

The Pennsylvania State University  
The J. Jeffrey and Ann Marie Graduate School

**HUMAN PERCEPTION AND INFORMATION ENCODING IN  
WEARABLE HAPTIC FEEDBACK SYSTEMS**

A Dissertation in  
Mechanical Engineering  
by  
Zachary Logan

© 2026 Zachary Logan

Submitted in Partial Fulfillment  
of the Requirements  
for the Degree of

Doctor of Philosophy

May 2026

The dissertation of Zachary Logan was reviewed and approved by the following:

Katie Fitzsimons  
Professor of Mechanical Engineering  
Dissertation Advisor  
Chair of Committee

Xiaogang Hu  
Professor of Mechanical Engineering

Anne Martin  
Professor of Mechanical Engineering

Alan Wagner  
Professor of Aerospace Engineering

Robert Kunz  
Director of Graduate Studies



# Abstract

Haptic wearables have the potential to enhance user experiences, improve the realism of AR/VR, reduce collisions, and augment feedback for motor learning by providing an additional channel to convey information. However, the success of such systems relies on human perceptual capabilities, the strategies used to encode tactile cues, and environmental perturbations. This thesis investigates the interplay between human perception and information encoding to optimize the performance of wearable haptic systems in terms of accuracy, intuitiveness, and reliability.

I first examine how voluntary motion affects multi-sensory haptic perception. Contrary to expectations, tactile suppression did not have a substantial impact on performance; instead, stimulus location had the most significant impact on accuracy. Data-driven, sensitivity-adjusted perception models reduced the effect of stimulus location across a wide range of people, highlighting the importance of accounting for human sensory sensitivities. I then explore visuomotor adaptation with a myoelectrically controlled virtual elbow under gain-based perturbations. Participants were able to adapt to the perturbation; however, substitutionary haptic feedback did not improve performance, underscoring the need for feedback to align with users' sensorimotor expectations to encode information effectively.

Finally, I explored the use of reinforcement learning to obtain individualized haptic feedback policies in real time. Using a pre-trained neural network and proximal policy optimization, we found that the framework could adapt to each user and to environmental changes. Collectively, these studies demonstrate how wearable haptic systems can leverage human perception to efficiently encode information, adapt to individual differences, and support intuitive control.

# Table of Contents

|                                                                                                                                            |          |
|--------------------------------------------------------------------------------------------------------------------------------------------|----------|
| List of Figures                                                                                                                            | vii      |
| List of Tables                                                                                                                             | xii      |
| Acknowledgments                                                                                                                            | xiv      |
| Preface                                                                                                                                    | xv       |
| <b>Chapter 1</b>                                                                                                                           |          |
| <b>Introduction</b>                                                                                                                        | <b>1</b> |
| 1.0.1 Haptic Feedback During Motion . . . . .                                                                                              | 2        |
| 1.0.2 Wearable Haptic Devices . . . . .                                                                                                    | 3        |
| 1.0.3 Perception Models . . . . .                                                                                                          | 4        |
| 1.1 Main Contributions . . . . .                                                                                                           | 5        |
| <b>Chapter 2</b>                                                                                                                           |          |
| <b>Evaluating factors affecting the perception of multi-sensory vibration<br/>        and skin-squeeze cues during voluntary movement.</b> | <b>8</b> |
| 2.1 Research Questions . . . . .                                                                                                           | 9        |
| 2.2 Introduction . . . . .                                                                                                                 | 9        |
| 2.3 Methods . . . . .                                                                                                                      | 10       |
| 2.3.1 Haptic Platform . . . . .                                                                                                            | 10       |
| 2.3.2 Robotic Data Collection Platform . . . . .                                                                                           | 12       |
| 2.3.3 Visual Display . . . . .                                                                                                             | 12       |
| 2.3.4 Experimental Procedure . . . . .                                                                                                     | 13       |
| 2.3.5 Data Analysis . . . . .                                                                                                              | 14       |
| 2.3.6 Sensitivity-Adjusted Perceptual Models . . . . .                                                                                     | 15       |
| 2.4 Results . . . . .                                                                                                                      | 17       |
| 2.4.1 Sample Response . . . . .                                                                                                            | 17       |
| 2.4.2 Statistical Results . . . . .                                                                                                        | 17       |
| 2.5 Discussion . . . . .                                                                                                                   | 23       |
| 2.6 Conclusions . . . . .                                                                                                                  | 25       |

## Chapter 3

|                                                                            |           |
|----------------------------------------------------------------------------|-----------|
| <b>Substitutionary somatosensory haptic feedback in EMG-Driven control</b> | <b>26</b> |
| 3.1 Research Questions . . . . .                                           | 27        |
| 3.2 Introduction . . . . .                                                 | 27        |
| 3.3 Background . . . . .                                                   | 28        |
| 3.4 Methods . . . . .                                                      | 29        |
| 3.4.1 Virtual Reality Apparatus . . . . .                                  | 30        |
| 3.4.2 Representative Prosthetic Elbow . . . . .                            | 30        |
| 3.4.3 Somatosensory Feedback Platform . . . . .                            | 32        |
| 3.4.4 Motor Adaptation Protocol . . . . .                                  | 33        |
| 3.4.5 Isometric Calibration Protocol . . . . .                             | 34        |
| 3.4.6 Data Analysis . . . . .                                              | 36        |
| 3.5 Results . . . . .                                                      | 37        |
| 3.5.1 Sample Responses . . . . .                                           | 37        |
| 3.5.2 Statistical Results . . . . .                                        | 37        |
| 3.5.2.1 Average Velocity . . . . .                                         | 37        |
| 3.5.2.2 Distance Error . . . . .                                           | 39        |
| 3.6 Discussion . . . . .                                                   | 43        |
| 3.7 Conclusion . . . . .                                                   | 46        |

## Chapter 4

|                                                                                                      |           |
|------------------------------------------------------------------------------------------------------|-----------|
| <b>Using machine learning to generate an online haptic feedback policy.</b>                          | <b>47</b> |
| 4.1 Research Questions . . . . .                                                                     | 48        |
| 4.2 Introduction . . . . .                                                                           | 48        |
| 4.3 Methods . . . . .                                                                                | 50        |
| 4.3.1 Policy Pre-training . . . . .                                                                  | 50        |
| 4.3.2 Online Learning Framework . . . . .                                                            | 53        |
| 4.3.3 Experimental Platform . . . . .                                                                | 54        |
| 4.3.4 Experimental Protocol . . . . .                                                                | 55        |
| 4.3.5 Data Analysis . . . . .                                                                        | 56        |
| 4.4 Results . . . . .                                                                                | 58        |
| 4.4.1 Policy Convergence . . . . .                                                                   | 58        |
| 4.4.2 Objective and Subjective Policy Performance . . . . .                                          | 59        |
| 4.4.3 Comparison with Pre-trained Baseline Policy . . . . .                                          | 60        |
| 4.4.4 Comparison Between Participants' Post-personalization and Post-perturbation Policies . . . . . | 62        |
| 4.4.5 Comparison between participants . . . . .                                                      | 64        |
| 4.5 Discussion . . . . .                                                                             | 64        |
| 4.6 Conclusion . . . . .                                                                             | 68        |

## Chapter 5

|                                    |           |
|------------------------------------|-----------|
| <b>Conclusions and Future Work</b> | <b>70</b> |
|------------------------------------|-----------|

|                                                           |    |
|-----------------------------------------------------------|----|
| Appendix A                                                |    |
| Visuomotor Adaptation Pairwise T-test Comparison Tables   | 73 |
| Appendix B                                                |    |
| Average Jensen-Shannon Divergence comparison Tables       | 75 |
| Appendix C                                                |    |
| Proximal Policy Optimization Hyperparameters and Packages | 78 |
| Bibliography                                              | 82 |

# List of Figures

|     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |    |
|-----|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 2.1 | The haptic platform was comprised of eight vibrating mini-motors arranged in a circular array around the circumference of the forearm and one band of nylon webbing around the upper arm, whose circumference was controlled by a servo motor. The shortening or lengthening of the band produced a squeezing sensation around the bicep. The tactors and servo were actuated simultaneously to indicate the desired motion, with the tactors providing the target angle and the servo providing the target speed. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | 11 |
| 2.2 | This schematic shows how the set of unique trial conditions was set, duplicated 5 times such that each participant completed 5 trials per combination. The order of the list of duplicated trials was shuffled with a random number generator. For voluntary trials, half of the trials were randomly assigned to have the participant track a figure-8, and in the other half, participants tracked a circular trajectory. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 14 |
| 2.3 | The top plot shows the participant response during a static trial, where they would remain at rest until the presentation of the haptic stimuli. The bottom plot denotes the participant response during a voluntary motion trial, where they would be following a figure-8 trajectory as indicated on the visual display. The haptic stimulus was applied at $t = 0$ seconds, as indicated by the gray bar. In both trials the haptic stimulus indicated a direction of fifty degrees at the same speed. In the static response, the participant begins responding to the haptic cue around 0.5 seconds into the trial. The start of their response is clearly denoted by shift to a non-zero velocity. In the voluntary motion response, the participant begins responding to the haptic cue at around 0.3 seconds. The start of the voluntary motion response can be visualized by the point where the y-position diverges from the z-position, indicating a divergence from the figure-8 trajectory. . . . . | 18 |

|     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |    |
|-----|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 2.4 | The top plots show the position and velocity of a typical response during a static trial in the YZ-Plane, while the bottom plots show the typical response during a voluntary motion trial. The black arrow in the plots indicates the haptically cued direction. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | 19 |
| 2.5 | <b>Angle error vs. Target direction:</b> Participants were able to more accurately interpret and execute target angles that were in the direction of up and to the right. The accuracy of participants interpretation and execution of target angles was not affected by their state of motion. . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | 20 |
| 2.6 | <b>Speed responses vs Target speed:</b> Participants were able to differentiate the different levels of squeeze and increased the speed of their movement when the intensity of the squeeze increased. However, participant responses tended to not closely match the true desired speed value and would overshoot it. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | 21 |
| 2.7 | <b>Speed responses vs Target angle:</b> Generally, as the speed intensity increased, participants increased their speed. For target angles between 255 to 310 degrees, participants' speed responses were generally slower than the other directions. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 22 |
| 3.1 | Participants were sat in the Kinreach virtual reality and motion tracking system, where a visual display was presented on a glass mirror from an inverted HD TV (Top Left). Participants performed natural reaching motions using elbow flexion and extension or virtual reaching using isometric contractions to control EMG-driven elbow (Bottom Right). Participants controlled the EMG-driven elbow using isometric contractions. The EMG signal was passed through an analog amplification and rectification circuit before reaching the Arduino microcontroller used to modulate the angular position of the servo. The Participants will perform 80 baseline reaches, 80 reaching movements with a visual gain perturbation and 80 reaching movements with the visual gain perturbation turned off. . . . . | 31 |
| 3.2 | a) <b>Natural reaching task:</b> The participant's upper arm was secured in a brace to restrict shoulder motion and limit movement to only elbow flexion and extension. The forearm was placed in a splint to immobilize the wrist and fingers and rested on a low-friction air sled to minimize the effect of friction. b) <b>Isometric virtual reaching task:</b> Both the upper arm and forearm were secured using custom braces, entirely restricting limb movement and enforcing isometric contractions. . . . .                                                                                                                                                                                                                                                                                              | 32 |

|      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |    |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 3.3  | The Snaptics platform is a low-cost, open-source wearable haptics platform consisting of three feedback modules: vibration, twist, and stretch. Although the vibration and stretch modules are depicted, they remained inactive during the experiment. Only the twist module was used to provide haptic feedback. The angular position of the twist module was mapped in real-time to the virtual elbow's joint angle. . . . .                                                                                                                            | 33 |
| 3.4  | The Kinereach virtual reality system was used to display targets to the participants. The green dot represents the start position. The gray circle with a central blue dot indicates the target. The circular icon with a crosshair represents the cursor. . . . .                                                                                                                                                                                                                                                                                        | 34 |
| 3.5  | Sample trajectories showing the path traveled by the cursor for a participant from the Control group, the Isometric (no feedback) group, and the Isometric + Haptic Feedback group, during the Late Baseline, Early Adaptation, and Early Washout Epochs . . . . .                                                                                                                                                                                                                                                                                        | 38 |
| 3.6  | Sample EMG response from a Participant in the Virtual Reaching Isometric (no feedback) group during Late Baseline, Early Adaptation, and Early Washout. . . . .                                                                                                                                                                                                                                                                                                                                                                                           | 38 |
| 3.7  | Sample EMG response from a Participant in the Virtual Reaching Isometric + Haptic Feedback group during Late Baseline, Early Adaptation, and Early Washout. . . . .                                                                                                                                                                                                                                                                                                                                                                                       | 39 |
| 3.8  | Plot showing the Average Velocity (meters per second) averaged over a set of epochs containing ten trials. The shaded regions represent the standard error. The gain perturbation was introduced after epoch 8 and removed after epoch 16, indicated by the two vertical lines. Participants in each group performed the reaching motions at a similar average velocity throughout the experiment. Participants in each group also reduced their average velocity during the adaptation phase to account for the effect of the gain perturbation. . . . . | 39 |
| 3.9  | Plot of the Distance Error (meters) averaged over a set of epochs containing ten trials. The shaded regions represent the standard error. The gain perturbation was introduced after epoch eight and removed after epoch 16, indicated by the two vertical lines. . . . .                                                                                                                                                                                                                                                                                 | 40 |
| 3.10 | Initial Finalized Bar plot for EMG Results . . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | 43 |

|     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |    |
|-----|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 4.1 | Overview of the PPO-based training loop for the haptic feedback policy. At the beginning of each trial, the current policy receives a desired direction for the joystick and outputs motor commands corresponding to the vibration intensity. These commands are applied as haptic stimuli via the tactors to the participant’s arm. The participant interprets the haptic stimuli as instructions and moves the joystick toward the most intense vibration. Their response is recorded and compared to the desired direction to compute a reward value. This response, the original desired direction, and the current policy distribution are used to calculate the PPO loss and update the policy parameters. The process repeats across trials to refine the haptic feedback policy based on the user’s behavior. .                                                                                                                                                                        | 50 |
| 4.2 | The experimental platform consisted of a haptic armband, an analog joystick, and a host computer. The haptic platform featured four eccentric rotating mass (ERM) mini-motors arranged in a circular array around the participant’s forearm, aligned with the four cardinal directions (Up, Down, Left, and Right). These motors were controlled by an Arduino Uno connected to the host computer. The analog joystick was a two-axis joystick from Elegoo, which measured thumb movements via two potentiometers mounted on a gimbal mechanism. The joystick signals were read by an Arduino Mega and transmitted to the host computer. During each trial, the host computer sent vibration intensity commands to the Arduino Uno, prompting the motors to activate. Participants interpreted the resulting haptic cues as directional instructions and moved the joystick accordingly. The Arduino Mega then relayed the joystick position back to the host computer for evaluation. . . . . | 55 |
| 4.3 | Overall Average Jensen Shannon Divergence computed across all directions and actuators between each participant’s post-personalization policy and post-perturbation policy with the pre-trained policy . . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | 62 |
| 4.4 | Average Jensen-Shannon Divergence computed per actuator between each participant’s post-personalization policy and the pre-trained policy, computed per motor and averaged across all joystick directions. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 62 |
| 4.5 | Average Jensen-Shannon Divergence computed per actuator between each participant’s post-perturbation policy and the pre-trained policy, computed per motor and averaged across all joystick directions. . . . .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | 63 |



|     |                                                                                                                                                                                                                                                                                                                                          |    |
|-----|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 4.6 | Average Jensen–Shannon Divergence (JSD) between each pair of participants’ learned baseline policies across all four motors. Each cell in the matrix represents the overall JSD for the corresponding participant pair, with lighter colors indicating lower divergence and darker colors indicating higher divergence. . . . .          | 65 |
| 4.7 | Average Jensen–Shannon Divergence (JSD) between each pair of participants’ learned post-perturbation policies across all four motors. Each cell in the matrix represents the overall JSD for the corresponding participant pair, with lighter colors indicating lower divergence and darker colors indicating higher divergence. . . . . | 66 |

# List of Tables

- 2.1 This table contains the mean and standard deviation of the angle error for each model trained using the combined data from all participants. All the models have the capability of reducing the response error. The linear model trained using the combined static and voluntary motion data had the greatest reduction in the angle portion of the absolute decoding error. 21
  
- 2.2 This table contains the mean and standard deviation of the speed response for each model trained using the combined data from all participants. All the models have the capability of reducing the response error. The power model trained using the combined static and voluntary motion data had the greatest reduction in the speed portion of the absolute decoding error. 22
  
- 2.3 This table shows the resulting p-values for the non-parametric Wilcoxon-signed rank test comparing the absolute decoding error of the different models trained using the combined data from all participants. The absolute decoding errors  $\varepsilon$ , was denoted by the training data: S (static trials), V (voluntary motion trials), or C (both) and the model: L (Linear) or P (Power). There is a statistically significant difference in the absolute decoding error between the static linear and static power models, the static linear and dynamic power, the static power and dynamic power, the dynamic linear and dynamic power, the static linear and linear, the static linear and power, the static power and linear, the static power and power, the dynamic linear and linear, the dynamic linear and power, the dynamic power and linear, the dynamic power and power, and the linear and power models. . . . . 23

|     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |    |
|-----|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----|
| 4.1 | Absolute value of the final policy loss slopes computed over the last 10 trials for each participant in both the post-personalization and post-perturbation conditions. Slopes with absolute value below 0.005 indicate convergence. The loss rates for the policies of participants 1-4, and 6-11 were below the threshold, indicating successful convergence, and the loss rate for participant 5 was above the threshold indicating the policy had not fully converged. . . . . | 59 |
| 4.2 | Average Jensen-Shannon Divergence between each participant's post-personalization policy and the pre-trained policy, computed per motor and averaged across all joystick directions. . . . .                                                                                                                                                                                                                                                                                       | 61 |
| 4.3 | Average Jensen-Shannon Divergence between each participant's post-perturbation policy and the pre-trained policy, computed per motor and averaged across all joystick directions. . . . .                                                                                                                                                                                                                                                                                          | 61 |
| 4.4 | Average Jensen-Shannon Divergence (JSD) between each participant's post-perturbation policy and their post-personalization policy, computed per motor and averaged across all joystick directions. . . . .                                                                                                                                                                                                                                                                         | 63 |

# Acknowledgments

This material is based upon work supported by the U.S. National Science Foundation under Award No. 2319139.

Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author and do not necessarily reflect the views of the U.S. National Science Foundation.

# Preface

**Chapter 2** A version of this work has been published in the Transactions on Haptics Journal, Logan, Z., Q. Deitrick, and K. Fitzsimons (2025) “Evaluating factors affecting the perception of multi-sensory vibration and skin-squeeze cues during voluntary movement,” IEEE Transactions on Haptics, pp. 1–8, doi:10.1109/TOH.2025.3572862. My key roles were the design and development of the test apparatus, data collection, analysis, and manuscript writing, as the supervisory author, Dr. Katie Fitzsimons, was involved throughout the project from the initial formulation to the final manuscript edits. Quinn Deitrick supported data collection and analysis and provided feedback on manuscript drafts and revisions.

In reference to IEEE copyrighted material which is used with permission in this thesis, the IEEE does not endorse any of The Pennsylvania State University’s products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing IEEE copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to [http://www.ieee.org/publications\\_standards/publications/rights/rights\\_link.html](http://www.ieee.org/publications_standards/publications/rights/rights_link.html) to learn how to obtain a License from RightsLink. If applicable, University Microfilms and/or ProQuest Library, or the Archives of Canada may supply single copies of the dissertation.

**Chapter 3 Author Contributions** A version of this work is currently under revision. The project was conducted jointly with Quentin Anderson-Watson and Sophie Walsh under the supervision of Dr. Katie Fitzsimons in the Human Centered Robotics Laboratory at Penn State and with Dr. Nick Kitchen under the supervision of Dr. Robert Sainburg in the Movement Neuroscience Laboratory. I am the lead author of the paper, and my key contributions included developing the EMG controller and recording pipeline, designing the mount for the virtual elbow and the arm support brace, conducting the experiment, performing the data analysis, and writing and revising the manuscript. Dr. Nick Kitchen contributed to developing the physical reaching components of the task, assisted with data collection and analysis, and created several of the manuscript figures. As supervisory authors, Dr. Katie Fitzsimons and Dr. Robert Sainburg were involved from the initial conceptualization of the experiment through results validation and final manuscript editing.

**Chapter 4 Author Contributions** A version of this work has been submitted to the 2026 IEEE International Conference on Biomedical Robotics and Biomechatronics. My primary contributions included conceptualizing the experiment, validating the reinforcement learning algorithm, developing the whole training framework, implementing the model, conducting the experiment, performing the data analysis, and writing and revising the manuscript. Sophie Walsh assisted with conducting the experiment, contributed to the data analysis, and provided feedback on manuscript drafts. Quentin Anderson-Watson supported data collection and contributed to manuscript editing. As the supervisory author, Dr. Katie Fitzsimons was involved throughout the project from the initial formulation to the final manuscript revisions.

# Chapter 1 |

## Introduction

Wearable haptic interfaces provide an opportunity to communicate information from an autonomous system to a person without adding to the native auditory and visual stimuli of an environment. Even low-bandwidth vibrotactile feedback with a simple control scheme can enhance user experiences with consumer electronics [1]. With a more complex set of symbols, a haptic interface can provide information about spatial navigation [2], substitute for kinesthetic feedback [3], improve motor control [4], reduce powered wheelchair collisions [5], improve virtual medical training [6], and support drive-by-wire control platforms for automobiles [7].

Given that the communication bandwidth of wearable devices may be limited by presentation rate [8] or number of unique symbols [9], several groups have developed multi-sensory wearable devices capable of displaying more information more quickly via additional actuators and stimuli modalities [10–12]. The challenge to implementing these in the wild is that the dynamic nature of the environments and activities that individuals engage in may degrade their ability to correctly perceive and act upon the haptic symbols. Some studies have noted a reduction in perception during motion that could be due to a phenomenon called tactile suppression (proprioceptive interference) [13–16], tactile masking (stimuli interference) [17–20], cognitive load [21,22], intuitiveness of the cue [23], user preferences [2, 24–26], or some combination of these factors. Even with perfect perception, correctly acting on the haptic symbols displayed requires either a high amount of training to recognize the meaning of each cue or an intuitive mapping of haptic cues to desired actions [8].

The overarching goal of this research is to develop a framework for wearable haptic feedback devices that adapt to the user, rather than requiring the user to adapt to the device. Achieving this goal requires a deeper understanding of how sub-conscious interference (e.g. tactile suppression) and perceptual interference between stimuli (tactile

masking), affect haptic perception during dynamic, real world conditions. Building on prior studies, the first objective investigates how voluntary motion influences the perception of haptic cues, disentangling the role of stimulus location and movement-related suppression. The second objective examines the degree to which artificial haptic feedback can support adaptation to visuomotor perturbations in EMG-controlled systems that lack intrinsic proprioceptive feedback.

Finally the third objective investigates the potential of developing a personalized haptic feedback policy that can adapt to perturbations of the feedback device and environment using reinforcement learning. These studies will lay the groundwork for the development of intelligent and adaptive haptic feedback devices that can be leveraged to improve the performance of prosthetics, teleoperated systems, rehabilitation interfaces and virtual reality training.

### **1.0.1 Haptic Feedback During Motion**

Tactile suppression—the phenomenon where task-related motion reduces our ability to perceive haptic stimuli—poses a significant challenge to the design of effective haptic feedback devices and encodings. This phenomenon consistently emerges during a range of motor activities, including simple finger motions [27, 28], juggling [16], and back-bending actions [29]. During bi-manual reaching tasks prompted by visual cues, participants show a marked decline in their ability to detect vibratory stimuli on their index fingers, especially within the 0-50 millisecond window before movement onset [13]. Fast, multi-joint movements (e.g., 3-DOF arm trajectories) further impair haptic perception relative to slower, constrained motions [25].

Not all forms of haptic feedback experience suppression equally. The placement of haptic feedback on the body influences the presence in a speeded reach-and-grasp task [14]. Tactile suppression is less pronounced at limb locations that are anticipated to receive feedback as part of the task—for instance, the fingers during grasping [14]. When asked to report the presence and location of vibration after movement, participants showed improved detection when the stimuli occurred at functionally relevant sites. This trend becomes even more apparent as the hand nears a grasp target, where suppression on the index finger decreases substantially [15]. Further, when haptic cues were presented on fingers and wrists perception rates were poor [13, 14, 16]. While perception rates were high in studies where the feedback interface was on the dorsal and palmar sides of the hand [2], the shoulders [30], and thighs [31].

Tactile suppression also interferes with tasks beyond motor execution, affecting



perceptual tasks such as force estimation [32,33] and Braille character reading [34]. These collective insights underscore the importance for haptic interface designers to thoughtfully consider the timing, location, and contextual relevance of feedback signals. Without accounting for tactile suppression, devices risk delivering feedback that users may fail to perceive during dynamic interactions.

While tactile suppression is frequently observed at the onset of motion it is not clear whether this phenomenon is relevant to continuously guided, dynamic activities. Low to moderate levels of physical activity and sitting versus standing do not have any significant impact on participant’s ability to recognize the spatio-temporal patterns [22]. Therefore, tactile suppression may not be present, or its effects are highly diminished after a user has remained in motion for an extended duration beyond the onset of motion. When using haptic feedback for guided physical activity/exercise, tactile parameters of duration, vibration frequency, amplitude, and location are effective presentation forms [35]. Elderly participants in a study of haptic guided exercise missed no cues and had favorable impressions of the haptic instructions. Based on the responses from a UX questionnaire survey, they found it easy to learn and pleasant to use [35]. The major shortcoming of the haptic instructions was the delay in the system between detecting leg movement and supplying the feedback. Tactile feedback has also been successfully implemented into other dynamically complex activities such as snowboarding [30], ball-room dancing [31], and general spatial navigation [2] with high levels of perception and recognition of the haptic cues.

## **1.0.2 Wearable Haptic Devices**

Channel capacity refers to the maximum rate at which information can be transmitted through a given sensory modality [8]. In the context of haptic interfaces, this capacity is constrained by users’ perceptual resolution and reaction time. Empirical studies have shown that optimal cue presentation rates for vibrotactile feedback are approximately 2-3 cues per second and that information transfer is maximized by increasing the number of distinct, perceivable cues [36,37]. However, generating many distinct cues primarily through variation in spatial and temporal vibration patterns can strain perceptual limits, often requiring extensive user training and resulting in delays between stimulus presentation and user response [8]. For more complex feedback or informational tasks, vibrotactile interfaces are therefore limited by the user’s ability to perceive and accurately interpret these subtle variations.

