interface for continuous control. *IFAC-PapersOnLine*, 55(41):95–100, January 2022.

Maneeshika M Madduri, Samuel A Burden, and Amy L Orsborn. Biosignal-based co-adaptive user-machine interfaces for motor control. *Current Opinion in Biomedical Engineering*, 27: 100462, September 2023.

Maneeshika M Madduri, Momona Yamagami, Si Jia Li, Sasha Burckhardt, Samuel A Burden, and Amy L Orsborn. Predicting and shaping human-machine interactions in closed-loop, co-adaptive neural interfaces. *bioRxiv*, page 2024.05.23.595598, May 2024.

Maneeshika M Madduri, Momona Yamagami, Si Jia Li, Sasha Burckhardt, Samuel A Burden, and Amy L Orsborn. Computational framework to predict and shape human–machine interactions in closed-loop, co-adaptive neural interfaces. *Nat. Mach. Intell.*, 8(3):372–387, March 2026.

Musa Mahmood, Deogratias Mzurikwao, Yun-Soung Kim, Yongkuk Lee, Saswat Mishra, Robert Herbert, Audrey Duarte, Chee Siang Ang, and Woon-Hong Yeo. Fully portable and wireless universal brain–machine interfaces enabled by flexible scalp electronics and deep learning algorithm. *Nat. Mach. Intell.*, 1(9):412–422, 2019.

Babak Mahmoudi and Justin C Sanchez. A symbiotic brain-machine interface through value-based decision making. *PLoS One*, 6(3):e14760, March 2011.

Regan L Mandryk, M Stella Atkins, and Kori M Inkpen. A continuous and objective evaluation of emotional experience with interactive play environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, pages 1027–1036, New York, NY, USA, April 2006. Association for Computing Machinery.

Albert C Manero, Shea L McLinden, John Sparkman, and Björn Oskarsson. Evaluating surface EMG control of motorized wheelchairs for amyotrophic lateral sclerosis patients. *J. Neuroeng. Rehabil.*, 19(1):88, August 2022.

Laura Marchal-Crespo and David J Reinkensmeyer. Review of control strategies for robotic movement training after neurologic injury. *J. Neuroeng. Rehabil.*, 6(1):1–15, 2009.

Franck Mars, Mathieu Deroo, and Jean-Michel Hoc. Analysis of human-machine cooperation when driving with different degrees of haptic shared control. *IEEE Trans. Haptics*, 7(3): 324–333, July 2014.

Julio C Mateo, Javier San Agustin, and John Paulin Hansen. Gaze beats mouse: hands-free selection by combining gaze and emg. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '08, pages 3039–3044, New York, NY, USA, April 2008. ACM.

Yoshio Matsumotot, Tomoyuki Ino, and Tsukasa Ogsawara. Development of intelligent wheelchair system with face and gaze based interface. In *Proceedings 10th IEEE International Workshop on Robot and Human Interactive Communication. ROMAN 2001 (Cat. No.01TH8591)*, pages 262–267. IEEE, 2001.

Jonathan Samir Matthis, Jacob L Yates, and Mary M Hayhoe. Gaze and the control of foot placement when walking in natural terrain. *Curr. Biol.*, 28(8):1224–1233.e5, April 2018.

Craig G McDonald, Jennifer L Sullivan, Troy A Dennis, and Marcia K O'Malley. A myoelectric control interface for upper-limb robotic rehabilitation following spinal cord injury. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 28(4):978–987, April 2020.

Jess McIntosh, Charlie McNeill, Mike Fraser, Frederic Kerber, Markus Löchtefeld, and Antonio Krüger. EMPress: Practical hand gesture classification with wrist-mounted EMG and pressure sensing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, page 2332–2342, San Jose California USA, May 2016. ACM.

D T McRuer and H R Jex. A review of quasi-linear pilot models. *IEEE Transactions on Human Factors in Electronics*, HFE-8(3):231–249, September 1967.

McRuer, Duane T and Jex, Henry R. A review of quasi-linear pilot models. *IEEE transactions on human factors in electronics*, 8(3):38–51, 1967.

Josh S Merel, Roy Fox, Tony Jebara, and Liam Paninski. A multi-agent control framework for co-adaptation in brain-computer interfaces. *Neural Inf Process Syst*, 26:2841–2849, December 2013.

