Optimization of Myoelectric Control via Context-Informed Incremental Learning

by

Evan Campbell

MScE of Electrical Engineering, UNB, 2021
BScE of Electrical Engineering, UNB, 2019

A Dissertation Submitted in Partial Fulfilment of
the Requirements for the Degree of

Doctor of Philosophy

In the Graduate Academic Unit of Electrical and Computer Engineering

Supervisor:	Erik Scheme, PhD, Electrical and Computer Engineering
Examining Board:	Jon Sensinger, PhD, Electrical and Computer Engineering Ahmed Shehata, PhD, Electrical and Computer Engineering Michael Fleming, PhD, Computer Science
External Examiner:	Christian Antfolk, PhD, Biomedical Engineering, Lund University

THE UNIVERSITY OF NEW BRUNSWICK

June, 2026

© Evan Campbell, 2026

Abstract

Myoelectric control systems translate muscle activity into control signals for external devices, offering intuitive interfaces for assistive technologies and human–computer interaction. However, these systems are highly sensitive to distributional shift—the mismatch between calibration and deployment conditions—caused by changes in limb posture, electrode placement, and, inescapably, user behaviour. This thesis introduces and evaluates Context-Informed Incremental Learning (CIIL), a generalizable adaptation framework that leverages situational task context to guide real-time model updates without requiring explicit supervision.

CIIL is validated across classification and regression control schemes, statistical models, static neural networks, and temporal neural network models, and a range of closed-loop control tasks, including target acquisition and object manipulation in virtual reality. Through experiments involving over 60 participants, CIIL is shown to outperform conventionally calibrated models and unsupervised high-confidence adaptation both in favourable conditions and under severe distributional changes, such as electrode shifts. The framework enables stable adaptation by assigning pseudo-labels based not on model confidence but on the alignment between the user's behaviour and the inferred task context in a more tolerant strategy than existing environment-dependent incremental learning approaches.

A systematic review of prior incremental learning approaches is provided, and CIIL's

compatibility with diverse model classes—including feedforward and temporal architectures—is empirically demonstrated. Extensions to velocity-based control and proportional regression further highlight CIIL's flexibility. The thesis also positions CIIL within existing context-derived adaptation literature and discusses its unique suitability for non-unique control tasks.

Limitations are acknowledged, including the exclusive use of able-bodied participants. The work concludes by outlining future directions, including principled system analysis using control theory, evaluation of user psychological metrics such as agency, and clinical deployment in stroke or neuromuscular populations. Together, this thesis advances the field by demonstrating that situational context can serve not only as a behavioural cue but also as a powerful supervisory signal for stable, scalable, and user-aligned myoelectric adaptation.

Dedication

To my wife, Nicole,

who on our long grad school journey

somehow offered me both

unwavering compassion and unrelenting competition.

Acknowledgements

There are a lot of people to acknowledge who supported me during my degree; below are just a few mentioned.

First and foremost, Nicole, thank you for your constant support over this long journey. While you yourself were also exhausted with your own studies, being a student advocate, being our wedding planner, working full-time, and being a cat-mom to two, you always made me a priority and made sure I didn't spiral uncontrollably. Next, thank you to my family for their belief that I had it in me to finish, even as the degree dragged on. Thank you, Dad, I can finally appreciate the perseverance it took to get through all your training, but your utter lack of complaints might've made me believe walking in your footsteps would've been less gruelling; and Mom, whose unwavering and often over-the-top belief in my ability to finish never faltered.

Erik, I am thankful you accepted me as your PhD student after seeing firsthand how stubborn I can be during my MSc. Thank you for your expert guidance along both these journeys and life more generally. Ethan, thank you for coming along and making research fun. Thank you as well for drawing out my social side and showing me that a little spiralling can actually be a lot of fun. Ulysse, thank you for being a mentor, a friend, and an incredibly generous host throughout the two months I spent in Norway. Finally, thank you to the collaborators I have had the fortune to work with, including those at ULaval, UiO, FAU, and the ICNR workshop.

Table of Contents

Abstract **ii**

Dedication **iv**

Acknowledgments **v**

Table of Contents **vi**

List of Tables **xvi**

List of Figures **xvii**

1 Introduction **1**

1.1 Objectives 6

1.2 Thesis Structure 6

1.3 Statement of Authorship 7

1.4 Author's Declaration of Generative AI Use 8

Bibliography 8

2 Justification and Feasibility **13**

2.1 Foreword 13

2.2 Incremental Learning Justification 15

2.3 Taxonomy of Incremental Learning Strategies 18

2.3.1 Dedicated On-Demand Recalibration (DODR) 20

2.3.2 Unsupervised Incremental Learning (UIL) 21
2.3.3 Predictor-Dependent Incremental Learning (PDIL) 24
2.3.4 Environmental-Dependent Incremental Learning (EDIL) 27
2.4 Taxonomy of Incremental Learning Goals 31
2.4.1 Domain-Incremental Continual Learning: Variations of Existing Patterns 33
2.4.2 Task-Incremental Continual Learning: New Patterns Triggering Existing Actions 34
2.4.3 Class-Incremental Continual Learning: New Patterns Triggering New Actions 35
2.5 User acceptability and user-centered design 35
2.6 Consideration of Open Challenges 37
2.6.1 Sensitivity to Noise vs. User-Behaviours 38
2.6.2 Model Complexity vs. Data Requirements 38
2.6.3 User vs. Model Adaptation 39
2.7 Establishing CIIL as a Viable and Beneficial Learning Strategy 41
Bibliography 42

3 Context-Informed Incremental Learning Improves Both the Performance and Resilience of Myoelectric Control 57
3.1 Foreword 57
3.2 Reference 58
3.3 Abstract 59
3.4 Introduction 60
3.5 Related Works 62
3.5.1 Incremental Learning in Myoelectric Control 62

3.5.2 Reinforcement Learning 64
3.5.3 Context 67
3.6 Methods 68
3.6.1 A Proxy Environment for Online Adaptation 70
3.6.2 Control Scheme 70
3.6.3 Incremental Learning Strategies 71
3.6.3.1 Negative Context-Informed Incremental Learning (N-CIIL) 73
3.6.3.2 Positive Context-Informed Incremental Learning (P-CIIL) 74
3.6.3.3 Positive + Negative Context-Informed Incremental Learning (P+N-CIIL) 74
3.6.3.4 Unsupervised High-Confidence (UHC) 74
3.6.4 Outcome Measures 75
3.6.4.1 Offline Metrics 75
3.6.4.2 Online Metrics 75
3.7 Experimental Design 76
3.7.1 Experiment Scenario 1: Minimal Training Data 77
3.7.1.1 Velocity 77
3.7.1.2 UHC Confidence Threshold 78
3.7.1.3 Negative CIIL 78
3.7.1.4 Order 78
3.7.1.5 Gameplay Time 78
3.7.2 Experiment Scenario 2: Electrode Shift 79
3.7.2.1 Velocity 79
3.7.2.2 UHC Confidence Threshold 80

3.7.2.3 Negative CIIL . . . 80
3.7.2.4 Order . . . 80
3.7.2.5 Gameplay Time . . . 80
3.7.2.6 Questionnaire . . . 80
3.7.3 Statistical Analyses . . . 81
3.8 Results . . . 81
3.9 Discussion . . . 86
3.9.1 Adaptation Dynamics . . . 86
3.9.2 Qualitative Feedback . . . 90
3.9.3 Offline vs. Online Performance . . . 90
3.9.4 Practicality of CIIL . . . 92
3.9.5 Future Work . . . 94
3.10 Conclusion . . . 95
3.11 Declarations . . . 96
3.11.1 Ethics Approval and Consent to Participate . . . 96
3.11.2 Availability of Data and Materials . . . 96
3.11.3 Funding . . . 96
3.11.4 Author Contributions . . . 97
Bibliography . . . 97

4 Screen Guided Training Does Not Capture Goal-Oriented Behaviours: Learning Myoelectric Control Mappings From Scratch Using Context-Informed Incremental Learning 106

4.1 Foreword . . . 107
4.2 Reference . . . 107
4.3 Abstract . . . 108
4.4 Introduction . . . 108

4.5 Methods . . . 112
4.5.1 Participants . . . 112
4.5.2 Hardware and Data Processing . . . 112
4.5.3 Experimental Environments . . . 115
4.5.3.1 Screen Guided Training . . . 115
4.5.3.2 Context Available Game Environment . . . 115
4.5.3.3 Fitts' Law Environment . . . 116
4.5.4 Pipelines . . . 116
4.5.4.1 Initial Training . . . 117
4.5.4.2 Online Adaptation . . . 117
4.5.5 Context Informed Incremental Learning . . . 119
4.5.5.1 Label Spreading Relabeling . . . 120
4.5.5.2 Active/Non-Active Separation via a GMM . . . 120
4.5.5.3 Adaptive Proportional Control . . . 121
4.5.6 Evaluation Criteria . . . 124
4.6 Results & Discussion . . . 125
4.6.1 Offline Metrics Only Tell a Small Part of the Story . . . 125
4.6.2 Online Performance Tells the Full Story . . . 128
4.6.3 Fast or Slow With Very Little In-Between . . . 129
4.6.4 Future Directions . . . 131
Bibliography . . . 134

5 Context-Informed Incremental Learning Improves Throughput and Reduces Drift in Regression-Based Myoelectric Control 143

5.1 Foreword . . . 143
5.2 Reference . . . 144
5.3 Abstract . . . 145

5.4 Introduction . . . 146
5.5 Background . . . 148
5.5.1 Adaptive Learning . . . 148
5.5.2 Context Informed Incremental Learning . . . 150
5.6 Methods . . . 151
5.6.1 Participants . . . 151
5.6.2 Signal Processing . . . 152
5.6.3 Data Acquisition . . . 152
5.6.4 Models . . . 154
5.6.4.1 W-SGT . . . 155
5.6.4.2 C-SGT . . . 155
5.6.4.3 O-CIIL . . . 155
5.6.4.4 T-CIIL . . . 156
5.6.5 Online Validation Task . . . 160
5.6.5.1 Throughput . . . 162
5.6.5.2 Path Efficiency . . . 162
5.6.5.3 Overshoots . . . 163
5.6.5.4 Completion Rate . . . 163
5.6.5.5 Number of Trials . . . 163
5.6.5.6 Action Interference . . . 164
5.6.5.7 Drift . . . 164
5.6.5.8 Simultaneity Gain . . . 165
5.6.6 Statistics . . . 166
5.7 Results . . . 166
5.7.1 Adaptive Models Outperform SGT Models . . . 166
5.7.2 The Effect of User Compliance Adjustments in an Adaptive Setting . . . 171

5.7.3 The Cost of Simultaneity . . . 173
5.7.4 Participants Preferred Adaptive Models . . . 173
5.8 Discussion . . . 174
Bibliography . . . 180

6 Transformer-Based Context-Informed Incremental Learning with Soft-DTW Alignment Unlocks Fast and Precise Regression-Based Myoelectric Control 190

6.1 Foreword . . . 190
6.2 Reference . . . 191
6.3 Abstract . . . 192
6.4 Introduction . . . 193
6.4.1 Shortcomings of Regression . . . 193
6.4.2 Incremental Learning for Closed-Loop Behaviours . . . 194
6.4.3 Temporal Models for Temporal Dynamics . . . 197
6.4.4 Soft Dynamic Time Warping for Pseudo-label Alignment . . . 198
6.4.5 Proposed Solution . . . 198
6.5 Methods . . . 199
6.5.1 Data Acquitision . . . 199
6.5.2 Experimental Design . . . 200
6.5.2.1 Screen-Guided Training . . . 200
6.5.2.2 User-/Co-Adaptation Phase . . . 202
6.5.2.3 Evaluation Phase . . . 204
6.5.3 Intent Recognition Pipelines . . . 206
6.5.3.1 Model Architectures . . . 207
6.5.3.2 Initial Calibration . . . 208
6.5.3.3 Incremental Learning . . . 208

6.5.4 Statistical Analyses . . . 212
6.6 Results . . . 213
6.6.1 Success Rate . . . 214
6.6.2 Throughput . . . 215
6.6.3 Path Efficiency . . . 217
6.6.4 Simultaneity Gain . . . 217
6.7 Discussion . . . 220
6.7.1 Speed-Precision Trade-off . . . 220
6.7.2 CIIL with Temporal Insights Enable Lower Contraction Intensity . . . 220
6.7.3 Closed-Loop Performance of Transformers . . . 224
6.7.4 Model Latent Analysis . . . 225
6.7.5 Future Work . . . 226
6.8 Conclusion . . . 228
Bibliography . . . 228

7 General Discussion **235**
7.1 Addressing Thesis Objectives . . . 235
7.1.1 Objective 1: To establish CIIL as a viable adaptation strategy across both favourable and adverse deployment regimes—those where unsupervised high-confidence adaptation is expected to succeed and those where it is expected to fail. . . . 235
7.1.2 Objective 2: To investigate the deployment of CIIL across a spectrum of model classes, including Bayesian classifiers, static neural networks, and temporal neural networks. . . . 236

7.1.3 Objective 3: To investigate CIIL in both classification-based and regression-based myoelectric control paradigms. 237
7.2 Further Contributions 238
7.2.1 Contribution 1: To establish an incremental learning strategy for nonlinear, adaptive proportional control within classification-based myoelectric control. 238
7.2.2 Contribution 2: To examine CIIL in conventional target acquisition environments and extend its application to more complex interaction arenas. 239
7.2.3 Contribution 3: To compile a comprehensive review of state-of-the-art incremental learning approaches for myoelectric control, with emphasis on how different strategies address distribution shift between calibration and deployment. 240
7.2.4 Contribution 4: To make incremental learning more accessible for the community by open-sourcing software into LibEMG. . 241
7.3 Positioning CIIL Within Existing Context-Derived Adaptation Strategies 244
7.3.1 Regression-Based Context Leveraging 244
7.3.2 Classification-Based Context Leveraging 245
7.4 Future Work and Limitations 247
7.4.1 Principled Analyses of Adaptation Dynamics 247
7.4.2 Psychological Aspects and Human Factors in Incremental Learning 248
7.4.3 Benchmarking CIIL Against Reinforcement Learning 248
7.4.4 Developing CIIL Strategies for Additional Myoelectric Control Arenas 249

7.4.5 Translation to a Clinical Population 250
7.5 Conclusion 251
Bibliography 252

Vita

List of Tables

2.1 Collection of incremental learning studies and general outcomes. . . . 19

6.1 Success Rates Rate (%) Across Different Target Sizes (top) with Repeated Measures ANOVA Results (bottom). 215

6.2 Throughput (bits/s) Across Different Target Sizes (top) with Repeated Measures ANOVA Results (bottom). 216

6.3 Efficiency (%) Across Different Target Sizes (top) with Repeated Measures ANOVA Results (bottom). 218

6.4 Simultaneity Gain Across Different Target Sizes (top) with Repeated Measures ANOVA Results (bottom). 219

List of Figures

1.1 A diagram highlighting the pseudo-labelling process of CIIL across 8 timesteps. In a gamified environment, a goal target is provided as a purple planet to which an orange crosshair must navigate using a gesture mapping, then remain inside to acquire it. The output of the intent recognition model is given in the class field which dictates how the crosshair moves between timesteps (hand open:up, wrist extension:right, hand close:down, wrist flexion:left, no motion:still). The context suitability is provided as positive if the model's output reduces the distance to the goal location. The goal-directed options are the set of classes which would have brought the crosshair nearer to the goal given the gesture mapping. CIIL will either include the output of the intent recognition model as the pseudo-label when the suitability is positive or override the pseudo-label with a goal-directed option when the suitability is negative. 5

2.1 Illustration of a user-evoked screen-guided training session within a myogame environment typical of a dedicated on-demand recalibration incremental learning style. 21

2.2 Illustration of basis vectors established from calibration and subsequent adjustments following collection of closed-loop signals. 23

2.3 Consider a hypothetical ten-window classification horizon for an environment aware myoelectric prosthesis capable of performing hand open (HO), hand closed (HC), wrist pronation (WP), wrist supination (WS) and no motion (NM). In this scenario, the user must first open the hand to increase the aperture and then close the hand to secure contact with an object. The top portion of the diagram illustrates the sequence of classifier predictions, associated probabilities, environmental context, and the corresponding environmental reward signal during this task. The bottom portion of the diagram shows how several predictor-dependent incremental learning algorithms (UHC, ULC, BLIP) and environment-dependent algorithms (CIIL, RL) process each sample. For UHC, ULC, BLIP and CIIL, the bottom figure illustrates the pseudo-labels assigned to each window or, when applicable, the discarding of the sample (∅). For the RL strategy, the figure indicates whether the prediction for similar EMG patterns would be reinforced (↑) or dissuaded (↓) based on the environmental reward. 26

2.4 Illustration of sources of environmental feedback suitable for environment-dependent incremental learning (EDIL). The scenes depict how cues from game environments, gaze recognition, and geometry of proximal objects can provide rich contextual information to guide model adaptation. In CIIL, such context is used to determine the set of eligible classes and assign corresponding pseudo-labels, whereas RL uses the same information to generate a continuous reward signal reflecting the suitability of the performed action. 30

2.5 Illustration of an abstract feature space representing an initial model configuration of five well-separated motion classes shown as different coloured ellipses. Domain-incremental continual learning is represented by shifts in the location of each ellipse, introduced by changes in the model inputs. Task-incremental continual learning is represented by the appearance of a new red ellipse in a novel location when we already have a red class defined, suggesting a change in control strategy. Class-incremental continual learning is represented by a previously unseen grey class appearing in a new location, implying the addition of an entirely new action. 32

3.1 Depictions of various myoelectric control system designs. Solid lines indicate an output from a node or an input to a node, and a diamond indicates a signal that modifies the weights of a node. (a) Myoelectric control systems designed using traditional practices (optimized for open-loop control, user perception closes the loop). (b) Myoelectric control systems that are context-aware (CA) and modulate the intent relayed to the controller subject to environmental information. (c) Myoelectric control systems optimized via traditional incremental learning (IL) (these systems rely on the high confidence predictions of the classifier to incorporate incoming data to the adaptation set). (d) Myoelectric control systems that are supplied contextual feedback based on the suitability of the performed action in the environment and use this information to govern incremental learning (CIIL). . . . 65

3.2 Experimental setup used throughout the study. (a) The Thalmic Labs Myo. (b) The calibration screen used in Scenario 1 collects one second of EMG data per motion class. (c) The screen-guided training interface for collecting offline data. (d) The gamified proxy environment, Myoshoot, where adaptation occurred. The user's position was represented by a crosshair and the goal was represented by a planet. (e) The Fitts' law environment where online usability was ultimately tested. 69

3.3 A representation of a hypothetical sequence of cursor movements in the control task. The crosshair indicated the cursor position at each timestep with arrows indicating the direction of travel for the next timestep. An interpretation of the steps along the trajectory is given in the right table, where the decision (class and probability) and context suitability (was the performed action correct (P or N for positive or negative) and what were the correct options under the context) served as inputs to the inclusion criteria of different algorithms. The effects of the different algorithms are given by checkmarks (include with classifier's prediction), hyphens (override the decision using context), and x's (discard the sample). 72

3.4 Box plots, summary statistics, and critical difference diagrams for offline metrics (accuracy, active error, instability), and online metrics (overshoots, efficiency, throughput) in both evaluated scenarios (minimal training data, electrode shift). Ideal trends for values and normalized ranks of these metrics are high accuracy, low active error, low instability, low overshoots, high efficiency, and high throughput. Blue bars linking groups in the critical difference diagrams signifies a non-significant difference within the multiple comparison tests. Note: when subjects did not recover from the introduced electrode shift and could not complete the online assessment, their data was removed for the overshoots metric as a placeholder value was not appropriate. . . 83

3.5 Box plots, summary statistics, and critical difference diagrams for subjective ratings of the different adaptation methods (frustrating, performance, physical effort, and overall effort). Ideal ratings and normalized ranks for these categories are low frustration, high performance, low physical demand, and low overall effort. Blue bars linking groups in the critical difference diagrams signifies a non-significant difference between linked strategies. 85

3.6 Representation of the number of context windows used by the CIIL pipelines the electrode shift scenario. Standard error across subjects is shown by a shaded region. Each point represents a three second time frame. 89

3.7 Scatterplot of associative relationships between offline accuracy and online throughput from (a) Scenario 1 (Minimal Training Data) and (b) Scenario 2 (Electrode Shift). Grouped Spearman's correlation coefficients are presented in the legend. Global correlations, as calculated using all data irrespective of model, were $\rho = 0.74, 0.62$, for Scenario 1 and 2, respectively. 91

4.1 The 8-channel, 1000 Hz, *OYMotion gForcePro+* EMG armband that was used in this study and the raw EMG from a 3-second recording. . 113

4.2 **(a)** The Unity-built scene used for the game. The white hemisphere represented the playable region which was centred around the last meteor with 170 degree range on both axes. The ship did not move forward or backward; all movements only changed the ship orientation.**(b)** The cockpit view from the Myo-Shoot 2 game. The arrow in the center indicates the direction of the planet. **(c)** The ISO 9241-9 Fitts' law testing environment. 114

4.3 Histogram of the proportional control inputs from a pilot gameplay session colored by predicted motion class. 122

4.4 The mapping between the normalized contraction intensity and cursor speed (represented as a % of the maximum gain). Both the traditional linear mapping and the proposed sigmoid mapping are shown. 123

4.5 Classifier dependent results. **(a)** Classification accuracy across different models. **(b)** Throughput progression during gameplay for SGT, SGT-A, and ZS-A models. **(c)** Fitts' Law performance comparison between models. **(d)** Confusion matrices for SGT, SGT-A (pre and post), and ZS-A (post) models. **(e)** t-SNE visualization of SGT-A model predictions before (left) and after (right) adaptation on the gameplay data. 126

4.6 Analysis of proportional control in the ZS-A model. (a) Heatmap showing the distribution of proportional control speeds over 15-second game intervals. (b) Overall speed distribution throughout the game, with frequency histogram. (c) Class-specific proportional control threshold margins over time. Vertical dashed lines indicate activation of adaptive proportional control and GMM/label spreading at 30 and 120 seconds, respectively. (d) Time series of proportional control input for a single subject over the course of the game. 130

5.1 Illustration of data acquisition setup. EMG electrodes were placed on the participant's dominant forearm. Participants sat in front of a computer screen that displayed relevant information for each task. . . 153

5.2 Illustration of pseudo-label creation for O-CIIL. Predictions are represented by dotted arrows and the corresponding pseudo-label is represented as a solid black arrow. The length of the vector, p, is determined using Equation 5.1. 157

5.3 Illustration of pseudo-label creation for T-CIIL. Predictions are represented by dotted arrows and corresponding pseudo-labels are represented by solid arrows in the same colour. The red dot, labelled "C," represents the current cursor position, and the black circle, labelled "Target," shows the current goal target. Arrows marked with the "∼" symbol represent predictions that were considered poor for adaptation and have been ignored. A) Example of main pseudo-label decision flow. Predictions in the correct quadrant (green) are relabelled to point in the optimal direction to the target. Predictions for which one component is correct (blue and purple) are relabelled to point in the direction of the correct component. Predictions where neither component is in the correct direction (orange) are ignored. The magnitude of the corresponding pseudo-label vector is determined based on the distance from the cursor to the target. B) Illustration of target runways. In the case where the cursor is in line with the target (i.e., they can reach the target by moving in a single component), only the component that requires movement to reach the target is reinforced. C) Visualization of activation threshold. For all decisions, pseudo-labels are only created for adaptation if the model input is above the activation threshold (determined based on no motion data from calibration). 158

5.4 Online ISO 9241-9 inspired Fitts' Law task [48]. Targets are drawn in black and the cursor is drawn in orange. The current goal target is filled in, while potential target locations are outlined. Green arrows annotate the two DoFs, but this was not shown to the user during the experiment. As shown by the vertical arrow, hand open/close increases and decreases aperture size, respectively. effectively moving the cursor in the vertical direction. As shown by the curved arrow, rotating the wrist rotates the cursor in the corresponding rotational direction, where the direction of the mapping was dependent on the handedness of the participant. 161

5.5 Summary of online usability metrics over the validation period. The first 20 trials were excluded as they were deemed 'user learning' trials. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. 167

5.6 The L1 loss for each model. The mean across participants is shown as a solid line, and 95% confidence intervals are shown as shaded regions. Loss after the dotted line represents the loss during the adaptation period. The sudden changes in the mean and confidence intervals at the end of the loss curves for the adaptive models are simply a byproduct of the different number of epochs per participant during the adaptation period. 169

5.7 A rolling average (N = 60 seconds) of throughput over time during the adaptation (left) and validation periods (right). The mean across participants is shown as a solid line, and the 95% confidence intervals are shown as shaded regions. 170