To improve spatial resolution, some studies have employed illusory movement effects

by delivering vibrotactile stimuli at multiple locations in a controlled sequence [38]. This stimulation technique can enhance the perceptual spatial resolution of relatively sparse tactor arrays [39]. However, this improvement comes at the cost of significantly increased reaction times, with responses to saltatory effects typically ranging from 2.5-4.5 seconds, compared to 200-500 milliseconds for isolated vibrotactile cues [25, 40]. This delay may be further exacerbated by perceptual confusion for vibrotactile stimuli presented on a moving body part [25, 41]. One potential solution to increase the information bandwidth of wearable haptic devices without sacrificing reaction time and permeability is to incorporate other cutaneous feedback modalities, such as skin stretch and skin squeeze.

Wearable haptic devices that combine multiple feedback modalities are known as multisensory haptic devices. Multisensory haptic devices that combine vibration with skin-squeeze or skin stretch have been used to support therapy [42] and enhance the immersive experience of virtual reality [12]. This type of multisensory feedback increases the number of distinct cues that a person can accurately interpret from cutaneous feedback [10, 43]—making them well-suited to communicate more complex information than vibrotactile stimuli alone.

Though multisensory haptic feedback can improve the range and complexity of perceivable cues, simultaneous presentation of different tactile modalities can lead to perceptual interference, a phenomenon known as tactile masking [17–20, 44]. Tactile masking occurs when one type of stimulus reduces or blocks the perception of another, resulting in degraded performance in tasks that involve multisensory haptic feedback [10, 43, 45]. For instance, a study by Zook et al. [19] presented skin-squeeze cues and skin-stretch cues to a participant’s forearm and found that skin-squeeze cues were capable of masking the perception of skin-stretch cues. This result indicates that the combination of different simultaneous sensory stimuli methods plays an essential role in the performance of haptic feedback.

### 1.0.3 Perception Models

Tactile suppression can be measured as a percentage of correct responses to probe stimuli or captured by a sensitivity calculation [14], therefore we could view tactile suppression as a reduction in perceptual sensitivity. Perceptual-sensitivity plays an important role in the perceptual illusion known as phantom sensation. Phantom sensation is a vibrotactile illusion that occurs when two adjacent actuators are activated simultaneously, yet the user perceives the stimuli at a single location between the two tactors [46]. The position

of the phantom actuator can be shifted by modulating the relative intensities of the two source actuators. Vibrotactile devices will often utilize a perceptual model to control and predict where the user will feel the phantom sensation [40, 47]. These perception models are based around three mathematical models [46, 47]. Early iterations of all three models, the linear model, the log model and the power model, assume that perceptual sensitivity is constant across users and body locations. However, Cholewiak et al. [48] determined that perceptual sensitivity depends on the location of the stimulus, which is consistent with physiological differences in mechanoreceptors across the human body. Personalized models significantly increased the perception and correct interpretation of vibrotactile cues using phantom sensation when participants were at rest [49]. Then it may be possible to use data-driven perception models to reduce or even mitigate the accuracy reduction.

## 1.1 Main Contributions

In this thesis, we explored three critical aspects of wearable haptic feedback: (1) how human perception interacts with multisensory cues during voluntary motion, (2) how users adapt to perturbations when controlling a device through myoelectric input, and (3) how reinforcement learning can be used to generate adaptive, individualized haptic feedback policies.

**Chapter 2: Evaluating factors affecting the perception of multisensory vibration and skin-squeeze cues during voluntary movement.** Prior work shows that tactile suppression, perceptual interference, and multisensory complexity can degrade the accuracy of haptic perception—effects that become even more pronounced during voluntary movement. These constraints directly affect how reliably wearable systems can encode directional or task-relevant information. In the second chapter, I examine how voluntary motion alters the perception of combined vibrotactile and skin-squeeze cues by having participants perform controlled reaching movements in the direction indicated by vibrotactile feedback and at the speed indicated by skin-squeeze feedback. Participants reported the direction and speed by manipulating the Kuka iiwa LBR robotic arm, which acted as a three-degree-of-freedom joystick. I then evaluate whether data-driven sensitivity-adjusted perception models can better interpret these cues across users and capture the motion-related distortions observed in the study. **The contributions of this chapter are as follows:**

- I demonstrated that voluntary motion—and by extension, tactile suppression—is

not the dominant factor limiting perception of multisensory haptic cues.

- I showed that for voluntary motion, stimulus location was the dominant factor affecting multisensory devices.
- I implemented and evaluated data-driven, sensitivity-adjusted perception models that improved decoding of multisensory cues, demonstrating the potential for adaptive feedback strategies in wearable devices.

**Chapter 3: Substitutionary somatosensory haptic feedback in EMG-Driven control** Motor adaptation is essential for maintaining accurate control despite changes in the environment or the body. In assistive and myoelectric devices, users rely heavily on visual feedback, which is slower and cognitively demanding than natural proprioception. Substitutionary tactile feedback—such as vibrotactile or skin-stretch cues—can augment or replace missing sensory channels, potentially reducing cognitive load and improving control. Effective feedback design is especially critical for upper-limb prostheses, where lack of somatosensory information limits precision, increases error, and can worsen phantom limb pain.

In the third chapter, I examine how substitutionary tactile feedback affects visuomotor adaptation in an EMG-controlled virtual elbow. Participants performed reaching tasks under three conditions: natural reaching with their real arm, virtual reaching with a myoelectrically controlled virtual elbow without feedback, and virtual reaching with a myoelectrically controlled virtual elbow with haptic feedback provided via a wearable haptic twist module. To examine the adaptability of the haptic feedback, we implemented a visuomotor gain perturbation that systematically altered the relationship between their movements and the visual cursor.

**The contributions of this chapter are as follows:**

- I implemented a novel paradigm, adapted from visuomotor gain perturbations, for a myoelectrically controlled virtual elbow.
- I demonstrated that participants can adapt to a visuomotor gain perturbation while controlling a myoelectrically driven virtual elbow.
- I showed that haptic feedback must be carefully matched to the task’s structure to be beneficial.

**Chapter 4: Using machine learning to generate an online haptic feedback policy.** Haptic wearables can improve spatial awareness, reduce collisions, and support

skill development in virtual environments. Despite this potential, the literature is inconsistent about which feedback paradigms work best. As a result, many systems rely on ad-hoc, handcrafted mappings that are slow to design, brittle across users, and challenging to scale. Perception of vibrotactile cues varies widely with stimulus location, actuator interactions, and individual user sensitivity, making a one-size-fits-all approach unrealistic.

To address these limitations, we investigate whether reinforcement learning can generate a personalized feedback policy tailored to each user’s perceptual tendencies. Our first goal was to test feasibility: can an RL policy actually converge to a stable, individualized solution given noisy human responses? Our second goal was to test robustness. After establishing a baseline policy, we intentionally perturbed the haptic platform and examined whether the algorithm could adapt and converge to a new policy under altered sensory conditions. **The Contributions of this chapter are as follows:**

- I developed a reinforcement learning framework using a pre-trained neural network and proximal policy optimization to adapt haptic feedback policies online for joystick control.
- I demonstrated that the framework can obtain policies that converge to user-specific solutions and adapt to environmental perturbations.

## Chapter 2 | **Evaluating factors affecting the perception of multi-sensory vibration and skin-squeeze cues during voluntary movement.**

Wearable haptic interfaces increasingly employ multi-sensory feedback to convey complex information, such as the direction and magnitude of motion. While combining modalities such as vibration and skin-squeeze can increase the number of distinct cues, it also introduces additional perceptual challenges. Prior work suggests that phenomena such as tactile masking and tactile suppression—particularly during voluntary motion—may degrade a user’s ability to perceive and interpret these cues accurately. However, it remains unclear whether motion-induced effects are the dominant factor limiting perception in multi-sensory haptic feedback. In particular, spatial variation in tactile sensitivity across the body may play a more significant role, especially for wearable devices that distribute stimuli across the skin.

This chapter addresses this question by systematically evaluating how voluntary motion, stimulus location, and cue magnitude influence the perception of multi-sensory haptic feedback during a motion-guidance task. In addition, it investigates whether data-driven, sensitivity-adjusted perceptual models can compensate for spatial variability in perception and improve decoding accuracy across users and conditions.

## 2.1 Research Questions

The study is structured around three core research questions that guide the experimental design and interpretation of results:

- Does voluntary motion significantly affect the perception and execution of multi-sensory haptic cues?
- How do stimulus location and cue magnitude influence perceptual error relative to motion state?
- Can data-driven perceptual models reduce decoding error in multi-sensory haptic feedback?

## 2.2 Introduction

The effectiveness of wearable haptic feedback depends on the user’s ability to reliably perceive and interpret tactile cues during active interaction. When feedback is presented during voluntary motion [16, 27–29] or through multiple sensory modalities [17–20], perceptual performance can degrade, limiting the reliability of the conveyed information.

One proposed explanation for this degradation is tactile suppression, where sensitivity to externally applied stimuli is reduced during motion. This effect has been observed when presenting vibrotactile-only stimuli during simple finger motions [27, 28], juggling [16], and back-bending motions [29], and is believed to be a primary limitation when presenting haptic feedback during active control. Additionally, multi-sensory feedback introduces potential interference effects such as tactile masking, where one or more stimuli mask or overpower the sensations created by other stimuli [10, 17–20, 43, 45].

However, perceptual performance is also influenced by spatial sensitivity, which varies across different locations on the body. Prior work has shown that suppression effects can depend on task relevance and location [14, 15], while more general changes in activity level have minimal impact on perception [22]. These findings suggest that spatial variability and individual differences may play a greater role than motion state alone.

In this chapter, we evaluated whether tactile suppression had a significant impact on perception of multi-sensory cues during voluntary motion and found that our main hypothesis was not supported. That is, tactile suppression is not the primary factor affecting perception in this context. This implies that design of wearable haptic feedback

systems should first account for variation in sensitivity across feedback locations rather than an individual’s current state of motion. This chapter also evaluated whether sensitivity-adjusted perception models can support effective multi-sensory feedback by reducing the impact of stimulus location on perception.

## 2.3 Methods


In this study, we investigate the effect of voluntary motion on the perception of multi-sensory haptic cues used for motion guidance by comparing how participants perceive, interpret, and respond to multi-sensory haptic cues presented when they are at rest versus when they are already in motion. The experimental setup used a haptic feedback platform to provide motion guidance cues indicating a desired movement direction and speed, a robotic data collection platform, and a visual display. We hypothesized that participants presented with multi-sensory haptic cues while already in motion would have difficulty interpreting the cue, resulting in greater deviation from the direction and speed indicated by the motion guidance cue. As with previous studies, the location of the tactors may also affect perception.

### 2.3.1 Haptic Platform

We used a haptic platform comprised of eight eccentric rotating mass (ERM) mini-motors (Adafruit)—10mm in diameter and 2.7mm thickness operating at 183 Hz at 5V—arranged in a circular array around the circumference of the forearm, similar to [31, 50, 51], shown in Figure 2.1. The squeeze mechanism consisted of one band of nylon webbing around the upper arm whose circumference was controlled by a 270-degree high torque servo motor (TianKongRC TD-8120MG) as seen in Figure 2.1. The shortening of the band by the servo motor produced a squeezing sensation around the upper arm. The servo and tactors were controlled using four DRV 8883 motor drivers and an Arduino Uno that communicated with a host computer to synchronize the cues with the kinematic data collected.

The tactors and servo were activated simultaneously to indicate the desired motion. The activation pattern of the tactor array indicated direction. Participants were asked to move towards the direction where they felt the highest vibration intensity. A target angle of 0 degrees would activate the tactor on the right side of the forearm, while a target angle of 110 degrees would activate tactors on the upper-left surface of the forearm, as





**Figure 2.1.** The haptic platform was comprised of eight vibrating mini-motors arranged in a circular array around the circumference of the forearm and one band of nylon webbing around the upper arm, whose circumference was controlled by a servo motor. The shortening or lengthening of the band produced a squeezing sensation around the bicep. The tactors and servo were actuated simultaneously to indicate the desired motion, with the tactors providing the target angle and the servo providing the target speed.

shown in Figure 2.1. The angles used in the present study were 15, 22.5, 50, 75, 105, 112.5, 140, 165, 195, 202.5, 220, 255, 285, 295.5, 310, 345 degrees, and they did not directly correspond to the location of any of the tactors. Therefore, two adjacent tactors were activated simultaneously to create an illusionary stimulus. This phantom sensation would feel as if the vibration was coming from a point between two tactors, at the desired angles. To render these phantom sensations, we used the linear relationship presented by Alles [46] to determine the appropriate intensity of each tactors. As the motors used in this study were ERM vibration motors, the vibration frequency was directly proportional to the intensity of the vibration.

The intensity of the squeeze indicated the desired movement speed that a participant should achieve over the next two seconds. The squeeze level produced by the servo was broken up into four levels corresponding to the servo angle. We set the maximum servo angle to ensure the squeeze band achieved the desired tension without discomfort to the participant. The maximum speed corresponded to the maximum servo angle (and minimum band length). The other three levels, seventy-five percent, fifty percent, and twenty-five percent of the maximum servo angle, signaled to slower speed values respectively.

### 2.3.2 Robotic Data Collection Platform

The study used a Kuka LBR iiwa as a three-degree-of-freedom joystick to collect user inputs and to constrain the movement of the participants wrist so that the wrist and world coordinate system remained aligned. The iiwa is a seven-degree-of-freedom collaborative robot with a 14 kg payload and has a maximum reach of 820 millimeters. While in use, participants are seated in a chair facing the robot with a visual display and are asked to grasp the robot’s end-effector. The robot was operated in cartesian impedance mode where the impedance value in one direction was set to 1,000 N/m, while the impedance values in the other two directions were set to 20 N/m—restricting the motion of the end-effector to the vertical plane with low impedance to in-plane movements. The end-effector’s position was computed using forward kinematics and recorded throughout each trial.

### 2.3.3 Visual Display

Along with the haptic platform, participants received feedback from a visual display during each of the trials. The display was used to prompt participants to remain at rest or perform a voluntary movement and signaled when the trial started and stopped. Text indicated stimulus onset and when data collection ended. For static condition trials, text prompted participants to hold still. For voluntary motion conditions, a purple dot was used to trace a desired trajectory, while a yellow dot tracked the participant’s actual position as shown in Figure 2.1. Participants were instructed to follow the purple dot as closely as possible. The desired trajectories randomly switched between circular motion and figure-8 motion, as shown in Figure 2.2.

After following the visual cue for a random amount of time, between 1.5 and 3.5 seconds, the visual tracking component would disappear and the tactile cue would be


applied to indicate a direction and distance to move. Two seconds after the initial haptic cue was first displayed, the haptic stimuli were removed, and the participant was instructed to rest before starting the next trial.

### **2.3.4 Experimental Procedure**

Twenty-two participants (10 female, 12 male) aged 18 to 40 were recruited for this study and provided informed written consent before participation. The protocol for this research study was approved by the Pennsylvania State University Institutional Review Board, under IRB number STUDY00020703.

All participants completed 5 trials for each unique combination of the sixteen illusory angles indicating the direction of motion, four velocity magnitudes, and two movement conditions—completing a total of 640 trials in a session. The order in which these combinations of independent variable tested was determined by generating the list of 128 possible combinations and repeating these combinations five times. Using a random number generator, the order of the trials was randomly selected—effectively shuffling the ordered list of trial conditions that was generated for each participant. This process is outlined in Figure 2.2. This randomization allowed us to avoid capturing effects due to increasing familiarity with the stimulus and test setup or confounding the effect of fatigue with the impact of the trial conditions. The timing of the stimulus was randomized to occur between 1.5 and 3.5 seconds after the start of the trial to avoid anticipatory movement and would remain active while participants executed the indicated motion. The total trial length ranged from 3.5 to 5.5 seconds. Participants completed trials over approximately a 1 hour time window, though only about 30 minutes of that time was actual data collection. The remainder of the time was breaks between trials while the experiment was reset and any additional breaks requested by the participant.

Before data collection began, participants were familiarized with the testing setup. The speed and angle cues were demonstrated individually. For the speed cue, participants performed several practice movements with feedback on their speed error. During testing each stimuli is presented at a constant intensity for two seconds. All participants responded to the haptic cue within the two second time limit. Participants received short breaks between trials and were offered additional breaks frequently to avoid fatigue.



**Figure 2.2.** This schematic shows how the set of unique trial conditions was set, duplicated 5 times such that each participant completed 5 trials per combination. The order of the list of duplicated trials was shuffled with a random number generator. For voluntary trials, half of the trials were randomly assigned to have the participant track a figure-8, and in the other half, participants tracked a circular trajectory.

### 2.3.5 Data Analysis

A polar velocity vector was calculated by taking the numerical derivative of the Cartesian positions recorded during the last 0.5 seconds of each trial. The velocity vector was converted to polar form, with the radius representing the speed and the angle being the direction. Only the latter part of the movement was used to avoid differences due to varying reaction times under static and voluntary movement conditions. The error

between the angle and magnitude of the velocity vector and the angle and magnitude of the desired motion was computed. A non-parametric, multi-factor, repeated measures Aligned Rank Transform (ART) ANOVA [52] was used to analyze the effect of starting condition, desired angle, and desired magnitude on each of the error measures. Post Hoc corrections to the results of the ART ANOVA were performed using Benjamini-Hochberg [53] false discovery rate method to adjust the p-values.

### 2.3.6 Sensitivity-Adjusted Perceptual Models

The perception of haptic cues can be affected by a variety of factors including tactile masking, tactile suppression, level of skin contact and presentation location on the body. While these factors can all play a role in our perception, their actual effect can vary greatly from person to person, reducing the quality of haptic feedback. We assess the possibility of using a data-driven approach to mitigate various factors by fitting perceptual models to the data and computing the angle and speed errors based on the data-driven models.

Using the models presented by Luzhnica et al. [49], the cued angles,  $v \in [0, 1]$ , and user responses,  $y \in [0, 1]$ , were normalized. The intensities,  $I_i$  of the tactors were normalized using the maximum and minimum amplitude of the PWM signal. We extend Luzhnica’s model to include skin squeeze as well as vibrotactile sensations. The squeeze intensity,  $I_{sq}$ , was normalized by the maximum allowed servo angle.

Given a pair of tactors  $i \in \{a, b\}$  actuated at intensities  $I_i$  at a known location  $x_i = i$ , the location of the phantom sensation can be described as  $x_p \in [0, 1]$ . In this experiment, we consider linear [46] and power [47] models. In the linear model, the phantom sensation location,  $x_p$ , is proportional to the relative stimulation intensity  $I_a$ . While in the power model, the phantom sensation location,  $x_p$ , is proportional to the square of the intensities. These models can be generalized for vibrotactile displays containing more than two tactors with perceptual sensitivity changes depending on the location of the stimulus [48]. For vibrotactile displays with  $N$  tactors  $0 < i \leq N$  and  $M$  segments between the tactors and known intensities  $I_i$  and  $I_{i+1}$ , we can estimate the within-interval phantom sensation location  $x_p$  between tactors  $i$  and  $i + 1$  using the linear model, (eq. 2.1) and power model, (eq. 2.3). Using (eq. 2.5), estimated value of  $x_p$ , and number of segments between tactors  $M$ , we can compute the estimated location of the phantom sensation across all tactors as  $v \in \{0, 1\}$ .

The spatial variation of perceptual sensitivity can be accounted for by scaling the intensities  $I_i$  using location specific factors  $s_i \geq 0.5$  and can be applied to both models as shown in (eq. 2.1) and (eq. 2.3). We extend both models to predict the speed response,

$x_s$ , based on the squeeze intensity,  $I_{sq}$ , of the  $N + 1$  actuator using a similar relationship. We set up both a linear and power relationship between the speed response and the squeeze intensity as shown in eq. 2.2 and eq. 2.4. To account for the interactions between the angle and speed responses, we included an interaction term composed of the product between angle and speed response and a weight,  $W$ , as shown in eq. 2.6.

$$x_p^{ls} = \frac{s_{i+1}I_{i+1}}{s_iI_i + s_{i+1}I_{i+1}} \quad (2.1)$$

$$x_s^l = s_{N+1}F + s_{N+2} \quad (2.2)$$

$$x_p^{ps} = \frac{s_{i+1}^2I_{i+1}^2}{s_i^2I_i^2 + s_{i+1}^2I_{i+1}^2} \quad (2.3)$$

$$x_s^l = s_{N+1}^2 * I_{sq}^2 + s_{N+2} \quad (2.4)$$

$$v = \frac{x_p - i}{M} \quad (2.5)$$

$$interaction = Wx_px_s \quad (2.6)$$

The optimal values for these sensitivity factors were found by minimizing the mean squared error between the user response and the predicted response as in (2.7), where  $y_i$  is the user response at stimulation intensities  $I_1^i, \dots, I_N^i, I_{sq}$  and  $v(S, I_1^i, \dots, I_N^i, I_{sq})$  being the response location estimated by the models at intensities  $I_1^i, \dots, I_N^i, I_{sq}$ .

$$S_0 = argmin \sum_{i=1}^D (y_i - v(S, I_1^i, \dots, I_N^i, I_{sq}))^2, s_i \geq 0.5 \quad (2.7)$$

Two sets of models were generated with one set being trained against the combined data of every participant and the other set was trained on the data for each individual participant. Each set contained six total models, three were based on the linear equation 2.1 and the other three were based on the power model equation 2.3. The three different models were trained on the data containing only the static, only the voluntary motion trials, and on combined data of the static and voluntary movement trials. The data-driven models were evaluate using using K-Fold cross-validation with 10 splits. The model performance was evaluated using the absolute decoding error, which was calculated as  $\varepsilon = |y - v|$ , where  $y \in [0, 1]$  is the user response and  $v \in [0, 1]$  is the phantom sensation location as predicted by the perception models. Comparison of the models was performed using the non-parametric Wilcox-signed rank test, to determine which model, if any, was capable of accurately predicting participant responses.

## 2.4 Results

### 2.4.1 Sample Response


Figure 2.3 shows sample responses to the haptic stimuli under static and voluntary motion conditions. The stimulus is applied in both plots at time  $t = 0$  seconds, indicated by the gray bar. In both trials participants received a haptic cue indicating them to move in a direction of fifty degrees at a speed of 75% of the maximum speed. Under that static condition, a delay of about 0.5 seconds was frequently observed between the display of the haptic cue and the initiation of the indicated movement. In the voluntary motion condition, a delay of 0.3 seconds in movement initiation was observed as indicated by the point in time where the y and z positions diverge. While the voluntary motion condition had a shorter delay in response than the static condition response, we did not have to perform any corrections for this difference as we only examined the last 0.5 seconds of the motions.

Figure 2.4 shows a typical position and velocity response to a fifty degree direction and seventy-five percent speed cue for both a static trial and voluntary motion trial. In the top plots for the static trial, the participants response starts near the zero starting position with an initial velocity of zero in both the y and z directions. Once the haptic cue is presented, the participants begin moving in the direction indicated by the black arrow with an increasing velocity. In the bottom voluntary motion plots, the participants start near the zero position and begin following the voluntary motion pattern until they are presented with the haptic stimuli. Once presented with the haptic stimuli, participants begin attempting to move in the presented direction as indicated by the black arrow. In the static response, participants responded in the same general direction of the desired trajectory. In the voluntary motion response, participants tracked the desired vector with greater precision.

### 2.4.2 Statistical Results

The angle error between the velocity angle and the target angle is shown in Figure 2.5. A non-parametric, three-factor, repeated measures Aligned Rank Transform ANOVA was performed using R [54], to assess the impact of the starting condition, target angle and target speed on angle error.

**Angle Error** The main effect of target angle ( $p = 1.67 \times 10^{-21}$ ,  $F = 12.120$ ) was statistically significant. Figure 2.5 shows that in general participants tended to




**Figure 2.3.** The top plot shows the participant response during a static trial, where they would remain at rest until the presentation of the haptic stimuli. The bottom plot denotes the participant response during a voluntary motion trial, where they would be following a figure-8 trajectory as indicated on the visual display. The haptic stimulus was applied at  $t = 0$  seconds, as indicated by the gray bar. In both trials the haptic stimulus indicated a direction of fifty degrees at the same speed. In the static response, the participant begins responding to the haptic cue around 0.5 seconds into the trial. The start of their response is clearly denoted by shift to a non-zero velocity. In the voluntary motion response, the participant begins responding to the haptic cue at around 0.3 seconds. The start of the voluntary motion response can be visualized by the point where the y-position diverges from the z-position, indicating a divergence from the figure-8 trajectory.

underestimate the target angle resulting in negative error for trajectories that were rotated clockwise from the target angle. Figure 2.5 also shows that the angle error was lower for directions in the upper-right quadrant, which indicates that participants were more successful at interpreting cues presented in that location. The interaction effect between the movement condition and target speed ( $p = 0.009$ ,  $F = 4.961$ ) was statically significant. There was no difference between static and voluntary motion conditions when the target speed was 25% and 75% of maximum speed. For target speeds of 50% and 100% of maximum speed, the angle error was lower when participants received the cue during motion.

The interaction effect between target angle and target speed ( $p = 0.013$ ,  $F = 1.553$ ) was statistically significant. On the left side of the arm and lower side of the arm (angles between 90 and 310 degrees) there were only small differences in angle error at different target speeds. However, in the regions where angle error was larger overall there were larger differences across target speed. This suggests that the target speed may






**Figure 2.4.** The top plots show the position and velocity of a typical response during a static trial in the YZ-Plane, while the bottom plots show the typical response during a voluntary motion trial. The black arrow in the plots indicates the haptically cued direction.

have further reduced perception of the angle cue on the right side of the participant's arm. The main effects of movement condition ( $p = 0.087, F = 4.171$ ) and target speed ( $p = 0.546, F = 0.722$ ) were not significant. The interaction effect between movement condition and target angle ( $p = 0.387, F = 1.068$ ) and between movement condition, target angle, and target speed ( $p = 0.278, F = 1.118$ ) were also not significant.