Jenifer Miehlbradt, Alexandre Cherpillod, Stefano Mintchev, Martina Coscia, Fiorenzo Artoni, Dario Floreano, and Silvestro Micera. Data-driven body-machine interface for the accurate control of drones. *Proc. Natl. Acad. Sci. U. S. A.*, 115(31):7913–7918, July 2018.

Mohamed El Mistiri, Owais Khan, César A Martin, Eric Hekler, and Daniel E Rivera.

Data-driven mobile health: System identification and hybrid model predictive control to deliver personalized physical activity interventions. *IEEE Open J. Control Syst.*, 4:83–102, February 2025.

Inhyuk Moon, Myungjoon Lee, Jeicheong Ryu, and Museong Mun. Intelligent robotic wheelchair with EMG-, gesture-, and voice-based interfaces. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003) (Cat. No.03CH37453)*, volume 4, pages 3453–3458 vol.3. IEEE, 2003.

Luca Morelli, Simone Guadagni, Gregorio Di Franco, Matteo Palmeri, Giulio Di Candio, and Franco Mosca. Da vinci single site© surgical platform in clinical practice: a systematic review. *Int. J. Med. Robot.*, 12(4):724–734, December 2016.

Silvia Muceli and Dario Farina. Simultaneous and proportional estimation of hand kinematics from EMG during mirrored movements at multiple degrees-of-freedom. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 20(3):371–378, May 2012.

Emily M Mugler, Goran Tomic, Aparna Singh, Saad Hameed, Eric W Lindberg, Jon Gaide, Murad Alqadi, Elizabeth Robinson, Katherine Dalzotto, Camila Limoli, Tyler Jacobson, Jungwha Lee, and Marc W Slutzky. Myoelectric computer interface training for reducing co-activation and enhancing arm movement in chronic stroke survivors: A randomized trial. *Neurorehabil. Neural Repair*, 33(4):284–295, April 2019.

R R Murphy. Human–robot interaction in rescue robotics. *IEEE Trans. Syst. Man Cybern. C Appl. Rev.*, 34(2):138–153, May 2004.

Jan Saputra Müller, Carmen Vidaurre, Martijn Schreuder, Frank C Meinecke, Paul von Bünau, and Klaus-Robert Müller. A mathematical model for the two-learners problem. *J. Neural Eng.*, 14(3):036005, June 2017.

Gernot R Müller-Putz, Christian Breitwieser, Febo Cincotti, Robert Leeb, Martijn Schreuder, Francesco Leotta, Michele Tavella, Luigi Bianchi, Alex Kreilinger, Andrew Ramsay, Martin Rohm, Max Sagebaum, Luca Tonin, Christa Neuper, and José Del R Millán. Tools for brain-computer interaction: A general concept for a hybrid BCI. *Frontiers in neuroinformatics*, 5:30, November 2011.

I Nahum-Shani, E Hekler, and D Spruijt-Metz. Building health behavior models to guide the development of just-in-time adaptive interventions: A pragmatic framework. *Health Psychol.*, 34S:1209–1219, December 2015.

Inbal Nahum-Shani, Shawna N Smith, Bonnie J Spring, Linda M Collins, Katie Witkiewitz, Ambuj Tewari, and Susan A Murphy. Just-in-time adaptive interventions (JITAIs) in mobile health: Key components and design principles for ongoing health behavior support. *Ann. Behav. Med.*, 52(6):446–462, May 2018.

Yorie Nakahira, Quanying Liu, Terrence J Sejnowski, and John C Doyle. Diversity-enabled sweet spots in layered architectures and speed-accuracy trade-offs in sensorimotor control. *Proc. Natl. Acad. Sci. U. S. A.*, 118(22), June 2021.

J F Nash. Equilibrium points in N-Person games. *Proceedings of the National Academy of Sciences (PNAS)*, 36(1):48–49, January 1950.

J-Emeterio Navarro-Barrientos, Daniel E Rivera, and Linda M Collins. A dynamical model for describing behavioural interventions for weight loss and body composition change. *Math. Comput. Model. Dyn. Syst.*, 17(2):183–203, January 2011.

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.

W D Nothwang, M J McCourt, R M Robinson, S A Burden, and J W Curtis. The human should be part of the control loop? In *IEEE Resilience Week (RWS)*, pages 214–220, 2016.