5.8 Visualization of the use of simultaneity per model based on model predictions during the validation period. The heatmaps show the frequency of each activation level for both DoFs at once. The total % simultaneity (see Equation 5.9) for each model is annotated in a text box below the corresponding heatmap. A colour bar, using a logarithmic scale, maps colours to the percentage of total predictions made by each model. 172

5.9 Results of the survey posed after each model was used. The distribution of answers for each question are shown. 174

6.1 (A) Experimental setup and open-loop, screen-guided training interface with Cartesian animation displayed. (B) Closed-loop target acquisition environment used for user and model adaptation. (C) ISO 9241-9 Fitts' Law evaluation environment for assessing control performance. The largest of three target sizes is shown in this example. 196

6.2 Visualization of one representative repetition of screen-guided training. The top two time series represent the two-degree-of-freedom (2-DoF; wrist flexion/extension and wrist abduction/adduction) choreography presented via the Cartesian animation. The heat maps (bottom) represent the corresponding EMG activation, as shown by the WENG feature obtained over this 60-second period across 8 channels. 201

6.3 Visualization of the transformer pipeline. A horizon of 10 consecutive feature windows was independently embedded through a linear layer to 16 dimensions. Relative positioning information was injected over the horizon using SAPE. Subsequently, a 4-layer multi-headed attention network extracted temporal dynamics, which ultimately was processed by a single linear layer to render a 2-DoF prediction. . . . 205

6.4 Visualization of R matrix of sDTW loss function. Blue and green lines represent the two degrees of freedom of the control space. An ideal label trajectory (obtained via screen-guided training) is shown on the vertical axis, and emulated predictions from a participant eliciting "all-or-nothing" behaviours (where they struggle with the continuous transition between directions of the DOF) are shown on the horizontal axis. 211

6.5 Success rate, throughput, efficiency, and simultaneity gain outcome measures from the evaluation phase presented as box plots. Each sample represents the average metric value obtained across the 16 target acquisitions at each target size per subject. 213

6.6 Throughput over User-/Co-Adaptation and Evaluation Phases across all subjects. The model's throughput was represented via the average throughput for each trial across subjects, where the shaded region represents the standard error. 215

6.7 Paths traversed during the evaluation phase for one representative subject for all four models and three difficulties. Longer distances between neighbouring points represent higher speeds. 219

6.8 Trajectories of mean absolute values (MAV) from windows observed within the user-/co-adaptation phase of four models. MAVs were transformed into subject-specific z-scores with statistics computed from the data pooled from all models. User-/Co-adaptation phase time was not uniform between models or subjects, so each pass through was normalized to 100 temporal bins. A horizontal dashed line at 0 depicts the mean intensity of the subjects across all models. A black line shows the average MAV for time segments across the adaptation phase. 222

6.9 Quiver plots visualizing the latent space of MLP-CIIL and T-sDTW-CIIL models from the same representative subject. Each latent space was taken from the penultimate layer output of the corresponding model. The direction and length of the arrows depict the prediction of the model for that input. Plots that are the same colour represent the same data processed by different models (i.e, column 2 represents looking back on SGT data after the model was adapted, and column 3 represents looking ahead to co-learning phase data without having performed adaptation). 223

7.1 Overview of the incremental learning framework integrated into LibEMG. The system comprises four asynchronous processes: the Online Predictor, Environment, Memory Manager, and Adaptation Manager. The Online Predictor performs real-time inference and dynamically swaps to updated models when the Adapt Flag is triggered. The Environment consumes model outputs to drive task dynamics (e.g., cursor control) and returns Environment Feedback used for context-informed incremental learning or reinforcement learning. The Memory Manager associates model inputs with environment feedback and persists input–feedback pairs as adaptation-ready data. The Adaptation Manager incrementally updates the model using newly accumulated memory chunks, signals model availability to the Online Predictor, and coordinates data collection during adaptation. 243

Chapter 1

Introduction

Surface electromyography (sEMG) is a biosignal modality that captures electrical signals that arise during voluntary muscle contractions via non-invasive recording, providing a rich physiological signal from which user intent can be inferred [1]. Leveraging these muscle potentials, myoelectric control provides an intuitive and flexible human–computer interface with applications spanning assistive technologies—such as upper- and lower-limb prostheses [2], powered wheelchairs [3], and other assistive or rehabilitation devices—to broader ubiquitous-computing settings, including gesture-based interaction and hands-free control of digital systems [4, 5]. The performance of any myoelectric interface is shaped by hardware design choices, including electrode configuration and spacing [6], the extent of spatial coverage across functional muscle groups [7], instrumentation bandwidth and sampling rate [8], and the quality of electrode–skin contact [9]. The hardware of available sEMG devices is quite variable due to the different cost-benefit tradeoffs associated with their respective use cases [10, 11, 12, 13, 14]; however, all systems ultimately strive for the same thing: accurate and reliable intent recognition. Today, this decoding is accomplished most often using machine learning models that learn the patterns of

sEMG signals associated with different muscle activations, such as forearm, wrist, and hand gestures [15]. Effective myoelectric control requires these models to be properly calibrated to the user's unique characteristic muscle patterns and that they remain robust when translated to the real-world, closed-loop settings where these systems operate.

The calibration of these systems follows a broadly similar structure across both classification-based approaches, which map sEMG recordings to categorical motion classes (e.g., wrist flexion) [16], and regression-based approaches, which output continuous control signals indicating the intensity of movement along one or more degrees of freedom (e.g., $(\text{flexion/extension}, \text{abduction/adduction}) = (0.95, -0.2)$) [17]. In each case, these systems are typically trained on a dataset collected while the user performs a prescribed choreography of movements, with myoelectric signals recorded alongside labels that reflect either the intended motion class or the target contraction intensity. These labels may be generated through screen-guided training, in which visual cues instruct the user which motion to perform and at what level [18], or through prosthesis-guided training, where the controllable device itself moves through pre-programmed patterns and the user attempts to match its behaviour [19]. However, the data gathered during these routines often contain temporal misalignments between muscle activation and labels [20], inconsistencies in how users track the required contraction intensity (a form of spatial error) [21], and behavioural artifacts such as all-or-nothing contractions or anticipatory activation that degrade model fidelity [18]. Even nominally "static" models used in classification settings depend on transient muscle dynamics, as transitions between motion classes frequently contain information essential for accurate intent recognition [22]. Importantly, mitigating these issues during calibration does not eliminate the broader challenge: once the user enters the closed-loop control environment, where the model's outputs are being shown to the user as feedback, new and un-

foreseen muscle activation patterns emerge, producing out-of-distribution samples that the calibrated model was never exposed to, and thereby undermining real-time performance [23, 24, 25].

The behavioural variability that emerges once a user actively controls a myoelectric device highlights a fundamental gap between supervised calibration data and the distribution of signals encountered during real-time operation [23, 25]. Rather than treating these novel behaviours merely as confounds that must be tolerated, they may be seen as an opportunity. By harvesting user-in-the-loop data and incorporating it into the model, the system can progressively reduce the mismatch between training and deployment conditions. Incremental learning has therefore emerged as a prominent strategy for adapting myoelectric controllers online [26]. Most existing approaches rely on predictor-dependent incremental learning, in which the model's own output is used as a pseudo-label, either directly or after filtering to retain only high-confidence predictions [27, 28]. This strategy enables rapid supervised updates that are effective when the model is already performing well; however, because its adaptation signal depends on the model's current predictions, it inherently fails to correct errors arising from out-of-distribution behaviours that the model misclassifies [29]. Reinforcement-learning-based approaches offer the opposite trade-off: by using environment-derived reward signals, they can eventually correct the model even when its predictions are initially unreliable [30]. Yet, reinforcement learning (RL) is notoriously sample-inefficient, and its reliance on repeated upweighting and downweighting of actions [31] is poorly suited to human-in-the-loop settings where slow convergence and exploratory behaviour impose significant burdens on the user. These limitations point to the need for an intermediate solution—one that couples the speed and stability of supervised incremental updates with an environment-derived signal that remains reliable even when the model's predictions are not, thereby enabling adaptation to novel behaviours without incurring the inefficiencies

of reinforcement learning.

To address this gap between calibration data and user-in-the-loop behaviour, this dissertation introduces and develops context-informed incremental learning (CIIL) as an incremental learning framework. CIIL is an environment-dependent strategy that exploits task context to derive weak supervisory signals during real-time control, combining the sample efficiency of supervised updates with the robustness of reinforcement learning-style feedback. In an instrumented environment (e.g., a target acquisition task, such as Fitts' Law or a game-based interface), each myoelectric prediction is evaluated against assumed task goals using simple geometric rules, yielding a binary notion of context suitability (positive when the action advances the cursor towards the target or stabilizes it within bounds; negative when it moves away or destabilizes control) and a set of goal-directed options (e.g., directions along the cursor–target diagonal). For each incoming EMG window, the system stores the features, the current model prediction, the context tags, and a pseudo-label in a memory buffer; when suitability is positive, the model's prediction is trusted as the pseudo-label, whereas for negative outcomes, the label is reassigned based on the goal-directed options. This process is briefly outlined in Figure 1.1, and explained in greater detail in subsequent chapters. Periodically, the intent recognition model is updated on this buffer using standard supervised optimization, thereby harvesting and organizing the very behaviours that challenge conventional calibration and progressively aligning the controller with the user's evolving strategies.

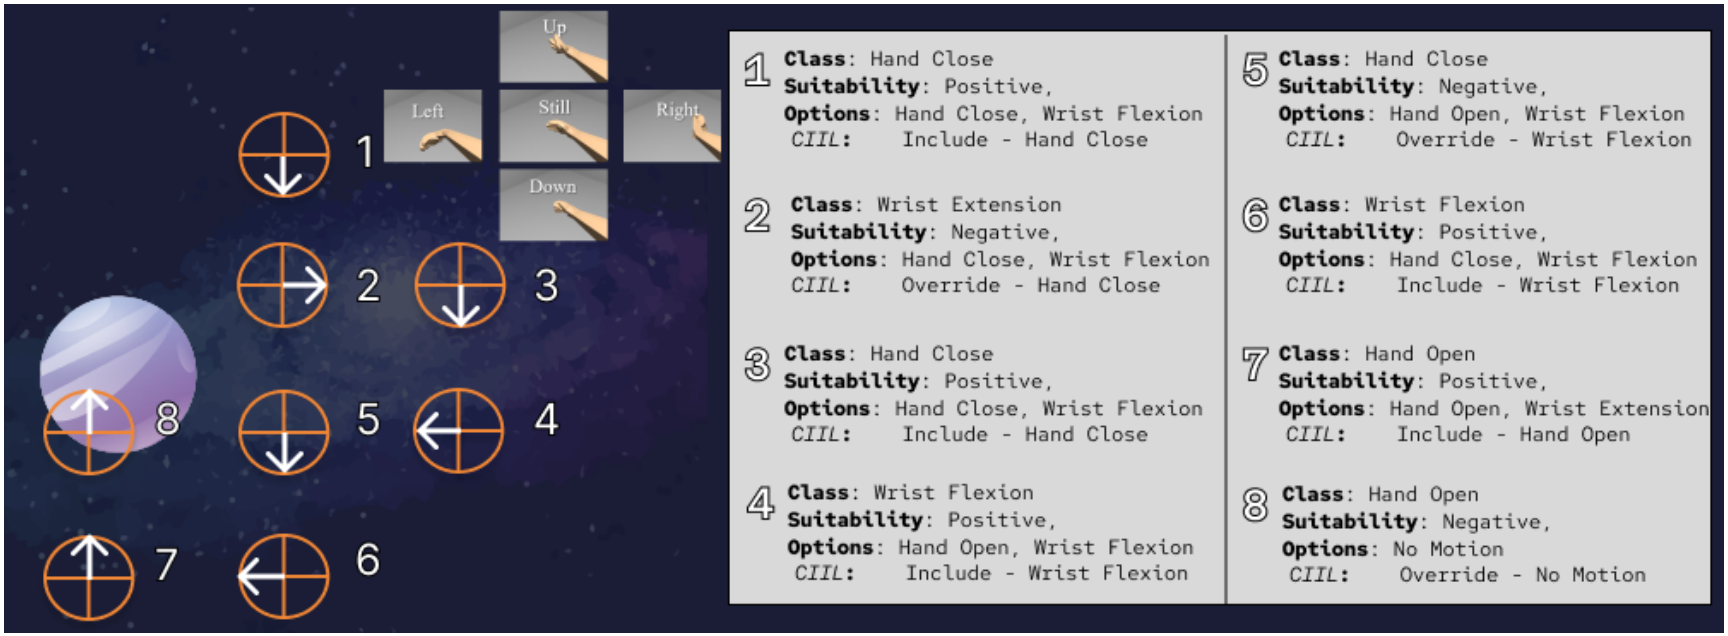


Figure 1.1: A diagram highlighting the pseudo-labelling process of CIIL across 8 timesteps. In a gamified environment, a goal target is provided as a purple planet to which an orange crosshair must navigate using a gesture mapping, then remain inside to acquire it. The output of the intent recognition model is given in the class field which dictates how the crosshair moves between timesteps (hand open:up, wrist extension:right, hand close:down, wrist flexion:left, no motion:still). The context suitability is provided as positive if the model's output reduces the distance to the goal location. The goal-directed options are the set of classes which would have brought the crosshair nearer to the goal given the gesture mapping. CIIL will either include the output of the intent recognition model as the pseudo-label when the suitability is positive or override the pseudo-label with a goal-directed option when the suitability is negative.

1.1 Objectives

This dissertation describes three interconnected objectives that outline a roadmap for establishing CIIL as a viable online learning strategy. These objectives were:

1. To establish CIIL as a viable adaptation strategy across both favourable and adverse deployment regimes—those where unsupervised high-confidence adaptation is hypothesized to succeed and those where it is hypothesized to fail. This is discussed in Chapter 3.

2. To investigate the deployment of CIIL across a spectrum of model classes, including Bayesian classifiers, static neural networks, and temporal neural networks. This is explored across Chapters 3, 4, 5, and 6.

3. To investigate context-informed incremental learning alongside both classification-based and regression-based myoelectric control paradigms. Classification-based paradigms are explored in Chapters 3 and 4, and regression-based control paradigms are explored in Chapters 5 and 6.

1.2 Thesis Structure

This dissertation is structured as a paper-based thesis.

- *Chapters 3 through 6* are published or under-review manuscripts that provide their own self-contained motivation and synthesis of findings. Each manuscript-based chapter includes a foreword that states the relevance of the work and the candidate's role within the research.

- *Chapter 2* provides an amended review of myoelectric incremental learning literature, which overviews the justification and feasibility of context-informed

incremental learning for myoelectric control, published in the IEEE Transactions on Neural Systems and Rehabilitation Engineering after the experiments of Chapters 3-5 were published.

- *Chapter 3*, published in the Journal of NeuroEngineering and Rehabilitation, presents the initial formulation of classification-based CIIL for statistical models and its effectiveness under both a limited calibration setting and an electrode shift setting.

- *Chapter 4*, published in the IEEE Transactions on Neural Systems and Rehabilitation Engineering, extends classification-based CIIL to neural network architectures in both a full calibration setting and a zero initialization setting.

- *Chapter 5*, published in the IEEE Transactions on Neural Systems and Rehabilitation Engineering, extends CIIL to the more challenging regression-based myoelectric control framework.

- *Chapter 6* extends CIIL in the regression-based myoelectric control setting to use temporal models (transformers) and more specialized loss functions, producing a drastically improved control experience. This work is currently under peer review.

- *Chapter 7* provides a synthesis of the results across the manuscripts, discusses its overall contributions, and acknowledges limitations and directions of future work.

1.3 Statement of Authorship

I was the primary author of the manuscripts of Chapters 2, 3, 4, and 6. For the review paper in Chapter 2, I co-organized the formative workshop at the 2024 Inter-

national Conference on NeuroRehabilitation and was responsible for the ideation, general structure, a majority of the manuscript sections, and revisions. For Chapters 3, 4, and 6, I was responsible for the problem statement, hypothesis, experimental design, data collection, data analysis, manuscript preparation, and revisions. For Chapter 5, I was involved in the problem statement, hypothesis, experimental design, and manuscript revisions; my role initially began as the primary researcher but transitioned to a supervisory capacity when the MSc student joined the project.

1.4 Author's Declaration of Generative AI Use

OpenAI's ChatGPT tool (versions 4.0-5.1) was used throughout the thesis to rewrite passages to improve clarity. All refined passages have been reviewed and edited to ensure that the final presentation reflects the original intent of the text.

Bibliography

[1] S. Lake, M. Bailey, and A. Grant, "Method and apparatus for analyzing capacitive EMG and IMU sensor signals for gesture control," Mar. 29 2016. US Patent 9,299,248.

[2] A. Fleming, N. Stafford, S. Huang, X. Hu, D. P. Ferris, and H. H. Huang, "Myoelectric control of robotic lower limb prostheses: a review of electromyography interfaces, control paradigms, challenges and future directions," *Journal of neural engineering*, vol. 18, no. 4, p. 041004, 2021.

[3] A. Phinyomark, P. Phukpattaranont, and C. Limsakul, "A review of control methods for electric power wheelchairs based on electromyography signals with

special emphasis on pattern recognition," *IETE Technical Review*, vol. 28, no. 4, pp. 316–326, 2011.

[4] T. S. Saponas, D. S. Tan, D. Morris, R. Balakrishnan, J. Turner, and J. A. Landay, "Enabling always-available input with muscle-computer interfaces," in *Proceedings of the 22nd annual ACM symposium on User interface software and technology*, pp. 167–176, 2009.

[5] P. Kaifosh and T. R. Reardon, "A generic non-invasive neuromotor interface for human-computer interaction," *Nature*, vol. 645, no. 8081, pp. 702–711, 2025.

[6] A. J. Young, L. J. Hargrove, and T. A. Kuiken, "The effects of electrode size and orientation on the sensitivity of myoelectric pattern recognition systems to electrode shift," *IEEE transactions on biomedical engineering*, vol. 58, no. 9, pp. 2537–2544, 2011.

[7] Y. Geng, X. Zhang, Y.-T. Zhang, and G. Li, "A novel channel selection method for multiple motion classification using high-density electromyography," *Biomedical engineering online*, vol. 13, no. 1, p. 102, 2014.

[8] A. Phinyomark, R. N. Khushaba, and E. Scheme, "Feature extraction and selection for myoelectric control based on wearable emg sensors," *Sensors*, vol. 18, no. 5, p. 1615, 2018.

[9] J. Lévesque, F. Chamberland, E. Scheme, and B. Gosselin, "Advancing myoelectric prostheses: Efficacy of gold-plated 3d-printed thermoplastic dry electrodes," *IEEE Transactions on Biomedical Engineering*, 2024.

[10] Coapt, LLC, "Handbook (v6.1) complete control system gen2." `https://coaptengineering.com/wp-content/uploads/Coapt-Gen2-Handbook.pdf`. Accessed: Feb 2025. See also `https://coaptengineering.com`.

[11] OyMotion, "Oymotion gforcepro+ emg armband." https://www.oymotion.com/en/product32/149. Accessed: November 2025.

[12] Delsys, "Delsys avanti sensor." https://delsys.com/product/trigno-avanti/. Accessed: November 2025.

[13] Pison Technology, "Pison timex neural wristband." https://pison.com/science/. Accessed: November 2025.

[14] Meta, "Emg wristbands and technology — meta." https://www.meta.com/emerging-tech/emg-wearable-technology/?srsltid=AfmBOoo49OHLQb89n7givKIz8Tp1q2wqTEo607D0LDrlrsFScssn5lCg. Accessed: November 2025.

[15] A. J. Young, L. H. Smith, E. J. Rouse, and L. J. Hargrove, "A comparison of the real-time controllability of pattern recognition to conventional myoelectric control for discrete and simultaneous movements," *Journal of neuroengineering and rehabilitation*, vol. 11, no. 1, p. 5, 2014.

[16] E. J. Scheme and K. B. Englehart, "Validation of a selective ensemble-based classification scheme for myoelectric control using a three-dimensional fitts' law test," *IEEE transactions on neural systems and rehabilitation engineering*, vol. 21, no. 4, pp. 616–623, 2012.

[17] A. Ameri, M. A. Akhaee, E. Scheme, and K. Englehart, "Regression convolutional neural network for improved simultaneous emg control," *Journal of neural engineering*, vol. 16, no. 3, p. 036015, 2019.

[18] C. Morrell, E. Campbell, and E. Scheme, "Exploring user compliance in the training of regression-based myoelectric control," in *Myoelectric Controls Symposium*, 2024.

[19] B. Lock, A. M. Simon, K. Stubblefield, and L. J. Hargrove, "Prosthesis-guided training for practical use of pattern recognition control of prostheses," in *Proceedings of the Myoelectric Controls Symposium (MEC)*, pp. 61–64, 2011.

[20] A. T. Wang, C. D. Olsen, W. C. Hamrick, and J. A. George, "Correcting Temporal Inaccuracies in Labeled Training Data for Electromyographic Control Algorithms," *IEEE International Conference on Rehabilitation Robotics*, 2023. Publisher: IEEE Computer Society ISBN: 9798350342758.

[21] T. N. Tully, C. J. Thomson, G. A. Clark, and J. A. George, "Validity and impact of methods for collecting training data for myoelectric prosthetic control algorithms," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2024.

[22] A. Hagengruber, U. Leipscher, B. M. Eskofier, and J. Vogel, "A New Labeling Approach for Proportional Electromyographic Control," *Sensors 2022, Vol. 22, Page 1368*, vol. 22, p. 1368, Feb. 2022. Publisher: Multidisciplinary Digital Publishing Institute.

[23] E. Campbell, E. Eddy, X. Isabel, S. Bateman, B. Gosselin, U. Cote-Allard, and E. Scheme, "Screen guided training does not capture goal-oriented behaviours: Learning myoelectric control mappings from scratch using context informed incremental learning," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2024.

[24] D. Yeung, I. M. Guerra, I. Barner-Rasmussen, E. Siponen, D. Farina, and I. Vujaklija, "Co-adaptive control of bionic limbs via unsupervised adaptation of muscle synergies," *IEEE Transactions on Biomedical Engineering*, vol. 69, no. 8, pp. 2581–2592, 2022.

[25] J. W. Sensinger and S. Dosen, "A review of sensory feedback in upper-limb prostheses from the perspective of human motor control," *Frontiers in neuroscience*, vol. 14, p. 345, 2020.

[26] E. Campbell, F. Egle, M. Oßwald, U. Côté-Allard, P. M. Pilarski, N. Boccardo, R. Meattini, I. Vujaklija, L. Hargrove, M. Canepa, *et al.*, "(un) supervised (co) adaptation via incremental learning for myoelectric control: Motivation, review, and future directions," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2025.

[27] J. W. Sensinger, B. A. Lock, and T. A. Kuiken, "Adaptive pattern recognition of myoelectric signals: exploration of conceptual framework and practical algorithms," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 17, no. 3, pp. 270–278, 2009.

[28] X. Jiang, C. Ma, and K. Nazarpour, "Posture-invariant myoelectric control with self-calibrating random forests," *Frontiers in Neurorobotics*, vol. 18, p. 1462023, 2024.

[29] K. Szymaniak, A. Krasoulis, and K. Nazarpour, "Recalibration of myoelectric control with active learning," *Frontiers in Neurorobotics*, vol. 16, p. 1061201, 2022.

[30] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. The MIT Press, second ed., 2018.

[31] R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," *Machine learning*, vol. 8, no. 3, pp. 229–256, 1992.

Chapter 2

Justification and Feasibility

2.1 Foreword

This chapter outlines the background necessary to contextualize the dissertation in the broader research field. It does so by expanding upon and drawing on sections from a review paper on the state of incremental learning in myoelectric control that was organized following a workshop hosted at the 2024 International Conference on NeuroRehabilitation (ICNR). As a co-organizer of the workshop and lead author of the paper, my role was recruiting the co-authors, establishing the framing and structure of the manuscript, producing sections 2.2-2.4.3, coordinating sub-sections with the co-authors, and conducting all manuscript revisions. Although this chapter draws heavily from that paper, some concepts not related to incremental learning have been removed for clarity, whereas additional diagrams have been added to support incremental learning concepts.

The original peer-reviewed article was published under ©2025 IEEE:

E. Campbell, F. Egle, M. Oßwald, U. Côté-Allard, P. Pilarski, N. Boccardo, R. Meat-

tini, I. Vuljaklija, L. Hargrove, M. Canepa, E. Eddy, A. Del Vecchio, C. Castellini and E. Scheme, *(Un)supervised (Co)adaptation via Incremental Learning for Myoelectric Control: Motivation, Review, and Future Directions,* Transactions on Neural Systems and Rehabilitation Engineering, vol. 33, no. 1, p. 3565, 2025.