**Speed Response** A non-parametric, three-factor, repeated measures Aligned Rank Transform ANOVA was also performed on the speed response. The main effect of



**Figure 2.5. Angle error vs. Target direction:** Participants were able to more accurately interpret and execute target angles that were in the direction of up and to the right. The accuracy of participants interpretation and execution of target angles was not affected by their state of motion.

target speed ( $p = 7.72 \times 10^{-8}$ ,  $F = 24.582$ ) was statistically significant. In Figure 2.6, as the intensity of the speed cue increased, participants moved faster, indicating that participants could accurately perceive and interpret the different squeeze intensities and respond accordingly. The interaction between movement condition and target speed ( $p = 0.008$ ,  $F = 5.021$ ) was also significant. Figure 2.6 shows that generally the speed response for the voluntary movement condition was larger. The larger response indicates that when participants were already in motion, they tended to overshoot the target speed compared to trials where they started at rest except in the case of the max speed. Figure 2.6, also shows that participants had a strong tendency to overshoot the desired speed value.

The main effect of target angle ( $p = 0.011$ ,  $F = 2.091$ ) was significant. Figure 2.7 shows that the speed response was slower for target angles of 255 to 310 degrees, indicating that in general for direction cues pointing down and to the right, participants were less successful at replicating their response compared to their response in other directions for the same speed cue. The interaction between target angle and target speed ( $p = 0.044$ ,  $F = 1.403$ ) was not significant. However, one can see that in Figure 2.7, generally, the higher speeds corresponded to higher target speeds with a few cases where



**Figure 2.6. Speed responses vs Target speed:** Participants were able to differentiate the different levels of squeeze and increased the speed of their movement when the intensity of the squeeze increased. However, participant responses tended to not closely match the true desired speed value and would overshoot it.


a pair of adjacent target speeds did not result in different speed responses as can be seen in target angles, 22.5, 75, 112.5, and 220. The interaction effect between movement condition, target angle and target speed ( $p = 0.008$ ,  $F = 1.667$ ) was significant. The main effect of movement condition ( $p = 0.237$ ,  $F = 1.724$ ) and the interaction effect of movement condition and target angle ( $p = 0.156$ ,  $F = 1.382$ ) were not significant.

We compared the absolute decoding error of linear and power models trained on the data aggregated across all participants using eight 2-sample t-tests. In general, we found that both the linear and power models could reduce user response error. We also

**Table 2.1.** This table contains the mean and standard deviation of the angle error for each model trained using the combined data from all participants. All the models have the capability of reducing the response error. The linear model trained using the combined static and voluntary motion data had the greatest reduction in the angle portion of the absolute decoding error.

| Model  | Static |        | Voluntary |        | All    |        |
|--------|--------|--------|-----------|--------|--------|--------|
|        | $\mu$  | SD     | $\mu$     | SD     | $\mu$  | SD     |
| Linear | 15.44° | 10.92° | 15.33°    | 11.59° | 8.723° | 11.04° |
| Power  | 18.51° | 11.41° | 19.18°    | 11.95° | 11.39° | 10.04° |

compared linear and power models trained using only static trial data, only voluntary



**Figure 2.7. Speed responses vs Target angle:** Generally, as the speed intensity increased, participants increased their speed. For target angles between 255 to 310 degrees, participants' speed responses were generally slower than the other directions.

motion trial data, and data from all trials, summarized in Tables 2.1, 2.2, 2.3. In all cases, both linear and power models showed promise at reducing the perception error; however, the linear model was more effective at reducing the error for vibrotactile cues and the power model was more effective at reducing the error for squeeze cues. The significant differences between the various models shown in Table 2.3, show that training the models using only static or only voluntary motion results in less effective compared to training the models using both the static and voluntary motion data. The statistical comparisons also show that there is a significant difference between the linear and power models, but no difference between training with static motion data or voluntary motion data for the linear model.

**Table 2.2.** This table contains the mean and standard deviation of the speed response for each model trained using the combined data from all participants. All the models have the capability of reducing the response error. The power model trained using the combined static and voluntary motion data had the greatest reduction in the speed portion of the absolute decoding error.

| Model  | Static    |           | Voluntary |           | All       |           |
|--------|-----------|-----------|-----------|-----------|-----------|-----------|
|        | $\mu$     | SD        | $\mu$     | SD        | $\mu$     | SD        |
| Linear | 2.597 m/s | 1.510 m/s | 2.822 m/s | 1.678 m/s | 7.220 m/s | 1.512 m/s |
| Power  | 4.480 m/s | 2.835 m/s | 4.549 m/s | 2.880 m/s | 3.710 m/s | 2.054 m/s |

**Table 2.3.** This table shows the resulting p-values for the non-parametric Wilcoxon-signed rank test comparing the absolute decoding error of the different models trained using the combined data from all participants. The absolute decoding errors  $\varepsilon$ , was denoted by the training data: S (static trials), V (voluntary motion trials), or C (both) and the model: L (Linear) or P (Power). There is a statistically significant difference in the absolute decoding error between the static linear and static power models, the static linear and dynamic power, the static power and dynamic power, the dynamic linear and dynamic power, the static linear and linear, the static linear and power, the static power and linear, the static power and power, the dynamic linear and linear, the dynamic linear and power, the dynamic power and linear, the dynamic power and power, and the linear and power models.

| Model              | $\varepsilon_{SL}$ | $\varepsilon_{SP}$ | $\varepsilon_{VL}$ | $\varepsilon_{VP}$ | $\varepsilon_{CL}$ |
|--------------------|--------------------|--------------------|--------------------|--------------------|--------------------|
| $\varepsilon_{SP}$ | 0.002              |                    |                    |                    |                    |
| $\varepsilon_{VL}$ | 0.432              | 0.002              |                    |                    |                    |
| $\varepsilon_{VP}$ | 0.002              | 0.037              | 0.002              |                    |                    |
| $\varepsilon_{CL}$ | 0.002              | 0.002              | 0.002              | 0.002              |                    |
| $\varepsilon_{CP}$ | 0.012              | 0.002              | 0.002              | 0.002              | 0.002              |

## 2.5 Discussion

Prior studies show that increasing the complexity of haptic cues and using multiple cutaneous stimuli can increase the number of distinct cues with a wearable haptic device—potentially enabling higher information transmission via cutaneous stimuli. However, the complexity of such cues, interference between stimuli, and tactile suppression may limit such haptic interfaces, especially when designed to guide motion or improve communication between a robot and human during voluntary motion tasks. Tactile suppression in its simplest definition is a phenomenon that degrades our ability to accurately perceive and interpret tactile stimuli during movement [55]. This phenomenon dominates from shortly before the beginning of a movement and continuing until the motions end, with a peak attenuation at the exact point motion is initiated [13, 15, 33]. Following the results of existing literature, we should expect that in our experiment, participant response accuracy to the haptic cues should be poor or inconsistent under the voluntary motion starting condition. Instead, we found that our main hypothesis was not supported and that in general participant response accuracy was unchanged by voluntary motion. In goal-directed motions, the effects of tactile suppression were absent in areas where tactile information is relevant to the motion [14], which could explain the greater quality of the responses. However, all the tactile stimuli supplied in our study was presented on the participant’s forearm, and according to the study by Colino et. al [14], vibrotactile stimuli applied to the forearm would consistently be

affected by tactile suppression. These findings suggest that tactile suppression is not a dominant factor in haptic perception accuracy for motion guidance. Our statistical results indicate that the cued angle was the dominant factor affecting perception errors. In this experiment, the indicated angle coincides with unique stimulus locations on the body. This suggests that the stimulus location is among the most important factors affecting haptic perception. This is consistent with other work on tactile perception [22]. This change in perceptual accuracy due to stimulus location is often attributed to varying densities of mechanoreceptors on different arm surfaces [56], or to the spatial differences in perceptual sensitivity around the forearm [48].

Perceptual interference (also known as tactile masking), is the phenomenon where our perception accuracy of one stimuli is degraded when in the presence of another [17]. This degraded perceptual accuracy has a greater chance of occurring in multi-sensory devices as they supply multiple types of stimuli simultaneously [57]. In our results, participants responded with a slower speed when presented with directional cues between target angles of 255 to 310 degrees, showing a possible confusion of the intensity of the presented stimuli. In our angle error results for target angles pointing to the right, participants had increased variation in their responses. This increased variation as compared to other directions could be caused by perceptual interference from the interaction of the vibration and squeeze cue. This result falls in line with work done in [45], where a loss of accuracy was found even though vibration and the other stimuli are not being applied at the exact same location.

This study used two data-driven perception models to analyze the possible improvement in perception for the vibrotactile and skin-squeeze feedback in the voluntary motion conditions. Our analysis of the data-driven, sensitivity-adjusted perception models found that there was a significantly better result when training the models using the combined static and voluntary data compared only static or only voluntary motion trials.

While the models used in this paper have shown a potential for great improvement. They may be challenging to implement online in a wearable device due to computation and storage limitations on portable micro-controllers like Arduino or ESP32. The feedback will also be limited by the performance characteristics of the hardware. In the present study, we use ERM tactors, where vibration amplitude and frequency cannot be controlled independently. The quality of haptic feedback might be improved with hardware that can independently control the amplitude and frequency of vibration such as linear resonant actuators. This change would increase perceptual model complexity and likely increase the amount of training data needed. Without appropriate training

data, the perceptual model used here could result in poor model performance. More sophisticated learning algorithms might enable online learning with relatively few data points. Further, using more general models—including one capable of switching between attractive and repulsive feedback—may allow for learning and adapting to user preferences and as well as perceptual differences.

## 2.6 Conclusions

This chapter investigates whether tactile suppression plays a role in the perception of multi-sensory haptic feedback. The feedback information was conveyed using a vibrotactile array for a desired direction and a squeeze band for a desired speed. Our results demonstrated that our main hypothesis was not supported and our ability to perceive haptic stimuli on the forearm accurately is not significantly affected by state of motion, instead stimulus location seems to be the most important factor. Our results also show that sensitivity-adjusted perception models are capable of reducing the impact of stimulus location for a wide range of people. These results provide insight into factors to consider in the design of wearable haptic devices to provide more effective feedback. Specifically, our results demonstrate that voluntary movement need not be explicitly considered in device and control design, but that it should be included in data-driven control methods. By considering these factors, improving and creating wearable haptic feedback devices to further enrich our experience with AR/VR and facilitate effective human-robot interaction in teleoperation is possible.

In addition to compensating for relative decreases in tactile perception across participants or body locations, data-driven control methods could enable more sophisticated adaptation of tactor control to user preferences and task performance. With appropriate training data and algorithmic updates to the tactor control policy, it is possible that for different users or tasks, a unique optimal policy could be learned. For instance, the literature is mixed on whether repulsive or attractive vibrotactile feedback is more effective or if it simply comes down to user preference. Future work will include customization of the tactor control using data-driven approaches similar to the perceptual models we investigate in the present study. Such data-driven approaches may be able to handle both body-location dependencies and preferences with sufficient model complexity. However, the challenge will be developing data-efficient strategies to train these models and adapt them to perturbations to the interface between the user and wearable device.

# Chapter 3 |

## Substitutionary somatosensory haptic feedback in EMG-Driven control

Modern prosthetic and assistive interfaces increasingly rely on indirect control signals such as surface electromyography (EMG), abstract mappings, and algorithmic inference rather than direct biomechanical actuation. While these systems can restore a degree of functional control, they fundamentally alter the sensorimotor loop: motor intent is filtered through noisy signals, device dynamics differ from biological limbs, and proprioceptive feedback is limited or absent. Under these conditions, it is not obvious whether classical principles of motor learning—derived largely from studies of intact limb control—still apply.

A central challenge for practical prosthetic control is stability over time. Real-world systems experience calibration drift, fatigue-related changes in EMG, socket shifts, and mechanical wear, all of which perturb the mapping between user intent and device behavior. If users cannot reliably adapt to such changes, long-term usability suffers. Conversely, if adaptation mechanisms remain functional under artificial control, then structured training paradigms and appropriately designed feedback may significantly improve performance and robustness.

This chapter addresses that gap by examining whether a well-established experimental framework from motor neuroscience—visuomotor perturbation—can be meaningfully applied to a myoelectrically controlled virtual elbow, and whether supplemental haptic feedback supports adaptation in this artificial control loop.



## 3.1 Research Questions

The study is structured around three core research questions that guide the experimental design and interpretation of results:

- Can users adapt to a systematic visuomotor gain perturbation when controlling a joint through EMG rather than their biological limb?
- How does adaptation under myoelectric control differ from adaptation in natural limb control?
- Does haptic feedback improve adaptation?

## 3.2 Introduction

Biological movement control relies on continuous sensory input from proprioceptive sensors to mediate high-precision, adaptive actions, which often occur without significant conscious effort. In contrast, operating assistive robots or prosthetics typically demands a high degree of cognitive attention with reduced intuitiveness [58], which can result from a lack of feedback about specific degrees of freedom. Biologically driven robotic prosthetics [59], incorporating somatosensory feedback [60], have begun to address this issue; however, the development of sophisticated, lifelike prosthetics is often prioritized over understanding fundamental issues in human control and adaptation, which are made difficult by the complexity of the device. We seek to improve our understanding of the role of haptic somatosensory feedback in controlling myoelectric prostheses.

In real-world tasks, visual information plays a key role in controlling biologically driven devices, especially in uncertain or adaptive environments. To understand how haptic feedback can improve our ability to adapt to environmental changes, we will evaluate its impact on a simple visuomotor adaptation paradigm by implementing a gain shift. Visuomotor adaptations are a form of sensorimotor learning that occurs when people are trying to adapt and correct for an external visuomotor perturbation [61]. Methods for implementing these adaptations can involve modifying the movement direction by rotating the displayed path or changing the movement amplitude by adjusting the gain between movement amplitude and visual feedback [62].

We hypothesize that adaptation to visuomotor perturbations in EMG-controlled prosthetics will be slower and less accurate than adaptation with the actual arm. However,

with the addition of haptic feedback, adaptation will improve, more closely resembling that in a healthy limb.

### 3.3 Background

Motor adaptation is a fundamental aspect of human motor control, enabling individuals to maintain accurate performance despite changes in the environment or the body. Visuomotor perturbation paradigms induce motor adaptation by introducing a systematic visual error, such as altering the gain between the limb’s movement and the displayed cursor position or rotating the axes on which the cursor moves. In gain perturbation paradigms, the gain determines how much the visual cursor responds to movement by scaling the displayed distance relative to the actual movement. Rotation perturbations apply an angular shift to cursor motion—for example, a straight forward reach in the sagittal plane may appear as a diagonal movement at a  $45^\circ$  angle. For adaptation to occur, the visual error must be systematic rather than random to justify a change in behavior [63,64], and access to concurrent visual and proprioceptive information is required for full adaptation [65].

Translating these paradigms to teleoperated systems and assistive devices, such as prostheses, presents additional challenges, particularly when control is mediated through electromyography (EMG). Although EMG-based control provides a direct interface to muscle activity, many such systems rely heavily on visual servoing [66], which can be slow and cognitively demanding. EMG signals are also noisy and variable, with their quality affected by factors such as skin impedance and electrode placement [67]. Together, these issues make sole reliance on vision in real-world settings both unrealistic and impractical [68].

To mitigate these challenges, various forms of feedback have been explored to enhance user performance, generally falling into two categories: augmentation and substitutionary. Augmentation feedback methods provide additional sensory information beyond what is naturally available and tend to provide modest benefits for control and motor learning [69,70]. Substitutionary feedback methods are feedback methods that supply sensory information to the body using a different channel than would normally be used or through the standard channel in a different form (such as temperature in place of pressure) [71]. Sensory substitution feedback has been implemented with vibrotactile feedback [72], electro-tactile feedback [73], auditory feedback [74] and skin stretching [75].

Substitutionary feedback remains challenging to implement, partly due to limited

actuator hardware and the difficulty of designing intuitive, high-fidelity encodings [76–78]. It can also introduce unintended consequences; for example, poorly tuned force-based feedback has been shown to destabilize human-in-the-loop systems [79]. Nevertheless, substitutionary feedback can outperform visual feedback alone, as vision is slower and imposes a higher cognitive load compared to proprioceptive channels. When designed with intuitive, straightforward mappings, substitutionary tactile feedback can meaningfully improve control [60, 80].

The importance of effective feedback is particularly evident in upper-limb prosthetics. Many commercially available prostheses require constant visual monitoring for accurate control, particularly single-degree-of-freedom devices [66]. Without some form of tactile or proprioceptive feedback, amputees cannot accurately control grasping force and perform fine manipulation [81]. Furthermore, a lack of meaningful interactions and tactile/proprioceptive feedback can exacerbate phantom limb pain, increasing cognitive effort and further diminishing the usability and acceptance of the prosthesis [82].

Given these limitations, substitutionary tactile feedback offers a promising avenue for improving and supporting myoelectrically driven devices. Intuitive, straightforward mappings can reduce cognitive load and enhance performance; however, they must be designed with this in mind. This study uses a linear mapping of elbow rotation to the rotation of a wearable haptic twist module to examine differences in visuomotor adaptation with natural, no, and substitutionary feedback, with the aim of developing feedback interfaces that support effective control of robotics via myoelectric interfaces.

### 3.4 Methods

This study examines the impact of somatosensory feedback on motor adaptation in a bidirectional human-machine interface involving elbow flexion and extension. Specifically, we compared motor adaptation performance across three groups during an elbow reaching task in a virtual reality apparatus: a control group performing natural reaching with their physical arm, an isometric group that mimicked elbow motion by using EMG signals from the biceps and triceps to control a virtual elbow joint via a servo, and a group using the same EMG-driven control with continuous haptic feedback encoding joint position.


### 3.4.1 Virtual Reality Apparatus

In this study, we used the Kinereach virtual reality and motion tracking system, shown in Figure 3.1. Participants were seated in a height-adjustable chair facing a glass table-top with a mirrored screen placed at chin height, which occluded the view of the hands, forearm, and EMG-driven elbow. For participants performing natural reaching motions, the upper arm was supported and fixed in place using a custom-built brace, shown in Figure 3.2a, and the forearm was supported using air sleds to reduce the effects of friction and gravity on arm movement. Participants performing virtual reaching using EMG-driven elbow, had their upper and forearm supported by a custom brace to prevent any movement of the arm, as shown in Figure 3.2b. We recorded the index finger position and the virtual elbow position at 116 Hz using a trakSTAR (Ascension Technology) six-degree-of-freedom motion-tracking system. An HD TV with a 1080p display resolution and a 60 Hz refresh rate was used to project the experimental task onto the mirror screen.

### 3.4.2 Representative Prosthetic Elbow

A prosthetic elbow joint was modeled as a rectangular rod attached to a TowerPro MG995R servo, controlled by an Arduino Uno. The servo was controlled myoelectrically by participants using two electromyography (EMG) electrodes. The EMG electrodes were placed on the belly of the participant’s bicep and on the short head of the participant’s tricep. The signals from the EMG electrodes were amplified with a gain of 13.89, rectified using a passive full-wave rectifier, and then passed through a voltage protection circuit before reaching the Arduino’s analog-to-digital converter. We sampled the EMG signals at 1000 Hz and filtered the incoming signals using a tunable digital low-pass filter with a cutoff frequency ranging between 8 and 16 Hz on the Arduino Uno. Before the task, peak EMG values were recorded while participants performed isometric contractions to generate flexion and extension torques at approximately 50% of their maximum effort. Similar to the methods described in Vargas et al. [83], we used the peak EMG to normalize the participant’s EMG data as indicated in Equation 3.1.

$$NormEMG[t] = \frac{EMG[t]}{0.25 \times EMG_{peak}} \quad (3.1)$$



**Figure 3.1.** Participants were sat in the Kinreach virtual reality and motion tracking system, where a visual display was presented on a glass mirror from an inverted HD TV (Top Left). Participants performed natural reaching motions using elbow flexion and extension or virtual reaching using isometric contractions to control EMG-driven elbow (Bottom Right). Participants controlled the EMG-driven elbow using isometric contractions. The EMG signal was passed through an analog amplification and rectification circuit before reaching the Arduino microcontroller used to modulate the angular position of the servo. The Participants will perform 80 baseline reaches, 80 reaching movements with a visual gain perturbation and 80 reaching movements with the visual gain perturbation turned off.

At each discrete timestep,  $t$ , the difference between the normalized EMG signal of the extensor and flexor was computed using Equation 3.2.

$$\Delta EMG[t] = NormEMG_{Extensor}[t] - NormEMG_{Flexor}[t] \quad (3.2)$$

The angular position of the servo motor was calculated using the difference between the normalized EMG signals, as shown in Equation 3.3,

$$\theta_{servo}[t] = \begin{cases} \theta_{servo}[t-1] + G_{speed} \Delta EMG \\ \quad - G_{damping} (\theta_{servo}[t-1] - \theta_{servo}[t-2]) & |\Delta EMG| > D \\ \theta_{servo}[t-1] & |\Delta EMG| \leq D \end{cases} \quad (3.3)$$




**Figure 3.2.** a) **Natural reaching task:** The participant’s upper arm was secured in a brace to restrict shoulder motion and limit movement to only elbow flexion and extension. The forearm was placed in a splint to immobilize the wrist and fingers and rested on a low-friction air sled to minimize the effect of friction. b) **Isometric virtual reaching task:** Both the upper arm and forearm were secured using custom braces, entirely restricting limb movement and enforcing isometric contractions.

where  $G_{speed}$  is a scalar gain that maps the normalized EMG difference to an incremental change in servo angle per discrete timestep,  $G_{damping}$  is a scalar gain that attenuates rapid changes in servo motion by penalizing inter-sample angular velocity, and  $D$  is a deadzone threshold defining the minimum absolute value of  $\Delta EMG$  required to induce motion. All controller parameters are implemented directly in discrete time on the Arduino Uno and therefore have units of either bits, or in the case of  $\Delta EMG$  and  $D$ , they are a normalized ratio of the bit values.  $\theta_{servo}[t]$  is represented internally as a servo command index updated at the controller sampling rate. The nominal value of  $G_{speed}$  was set to 0.75. The value of  $G_{damping}$  was established for participants during the initial calibration period and ranged from 0.0 to 0.25 with a nominal value of 0.10. The deadzone threshold was established for each participant based on their individual muscle background activity, ensuring that any difference in EMG signals was due to an intentional contraction of the muscle. The servo was mounted with a trackSTAR sensor, which served as the input to the Kinereach system.

### 3.4.3 Somatosensory Feedback Platform

The somatosensory haptic feedback was supplied using the Snaptics [84], a low-cost, open-source wearable haptics platform. The platform consisted of one twist module shown in Figure 3.3, providing somatosensory feedback for the angular position of the virtual elbow using the twist module, where the maximum amount of motor rotation will be analogous to a fully contracted elbow. This haptic feedback was presented to the

participant continuously during the experiment.



**Figure 3.3.** The Snaptics platform is a low-cost, open-source wearable haptics platform consisting of three feedback modules: vibration, twist, and stretch. Although the vibration and stretch modules are depicted, they remained inactive during the experiment. Only the twist module was used to provide haptic feedback. The angular position of the twist module was mapped in real-time to the virtual elbow’s joint angle.


#### 3.4.4 Motor Adaptation Protocol

Thirty participants were recruited for this study. Participants were randomly assigned to one of three groups: Control, Isometric (no feedback), and Isometric + Haptic Feedback. All participants completed 240 trials in a single one-and-a-half-hour session, structured into three blocks: Baseline (80 trials), Perturbation (80 trials), and Washout (80 trials). We introduced a gain shift during the perturbation block and removed it during the washout. The protocol for this research study was approved by the Pennsylvania State University Institutional Review Board, under IRB number STUDY00027479.

The task involved a one-degree-of-freedom elbow flexion/extension reaching movement. Participants either moved their physical elbow (Control group) or used EMG activation to control the virtual elbow (Isometric groups). The virtual cursor representing elbow position was aligned with one of two randomly displayed targets. Participants were instructed to reach the target in a smooth, single motion upon hearing an auditory cue. Feedback was provided as a score per trial: 10 points for perfect alignment, 3 for moderate accuracy, and 1 for near misses. These scores serve only as motivational feedback and were not used in data analysis.

Before testing, participants underwent a familiarization phase of up to forty-eight guided practice trials and an EMG calibration session (for isometric groups). The EMG calibration ensured signal quality and established the normalization, deadzone, low-pass filter, and damping parameters. During the task, the arm was either immobilized (in

isometric groups) or supported with air sleds (in the control group) to maintain consistent joint isolation. Target appearance and the movement cue were visually presented on the Kinereach interface (Figure 3.4). Cursor control was either via index-finger movement or a servo-driven virtual elbow, depending on the group.



**Figure 3.4.** The Kinereach virtual reality system was used to display targets to the participants. The green dot represents the start position. The gray circle with a central blue dot indicates the target. The circular icon with a crosshair represents the cursor.

### 3.4.5 Isometric Calibration Protocol

Before connecting the EMG signals to the virtual elbow, participants performed a set of flexor and extensor contractions to obtain the necessary calibration parameters. During these contractions, participants were instructed to perform the contraction at 50% of their maximum effort. Using an oscilloscope, we recorded the peak amplified and rectified EMG signal for each muscle and mapped this signal to the Arduino Uno’s 10-bit analog-to-digital converter. The resulting peak value,  $EMG_{peak}$ , was therefore expressed in bits (0–1023). The scaling factors were incrementally adjusted down until the normalized values were near zero. After connecting the EMG signals to the virtual elbow, an instructional session was used to train the participants and tune the other



controller parameters used in Eq. 3.3.

In the first part of the instructional session, participants performed isometric flexor and extensor contractions at 50% of their maximum effort until the virtual elbow reached the position corresponding to full flexion or full extension. The movement of the virtual elbow was observed to identify asymmetries in the movement speed between flexion and extension. This required adjusting the values of  $EMG_{peak}$ . If an asymmetry was identified, the  $EMG_{peak}$  of the slower moving direction was reduced until the asymmetry was resolved or the value was less than 150 bits, since lowering it further could result in instability due to signal noise. In that case, the value of  $EMG_{peak}$  for the faster movement direction was increased to resolve the asymmetry. Iterative adjustments in this manner resulted in similar movement speeds for flexion and extension when the same perceived effort was applied.

The deadzone,  $D$ , was tuned by asking participants to gradually increase their effort from 0 to 50%, when performing flexor and extensor contractions, which allowed us to observe how smoothly the virtual elbow responded to incremental changes in muscle activation. Values of  $D$ —the normalized ratio of the difference between flexor and extensor EMG signals—could be set between 0 and 1. However, we use a set of discrete values to anchor the tuning process. When the motion of the virtual elbow was not smooth, the deadzone was incrementally increased from 0.15, where jerks of the virtual elbow were minimal up to 0.25 for moderate and 0.30, for severely non-smooth movements. Participant perception of effort informed more minute adjustments, particularly for cases where minor co-contraction of the antagonist muscle was present.