Domen Novak, Ximena Omlin, Rebecca Leins-Hess, and Robert Riener. Effectiveness of different sensing modalities in predicting targets of reaching movements. *Conf. Proc. IEEE Eng. Med. Biol. Soc.*, 2013:4255–4258, 2013.

Emily R Oby, Matthew D Golub, Jay A Hennig, Alan D Degenhart, Elizabeth C Tyler-Kabara, Byron M Yu, Steven M Chase, and Aaron P Batista. New neural activity patterns emerge with long-term learning. *Proc. Natl. Acad. Sci. U. S. A.*, 116(30):15210–15215, July 2019.

Katsuhiko Ogata. *Modern Control Engineering*. Pearson, Upper Saddle River, NJ, 5 edition, August 2009.

G Onose, C Grozea, A Anghelescu, C Daia, C J Sinescu, A V Ciurea, T Spircu, A Mirea, I Andone, A Spânu, C Popescu, A-S Mihăescu, S Fazli, M Danóczy, and F Popescu. On the feasibility of using motor imagery EEG-based brain-computer interface in chronic tetraplegics for assistive robotic arm control: a clinical test and long-term post-trial follow-up. *Spinal Cord*, 50(8):599–608, August 2012.

Amy L Orsborn and Jose M Carmena. Creating new functional circuits for action via brain-machine interfaces. *Front. Comput. Neurosci.*, 7:157, November 2013.

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 Trans. Neural Syst. Rehabil. Eng.*, 20(4):468–477, July 2012.

Amy L Orsborn, Helene G Moorman, Simon A Overduin, Maryam M Shanechi, Dragan F Dimitrov, and Jose M Carmena. Closed-loop decoder adaptation shapes neural plasticity for skillful neuroprosthetic control. *Neuron*, 82(6):1380–1393, June 2014.

Yi Ouyang, Hamidreza Tavafoghi, and Demosthenis Teneketzis. Dynamic games with asymmetric information: Common information based perfect bayesian equilibria and sequential decomposition. *IEEE Trans. Automat. Contr.*, 62(1):222–237, January 2017.

S J Page, D R Gater, and P Bach-y Rita. Reconsidering the motor recovery plateau in stroke rehabilitation. *Archives of physical medicine and rehabilitation*, 85(8):1377–1381, 2004.

Yun Suen Pai, Tilman Dingler, and Kai Kunze. Assessing hands-free interactions for VR using eye gaze and electromyography. *Virtual Real.*, 23(2):119–131, June 2019.

Lizhi Pan, Dustin L Crouch, and He Helen Huang. Myoelectric control based on a generic musculoskeletal model: Towards a multi-user neural-machine interface. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 26(7):1435–1442, May 2018.

Chethan Pandarinath, Paul Nuyujukian, Christine H Blabe, Brittany L Sorice, Jad Saab, Francis R Willett, Leigh R Hochberg, Krishna V Shenoy, and Jaimie M Henderson. High performance communication by people with paralysis using an intracortical brain-computer interface. *Elife*, 6:e18554, February 2017.

M Pantic and L J M Rothkrantz. Toward an affect-sensitive multimodal human-computer interaction. *Proceedings of the IEEE*, 91(9):1370–1390, September 2003.

Ali Utku Pehlivan, Dylan P Losey, and Marcia K O'Malley. Minimal assist-as-needed controller for upper limb robotic rehabilitation. *IEEE Trans. Robot.*, 32(1):113–124, February 2016.

Serafeim Perdikis and Jose del R Millan. Brain-Machine interfaces: A tale of two learners. *IEEE Syst. Man Cybern. Mag.*, 6(3):12–19, July 2020.

Juan C Perez-Ibarra, Adriano A G Siqueira, Marcela A Silva-Couto, Thiago L de Russo, and Hermano I Krebs. Adaptive impedance control applied to robot-aided neuro-rehabilitation of the ankle. *IEEE Robot. Autom. Lett.*, 4(2):185–192, April 2019.

Robert J Peterka. Sensory integration for human balance control. *Handbook of clinical neurology*, 159:27–42, 2018.

Lauren N Peterson, Amber H Y Chou, Samuel A Burden, and Momona Yamagami. Assessing human feedback parameters for disturbance-rejection. *IFAC-PapersOnLine*, 55(41):1–6, January 2022.