2.2 Incremental Learning Justification

Upper-limb myoelectric control has a rich and enduring history extending back to the 1940s [1]. Initially, these systems used a conventional amplitude control paradigm (often referred to as direct control), wherein the rectified and low-pass filtered signal from one or two channels of bipolar surface electromyography (EMG) was used to modulate the aperture of a prosthetic hand [2]. As the field evolved, manual mode switching was integrated—achieved either by a button press or the simultaneous co-contraction of two EMG channels—to facilitate the control of multiple degrees of freedom, a strategy that continues to be widely prescribed in clinical settings. Despite the pervasive use of these direct control schemes, the advent of pattern recognition-based approaches, which are now available commercially and are covered under several national insurance programs, has marked a milestone in prosthetic technology, pursuing a more physiologically appropriate control strategy by leveraging the synergies of a larger number of residual muscles [3, 4, 5, 6, 7].

Pattern recognition systems offer more versatility than simple direct control methods by facilitating the direct selection of any available movement–often referred to as a motion class– at each time step. However, the performance of these systems is highly dependent on their initial calibration procedures, during which users provide a choreographed sequence of motions guided by on-screen instructions or prosthesis-guided training [8] to establish a robust training dataset mapping EMG patterns to intended motions or gestures. These calibration sessions are critical, as they form baseline models that drive the subsequent control and daily operation of the prosthetic device. When operating conditions closely match those observed during calibration, the models routinely achieve high accuracy (often in the high 90% range) thanks to advancements in feature engineering [9] and instrumentation [10]. Since the early 2000s, significant efforts have been directed toward mitigating the

influence of confounding factors such as limb position [11], contraction intensity [12], electrode shift [13], and fatigue [14], all of which can contaminate the EMG signal and degrade the laboratory assessed performance of both pattern recognition and conventional direct control systems. Through dedicated feature engineering [15] and specialized calibration procedures [16], many of these sources of degradation have been substantially improved, a feat that conventional direct control strategies lack the sophistication to accommodate.

Despite these advancements, the utility of pattern recognition-based approaches has been questioned in terms of functional outcomes [17, 18, 19, 20, 21]. Furthermore, the abandonment rates of myoelectric prostheses, regardless of how they are controlled, remain undesirably high [22]. Periodic retraining is required to accommodate non-stationarities introduced by physiological factors such as blood flow [23], as well as environmental variations including changes in socket fit [24] and skin-electrode contact characteristics [25], which introduce challenges for models to remain usable over time [26]. Further, inadequate spatial and temporal compliance with the prompts, common during regressor training, can manifest in models that drift at low intensities and elicit inadvertent off degree-of-freedom activity [27, 28]. Moreover, users often generate much more varied patterns of EMG signals that diverge from structured, choreographed patterns captured during calibration, especially when focused on completing real-world goal-oriented tasks rather than simply replicating prompted training gestures [29]. Importantly, optimizing classification accuracy in offline settings serves only as a reasonable proxy for functional performance during the early stages of model improvement; after high accuracy is achieved, further improvements in classification accuracy do not necessarily translate to better functional outcomes. This disconnect has fueled ongoing debates, with studies revealing both weak [30], strong [31], and coarse [32] correlations between offline and online performance metrics. Consequently, the field is left searching for strategies that

more directly optimize functional performance, leveraging data that naturally arises from user-in-the-loop scenarios to update and refine myoelectric control systems.

Recognizing the limitations of static pattern recognition (wherein a model is trained, frozen, and then applied), incremental learning strategies [33] have attracted growing attention as a means to continuously adapt myoelectric models to the evolving nature of user behaviours and signal variability. Incremental learning represents a paradigm shift: rather than relying on a fixed dataset collected during a dedicated training session, these methods allow models to be updated, refined, and improved as the user interacts with the system. Over time, such adaptive approaches may maintain more robust intent recognition accuracy despite shifts in sensor signals [34], user behaviours [29], and external environments [35]. As a result, incremental learning has become a core focus in efforts to develop long-lasting, user-centric, and adaptable myoelectric control systems [36].

This manuscript aims to provide a unifying taxonomy and perspective of incremental learning methods for myoelectric control. We categorize existing incremental learning strategies into four key categories: (1) dedicated on-demand recalibration (DODR), (2) unsupervised incremental learning (UIL), (3) predictor-dependent incremental learning (PDIL), and (4) environment-dependent incremental learning (EDIL). We also present a taxonomy of incremental learning studies based on the intended outcome, consisting of: (1) domain-incremental continual learning, (2) task-incremental continual learning, and (3) class-incremental continual learning. Finally, we share our perspective on several open challenges associated with incremental learning, namely, sensitivity to signal changes from noise vs. user behaviours, model complexity vs. data requirements, and user learning vs. model adaptation. Ultimately, this chapter identifies emerging directions that promise to shape the future of incremental learning in myoelectric prosthesis control.

2.3 Taxonomy of Incremental Learning Strategies

A taxonomy of existing incremental learning literature resulted, yielding the following categories: (1) dedicated on-demand recalibration, (2) unsupervised incremental learning, (3) predictor-dependent incremental learning, and (4) environment-dependent incremental learning. A summary of general details from exemplar incremental learning studies is given in Table 2.1.

Table 2.1: Collection of incremental learning studies and general outcomes.

Dedicated On-Demand Recalibration (DODR)

Ref	Algorithm	Frequency	Subjects	Test	Outcome	Compared Against
[37]	SH, SL, SHL, SA	BT	7-AB, 4-LD	OA	Reduced error by at least 26% compared to non-adapting classifier	NAC
[38]	iRR-RFF	BT	10-AB	OA, (RTC)	iRR-RFF outperformed iRR, incremental updates support a stable level of performance	iRR, KRR
[39]	iRR-RFF	BS	1-LD	RTC	Mastering increasingly complex ADLs over 13 months	-
[40, 41]	iRR	BT	6-AB, 5-LD	TAC	Improving or stable objective and subjective myocontrol scores over four sessions	-
[36]	RF-RLSC	BS	6-AB, 1-LD	OA	Maintained stable accuracy over months	LDA, kNN, RF-SVM, RF, MLP
[42]	SA	BT	3-LD	OA	Users needed to recalibrate less often and reported improved quality	NAC

Unsupervised Incremental Learning (UIL)

Ref	Algorithm	Frequency	Subjects	Test	Outcome	Compared Against
[43]	NNMF	N/A	7-AB, 7-LD	TAC	NNMF outperformed mode-switching pattern recognition (MSRPR)	MSPR
[44]	sNNMF	N/A	8-AB, 2-LD	TAC	Higher throughput and shorter completion time than NNMF or LR	NMF, LR
[45]	iNNMF	C	10-AB, 2-LD	TAC, RCT	Mitigated degradation due to electrode shift, comparable robustness to DODR	NAC, DODR
[46]	iNNMF	C	8-AB	TAC	Comparable performance with DODR, but higher workload	iRR-RFF
[47]	piNNMF	C	10-AB, 2-LD	TAC	Possibility to progressively activate new actions	iNNMF
[48]	iNNMF	C	8-AB	TAC	Up to six controllable muscle synergies utilizing higher-density sEMG	-
[49]	SotF	C	3-AB	OA	SotF reduced learning time by 89% while maintaining accuracy	NAC

Predictor-Dependent Incremental Learning (PDIL)

Ref	Algorithm	Frequency	Subjects	Test	Outcome	Compared Against
[50]	UB	C	2-AB, 1LD	RTC	UB improved stability	NAC
[37]	UH, UL, UB	BT	7-AB, 4-LD	OA	Provided small reductions in error, UH showed promise	NAC
[51]	MV	BT	DS	OA	Improved performance over no recalibration	NAC
[52]	SCADANN	BT	DS	OA, TAC	Outperformed other methods across weeks, but performance below supervised recalibration	NAC, MV, AdaBN, UDA-M, SR
[53]	scRF	P	66-AB	OA	scRF gradually improved accuracy over 5 weeks, visual feedback reduced muscle effort	NAC, RF, LDA
[54]	dfRLS	C	5-AB	TAC	Improved adaptation stability over exponential forgetting	efRLS, User Learning
[34]	SS-BSL	C	6-AB	OA	Maintained classification accuracy over time	aLDA, ST-S3VM, GCM
[55]	LCU, SMU, ER	BT	12-AB, 2-LD	OA	Improved decoder adaptation with only 3.2 minutes of selected data	Random Sampling

Environment-Dependent Incremental Learning (EDIL)

Ref	Algorithm	Frequency	Subjects	Test	Outcome	Compared Against
[56]	CCLR	C	10-AB, 2-LD	RTC	Improved simultaneity	NAR
[57]	Closed-Loop Feedback	C	20-AB, 4-LD	TAC	Improved path efficiency and success rate for amputee	NAR
[58]	Environment Feedback	C (5m)	16-AB, 4-LD	TAC	Improved classifier performance	NAC
[35]	CIIL	C (2m/5m)	32-AB	TAC	Robust adaptation to catastrophic changes (electrode shift)	UHC, NAC
[29]	CIIL	C (10m)	16-AB	TAC	Zero-shot models outperformed SGT models within 200s; better proportional control.	NAC
[59]	CIIL	C	9-AB	RTC	Reduce transport time by ∼50%	NAC
[60]	tCIIL, oCIIL	C	16-AB	TAC	tCIIL reduced drift and action interference over oCIIL, NAR	NAR
[61]	ACRL	C	1-AB	RTC	Improved control accuracy, also cf. [62]	NAR
[63]	ACRL	C (300s)	3-AB	RTC	Non-amputated arm provided guidance for training, also cf. [64]	-
[65]	General Value Functions	C (3x3m)	4-AB, 1-LD	RTC	Reduced switching effort and improved control efficiency, also cf. [66, 67, 68]	Fixed switching order
[69]	AWAC	C	15-AB	OA	AWAC provided ∼2x increase in decoding accuracy compared to user learning	NAC
[70]	Adaptive Wiener Filter	C	14-AB	TAC	Slower adaptation improved user adaptation	Wiener Filter
[71]	Gaze-Guided Labels	C	5-AB	TAC	Eliminated the need for instrumented environment though an instrumented user	NAR

Algorithm Terms: supervised high confidence (SH), supervised low confidence (SL) supervised high and low confidence (SHL), supervised all (SA), non-adapting classifier (NAC), incremental Ridge Regression with Random Fourier Features (iRR-RFF), random features regularized least squares for classification (RF-RLSC), linear discriminant analysis (LDA), k-nearest neighbors (kNN), random feature support vector machine (RF-SVM), random forest (RF), multi-layer perceptron (MLP), sparse non-negative matrix factorization (sNNMF), *progressive* incremental non-negative matrix factorization (*pi*NNMF), survival of the fittest (SotF), unsupervised high confidence (UHC), unsupervised low confidence (ULC), unsupervised blip (UB), self-correcting random forest (scRF), directional-forgetting recursive least squares (dfRLS), exponential-forgetting recursive least squares (ef-RLS), semi-supervised Bayesian sequential learning (SS-BSL), adaptive linear discriminant analysis (aLDA), Self-Calibrating Asynchronous Domain Adversarial Neural Network (SCADANN), Unsupervised Domain Adversarial-Methods (UDA-M), Multi-votes (MV), Supervised Recalibration (SR), self-training semi-supervised support vector machine (ST-S3VM), Gaussian classification model (GCM), least confidence uncertainty (LCU), smallest margin uncertainty (SMU), entropy reduction (ER), co-adaptive closed loop regression (CCLR), context informed incremental learning (CIIL), tolerant context informed incremental learning (tCIIL), oracle context informed incremental learning (oCIIL), actor-critic reinforcement learning (ACRL), advantage weighted actor-critic (AWAC), non-adapting regressor (NAR). **Frequency Terms**: between trials (BT), between sessions (BS), continuous (C), periodic (P), minutes (m), seconds (s). **Subject Terms**: able-bodied (AB), limb difference (LD), dataset (DS). **Test Terms**: offline accuracy (OA), target achievement control (TAC), real-time control (RTC).

2.3.1 Dedicated On-Demand Recalibration (DODR)

On-demand recalibration is possibly the simplest form of incremental learning. It involves performing a full or partial data collection each time the system performance is deemed to have become suboptimal due to natural changes in the user's signals or user feedback. Alternatively, this approach can be useful to extend the capabilities of a model when the user attempts to control the prosthesis in situations or environments that may not have been represented in the lab or clinic. This strategy is enforced by prompting the user to perform one or more rounds of the actions of interest (e.g., power grasp, wrist pronation, etc.). The stimulus can be provided either graphically on a laptop or smartphone (Figure 2.1), including the visualization of the signal patterns, or by having the prosthesis itself guide the actions following a pre-programmed regimen. Both commercially available myoelectric control systems, the Coapt Gen2 by Coapt LLC [72] and the Myo Plus by Ottobock [73], provide such a (re)calibration procedure, which the user or attending clinician can request at any time. It should be noted that the Coapt Gen2 system also incorporates a continuous learning capability (Adaptive Advance® [72]), enabling ongoing refinement of the control model during use. By contrast, the Myo Plus system discards previous models and fully re-trains during each recalibration.

Dedicated on-demand incremental learning approaches build on previously trained models [38, 39, 40] and allow the model to be updated sample by sample without the need to buffer or store previous data. The very low computational and storage requirements of these models enable a very lightweight implementation with real-time updates while the prosthesis is in operation. Although simple, such a procedure provides a way to keep the model up-to-date with the user's latest signals. Although preliminary data are promising [42], more evidence is needed to validate the approach. To the best of our knowledge, no large-scale comparative studies have

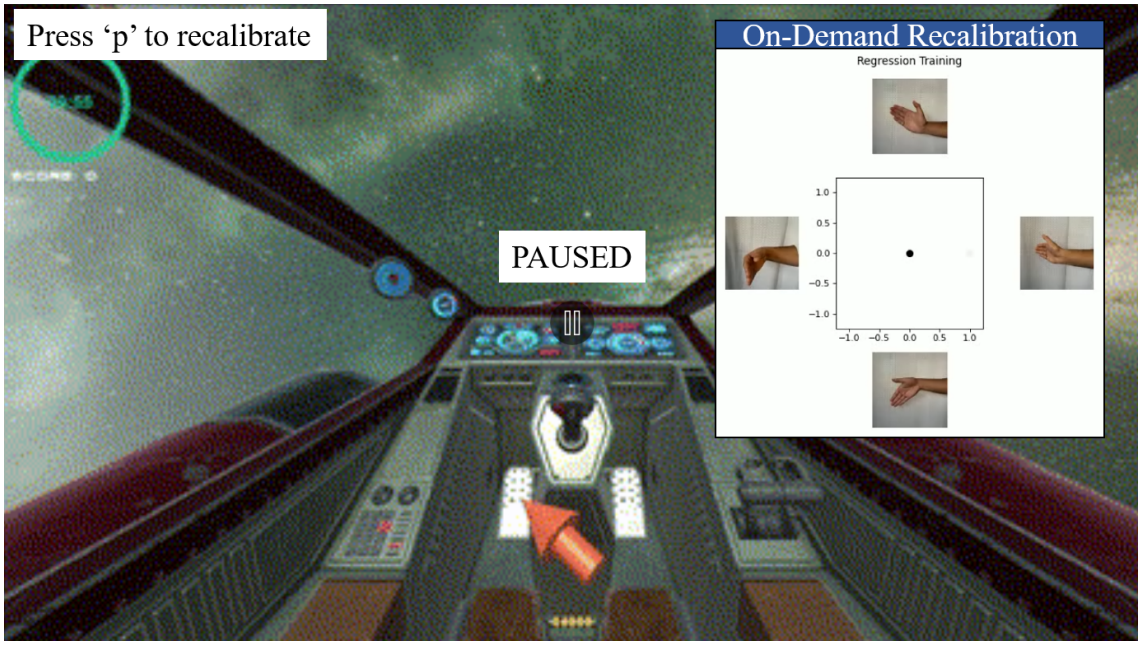


Figure 2.1: Illustration of a user-evoked screen-guided training session within a myogame environment typical of a dedicated on-demand recalibration incremental learning style.

assessed data quality after each recalibration with respect to data collected in the beginning.

2.3.2 Unsupervised Incremental Learning (UIL)

Despite promising results, systems that support on-demand recalibration face important drawbacks. Offline data collection disrupts normal user interaction and provides no immediate feedback, leaving users uncertain about whether their inputs are being interpreted correctly [37, 44]. Furthermore, supervised learning-based adaptation requires tight synchronization between the user and the training prompts, which can be difficult for users with inconsistent neural activation patterns [56] and frequently leads to spatial or temporal mislabelling of training data [74]. These approaches also tend to rely on specialized prompting equipment, complicating commercial integration.

To address these challenges, Unsupervised Incremental Learning (UIL) methods have been proposed. These techniques adapt the model online without explicit user prompting by identifying and updating structure within the EMG data stream. Typically, UIL approaches begin by discovering modes or basis vectors within an initial calibration dataset and then adjusting these representations during closed-loop use, as illustrated in Figure 2.2. A representative example is the work presented in [75], which introduced an adaptive decoder that incrementally updates its internal model based on unlabelled EMG observations. A related line of work advances bio-inspired regression-based decoding, which establishes a continuous mapping from EMG signals to intended movements [43]. These approaches draw on principles of neuromotor control, leveraging the observation that muscle activations during myoelectric control exhibit modular structure similar to natural motor tasks [76]. In particular, muscle synergies that reflect locally time-invariant patterns that capture coordinated activation across muscle groups [77], are often used as an underlying basis for model design. By exploiting these synergies, the model can represent user intent within a low-dimensional, physiologically meaningful space [44, 43]. This synergy-based formulation enables multifunctional control and supports motor learning by aligning model behaviour with the user's intrinsic motor modules [44]. However, most studies of these unsupervised or synergy-based learning strategies have been in emulated, non-adaptive settings with dedicated calibration sessions. Rigorous demonstrations of their efficacy during real-time, user-in-the-loop operation remain lacking.

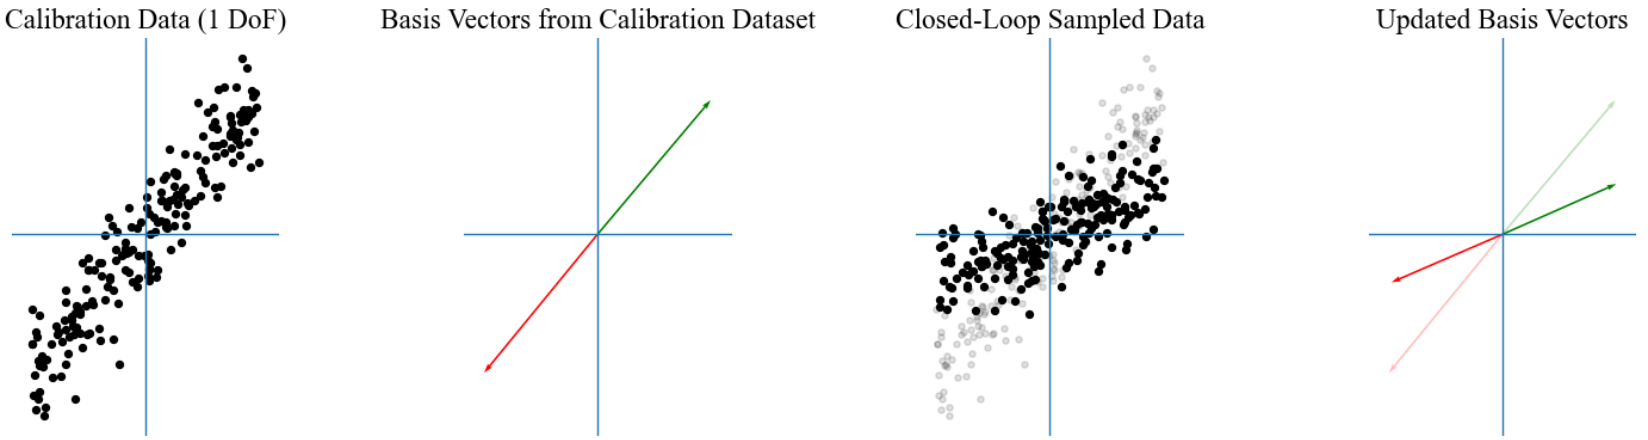


Figure 2.2: Illustration of basis vectors established from calibration and subsequent adjustments following collection of closed-loop signals.

Nevertheless, some demonstrations of UIL have been conducted in genuinely co-adaptive settings [47, 46, 45, 48], where models are adapted using data gathered directly during user-in-the-loop operation. For example, in [47], myoelectric control actions were progressively introduced, allowing the user to adapt gradually as task complexity increased. As discussed above, when reliable labels cannot be obtained during natural interaction, automated extraction of prominent muscle synergies and their assignment to DOFs can help to circumvent this labelling challenge. However, such arbitrary motor mappings may pose a challenge to the user because the resulting control may not align with natural or anthropomorphic movement patterns, making intuitive control more difficult.

2.3.3 Predictor-Dependent Incremental Learning (PDIL)

Although UIL can update basis vectors—thereby addressing distributional shifts over time—this approach cannot attribute incoming samples to specific classes or regression targets. In contrast, predictor-dependent incremental learning strategies leverage the model's own prediction outputs to guide further adaptation, directly supplying a categorical output or regression target as a pseudo-label. While initially relying on the model calibrated during the training phase to generate these pseudo-labels, predictor-dependent methods gradually become more self-consistent, aligning the adaptation process with the user's evolving internal model of the control scheme and alleviating conflict caused by fresh recalibration each session [78].

As pseudo-labels produced by a classifier can be incorrect, successful predictor-dependent incremental learning strategies have employed a variety of filtering mechanisms to enhance reliability. An illustration of the pseudo-labels assigned for various filters is given in Figure 2.3. The most common filter is the unsupervised high-confidence filter, which incorporates pseudo-labels into subsequent adaptation only

when the model exhibits sufficient certainty in its decision (e.g., $P(y_i|\bar{x}_t) > 0.99$) [37]. This approach ensures that samples located in well-described regions of feature space—where errors are less likely—are leveraged to gradually expand and refine the known space conservatively, ultimately resulting in improvements in both accuracy [37] and throughput [35] compared to initial calibration. However, classifiers can sometimes provide deceptively high confidences, particularly when a sample lies far from all known classes yet is marginally closer to one, prompting the exploration of augmentations such as the smallest margin uncertainty filter [55]. This method identifies samples near decision boundaries by searching for low differences between the highest and second-highest probability classes. In addition, the unsupervised low confidence filter [37] and the entropy reduction filter [55] have been proposed, correcting samples with the greatest ambiguity by suggesting a pseudo-label reflecting the prediction of adjacent time windows if that class ranked the second most likely. Finally, a blip filter supplies pseudo-labels reflecting nearby classes when an isolated prediction is surrounded by other predictions of a single class—operating under the assumption that such momentary blips are physiologically implausible [50]. In general, unless active learning is employed, where it is possible to retrieve ground truth labels upon a query [55], approaches that rely on low confidence predictions tend to be ineffective at improving the performance [37]. In contrast, high confidence-based approaches are generally successful, provided the system has not undergone a catastrophic shift [11, 37].