Finally, participants performed up to 48 reaching trials using the same visual feedback and task as the baseline session. During these trials, the same decision process was used to fine tune the value of  $EMG_{peak}$  and  $D$  to achieve a smooth and balanced control of the actual cursor that provided visual feedback to the participants. If the servo’s movements exhibited dynamic instabilities during this phase, such as rapid start-stop behavior or oscillations, we activated the damping term and tuned the damping coefficient,  $G_{damping}$ . Increasing  $G_{damping}$  consistently introduced additional effort, but this effect was nonlinear: small increases improved stability with only marginal changes in perceived effort, whereas exceeding an individual-specific threshold produced a sharp rise in effort without proportional benefit to control quality. To identify an appropriate damping value,  $G_{damping}$  was incrementally increased in steps of 0.05 until the instability was sufficiently attenuated or the participant reported that further increases imposed a substantial effort cost. Based on this criterion,  $G_{damping}$  was capped at 0.30, as values above this threshold

consistently required participants to exert a substantial amount of effort (often reported as 7-9 out of 10) to produce only a moderate amount of movement.

### 3.4.6 Data Analysis

The actual position of the virtual elbow was measured using the trakSTAR sensor attached to the end of the rectangular rod.

To determine participant performance, we first identified the start and end times of each movement. We obtained the velocity profile for each trial by applying an 8 Hz low-pass filter to the Cartesian position data of either the virtual elbow or the participant’s hand. The velocity at each time step was computed from the numerical derivative of the Cartesian position. After obtaining the velocity profile, we further filtered it using a Savitzky–Golay filter [85] with a 35-point sample window to smooth out rapid velocity changes that occur during rapid corrections.

The start of the movement was defined as the Cartesian position corresponding to the minimum velocity immediately preceding the velocity peak closest to the actual start position. The end of the movement was defined as the Cartesian position corresponding to the minimum velocity, occurring after a velocity peak, for which the velocity remained below 8% of the maximum velocity for more than 200 milliseconds. After obtaining the start and end points of the movement, we computed the distance between them.

To evaluate how participants modulated control of their physical arm or the EMG-driven virtual elbow, we computed the average velocity for each trial, defined as the average tangential velocity of the virtual elbow or the participant’s hand. A key indicator of adaptation is velocity modulation, where participants should have a lower average velocity when the gain is increased during the perturbation period. Following the removal of the perturbation in which the relative gain is decreased back to baseline, the average velocity should grow back to baseline levels during the post-perturbation period. We compared the average velocity from the last ten baseline trials, first and last ten gain shift trials, and the first and last ten post gain shift trials, called the Late Baseline (LB), Early Adaptation (EA), Late Adaptation (LA), Early Washout (EW) and Late Washout (LW) epochs respectively, using a two-way mixed model ANOVA, with group as the between-subjects factor and epoch as the within-subject factor.

Participant performance was also quantified using distance error, computed as the difference between the required travel distance to reach the target from the start position and the actual distance traveled by the cursor. Distance errors were then averaged over epochs of ten trials. To assess adaptation and performance among subject groups, we

compared the distance errors over the same epochs, using the same two-way mixed model as for average velocity. Post hoc comparisons were conducted using pairwise Student's t-tests and corrections were performed using the Benjamini-Hochberg [53] false discovery rate method to adjust the p-values.

## 3.5 Results

### 3.5.1 Sample Responses


Figure 3.5 shows sample trajectories for two trials during the Late Baseline, Early Adaptation, and Early Washout Epochs from a participant in each group. Each pair of trials contains both an outwards reaching motion shown on the left set of trajectories for the Tricep and an inwards reaching motion shown on the right set of trajectories for the Bicep. During the Late Baseline Epoch, we can see that participants move the cursor from the start position using either the EMG-driven virtual elbow or their real arm, and line it up with the target's center. After the gain shift is introduced in the early Adaptation Epoch, participants in each group now overshoot the target. Once the gain shift is removed during the Early Washout epoch, participants in each group undershoot the target, indicating that they had adapted their movements to account for the change in gain.

Figure 3.6 shows the corresponding EMG signals from the Triceps and Biceps and the servo angle corresponding to the sample trajectories shown in Figure 3.5 for the participant in the Isometric (no feedback) group. Figure 3.7 shows the corresponding EMG signals from the Triceps and Biceps and the servo angle corresponding to the sample trajectories shown in Figure 3.5 for the participant in the Isometric + Haptic feedback group.


### 3.5.2 Statistical Results

#### 3.5.2.1 Average Velocity

A two-way mixed ANOVA with group as the between-subjects factor and epoch as the within-subjects factor was conducted over the average velocity. The main effect of group ( $p = 0.230, F = 1.560$ ) was not significant, indicating that each group used a similar velocity during reaching movements, including the control group using their actual arm. The main effect of epoch ( $p = 7.5 \times 10^{-12}, F = 20.482$ ) was statistically significant.




**Figure 3.5.** Sample trajectories showing the path traveled by the cursor for a participant from the Control group, the Isometric (no feedback) group, and the Isometric + Haptic Feedback group, during the Late Baseline, Early Adaptation, and Early Washout Epochs




**Figure 3.6.** Sample EMG response from a Participant in the Virtual Reaching Isometric (no feedback) group during Late Baseline, Early Adaptation, and Early Washout.

As expected, we can see that each group significantly reduced their average velocity by approximately 0.05 meters per second during the adaptation phase in Figure 3.8 in response to the gain increase perturbation. The interaction effect of group and epoch ( $p = 0.230, F = 1.493$ ) was not significant, indicating that each group modulated its velocity similarly when exposed to the gain perturbation and when the gain was returned to baseline.



**Figure 3.7.** Sample EMG response from a Participant in the Virtual Reaching Isometric + Haptic Feedback group during Late Baseline, Early Adaptation, and Early Washout.




**Figure 3.8.** Plot showing the Average Velocity (meters per second) averaged over a set of epochs containing ten trials. The shaded regions represent the standard error. The gain perturbation was introduced after epoch 8 and removed after epoch 16, indicated by the two vertical lines. Participants in each group performed the reaching motions at a similar average velocity throughout the experiment. Participants in each group also reduced their average velocity during the adaptation phase to account for the effect of the gain perturbation.

### 3.5.2.2 Distance Error

Figure 3.9 shows the distance error over the course of the experiment, with each epoch representing the mean error over ten consecutive trials. In this Plot, we can see that, for

each group, participants decreased their error across the baseline trials as they became more familiar with natural reaching or with controlling the EMG-driven virtual elbow. After the gain perturbation was applied, we see that each group experienced a positive increase in error, indicating that they are now overshooting the target. As participants continued to experience the gain shift, the errors decreased as they adapted to the perturbation. After the gain shift is removed, participants in each group undershoot the target, resulting in a negative error. Over the remainder of the trials, participants in each group adjust to the change, and their errors trend towards their baseline performance. These changes in the distance error and plot shape share all the traits expected in a standard visuomotor gain perturbation experiment. To further analyze and compare participants' performance across groups and epochs, we conducted a two-way mixed ANOVA with group as the between-subjects factor and epoch as the within-subjects factor.



**Figure 3.9.** Plot of the Distance Error (meters) averaged over a set of epochs containing ten trials. The shaded regions represent the standard error. The gain perturbation was introduced after epoch eight and removed after epoch 16, indicated by the two vertical lines.

The main effect of group ( $p = 0.039$ ,  $F = 3.721$ ) was statistically significant. Figure 3.9 and Figure 3.10 both show that, in general, participants in the control group had the lowest distance error. In contrast, participants in the virtual reaching Isometric + Haptic Feedback group had the largest distance error. Participants receiving haptic feedback appear to have higher error rates than those in the other two groups, particularly in the early baseline trials and early adaptation.

The main effect of epoch ( $p = 8.845 \times 10^{-18}$ ,  $F = 33.631$ ) was significant. As shown in Figure 3.10, error was higher in early baseline, early adaptation, and early washout when participants were initially learning to control the prosthetic and when gain changes were applied to the system. The interaction effect of group and epoch ( $p = 2.63 \times 10^{-4}$ ,  $F = 4.142$ ) was statistically significant, which indicates that the effect of the gain perturbation on the distance error was different depending on the group. In Figure 3.10, we can see that the distance error for the Isometric + Haptic Feedback group ( $p = 4.02 \times 10^{-6}$ ,  $t = -5.47$ ) and the Isometric (no feedback) ( $p = 0.006$ ,  $t = -3.30$ ) group was significantly larger than the Control at the Early Adaptation Epoch, which indicates that the introduction of the gain perturbation had a more pronounced impact on the participants in both Isometric groups.

To further characterize the differences across epochs, we conducted pairwise t-tests for each experimental group. The complete list of pairwise comparisons is provided in Appendix A.

**Control:** Within the Control group, we found no significant differences between any consecutive epochs. Specifically, we did not find significant differences in response to the gain shifts from Late Baseline (Epoch 8) to Early Adaptation (Epoch 9) ( $p = 0.345$ ,  $t = -1.32$ ) or Late Adaptation (Epoch 16) and Early Washout (Epoch 17) ( $p = 0.361$ ,  $t = 1.22$ ) or throughout the adaptation periods from Early Adaptation to Late Adaptation (Epoch 16) ( $p = 0.514$ ,  $t = 0.91$ ) or Early Washout to Late Washout (Epoch 24) ( $p = 0.533$ ,  $t = -0.85$ ). Although Figure 3.10 shows some increase in error at Early Adaptation, this change was not statistically significant. This result indicates that the gain perturbation had a minimal effect on distance error in the Control Group, possibly because they adapted within only a few trials of the first 10 post-perturbation trials in Epoch 9.

**Comparison of Late Baseline and Early Adaptation Epochs:** When comparing the Late Baseline (Epoch 8) to the Early Adaptation (Epoch 9), we found there was a significant difference in the distance error for the Isometric (no feedback) ( $p = 1.84 \times 10^{-5}$ ,  $t = -4.95$ ) and Isometric + Haptic Feedback ( $p = 2.08 \times 10^{-9}$ ,  $t = -7.26$ ) groups with similar effect sizes of 0.24 and 0.16, respectively. In Figure 3.10, we can see that the distance error increases for both groups, with the Isometric + Haptic Feedback group showing a larger increase in error. However, this difference between the two groups was not significant ( $p = 0.105$ ,  $t = -2.08$ ). **Adaptation:** To determine whether participants had adapted to the gain perturbation, we compared the distance error at Early Adaptation (Epoch 9) and Late Adaptation (Epoch 16), and between the Late

Baseline (Epoch 7) and Late Adaptation (Epoch 16). For both the Isometric (no feedback) ( $p = 1.38 \times 10^{-4}, t = 4.42$ ) and the Isometric + Haptic Feedback ( $p = 1.81 \times 10^{-6}, t = 5.55$ ), we found that there was significant decrease in the distance error, which indicates that participants were adapting their control in response to the perturbation and improving their accuracy. If the adaptation were complete, we would expect the distance error to decrease to a level similar to baseline performance. We found that there was no significant difference between the distance error at Late Baseline (Epoch 8) and Late Adaptation (Epoch 16) for both the Isometric (no feedback) ( $p = 0.725, t = -0.53$ ) and for the Isometric + Haptic Feedback ( $p = 0.203, t = -1.70$ ) groups. This result confirms that participants have successfully adapted to the gain perturbation.


**Comparison between Late Adaptation and Early Washout:** In a standard visuomotor perturbation experiment, after the perturbation is removed, we would expect the adaptation to persist, with the error sign reversed relative to early adaptation. In this case, returning from the increased gain to the initial gain would result in a negative error. For both the Isometric (no feedback) ( $p = 0.008, t = 3.18$ ) and for Isometric + Haptic Feedback ( $p = 5.30 \times 10^{-4}, t = 4.02$ ) we find a significant change in the distance error between Late Adaptation (Epoch 16) and Early Washout (Epoch 17). As shown in Figure 3.10, this significant change resulted in a negative distance, indicating that participants in both groups were now undershooting the target. This result confirms that adaptation was maintained in both groups even after the perturbation was removed.

**Washout:** During the washout phase, we expect participants to return to the baseline motor control strategy and performance. Once they have fully unlearned adaptation, we would expect their distance error at Late Washout (Epoch 24) to be similar to their Late Baseline (Epoch 8). In Figure 3.10, we can see a significant decrease in the magnitude of the distance error between Early Washout (Epoch 17) and Late Washout (Epoch 24) for the Isometric + Haptic Feedback group ( $p = 0.032, t = -2.64$ ), indicating that the participants are returning to their previous control methods. However, the error decreased for the Isometric (no feedback) group ( $p = 0.199, t = -1.71$ ), but the difference was not significant. To confirm if participants have fully unlearned the adaptation, we compare the Late Baseline with the Late Washout. We found that the difference in distance error between Late Baseline (Epoch 8) and Late Washout (Epoch 24) was not significant for Isometric (no feedback) ( $p = 0.514, t = -0.93$ ) and Isometric + Haptic Feedback ( $p = 0.829, t = -0.31$ ), indicating that the participants have returned to baseline performance.

**Pairwise Comparisons of No Feedback and Haptic Feedback:** While pairwise



comparisons of epochs within each group demonstrate that the control group did not need to adapt and that both Isometric groups were able to adapt to the perturbation during the experimental protocol, we compared the two feedback groups using pairwise comparisons within each epoch. Across all epochs examined, there was no significant difference in the distance between Isometric (no feedback) and Isometric + Haptic Feedback. In particular, during the Late Baseline ( $p = 0.725, t = -0.52$ ), Early Adaptation ( $p = 0.105, t = -2.08$ ), Late Adaptation ( $p = 0.312, t = -1.39$ ), Early Washout ( $p = 0.528, t = -0.87$ ), and Late Washout ( $p = 0.272, t = -1.49$ ) epochs, the two groups distance error was not significantly different. These results show that although the perturbation disrupted both Isometric groups similarly, the addition of haptic feedback did not produce a statistically detectable difference in distance error at any individual epoch.



**Figure 3.10.** Initial Finalized Bar plot for EMG Results

### 3.6 Discussion

In this study, we applied a visuomotor perturbation paradigm, typically used to study intrinsic human motor control, to a myoelectrically driven virtual elbow and demonstrated that participants can adapt to this perturbation. The effects of the perturbation differed between control of the real arm and control of the virtual elbow, where proprioceptive feedback was limited. We analyzed both average velocity and distance error to evaluate adaptation and to characterize performance differences across groups.

We found that participants in all groups significantly decreased their average velocity during the adaptation phase when the gain was increased, demonstrating that a systematic change in control strategies were made in response to the perturbation. Importantly, average velocity did not differ between groups, indicating that the observed group differences in distance error cannot be attributed to slower movements or to greater time available for feedback corrections.

The typical reaching speeds for healthy adults are on the order of 0.35 to 0.71m/s [86–88]. Our results showed that participants in all groups reduced their average velocity by approximately 0.05m/s, corresponding to 7 to 14% decrease. This magnitude of reduction is comparable to that reported in individuals with mild to moderate stroke, who often exhibit speed decreases of 0.04 to 0.22m/s during upper-limb reaching tasks [86–88]. The observed reduction in reaching speed demonstrates that participants adapted to the systematic gain perturbation when controlling the joint via EMG, confirming that users can modify their motor output even when the biological limb is not directly engaged.

Interestingly, the Control group did not show the expected transient increase in error following the introduction or removal of the perturbation. Participants in this group maintained consistent distance error across all epochs, with no significant differences between the baseline and adaptation epochs. One likely explanation is the Control participants—who had access to intact proprioceptive feedback and natural limb dynamics—adapted to the relatively simple gain perturbation extremely rapidly, such that any transient deviations occurred only in the first few trials and were therefore not captured by averaging across ten-trial epochs.

In contrast, both groups controlling the virtual prosthetic exhibited a clear and systematic adaptation response. Distance error increased significantly following the introduction of the gain perturbation, which then progressively decreased over the course of the adaptation phase, returning to near baseline levels by Late Adaptation. Pairwise comparisons confirmed significantly elevated error at Early Adaptation and no difference between Late Baseline and Late Adaptation, indicating that participants in both groups fully adapted to the perturbation. This pattern is consistent with classical visuomotor gain adaptation, where individuals recalibrate the mapping between motor commands and sensory consequences.

In typical visuomotor gain perturbation studies using natural reaching, initial endpoint errors are on the order of 0.001 to 0.03m [89–91]. In our study, the Control group had an increase in error of approximately 0.01 to 0.015m from Late Baseline to Early Adaptation, which falls within this expected range. In contrast, the error of the two

isometric groups far exceeded this range with an approximate increase in error of 0.03 to 0.04m, demonstrating that the perturbation imposed a more challenging and behaviorally consequential disruption under isometric control. Despite this increased difficulty, participants were still able to fully adapt within a single session, providing strong evidence that meaningful sensorimotor adaptation can occur even under constrained, prosthesis-like control conditions.

Haptic feedback did not influence performance, contrary to our original hypothesis. Although Figures 3.9 and 3.10 visually suggest worse performance in the Isometric + Haptic Feedback group, these differences were not statistically significant. This result aligns with a substantial body of literature on prosthesis control, which reports that adding haptic feedback does not reliably translate into measurable performance gains in short-term laboratory tasks. Studies investigating vibrotactile, force, or position feedback for upper-limb prostheses often find that users either do not benefit or require extensive training before meaningful improvements emerge. Several explanations have been proposed: users may rely primarily on vision when it is available; the bandwidth, timing, or modality of the haptic signal may not align well with the motor task; and added sensory channels can increase cognitive load, particularly when the feedback is not naturally interpretable or does not map intuitively onto the user’s internal model of the limb.

In the present study, the implementation of haptic feedback may have limited its utility. Although the tactile stimulator rotated with the virtual elbow to provide an intuitive cue, the feedback was always mapped to the real-time position of the virtual joint and was not scaled with the visuomotor gain perturbation similar to how natural proprioceptive feedback would not be scaled. As a result, the haptic signal did not convey task-relevant error information and may have introduced sensory incongruence rather than facilitating correction. Prior prosthesis-control work suggests that such incongruent feedback can reduce the utility of supplemental cues and may even interfere with adaptation. In our experiment, the haptic feedback may therefore have increased cognitive workload rather than assisting performance. This could explain why performance in the Isometric + Haptic Feedback group was not better than that of the Isometric (no feedback) group, despite our initial hypothesis.

Several methodological factors may also have contributed to the lack of observable error fluctuations in the Control group. Movements were self-paced, which may have allowed participants to prioritize accuracy; however, the average velocity did not differ between groups indicating that this effect is not the cause. The movement endpoint

was defined using a velocity threshold of 8% of peak velocity and remained that way for 100 ms. This velocity threshold was chosen to capture the multi-peaked velocity profiles of the isometric groups. Participants in the Control group often executed rapid corrective sub-movements following overshoot or undershoot, which frequently occurred within this 100-ms window and were therefore included in the endpoint definition. This likely reduced the measured distance error in the Control group and further masked any transient effects.

Together, these findings demonstrate that participants can robustly adapt to visuomotor gain perturbations when controlling a limb via myoelectric input, even under constrained sensory conditions. They also underscore that artificial sensory feedback must be carefully aligned to the task’s structure to be beneficial. Consistent with broader prosthesis research, simply adding haptic feedback is insufficient; effective substitutionary feedback must integrate with the user’s sensorimotor predictions to meaningfully support learning and control.

### **3.7 Conclusion**

This chapter examined how participants adapted to a visuomotor gain perturbation under three feedback conditions: real-arm control, isometric control, and isometric control with added haptic feedback. The results confirmed that the perturbation successfully elicited robust adaptation in both isometric groups, whereas the Control group appeared largely unaffected. Despite the successful adaptation, haptic feedback did not enhance performance relative to isometric control alone, likely due to the lack of gain-matching between the haptic and visual signals and the potential for increased cognitive load. These findings highlight the challenge of designing and implementing substitutionary feedback that remains aligned with user expectations and the task environment.

Future work may extend beyond simple gain perturbations to incorporate more ecologically valid disturbances—such as transient force fields, EMG-amplification errors, or shared-control switching—as well as more realistic environments that include visual distractions or cognitive load. Such conditions would provide deeper insight into how users adapt their control of remotely operated devices and prosthetic systems. Future studies may also investigate real-time adaptation of the feedback policy through reinforcement learning or personalized calibration, enabling the system to both improve user accuracy and continuously adjust to changes in the environment or available feedback channels.

# Chapter 4 |

## Using machine learning to generate an online haptic feedback policy.

Wearable haptic interfaces offer a promising channel for conveying information in human–machine interaction; however, designing effective feedback mappings remains an open challenge. Existing approaches often rely on fixed, handcrafted encodings, which may not generalize across users or remain effective under changing conditions. In practice, both user interpretation and device behavior can vary over time, creating a co-adaptive system in which static mappings may degrade over time.

A key limitation of current methods is their inability to adapt online while maintaining stable and interpretable feedback. Although data-driven approaches can calibrate mappings offline, they do not account for ongoing interaction or changes in the environment. This raises the question of whether a feedback policy can be learned and updated online in a way that remains stable for the user while accommodating individual differences and perturbations.

This chapter addresses this challenge by investigating a reinforcement learning framework for online adaptation of haptic feedback. A pre-trained neural network policy is used to initialize learning, and proximal policy optimization is applied to update the policy during human interaction. Participants perform a joystick-based task using vibrotactile directional cues, allowing the system to adapt the feedback mapping based on user performance. To evaluate both feasibility and robustness, the framework is tested under nominal conditions and following a controlled perturbation to the haptic interface.

## 4.1 Research Questions

The study is structured around three core research questions that guide the experimental design and interpretation of results:

- Can a reinforcement learning-based framework converge to stable, individualized haptic feedback policies during human-in-the-loop interaction?
- Can the reinforcement learning framework maintain or re-establish stable policies following a change in the environment?
- Do individualized haptic feedback policies lead to differences in task performance or perceived usability?

## 4.2 Introduction

Most existing vibrotactile paradigms employ fixed mappings (e.g., encoding error as vibration intensity). These schemes are easy to implement but yield inconsistent results: some studies find repulsive encodings more effective [3, 24], while others favor attractive mappings [2]. Prior work has also shown that haptic perception is influenced by multiple factors, including tactile suppression during motion [16, 25, 27, 28], masking between simultaneous stimuli [19], cognitive load [21, 22], and the intuitiveness of the encoding itself [23], which in practice, generates a significant amount of variability across users and contexts. This variability make it difficult to design a single fixed mapping that is consistently interpretable. While data-driven approaches such as sensitivity-adjusted regression models can reduce perceptual error [49, 92], they still assume static encodings, cannot adapt online, and fail to capture higher-level interpretation differences.

Parameterized perceptual models could be updated incrementally using gradient-based methods rather than relying on fixed encodings, but this would allow for only limited personalization as they are typically simple linear regression-based models. As a result, incremental updates are unlikely to capture complex co-adaptive behaviors that arise when both the user and the system change over time.


Reinforcement learning offers a natural framework for addressing nonlinear, co-adaptive interaction because it updates policies based on trial-by-trial evaluative feedback rather than requiring predefined supervisory targets for each stimulus–response pair. However, not all reinforcement learning methods are appropriate for human-in-the-loop

settings. While direct policy optimization and bandit-based methods are capable of responding to changes in the environment in real time, they permit unconstrained parameter updates. For example, policy gradient methods such as REINFORCE [93], Actor-Critic [94] and Asynchronous Advantage Actor-Critic [95] can perform online updates to their policies in real-time; however, they do not impose any constraints on the magnitude of policy updates, making them less suitable for human-in-the-loop systems where feedback stability is critical. Similarly bandit algorithms such as,  $\epsilon$ -greedy [96], Upper Confidence Bounds [97], and Bayesian bandits [98] update action-selection strategies online but allow abrupt shifts in behavior following individual reward observations. In human-in-the-loop systems, such unconstrained updates can cause sudden and unintelligible changes in feedback, destabilizing human learning and increasing response variability. The limitations imposed by unconstrained updates directly motivated the development of trust-region policy methods, which explicitly constrain how much a policy may change between updates in order to ensure stable and predictable behavior during learning.

Trust-region policy methods were developed to address instability caused by large policy updates by explicitly constraining how much a policy may change between iterations. Trust Region Policy Optimization (TRPO) [99], enforces a strict divergence constraint but relies on second-order optimization and line search procedures, introducing substantial computational overhead. Such overhead limits practicality for real-time human-in-the-loop learning. Proximal Policy Optimization (PPO) [100], approximates trust-region behavior using a clipped objective function, enabling stable and incremental policy updates while remaining computationally efficient. Therefore, we propose an online adaptation framework for haptic feedback using PPO. Specifically, we outline a method for online adaptation of haptic feedback using proximal policy optimization. We empirically demonstrate that this method can converge in a reasonable time frame to a unique policy for each user and can adapt to perturbations in the feedback mechanism. Such individualized adaptation may be particularly valuable in human-robot interaction scenarios where haptic feedback conveys robot state or navigation information to a human operator, including teleoperation systems and upper-limb prosthetic devices in which vibrotactile cues are used to communicate task-relevant information. In these contexts, user-specific perceptual differences can significantly influence feedback interpretability, motivating adaptive personalization strategies.

## 4.3 Methods

In this study, we developed and evaluated an online learning framework that adapts a haptic feedback mapping to individual users and perturbations using proximal policy optimization (PPO), as shown in Figure 4.1. To accomplish this, we used a pretraining step to learn an initial policy that represents a generic haptic feedback mapping. We defined a reward function for online training and evaluated the adaptation framework in a study with eleven participants. We assessed policy convergence, task performance, subjective usability, and the uniqueness of the learned policies across users and experimental phases.



**Figure 4.1.** Overview of the PPO-based training loop for the haptic feedback policy. At the beginning of each trial, the current policy receives a desired direction for the joystick and outputs motor commands corresponding to the vibration intensity. These commands are applied as haptic stimuli via the tactors to the participant’s arm. The participant interprets the haptic stimuli as instructions and moves the joystick toward the most intense vibration. Their response is recorded and compared to the desired direction to compute a reward value. This response, the original desired direction, and the current policy distribution are used to calculate the PPO loss and update the policy parameters. The process repeats across trials to refine the haptic feedback policy based on the user’s behavior.

### 4.3.1 Policy Pre-training

Human-in-the-loop trials are inherently costly and constrained, as each trial requires sustained user attention, active interpretation of haptic cues, and repeated motor responses. Extended sessions are further limited by mental fatigue, degradation of local sensitivity from repeated exposure, and learning effects, all of which can degrade the quality of



the data. To reduce the number of human-in-the-loop trials needed for convergence, we pre-trained a multilayer perceptron (MLP) using a simple mapping from the desired joystick direction to the intensities of the haptic feedback motors. Given that we have a four-dimensional continuous action space, and the requirement to have an interpretable and spatially coherent haptic patterns, successfully learning this encoding through uninformed exploration could take upwards of thousands of human interaction trials—far exceeding what can be reasonably be obtained from a participant in a single session.