Leilaalsadat Pezeshki, Hamid Sadeghian, Mehdi Keshmiri, Xiao Chen, and Sami Haddadin. Cooperative assist-as-needed control for robotic rehabilitation: A two-player game approach. *IEEE Robot. Autom. Lett.*, 8(5):2852–2859, May 2023.

Guido G Peña, Leonardo J Consoni, Wilian M dos Santos, and Adriano A G Siqueira. Feasibility of an optimal EMG-driven adaptive impedance control applied to an active knee orthosis. *Rob. Auton. Syst.*, 112:98–108, February 2019.

Annika Pfister, Maneeshika M Madduri, Amber H Y Chou, and Sam A Burden. Matching user and machine learning rates in co-adaptive closed-loop myoelectric interfaces. IEEE Conference on Neural Engineering and Rehabilitation, April 2023.

Ivan Phelan, Madelynne Arden, Maria Matsangidou, Alicia Carrion-Plaza, and Shirley Lindley. Designing a virtual reality myoelectric prosthesis training system for amputees. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, number Article 49 in CHI EA '21, pages 1–7, New York, NY, USA, May 2021. ACM.

Charles L Phillips, John M Parr, and Eve A Riskin. *Signals, systems, and transforms.* Prentice Hall, Old Tappan, NJ, 2003.

Patrick M Pilarski, Michael R Dawson, Thomas Degris, Farbod Fahimi, Jason P Carey, and Richard S Sutton. Online human training of a myoelectric prosthesis controller via actor-critic reinforcement learning. *IEEE Int. Conf. Rehabil. Robot.*, 2011:5975338, 2011.

Rajiv Ranganathan, Jon Wieser, Kristine M Mosier, Ferdinando A Mussa-Ivaldi, and Robert A Scheidt. Learning redundant motor tasks with and without overlapping dimensions: facilitation and interference effects. *Journal of Neuroscience*, 34(24):8289–8299, June 2014.

Linda Resnik, Matthew Borgia, and Melissa Clark. Function and quality of life of unilateral major upper limb amputees: Effect of prosthesis use and type. *Arch. Phys. Med. Rehabil.*, 101(8):1396–1406, August 2020.

Mattia Rigotti-Thompson, Samuel R Nason-Tomaszewski, Payton Bechefsky, Alexander Acosta, Nick Hahn, Donald Avansino, Brice Richards, Claire Nicolas, Yahia H Ali, Jaimie M Henderson, Leigh R Hochberg, Nicholas AuYong, and Chethan Pandarinath. Preparatory encoding of diverse features of intended movement in the human motor cortex. *bioRxivorg*, September 2025.

F Rizzoglio, M Casadio, Santis D De, and F A Mussa-Ivaldi. Building an adaptive interface via unsupervised tracking of latent manifolds. *Neural Networks*, 137:174–187, 2021.

Fabio Rizzoglio, Francesca Sciandra, Elisa Galofaro, Luca Losio, Elisabetta Quinland, Clara Leoncini, Antonino Massone, F A Mussa-Ivaldi, and Maura Casadio. A myoelectric computer interface for reducing abnormal muscle activations after spinal cord injury. *IEEE Int. Conf. Rehabil. Robot.*, 2019:1049–1054, June 2019.

Fabio Rizzoglio, Camilla Pierella, Dalia De Santis, Ferdinando Mussa-Ivaldi, and Maura Casadio. A hybrid body-machine interface integrating signals from muscles and motions. *J. Neural Eng.*, 17(4):046004, July 2020.

R Rosenberg. The biofeedback pointer: EMG control of a two dimensional pointer. In *Digest of Papers. Second International Symposium on Wearable Computers (Cat. No.98EX215)*, pages 162–163. IEEE, 1998.

Eatai Roth, Katie Zhuang, Sarah A Stamper, Eric S Fortune, and Noah J Cowan. Stimulus predictability mediates a switch in locomotor smooth pursuit performance for eigenmannia virescens. *J. Exp. Biol.*, 214(Pt 7):1170–1180, April 2011.

Eatai Roth, Robert W Hall, Thomas L Daniel, and Simon Sponberg. Integration of parallel mechanosensory and visual pathways resolved through sensory conflict. *Proceedings of the National Academy of Sciences*, October 2016.