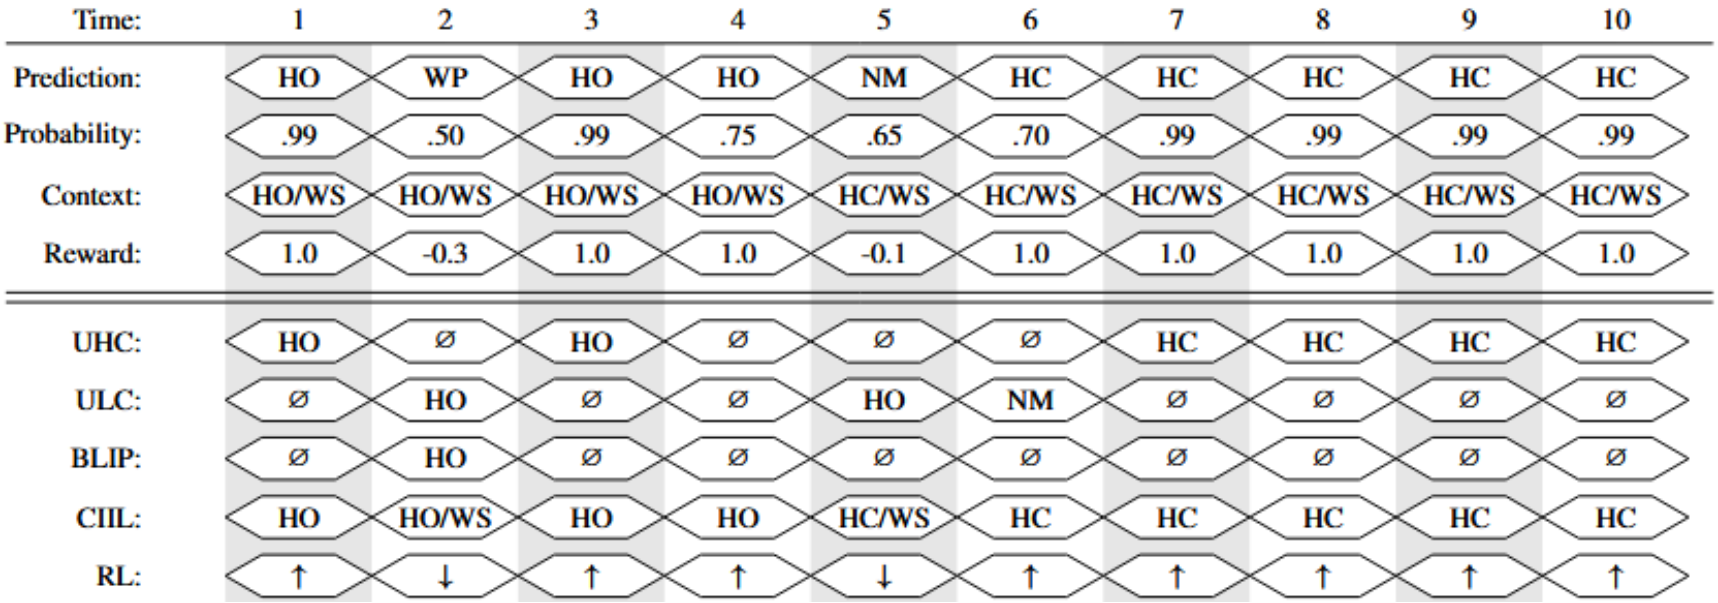


Figure 2.3: Consider a hypothetical ten-window classification horizon for an environment aware myoelectric prosthesis capable of performing hand open (HO), hand closed (HC), wrist pronation (WP), wrist supination (WS) and no motion (NM). In this scenario, the user must first open the hand to increase the aperture and then close the hand to secure contact with an object. The top portion of the diagram illustrates the sequence of classifier predictions, associated probabilities, environmental context, and the corresponding environmental reward signal during this task. The bottom portion of the diagram shows how several predictor-dependent incremental learning algorithms (UHC, ULC, BLIP) and environment-dependent algorithms (CIIL, RL) process each sample. For UHC, ULC, BLIP and CIIL, the bottom figure illustrates the pseudo-labels assigned to each window or, when applicable, the discarding of the sample (∅). For the RL strategy, the figure indicates whether the prediction for similar EMG patterns would be reinforced (↑) or dissuaded (↓) based on the environmental reward.

2.3.4 Environmental-Dependent Incremental Learning (EDIL)

As discussed above, predictor-dependent incremental learning strategies rely on the current model's decision stream, which can lead to performance degradation when an unaddressed covariate shift occurs (such as a previously unseen limb rotation) even if the shift is momentary. In contrast, environment-dependent incremental learning leverages an external oracle or environmental feedback to assess and validate actions with no dependence on how well the model performs. As such, environment-dependent approaches have been shown to result in strong online models without any initial calibration data provided [29, 65]. Currently, there are two primary strategies for environment-dependent incremental learning: reinforcement learning (RL) and context-informed incremental learning (CIIL).

RL is a paradigm where an agent learns a control policy by interacting with its environment while maximizing cumulative reward assigned based on model behaviour. Simply put, RL is an appealing approach when it is more practical to critique the general suitability of a model's behaviour than it is to explicitly define the desired behaviour using labels for each time step. In myoelectric control, RL methods derive mappings directly from EMG signals and a reward signal, without relying on choreographed supervised training protocols [63], potentially flawed pseudo-labels, or explicit biomechanics models [79]. When compared with supervised models, RL-based decoders provide significant performance improvements [69], enhance robustness to various factors (e.g., movement speed [80]), and offer more intuitive or more flexible training approaches. Examples include demonstrating desired behaviours through contra-lateral, intact limb movements (learning from demonstration [63, 64]), learning joint synergies based on reward or predicted future actuator changes [62, 68], providing learned vibrotacticle or auditory feedback signals [67, 81, 66], and even

delivering real-time human-provided reward signals relating to correct or invalid decoder outputs (e.g., directly training a multi-joint myoelectric regression controller using reward given to the system via an up/down key press [61]). Further, techniques from the field of RL have been shown to enable beneficial adaptive switching in prosthesis control by applying temporal-difference learning to both optimize and, in some cases, replace manual human switching between prosthetic modes and functions [65, 66]. Unfortunately, sample efficiency can be poor when the RL policy performs poorly or is initialized poorly. If the model experiences a catastrophic shift or starts with no calibration, RL approaches have been shown to require extensive exploration to converge or re-converge to effective policies. However, with a good initial calibration, dimensionality reduction, reward shaping, or training scaffolding, RL has been found to rapidly converge to an optimal (or user-functional) policy.

Instead of relying solely on exploration through continued interaction, CIIL incorporates richer environmental feedback that directly informs the selection of suitable classes or regression targets. By leveraging this additional context, CIIL is capable of jump-starting training even when the model is performing poorly, while behaving similarly to RL once the model begins to produce appropriate outputs (see Figure 2.3). The environmental feedback used by CIIL can come from various sources, such as those illustrated in Figure 2.4.

As will be discussed in subsequent chapters, CIIL has been shown to capitalize on environmental cues from interactive game environments to improve model performance even when only sparse calibration data are available [35]. It has also demonstrated robustness following large catastrophic electrode shifts [35] and enabled the simultaneous learning of proportional control mappings and classification boundaries from scratch [29]. Beyond classification, CIIL has been extended to regression

models, where it has yielded higher throughput, reduced drift, and decreased action interference than models trained for the same durations using screen-guided training [60].

However, environmental feedback need not be limited to game-derived signals. For example, eye gaze has been used to derive contextual cues during target achievement tests [71]. CIIL has also been applied in more realistic prosthesis use scenarios, such as an object transport task in virtual reality, where context derived from nearby objects decreased transport time per item by approximately 50% compared to a well-calibrated, non-adapting classifier [59]. Despite these advantages, a key challenge for EDIL in prosthesis applications, particularly for CIIL, which requires richer feedback, is the reliable deployment of perception models capable of providing accurate, low-latency context as feedback. Addressing this challenge will be essential for enabling CIIL methods to transition from controlled environments to practical prosthesis use.

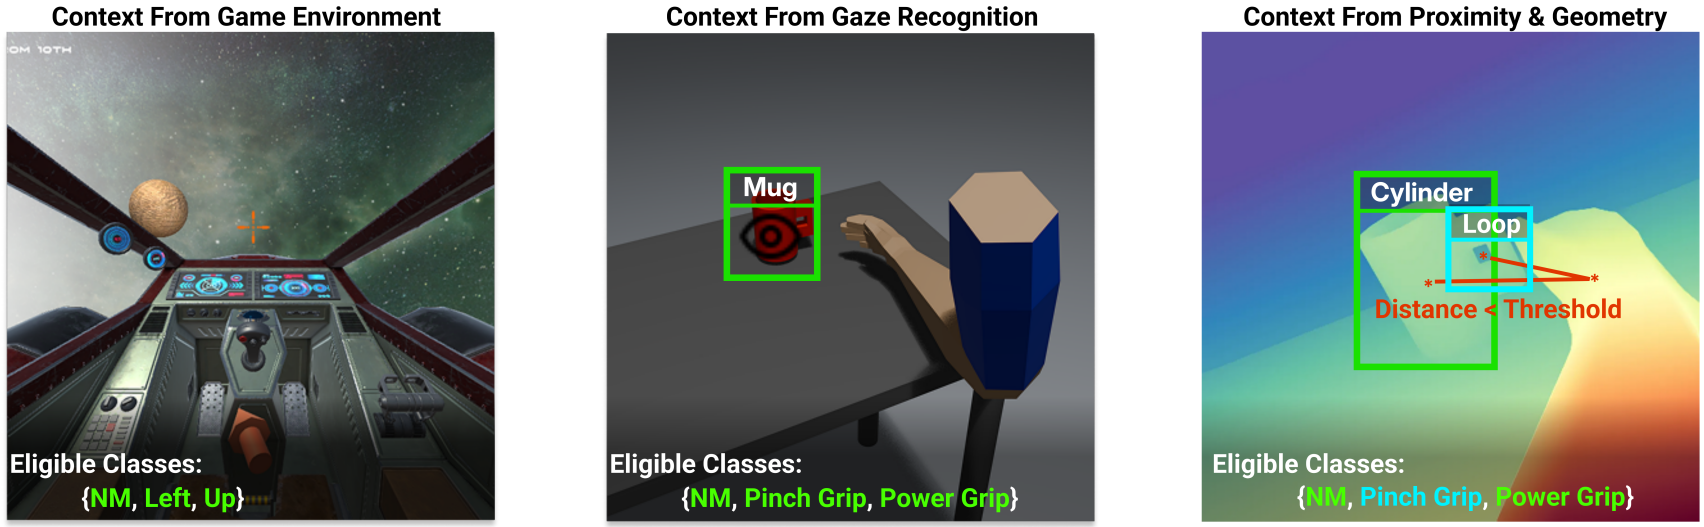


Figure 2.4: Illustration of sources of environmental feedback suitable for environment-dependent incremental learning (EDIL). The scenes depict how cues from game environments, gaze recognition, and geometry of proximal objects can provide rich contextual information to guide model adaptation. In CIIL, such context is used to determine the set of eligible classes and assign corresponding pseudo-labels, whereas RL uses the same information to generate a continuous reward signal reflecting the suitability of the performed action.

2.4 Taxonomy of Incremental Learning Goals

The capacity to incorporate new patterns of muscle activation depends on the intended adaptation goal, whether that is refining known functions, repurposing new inputs, or expanding the set of control actions. Each scenario requires different architectural and algorithmic capabilities. These needs are not strictly aligned with any single category of the myoelectric incremental learning taxonomy, but can be phrased more generally as three continual learning challenges: domain-incremental continual learning, task-incremental continual learning, and class-incremental continual learning. Figure 2.5 provides a depiction of these three continual learning problems.

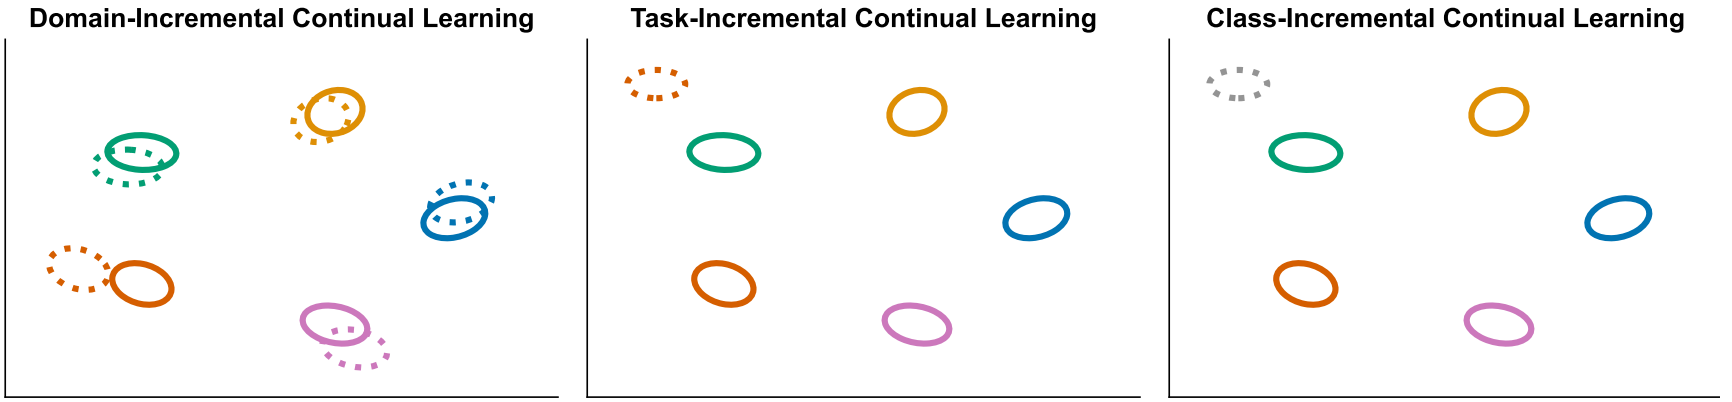


Figure 2.5: Illustration of an abstract feature space representing an initial model configuration of five well-separated motion classes shown as different coloured ellipses. Domain-incremental continual learning is represented by shifts in the location of each ellipse, introduced by changes in the model inputs. Task-incremental continual learning is represented by the appearance of a new red ellipse in a novel location when we already have a red class defined, suggesting a change in control strategy. Class-incremental continual learning is represented by a previously unseen grey class appearing in a new location, implying the addition of an entirely new action.

2.4.1 Domain-Incremental Continual Learning: Variations of Existing Patterns

Domain-incremental continual learning focuses on maintaining intent-recognition performance in the presence of natural variations in muscle signals for the existing motion classes and control outputs. These variations can arise from factors such as muscle fatigue, limb position changes, perspiration, or minor electrode shift. Each of these factors causes the input distribution to drift over time and, if left unaddressed, degrades control performance. A number of incremental learning strategies have already proven effective for mitigating this form of within-class drift. For example, Nishikawa et al. [82] used a DODR-based approach to periodically recalibrate the intent-recognition model and maintain performance over prolonged use of myoelectric control. Similarly, as described in the following chapters, EDIL approaches have been used to recover from a substantial electrode shift that had severely degraded model accuracy [35].

In general, DODR and EDIL strategies offer guarantees of recovery provided two conditions hold: 1) the feedback from the environment is reliable, and 2) the user continues to behave consistently with task expectations (e.g following a structured routine or pursuing an identifiable goal). Under these circumstances, the model can incrementally re-align itself with the shifted distribution after any domain shift.

When an incoming muscle pattern is merely a mild variation of a well-described gesture (i.e., it is still within previous decision boundaries), any incremental learning strategy should incorporate the incoming pattern beneficially. However, if the variation is large or sudden (e.g., a 45°electrode shift), strategies without reliable labels struggle, often mislabeling the new signals, and may worsen performance through subsequent model updates [35]. Using safeguards like the filtering strategies in PDIL

and regularization in UIL can help to provide more robustness to these input distribution shifts, but there is currently no consensus on how big a distributional shift can be handled by PDIL and UIL strategies.

2.4.2 Task-Incremental Continual Learning: New Patterns Triggering Existing Actions

Task-incremental continual learning considers the scenario where a user develops a new muscle activation pattern to achieve an already-supported function. For example, imagine the system was initially calibrated to move a cursor downward using wrist pronation, but over time, the user finds ulnar deviation more desirable to achieve the same action. The intended control output (move cursor down) hasn't changed, but they are now producing a muscle pattern that the system has never seen for that action. The challenge introduced here is not model failure but a natural evolution in user preference or motor strategy, which leads to two or more distinct inputs corresponding to the same output command.

Unimodal classifiers like LDA and linear SVM often struggle with this type of remapping. Their structure assumes a single, compact representation per class, making it difficult to accommodate multiple modes without degrading decision boundaries. In contrast, models like KNN or deep neural networks can support multi-modal mappings, making them more robust when users adapt their control strategies over time. Among incremental strategies, EDIL (CIIL) and DODR are well-suited to this case, as they can establish the novel mappings using explicit labels; whereas, UIL and PDIL would struggle to associate the new pattern with the correct output.

2.4.3 Class-Incremental Continual Learning: New Patterns Triggering New Actions

Class-incremental continual learning is required when a new muscle activation pattern associated with a previously unsupported function is introduced. The system must expand its output space, requiring not just parameter updates but structural changes to the intent recognition model. Among current approaches, DODR remains the most reliable strategy for this class-incremental scenario, as users can explicitly label new gestures to integrate them into the model. When an individual with a limb difference learns to use pattern recognition in a clinic, DODR is sometimes used, first introducing 1-3 motions until reliable control can be achieved, before expanding the model to support more motions [83]. UIL, likewise, has shown the ability to progressively add control outputs for novel input patterns upon request using piNNMF [47], which similarly adds a new control action when reliable control is achieved. EDIL may also support this case, provided that contextual signals or environmental cues are available to infer new class boundaries; however, no study has validated this. In contrast, PDIL is constrained by its reliance on existing model outputs and lacks the mechanisms needed to establish entirely new control actions. While new muscle activation patterns have been tangentially addressed in open-set myoelectric control (i.e., explicit rejection of activities of daily living [84]), harnessing emergent distinguishable motions from the user-in-the-loop setting for class-incremental continual learning has been a relatively neglected area.

2.5 User acceptability and user-centered design

Incremental learning could foster the transition of a prosthetic system from a tool to a lifelong companion by enabling the user to update the myoelectric control

system on demand. This can involve adding new data (both to improve the existing control and add new actions), deleting previous data (to correct for wrong updates), changing the emphasis on existing data, and more. At the same time, such a system could call for updates. For example, if the confidence in its prediction is deemed insufficient or if the sensors aboard the prosthesis report instability in the grasp, the model itself could prompt the user for intervention. While this could lead to improved performance, this process may be demanding for the user, leading to frustration and, potentially, rejection. Conversely, if these updates were to happen completely autonomously, these strategies could potentially reduce the user's feeling of agency over the system. It is therefore crucial that incremental learning solutions are designed to account for these factors.

Many other solutions have been proposed with an emphasis on overcoming the training burden that could potentially be associated with incremental learning. For example, serious games, which have had success in skill acquisition of myoelectric control [85, 86], could be used as arenas for calibration. This is particularly interesting because these environments can use gamification elements (e.g., difficulty level, target size, cursor speed, etc.) to capture different user behaviours that are often not accounted for during typical screen-guided training [87, 58]. Additionally, these games can improve motivation and adherence to training protocols [88]. Coapt [72], for example, has built-in games for user training that are being extended to arenas for incremental learning[42]. Alternatively, the training burden could be almost eliminated by leveraging autonomous incremental learning through additional sensors or other forms of context during regular everyday use. In both cases, however, the user should always maintain agency over the system by being able to revert to older models or update a model that is not behaving as desired.

Compelling evidence exists that myogames can help to improve the myocontrol al-

gorithm, but it remains unclear to what extent they can train the user and thereby foster both user and model co-adaptation [89, 90]. A growing body of literature has shown skill transfer from a virtual to a real environment, as long as care is taken in the game design [91, 92, 86]. Therefore, environments that are physically accurate and sufficiently realistic in robotic simulation and interaction could be beneficial to train and assess future co-adaptive myocontrol paradigms. In all cases, the *user interface* must be carefully designed to promote user-friendliness and motivation, leading to tighter and tighter co-adaptation and, ultimately, to optimal user control. Visualization of the user's own data in real-time is paramount, together with a way of quantifying the quality of the provided data. Moreover, the timing of data collection must be carefully controlled to avoid capturing noise or signals unrelated to the user's intent. The challenges can be addressed by drawing on using principles from *Interaction Design* [93] and constructivist psychology [94, 95], both of which emphasize aligning system behavior with the user's cognitive processes and patterns of real-time interaction. Recent systems have begun incorporate these principles. For example, the *Infinite biomedical technologies* system [96] enforces selective action-specific data capture and provides multiple real-time visualizations of user signals to support clearer, more intentional data collection.

2.6 Consideration of Open Challenges

Building on the taxonomy of incremental learning strategies for myoelectric control, the following discussion explores open challenges and limitations that span broadly across incremental learning strategies. Key topics include the ability to distinguish between signal changes caused by non-stationarities versus those reflecting genuine changes in user behaviour, the quality and quantity of data required for different strategies to function robustly, and user adaptation versus model adaptation.

2.6.1 Sensitivity to Noise vs. User-Behaviours

A key limitation across many incremental learning approaches for myoelectric control is the inability to explicitly distinguish between meaningful shifts in user behaviour and non-stationarities arising from noise or sensor drift. PDIL and UIL rely on their own predictions or latent representations to update the model, and because they lack any external guidance for class labels, they can easily mistake transient noise, drift, or other non-stationarities as genuine behavioural change and end up selecting for this variation as they adapt. While forgetting mechanisms [54] offer one means to mitigate the accumulation of noise-selective updates, few methods actively attempt to attribute signal changes to behaviour versus noise. Nonetheless, this challenge underscores the importance of carefully selecting adaptation mechanisms and incorporating safeguards—such as confidence thresholds, regularization, or controlled forgetting—especially in unsupervised or predictor-dependent frameworks. Other methods also offer stronger implicit safeguards. EDIL and DODR, for example, draw on external information (e.g., environmental pseudo-labels, reward signals, or explicitly supervised calibration data) to verify that each update is functionally beneficial without needing to attribute the change as either noise or behavioural. This additional validation layer offers more robust guarantees that learning remains directed to meaningful control improvements.

2.6.2 Model Complexity vs. Data Requirements

The quantity of data required for effective incremental learning is influenced more by the underlying intent recognition model and control complexity than by the incremental learning strategy itself. For instance, simple linear models can converge with minimal training data, while higher capacity architectures such as transformers typically require orders of magnitude more data to perform reliably, regardless of

which incremental learning strategy is deployed. Additional factors, such as the number of degrees of freedom being controlled and the variability in user behaviour, further impact data requirements. Nevertheless, learning paradigms that rely on externally validated updates, such as EDIL and DODR, tend to be more data-efficient due to their access to high-quality labels or reward signals. In contrast, PDIL approaches rely on their own predictions to generate pseudo-labels. While this allows them to function with limited data, it also makes them highly sensitive to sudden perturbations in recording conditions. To mitigate erroneous updates, these methods often use filtering techniques (e.g., high-confidence or margin-based filters) that deliberately reduce the number of training samples to prioritize their reliability. Similarly, UIL approaches do not require extensive data but depend heavily on consistent collection conditions. For instance, if a low-dimensional representation is learned using an incremental form of NNMF, the emergence of a new, more separable muscle potential during operation could shift the loading vectors, thereby disrupting the control mapping that the user has already internalized.

2.6.3 User vs. Model Adaptation

In co-adaptive prosthetic learning paradigms, the interaction between the user and the model forms a two-learner problem. Prior work has shown that users can compensate for model errors [97], and even adapt to arbitrary motor mappings [47]. An interesting future research direction could be to prevent premature convergence to a suboptimal state, e.g., when a user adapts quickly to a poorly performing model, leading to input patterns that further reinforce the model's unwanted behaviour. It remains unclear which co-adaptation protocols best guide both learners toward global convergence. Previous research has shown that the model's learning rate plays a critical role. Müller et al. [98] demonstrated that while overly slow learning

rates lead to stagnation and overly fast rates destabilize user learning, there exists a broad intermediate range that supports stable co-adaptive convergence. Subsequent research [99] has added that a matching adaptation between user and system supports the fastest and most stable convergence. When users interact with a prosthesis based on internal expectations of system behaviour, their ability to adapt depends on how well the (visual) feedback aligns with these expectations [100]. Following this line of thought, human behaviour is not static or deterministic in response to a given feedback. Rather, it can be actively managed by providing instructions and information that influence expectations and behaviour. For example, during the early training stages of a high-capacity model (with large data requirements and slow adaptation), the user may be asked to persist in performing the intended behaviour, even with unsatisfactory feedback, to avoid prematurely adapting to poor mappings. Similarly, the right balance between open-loop training (e.g., screen-guided training without feedback) and closed-loop co-adaptation remains an open question. Open-loop phases may allow for data collection without human adaptation and compensation, while testing sessions with static models could support dedicated human adaptation and expectation management. Similarly, actively instructing users or injecting small perturbations into the model could promote exploration and help escape local optima. The applicability and necessity of curated training protocols may vary by incremental learning learning category. In DODR, model and user adaptation are inherently staged sequentially, limiting the option of co-adaptation phases where both learners are active simultaneously. Whereas, in some of the UIL models (e.g. iNNMF), the user could be instructed to either reinforce muscle synergies discovered by the model or explore new activation patterns. Given that PDIL is self-referential, it can converge promptly with little risk of destabilizing. Furthermore, EDIL inherently reduces the risk of local optima by providing external context or task success as neutral guidance.