The goal of this pre-training step was not to generate an initial model of human behavior and perception, but to encode a simple haptic feedback policy based on a linear mapping between vibration intensity and perceived stimuli location. Alternative surrogate datasets could include simulated user responses or data collected from users who participated in similar studies; however, both approaches would require the introduction of some assumptions on perception, interpretation strategies, motor responses that would be highly user-specific and not generalizable. Whereas the chosen synthetic dataset used here encodes only the geometrical relationship between which tactors would correspond to which direction without embedding any assumptions about user interpretations or learning. This setup allows the pre-training stage to generate a spatially coherent and simple interpretable mapping that can be adapted to each individual user.

The mapping assumes a linear relationship between the relative intensities of two adjacent tactors to generate a phantom sensation between the two. The MLP was initialized as a fully connected, feedforward neural network with one input node, two hidden layers of 64 nodes each, and four output nodes, using the ReLU activation function. The input was the desired joystick directions normalized to be between -1 and 1, and the four output nodes were the mean intensity for each of the four vibration motors, normalized from 0 to 1. We generated the synthetic training data set composed of 5000 data points by mapping the desired joystick direction to a generalized set of motor intensities,  $\mathbf{I}_d$ , using equation 4.1, where  $\phi_d$  corresponds to the desired directions of joystick movement.

Equation 4.1, is a simple linear mapping based on the phantom sensation haptic phenomenon [47], where each joystick direction  $\phi_d$ , is represented by activating either a

single tactor or a pair of tactors.

$$\mathbf{I}_d(\phi_d) = \begin{cases} [0.0, 0.0, 0.0, 0.0] & \phi_d = 0 \\ [1.0, 0.0, 0.0, 0.0] & \phi_d = 1 \\ [0.5, 0.5, 0.0, 0.0] & \phi_d = 2 \\ [0.0, 1.0, 0.0, 0.0] & \phi_d = 3 \\ [0.0, 0.5, 0.5, 0.0] & \phi_d = 4 \\ [0.0, 0.0, 1.0, 0.0] & \phi_d = 5 \\ [0.0, 0.0, 0.5, 0.5] & \phi_d = 6 \\ [0.0, 0.0, 0.0, 1.0] & \phi_d = 7 \\ [0.5, 0.0, 0.0, 0.5] & \phi_d = 8 \end{cases} \quad (4.1)$$

For initialization, the mapping is specified for nine discrete target directions; however, the underlying motor control space remains continuous, and intermediate directions can be represented through arbitrary intensity combinations. When a desired direction is aligned with a motor, the activation of that motor is maximum. For directions between motors, two adjacent motors are activated such that the level of activation is proportional to their distance from the desired direction. For example a direction halfway between two tactors would be displayed by 50% intensity of the two adjacent tactors to generate a phantom point of vibration between the two tactors. This equation reflects previous work in tactile perception that found that relative intensity across adjacent actuators is interpreted as a continuous spatial cue rather than two separate points and provides a simple baseline mapping between joystick direction and perceived haptic direction.

The MLP was trained on this data using the mean squared error loss between the predicted and actual motor intensities supplied by our mathematical mapping, as shown in equation 4.2, with a learning rate of 0.0003, 1000 epochs, and the ADAM optimizer.

$$Loss = \frac{1}{n} * \sum \mathbf{I}_{actual} - \mathbf{I}_d \quad (4.2)$$

Training of the MLP continued until either the MLP had finished 5000 epochs or until the error between the predicted motor intensities and the actual motor intensities was within a tolerance of  $1 \times 10^{-5}$ . After training was complete, the MLP's weights and biases were used to initialize the policy for online adaptation of the haptic feedback.

### 4.3.2 Online Learning Framework

We selected the clipped object variant of PPO [100], implemented using the Advantage Actor-Critic framework [101, 102], to explicitly limit how much the magnitude the policy can change during each update. This limitation ensures gradual, stable changes in haptic feedback that preserves interpretability while still allowing adaptation to user-specific perception and evolving task dynamics.

In this formulation, the interaction is modeled as an approximate Markov Decision Process. At each trial  $t$ , the policy receives the desired joystick direction,  $\phi_d$ , and outputs a continuous four-dimensional vector of motor intensities as the action. The true environment state, however, includes not only the externally specified direction but also the participant’s internal perceptual mapping between vibration patterns and interpreted direction, which is not directly observable and may evolve over time. As a result, the interaction is more accurately characterized as a partially observable and potentially nonstationary process. We nevertheless employ PPO due to its empirical robustness to partial observability and changing dynamics, and its ability to optimize perception–action coupling over batches of data rather than relying on locally linear, pointwise adaptation. The reward at trial  $t$  therefore depends on both the current action and prior policy updates that have influenced user interpretation.

The PPO policy was initialized with the weights and biases from the pre-trained MLP, and the PPO value function was initialized with randomly generated weights and biases using the Stable Baselines3 library [102]. To limit rapid deviation from the pre-trained policy, we set a small learning rate of 0.0001, which constrained the magnitude of the policy updates. We also initialized the logarithmic standard deviation matrix of the policy’s action distribution to values between -5 and -1; preliminary testing showed that these values preserved the pre-trained mapping during early training while allowing limited exploration. To encourage exploration around the pre-trained policy’s mean actions—without overwhelming the pretrained behavior—we set the entropy coefficient,  $c_2$ , to 0.01. Due to limited data, we set the batch size to 10 to ensure adequate sample efficiency and the number of epochs to 15 to balance training performance and the risk of overfitting. The value function coefficient,  $c_1$  was set to 0.0 and the clipping ratio,  $\varepsilon$ , to 0.2 based on the typical values used in previous PPO implementations [100, 103].

In our experimental setup, each trial yields a single user response, resulting in sparse human-generated rewards. We designed a piecewise reward function based on the

normalized absolute value of the wrapped direction error (equation 4.3).

$$\rho_\phi = \frac{\min(|\phi_d - \phi_u|, 8 - |\phi_d - \phi_u|)}{8} \quad (4.3)$$

In equation 4.4, when the absolute value of the normalized error is equal to zero, the policy receives a reward of +1 for that action, which helps increase the rate at which desirable actions are attributed to the correct joystick directions,  $\phi_d$ . For non-zero errors, the reward is set to  $-2$  times the error magnitude,  $\rho_\phi$ . This results in a penalty range of  $-1$  to  $0$ , which aligns with the normalized scales used for joystick input and motor output.


$$Reward = \begin{cases} +1 & \text{if } \rho_\phi = 0 \\ -2 * \rho_\phi & \text{if } \rho_\phi \neq 0 \end{cases} \quad (4.4)$$

### 4.3.3 Experimental Platform

We used a custom open-source haptic platform, the Snaptics system [84]. This low-cost open-source wearable device is designed to deliver tactile feedback to the skin. The platform consisted of two armbands worn on the forearm. The first armband contained four eccentric rotating mass (ERM) mini-motors (MicroDCMotors), each measuring 4 mm in diameter and 11 mm in length, and operating at 183 Hz under a 3V input. These motors were arranged equidistantly around the circumference of the forearm in a circular layout, similar to previous designs that leverage phantom sensations [31, 50, 51], as illustrated in Figure 4.2. The second armband housed the control electronics: two DRV8883 motor drivers and an Arduino Uno microcontroller. This microcontroller communicated with a host computer to synchronize haptic cues with collected kinematic data during the trials. Both armbands were secured to the participant’s right forearm using adjustable nylon straps. The four tactors were positioned at the cardinal directions around the forearm, and intermediate joystick directions were represented through relative intensity modulation across adjacent actuators.

We utilized an analog gaming joystick, similar to those in video game controllers sourced from Elegoo, as shown in Figure 4.2. The joystick contains two potentiometers that produce analog electrical signals corresponding to the joystick’s horizontal and vertical displacement. During trials, participants sat in a chair, held the joystick in their right hand, and moved the joystick with their thumb. An Arduino Mega was used to read the two analog signals and transmit them to the host computer, where the joystick’s

directional input and displacement from center were computed.



**Figure 4.2.** The experimental platform consisted of a haptic armband, an analog joystick, and a host computer. The haptic platform featured four eccentric rotating mass (ERM) mini-motors arranged in a circular array around the participant’s forearm, aligned with the four cardinal directions (Up, Down, Left, and Right). These motors were controlled by an Arduino Uno connected to the host computer. The analog joystick was a two-axis joystick from Elegoo, which measured thumb movements via two potentiometers mounted on a gimbal mechanism. The joystick signals were read by an Arduino Mega and transmitted to the host computer. During each trial, the host computer sent vibration intensity commands to the Arduino Uno, prompting the motors to activate. Participants interpreted the resulting haptic cues as directional instructions and moved the joystick accordingly. The Arduino Mega then relayed the joystick position back to the host computer for evaluation.

#### 4.3.4 Experimental Protocol

Eleven participants (six males and four females aged 18-59) were recruited for this study and provided informed written consent before participation. The protocol for this research study was reviewed and approved by the Penn State Institutional Review Board under STUDY00026482.

Each participant completed a single experimental session consisting of two sequential 270-trial blocks. Each block included 30 repetitions of the nine possible directions. Upon

enrollment, participants are assigned a randomized lists of joystick directions each was generated using a random number generator to ensure even sampling across directions.

Prior to the online learning session, participants completed 10 practice trials using the pre-trained policy to familiarize themselves with the system. Immediately following these practice trials, participants completed a subjective usability questionnaire assessing their experience with the baseline feedback condition. The questionnaire included items from the System Usability Scale (SUS) [104], along with three open-ended questions. The open-ended prompts asked participants: (1) whether they preferred any particular version or presentation of the feedback, (2) what, if anything, they liked about the feedback, and (3) what, if anything, they disliked about the feedback.

After completion of the baseline survey, we initialized the PPO policy using a pre-trained multilayer perceptron. During each trial, participants were instructed to move the joystick in the direction indicated by the haptic feedback platform, corresponding to where they felt the most intense vibration. After each trial, we passed the collected data to the proximal policy optimization algorithm to evaluate the performance of the current haptic feedback mapping and optimize the mapping by performing a policy update.

Following the completion of the first set of 270 trials, participants were presented with the same subjective usability questionnaire, evaluating the adapted, user-specific policy.

A mechanical perturbation was then introduced by rotating the haptic armband by 2 cm counterclockwise relative to its original position on the forearm, thereby altering the spatial alignment between motor location and perceived direction. Participants were instructed to continue moving the joystick in the direction indicated by the haptic platform (toward the strongest perceived vibration). Participants then completed the second set of 270 trials with the perturbed configuration. Following the perturbation block (post-perturbation phase), participants completed the SUS questionnaire for a third time.

### 4.3.5 Data Analysis

Policy convergence was assessed separately for the post-personalization and post perturbation blocks. For each participant and each 270-trial block, the policy loss rate was computed by applying a moving-average slope method to the policy loss and fitting a linear model to the last 10 trials. A policy was considered converged if the absolute value of this slope was below 0.005 or if the change was less than 0.5%, indicating that the loss had stabilized.

Policy performance was quantified using directional error, defined as the absolute difference between the instructed direction and the participant’s joystick response (wrapped to the minimum circular distance). Directional error was averaged within sliding windows across trials to assess trends during learning. For each participant, a linear regression was performed on directional error over trials within each 270-trial block (post-personalization and post-perturbation). The slope of the regression and corresponding p-value were used to evaluate whether performance changed significantly over time. A negative slope was interpreted as a reduction of directional error, indicating that the adapted policy improved the interpretability of the haptic encoding.

Subjective usability was quantified using the System Usability Survey (SUS) [104]. Individual item responses were scored according to standard SUS procedures and converted to a composite score ranging from 0–100. Mean SUS scores were computed for three conditions: baseline (pre-trained policy), post-personalization, and post-perturbation. Paired t-tests were conducted to evaluate whether the learned, user-specific policy differed from the baseline policy in perceived usability and whether the usability changed following the introduction of the perturbation. Comparisons were performed between baseline and post-personalization, baseline and post-perturbation and post-personalization and post-perturbation conditions. P-values were corrected for multiple comparisons using the Benjamini-Hochberg [53] false discovery rate method. Open-ended responses were analyzed using qualitative thematic analysis. Responses were reviewed and coded to identify recurring themes related to clarity, comfort, interpretability, and perceived differences between feedback versions.

A policy uniqueness metric was quantified by computing the Jensen-Shannon divergence (JSD) [105] between the action space distributions of pairs of policies. For each actuator, we computed the JSD between the two policies’ action distributions at each of the nine joystick directions. These nine JSD values were then averaged to obtain that actuator’s Average JSD. This procedure was repeated for all four actuators. The overall Average JSD for a participant was computed as the mean of the four actuator-specific Average JSD values.

This measure reflects how distinct or unique one policy is relative to another. In particular, when computed between a participant’s learned policy and the pre-trained policy, the Average JSD quantifies the extent to which the learned policy has diverged from the baseline mapping. Within each participant, JSD comparisons were performed between (1) the pre-trained policy and the post-personalization policy, (2) the pre-trained policy and the post-perturbation policy, and (3) the post-personalization and the post-perturbation

policies. These comparisons quantify both the degree of personalization relative to the baseline mapping and the magnitude of adaptation following the perturbation.

To evaluate whether divergence from the pre-trained policy differed across participants, a one-way analysis of variance (ANOVA) was performed with the participant as the factor. A significant result would indicate that the degree to which participant-specific policies diverged from the pre-trained policy varied by participant.

To further examine participant-specific differences, pairwise Student’s paired t-tests were performed on the Average JSD values between each possible pairing of participants. The null hypothesis for each comparison was that there was no significant difference in the the degree of divergence between two participants’ learned policies. Rejecting the null hypothesis would indicate that participants developed policies that diverged from the pre-trained policy in ways specific to each participant, supporting the idea that policy learning was individualized.

The same computational procedure was used to compute pairwise JSD values between participants’ post-personalization policies and separately between participants’ post-perturbation policies, allowing for comparison of inter-participant similarity. All post hoc corrections were performed with the Benjamini-Hochberg [53] false discovery rate method to adjust the p-values.

## 4.4 Results

### 4.4.1 Policy Convergence

To evaluate the convergence of both the baseline and post-perturbation participant policies in the PPO training process, we analyzed the loss rate of the policy loss over the last 10 trials for each participant, as shown in Table 4.1. For the post-personalization policies, the final loss rates for participants 1-4 and 6-11 ranged from 0.0001 to 0.0036, all below the threshold, suggesting that the policies have successfully converged to a solution. However, for participant 5, the final loss rate was 0.0071, which is above the threshold, indicating significant variability between updates and that the policy has not fully converged. For the post-perturbation policies, the final loss rates for participants 1-11 were all below 0.005, suggesting that these policies have converged to a solution.



**Table 4.1.** Absolute value of the final policy loss slopes computed over the last 10 trials for each participant in both the post-personalization and post-perturbation conditions. Slopes with absolute value below 0.005 indicate convergence. The loss rates for the policies of participants 1-4, and 6-11 were below the threshold, indicating successful convergence, and the loss rate for participant 5 was above the threshold indicating the policy had not fully converged.

| Participant | Post-Personalization<br>Final Loss Slope | Post-Perturbation Final<br>Loss Slope |
|-------------|------------------------------------------|---------------------------------------|
| 1           | 0.0021                                   | 0.0016                                |
| 2           | 0.0031                                   | 0.0012                                |
| 3           | 0.0023                                   | 0.0041                                |
| 4           | 0.0001                                   | 0.0004                                |
| 5           | 0.0071                                   | 0.0002                                |
| 6           | 0.0024                                   | 0.0038                                |
| 7           | 0.0005                                   | 0.0017                                |
| 8           | 0.0036                                   | 0.0003                                |
| 9           | 0.0023                                   | 0.0017                                |
| 10          | 0.0006                                   | 0.0003                                |
| 11          | 0.0011                                   | 0.0040                                |

#### 4.4.2 Objective and Subjective Policy Performance

Directional error was analyzed using overlapping sliding window averages across trials to assess changes in task performance over time. For each participant, a linear regression was performed on directional error across trials within each 270-trial block (post-personalization and post-perturbation).

For the post-personalization policies, initial mean error ranged from 0.4 to 2.8 across participants, while final mean error ranged from 0.2 to 1.8. Although several participants exhibited reductions in directional error, the regression analysis found that the reductions were not statistically significant, indicating that performance remained stable over the course of training. However, participant 7 had significant positive slope ( $p = 2.00 \times 10^{-5}$ , indicating an increase in direction error).

For the post-perturbation policies, initial window means ranged from 0.4 to 2.2 and final window means ranged from 0.2 to 1.6 across participants. Regression analysis identified significantly significant changes for two participants. Participant 7 demonstrated a significant negative slope ( $p = 0.035$ ), indicating a reduction in directional error, while participant 4 had a significant positive slope ( $p = 7.21 \times 10^{-4}$ ), indicating an increase in directional error. For the remaining participants, no statistically significant changes were observed, indicating stable performance across the training period.

Overall, directional error remained largely stable throughout both the post-personalization

and post-perturbation training periods, with only a small number of participants exhibiting individual changes.

Survey responses were analyzed to assess perceived usability across conditions. Mean SUS scores were  $\mu = 76.00$   $SD = 13.55$  at baseline,  $\mu = 70.50$   $SD = 16.24$  following personalization, and  $\mu = 72.78$   $SD = 18.24$  following the perturbation condition. Pairwise comparisons using paired t-tests with Benjamini–Hochberg false discovery rate corrected p-values, revealed no significant differences between conditions, Baseline vs. Post-Personalization  $t = 2.31$ ,  $p = 0.139$ , Baseline vs. Post-Personalization  $t = 2.31$ ,  $p = 0.139$ , and Baseline vs. Post-Perturbation  $t = 1.18$ ,  $p = 0.393$ . These results suggest that perceived system usability did not significantly differ across the baseline, post-personalization, and post-perturbation conditions.

#### 4.4.3 Comparison with Pre-trained Baseline Policy

Table 4.2 summarizes the Jensen–Shannon Divergence (JSD) between each participant’s learned post-personalization policy and our pretrained policy and shows that the divergence between participants’ post-personalization policies and the pretrained policy varied widely across individuals. Figure 4.3 shows the overall average JSD computed across all directions and all actuators, while Figure 4.4 shows the average JSD computed across all directions for each actuator. As shown in Figure 4.3, participants 1 and 2 showed the smallest deviations (0.1336 and 0.1318), while participants 7 and 8 diverged the most from the pre-trained policy (0.2926 and 0.2781). The JSD values met the normality assumption ( $W = 0.954$ ,  $p = 0.077$ ) and the homogeneity-of-variance assumption ( $F = 0.488$ ,  $p = 0.886$ ), allowing a one-way ANOVA with participant as the factor. The analysis shows that the effect of participant ( $F = 0.884$ ,  $p = 0.557$ ) was not statistically significant. These results indicate that although the learned policies diverged from the pretrained initialization, the magnitude of divergence did not significantly differ across participants.

Post-perturbation results (Table 4.3) showed a similar pattern. Average JSD values ranged from 0.1639 to 0.3753 across participants. Figure 4.5 illustrates the Jensen–Shannon Divergence computed for each actuator, highlighting how the divergence between the learned and pretrained policies varied across motors for each participant. As shown in Figure 4.3, average JSD values ranged from 0.1639 to 0.3753 across participants. Participants 2 and 4 showed the smallest deviations (0.1931 and 0.1639), and participants 9 and 11 showed the largest deviation (0.3238 and 0.3753), as shown in Figure 4.3. The Shapiro–Wilk test indicated a deviation from normality ( $W = 0.927$ ,  $p = 0.008$ , while


Levene’s test confirmed homogeneity of variance ( $F = 0.7034$ ,  $p = 0.714$ ). As ANOVA is generally robust to modest violations of normality, we proceeded with a one-way ANOVA using participant as the factor. The ANOVA results show that the effect of participant ( $F = 0.621$ ,  $p = 0.785$ ) was not statistically significant. These results suggest that even after the perturbation manipulation, the magnitude of policy divergence remained broadly similar across participants.

**Table 4.2.** Average Jensen-Shannon Divergence between each participant’s post-personalization policy and the pre-trained policy, computed per motor and averaged across all joystick directions.


| Participant | Motor 1 | Motor 2 | Motor 3 | Motor 4 | Avg. JSD |
|-------------|---------|---------|---------|---------|----------|
| 1           | 0.1406  | 0.0616  | 0.2468  | 0.0854  | 0.1336   |
| 2           | 0.1430  | 0.0926  | 0.0524  | 0.2391  | 0.1318   |
| 3           | 0.4569  | 0.2779  | 0.0648  | 0.1920  | 0.2479   |
| 4           | 0.1216  | 0.2639  | 0.2565  | 0.2051  | 0.2117   |
| 5           | 0.0817  | 0.2887  | 0.2987  | 0.1107  | 0.1950   |
| 6           | 0.1222  | 0.2358  | 0.0782  | 0.4286  | 0.2162   |
| 7           | 0.3042  | 0.3334  | 0.1678  | 0.3651  | 0.2926   |
| 8           | 0.3774  | 0.1534  | 0.3069  | 0.2748  | 0.2781   |
| 9           | 0.1510  | 0.1837  | 0.1981  | 0.4849  | 0.2544   |
| 10          | 0.2640  | 0.0810  | 0.2727  | 0.2005  | 0.2045   |
| 11          | 0.2446  | 0.1468  | 0.0663  | 0.2634  | 0.1803   |

**Table 4.3.** Average Jensen-Shannon Divergence between each participant’s post-perturbation policy and the pre-trained policy, computed per motor and averaged across all joystick directions.

| Participant | Motor 1 | Motor 2 | Motor 3 | Motor 4 | Avg. JSD |
|-------------|---------|---------|---------|---------|----------|
| 1           | 0.0753  | 0.2040  | 0.5946  | 0.1274  | 0.2503   |
| 2           | 0.2728  | 0.1623  | 0.0935  | 0.2438  | 0.1931   |
| 3           | 0.5172  | 0.0526  | 0.1098  | 0.2243  | 0.2260   |
| 4           | 0.0803  | 0.0794  | 0.3576  | 0.1384  | 0.1639   |
| 5           | 0.1775  | 0.3577  | 0.3353  | 0.0917  | 0.2406   |
| 6           | 0.1637  | 0.3441  | 0.0743  | 0.2405  | 0.2056   |
| 7           | 0.1783  | 0.4305  | 0.4705  | 0.1490  | 0.3071   |
| 8           | 0.3821  | 0.1785  | 0.2339  | 0.3134  | 0.2770   |
| 9           | 0.2022  | 0.2048  | 0.3234  | 0.5647  | 0.3238   |
| 10          | 0.3601  | 0.2010  | 0.3491  | 0.2127  | 0.2807   |
| 11          | 0.2702  | 0.6038  | 0.1175  | 0.5099  | 0.3753   |




**Figure 4.3.** Overall Average Jensen Shannon Divergence computed across all directions and actuators between each participant’s post-personalization policy and post-perturbation policy with the pre-trained policy



**Figure 4.4.** Average Jensen-Shannon Divergence computed per actuator between each participant’s post-personalization policy and the pre-trained policy, computed per motor and averaged across all joystick directions.

#### 4.4.4 Comparison Between Participants’ Post-personalization and Post-perturbation Policies

Table 4.4 summarizes the Jensen–Shannon divergence (JSD) between each participant’s learned post-personalization policy and their learned post-perturbation policy. As with the comparisons to the pre-trained baseline, we first evaluated the distributional assumptions of the JSD values. A Shapiro–Wilk test indicated a deviation from normality ( $W = 0.933$ ,  $p = 0.014$ ), while Levene’s test confirmed homogeneity of variance ( $F =$



**Figure 4.5.** Average Jensen-Shannon Divergence computed per actuator between each participant’s post-perturbation policy and the pre-trained policy, computed per motor and averaged across all joystick directions.

0.654,  $p = 0.757$ ). Because the groups met the equal-variance assumption and ANOVA is generally robust to modest violations of normality, we proceeded with a one-way ANOVA using participant as the factor. The ANOVA results show that the effect of participant ( $F = 1.852$ ,  $p = 0.090$ ) is not significant, indicating that participants’ post-perturbation policies diverged from their baseline policies in a participant-specific manner.

**Table 4.4.** Average Jensen-Shannon Divergence (JSD) between each participant’s post-perturbation policy and their post-personalization policy, computed per motor and averaged across all joystick directions.

| Participant | Motor 1 | Motor 2 | Motor 3 | Motor 4 | Avg. JSD      |
|-------------|---------|---------|---------|---------|---------------|
| 1           | 0.1719  | 0.1941  | 0.3978  | 0.1280  | <b>0.2229</b> |
| 2           | 0.1621  | 0.0774  | 0.1443  | 0.0594  | <b>0.1108</b> |
| 3           | 0.2043  | 0.3105  | 0.1626  | 0.1625  | <b>0.2100</b> |
| 4           | 0.0683  | 0.2395  | 0.1268  | 0.0892  | <b>0.1309</b> |
| 5           | 0.2229  | 0.0845  | 0.0729  | 0.1022  | <b>0.1206</b> |
| 6           | 0.2041  | 0.3026  | 0.1367  | 0.3217  | <b>0.2413</b> |
| 7           | 0.2039  | 0.2128  | 0.3489  | 0.3381  | <b>0.2759</b> |
| 8           | 0.2581  | 0.0813  | 0.2472  | 0.1179  | <b>0.1761</b> |
| 9           | 0.2313  | 0.0358  | 0.1457  | 0.1971  | <b>0.1525</b> |
| 10          | 0.1565  | 0.2499  | 0.1369  | 0.1111  | <b>0.1636</b> |
| 11          | 0.2993  | 0.5473  | 0.0727  | 0.3289  | <b>0.3120</b> |


#### 4.4.5 Comparison between participants

Though the pairwise t-tests demonstrate differences in divergence from the pre-trained baseline, they do not capture divergence between individual policies. Figures 4.6 summarize the JSD between every pair of participants’ learned baseline policies. Average JSD values varied across pairs: participants 6 and 9 ( $AverageJSD = 0.1294$ ) showed minimal differences in policy behavior, whereas participants 7 and 8 ( $AverageJSD = 0.4492$ ) showed relatively large differences. Figure 4.7 shows the same pairwise comparison for the post-perturbation policies. Similar trends appeared: some participant pairs remained highly similar—for example, participants 3 and 10 ( $AverageJSD = 0.2191$ ) and participants 1 and 5 ( $AverageJSD = 0.2292$ ) while others, such as participants 9 and 11 ( $AverageJSD = 0.5167$ ) and participants 7 and 11 ( $AverageJSD = 0.5061$ ), showed significant differences. These results indicate that both participants’ post-personalization and post-perturbation policies varied between individuals, reflecting differences in how each person interpreted the pre-trained feedback, learned from experience, and adapted to perturbations. The observed variability highlights the importance of individualized policy assessment and suggests that even under the same training conditions, users can develop unique strategies for interpreting haptic feedback. The full list of pairwise comparisons is provided in Appendix B.