Uta Sailer, J Randall Flanagan, and Roland S Johansson. Eye–Hand coordination during learning of a novel visuomotor task. *J. Neurosci.*, 25(39):8833–8842, September 2005.

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*, CHI '08, pages 515–524, New York, NY, USA, April 2008. Association for Computing Machinery.

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*, UIST '09, pages 167–176, New York, NY, USA, October 2009. Association for Computing Machinery.

R A Schmidt, T D Lee, C Winstein, G Wulf, and H N Zelaznik. Motor control and learning: A behavioral emphasis. 2018.

J P Scholz and G Schöner. The uncontrolled manifold concept: identifying control variables for a functional task. *Experimental brain research*, 126(3):289–306, 1999.

Julia Schwarz, Charles Claudius Marais, Tommer Leyvand, Scott E Hudson, and Jennifer Mankoff. Combining body pose, gaze, and gesture to determine intention to interact in vision-based interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 3443–3452, New York, NY, USA, April 2014. Association for Computing Machinery.

S Sefati, I D Neveln, E Roth, T R T Mitchell, J B Snyder, M A MacIver, E S Fortune, and N J Cowan. Mutually opposing forces during locomotion can eliminate the tradeoff between maneuverability and stability. *Proceedings of the National Academy of Sciences*, November 2013.

Jessica C Selinger, Jeremy D Wong, Surabhi N Simha, and J Maxwell Donelan. How humans initiate energy optimization and converge on their optimal gaits. *J. Exp. Biol.*, 222(Pt 19), October 2019.

Maryam M Shanechi, Amy L Orsborn, and Jose M Carmena. Robust brain-machine interface design using optimal feedback control modeling and adaptive point process filtering. *PLoS Comput. Biol.*, 12(4):e1004730, April 2016.

A J S Sheffler, S A S Mousavi, E Hellström, M Jankovic, M A Santillo, T M Seigler, and J B

Hoagg. Effects of reference-command preview as humans learn to control dynamic systems. In *2019 IEEE International Conference on Systems, Man and Cybernetics (SMC)*, pages 4199–4204, October 2019.

Krishna V Shenoy and Jose M Carmena. Combining decoder design and neural adaptation in brain-machine interfaces. *Neuron*, 84(4):665–680, November 2014.

Daniel B Silversmith, Reza Abiri, Nicholas F Hardy, Nikhilesh Natraj, Adelyn Tu-Chan, Edward F Chang, and Karunesh Ganguly. Plug-and-play control of a brain–computer interface through neural map stabilization. *Nat. Biotechnol.*, 39(3):326–335, September 2020.

Vasileios Skaramagkas, Giorgos Giannakakis, Emmanouil Ktistakis, Dimitris Manousos, Ioannis Karatzanis, Nikolaos Tachos, Evanthia Tripoliti, Kostas Marias, Dimitrios I Fotiadis, and Manolis Tsiknakis. Review of eye tracking metrics involved in emotional and cognitive processes. *IEEE Rev. Biomed. Eng.*, 16:260–277, January 2023.

Patrick Slade, Mykel J Kochenderfer, Scott L Delp, and Steven H Collins. Personalizing exoskeleton assistance while walking in the real world. *Nature*, 610(7931):277–282, October 2022.

Patrick Slade, Christopher Atkeson, J Maxwell Donelan, Han Houdijk, Kimberly A Ingraham, Myunghee Kim, Kyoungchul Kong, Katherine L Poggensee, Robert Riener, Martin Steinert, Juanjuan Zhang, and Steven H Collins. On human-in-the-loop optimization of human-robot interaction. *Nature*, 633(8031):779–788, September 2024.

Lauren H Smith, Todd A Kuiken, and Levi J Hargrove. Evaluation of linear regression simultaneous myoelectric control using intramuscular EMG. *IEEE Trans. Biomed. Eng.*, 63(4):737–746, April 2016.

Hyungeun Song, Tsung-Han Hsieh, Seong Ho Yeon, Tony Shu, Michael Nawrot, Christian F Landis, Gabriel N Friedman, Erica A Israel, Samantha Gutierrez-Arango, Matthew J Carty, Lisa E Freed, and Hugh M Herr. Continuous neural control of a bionic limb restores biomimetic gait after amputation. *Nat. Med.*, 30(7):2010–2019, July 2024.