2.7 Establishing CIIL as a Viable and Beneficial Learning Strategy

This dissertation adopts environment-dependent incremental learning (EDIL) and introduces context-informed incremental learning (CIIL) as a viable adaptation strategy. As outlined in the following chapters, CIIL provides a principled mechanism for adapting models to systematic changes in user behaviour that arise in closed-loop control, rather than to transient variability or noise. By conditioning pseudo-labels on task context and behavioural suitability, CIIL preferentially incorporates pseudo-labels that reflect stable, goal-directed user strategies, while down-weighting samples that are inconsistent with the task objective. CIIL is also data efficient relative to the alternatives outlined in this chapter. By constraining adaptation to contextually valid samples and by converting weak environmental feedback into supervised targets, it reduces the volume of data required to achieve stable improvements, which is particularly important for higher-capacity models operating within online, human-in-the-loop constraints. Finally, CIIL alleviates concerns of unproductive or antagonistic co-adaptation. Because pseudo-labels are generated under the assumption of goal-seeking behaviour, model updates tend to reinforce mappings that are aligned with the user's emerging internal control model, rather than amplifying model errors or inducing exploratory dynamics that may interfere with user learning and trust.

The central challenge addressed in the remaining chapters of this thesis is the out-of-distribution mismatch between calibration and deployment, which is treated as a form of domain-incremental continual learning, where the set of intended actions remains fixed but the input distribution shifts due to behavioural and contextual changes in closed-loop use. The subsequent chapters investigate CIIL un-

der both low-difficulty and high-difficulty domain-incremental continual learning settings, corresponding to realistic/modest to elevated/severe mismatches between calibration and deployment conditions. CIIL is a model-agnostic pseudo-labelling strategy, so in this work, its performance is systematically evaluated across Bayesian classifiers, static neural networks, and temporal neural networks, as well as across both classification-based and regression-based control paradigms, to establish the generality of its benefits across model classes and control formulations.

Bibliography

[1] D. S. Childress, "Historical aspects of powered limb prostheses," *Clin Prosthet Orthot*, vol. 9, no. 1, pp. 2–13, 1985.

[2] T. W. Williams III, "Controlling powered upper extremity prostheses now and in the future," in *Proceedings of the Myoelectric Symposium 2005*, 2005.

[3] K. Englehart and B. Hudgins, "A robust, real-time control scheme for multifunction myoelectric control," *IEEE transactions on biomedical engineering*, vol. 50, no. 7, pp. 848–854, 2003.

[4] C. Castellini, P. Artemiadis, M. Wininger, A. Ajoudani, M. Alimusaj, A. Bicchi, B. Caputo, W. Craelius, S. Dosen, K. Englehart, D. Farina, A. Gijsberts, S. B. Godfrey, L. Hargrove, M. Ison, T. A. Kuiken, M. Markovic, P. M. Pilarski, R. Rupp, and E. Scheme, "Proceedings of the first workshop on Peripheral Machine Interfaces: going beyond traditional surface electromyography," *Frontiers in Neurorobotics*, vol. 8, no. 22, 2014.

[5] R. Merletti, A. Botter, C. Cescon, M. A. Minetto, and T. M. Vieira, "Advances in surface emg: Recent progress in clinical research applications," *Critical Reviewstrade; in Biomedical Engineering*, vol. 38, no. 4, pp. 347–379, 2010.

[6] P. Beckerle, C. Castellini, and B. Lenggenhager, "Robotic interfaces for cognitive psychology and embodiment research: a research roadmap," *Wiley Interdisciplinary Reviews - Cognitive Science*, p. e1486, 2018.

[7] A. Shehata *et al.*, "Machine learning for the control of prosthetic arms: Using electromyographic signals for improved performance," *IEEE Signal Process. Mag.*, vol. 38, pp. 46–53, 2021.

[8] B. Lock, A. M. Simon, K. Stubblefield, and L. J. Hargrove, "Prosthesis-guided training for practical use of pattern recognition control of prostheses," in *Proceedings of the Myoelectric Controls Symposium (MEC)*, pp. 61–64, 2011.

[9] A. Phinyomark, P. Phukpattaranont, and C. Limsakul, "Feature reduction and selection for emg signal classification," *Expert systems with applications*, vol. 39, no. 8, pp. 7420–7431, 2012.

[10] J. Lévesque, F. Chamberland, E. Scheme, and B. Gosselin, "Advancing myoelectric prostheses: Efficacy of gold-plated 3d-printed thermoplastic dry electrodes," *IEEE Transactions on Biomedical Engineering*, 2024.

[11] E. Campbell, A. Phinyomark, and E. Scheme, "Current trends and confounding factors in myoelectric control: Limb position and contraction intensity," *Sensors*, vol. 20, no. 6, p. 1613, 2020.

[12] E. Scheme and K. Englehart, "Training strategies for mitigating the effect of proportional control on classification in pattern recognition–based myoelectric control," *JPO: Journal of Prosthetics and Orthotics*, vol. 25, no. 2, pp. 76–83, 2013.

[13] A. J. Young, L. J. Hargrove, and T. A. Kuiken, "The effects of electrode size and orientation on the sensitivity of myoelectric pattern recognition systems to electrode shift," *IEEE transactions on biomedical engineering*, vol. 58, no. 9, pp. 2537–2544, 2011.

[14] I. Kyranou, S. Vijayakumar, and M. S. Erden, "Causes of performance degradation in non-invasive electromyographic pattern recognition in upper limb prostheses," *Frontiers in neurorobotics*, vol. 12, p. 58, 2018.

[15] R. N. Khushaba, M. Takruri, J. V. Miro, and S. Kodagoda, "Towards limb position invariant myoelectric pattern recognition using time-dependent spectral features," *Neural networks*, vol. 55, pp. 42–58, 2014.

[16] F. Chamberland, É. Buteau, S. Tam, E. Campbell, A. Mortazavi, E. Scheme, P. Fortier, M. Boukadoum, A. Campeau-Lecours, and B. Gosselin, "Novel wearable hd-emg sensor with shift-robust gesture recognition using deep learning," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 17, no. 5, pp. 968–984, 2023.

[17] L. Resnik, H. Huang, A. Winslow, D. L. Crouch, F. Zhang, and N. Wolk, "Evaluation of emg pattern recognition for upper limb prosthesis control: a case study in comparison with direct myoelectric control," *Journal of neuroengineering and rehabilitation*, vol. 15, pp. 1–13, 2018.

[18] L. Resnik, M. Borgia, J. Cancio, J. Heckman, J. Highsmith, C. Levy, S. Phillips, and J. Webster, "Dexterity, activity performance, disability, quality of life, and independence in upper limb veteran prosthesis users: a normative study," *Disability and Rehabilitation*, pp. 1–12, 2020.

[19] T. A. Kuiken, L. A. Miller, K. Turner, and L. J. Hargrove, "A comparison of pattern recognition control and direct control of a multiple degree-of-freedom

transradial prosthesis," *IEEE journal of translational engineering in health and medicine*, vol. 4, pp. 1–8, 2016.

[20] M. M. White, W. Zhang, A. T. Winslow, M. Zahabi, F. Zhang, H. Huang, and D. B. Kaber, "Usability comparison of conventional direct control versus pattern recognition control of transradial prostheses," *IEEE Transactions on Human-Machine Systems*, vol. 47, no. 6, pp. 1146–1157, 2017.

[21] A. J. Young, L. H. Smith, E. J. Rouse, and L. J. Hargrove, "A comparison of the real-time controllability of pattern recognition to conventional myoelectric control for discrete and simultaneous movements," *Journal of neuroengineering and rehabilitation*, vol. 11, pp. 1–10, 2014.

[22] S. Salminger, H. Stino, L. H. Pichler, C. Gstoettner, A. Sturma, J. A. Mayer, M. Szivak, and O. C. Aszmann, "Current rates of prosthetic usage in upper-limb amputees–have innovations had an impact on device acceptance?," *Disability and Rehabilitation*, pp. 1–12, 2020.

[23] M. B. I. Reaz, M. S. Hussain, and F. Mohd-Yasin, "Techniques of emg signal analysis: detection, processing, classification and applications," *Biological procedures online*, vol. 8, pp. 11–35, 2006.

[24] J. Sanford, R. Patterson, and D. O. Popa, "Concurrent surface electromyography and force myography classification during times of prosthetic socket shift and user fatigue," *Journal of Rehabilitation and Assistive Technologies Engineering*, vol. 4, p. 2055668317708731, 2017.

[25] S. Day, "Important factors in surface emg measurement," *Bortec Biomedical Ltd publishers*, pp. 1–17, 2002.

[26] A. Marinelli, N. Boccardo, F. Tessari, D. Di Domenico, G. Caserta, M. Canepa, G. Gini, G. Barresi, M. Laffranchi, L. De Michieli, and M. Semprini, "Ac-

tive upper limb prostheses: a review on current state and upcoming breakthroughs," *Progress in Biomedical Engineering*, vol. 5, p. 012001, Jan. 2023.

[27] C. Morrell, E. Campbell, and E. Scheme, "Exploring user compliance in the training of regression-based myoelectric control," in *Myoelectric Controls Symposium*, 2024.

[28] A. T. Wang, C. D. Olsen, W. C. Hamrick, and J. A. George, "Correcting Temporal Inaccuracies in Labeled Training Data for Electromyographic Control Algorithms," *IEEE International Conference on Rehabilitation Robotics*, 2023. Publisher: IEEE Computer Society ISBN: 9798350342758.

[29] E. Campbell, E. Eddy, X. Isabel, S. Bateman, B. Gosselin, U. Côté-Allard, and E. Scheme, "Screen guided training does not capture goal-oriented behaviours: Learning myoelectric control mappings from scratch using context informed incremental learning," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2024.

[30] B. Lock, K. Englehart, and B. Hudgins, "Real-time myoelectric control in a virtual environment to relate usability vs. accuracy," in *MyoElectric Controls Symposium*, pp. 122–127, Citeseer, 2005.

[31] J. Nawfel, K. Englehart, and E. Scheme, "A multi-variate approach to predicting myoelectric control usability," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, p. 1312–1327, 2021.

[32] R. M. Hinson, J. Berman, W. Filer, D. Kamper, X. Hu, and H. Huang, "Offline evaluation matters: investigation of the influence of offline performance on real-time operation of electromyography-based neural-machine interfaces," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 31, pp. 680–689, 2022.

[33] C. Castellini, "Incremental learning of muscle synergies: from calibration to interaction," in *Human and robot hands: sensorimotor synergies to bridge the gap between neuroscience and robotics* (M. Bianchi and A. Moscatelli, eds.), Touch and haptic systems, pp. 171–193, Springer International Publishing, 2016. Section: 11.

[34] S. Yoneda and A. Furui, "Bayesian approach for adaptive emg pattern classification via semi-supervised sequential learning," in *2023 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 3310–3315, IEEE, 2023.

[35] E. Campbell, E. Eddy, S. Bateman, U. Côté-Allard, and E. Scheme, "Context-informed incremental learning improves both the performance and resilience of myoelectric control," *Journal of NeuroEngineering and Rehabilitation*, vol. 21, no. 1, p. 70, 2024.

[36] D. Di Domenico, N. Boccardo, A. Marinelli, M. Canepa, E. Gruppioni, M. Laffranchi, and R. Camoriano, "Long-term upper-limb prosthesis myocontrol via high-density semg and incremental learning," *IEEE Robotics and Automation Letters*, 2024.

[37] J. W. Sensinger, B. A. Lock, and T. A. Kuiken, "Adaptive pattern recognition of myoelectric signals: exploration of conceptual framework and practical algorithms," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 17, no. 3, pp. 270–278, 2009.

[38] A. Gijsberts, R. Bohra, D. Sierra González, A. Werner, M. Nowak, B. Caputo, M. A. Roa, and C. Castellini, "Stable myoelectric control of a hand prosthesis using non-linear incremental learning," *Frontiers in Neurorobotics*, vol. 8, 2014.

[39] M. Nowak, R. M. Bongers, C. K. van der Sluis, A. Albu-Schäffer, and C. Castellini, "Simultaneous assessment and training of an upper-limb amputee using incremental machine-learning-based myocontrol: a single-case experimental design," *Journal of NeuroEngineering and Rehabilitation*, vol. 20, p. 39, Apr. 2023.

[40] F. Egle, D. Di Domenico, A. Marinelli, N. Boccardo, M. Canepa, M. Laffranchi, L. De Michieli, and C. Castellini, "Preliminary Assessment of Two Simultaneous and Proportional Myocontrol Methods for 3-DoFs Prostheses Using Incremental Learning," in *2023 International Conference on Rehabilitation Robotics (ICORR)*, pp. 1–6, IEEE, 2023.

[41] D. Di Domenico, F. Egle, A. Marinelli, M. Canepa, E. Gruppioni, N. Boccardo, M. Laffranchi, and C. Castellini, "Towards Bioinspired Co-Adaptive Control of a 3-DoFs Upper-Limb Prosthesis via Incremental Learning," Mar. 2025.

[42] Z. Wright and B. Lock, "Adaptive EMG Pattern Recognition Reduces Frequency and Improves Quality of At-home Prosthesis Training For Upper Limb Myoelectric Prosthesis Wearers," *Myoelectric Controls Symposium*, Aug. 2022.

[43] N. Jiang, H. Rehbaum, I. Vujaklija, B. Graimann, and D. Farina, "Intuitive, online, simultaneous, and proportional myoelectric control over two degrees-of-freedom in upper limb amputees," *IEEE transactions on neural systems and rehabilitation engineering*, vol. 22, no. 3, pp. 501–510, 2013.

[44] C. Lin, B. Wang, N. Jiang, and D. Farina, "Robust extraction of basis functions for simultaneous and proportional myoelectric control via sparse non-negative matrix factorization," *Journal of neural engineering*, vol. 15, no. 2, p. 026017, 2018.

[45] D. Yeung, I. M. Guerra, I. Barner-Rasmussen, E. Siponen, D. Farina, and I. Vujaklija, "Co-adaptive control of bionic limbs via unsupervised adaptation of muscle synergies," *IEEE Transactions on Biomedical Engineering*, vol. 69, no. 8, pp. 2581–2592, 2022.

[46] A. Gigli, A. Gijsberts, and C. Castellini, "Unsupervised myocontrol of a virtual hand based on a coadaptive abstract motor mapping," in *2022 International Conference on Rehabilitation Robotics (ICORR)*, pp. 1–6, IEEE, 2022.

[47] A. Gigli, A. Gijsberts, M. Nowak, I. Vujaklija, and C. Castellini, "Progressive unsupervised control of myoelectric upper limbs," *Journal of Neural Engineering*, vol. 20, no. 6, p. 066016, 2023.

[48] A. Schulz, F. Egle, M. Osswald, A. Del Vecchio, and C. Castellini, "Towards Unsupervised Incremental and Proportional Myocontrol Based on Higher-Density Surface Electromyography," in *2025 International Conference On Rehabilitation Robotics (ICORR)*, pp. 1106–1111, May 2025.

[49] Y. Yamanoi, S. Togo, Y. Jiang, and H. Yokoi, "Learning data correction for myoelectric hand based on "survival of the fittest"," *Cyborg and Bionic Systems*, 2021.

[50] R. Kato, T. Fujita, H. Yokoi, and T. Arai, "Adaptable emg prosthetic hand using on-line learning method-investigation of mutual adaptation between human and adaptable machine," in *ROMAN 2006-The 15th IEEE International Symposium on Robot and Human Interactive Communication*, pp. 599–604, IEEE, 2006.

[51] X. Zhai, B. Jelfs, R. H. Chan, and C. Tin, "Self-recalibrating surface emg pattern recognition for neuroprosthesis control based on convolutional neural network," *Frontiers in neuroscience*, vol. 11, p. 379, 2017.

[52] U. Côté-Allard, G. Gagnon-Turcotte, A. Phinyomark, K. Glette, E. J. Scheme, F. Laviolette, and B. Gosselin, "Unsupervised domain adversarial self-calibration for electromyography-based gesture recognition," *IEEE Access*, vol. 8, pp. 177941–177955, 2020.

[53] X. Jiang, C. Ma, and K. Nazarpour, "Plug-and-play myoelectric control via a self-calibrating random forest common model," *Journal of Neural Engineering*, 2024.

[54] D. Yeung, D. Farina, and I. Vujaklija, "Directional forgetting for stable co-adaptation in myoelectric control," *Sensors*, vol. 19, no. 9, p. 2203, 2019.

[55] K. Szymaniak, A. Krasoulis, and K. Nazarpour, "Recalibration of myoelectric control with active learning," *Frontiers in Neurorobotics*, vol. 16, p. 1061201, 2022.

[56] J. M. Hahne, S. Dähne, H.-J. Hwang, K.-R. Müller, and L. C. Parra, "Concurrent adaptation of human and machine improves simultaneous and proportional myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 23, no. 4, pp. 618–627, 2015.

[57] C. Igual, A. Castillo, and J. Igual, "An interactive training model for myoelectric regression control based on human–machine cooperative performance," *Computers*, vol. 13, no. 1, p. 29, 2024.

[58] R. B. Woodward and L. J. Hargrove, "Adapting myoelectric control in real-time using a virtual environment," *Journal of neuroengineering and rehabilitation*, vol. 16, pp. 1–12, 2019.

[59] G. Gagné, A. Azad, T. Labbé, E. Campbell, X. Isabel, E. Scheme, U. Côté-Allard, and B. Gosselin, "Context informed incremental learning improves

myoelectric control performance in virtual reality object manipulation tasks," *arXiv preprint arXiv:2505.06064*, 2025.

[60] C. Morrell, E. Campbell, E. Eddy, and E. Scheme, "Context-informed incremental learning improves throughput and reduces drift in regression-based myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, pp. 1–1, 2025.

[61] P. M. Pilarski, M. R. Dawson, T. Degris, F. Fahimi, J. P. Carey, and R. S. Sutton, "Online human training of a myoelectric prosthesis controller via actor-critic reinforcement learning," in *2011 IEEE international conference on rehabilitation robotics*, pp. 1–7, IEEE, 2011.

[62] P. M. Pilarski, T. B. Dick, and R. S. Sutton, "Real-time prediction learning for the simultaneous actuation of multiple prosthetic joints," in *2013 IEEE 13th International Conference on Rehabilitation Robotics (ICORR)*, pp. 1–8, 2013.

[63] G. Vasan and P. M. Pilarski, "Learning from demonstration: Teaching a myoelectric prosthesis with an intact limb via reinforcement learning," in *2017 International Conference on Rehabilitation Robotics (ICORR)*, pp. 1457–1464, IEEE, 2017.

[64] G. Vasan and P. M. Pilarski, "Context-aware learning from demonstration: Using camera data to support the synergistic control of a multijoint prosthetic arm," in *7th IEEE RAS EMBS Int. Conf. Biomed. Rob. Biomechatronics (BioRob)*, pp. 199–206, 2018.

[65] A. L. Edwards, M. R. Dawson, J. S. Hebert, C. Sherstan, R. S. Sutton, K. M. Chan, and P. M. Pilarski, "Application of real-time machine learning to my-

oelectric prosthesis control: A case series in adaptive switching," *Prosthetics and orthotics international*, vol. 40, no. 5, pp. 573–581, 2016.

[66] A. L. Edwards, J. S. Hebert, and P. M. Pilarski, "Machine learning and unlearning to autonomously switch between the functions of a myoelectric arm," in *2016 6th IEEE International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, pp. 514–521, 2016.

[67] A. S. R. Parker, A. L. Edwards, and P. M. Pilarski, "Exploring the impact of machine-learned predictions on feedback from an artificial limb," in *2019 IEEE-RAS-EMBS International Conference on Rehabilitation Robotics (ICORR)*, (Toronto, Canada), pp. 1239–1246, IEEE, June 2019.

[68] C. Sherstan, J. Modayil, and P. M. Pilarski, "A collaborative approach to the simultaneous multi-joint control of a prosthetic arm," in *2015 IEEE International Conference on Rehabilitation Robotics (ICORR)*, pp. 13–18, 2015.

[69] K. Freitag, Y. Karayiannidis, J. Zbinden, and R. Laezza, "Fine-tuning myoelectric control through reinforcement learning in a game environment," *arXiv preprint arXiv:2411.13327*, 2024.

[70] M. M. Madduri, M. Yamagami, S. J. Li, S. Burckhardt, S. A. Burden, and A. L. Orsborn, "Predicting and shaping human-machine interactions in closed-loop, co-adaptive neural interfaces," *bioRxiv*, pp. 2024–05, 2024.

[71] A. H. Chou, M. Madduri, S. J. Li, J. Isa, A. Christensen, F. L. Hutchison, S. A. Burden, and A. L. Orsborn, "Using eye gaze to train an adaptive myoelectric interface," *bioRxiv*, pp. 2024–04, 2024.

[72] Coapt, LLC, "Handbook (v6.1) complete control system gen2." `https://coaptengineering.com/wp-content/uploads/Coapt-Gen2-Handbook.pdf`. Accessed: Feb 2025. See also `https://coaptengineering.com`.

[73] Otto Bock HealthCare Deutschland GmbH, "Myo plus instructions for use (user)." https://www.ottobock.com/de-de/product/13E520#technical-data. Accessed: Feb 2025. See also https://www.ottobock.com.

[74] T. N. Tully, C. J. Thomson, G. A. Clark, and J. A. George, "Validity and impact of methods for collecting training data for myoelectric prosthetic control algorithms," *IEEE transactions on neural systems and rehabilitation engineering*, vol. 32, pp. 1974–1983, 2024.

[75] N. Jiang, S. Dosen, K.-R. Muller, and D. Farina, "Myoelectric control of artificial limbs—is there a need to change focus?[in the spotlight]," *IEEE Signal Processing Magazine*, vol. 29, no. 5, pp. 152–150, 2012.

[76] M. Ison and P. Artemiadis, "The role of muscle synergies in myoelectric control: trends and challenges for simultaneous multifunction control," *Journal of neural engineering*, vol. 11, no. 5, p. 051001, 2014.

[77] A. d'Avella and E. Bizzi, "Shared and specific muscle synergies in natural motor behaviors," *Proceedings of the national academy of sciences*, vol. 102, no. 8, pp. 3076–3081, 2005.

[78] I. Al-Tahir, J. Sensinger, and E. Scheme, "A better framework for the assessment of performance and stability of co-adaptive myoelectric systems," in *2022 International Conference on Rehabilitation Robotics (ICORR)*, pp. 1–5, IEEE, 2022.

[79] W. Wu, K. R. Saul, and H. Huang, "Using reinforcement learning to estimate human joint moments from electromyography or joint kinematics: An alternative solution to musculoskeletal-based biomechanics," *Journal of Biomechanical Engineering*, vol. 143, no. 4, p. 044502, 2021.

[80] J. Berman, R. Hinson, and H. Huang, "Comparing reinforcement learning agents and supervised learning neural networks for emg-based decoding of continuous movements," in *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pp. 6297–6300, IEEE, 2021.

[81] A. S. R. Parker, *Learning to Partner: Exploring Real-Time Adaptive Feedback via Temporal-Difference Machine Learning for Improved Human-Prosthesis Collaboration*. Phd thesis, Faculty of Rehabilitation Medicine, University of Alberta, 2024.

[82] D. Nishikawa, W. Yu, M. Maruishi, I. Watanabe, H. Yokoi, Y. Mano, and Y. Kakazu, "On-line learning based electromyogram to forearm motion classifier with motor skill evaluation," *JSME International Journal Series C Mechanical Systems, Machine Elements and Manufacturing*, vol. 43, no. 4, pp. 906–915, 2000.

[83] M. A. Powell and N. V. Thakor, "A training strategy for learning pattern recognition control for myoelectric prostheses," *JPO: Journal of Prosthetics and Orthotics*, vol. 25, no. 1, pp. 30–41, 2013.

[84] E. Eddy, E. Campbell, S. Bateman, and E. Scheme, "Human-machine interaction using discrete myoelectric control: Contrastive learning reduces false activations during activities of daily living," in *2024 10th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)*, pp. 128–133, IEEE, 2024.

[85] A. Tabor, S. Bateman, E. Scheme, D. R. Flatla, and K. Gerling, "Designing game-based myoelectric prosthesis training," in *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 1352–1363, 2017.

[86] B. Maas, J. Tchimino, B. van Dijk, A. Murgia, C. K. van der Sluis, and R. M. Bongers, "Skill transfer of upper limb prosthesis control after training in a virtual reality environment," *Technology and Disability*, p. 10554181241311685, 2025.

[87] R. Meredith, E. Eddy, S. Bateman, and E. Scheme, "Comparing online wrist and forearm emg-based control using a rhythm game-inspired evaluation environment," *Journal of Neural Engineering*, vol. 21, no. 4, p. 046057, 2024.

[88] F. M. Alfieri, C. da Silva Dias, N. C. de Oliveira, and L. R. Battistella, "Gamification in musculoskeletal rehabilitation," *Current reviews in musculoskeletal medicine*, vol. 15, no. 6, pp. 629–636, 2022.