### 4.5 Discussion

Prior research shows that user performance differs even for simple motions when different haptic encodings are used. For instance, the effectiveness of haptic cues is sometimes better when repulsive cues are used [3,25], while other studies demonstrate attractive [2] cues are better. This may be attributed to the fact that some users prefer one mapping over the other [40]. However, many haptic feedback systems rely on fixed, hand-crafted mappings, limiting their ability to adapt to different users and environments. Our results demonstrate that proximal policy optimization can learn haptic mappings that adapt to individual users.


The policies for most participants successfully converged, indicating that this approach is feasible in a human-in-the-loop setting. Specifically, participant 5’s policy did not converge during the personalization phase but did converge after the perturbation phase. This lack of convergence could be due to a variety of factors, such as initial overfitting, extremely low or high variability in human response, the selection of hyperparameters,



**Figure 4.6.** Average Jensen–Shannon Divergence (JSD) between each pair of participants’ learned baseline policies across all four motors. Each cell in the matrix represents the overall JSD for the corresponding participant pair, with lighter colors indicating lower divergence and darker colors indicating higher divergence.

or insufficient trials. These results indicate that policy convergence is also dependent on individual factors. While our selection of hyperparameters and training time worked for most participants, some users may require different parameters or additional training. Successful online learning also depended on the availability of a reasonable pretrained initialization, which provided a stable starting point for adaptation. These findings suggest that while reinforcement learning can personalize haptic mappings online, careful initialization and parameter selection remain critical for reliable deployment.

Although policy convergence was observed among most participants, this adaptation did not result in statistically significant changes in directional error during either the personalization or perturbation phases for most participants. Directional error remained stable across trials, and subjective usability scores did not differ significantly across baseline, post-personalization, and post-perturbation conditions. A small subset of participants exhibited significant changes in error following perturbation, with some



**Figure 4.7.** Average Jensen–Shannon Divergence (JSD) between each pair of participants’ learned post-perturbation policies across all four motors. Each cell in the matrix represents the overall JSD for the corresponding participant pair, with lighter colors indicating lower divergence and darker colors indicating higher divergence.

participants showing improvement, while one participant exhibited an increase in error. These results show that personalization and adaptation of the policy did not degrade task performance or perceived usability, nor did it produce measurable improvements within the scope of this task. The simplistic nature of the task may have had a ceiling effect on performance improvements with personalization. Additionally, participants operated in a controlled setting with explicit knowledge of the feedback structure. In more ecological scenarios—such as naive users, more complex control spaces, prolonged exposure, or less intuitive feedback mappings—greater performance gains may emerge, particularly in situations where perceptual ambiguity is higher or users have not been preconditioned to interpret the feedback.

Wearable haptic interfaces are subject to a range of real-world factors, including device placement shifts, variations in skin contact, and changes in muscle activity, all of which can alter how feedback is perceived. The perturbation condition was designed to emulate



such changes. The largely stable performance observed across participants suggests that the learned policies maintained stable performance under moderate changes in feedback conditions. However, the variability observed in the subset of participants who exhibited significant changes indicates that this robustness is not uniform. Instead, adaptation to altered feedback appears to occur in a participant-specific manner, reinforcing the need for systems that can accommodate individual differences under dynamic conditions.

When comparing each participant’s learned policy to the pretrained initialization, no statistically significant differences were found in the magnitude of divergence across participants in either the post-personalization or post-perturbation phases. Although numerical differences in average JSD were observed, the ANOVA and pairwise tests indicated that the overall degree of deviation from the pretrained policy was comparable across individuals during both phases. This shows that, in terms of overall deviation from initialization, participants required similar magnitudes of adaptation, even though the resulting policies differed in structure. This consistency in divergence magnitude, combined with the absence of significant changes in directional error and usability scores, suggests that similar levels of adaptation can produce comparable performance and perceived usability across participants.

To examine whether the learned policies differed across participants, we computed pairwise JSD values between each participant’s personalized policy. While divergence from the pretrained policy was statistically comparable across participants, this measure does not capture differences between their personalized policies. Our results for the pairwise JSD revealed that the personalized policies developed unique structures. For example, participants 1 and 2 exhibited a low pairwise JSD value, indicating that their learned policies were comparable. In contrast, participants 4 and 5 showed a substantially larger pairwise JSD, suggesting that their learned mappings differed considerably. These pairwise differences confirm that the learned policies reflected individualized adaptations rather than uniform deviations from the pretrained initialization. The comparison between post-personalization and post-perturbation policies further indicates that perturbation did not drive convergence toward a common mapping; instead, individualized policy structures were preserved. Although these individualized adaptations did not correspond to measurable improvements in task accuracy, and were not reflected in usability scores, their presence indicates that participants did not converge to a single shared mapping; instead, the reinforcement learning framework adapted the feedback policy in a participant-specific manner. This results indicates that multiple distinct policy structures can yield equivalent task performance and user-perceived usability, which highlights the potential

value of personalized adaptation in haptic feedback systems, particularly for more complex tasks or feedback mappings that are less intuitive.

Generalized haptic feedback policies are attractive from a deployment perspective, as they require minimal calibration and exhibit predictable behavior across users. However, such fixed mappings may not account for individual perceptual differences. In contrast, individualized policies can adapt to each user’s interpretation of the cues, potentially improving robustness to perceptual variability. This personalization introduces additional complexity and requires careful initialization to avoid overfitting. The variability observed between participants in this study highlights the extent to which perceptual strategies differ, and motivates the use of reinforcement learning and data-driven approaches to capture these differences within the learning process.

These results show that participants did not converge to a single shared strategy, even when exposed to the same perturbation, and that the reinforcement learning framework maintained distinct user-specific mappings across conditions. Although these differences did not correspond to measurable changes in directional error or usability scores, the persistence of individualized policies suggests that multiple mappings can support similar task performance. This indicates that the learning process captures differences in how users interpret and respond to the feedback, even when those differences are not reflected in aggregate performance metrics.

## 4.6 Conclusion

This chapter investigates the use of proximal policy optimization to implement a framework for the online adaptation of a haptic feedback policy. Feedback was conveyed via an array of 4 vibrotactile motors indicating the desired direction of joystick movement. Our results demonstrate that proximal policy optimization can be implemented in a human-in-the-loop wearable system and can produce individualized feedback policies across users without degrading task performance or perceived usability. Although, measurable improvements in directional error were not observed, the learned policies converged for most participants and remained stable after the introduction of a perturbation to the haptic platform.

Our results showed that proximal policy optimization produced stable, individualized haptic feedback policies that persisted across both the personalization and perturbation phases without degrading task performance or perceived usability. These results also showed that the final policy structures differed between individuals, indicating that users

develop distinct mappings that persist even under changing conditions. The consistency of directional error and usability scores across conditions further suggests that multiple distinct haptic mappings can yield equivalent behavioral and subjective outcomes.

Future work may begin by examining how different pretrained policy initializations, such as attractive versus repulsive mappings, influence learning dynamics and final policy structure. Building on this, future studies can incorporate user-specific factors, including prior exposure, learning dynamics, and subjective preferences such as perceived comfort or intuitiveness, to evaluate policy quality beyond objective performance alone.

Extending this framework to a more demanding and meaningful task is a necessary next step. Rather than relying on discrete directional responses, such tasks would require continuous control and coordination across multiple stages of movement. A virtual prosthetic manipulation task—where users control a simulated hand to reach, grasp, and move objects—provides one such example. This type of task introduces sequential decision-making and sustained control demands, allowing performance to be evaluated through task completion time, grasp success rate, and control stability. These conditions provide a more direct test of how policy initialization, user-specific factors, and learned feedback mappings influence behavior as task complexity increases.

# Chapter 5 |

## Conclusions and Future Work

This thesis investigated how human perception, motor adaptation, and learning-based methods shape the design of wearable haptic systems. Across these studies, a consistent observation emerged: the effectiveness of haptic feedback is not determined by actuator fidelity or the presence of feedback, but by how well the feedback aligns with user-specific perceptual characteristics and the structure of the task. As a result, this work challenges the common assumption that a single, fixed mapping can effectively convey information across users and conditions.

In Chapter 2, we investigated the role of voluntary motion and tactile perception in haptic perception. While tactile suppression is often cited as a primary factor inhibiting the perception of haptic feedback during motion, the results demonstrated that it was not the dominant source of error in this context. Instead, stimulus location consistently accounted for the largest variation in error, regardless of movement condition, indicating that spatial sensitivity is the primary constraint. This finding shifts the design focus away from compensating for motion dependent effects and towards accounting for spatial variability and individual differences. The use of sensitivity-adjusted, data-driven perceptual models further demonstrated that these differences can be systematically captured and used to reduce decoding error, supporting the need for personalized feedback encodings.

In Chapter 3, we examined adaptation to a visuomotor perturbation using a myoelectrically controlled system and evaluated whether substitutionary haptic feedback improved performance. Participants were able to reduce distance error over time despite the absence of natural proprioceptive feedback, demonstrating that our adaptation mechanisms for visuomotor perturbations remain active even under indirect control. However, the addition of haptic feedback did not improve performance, adaptation quality, or the speed of adaptation, particularly during early exposure to the perturbation. This lack of

impact was attributed to the mismatch between feedback signal and the imposed gain, indicating that haptic feedback must be aligned with the underlying task dynamics and the user’s internal model to be effective. These results demonstrate that while users can adapt to altered control dynamics, this adaptation does not overcome feedback that is misaligned with the task.

In Chapter 4, we extended this framework by introducing reinforcement learning to adapt haptic feedback policies in a human-in-the-loop system. The results demonstrated that proximal policy optimization could converge to stable, individualized policies within a reasonable time frame and maintain performance following a perturbation to the haptic platform. Policy adaptation did not lead to improvements in task performance or perceived utility. Instead, the policies diverged between participants while still producing comparable behavioral outcomes, indicating that multiple feedback mappings can exist for the same task. This suggests that while optimal encodings may exist, they are likely user-specific and task-dependent, particularly as task complexity increases.

Taken together, these findings support a shift in how information is encoded in wearable haptic systems. Rather than optimizing a fixed mapping and expecting users to adapt, perception and interpretation should be treated as inherently user-dependent and continuously evolving. To achieve effective haptic feedback, encoding strategies must account for spatial sensitivity, individual differences, and task dynamics. Adaptive, data-driven methods provide a practical approach for capturing these factors and personalizing the feedback to maintain interpretability.

The work presented in this thesis was conducted under controlled laboratory conditions with simplified tasks and low-dimensional feedback in order to isolate key aspects of perception, adaptation and information encoding. These constraints allowed for systematic evaluation of spatial sensitivity, adaptation under altered control dynamics, and the feasibility of learning personalized feedback mappings. However, these constraints also limit the extent to which the results generalize to real-world applications. The reinforcement learning framework was reactive, adapting feedback only after the fact. It did not incorporate task context or predictive elements and the adaptation was conducted over a relatively short amount of time. As a result, these studies do not fully capture the influence of long-term interaction, user state, or environmental variability on perception and performance.

Future work can extend this research by relaxing the simplifying constraints used in this thesis and moving toward more complex, realistic interaction scenarios. Evaluating haptic feedback strategies under ecologically valid conditions—such as transient distur-

bances, changing control mappings, and environments with increased cognitive or sensory load—would provide a more complete understanding of how users interpret and adapt to feedback in systems that more closely resemble real-world prosthetic or teleoperation tasks.

A more challenging advancement lies in improving the formulation of the learning problem itself. Developing more informative reward functions, more efficient data collection methods, and policy representations that can adapt to changes in user perception or feedback presentation, would enable a system to move beyond static personalization and towards continuous, user-aware adaptations. This shift would allow for the feedback to evolve alongside the user, rather than converging to a single fixed mapping.

Taken together, while the work presented here was conducted under simplified conditions, it represents an important step toward understanding haptic perception and adaptation in more complex, dynamic scenarios, providing a foundation for future studies that examine richer, multi-stage tasks and user-specific feedback strategies.

# **Appendix A |**

## **Visuomotor Adaptation Pairwise T-test Comparison Tables**

Table A.1. Pairwise T-Test comparisons with FDR-adjusted p-values

| Group 1 | Phase 1    | Group 2 | Phase 2    | Diff.    | SE       | t     | Raw p          | 95% CI Lower | 95% CI Upper | FDR Adj p            |
|---------|------------|---------|------------|----------|----------|-------|----------------|--------------|--------------|----------------------|
| RealArm | LateBL     | RealArm | EarlyAdapt | -0.00944 | 0.007157 | -1.32 | 0.1904         | -0.02364     | 0.004764     | 0.344657374          |
| RealArm | LateBL     | RealArm | LateAdapt  | -0.00295 | 0.007157 | -0.41 | 0.6808         | -0.01715     | 0.011247     | 0.768699833          |
| RealArm | LateBL     | RealArm | EarlyWash  | 0.00580  | 0.007157 | 0.81  | 0.4201         | -0.00841     | 0.019994     | 0.558319528          |
| RealArm | LateBL     | RealArm | LateWash   | -0.00031 | 0.007157 | -0.04 | 0.9655         | -0.01451     | 0.009889     | 0.974813228          |
| RealArm | LateBL     | EMG     | LateBL     | -0.00470 | 0.009303 | -0.51 | 0.6147         | -0.02323     | 0.013830     | 0.725247116          |
| RealArm | LateBL     | EMG     | EarlyAdapt | -0.04015 | 0.009303 | -4.32 | 0.000028       | -0.05869     | -0.02162     | <b>0.000251637</b>   |
| RealArm | LateBL     | EMG     | LateAdapt  | -0.00850 | 0.009303 | -0.91 | 0.3636         | -0.02704     | 0.010029     | 0.514305034          |
| RealArm | LateBL     | EMG     | EarlyWash  | 0.01423  | 0.009303 | 1.53  | 0.1303         | -0.00430     | 0.032763     | 0.258099678          |
| RealArm | LateBL     | EMG     | LateWash   | 0.00195  | 0.009303 | 0.21  | 0.8347         | -0.01658     | 0.020481     | 0.885266898          |
| RealArm | LateBL     | Haptic  | LateBL     | -0.00947 | 0.009094 | -1.04 | 0.3011         | -0.02758     | 0.008645     | 0.464980966          |
| RealArm | LateBL     | Haptic  | EarlyAdapt | -0.05902 | 0.009067 | -6.51 | 0.00000010     | -0.07709     | -0.04096     | <b>0.000000100</b>   |
| RealArm | LateBL     | Haptic  | LateAdapt  | -0.02111 | 0.009094 | -2.32 | 0.0229         | -0.03923     | -0.00300     | 0.005104534          |
| RealArm | LateBL     | Haptic  | EarlyWash  | 0.00634  | 0.009067 | 0.70  | 0.4863         | -0.01172     | 0.024407     | 0.615223578          |
| RealArm | LateBL     | Haptic  | LateWash   | -0.01158 | 0.009067 | -1.28 | 0.2057         | -0.02964     | 0.006487     | 0.354005853          |
| RealArm | EarlyAdapt | RealArm | LateAdapt  | 0.00648  | 0.007157 | 0.91  | 0.3672         | -0.00772     | 0.020683     | 0.514305034          |
| RealArm | EarlyAdapt | RealArm | EarlyWash  | 0.01523  | 0.007157 | 2.13  | 0.0358         | 0.00103      | 0.029430     | 0.096370354          |
| RealArm | EarlyAdapt | RealArm | LateWash   | 0.00913  | 0.007157 | 1.28  | 0.2052         | -0.00507     | 0.023325     | 0.354005853          |
| RealArm | EarlyAdapt | EMG     | LateBL     | 0.00473  | 0.009303 | 0.51  | 0.6123         | -0.01380     | 0.023266     | 0.725247116          |
| RealArm | EarlyAdapt | EMG     | EarlyAdapt | -0.03072 | 0.009303 | -3.30 | 0.0015         | -0.04925     | -0.01219     | <b>0.005948637</b>   |
| RealArm | EarlyAdapt | EMG     | LateAdapt  | 0.00093  | 0.009303 | 0.10  | 0.9204         | -0.01760     | 0.019465     | 0.947503141          |
| RealArm | EarlyAdapt | EMG     | EarlyWash  | -0.00862 | 0.009303 | -0.93 | 0.3446         | -0.02028     | 0.016783     | <b>0.039827729</b>   |
| RealArm | EarlyAdapt | EMG     | LateWash   | 0.01138  | 0.009303 | 1.22  | 0.2249         | -0.00715     | 0.029916     | 0.361086770          |
| RealArm | EarlyAdapt | Haptic  | LateBL     | -0.00003 | 0.009094 | 0.00  | 0.9972         | -0.01815     | 0.018081     | 0.997188741          |
| RealArm | EarlyAdapt | Haptic  | EarlyAdapt | -0.04959 | 0.009067 | -5.47 | 0.00000402     | -0.06765     | -0.03152     | <b>0.000004020</b>   |
| RealArm | EarlyAdapt | Haptic  | LateAdapt  | -0.01168 | 0.009094 | -1.28 | 0.2030         | -0.02979     | 0.006436     | 0.354005853          |
| RealArm | EarlyAdapt | Haptic  | EarlyWash  | 0.01578  | 0.009067 | 1.74  | 0.0859         | -0.00228     | 0.033842     | 0.196112447          |
| RealArm | EarlyAdapt | Haptic  | LateWash   | -0.00214 | 0.009067 | -0.24 | 0.8141         | -0.02020     | 0.015923     | 0.872207560          |
| RealArm | LateAdapt  | RealArm | EarlyWash  | 0.00875  | 0.007157 | 1.22  | 0.2245         | -0.00545     | 0.022946     | 0.361086770          |
| RealArm | LateAdapt  | RealArm | LateWash   | 0.00264  | 0.007157 | 0.37  | 0.7128         | -0.01156     | 0.016842     | 0.796178855          |
| RealArm | LateAdapt  | EMG     | LateBL     | -0.01715 | 0.009303 | -1.91 | 0.0613         | -0.02028     | 0.016783     | 0.893800112          |
| RealArm | LateAdapt  | EMG     | EarlyAdapt | -0.03720 | 0.009303 | -4.00 | 0.000102       | -0.05573     | -0.01867     | <b>0.000675219</b>   |
| RealArm | LateAdapt  | EMG     | LateAdapt  | -0.00555 | 0.009303 | -0.60 | 0.5525         | -0.02408     | 0.012981     | 0.690616555          |
| RealArm | LateAdapt  | EMG     | EarlyWash  | 0.01718  | 0.009303 | 1.85  | 0.0687         | -0.00135     | 0.035716     | 0.163881323          |
| RealArm | LateAdapt  | EMG     | LateWash   | 0.00490  | 0.009303 | 0.53  | 0.5999         | -0.01363     | 0.023433     | 0.725247116          |
| RealArm | LateAdapt  | Haptic  | LateBL     | -0.00652 | 0.009094 | -0.72 | 0.4759         | -0.02463     | 0.011598     | 0.609383733          |
| RealArm | LateAdapt  | Haptic  | EarlyAdapt | -0.05607 | 0.009067 | -6.18 | 0.00000032     | -0.07413     | -0.03801     | <b>0.000000318</b>   |
| RealArm | LateAdapt  | Haptic  | LateAdapt  | -0.01816 | 0.009094 | -2.00 | 0.0494         | -0.03627     | -0.00005     | 0.123535910          |
| RealArm | LateAdapt  | Haptic  | EarlyWash  | 0.00930  | 0.009067 | 1.03  | 0.3085         | -0.00877     | 0.027359     | 0.469528726          |
| RealArm | LateAdapt  | Haptic  | LateWash   | -0.00862 | 0.009067 | -0.96 | 0.3446         | -0.02669     | 0.009440     | 0.506589707          |
| RealArm | EarlyWash  | RealArm | LateWash   | -0.00611 | 0.007157 | -0.85 | 0.3957         | -0.02030     | 0.008095     | 0.532696225          |
| RealArm | EarlyWash  | EMG     | LateBL     | -0.01050 | 0.009303 | -1.13 | 0.2628         | -0.02903     | 0.008036     | 0.411820953          |
| RealArm | EarlyWash  | EMG     | EarlyAdapt | -0.04595 | 0.009303 | -4.94 | 0.00000151     | -0.06448     | -0.02742     | <b>0.000027088</b>   |
| RealArm | EarlyWash  | EMG     | LateAdapt  | -0.01430 | 0.009303 | -1.54 | 0.1285         | -0.03283     | 0.004234     | 0.258099678          |
| RealArm | EarlyWash  | EMG     | EarlyWash  | 0.00844  | 0.009303 | 0.91  | 0.3674         | -0.01010     | 0.026969     | 0.514305034          |
| RealArm | EarlyWash  | EMG     | LateWash   | -0.00385 | 0.009303 | -0.41 | 0.6804         | -0.02238     | 0.014686     | 0.768699833          |
| RealArm | EarlyWash  | Haptic  | LateBL     | -0.01526 | 0.009094 | -1.68 | 0.0974         | -0.03338     | 0.002851     | 0.208734275          |
| RealArm | EarlyWash  | Haptic  | EarlyAdapt | -0.06432 | 0.009067 | -7.15 | 0.000000001    | -0.08288     | -0.04676     | <b>0.00000000863</b> |
| RealArm | EarlyWash  | Haptic  | LateAdapt  | -0.02691 | 0.009067 | -2.99 | 0.0041         | -0.05602     | -0.00880     | <b>0.014913722</b>   |
| RealArm | EarlyWash  | Haptic  | EarlyWash  | 0.00055  | 0.009067 | 0.06  | 0.9519         | -0.01751     | 0.018612     | 0.970350779          |
| RealArm | EarlyWash  | Haptic  | LateWash   | -0.01737 | 0.009067 | -1.92 | 0.0592         | -0.03543     | 0.000693     | 0.144595624          |
| RealArm | LateWash   | EMG     | LateBL     | -0.00439 | 0.009303 | -0.47 | 0.6382         | -0.02292     | 0.014140     | 0.736408949          |
| RealArm | LateWash   | EMG     | EarlyAdapt | -0.03984 | 0.009303 | -4.28 | 0.000033       | -0.05838     | -0.02131     | <b>0.000270338</b>   |
| RealArm | LateWash   | EMG     | LateAdapt  | -0.00819 | 0.009303 | -0.88 | 0.3813         | -0.02673     | 0.010339     | 0.526754773          |
| RealArm | LateWash   | EMG     | EarlyWash  | 0.01454  | 0.009303 | 1.56  | 0.1222         | -0.00399     | 0.033074     | 0.251665236          |
| RealArm | LateWash   | EMG     | LateWash   | 0.00226  | 0.009303 | 0.24  | 0.8089         | -0.01627     | 0.020791     | 0.872207560          |
| RealArm | LateWash   | Haptic  | LateBL     | -0.00916 | 0.009303 | -1.01 | 0.3171         | -0.02727     | 0.008955     | 0.475695543          |
| RealArm | LateWash   | Haptic  | EarlyAdapt | -0.05871 | 0.009067 | -6.48 | 0.00000010     | -0.07678     | -0.04065     | <b>0.000000103</b>   |
| RealArm | LateWash   | Haptic  | LateAdapt  | -0.02080 | 0.009094 | -2.29 | 0.0250         | -0.03892     | -0.00269     | 0.068953939          |
| RealArm | LateWash   | Haptic  | EarlyWash  | 0.00665  | 0.009067 | 0.73  | 0.4653         | -0.01141     | 0.024717     | 0.603217345          |
| RealArm | LateWash   | Haptic  | LateWash   | -0.01127 | 0.009067 | -1.24 | 0.2179         | -0.02933     | 0.006797     | 0.361086770          |
| EMG     | LateBL     | EMG     | EarlyAdapt | -0.03545 | 0.007157 | -4.95 | 0.00000205     | -0.04965     | -0.02125     | <b>0.000018384</b>   |
| EMG     | LateBL     | EMG     | LateAdapt  | -0.00380 | 0.007157 | -0.53 | 0.5965         | -0.01800     | 0.010398     | 0.725247116          |
| EMG     | LateBL     | EMG     | EarlyWash  | 0.01893  | 0.007157 | 2.65  | 0.0095         | 0.00473      | 0.033132     | <b>0.031638799</b>   |
| EMG     | LateBL     | EMG     | LateWash   | 0.00665  | 0.007157 | 0.93  | 0.3550         | -0.00755     | 0.020849     | 0.514305034          |
| EMG     | LateBL     | Haptic  | LateBL     | -0.00477 | 0.009094 | -0.52 | 0.6018         | -0.02288     | 0.013347     | 0.725247116          |
| EMG     | LateBL     | Haptic  | EarlyAdapt | -0.05432 | 0.009067 | -5.99 | 0.00000059     | -0.07238     | -0.03626     | <b>0.000000592</b>   |
| EMG     | LateBL     | Haptic  | LateAdapt  | -0.01641 | 0.009094 | -1.80 | 0.0751         | -0.03452     | 0.001702     | 0.175246886          |
| EMG     | LateBL     | Haptic  | EarlyWash  | 0.01105  | 0.009067 | 1.22  | 0.2270         | -0.00702     | 0.029109     | 0.361086770          |
| EMG     | LateBL     | Haptic  | LateWash   | -0.00687 | 0.009067 | -0.76 | 0.4508         | -0.02494     | 0.011189     | 0.591627737          |
| EMG     | EarlyAdapt | EMG     | LateAdapt  | 0.03165  | 0.007157 | 4.42  | 0.000015       | 0.01745      | 0.045850     | <b>0.000137630</b>   |
| EMG     | EarlyAdapt | EMG     | EarlyWash  | 0.05439  | 0.007157 | 7.60  | 0.0000000005   | 0.04019      | 0.068584     | <b>0.00000000058</b> |
| EMG     | EarlyAdapt | EMG     | LateWash   | 0.04210  | 0.007157 | 5.88  | 0.00000052     | 0.02790      | 0.056301     | <b>0.000000516</b>   |
| EMG     | EarlyAdapt | Haptic  | LateBL     | 0.03069  | 0.009094 | 3.37  | 0.0012         | 0.01257      | 0.048799     | <b>0.004911528</b>   |
| EMG     | EarlyAdapt | Haptic  | EarlyAdapt | -0.01887 | 0.009067 | -2.08 | 0.0409         | -0.03693     | -0.00081     | 0.104619330          |
| EMG     | EarlyAdapt | Haptic  | LateAdapt  | 0.01904  | 0.009094 | 2.09  | 0.0396         | 0.00093      | 0.037154     | 0.104022730          |
| EMG     | EarlyAdapt | Haptic  | EarlyWash  | 0.04650  | 0.009067 | 5.13  | 0.00000146     | 0.02843      | 0.064560     | <b>0.000014601</b>   |
| EMG     | EarlyAdapt | Haptic  | LateWash   | 0.02858  | 0.009067 | 3.15  | 0.0023         | 0.01052      | 0.046641     | <b>0.008751850</b>   |
| EMG     | LateAdapt  | EMG     | EarlyWash  | 0.02273  | 0.007157 | 3.18  | 0.0020         | 0.00853      | 0.036934     | <b>0.007712503</b>   |
| EMG     | LateAdapt  | EMG     | LateWash   | 0.01045  | 0.007157 | 1.46  | 0.1473         | -0.00375     | 0.024651     | 0.281270106          |
| EMG     | LateAdapt  | Haptic  | LateBL     | -0.00097 | 0.009094 | -0.11 | 0.9158         | -0.01908     | 0.017149     | 0.947503141          |
| EMG     | LateAdapt  | Haptic  | EarlyAdapt | -0.05052 | 0.009067 | -5.57 | 0.00000284     | -0.06858     | -0.03246     | <b>0.000002840</b>   |
| EMG     | LateAdapt  | Haptic  | LateAdapt  | -0.01261 | 0.009094 | -1.39 | 0.1696         | -0.03072     | 0.005503     | 0.312463860          |
| EMG     | LateAdapt  | Haptic  | EarlyWash  | 0.01485  | 0.009067 | 1.64  | 0.1057         | -0.00322     | 0.032910     | 0.222032615          |
| EMG     | LateAdapt  | Haptic  | LateWash   | -0.00307 | 0.009067 | -0.34 | 0.7357         | -0.02114     | 0.014991     | 0.813108029          |
| EMG     | EarlyWash  | EMG     | LateWash   | -0.01228 | 0.007157 | -1.72 | 0.0892         | -0.02648     | 0.001916     | 0.199319485          |
| EMG     | EarlyWash  | Haptic  | LateBL     | -0.02370 | 0.009094 | -2.61 | 0.0110         | -0.04181     | -0.00559     | <b>0.035077830</b>   |
| EMG     | EarlyWash  | Haptic  | EarlyAdapt | -0.07325 | 0.009067 | -8.08 | 0.000000000450 | 0.09132      | -0.05519     | <b>0.00000000045</b> |
| EMG     | EarlyWash  | Haptic  | LateAdapt  | -0.03534 | 0.009094 | -3.89 | 0.0002         | -0.05346     | -0.01723     | <b>0.000946379</b>   |
| EMG     | EarlyWash  | Haptic  | EarlyWash  | -0.00789 | 0.009067 | -0.87 | 0.3871         | -0.02595     | 0.010176     | 0.527902520          |
| EMG     | EarlyWash  | Haptic  | LateWash   | -0.02581 | 0.009067 | -2.85 | 0.0057         | -0.04387     | -0.00774     | <b>0.019968781</b>   |
| EMG     | LateWash   | Haptic  | LateBL     | -0.01142 | 0.009094 | -1.26 | 0.2132         | -0.02953     | 0.006697     | 0.361077927          |
| EMG     | LateWash   | Haptic  | EarlyAdapt | -0.06097 | 0.009067 | -6.72 | 0.000000046    | -0.07903     | -0.04291     | <b>0.000000046</b>   |
| EMG     | LateWash   | Haptic  | LateAdapt  | -0.02306 | 0.009094 | -2.54 | 0.0133         | -0.04118     | -0.00495     | <b>0.039827729</b>   |
| EMG     | LateWash   | Haptic  | EarlyWash  | 0.00440  | 0.009067 | 0.48  | 0.6293         | -0.01367     | 0.022458     | 0.734131136          |
| EMG     | LateWash   | Haptic  | LateWash   | -0.01352 | 0.009067 | -1.49 | 0.1400         | -0.03159     | 0.004539     | 0.272258298          |
| Haptic  | LateBL     | Haptic  | EarlyAdapt | -0.04956 | 0.006825 | -7.26 | 0.00000000221  | -0.06309     | -0.03602     | <b>0.00000000221</b> |
| Haptic  | LateBL     | Haptic  | LateAdapt  | -0.01165 | 0.006861 | -1.70 | 0.0926         | -0.02525     | 0.001960     | 0.202604895          |
| Haptic  | LateBL     | Haptic  | EarlyWash  | 0.01581  | 0.006825 | 2.32  | 0.0225         | -0.00227     | 0.029350     | 0.065104534          |
| Haptic  | LateBL     | Haptic  | LateWash   | -0.00211 | 0.006825 | -0.31 | 0.7581         | -0.01565     | 0.011431     | 0.829165353          |
| Haptic  | EarlyAdapt | Haptic  | LateAdapt  | 0.03791  | 0.006825 | 5.55  | 0.0000181      | 0.02437      | 0.051448     | <b>0.000001810</b>   |
| Haptic  | EarlyAdapt | Haptic  |            |          |          |       |                |              |              |                      |