Rong Song and Kai Yu Tong. EMG and kinematic analysis of sensorimotor control for patients after stroke using cyclic voluntary movement with visual feedback. *J. Neuroeng. Rehabil.*, 10(1):18, February 2013.

Seungmoon Song and Steven H Collins. Optimizing exoskeleton assistance for faster self-selected walking. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 29:786–795, May 2021.

Janis Stolzenwald and Walterio W Mayol-Cuevas. Rebellion and obedience: The effects of intention prediction in cooperative handheld robots. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3012–3019. IEEE, November 2019.

David Sussillo, Sergey D Stavisky, Jonathan C Kao, Stephen I Ryu, and Krishna V Shenoy. Making brain–machine interfaces robust to future neural variability. *Nat. Commun.*, 7(1): 1–13, December 2016.

Martin Tall, Alexandre Alapetite, Javier San Agustin, Henrik H T Skovsgaard, John Paulin Hansen, Dan Witzner Hansen, and Emilie Møllenbach. Gaze-controlled driving. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '09, pages 4387–4392, New York, NY, USA, April 2009. Association for Computing Machinery.

Thijs Tankink, Juha M Hijmans, Raffaella Carloni, and Han Houdijk. Human-in-the-loop optimization of rocker shoes via different cost functions during walking. *J. Biomech.*, 166 (112028):112028, March 2024.

Dawn M Taylor, Stephen I Helms Tillery, and Andrew B Schwartz. Direct cortical control of 3D neuroprosthetic devices. *Science*, 296(5574):1829–1832, June 2002.

Jordan A Taylor, John W Krakauer, and Richard B Ivry. Explicit and implicit contributions to learning in a sensorimotor adaptation task. *Journal of Neuroscience*, 34(8):3023–3032, 2014.

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

Emanuel Todorov and Michael I Jordan. Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.*, 5(11):1226–1235, November 2002.

Maegan Tucker, Ellen Novoseller, Claudia Kann, Yanan Sui, Yisong Yue, Joel W Burdick, and Aaron D Ames. Preference-based learning for exoskeleton gait optimization. *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020.

A Tversky and D Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, September 1974.

Kasper van der El, Daan M Pool, Herman J Damveld, Marinus Rene M van Paassen, and Max Mulder. An empirical human controller model for preview tracking tasks. *IEEE Trans Cybern*, 46(11):2609–2621, November 2016.

Deepanshu Vasal and Achilleas Anastasopoulos. Signaling equilibria for dynamic LQG games with asymmetric information. *IEEE Trans. Control Netw. Syst.*, 8(3):1177–1188, September 2021.

Xin Wang, Yudong Chen, and Wenwu Zhu. A survey on curriculum learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(9):4555–4576, September 2022.

Rahul B Warrier and Santosh Devasia. Iterative learning from novice human demonstrations for output tracking. *IEEE Trans. Hum. Mach. Syst.*, 46(4):510–521, 2016.

Richard J Wasicko, Duane T McRuer, and Raymond E Magdaleno. Human pilot dynamic response in single-loop systems with compensatory and pursuit displays. Technical report, SYSTEMS TECHNOLOGY INC HAWTHORNE CA, 1966.

Francis R Willett, Brian A Murphy, Daniel R Young, William D Memberg, Christine H Blabe, Chethan Pandarinath, Brian Franco, Jad Saab, Benjamin L Walter, Jennifer A Sweet, Jonathan P Miller, Jaimie M Henderson, Krishna V Shenoy, John D Simeral, Beata Jarosiewicz, Leigh R Hochberg, Robert F Kirsch, and Abidemi Bolu Ajiboye. A comparison of intention estimation methods for decoder calibration in intracortical brain-computer interfaces. *IEEE Trans. Biomed. Eng.*, 65(9):2066–2078, September 2018.

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, May 2021.

Francis R Willett, Erin M Kunz, Chaofei Fan, Donald T Avansino, Guy H Wilson, Eun Young Choi, Foram Kamdar, Matthew F Glasser, Leigh R Hochberg, Shaul Druckmann, Krishna V Shenoy, and Jaimie M Henderson. A high-performance speech neuroprosthesis. *Nature*, 620(7976):1031–1036, August 2023.