[89] M. B. Kristoffersen, A. W. Franzke, C. K. van der Sluis, A. Murgia, and R. M. Bongers, "Serious gaming to generate separated and consistent EMG patterns in pattern-recognition prosthesis control," *Biomedical Signal Processing and Control*, vol. 62, p. 102140, Sept. 2020.

[90] C. A. Garske, M. Dyson, S. Dupan, G. Morgan, and K. Nazarpour, "Serious Games Are Not Serious Enough for Myoelectric Prosthetics," *JMIR Serious Games*, vol. 9, p. e28079, Nov. 2021.

[91] A. Boschmann, D. Neuhaus, S. Vogt, C. Kaltschmidt, M. Platzner, and S. Dosen, "Immersive augmented reality system for the training of pattern classification control with a myoelectric prosthesis," *Journal of NeuroEngineering and Rehabilitation*, vol. 18, p. 25, Feb. 2021.

[92] L. van Dijk, C. K. van der Sluis, H. W. van Dijk, and R. M. Bongers, "Task-Oriented Gaming for Transfer to Prosthesis Use," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 24, pp. 1384–1394, Dec. 2016.

[93] B. Moggridge, *Designing interactions*. MIT Press, 2007.

[94] M. Nowak, C. Castellini, and C. Massironi, "Applying Radical Constructivism to machine learning: a pilot study in assistive robotics," *Constructivist Foundations*, vol. 13, no. 2, pp. 250–262, 2018.

[95] M. C. Bettoni and C. Castellini, "Interaction in assistive robotics: a radical constructivist design framework," *Frontiers in Neurorobotics*, vol. 15, p. 67, 2021.

[96] Infinite Biomedical Technologies, LLC, "ibt." `https://www.i-biomed.com/`. Accessed: Jun 2025.

[97] J. M. Hahne, M. Markovic, and D. Farina, "User adaptation in Myoelectric Man-Machine Interfaces," *Scientific Reports*, vol. 7, p. 4437, June 2017.

[98] J. S. Müller, C. Vidaurre, M. Schreuder, F. C. Meinecke, P. von Bünau, and K.-R. Müller, "A mathematical model for the two-learners problem," *Journal of Neural Engineering*, vol. 14, p. 036005, Mar. 2017.

[99] D. De Santis, "A Framework for Optimizing Co-adaptation in Body-Machine Interfaces," *Frontiers in Neurorobotics*, vol. 15, p. 662181, Apr. 2021.

[100] C. Igual, J. Igual, J. Hahne, and K. Nazarpour, "User-prosthesis coadaptation," in *Control of Prosthetic Hands: Challenges and Emerging Avenues* (K. Nazarpour, ed.), pp. 159–174, Institution of Engineering and Technology, Dec. 2020.

Chapter 3

Context-Informed Incremental Learning Improves Both the Performance and Resilience of Myoelectric Control

Evan Campbell, Ethan Eddy, Scott Bateman, Ulysse Côté-Allard, Erik Scheme

3.1 Foreword

This chapter presents the first formulation of context-informed incremental learning. It was demonstrated in a classification-based myoelectric control system and compared against an established incremental learning strategy, unsupervised high-confidence adaptation. These approaches were compared using target acquisition tests with 1) a minimal calibration data initialization (where unsupervised high confidence adaptation was hypothesized to be beneficial), and 2) an electrode shift

introduced post-calibration (where unsupervised high confidence adaptation was hypothesized to fail). My contribution to this chapter spans ideation, experimental design, data collection, data analysis, manuscript preparation, and revision corrections.

3.2 Reference

This work was peer-reviewed and published under Creative Commons License BY-NC-ND 4.0 as:

E. Campbell, E. Eddy, S. Bateman, U. Côté-Allard, and E. Scheme, *Context-Informed Incremental Learning Improves Both the Performance and Resilience of Myoelectric Control,* Journal of NeuroEngineering and Rehabilitation, vol. 21, no. 70, p. 1, 2024.

3.3 Abstract

Despite its rich history of success in controlling powered prostheses and emerging commercial interests in ubiquitous computing, myoelectric control continues to suffer from a lack of robustness. In particular, EMG-based systems often degrade over prolonged use, resulting in tedious recalibration sessions, user frustration, and device abandonment. Unsupervised adaptation is one proposed solution that updates a model's parameters over time based on its own predictions during real-time use to maintain robustness without requiring additional user input or dedicated recalibration. However, these strategies can actually accelerate performance deterioration when they begin to classify (and thus adapt) incorrectly, defeating their own purpose. To overcome these limitations, we propose a novel adaptive learning strategy, Context-Informed Incremental Learning (CIIL), that leverages in situ context to better inform the prediction of pseudo-labels. In this work, we evaluate these CIIL strategies in an online target acquisition task for two use cases: (1) when there is a lack of training data and (2) when a drastic and enduring alteration in the input space has occurred. A total of 32 participants were evaluated across the two experiments. The results show that the CIIL strategies significantly outperform the current state-of-the-art unsupervised high-confidence adaptation and outperform models trained with the conventional screen-guided training approach, even after a 45-degree electrode shift ($p < 0.05$). Consequently, CIIL has substantial implications for the future of myoelectric control, potentially reducing the training burden while bolstering model robustness, and leading to improved real-time control.

3.4 Introduction

Following decades of success in controlling powered prostheses [1], myoelectric control is becoming an increasingly sought-after hands-free input modality for human-computer interaction (HCI) [2]. Its convenience, subtlety, and intuitiveness make it a particularly attractive solution for emerging ubiquitous applications, such as in mixed-reality scenarios where camera-based approaches are infeasible. Leveraging the electromyogram (EMG) signals generated during muscular contractions and associating them with gestures enables hands-free, always-available input for various potential tasks. For example, myoelectric control has been used for human-robot interaction [3, 4], cursor control [5], sign language recognition [6], and more recently mixed reality [7]. Its primary use has been, and continues to be, prosthesis control, where it is commercially offered by several vendors (e.g., Coapt[1], Infinite Biomedical[2], and Ottobock[3]). However, despite its inherent promise, myoelectric control continues to suffer from several factors that hinder its real-world viability [8, 9].

A major impediment to the widespread viability of myoelectric control systems is their inherent susceptibility to degradation over time, influenced by factors such as electrode displacement, user fatigue, and shifts in muscular contractions [8]. Although systems can be recalibrated by acquiring guided training data, this process is tedious, especially for general-purpose applications where systems should be as close to "ready to use" as possible. Furthermore, the data collected during these controlled calibration sessions frequently fails to accurately reflect the patterns of contractions elicited during real-time device usage [10]. An ideal system would be capable of learning realistic user behaviours while maintaining the model's robustness by continuously adapting to the user in real-time.

[1]https://coaptengineering.com
[2]https://www.i-biomed.com/
[3]https://www.ottobock.com/

Unsupervised adaptation is one proposed solution that adapts to user-in-the-loop data without requiring direct user input, solving many of the issues associated with supervised recalibration [11]. Because it uses the classifier's own outputs as pseudo-labels for adaptation, other fields have referred to this type of adaptation as *semi-supervised* [12]. This reliance on the classifier that is being adapted, however, leads to inherent limitations. For example, these approaches excel when preserving models that already perform well, producing pseudo-labels that are correct and confident. However, when faced with challenges such as a significant change in the input space – leading to reduced model accuracy – these unsupervised adaptation approaches may fail. As such, a crucial research objective is to combine the ground truth labels of supervised methods with the convenience offered by unsupervised methods.

To achieve the combined benefits of both methods, we introduce a novel approach called Context-Informed Incremental Learning (CIIL). This approach draws motivation from the demonstrated benefits of reinforcement learning in myoelectric control [13], while overcoming the user training burden introduced by the sample inefficiency of reinforcement learning approaches. Inspired by reinforcement learning paradigms, but remaining an unsupervised learning method, CIIL operates effectively in contexts where the user's true intentions can be inferred. CIIL monitors the consequences of actions within an instrumented environment to assess their appropriateness, categorizing actions as positive or negative. Before adapting the model, the categories of actions provide an opportunity to reassign more appropriate labels if the classifier-assigned actions do not align with the inferred goal of the environment. Further, the context-informed labels can be used to isolate samples the model assigned an incorrect label (negative) or correct label (positive). This method overcomes the limitations of unsupervised adaptation strategies that often rely on the classifier's confidence predictions being correct, as it focuses on the action's suitability based on environmental context.

We evaluate the use of several CIIL strategies in two separate scenarios that could benefit from adaptation: (1) when limited training data is available and (2) when a drastic shift in the input space has occurred (i.e., after a 45-degree electrode shift). By extracting environmental context in a simple target acquisition game, we show how this additional context can not only maintain model performance over time but improve the underlying model behaviour compared to dedicated screen-guided training approaches. In a departure from conventional methods, our innovative CIIL training approach not only outperforms current standards but also holds substantial promise for advancing the effectiveness of next-generation myoelectric control systems.

3.5 Related Works

3.5.1 Incremental Learning in Myoelectric Control

The practice of collecting offline data (i.e., where contractions are elicited with no active control or feedback provided to the user) to train a classifier is arguably an incomplete solution when optimizing for online, user-in-the-loop interactions. Correspondingly, there is often weak association between offline classification performance and usability during later online use [10]. While other offline metrics may provide better predictive power [14], the best indicator of online usability is evaluating classifiers in more realistic intermediate tasks, such as a simple virtual environment [15]. This is, at least in part, because the performance of myoelectric control is dependent not only on the model's fit but how well understood the user's behaviour is under all circumstances (i.e., when a user's mental bandwidth is split between adhering to consistent patterns as produced during training and performing their control task). As such, conventional guided calibration approaches, wherein a user

mimics a sequence of prompted gestures used as labels for supervised learning (such as screen-guided, or prosthesis-guided, training [16]), prioritize the ability of the classifier to distinguish motion classes within the capacity of the controlled training dataset; but not necessarily for a more variable user-in-the-loop setting. Given the importance of capturing these user-in-the-loop behaviours for model performance, the online adaptation of model parameters during real-time user interaction is an appealing alternative to guided calibration techniques.

Traditionally, online adaptation has been explored as a mechanism to combat a decline in performance due to gradual physical and behavioural changes [17, 18, 19]. In the work done by Zhang *et al.* , it was shown that unsupervised adaptation could improve model resilience to simulated random noise that was added to offline testing data [17]. More recently, it has been demonstrated that serious games provide an avenue for online adaptation to improve classifiers trained with screen-guided prompts [18, 20]. For example, Woodward *et al.* found that users elicited greater limb position variability during online use which, when accounted for through online adaptation, yielded significantly higher scores during functional tests for both intact-limbed and amputee populations. Furthermore, researchers have addressed performance decline in myoelectric control across days by developing recalibration strategies, such as the self-calibrating asynchronous domain adversarial neural network (SCADANN) introduced by Côté-Allard *et al.* This approach effectively combines established myoelectric domain adversarial methods [21, 22] with stability-focused heuristics and, notably, demonstrated superior robustness through gradual refinement of its representation using minimal unlabeled data [19].

Nevertheless, most unsupervised adaptation research has assumed the availability of a reasonably robust initial model. Comparatively little work has explored its performance in light of a catastrophic concept shift (such as doffing/donning of the

EMG recording device), when the previously trained classifier may yield ineffective pseudo-labels, or with little-to-no initial training. In such situations, the classifier's predictions may result in incremental learning updates, reinforcing the same flawed representation or, worse, driving the model in unknown and irreproducible directions. Although auxiliary information, such as informativeness [23], has been used alongside classifier confidence to determine valid samples to reincorporate, this may prevent further degradation of a flawed model but does not enable recovery or improvement. Correspondingly, there is a need for adaptive techniques that are less susceptible to such factors.

3.5.2 Reinforcement Learning

Unlike traditional supervised strategies that utilize pairs of inputs and labels, reinforcement learning uses a reward signal derived from the advantageous or disadvantageous outcomes of an agent's interactions. This reward signal serves as a *direct* measure of the agent's functional ability and can guide model updates independent of the classifier's accuracy in predicting the user's intent. Reinforcement learning has been explored in the myoelectric literature, although primarily for offline supervised strategies. Wu *et al.* demonstrated the effective learning of wrist and finger joint angles using proximal policy optimization, circumventing the drift typically encountered with supervised regression and forward dynamical simulations [24]. They defined the discrepancy between predicted and measured joint trajectories as a continuous-valued reward signal. Similarly, Vasconez *et al.* utilized deep Q-learning for gesture recognition, employing positive rewards for correct predictions and negative rewards for incorrect predictions based on data collected from screen-guided training [25]. Likewise, Edwards *et al.* used generalized value functions with EMG signals to elicit a switch command that changed the degree of

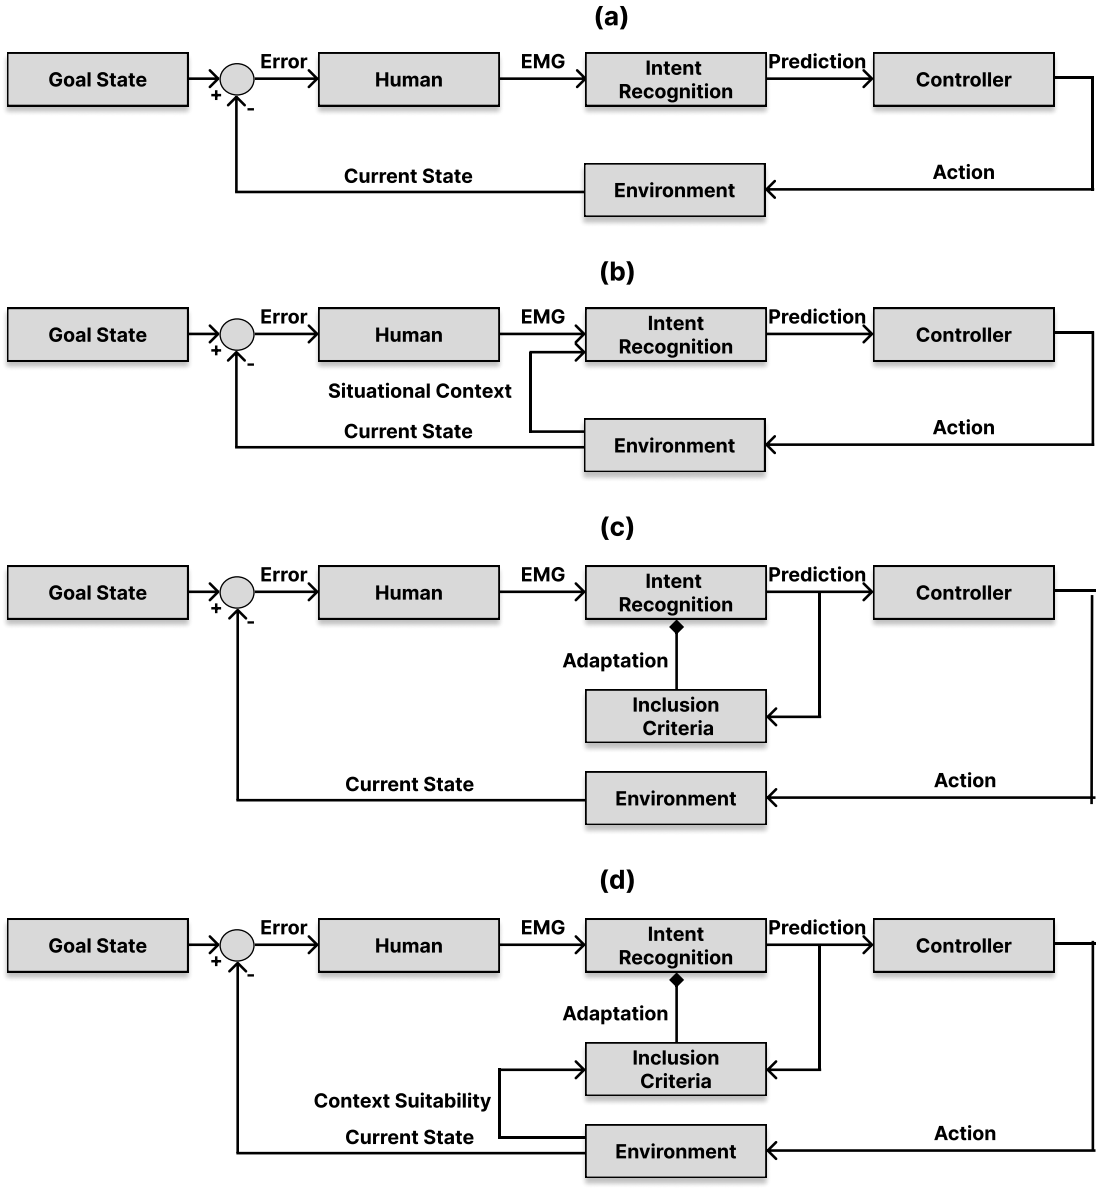


Figure 3.1: Depictions of various myoelectric control system designs. Solid lines indicate an output from a node or an input to a node, and a diamond indicates a signal that modifies the weights of a node. (a) Myoelectric control systems designed using traditional practices (optimized for open-loop control, user perception closes the loop). (b) Myoelectric control systems that are context-aware (CA) and modulate the intent relayed to the controller subject to environmental information. (c) Myoelectric control systems optimized via traditional incremental learning (IL) (these systems rely on the high confidence predictions of the classifier to incorporate incoming data to the adaptation set). (d) Myoelectric control systems that are supplied contextual feedback based on the suitability of the performed action in the environment and use this information to govern incremental learning (CIIL).

freedom being controlled, a mechanism typically used alongside direct sequentially controlled prostheses [26]. Although these approaches showcase the application of reinforcement learning in myoelectric control, the reward signals utilized were not direct measurements of performance in functional tasks.

There are relatively few studies that explore reinforcement learning for myoelectric control with a reward signal that directly quantifies functional performance from a user-in-the-loop setting. Pilarski *et al.* has demonstrated that actor critic reinforcement learning can optimize the control of a robotic arm from reward signals that quantify functional performance in two ways [13]. The first method uses a binary success metric, issuing a reward of +1 when within a region around the target trajectory and a reward of -0.5 when outside this region. The second method was making the user directly assign the rewards using arrow keys during use, as was pioneered by the TAMER approach [27]. Using both approaches, the user could precisely control two joints of the robotic limb to the desired target after sufficient training (<25 minutes and <10 minutes for the binary and human-assigned reward signals, respectively). Despite the success of this approach, reinforcement learning approaches are notoriously sample inefficient [28]. This relatively rapid training time was achieved by first developing an efficient EMG input representation and further tile coding the inputs of the continuous values into discrete bins. This compressed representation may not extend to more rich myoelectric control tasks when greater spatial and temporal resolution are needed to decipher user intent. As such, there is motivation to learn how to harness such contextual information, but to leverage it in a more sample efficient way.

3.5.3 Context

As characterized in human-computer interaction research, context is "any information that can be used to characterize a situation or entity" [29, 30]. By leveraging this often simple additional information, naive computer systems can become contextually aware, significantly improving the potential bandwidth of the particular human-computer input [31, 32]. For example, systems that leverage past input history as context can provide improved and tailored recommendations to users [33]. Alternatively, a user's location could be used as context to improve search results on the web [34]. In turn, these contextually aware systems have become an expectation for many users, reducing the need for user input to enhance the user experience.

Despite its history of success within HCI, the concept of context has only recently begun to surface in prosthetics. For example, some have used decision stream context like majority voting [35] and rejection [36] to improve control. Similarly, Sensinger *et al.* proposed a form of decision stream context during unsupervised adaptation by not adapting to changes in contractions that were not physically possible (i.e., they happened too quickly) [37]. While both approaches utilize the temporal context of the decision stream, they cannot be considered fully contextually aware systems. This is because their reliance on a potentially incorrect underlying model may result in a context that is isolated from, and possibly in disagreement with, the real-world environment. Others have sought to improve situational context by including additional sensors, such as cameras embedded within the prosthesis [38, 39, 40]. These are examples of context-aware myoelectric control systems that incorporate environmental information into their classifier inputs (see Figure 3.1b). Although such systems can greatly improve control by providing additional information to the underlying classifier [41, 39], they do not attempt to correct the model's behaviour. Moreover, these techniques may require the introduction of additional sensors, hard-

ware, or environmental constraints and may reduce the user's agency over the system – thus leading to abandonment of the device [42].

In this work, we lean on context to better derive pseudo-labels during unsupervised adaptation. In this novel paradigm, the context from the environment (i.e., the task's setting) acts as a supervisory layer to inform the active learning process and correct the underlying model rather than the output. This becomes particularly interesting for general-purpose myoelectric settings where context becomes more trivial to extract. For example, environmental context could be a successful or missed button click. Alternatively, context could be quick error corrections such as quickly opening and closing a menu in mixed reality. By leveraging this seemingly simple context to continuously refine and enhance a model's performance, the viability of myoelectric control incrementally improves, encompassing not only prosthetic control but also extending to mixed reality environments where the context might be even more immediately perceptible.

3.6 Methods

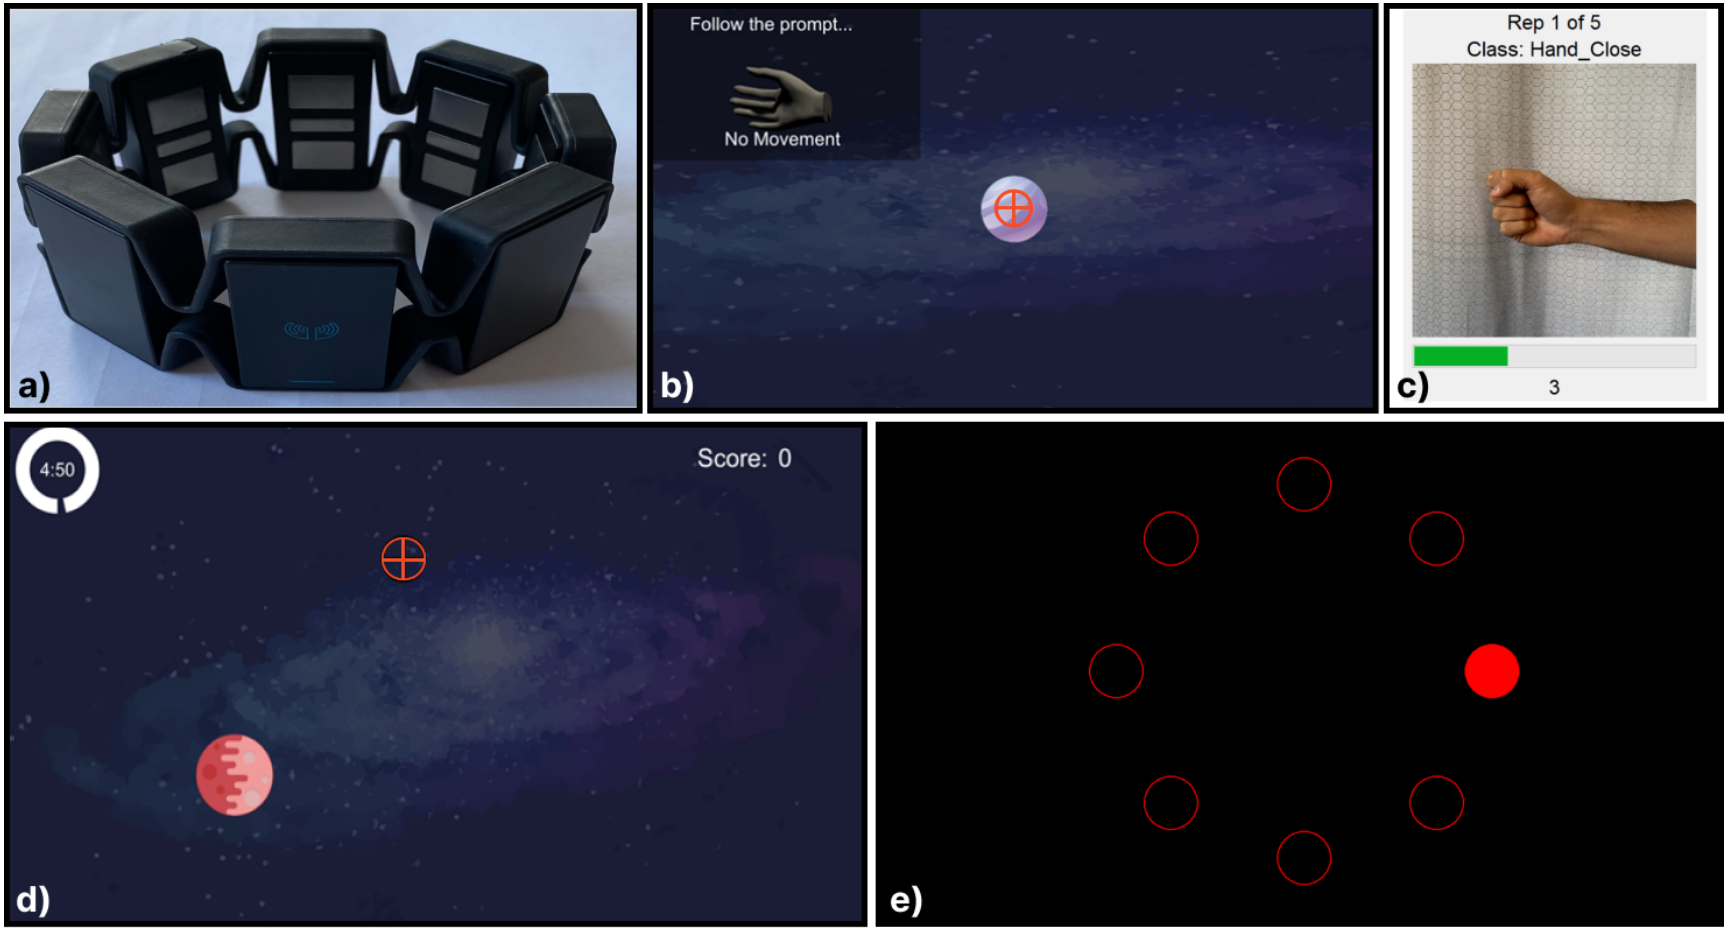


Figure 3.2: Experimental setup used throughout the study. (a) The Thalmic Labs Myo. (b) The calibration screen used in Scenario 1 collects one second of EMG data per motion class. (c) The screen-guided training interface for collecting offline data. (d) The gamified proxy environment, Myoshoot, where adaptation occurred. The user's position was represented by a crosshair and the goal was represented by a planet. (e) The Fitts' law environment where online usability was ultimately tested.