## **Appendix B |**

# **Average Jensen-Shannon Divergence comparison Tables**

**Table B.1.** Average Jensen–Shannon Divergence (JSD) between each pair of participants’ post-personalization policies across all four motors. The table shows both motor-specific JSD values and the overall average JSD for each participant pair.

| Pair  | Motor 0 | Motor 1 | Motor 2 | Motor 3 | Avg. JSD      |
|-------|---------|---------|---------|---------|---------------|
| 1–2   | 0.0818  | 0.1060  | 0.2524  | 0.2733  | <b>0.1784</b> |
| 1–3   | 0.5147  | 0.2702  | 0.2672  | 0.2233  | <b>0.3189</b> |
| 1–4   | 0.1848  | 0.2726  | 0.4824  | 0.2221  | <b>0.2905</b> |
| 1–5   | 0.1804  | 0.2927  | 0.3222  | 0.1077  | <b>0.2257</b> |
| 1–6   | 0.1214  | 0.2130  | 0.3054  | 0.4597  | <b>0.2749</b> |
| 1–7   | 0.3940  | 0.3317  | 0.3745  | 0.3920  | <b>0.3730</b> |
| 1–8   | 0.3889  | 0.1641  | 0.1698  | 0.3095  | <b>0.2581</b> |
| 1–9   | 0.1958  | 0.1675  | 0.4088  | 0.5092  | <b>0.3203</b> |
| 1–10  | 0.3447  | 0.0641  | 0.3912  | 0.2190  | <b>0.2548</b> |
| 1–11  | 0.2305  | 0.1652  | 0.2752  | 0.2385  | <b>0.2273</b> |
| 2–3   | 0.5223  | 0.3509  | 0.0327  | 0.2841  | <b>0.2975</b> |
| 2–4   | 0.1539  | 0.3376  | 0.3002  | 0.4042  | <b>0.2990</b> |
| 2–5   | 0.1917  | 0.2111  | 0.2521  | 0.2991  | <b>0.2385</b> |
| 2–6   | 0.0720  | 0.2773  | 0.1296  | 0.2599  | <b>0.1847</b> |
| 2–7   | 0.3728  | 0.3346  | 0.1541  | 0.1439  | <b>0.2513</b> |
| 2–8   | 0.3674  | 0.1839  | 0.2629  | 0.3992  | <b>0.3033</b> |
| 2–9   | 0.1627  | 0.1565  | 0.2402  | 0.2879  | <b>0.2118</b> |
| 2–10  | 0.3482  | 0.1616  | 0.2996  | 0.3392  | <b>0.2871</b> |
| 2–11  | 0.2410  | 0.2167  | 0.0603  | 0.4501  | <b>0.2420</b> |
| 3–4   | 0.4132  | 0.2160  | 0.2976  | 0.1378  | <b>0.2661</b> |
| 3–5   | 0.4047  | 0.4893  | 0.2497  | 0.2364  | <b>0.3450</b> |
| 3–6   | 0.5001  | 0.3266  | 0.1309  | 0.4452  | <b>0.3507</b> |
| 3–7   | 0.2830  | 0.4372  | 0.1317  | 0.4025  | <b>0.3136</b> |
| 3–8   | 0.6022  | 0.2406  | 0.2563  | 0.1430  | <b>0.3105</b> |
| 3–9   | 0.4303  | 0.3958  | 0.2389  | 0.4868  | <b>0.3880</b> |
| 3–10  | 0.2410  | 0.2160  | 0.3037  | 0.0884  | <b>0.2123</b> |
| 3–11  | 0.4464  | 0.1743  | 0.0777  | 0.2455  | <b>0.2360</b> |
| 4–5   | 0.1671  | 0.4684  | 0.4983  | 0.2517  | <b>0.3463</b> |
| 4–6   | 0.1217  | 0.3970  | 0.1990  | 0.5470  | <b>0.3162</b> |
| 4–7   | 0.3349  | 0.4658  | 0.3487  | 0.5095  | <b>0.4147</b> |
| 4–8   | 0.3661  | 0.1825  | 0.5261  | 0.1012  | <b>0.2940</b> |
| 4–9   | 0.2140  | 0.3827  | 0.1289  | 0.5953  | <b>0.3303</b> |
| 4–10  | 0.2495  | 0.2645  | 0.3748  | 0.0755  | <b>0.2411</b> |
| 4–11  | 0.2960  | 0.1660  | 0.2674  | 0.1552  | <b>0.2212</b> |
| 5–6   | 0.1706  | 0.2234  | 0.3646  | 0.4908  | <b>0.3124</b> |
| 5–7   | 0.2652  | 0.2133  | 0.2103  | 0.4154  | <b>0.2760</b> |
| 5–8   | 0.4152  | 0.3555  | 0.1691  | 0.3276  | <b>0.3169</b> |
| 5–9   | 0.1425  | 0.1380  | 0.4499  | 0.5384  | <b>0.3172</b> |
| 5–10  | 0.2261  | 0.3405  | 0.4815  | 0.2419  | <b>0.3225</b> |
| 5–11  | 0.2294  | 0.3917  | 0.2865  | 0.2491  | <b>0.2892</b> |
| 6–7   | 0.3358  | 0.1951  | 0.1971  | 0.2194  | <b>0.2368</b> |
| 6–8   | 0.3975  | 0.3476  | 0.3743  | 0.5419  | <b>0.4153</b> |
| 6–9   | 0.1396  | 0.1400  | 0.1292  | 0.1089  | <b>0.1294</b> |
| 6–10  | 0.3315  | 0.2129  | 0.2555  | 0.5046  | <b>0.3261</b> |
| 6–11  | 0.3021  | 0.3088  | 0.0882  | 0.6109  | <b>0.3275</b> |
| 7–8   | 0.6060  | 0.4433  | 0.2734  | 0.4742  | <b>0.4492</b> |
| 7–9   | 0.2434  | 0.2207  | 0.2918  | 0.2110  | <b>0.2417</b> |
| 7–10  | 0.2308  | 0.3238  | 0.3520  | 0.4527  | <b>0.3398</b> |
| 7–11  | 0.3649  | 0.4177  | 0.1791  | 0.5498  | <b>0.3779</b> |
| 8–9   | 0.4471  | 0.3050  | 0.4693  | 0.5492  | <b>0.4426</b> |
| 8–10  | 0.5104  | 0.1532  | 0.4477  | 0.1031  | <b>0.3036</b> |
| 8–11  | 0.3344  | 0.1023  | 0.3039  | 0.2335  | <b>0.2435</b> |
| 9–10  | 0.3063  | 0.2197  | 0.2722  | 0.5441  | <b>0.3356</b> |
| 9–11  | 0.2591  | 0.3110  | 0.2057  | 0.6519  | <b>0.3569</b> |
| 10–11 | 0.2714  | 0.1352  | 0.2631  | 0.1806  | <b>0.2126</b> |

**Table B.2.** Average Jensen–Shannon Divergence (JSD) between each pair of participants’ learned post-perturbation policies across all four motors. The table shows both motor-specific JSD values and the overall average JSD for each participant pair.

| Pair  | Motor 0 | Motor 1 | Motor 2 | Motor 3 | Avg. JSD      |
|-------|---------|---------|---------|---------|---------------|
| 1–2   | 0.2853  | 0.1062  | 0.6347  | 0.3317  | <b>0.3395</b> |
| 1–3   | 0.4926  | 0.1750  | 0.6442  | 0.1724  | <b>0.3710</b> |
| 1–4   | 0.1012  | 0.2273  | 0.7526  | 0.1163  | <b>0.2994</b> |
| 1–5   | 0.2225  | 0.1925  | 0.3561  | 0.1459  | <b>0.2292</b> |
| 1–6   | 0.1553  | 0.3577  | 0.5638  | 0.3214  | <b>0.3496</b> |
| 1–7   | 0.1957  | 0.3296  | 0.4288  | 0.1273  | <b>0.2704</b> |
| 1–8   | 0.4200  | 0.3351  | 0.5516  | 0.2644  | <b>0.3928</b> |
| 1–9   | 0.2355  | 0.0916  | 0.7273  | 0.6210  | <b>0.4188</b> |
| 1–10  | 0.3304  | 0.1289  | 0.7099  | 0.1392  | <b>0.3271</b> |
| 1–11  | 0.2527  | 0.6716  | 0.5569  | 0.4306  | <b>0.4779</b> |
| 2–3   | 0.5858  | 0.1560  | 0.0458  | 0.4162  | <b>0.3009</b> |
| 2–4   | 0.3165  | 0.2199  | 0.2888  | 0.3325  | <b>0.2894</b> |
| 2–5   | 0.2514  | 0.2211  | 0.4074  | 0.2293  | <b>0.2773</b> |
| 2–6   | 0.2788  | 0.3794  | 0.1492  | 0.1339  | <b>0.2353</b> |
| 2–7   | 0.2678  | 0.3923  | 0.4354  | 0.2251  | <b>0.3301</b> |
| 2–8   | 0.4899  | 0.2506  | 0.2159  | 0.4595  | <b>0.3540</b> |
| 2–9   | 0.3680  | 0.1567  | 0.2428  | 0.3921  | <b>0.2899</b> |
| 2–10  | 0.4901  | 0.1054  | 0.2985  | 0.3602  | <b>0.3136</b> |
| 2–11  | 0.4697  | 0.6665  | 0.1687  | 0.6375  | <b>0.4856</b> |
| 3–4   | 0.5330  | 0.0770  | 0.2733  | 0.1529  | <b>0.2591</b> |
| 3–5   | 0.6106  | 0.3199  | 0.4236  | 0.2883  | <b>0.4106</b> |
| 3–6   | 0.4243  | 0.3291  | 0.1737  | 0.4210  | <b>0.3370</b> |
| 3–7   | 0.4336  | 0.4096  | 0.4247  | 0.2418  | <b>0.3774</b> |
| 3–8   | 0.6313  | 0.1806  | 0.2306  | 0.1243  | <b>0.2917</b> |
| 3–9   | 0.4660  | 0.1831  | 0.2261  | 0.6675  | <b>0.3857</b> |
| 3–10  | 0.3049  | 0.1755  | 0.2893  | 0.1068  | <b>0.2191</b> |
| 3–11  | 0.3866  | 0.6311  | 0.2019  | 0.3255  | <b>0.3863</b> |
| 4–5   | 0.1883  | 0.3748  | 0.5872  | 0.1909  | <b>0.3353</b> |
| 4–6   | 0.1613  | 0.3326  | 0.4057  | 0.3269  | <b>0.3066</b> |
| 4–7   | 0.2158  | 0.4241  | 0.5387  | 0.1327  | <b>0.3278</b> |
| 4–8   | 0.4048  | 0.1938  | 0.4303  | 0.2376  | <b>0.3166</b> |
| 4–9   | 0.2411  | 0.2199  | 0.2015  | 0.6271  | <b>0.3224</b> |
| 4–10  | 0.3174  | 0.2256  | 0.4134  | 0.1622  | <b>0.2796</b> |
| 4–11  | 0.2255  | 0.5975  | 0.4284  | 0.4387  | <b>0.4225</b> |
| 5–6   | 0.3148  | 0.4841  | 0.2818  | 0.2624  | <b>0.3358</b> |
| 5–7   | 0.2876  | 0.3084  | 0.4532  | 0.1567  | <b>0.3015</b> |
| 5–8   | 0.4438  | 0.4147  | 0.3236  | 0.3684  | <b>0.3876</b> |
| 5–9   | 0.2420  | 0.2150  | 0.5602  | 0.5449  | <b>0.3905</b> |
| 5–10  | 0.4403  | 0.1906  | 0.5623  | 0.2610  | <b>0.3635</b> |
| 5–11  | 0.3206  | 0.7362  | 0.2581  | 0.5151  | <b>0.4575</b> |
| 6–7   | 0.1546  | 0.5422  | 0.4631  | 0.2090  | <b>0.3423</b> |
| 6–8   | 0.4449  | 0.3731  | 0.2530  | 0.3766  | <b>0.3619</b> |
| 6–9   | 0.3209  | 0.3549  | 0.3750  | 0.4015  | <b>0.3631</b> |
| 6–10  | 0.2733  | 0.4340  | 0.3900  | 0.3360  | <b>0.3583</b> |
| 6–11  | 0.2603  | 0.4966  | 0.0599  | 0.6097  | <b>0.3567</b> |
| 7–8   | 0.4461  | 0.5252  | 0.5898  | 0.2819  | <b>0.4607</b> |
| 7–9   | 0.2859  | 0.3180  | 0.4618  | 0.5480  | <b>0.4034</b> |
| 7–10  | 0.3395  | 0.4132  | 0.5416  | 0.1601  | <b>0.3636</b> |
| 7–11  | 0.3490  | 0.6966  | 0.4842  | 0.4946  | <b>0.5061</b> |
| 8–9   | 0.4171  | 0.3433  | 0.3579  | 0.6118  | <b>0.4325</b> |
| 8–10  | 0.5175  | 0.2779  | 0.3693  | 0.1418  | <b>0.3266</b> |
| 8–11  | 0.5291  | 0.5571  | 0.2133  | 0.3449  | <b>0.4111</b> |
| 9–10  | 0.3360  | 0.2055  | 0.2446  | 0.6217  | <b>0.3520</b> |
| 9–11  | 0.2400  | 0.6561  | 0.3930  | 0.7776  | <b>0.5167</b> |
| 10–11 | 0.2542  | 0.6854  | 0.3989  | 0.3905  | <b>0.4323</b> |

# Appendix C |

## Proximal Policy Optimization Hyperparameters and Packages

### System Overview:

The PPO training framework was implemented on **Ubuntu 24.04** using **ROS 2 Jazzy** and **Python 3.12**. The reinforcement learning agent used the **Stable Baselines3** library with **PyTorch**. The setup is human-in-the-loop: a user interprets haptic feedback from a 4-motor platform to control a joystick.

### Required Python Packages

The following Python packages are required to run the PPO codebase with Stable Baselines3:

- `stable-baselines3[extra]`
- `torch`
- `numpy`
- `gymnasium`
- `matplotlib`
- `pandas`
- `scipy`
- `pyserial`

## Environment

- **Train\_Model:**
  - Input: normalized joystick direction in  $[-1, 1]$
  - Output: 4-dimensional motor intensity commands (Box action space)
  - Trial logic and reward calculation based on directional error

## Neural Network Architecture and Hyperparameters

| Parameter                    | Value                                  |
|------------------------------|----------------------------------------|
| Hidden Layers                | 2                                      |
| Nodes per Layer              | 64                                     |
| Policy/Value Network         | Separate (no shared feature extractor) |
| Activation Function          | ReLU                                   |
| Policy Std                   | Full matrix                            |
| Learning Rate                | 0.0001                                 |
| Batch Size                   | 10                                     |
| Number of Steps              | 10                                     |
| Number of Epochs             | 15                                     |
| Entropy Coefficient          | 0.01                                   |
| Value Coefficient            | 0.00                                   |
| Discount Factor ( $\gamma$ ) | 0.99                                   |
| GAE Lambda                   | 0.95                                   |

## ROS 2 Nodes and Execution

The code requires four terminal windows:

1. **Joystick Serial Node:** Reads joystick Arduino Mega and publishes ‘/joystick\_values’.  
Command: `ros2 run ppo_joystick JoystickArduinoSerial_____`
2. **Haptic Serial Node:** Sends motor commands to Arduino Uno.  
Command: `ros2 run ppo_joystick HapticArduinoSerial_____`
3. **Main PPO Node:** Runs either ‘PPO\_Trainer’ (personalization) or ‘Adaptation\_PPO\_Trainer’.  
Command: `ros2 run ppo_joystick PPO_Trainer or Adaptation_PPO_Trainer_____`

4. **Participant Performance Data Logger:** Logs the information published to the /trial\_results ROS topic containing the trial number, goal direction, response direction, and policy reward

Command: `./Data_Logger.sh` \_\_\_\_\_

## Trial Structure

- Each trial sends a new joystick target to the policy
- Motor output is active up to 5 seconds, then auto-shuts off
- User responses recorded over 2-second windows
- Directional error between target and response determines reward

## Notes

- Pre-training on expert demonstrations recommended before PPO fine-tuning
- Environment handles motor shutdown and trial resets automatically
- Thread-safe shared data dictionaries manage joystick subscriptions
- Configure ‘save\_path’ in the PPO scripts and Data\_Logger to specify file locations and save directory.
  - New\_PPO\_control\_node.py: update **save\_path** to define where personalized training models are stored
    - \* save\_dir: Directory for saving log and model files
    - \* file\_name: name for the log file storing model training information
    - \* model\_name: name of the final outputted model to be saved as.
    - \* expert\_model\_path: file path pointing to the location of the pretrained policy.
  - New\_PPO\_control\_node\_Adaptation.py: update save\_dir, file\_name, old\_model\_name, and model\_name to define where adaptation training models are stored
    - \* save\_dir: Directory for saving log and model files
    - \* file\_name: name for the log file storing model training information

- \* `old_model_name`: Name of the PPO model file obtained from running `New_PPO_control_node.py`
- \* `model_name`: name of the final outputed model to be saved as.
- `Data_Logger.sh`: update `BASE_DIR` and `FILENAME` to define the data storage location and the name of the saved ROS 2 bag file.

# Bibliography

- [1] SINGHAL, T. and O. SCHNEIDER (2021) “Juicy Haptic Design: Vibrotactile Embellishments Can Improve Player Experience in Games,” in *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, Association for Computing Machinery, New York, NY, USA.  
URL <https://doi.org/10.1145/3411764.3445463>
- [2] GÜNTHER, S., F. MÜLLER, M. FUNK, J. KIRCHNER, N. DEZFULI, and M. MÜHLHÄUSER (2018) “TactileGlove: Assistive Spatial Guidance in 3D Space through Vibrotactile Navigation,” in *Proceedings of the 11th pervasive technologies related to assistive environments conference*, Association for Computing Machinery, New York, NY, USA, pp. 273—280.  
URL <https://doi.org/10.1145/3197768.3197785>
- [3] KAPUR, P., M. JENSEN, L. J. BUXBAUM, S. A. JAX, and K. J. KUCHENBECKER (2010) “Spatially distributed tactile feedback for kinesthetic motion guidance,” in *2010 IEEE Haptics Symposium*, pp. 519–526.
- [4] KRUEGER, A. R., P. GIANNONI, V. SHAH, M. CASADIO, and R. A. SCHEIDT (2017) “Supplemental vibrotactile feedback control of stabilization and reaching actions of the arm using limb state and position error encodings,” *Journal of neuroengineering and rehabilitation*, **14**, pp. 1–23.
- [5] DEVIGNE, L., M. AGGRAVI, M. BIVAUD, N. BALIX, C. S. TEODORESCU, T. CARLSON, T. SPRETERS, C. PACCHIEROTTI, and M. BABEL (2020) “Power wheelchair navigation assistance using wearable vibrotactile haptics,” *IEEE transactions on haptics*, **13**(1), pp. 52–58.
- [6] COLLAÇO, E., E. KIRA, L. H. SALLABERRY, A. C. QUEIROZ, M. A. MACHADO, O. CRIVELLO JR, and R. TORI (2021) “Immersion and haptic feedback impacts on dental anesthesia technical skills virtual reality training,” *Journal of Dental Education*, **85**(4), pp. 589–598.
- [7] GIL, J. J., I. DÍAZ, P. CIÁURRIZ, and M. ECHEVERRÍA (2013) “New driving control system with haptic feedback: Design and preliminary validation tests,” *Transportation Research Part C: Emerging Technologies*, **33**, pp. 22–36.



URL <https://www.sciencedirect.com/science/article/pii/S0968090X13000776>

- [8] TAN, H. Z., S. CHOI, F. W. LAU, and F. ABNOUSI (2020) “Methodology for maximizing information transmission of haptic devices: A survey,” *Proceedings of the IEEE*, **108**(6), pp. 945–965.
- [9] CHAN, A., K. MACLEAN, and J. MCGRENERE (2005) “Learning and identifying haptic icons under workload,” in *First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics Conference*, pp. 432–439.
- [10] DUNKELBERGER, N., J. SULLIVAN, J. BRADLEY, N. P. WALLING, I. MANICKAM, G. DASARATHY, A. ISRAR, F. W. Y. LAU, K. KLUMB, B. KNOTT, F. ABNOUSI, R. BARANIUK, and M. K. O’MALLEY (2018) “Conveying Language through Haptics: A Multi-Sensory Approach,” in *Proceedings of the 2018 ACM International Symposium on Wearable Computers, ISWC ’18*, Association for Computing Machinery, New York, NY, USA, p. 25–32.  
URL <https://doi.org/10.1145/3267242.3267244>
- [11] AGGRAVI, M., F. PAUSÉ, P. R. GIORDANO, and C. PACCHIEROTTI (2018) “Design and evaluation of a wearable haptic device for skin stretch, pressure, and vibrotactile stimuli,” *IEEE Robotics and Automation Letters*, **3**(3), pp. 2166–2173.
- [12] PEZENT, E., A. ISRAR, M. SAMAD, S. ROBINSON, P. AGARWAL, H. BENKO, and N. COLONNESE (2019) “Tasbi: Multisensory Squeeze and Vibrotactile Wrist Haptics for Augmented and Virtual Reality,” in *2019 IEEE World Haptics Conference (WHC)*, pp. 1–6.
- [13] BUCKINGHAM, G., D. P. CAREY, F. L. COLINO, J. DEGROSBOIS, and G. BINSTED (2010) “Gating of vibrotactile detection during visually guided bimanual reaches,” *Experimental brain research*, **201**, pp. 411–419.
- [14] COLINO, F. L., G. BUCKINGHAM, D. T. CHENG, P. VAN DONKELAAR, and G. BINSTED (2014) “Tactile gating in a reaching and grasping task,” *Physiological reports*, **2**(3), p. e00267.
- [15] COLINO, F. L. and G. BINSTED (2016) “Time course of tactile gating in a reach-to-grasp and lift task,” *Journal of Motor Behavior*, **48**(5), pp. 390–400.
- [16] JURAVLE, G. and C. SPENCE (2011) “Juggling reveals a decisional component to tactile suppression,” *Experimental Brain Research*, **213**, pp. 87–97.
- [17] VARDAR, Y., B. GÜÇLÜ, and C. BASDOGAN (2018) “Tactile Masking by Electro-vibration,” *IEEE Transactions on Haptics*, **11**(4), pp. 623–635.