H S Witsenhausen. A counterexample in stochastic optimum control. *SIAM Journal on Control and Optimization*, 6(1):131–147, February 1968.

Jacob O Wobbrock, Meredith Ringel Morris, and Andrew D Wilson. User-defined gestures for surface computing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '09, pages 1083–1092, New York, NY, USA, April 2009. Association for Computing Machinery.

Jonathan R Wolpaw. Brain-computer interfaces as new brain output pathways: BCIs as new brain outputs. *J. Physiol.*, 579(Pt 3):613–619, March 2007.

Richard B Woodward and Levi J Hargrove. Adapting myoelectric control in real-time using a virtual environment. *J. Neuroeng. Rehabil.*, 16(1):11, January 2019.

Jian Wu, Lu Sun, and Roozbeh Jafari. A wearable system for recognizing american sign language in real-time using IMU and surface EMG sensors. *IEEE journal of biomedical and health informatics*, 20(5):1281–1290, September 2016.

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(34), pages 61–66. Elsevier, 2018a.

Momona Yamagami, Keshia M Peters, Ivana Milovanovic, Irene Kuang, Zeyu Yang, Nanshu Lu, and Katherine M Steele. Assessment of dry epidermal electrodes for long-term electromyography measurements. *Sensors*, 18(4), April 2018b.

Momona Yamagami, Katherine M Steele, and Samuel A Burden. Decoding intent with control theory: Comparing muscle versus manual interface performance. *Proc. SIGCHI Conf. Hum. Factor. Comput. Syst.*, 2020:1–12, April 2020.

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. *bioRxiv*, page 2020.08.01.232454, March 2021a.

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*, September 2021b.

Momona Yamagami, Maneeshika M Madduri, Benjamin J Chasnov, Amber H Y Chou, Lauren N Peterson, and Samuel A Burden. Co-adaptation improves performance in a dynamic human-machine interface. *bioRxiv*, page 2023.07.14.549053, July 2023a.

Momona Yamagami, Alexandra A Portnova-Fahreeva, Junhan Kong, Jacob O Wobbrock, and Jennifer Mankoff. How do people with limited movement personalize upper-body gestures? considerations for the design of personalized and accessible gesture interfaces. In *Proceedings of the 25th International ACM SIGACCESS Conference on Computers and Accessibility*, number Article 1 in ASSETS '23, New York, NY, USA, October 2023b. ACM.

Jehan Yang, Kent Shibata, Douglas Weber, and Zackory Erickson. High-density electromyography for effective gesture-based control of physically assistive mobile manipulators. *Npj Robot.*, 3(1):2, January 2025.

Yoonsik Yang, Seungho Chae, Jinwook Shim, and Tack-Don Han. EMG sensor-based two-hand smart watch interaction. In *Adjunct Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology*, UIST '15 Adjunct, pages 73–74, New York, NY, USA, November 2015. Association for Computing Machinery.

Meiqi Yao, Jinhua Chen, Jiyong Jing, Han Sheng, Xing Tan, and Jingfen Jin. Defining the rehabilitation adherence curve and adherence phases of stroke patients: an observational study. *Patient Prefer. Adherence*, 11:1435–1441, August 2017.

Shaowei Yao, Yu Zhuang, Zhijun Li, and Rong Song. Adaptive admittance control for an ankle exoskeleton using an EMG-driven musculoskeletal model. *Front. Neurorobot.*, 12:16, April 2018.

Aaron J Young, Levi J Hargrove, and Todd A Kuiken. The effects of electrode size and orientation on the sensitivity of myoelectric pattern recognition systems to electrode shift. *IEEE Trans. Biomed. Eng.*, 58(9):2537–2544, September 2011.

Hayder A Yousif, Ammar Zakaria, Norasmadi Abdul Rahim, Ahmad Faizal Bin Salleh, Mustafa Mahmood, Khudhur A Alfarhan, Latifah Munirah Kamarudin, Syed Muhammad Mamduh, Ali Majid Hasan, and Moaid K Hussain. Assessment of muscles fatigue based on surface EMG signals using machine learning and statistical approaches: A review. *IOP Conf. Ser. Mater. Sci. Eng.*, 705(1):012010, November 2019.