3.6.1 A Proxy Environment for Online Adaptation

This work employs a simple target acquisition game to provide a context-aware environment from which real-time adaptation can be performed (see Figure 3.1 b). In this game, participants controlled a crosshair along two independent degrees of freedom (Left/Right and Up/Down) with the goal of "saving" as many planets as possible in the given amount of time. After successfully moving the cursor over the top of a planet and hovering within its bounds for three seconds, the planet spawned at a different location. This cycle continued for the game's duration. To keep consistency across participants and trials, the sequence of positions in which the planets were generated remained the same for all participants and in each trial. Although contrived, this environment serves as a proxy for any myoelectric control setting where context can be inferred.

3.6.2 Control Scheme

The Myo Armband, a previously commercially available dry-electrode EMG cuff that records data at 200 Hz across eight channels, was used in this work. To discriminate between inputs, data were split into windows of 200 ms with 100 ms increments corresponding to 40 samples per window with an increment of 20 samples between windows. Hudgins' time domain features were extracted from each window [43]. A linear discriminant analysis (LDA) classifier was then used to differentiate between the five gestural commands (Wrist Flexion - Left, Wrist Extension - Right, Hand Open - Up, Hand Close - Down, Neutral Gesture - No Motion). After training, the LDA model was adapted in real-time during the aforementioned target-acquisition game using ten-second batches of data. This interval was chosen empirically through pilot testing and was a trade-off between user adaptation to the changing control scheme and optimal adaptation performance [44]. The mean and covariance of each

class were computed and combined with the existing LDA parameters using a static adaptation rate of 0.1 and a scaling factor reflecting the effective sampling rate. The adaptation of the mean and covariance matrices is shown below in equations (1) and (2):

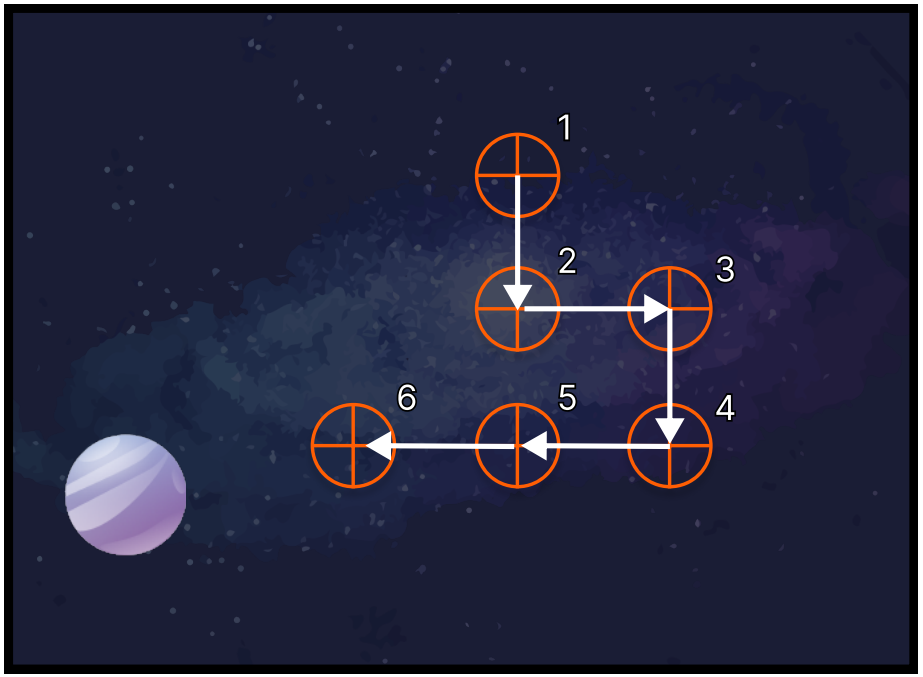
$$a = \frac{\alpha N_{b,c}}{N_c + \alpha N_{b,c}} \tag{3.1}$$

$$\mu_{c,i} = (1 - a)\mu_{c,i-1} + a\hat{\mu}_c \tag{3.2}$$

where $\mu_{c,i}$ is the mean of class c after adaptation, $\mu_{c,i-1}$ is the current mean of class c, $\hat{\mu}_c$ is the mean of class c of the batch, a is the adaptation rate, α is the static adaptation rate, $N_{b,c}$ is the number of new samples of class c, and N_c is the total number of samples in class c collected for adaptation thus far.

* Note that this work was built with LibEMG, an open-source Python library for designing and evaluating real-time myoelectric control systems [45]. The code and data can be found at *github.com/ECEEvanCampbell/CIIL_LDA*.

3.6.3 Incremental Learning Strategies



	Decision	Suitability	UHC	N	P	P+N
1	**Class:** Close **Prob:** 90%	P Close Flexion	✗	✗	✓	✓
2	**Class:** Extension **Prob:** 100%	N Close Flexion	✓	– Close	✗	– Close
3	**Class:** Close **Prob:** 95%	P Close Flexion	✗	✗	✓	✓
4	**Class:** Flexion **Prob:** 90%	P Close Flexion	✗	✗	✓	✓
5	**Class:** Flexion **Prob:** 60%	P Close Flexion	✗	✗	✓	✓
6	**Class:** Flexion **Prob:** 100%	P Close Flexion	✓	✗	✓	✓

Figure 3.3: A representation of a hypothetical sequence of cursor movements in the control task. The crosshair indicated the cursor position at each timestep with arrows indicating the direction of travel for the next timestep. An interpretation of the steps along the trajectory is given in the right table, where the decision (class and probability) and context suitability (was the performed action correct (P or N for positive or negative) and what were the correct options under the context) served as inputs to the inclusion criteria of different algorithms. The effects of the different algorithms are given by checkmarks (include with classifier's prediction), hyphens (override the decision using context), and x's (discard the sample).

To improve the usability of myoelectric control systems beyond the current status quo, we introduce the concept of **Context-Informed Incremental Learning (CIIL)**. In the case of myoelectric control, context can be acquired through task-specific constraints, knowledge, or additional sensors. The system then gets *feedback* based on its perceived contextually-defined performance, and using this feedback, the model can *learn* from its actions and make improvements (see Figure 3.1d). The model is then *incrementally* updated based on the contextual knowledge acquired, leading to more robust control systems that are defined by active in situ user behaviour. This work highlights three categories of CIIL systems (Negative, Positive, and Positive + Negative) and compares their performance to the state of the art adaptive system (unsupervised high-confidence). An overview of the these strategies is presented in Figure 3.3.

3.6.3.1 Negative Context-Informed Incremental Learning (N-CIIL)

N-CIIL attempts to correct incorrect (i.e., negative) model behaviour by updating the predicted label with the assumed "correct" one. In the target-acquisition environment, incorrect actions were flagged when the cursor moved away from the target, indicating a disagreement between the task objectives and user behaviour. As the cursor could only be controlled along one axis at a time, two directions could usually be deemed correct (i.e., the diagonal between the cursor and the target). The training cluster whose mean was closest in Euclidean distance to the newly elicited contraction was selected to differentiate between the two potentially correct labels. Note that all positive actions were ignored and not used to adapt the system.

3.6.3.2 Positive Context-Informed Incremental Learning (P-CIIL)

P-CIIL reinforces positive model behaviour by adapting the model with predictions that lead to positive within-task results. In the target-acquisition environment, positive actions were flagged when the cursor moved toward the target, indicating that the user was accomplishing the task objectives. Correspondingly, the assumption was that the model made a correct decision that should be reinforced. All negative actions were ignored and not used to adapt the system.

3.6.3.3 Positive + Negative Context-Informed Incremental Learning (P+N-CIIL)

P+N-CIIL uses a combination of positive and negative CIIL, where all of the model's predictions are leveraged for adaptation. Positive outcomes are adapted using P-CIIL, and negative outcomes are adapted using N-CIIL.

3.6.3.4 Unsupervised High-Confidence (UHC)

UHC is a popular extension of unsupervised adaptation where a model is updated with all decisions that have a confidence value above some predefined threshold [37]. Although this strategy uses the additional context of the classifier's confidence profile, it is not considered a CIIL strategy as it does not leverage context beyond the knowledge of the underlying model.

3.6.4 Outcome Measures

3.6.4.1 Offline Metrics

Three offline performance metrics were leveraged to assess the performance of each model: classification accuracy, active error, and instability [46]. Accuracy is a common metric that captures the number of correct predictions as a percentage of the total predictions. In contrast, active error is a specialized metric that focuses only on misclassifications that resulted in unwanted movement, and does not penalize misclassifications that were predicted as no motion. Instability is a metric that captures the number of class transitions seen in the decision stream. If analyzing data obtained through screen-guided-training, which typically does not record class transitions, the observed instability should ideally be zero. These metrics were computed on the screen guided training data acquired during each experiment.

3.6.4.2 Online Metrics

To evaluate the online performance of each adapted model, an ISO 9241-9 inspired Fitts' Law test evaluation (see Figure 3.2c) was used. Fitts' Law is broadly used for human-computer interaction studies and provides an avenue for testing pointing and selection tasks [47]. During this online usability test, participants were required to capture a set of eight targets arranged in a circle. Although similar to the target acquisition game, no model adaptation occurred during this task. Additionally, participants were told they could give up at any time if they felt they did not have the control required to complete the test. Three commonly used online metrics were assessed during this task: throughput, path efficiency, and overshoots [48]. Path efficiency is the ratio of the shortest possible path from the cursor to the target to the actual path taken. The overshoot metric – the number of times the cursor enters

and leaves a target before acquisition – is particularly important for measuring the stability of the control scheme during target acquisition. Finally, the throughput of the system is a measurement of the overall performance and is computed by the following equations:

$$TP = \frac{1}{T}\sum_{i=1}^{T}\frac{ID_i}{MT_i} \tag{3.3}$$

$$ID = \log_2(\frac{D}{W} + 1) \tag{3.4}$$

where T is the number of trials, ID is the index of difficulty, MT is the movement time, D is the distance to the target, W is the width of the target, and i is the trial number.

3.7 Experimental Design

In this work, we evaluated the CIIL strategies during two situations that could benefit from adaptation: (1) when there is minimal training data (e.g., when adapting to a new user) and (2) when the input space has catastrophically shifted (e.g., after a 45 degree electrode shift). A total of 32 participants took part in these online, user-in-the-loop experiments, as described in the following sections. At the start of the session, participants were asked to place the EMG cuff in a comfortable position on their right forearm. After giving the armband time to adjust to the participants' skin and humidity (approximately 5 minutes), the initial model (i.e., the baseline for all the adaptation trials) was trained. The participants then went through a progression of in-game adaptation on the baseline classifier, proceeded by an ISO Fitts' evaluation of the adapted model. Additionally, two repetitions

of steady-state, open-loop (no feedback) contractions were collected through screen guided training after each cycle to acquire a representative testing sample. This proceeded until all the adaptation strategies had been tested. All participants gave written informed consent prior to beginning, as approved by the University of New Brunswick's Research Ethics Board and are on file as REB (2022-122).

3.7.1 Experiment Scenario 1: Minimal Training Data

Experiment one consisted of a pilot study to evaluate the efficacy of the CIIL strategies on a sparse yet reasonably accurate starting classifier. The preliminary results of this study were described in [49]. Eleven participants were recruited for this initial work (4 female, 7 male; ages 20-56). To train the starting model, one second of training data for each of the five classes using the built-in data collection screen (see Figure 3.2) were acquired. The goal was to train a model whose feature space was correct but sparsely populated, possibly leading to unstable classifier behaviour. Although this limited amount of data is known to be suboptimal, commercial EMG systems have promoted the use of single repetition calibration protocols (e.g. Myo) to prioritize an expedited setup at the expense of performance. We also acquired five repetitions of steady-state contractions through screen guided training to evaluate the performance of a classifier trained on offline data (i.e., the optimal baseline). Correspondingly, the goal of the first experiment was an initial analysis of the newly proposed CIIL adaptation strategies for a relatively simple problem. The unique design constraints of the first experiment are as follows:

3.7.1.1 Velocity

To simplify the control system, we opted for simple constant-velocity control mapping meaning that the crosshair moved at a constant speed regardless of contraction

intensity.

3.7.1.2 UHC Confidence Threshold

The confidence threshold was empirically tuned during pilot studies to give the best outcome for UHC. Because the concept drift between training and testing is small and the confidence distribution of LDA is skewed high [50], the threshold was set to 100%. This conservative confidence threshold of 100% was chosen to adapt the model in the most cautious way possible. Even at this threshold, approximately 50% of within-game decisions were used to adapt the model.

3.7.1.3 Negative CIIL

In addition to the mechanisms of N-CIIL described above, we assumed any movement within the acquired target to be unwarranted and thus used it as negative context.

3.7.1.4 Order

The trial order was randomized across participants in this within-participant study. The six unique trials included the use of the screen guided training model (SGT), the initial one second model (initial), and the four unsupervised strategies (UHC, P, N, P+N).

3.7.1.5 Gameplay Time

The baseline model was adapted during a two-minute gameplay session – the approximate equivalent of acquiring five repetitions of offline SGT data.

3.7.2 Experiment Scenario 2: Electrode Shift

In a second, separate, phase of the study, the efficacy of the CIIL strategies was evaluated when a model's input space had been catastrophically shifted. An additional 21 participants were recruited for this experiment (8 female, 13 male; ages 18-41). Participants trained an initial model using five repetitions of ramp contractions, increasing from rest to their desired activation level over three seconds, for each input class. Then, after acquiring the baseline performance of the model in an ISO Fitts' test, participants were asked to take off and rotate the EMG cuff by 45 degrees clockwise from their perspective ($\sim$ one electrode). This simulated doffing and donning of the EMG cuff in a different location with the amount of shift likely being even more extreme than what might be reasonably expected in a prosthetic socket. While dealing with electrode shift through automatic approaches is an ongoing area of research [51], its role in this experiment was intended as a proxy for any time the input space has drastically shifted, such as after doffing and donning a device, between-day control, or even after cross-user adaptation [22]. The unique design constraints of the second experiment were as follows:

3.7.2.1 Velocity

A proportional control mapping was used, meaning harder contractions resulted in faster speeds. Subject's class-specific mean absolute value was used to determine 20% and 70% thresholds that formed a linear mapping for 0% to 100% of the available cursor speed, as per [52].

3.7.2.2 UHC Confidence Threshold

Given that the concept drift introduced by the electrode shift was large, a lower confidence threshold of 99% was found to be best for UHC from pilot studies. This selection allowed UHC to include some incoming samples to attempt to improve; whereas too few samples would have been gathered with a threshold of 100%.

3.7.2.3 Negative CIIL

The negative adaptation from the first experiment was modified so that only within-target movement that drifted away from the center was considered as negative context.

3.7.2.4 Order

The order of the trials (P, N, P+N, UHC) was balanced across participants using a Latin square. Each user started with the pre-shifted SGT model.

3.7.2.5 Gameplay Time

To observe a plateau in the adapted model that wasn't seen with the two minutes of adaptation from experiment one, an adaptation time of five minutes was selected.

3.7.2.6 Questionnaire

Participants completed a NASA TLX survey after the ISO Fitts' test to better gauge the perceived workload of controlling the cursor with each adapted model. The questionnaire queried the quality of the models *after* adaptation had occurred, and does not reflect the user's perception of the model during the gameplay session.

3.7.3 Statistical Analyses

Statistical analyses were conducted using the Statistical Tests for Algorithms Comparison (STAC) platform [53]. Analyses were performed independently for each experiment to determine differences between learning strategies from observed offline metrics (classification accuracy, active error, instability) and online metrics (overshoots, path efficiency, throughput). For each metric, normality and homogeneity of variance were tested and Friedman ranking tests were determined to be appropriate.

Following a significant outcome, post-hoc multiple comparison tests were conducted using Finner correction [54] with a confidence level of 0.05. The resulting multiple comparison tests were visualized with critical difference diagrams. Within these diagrams, the mean normalized rank of the pipelines were determined and plotted on a number line. The lowest ranked pipeline is indicative of the lowest values among the pipelines for the metric being analyzed, whereas the highest rank is indicative of the highest values among the pipelines for the metric being analyzed. As such, the most desirable (best performing) rank for the metrics analyzed in the study were highest for accuracy, lowest for active error, lowest for instability, lowest for overshoots, highest for efficiency, and highest for throughput. Finally, these diagrams also illustrate when pipelines are not significantly different by grouping them together with a horizontal bar.

3.8 Results

The quantitative outcomes of scenario 1 were compiled into the first column of Figure 3.4. For offline metrics, the SGT model was found to have significantly higher classification accuracy than the initial model, N-CIIL, and UHC adaptation

(91.28% vs. 65.75%, 72.47%, and 79.05%, respectively; $p < 0.05$), but not P+N-CIIL or P-CIIL (85.26% and 86.01%). For active error, the P+N-CIIL and SGT models significantly lower than the initial model, N-CIIL, and UHC (10.53% and 6.22% vs. 32.62%, 18.97%, and 18.23%, respectively; $p < 0.05$), but not P-CIIL (11.53%). Likewise, instability was significantly lower for SGT, P-CIIL, and P+N-CIIL compared to the initial model, N-CIIL, and UHC (5.00%, 7.77%, and 8.13% vs. 19.92%, 16.12%, and 12.08%, respectively; $p < 0.05$). For the online metrics, N-CIIL had significantly lower overshoots compared to the initial model (2.27 vs. 13.55; $p < 0.05$), but not P-CIIL, P+N-CIIL, SGT, and UHC (3.00, 3.64, 4.45, and 6.09, respectively). Path efficiency of SGT, N-CIIL, P-CIIL, and P+N-CIIL was significantly higher than the initial model (68.45%, 67.66%, 70.77%, and 70.14% vs. 42.55%, respectively; $p < 0.05$). Finally, the throughput of P-CIIL was significantly higher than the initial model, UHC, and N-CIIL (0.43 bit/s vs. 0.24 bit/s, 0.35 bit/s, and 0.35 bit/s, respectively; $p < 0.05$), but not SGT or P+N-CIIL (0.41 bit/s and 0.40 bit/s).

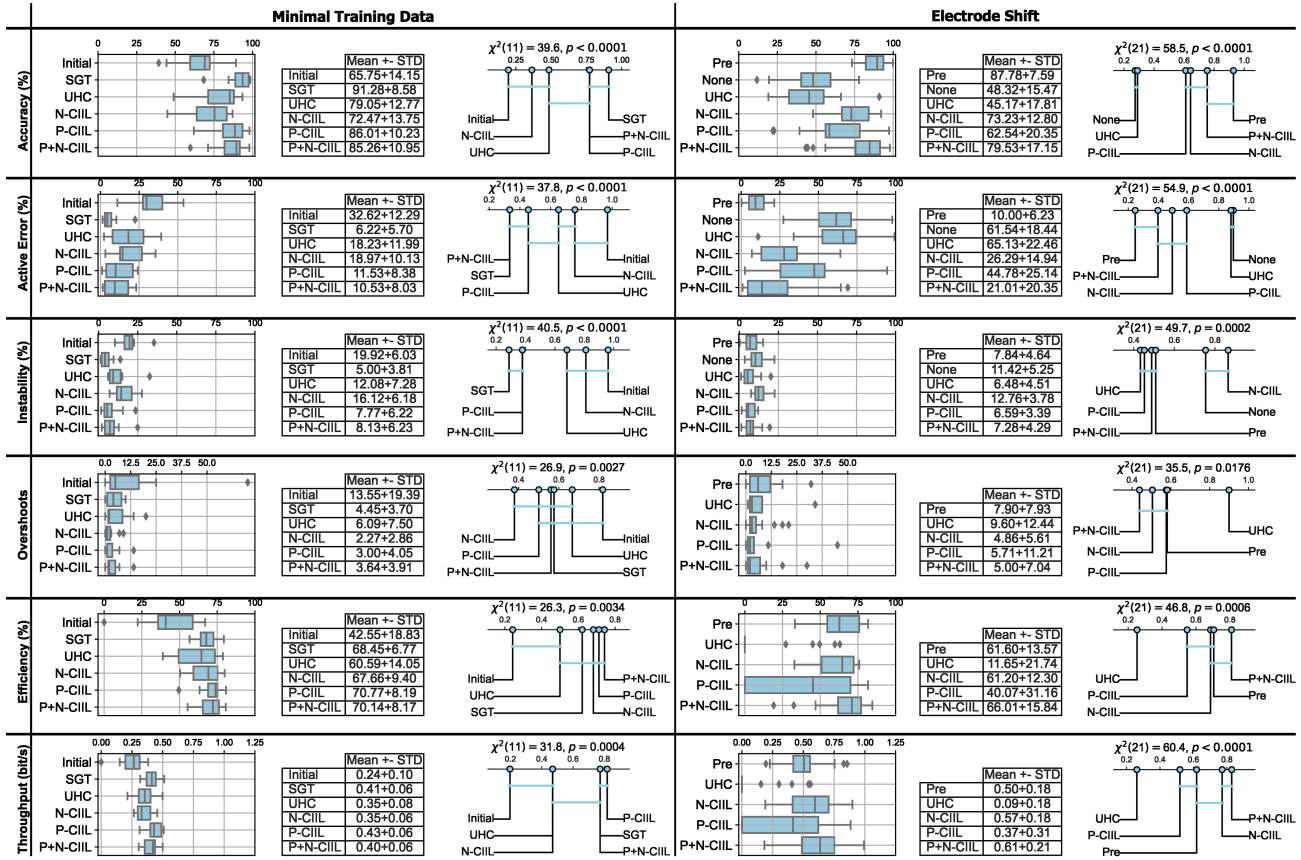


Figure 3.4: Box plots, summary statistics, and critical difference diagrams for offline metrics (accuracy, active error, instability), and online metrics (overshoots, efficiency, throughput) in both evaluated scenarios (minimal training data, electrode shift). Ideal trends for values and normalized ranks of these metrics are high accuracy, low active error, low instability, low overshoots, high efficiency, and high throughput. Blue bars linking groups in the critical difference diagrams signifies a non-significant difference within the multiple comparison tests. Note: when subjects did not recover from the introduced electrode shift and could not complete the online assessment, their data was removed for the overshoots metric as a placeholder value was not appropriate.