- [18] CRAIG, J. C. and P. M. EVANS (1995) “Tactile selective attention and temporal masking,” *Perception & Psychophysics*, **57**, pp. 511–518.  
URL <https://doi.org/10.3758/BF03213076>
- [19] ZOOK, Z. A., J. J. FLECK, T. W. TJANDRA, and M. K. O’MALLEY (2019) “Effect of Interference on Multi-Sensory Haptic Perception of Stretch and Squeeze,” in *2019 IEEE World Haptics Conference (WHC)*, pp. 371–376.
- [20] TAN, H. Z., C. M. REED, L. A. DELHORNE, N. I. DURLACH, and N. WAN (2003) “Temporal masking of multidimensional tactual stimuli,” *The Journal of the Acoustical Society of America*, **114**(6), pp. 3295–3308, [https://pubs.aip.org/asa/jasa/article-pdf/114/6/3295/8092903/3295\\_1\\_online.pdf](https://pubs.aip.org/asa/jasa/article-pdf/114/6/3295/8092903/3295_1_online.pdf).  
URL <https://doi.org/10.1121/1.1623788>
- [21] OAKLEY, I. and J. PARK (2008) “Did you feel something? Distracter tasks and the recognition of vibrotactile cues,” *Interacting with Computers*, **20**, pp. 354–363.
- [22] CHEN, Q., S. T. PERRAULT, Q. ROY, and L. WYSE (2018) “Effect of Temporality, Physical Activity and Cognitive Load on Spatiotemporal Vibrotactile Pattern Recognition,” in *Proceedings of the 2018 International Conference on Advanced Visual Interfaces, AVI ’18*, Association for Computing Machinery, New York, NY, USA.  
URL <https://doi.org/10.1145/3206505.3206511>
- [23] TAN, H. Z., C. M. REED, and N. I. DURLACH (2009) “Optimum information transfer rates for communication through haptic and other sensory modalities,” *IEEE Transactions on Haptics*, **3**(2), pp. 98–108.
- [24] BARK, K., P. KHANNA, R. IRWIN, P. KAPUR, S. A. JAX, L. J. BUXBAUM, and K. J. KUCHENBECKER (2011) “Lessons in using vibrotactile feedback to guide fast arm motions,” in *2011 IEEE World Haptics Conference*, IEEE, pp. 355–360.
- [25] BARK, K., E. HYMAN, F. TAN, E. CHA, S. A. JAX, L. J. BUXBAUM, and K. J. KUCHENBECKER (2015) “Effects of Vibrotactile Feedback on Human Learning of Arm Motions,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, **23**(1), pp. 51–63.
- [26] SALAZAR LUCES, J. V., K. OKABE, Y. MURAO, and Y. HIRATA (2018) “A Phantom-Sensation Based Paradigm for Continuous Vibrotactile Wrist Guidance in Two-Dimensional Space,” *IEEE Robotics and Automation Letters*, **3**(1), pp. 163–170.
- [27] BULTITUDE, J. H., G. JURAVLE, and C. SPENCE (2016) “Tactile gap detection deteriorates during bimanual symmetrical movements under mirror visual feedback,” *PloS one*, **11**(1), p. e0146077.

- [28] CHAPMAN, C. E. and E. BEAUCHAMP (2006) “Differential controls over tactile detection in humans by motor commands and peripheral refference,” *Journal of Neurophysiology*, **96**(3), pp. 1664–1675.
- [29] VAN HULLE, L., G. JURAVLE, C. SPENCE, G. CROMBEZ, and S. VAN DAMME (2013) “Attention modulates sensory suppression during back movements,” *Consciousness and cognition*, **22**(2), pp. 420–429.
- [30] SPELMEZAN, D., M. JACOBS, A. HILGERS, and J. BORCHERS (2009) “Tactile motion instructions for physical activities,” in *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 2243–2252.
- [31] GUTIERREZ, S. F. C., J. V. S. LUCES, and Y. HIRATA (2022) “Modality Influence on the Motor Learning of Ballroom Dance with a Mixed-Reality Human-Machine Interface,” in *2022 IEEE/SICE International Symposium on System Integration (SII)*, pp. 177–182.
- [32] SHERGILL, S. S., P. M. BAYS, C. D. FRITH, and D. M. WOLPERT (2003) “Two eyes for an eye: the neuroscience of force escalation,” *Science*, **301**(5630), pp. 187–187.
- [33] BAYS, P. M., D. M. WOLPERT, and J. R. FLANAGAN (2005) “Perception of the consequences of self-action is temporally tuned and event driven,” *Current Biology*, **15**(12), pp. 1125–1128.
- [34] ZIAT, M., V. HAYWARD, C. E. CHAPMAN, M. O. ERNST, and C. LENAY (2010) “Tactile suppression of displacement,” *Experimental brain research*, **206**, pp. 299–310.
- [35] LYLYKANGAS, J., J. HEIKKINEN, V. SURAKKA, R. RAISAMO, K. MYLLYMAA, and A. LAITINEN (2015) “Vibrotactile Stimulation as an Instructor for Mimicry-Based Physical Exercise,” *Advances in Human - Computer Interaction*, **2015**, copyright - Copyright © 2015 Jani Lylykangas et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited; Last updated - 2023-02-24.  
URL <https://ezaccess.libraries.psu.edu/login?url=https://www.proquest.com/scholarly-journals/vibrotactile-stimulation-as-instructor-mimicry/docview/1731737754/se-2>
- [36] FITTS, P. M. (1954) “The information capacity of the human motor system in controlling the amplitude of movement.” *Journal of experimental psychology*, **47**(6), p. 381.
- [37] PROCTOR, R. W. and T. G. REEVE (1989) *Stimulus-response compatibility: An integrated perspective*, Elsevier.

- [38] LEDERMAN, S. J. and L. A. JONES (2011) “Tactile and haptic illusions,” *IEEE Transactions on Haptics*, **4**(4), pp. 273–294.
- [39] CHOLEWIAK, R. W. and A. A. COLLINS (2000) “The generation of vibrotactile patterns on a linear array: Influences of body site, time, and presentation mode,” *Perception & Psychophysics*, **62**, pp. 1220–1235.
- [40] SALAZAR, J., Y. HIRATA, and K. KOSUGE (2016) “Motion guidance using haptic feedback based on vibrotactile illusions,” in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, pp. 4685–4691.
- [41] LIN, C.-Y., C.-M. TSAI, P.-C. SHIH, and H.-C. WU (2016) “Development of a novel haptic glove for improving finger dexterity in poststroke rehabilitation,” *Technology and Health Care*, **24**(s1), pp. S97–S103.
- [42] HERNANDEZ, H., I. POITRAS, L. FAY, A. KHAN, J.-S. ROY, and E. BIDDIS (2021) “A gaming system with haptic feedback to improve upper extremity function: A prospective case series,” *Technology and Disability*, **33**(3), pp. 195–206.
- [43] SULLIVAN, J. L., N. DUNKELBERGER, J. BRADLEY, J. YOUNG, A. ISRAR, F. LAU, K. KLUMB, F. ABNOUSI, and M. K. O’MALLEY (2019) “Multi-Sensory Stimuli Improve Distinguishability of Cutaneous Haptic Cues,” *IEEE Transactions on Haptics*, **13**(2), pp. 286–297.
- [44] ALEXANDER, S. A., R. GARCIA, and M. K. O’MALLEY (2021) “Enhancing Multi-Sensory Cue Salience and Perceptual Identification in a Wearable Haptic Device,” in *2021 IEEE World Haptics Conference (WHC)*, IEEE, pp. 900–905.
- [45] REZA MOTAMEDI, M., M. OTIS, and V. DUCHAINE (2017) “The Impact of Simultaneously Applying Normal Stress and Vibrotactile Stimulation for Feedback of Exteroceptive Information,” *Journal of Biomechanical Engineering*, **139**(6), p. 061004, [https://asmedigitalcollection.asme.org/biomechanical/article-pdf/139/6/061004/6095249/bio\\_139\\_06\\_061004.pdf](https://asmedigitalcollection.asme.org/biomechanical/article-pdf/139/6/061004/6095249/bio_139_06_061004.pdf).  
URL <https://doi.org/10.1115/1.4036417>
- [46] ALLES, D. S. (1970) “Information Transmission by Phantom Sensations,” *IEEE Transactions on Man-Machine Systems*, **11**(1), pp. 85–91.
- [47] ISRAR, A. and I. POUPYREV (2011) “Tactile Brush: Drawing on Skin with a Tactile Grid Display,” in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’11, Association for Computing Machinery, New York, NY, USA, p. 2019–2028.  
URL <https://doi.org/10.1145/1978942.1979235>
- [48] CHOLEWIAK, R. W. and A. A. COLLINS (2003) “Vibrotactile localization on the arm: Effects of place, space, and age,” *Perception & psychophysics*, **65**(7), pp. 1058–1077.

- [49] LUZHINICA, G., S. STEIN, E. VEAS, V. PAMMER, J. WILLIAMSON, and R. M. SMITH (2017) “Personalising vibrotactile displays through perceptual sensitivity adjustment,” in *Proceedings of the 2017 ACM International Symposium on Wearable Computers*, pp. 66–73.
- [50] LUCES, J. V. S., K. OKABE, Y. MURAO, and Y. HIRATA (2017) “A phantom-sensation based paradigm for continuous vibrotactile wrist guidance in two-dimensional space,” *IEEE Robotics and Automation Letters*, **3**(1), pp. 163–170.
- [51] LIAO, Z., J. V. S. LUCES, and Y. HIRATA (2020) “Human Navigation Using Phantom Tactile Sensation Based Vibrotactile Feedback,” *IEEE Robotics and Automation Letters*, **5**(4), pp. 5732–5739.
- [52] WOBBROCK, J. O., L. FINDLATER, D. GERGLE, and J. J. HIGGINS (2011) “The aligned rank transform for nonparametric factorial analyses using only anova procedures,” in *Proceedings of the SIGCHI conference on human factors in computing systems*, pp. 143–146.
- [53] BENJAMINI, Y. and Y. HOCHBERG (1995) “Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing,” *Journal of the Royal Statistical Society: Series B (Methodological)*, **57**(1), pp. 289–300, [https://academic.oup.com/jrsssb/article-pdf/57/1/289/49173396/jrsssb\\_57\\_1\\_289.pdf](https://academic.oup.com/jrsssb/article-pdf/57/1/289/49173396/jrsssb_57_1_289.pdf). URL <https://doi.org/10.1111/j.2517-6161.1995.tb02031.x>
- [54] R CORE TEAM (2023) *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>
- [55] JURAVLE, G., G. BINSTED, and C. SPENCE (2017) “Tactile suppression in goal-directed movement,” *Psychonomic Bulletin & Review*, **24**, pp. 1060–1076.
- [56] ACKERLEY, R., I. CARLSSON, H. WESTER, H. OLAUSSON, and H. BACKLUND WASLING (2014) “Touch perceptions across skin sites: differences between sensitivity, direction discrimination and pleasantness,” *Frontiers in behavioral neuroscience*, **8**, p. 54.
- [57] LOW, A. K., Z. A. ZOOK, J. J. FLECK, and M. K. O’MALLEY (2021) “Effects of Interfering Cue Separation Distance and Amplitude on the Haptic Detection of Skin Stretch,” *IEEE Transactions on Haptics*, **14**(2), pp. 254–259.
- [58] CORDELLA, F., A. L. CIANCIO, R. SACCHETTI, A. DAVALLI, A. G. CUTTI, E. GUGLIELMELLI, and L. ZOLLO (2016) “Literature review on needs of upper limb prosthesis users,” *Frontiers in neuroscience*, **10**, p. 209.
- [59] STEIN, J., K. NARENDHAN, J. McBEAN, K. KREBS, and R. HUGHES (2007) “Electromyography-controlled exoskeletal upper-limb-powered orthosis for exercise training after stroke,” *American journal of physical medicine & rehabilitation*, **86**(4), pp. 255–261.

- [60] WITTEVEEN, H. J., E. A. DROOG, J. S. RIETMAN, and P. H. VELTINK (2012) “Vibro-and electrotactile user feedback on hand opening for myoelectric forearm prostheses,” *IEEE transactions on biomedical engineering*, **59**(8), pp. 2219–2226.
- [61] ANGLIN, J., T. SUGIYAMA, and S.-L. LIEW (2017) “Visuomotor adaptation in head-mounted virtual reality versus conventional training,” *Scientific reports*, **7**(1), p. 45469.
- [62] PALLUEL-GERMAIN, R., S. A. JAX, and L. J. BUXBAUM (2011) “Visuo-motor gain adaptation and generalization following left hemisphere stroke,” *Neuroscience letters*, **498**(3), pp. 222–226.
- [63] BERNIKER, M. and K. KORDING (2008) “Estimating the sources of motor errors for adaptation and generalization,” *Nature neuroscience*, **11**(12), pp. 1454–1461.
- [64] COHEN, E. J., K. WEI, and D. MINCIACCHI (2019) “Visuomotor perturbation in a continuous circle tracing task: novel approach for quantifying motor adaptation,” *Scientific Reports*, **9**(1), p. 18679.
- [65] SHABBOTT, B. A. and R. L. SAINBURG (2010) “Learning a visuomotor rotation: simultaneous visual and proprioceptive information is crucial for visuomotor remapping,” *Experimental Brain Research*, **203**(1), pp. 75–87.
- [66] SEGURA, D., E. ROMERO, V. E. ABARCA, and D. A. ELIAS (2024) “Upper Limb Prostheses by the Level of Amputation: A Systematic Review,” *Prosthesis*, **6**(2), pp. 277–300.  
URL <https://www.mdpi.com/2673-1592/6/2/22>
- [67] FLEMING, A., N. STAFFORD, S. HUANG, X. HU, D. P. FERRIS, and H. H. HUANG (2021) “Myoelectric control of robotic lower limb prostheses: a review of electromyography interfaces, control paradigms, challenges and future directions,” *Journal of neural engineering*, **18**(4), p. 041004.
- [68] MA, T. and X. HU (2025) “Adaptive motor unit decomposition using a cross-validation-based update policy,” *Computers in Biology and Medicine*, **193**, p. 110479.
- [69] ANDERSON, D. I., R. A. MAGILL, A. M. MAYO, and K. A. STEEL (2019) “Enhancing motor skill acquisition with augmented feedback,” in *Skill acquisition in sport*, Routledge, pp. 3–19.
- [70] SIGRIST, R., G. RAUTER, R. RIENER, and P. WOLF (2013) “Augmented visual, auditory, haptic, and multimodal feedback in motor learning: a review,” *Psychonomic bulletin & review*, **20**(1), pp. 21–53.
- [71] ANTFOLEK, C., M. D’ALONZO, B. ROSÉN, G. LUNDBORG, F. SEBELIUS, and C. CIPRIANI (2013) “Sensory feedback in upper limb prosthetics,” *Expert review of medical devices*, **10**(1), pp. 45–54.

- [72] ANTFOLEK, C., M. D'ALONZO, M. CONTROZZI, G. LUNDBORG, B. ROSEN, F. SEBELIUS, and C. CIPRIANI (2012) "Artificial redirection of sensation from prosthetic fingers to the phantom hand map on transradial amputees: vibrotactile versus mechanotactile sensory feedback," *IEEE transactions on neural systems and rehabilitation engineering*, **21**(1), pp. 112–120.
- [73] KACZMAREK, K. A., J. G. WEBSTER, P. BACH-Y RITA, and W. J. TOMPKINS (2002) "Electrotactile and vibrotactile displays for sensory substitution systems," *IEEE transactions on biomedical engineering*, **38**(1), pp. 1–16.
- [74] GONZALEZ, J., H. SOMA, M. SEKINE, and W. YU (2012) "Psycho-physiological assessment of a prosthetic hand sensory feedback system based on an auditory display: a preliminary study," *Journal of neuroengineering and rehabilitation*, **9**(1), p. 33.
- [75] WHEELER, J., K. BARK, J. SAVALL, and M. CUTKOSKY (2010) "Investigation of rotational skin stretch for proprioceptive feedback with application to myoelectric systems," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, **18**(1), pp. 58–66.
- [76] MARTIN, J., J. SAVALL, I. DIAZ, J. HERNANTES, and D. BORRO (2008) "Evaluation of sensory substitution to simplify the mechanical design of a haptic wrist," in *2008 IEEE International Workshop on Haptic Audio visual Environments and Games*, pp. 132–136.
- [77] GARENFELD, M. A., M. STRBAC, N. JORGOVANOVIC, J. L. DIDERIKSEN, and S. DOSEN (2023) "Closed-Loop Control of a Multifunctional Myoelectric Prosthesis With Full-State Anatomically Congruent Electrotactile Feedback," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, **31**, pp. 2090–2100.
- [78] SENSINGER, J. W. and S. DOSEN (2020) "A review of sensory feedback in upper-limb prostheses from the perspective of human motor control," *Frontiers in neuroscience*, **14**, p. 345.
- [79] OKAMURA, A. M., L. N. VERNER, T. YAMAMOTO, J. C. GWILLIAM, and P. G. GRIFFITHS (2010) "Force feedback and sensory substitution for robot-assisted surgery," *Surgical Robotics: Systems Applications and Visions*, pp. 419–448.
- [80] ZHANG, D., H. XU, P. B. SHULL, J. LIU, and X. ZHU (2015) "Somatotopical feedback versus non-somatotopical feedback for phantom digit sensation on amputees using electrotactile stimulation," *Journal of neuroengineering and rehabilitation*, **12**(1), p. 44.
- [81] MATOS, J., P. CAPSI-MORALES, and C. PIAZZA (2025) "Autonomous slip control inspired by human physiology for improved shared control strategy," *Wearable Technologies*, **6**, copyright - © The Author(s), 2025. Published by Cambridge University Press. This work is licensed under the Creative Commons

Attribution License <http://creativecommons.org/licenses/by/4.0> (the “License”). Notwithstanding the ProQuest Terms and Conditions, you may use this content in accordance with the terms of the License; Last updated - 2025-06-09.

URL <https://ezaccess.libraries.psu.edu/login?url=https://www.proquest.com/scholarly-journals/autonomous-slip-control-inspired-human-physiology/docview/3216742549/se-2>

- [82] RASPOPOVIC, S., G. VALLE, and F. M. PETRINI (2021) “Sensory feedback for limb prostheses in amputees,” *Nature Materials*, **20**(7), pp. 925–939.
- [83] VARGAS, L., H. HUANG, Y. ZHU, and X. HU (2021) “Evoked tactile feedback and control scheme on functional utility of prosthetic hand,” *IEEE Robotics and Automation Letters*, **7**(2), pp. 1308–1315.
- [84] ZOOK, Z. A., O. O. OZOR-ILO, G. T. ZOOK, and M. K. O’MALLEY (2021) “Snaptics: Low-Cost Open-Source Hardware for Wearable Multi-Sensory Haptics,” in *2021 IEEE World Haptics Conference (WHC)*, pp. 925–930.
- [85] SAVITZKY, A. and M. J. GOLAY (1964) “Smoothing and differentiation of data by simplified least squares procedures.” *Analytical chemistry*, **36**(8), pp. 1627–1639.
- [86] THRANE, G., K. S. SUNNERHAGEN, and M. A. MURPHY (2020) “Upper limb kinematics during the first year after stroke: the stroke arm longitudinal study at the University of Gothenburg (SALGOT),” *Journal of NeuroEngineering and Rehabilitation*, **17**(1), p. 76.
- [87] HUSSAIN, N., M. ALT MURPHY, and K. S. SUNNERHAGEN (2018) “Upper limb kinematics in stroke and healthy controls using target-to-target task in virtual reality,” *Frontiers in neurology*, **9**, p. 300.
- [88] MURPHY, M. A., C. WILLÉN, and K. S. SUNNERHAGEN (2011) “Kinematic variables quantifying upper-extremity performance after stroke during reaching and drinking from a glass,” *Neurorehabilitation and neural repair*, **25**(1), pp. 71–80.
- [89] CRESSMAN, E. K. and D. Y. HENRIQUES (2015) “Generalization patterns for reach adaptation and proprioceptive recalibration differ after visuomotor learning,” *Journal of neurophysiology*, **114**(1), pp. 354–365.
- [90] MARKO, M. K., A. M. HAITH, M. D. HARRAN, and R. SHADMEHR (2012) “Sensitivity to prediction error in reach adaptation,” *Journal of neurophysiology*, **108**(6), pp. 1752–1763.
- [91] HERZFELD, D. J., P. A. VASWANI, M. K. MARKO, and R. SHADMEHR (2014) “A memory of errors in sensorimotor learning,” *Science*, **345**(6202), pp. 1349–1353.



- [92] LOGAN, Z., Q. DEITRICK, and K. FITZSIMONS (2025) “Evaluating factors affecting the perception of multi-sensory vibration and skin-squeeze cues during voluntary movement,” *IEEE Transactions on Haptics*, pp. 1–8, © 2026 IEEE. Reprinted, with permission, from Logan, Z., Q. Deitrick, and K. Fitzsimons (2025) “Evaluating factors affecting the perception of multi-sensory vibration and skin-squeeze cues during voluntary movement,” *IEEE Transactions on Haptics*, pp. 1–8.
- [93] WILLIAMS, R. J. (1992) “Simple statistical gradient-following algorithms for connectionist reinforcement learning,” *Machine learning*, **8**(3), pp. 229–256.
- [94] KONDA, V. and J. TSITSIKLIS (1999) “Actor-critic algorithms,” *Advances in neural information processing systems*, **12**.
- [95] MNIH, V., A. P. BADIA, M. MIRZA, A. GRAVES, T. LILLICRAP, T. HARLEY, D. SILVER, and K. KAVUKCUOGLU (2016) “Asynchronous methods for deep reinforcement learning,” in *International conference on machine learning*, PmLR, pp. 1928–1937.
- [96] SUTTON, R. S., A. G. BARTO, ET AL. (1998) *Reinforcement learning: An introduction*, vol. 1, MIT press Cambridge.
- [97] LI, L., W. CHU, J. LANGFORD, and R. E. SCHAPIRE (2010) “A contextual-bandit approach to personalized news article recommendation,” in *Proceedings of the 19th international conference on World wide web*, pp. 661–670.
- [98] SCOTT, S. L. (2010) “A modern Bayesian look at the multi-armed bandit,” *Applied Stochastic Models in Business and Industry*, **26**(6), pp. 639–658.
- [99] SCHULMAN, J., S. LEVINE, P. ABBEEL, M. JORDAN, and P. MORITZ (2015) “Trust region policy optimization,” in *International conference on machine learning*, PMLR, pp. 1889–1897.
- [100] SCHULMAN, J., F. WOLSKI, P. DHARIWAL, A. RADFORD, and O. KLIMOV (2017) “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*.
- [101] ACHIAM, J. (2018) “Spinning Up in Deep Reinforcement Learning,” *GitHub repository*.
- [102] RAFFIN, A., A. HILL, A. GLEAVE, A. KANERVISTO, M. ERNESTUS, and N. DORMANN (2021) “Stable-Baselines3: Reliable Reinforcement Learning Implementations,” *Journal of Machine Learning Research*, **22**(268), pp. 1–8.  
URL <http://jmlr.org/papers/v22/20-1364.html>
- [103] DHARIWAL, P., C. HESSE, O. KLIMOV, A. NICHOL, M. PLAPPERT, A. RADFORD, J. SCHULMAN, S. SIDOR, Y. WU, and P. ZHOKHOV (2017), “OpenAI Baselines,” <https://github.com/openai/baselines>.

- [104] BROOKE, J. ET AL. (1996) “SUS-A quick and dirty usability scale,” *Usability evaluation in industry*, **189**(194), pp. 4–7.
- [105] LIN, J. (1991) “Divergence measures based on the Shannon entropy,” *IEEE Transactions on Information theory*, **37**(1), pp. 145–151.

**Vita**  
**Zachary Logan**

## **EDUCATION**

---

|                                                                                                    |                     |
|----------------------------------------------------------------------------------------------------|---------------------|
| <b>The Pennsylvania State University</b>                                                           | University Park, PA |
| <i>College of Engineering</i>                                                                      |                     |
| <b>Ph.D. Candidate in Mechanical Engineering</b>                                                   | May 2026            |
| <i>Dissertation: Human Perception and Information Encoding in Wearable Haptic Feedback Systems</i> |                     |
| <b>The Pennsylvania State University</b>                                                           | University Park, PA |
| <i>College of Engineering</i>                                                                      |                     |
| <b>Master of Science in Mechanical Engineering</b>                                                 | August 2025         |
| <b>The Pennsylvania State University</b>                                                           | University Park, PA |
| <i>College of Engineering</i>                                                                      |                     |
| <b>Bachelor of Science in Mechanical Engineering</b>                                               | May 2021            |

## **RESEARCH PUBLICATIONS**

---

1. Z. Logan, Q. Deitrick, and K. Fitzsimons, “Evaluating factors affecting the perception of multi-sensory vibration and skin-squeeze cues during voluntary movement,” *IEEE Transactions on Haptics*, pp. 1–8, 2025. DOI: 10.1109/TOH.2025.3572862.
2. Z. Logan, K. Undieh, and M. Goli, “Autonomous integration of bench-top wet lab equipment,” in *2025 22nd International Conference on Ubiquitous Robots (UR)*, 2025, pp. 13–19. DOI: 10.1109/UR65550.2025.11078042.