B Yu, R B Gillespie, J S Freudenberg, and J A Cook. Human control strategies in pursuit tracking with a disturbance input. In *53rd IEEE Conference on Decision and Control*, pages 3795–3800. ieeexplore.ieee.org, December 2014a.

Mingxin Yu, Yingzi Lin, David Schmidt, Xiangzhou Wang, and Yu Wang. Human-robot interaction based on gaze gestures for the drone teleoperation. *J. Eye Mov. Res.*, 7(4), September 2014b.

Ajmal Zemmar, Andres M Lozano, and Bradley J Nelson. The rise of robots in surgical environments during COVID-19. *Nat. Mach. Intell.*, 2(10):566–572, 2020.

Jinhua Zhang, Baozeng Wang, Cheng Zhang, Yanqing Xiao, and Michael Yu Wang. An EEG/EMG/EOG-based multimodal human-machine interface to real-time control of a soft robot hand. *Front. Neurorobot.*, 13:7, March 2019.

Juanjuan Zhang, Pieter Fiers, Kirby A Witte, Rachel W Jackson, Katherine L Poggensee, Christopher G Atkeson, and Steven H Collins. Human-in-the-loop optimization of exoskeleton assistance during walking. *Science*, 356(6344):1280–1284, June 2017a.

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 2020a.

Xiaoyi Zhang, Harish Kulkarni, and Meredith Ringel Morris. Smartphone-based gaze gesture communication for people with motor disabilities. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, CHI '17, pages 2878–2889, New York, NY, USA, May 2017b. ACM.

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 Trans. Cybern.*, 48(2):543–555, February 2018.

Xingye Zhang, T Michael Seigler, and Jesse B Hoagg. The impact of nonminimum-phase zeros on human-in-the-loop control systems. *IEEE Trans Cybern*, PP, November 2020b.

Xuan Zhang, Xu Zhang, Le Wu, Chang Li, Xiang Chen, and Xun Chen. Domain adaptation with self-guided adaptive sampling strategy: Feature alignment for cross-user myoelectric pattern recognition. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 30:1374–1383, 2022.

Yuling Zhang, Tong Li, Haoran Tao, Fengchen Liu, Bingshan Hu, Minghui Wu, and Hongliu Yu. Research on adaptive impedance control technology of upper limb rehabilitation robot based on impedance parameter prediction. *Front. Bioeng. Biotechnol.*, 11:1332689, 2023.

Maozheng Zhao, Wenzhe Cui, I V Ramakrishnan, Shumin Zhai, and Xiaojun Bi. Voice and touch based error-tolerant multimodal text editing and correction for smartphones. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, UIST '21, pages 162–178, New York, NY, USA, October 2021. ACM.

Mingde Zheng, Michael S Crouch, and Michael S Eggleston. Surface electromyography as a natural Human–Machine interface: A review. *IEEE Sens. J.*, 22(10):9198–9214, May 2022.

Kemin Zhou, John Comstock Doyle, and Keith Glover. *Robust and optimal control*, volume 40. Prentice Hall, 1996.

Katie Z Zhuang, Nicolas Sommer, Vincent Mendez, Saurav Aryan, Emanuele Formento, Edoardo D'Anna, Fiorenzo Artoni, Francesco Petrini, Giuseppe Granata, and Giovanni Cannaviello. Shared human–robot proportional control of a dexterous myoelectric prosthesis. *Nat. Mach. Intell.*, 1(9):400–411, 2019a.

Katie Z Zhuang, Nicolas Sommer, Vincent Mendez, Saurav Aryan, Emanuele Formento, Edoardo D'Anna, Fiorenzo Artoni, Francesco Petrini, Giuseppe Granata, Giovanni Cannaviello, Wassim Raffoul, Aude Billard, and Silvestro Micera. Shared human–robot proportional control of a dexterous myoelectric prosthesis. *Nature Machine Intelligence*, 1(9): 400–411, September 2019b.

Muhammad Zia Ur Rehman, Asim Waris, Syed Omer Gilani, Mads Jochumsen, Imran Khan Niazi, Mohsin Jamil, Dario Farina, and Ernest Nlandu Kamavuako. Multiday EMG-based classification of hand motions with deep learning techniques. *Sensors*, 18(8), August 2018.

K J Åström and R Murray. Feedback systems: an introduction for scientists and engineers. 2021.