The quantitative outcomes for scenario 2 were compiled into the second column of Figure 3.4. For offline metrics, the pre-shift scenario had significantly higher classification accuracy than the shifted scenario with no adaptation, UHC, P-CIIL, and N-CIIL (87.78% vs. 48.32%, 45.17%, 62.54%, and 73.23%, respectively; $p < 0.05$), but no difference was found with P+N-CIIL (79.53%). The active error multiple comparison had similar outcomes, with the pre-shifted setting significantly outperforming the shifted scenario with no adaptation, UHC, P-CIIL, and N-CIIL (10.00% vs. 61.54%, 65.13%, 44.78%, and 26.29%, respectively; $p < 0.05$), but not P+N-CIIL (21.01%). As a consequence of not recovering all gestures during adaptation for most subjects, UHC had the lowest instability (manifested by failing to elicit all classes); however, no significant differences were found between UHC, P-CIIL, P+N-CIIL, and the pre-shifted setting (6.48%, 6.59%, 7.28%, and 7.84%, respectively) – of which P+N-CIIL and the pre-shifted setting could elicit all classes for all subjects. For the online metrics, UHC had significantly more overshoots than the pre-shifted setting, P-CIIL, N-CIIL, and P+N-CIIL (9.60 vs. 7.90, 5.71, 4.86, and 5.00, respectively; $p < 0.05$). Path efficiency of P+N-CIIL was significantly better than P-CIIL and UHC (66.01% vs. 40.07% and 11.65%, respectively; $p < 0.05$), and was marginally better than the pre-shifted setting and N-CIIL (61.60% and 61.20%, respectively). Finally and most importantly, the throughput of P+N-CIIL was significantly higher than all other methods including the pre-shifted setting (0.61 bit/s vs. 0.50 bit/s; $p < 0.05$), but was only marginally higher than N-CIIL (0.57 bit/s).

The qualitative outcomes for scenario 2, which were extracted from the NASA TLX survey, are shown in Figure 3.5. Reported frustration levels were lowest for P+N-CIIL, N-CIIL, and the pre-shift model with the P-CIIL and UHC models being significantly higher (23.81, 24.29, and 25.24 vs. 47.38 and 69.76, respectively; $p < 0.05$). In addition to P+N-CIIL having the highest throughput (ob-

jective performance measure), P+N-CIIL was also reported to be the significantly best-performing model (subjectively perceived performance measure, 82.62). The pre-shifted SGT, P+N-CIIL, and N-CIIL models had significantly lower physical demand than the P-CIIL and UHC models (23.81, 25.00, and 27.14 vs. 44.52 and 59.05, respectively; $p < 0.05$). Likewise, P+N-CIIL and the pre-screening SGT model had significantly lower overall effort compared to P-CIIL and UHC (32.62 and 34.52 vs. 52.86 and 70.00, respectively; $p < 0.05$), but not N-CIIL (38.57).

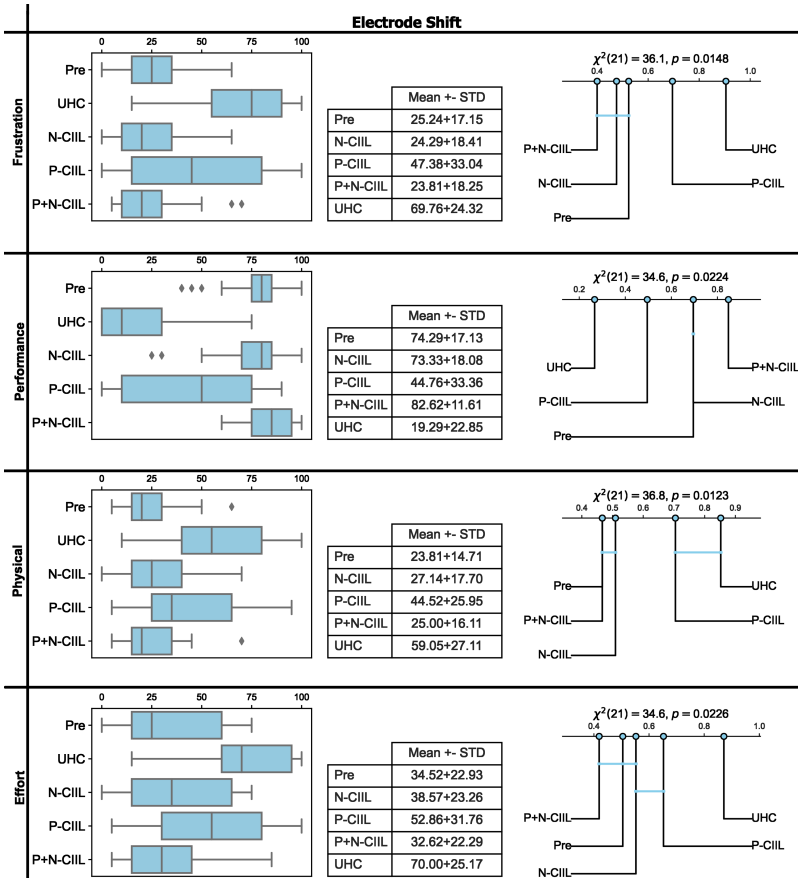


Figure 3.5: Box plots, summary statistics, and critical difference diagrams for subjective ratings of the different adaptation methods (frustrating, performance, physical effort, and overall effort). Ideal ratings and normalized ranks for these categories are low frustration, high performance, low physical demand, and low overall effort. Blue bars linking groups in the critical difference diagrams signifies a non-significant difference between linked strategies.

3.9 Discussion

3.9.1 Adaptation Dynamics

In evaluating the various myoelectric control adaptation strategies across the two experiments of the study, clear differences between approaches emerged.

Unsupervised high-confidence adaptation – the current standard for unsupervised learning – was beneficial when the system was already performing well, as demonstrated in the minimal training data experiment (i.e., Experiment 1). Because the class distributions were sufficiently informed, the UHC approach was able to continue to refine them appropriately. However, it does not yield any measurable benefits when the initial model is underperforming and would benefit the most from adaptation. Within the second experiment, only 5 of 21 subjects had sufficient recovery after the electrode shift to complete the usability test, and even in these cases throughput after UHC adaptation was drastically lower than the CIIL approaches. The failure of UHC to recover after catastrophic shift is a crucial drawback for its viability to maintain system robustness for myoelectric control where these perturbations are common. Nevertheless, UHC does not have the same contextual reliance as the other approaches, highlighting its usefulness and potential practicality for maintaining already robust models with minimal overhead when contextual information in unavailable.

P-CIIL demonstrated similarities with UHC, allowing the model to enhance its performance following successful actions. However, P adaptation solely focuses on positive outcomes, ignoring negative actions which are most commonly the result of misclassifications. It achieved significantly higher throughput than UHC when the initial model was reasonable (Experiment 1), and was more reliable than UHC when the initial model was performing poorly (Experiment 2). Its shortcomings

primarily emerge when the classifier aligns with a 'correct' contextual direction that was not intended by the user. In the environment used in this study, this could arise when the planet is diagonal from the crosshair (up and left), the user tries to go left but the classifier went up. In this scenario, P-CIIL ends up reinforcing this inadvertent direction, albeit less frequently than UHC thanks to the reliance of contextual suitability of the direction. P-CIIL also runs into difficulties when the user was completely unable to evoke any positive context (for example when they get stuck in a certain mode). Consequently, only 14 of 21 subjects were able to recover after the electrode shift to a high throughput model when relying on positive context alone.

N-CIIL offers a contrasting approach to P-CIIL by enabling model recovery in the presence of mistakes, regardless of classifier performance. Even when employing a simplistic tie-breaking system dependent on the classifier's best guess among the valid options, N-CIIL consistently demonstrated recovery capabilities, as evident in the electrode shift experiment. This holds true, even though the adaptation process had to choose between two potential behaviors during adaptation, inevitably leading to erroneous information being provided at times. Consequently, every subject was able to recover after the electrode shift to a usable state (i.e., they could complete the online Fitts' law task). However, in longer adaptation scenarios, its performance may plateau as errors decrease in frequency, leading to a weakened adaptation signal. A hybrid approach, combining positive and negative adaptation (P+N-CIIL), capitalizes on mistakes in the early stages and leverages positive interactions as they emerge. Generally, N-CIIL can recover a completely flawed model to the point where it can then benefit from P-CIIL. In doing so, this strategy harnesses the strengths of both adaptation approaches and exhibits the highest throughput in both experiments. This shift from negative to positive context for the P+N-CIIL pipeline is shown in Figure 3.6, where P+N-CIIL began by using more negative con-

text to accommodate the challenging setting, before transitioning to predominantly positive context after initial improvements. P+N-CIIL not only recovered to a usable level for every subject, but achieved a significantly higher throughput to the pre-shift screen guided training model. These findings underscore the significance of selecting an appropriate adaptation strategy based on the specific performance context and objectives in myoelectric control systems

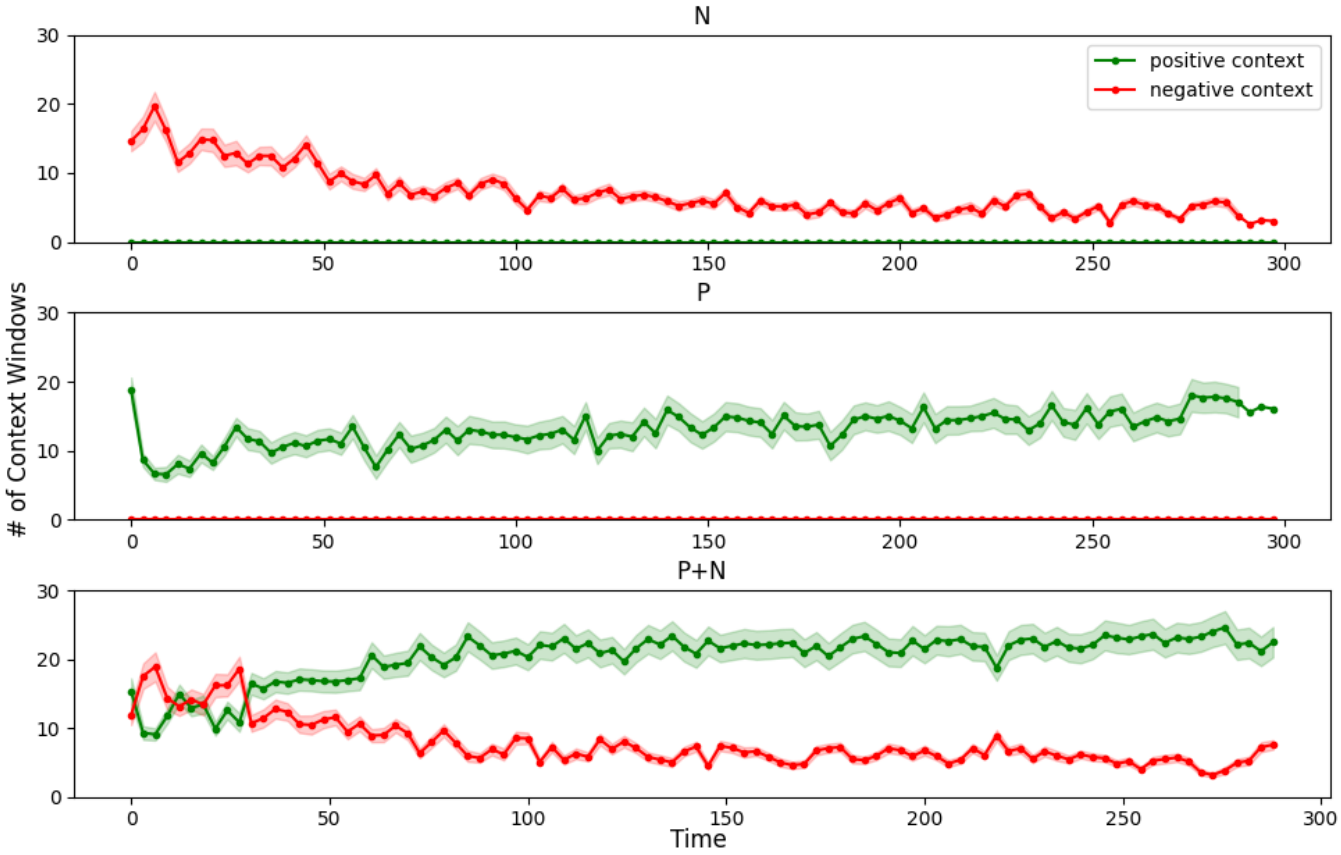


Figure 3.6: Representation of the number of context windows used by the CIIL pipelines the electrode shift scenario. Standard error across subjects is shown by a shaded region. Each point represents a three second time frame.

3.9.2 Qualitative Feedback

In addition to quantitative measures of performance (throughput) being better post-shift post-adaptation in Scenario 2, self-reported qualitative measures of performance (extracted from the NASA TLX) showed users preferred the P+N-CIIL adapted model over the pre-shift SGT model. From the reported performance metric, P+N-CIIL was significantly better than all other models. Further, P+N-CIIL marginally outperformed the pre-shift SGT model in regards to having lower perceived effort and lower frustration. Overall, these qualitative measures corroborate the quantitative results, further supporting P+N-CIIL as the superior option.

3.9.3 Offline vs. Online Performance

Offline accuracy is often used as an indicator for anticipated model performance during user-in-the-loop settings, however, the associative relationship between these two metrics is often debated. As highlighted in section 3.5.1, offline accuracy has a weak associative relationship to throughput and therefore should not be used to decide what algorithm would result in the best online control [10]. Recently, Hinson *et al.* demonstrated that offline metrics do correlate, albeit on a granular level, with online performance by training EMG to kinematics models to different levels of described variances (R^2 =0.8, 0.6, 0.4) and performing Fitts' law tests [55]. Hinson suggests the reason for the stronger relationship was due to experimental challenges in other works (altered outputs from controller gain [56], inadequate task duration and feedback [57]).

Despite having intentionally introduced minimal data and electrode shift as experimental challenges in this study, we still observed "strong" correlations between offline accuracy and online throughput for both scenarios (global correlations of 0.74

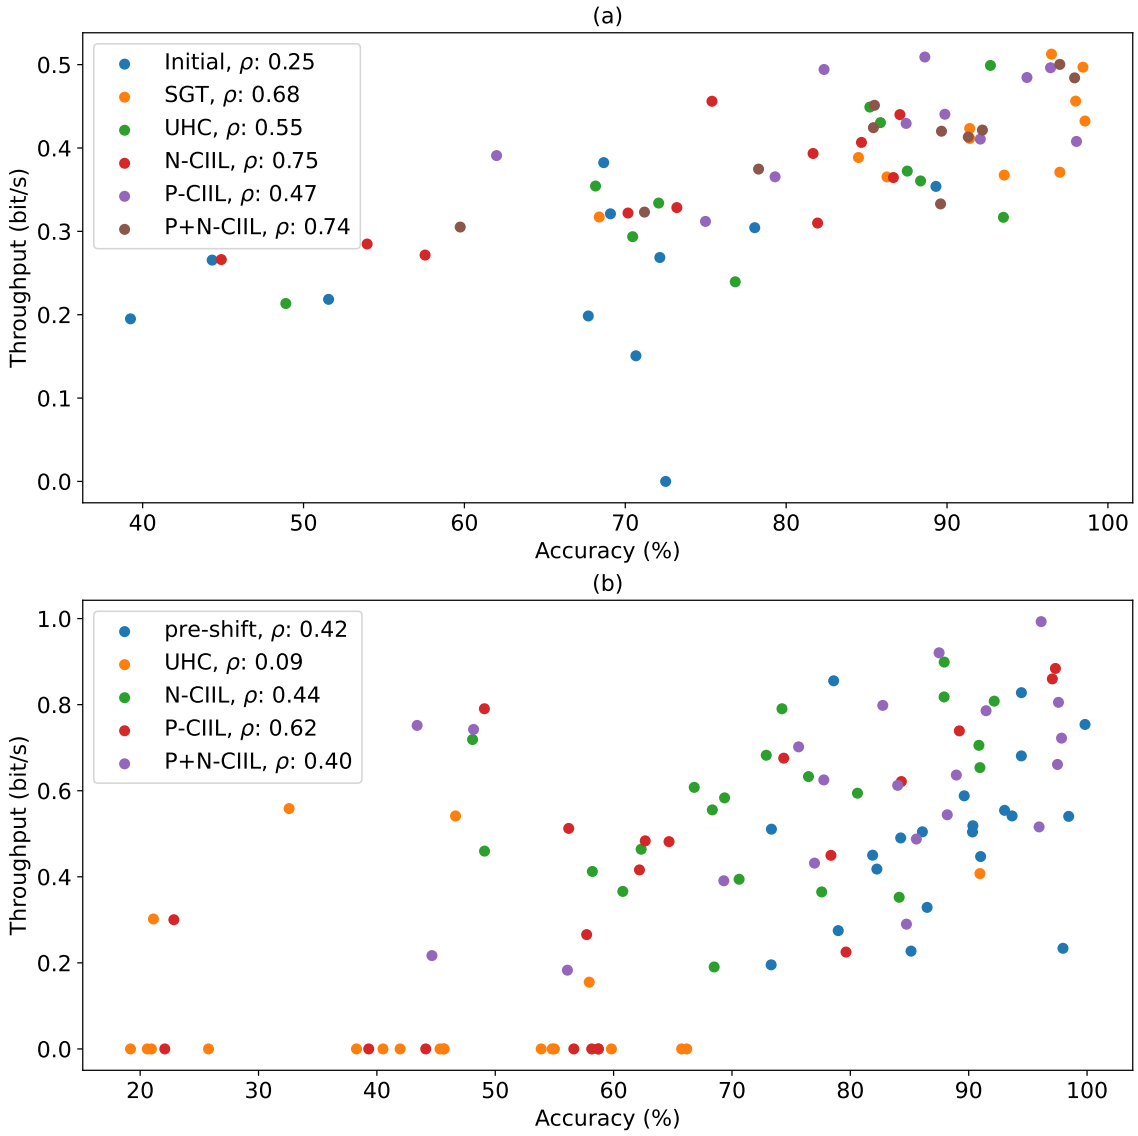


Figure 3.7: Scatterplot of associative relationships between offline accuracy and online throughput from (a) Scenario 1 (Minimal Training Data) and (b) Scenario 2 (Electrode Shift). Grouped Spearman's correlation coefficients are presented in the legend. Global correlations, as calculated using all data irrespective of model, were $\rho = 0.74, 0.62$, for Scenario 1 and 2, respectively.

and 0.62 for Scenarios 1 and 2, respectively, as shown in Figure 3.7). Nevertheless, for the electrode shift scenario, the model with the highest accuracy differed from that with the highest throughput. In fact, although SGT produced the highest accuracy, P+N-CIIL significantly outperformed it in terms of throughput ($p < 0.05$). This discrepancy is justified given the behavioural differences a user exhibits while in-the-loop, but it illustrates that subtle differences in offline classification accuracy cannot be used naively to justify model rankings.

Moreover, classification accuracy computed on steady-state or ramp contractions may become irrelevant when it comes to incremental learning approaches such as those proposed in this work. For example, there were multiple cases where users unintentionally combined the hand open and closed contractions with radial and ulnar deviation, respectively, as they found this more intuitive to control the up and down directions of the cursor. While this could lead to improved within-task outcomes (i.e., increased throughput), it may also degrade the classification accuracy. Correspondingly, this tradeoff between these two metrics must be considered as we continue to move toward adaptive systems, and a particular emphasis should be placed on how to best evaluate these systems.

3.9.4 Practicality of CIIL

While this study offers an important initial demonstration and validation of CIIL, the evaluation environment was relatively simple compared to the diverse settings where myoelectric control might be applied, such as dexterous manipulations of prosthetic limbs or mixed reality scenarios. Nevertheless, it is worth emphasizing that the potential benefits of CIIL may be most pronounced in these more complex and dynamic settings, if situational context can be appropriately measured or inferred.

In prosthetics, determining the suitability of control actions would require knowledge of appropriate gestures given the task being conducted. For example, near field communication tags could be used in the home to notify the device of an appropriate gesture when reaching for certain objects. Alternatively, additional sensors (e.g., cameras and force sensors) could be added to a prosthesis to provide context about the successful grasp or accidental release of an object.

For general human-computer interaction using myoelectric control, the greater availability of contextual information in digital environments could help to develop rich and intuitive control. Each successful click of a button, or the detection of a click when no button is available, could provide context with which to inform model adaptation. Quick corrective actions, such as playing and immediately pausing a music application or passing over then quickly returning to a selection menu, could provide the context of an inadvertent command. By assigning appropriate labels, the system could iteratively adapt and reduce similar classification errors from arising for similar user inputs.

Recognizing that the quality of context engineering impacts the effectiveness of CIIL, the instrumentation does not need to be ubiquitous. Even when only partial access to contextual information is available, such as within a user's home or during the introduction of new control actions as part of a curriculum, substantial benefits can still be gained. For example, a prosthesis could be trained using CIIL in a rehabilitation environment where context is easy to extract (e.g. using instrumented environments or tests). Furthermore, while CIIL assumes consistent goal-driven behavior from users, it acknowledges the real-world complexity where users may not always adhere to this paradigm. For instance, naturally-limbed individuals gesticulate while talking; however, the context for this scenario is foreseeably unobtainable and as a result CIIL would likely adapt the underlying activity to no motion. As

such, there is future work in exploring how to benefit from partial contextual feedback and how to not altogether remove creative expression.

3.9.5 Future Work

This work provides a foundational basis for CIIL approaches that maintain the accuracy of a classifier given a contextual environment; however, there still remain numerous avenues of research to explore and improve the capabilities of this approach.

Future work should investigate the efficacy of CIIL within more realistic, physical environments. Although the instrumentation of these settings is not trivial (vision based, NFC tags, etc.), integrating within existing smart home technologies or commonly used items could prove valuable for all populations that use myoelectric control (general HCI, stroke, spinal cord injury, muscular distrophy, amputees, etc.).

Although this work demonstrated robustness to two confounding factors (minimal training data and electrode shift), there are many other sources of variability that could be minimized with this incremental learning approach (limb position variability, contraction intensity variability, hands-busy interference, between-session differences, cross-subject differences). Electrode shift degradation is typically addressed with recalibration or by making representations of multiple electrode configurations [58]. However, this work demonstrated that an online incremental learning approach could offer a viable solution, adapting to the current electrode position rather than attempting to describe multiple positions.

Within this experiment, a virtual environment was used that had perfect situational context available for every myoelectric decision; however, this may not be the case for real-world systems, such as one relying on object detection for contextual infor-

mation. In such a scenario, it would be vital to understand the relationship between the accuracy of contextual information and the trajectory of the CIIL training process. Future work should evaluate this relationship to understand the lower-bounds of context accuracy and frequency required for CIIL systems to reliably improve myoelectric usability.

CIIL demonstrated it could improve the classifier's predictions over time given contextual suitability; however, it could be beneficial to adopt this approach within regression problems – as these are also prevalent in velocity-based myoelectric control applications. Similarly, CIIL should be extended beyond 2-DOF problems to determine if ambiguity among the potential options greatly increases the time required to reach a usable model. While this work demonstrated that CIIL is an effective framework for classification-based online learning that is invariant to initial model performance, future work could leverage structured Bayesian models [59], auto-regressive or nonlinear regression approaches [60, 61], or deep learning to extend its benefits to other tasks such as regression.

Perhaps most importantly, we hope to reduce the complexity involved in implementing CIIL approaches to facilitate the uptake of these methods. This work required constant communication between different processes performing user-in-the-loop classification, collecting the associated context, and periodically performing adaptation. We aim to provide the supporting infrastructure for this style of experiment in LibEMG in the future [45].

3.10 Conclusion

In this paper we presented Context-Informed Incremental Learning (CIIL), a new incremental learning framework for myoelectric control. In two online studies, CIIL

was shown to significantly outperform the state-of-the-art approach (unsupervised high-confidence adaptation) when beginning with a sparsely defined classifier trained with minimal data and when recovering from a catastrophic change in the input space due to electrode shift. Additionally, through the adoption of CIIL strategies, myoelectric control systems better incorporate user-in-the-loop behaviours which will in turn improve the ability of models to decipher user intent, resulting in better overall control.

3.11 Declarations

3.11.1 Ethics Approval and Consent to Participate

The procedures undertaken within this manuscript were approved by the University of New Brunswick Research Ethics Board (REB#2022-122). Prior to the subjects participating in the experiment, written consent was given.

3.11.2 Availability of Data and Materials

A link to the analysis code and data is provided within the manuscript (https://github.com/ECEEvanCampbell/CIIL_LDA).

3.11.3 Funding

This work was funded by the Natural Sciences and Engineering Research Council under grant 2020-04776.

3.11.4 Author Contributions

EC and EE contributed in all aspects of the experiment (conception, design, acquisition, analysis, interpretation, software creation, manuscript preparation (all text and figures), and editing). SB contributed to the interpretation and editing of the manuscript. UCA contributed to the acquisition, analysis, interpretation, and editing of the manuscript. ES contributed to the conception, design, analysis, interpretation, and editing of the manuscript.

Bibliography

[1] E. Scheme and K. Englehart, "Electromyogram pattern recognition for control of powered upper-limb prostheses: state of the art and challenges for clinical use," *Journal of Rehabilitation Research and Development*, vol. 48, no. 6, pp. 643–659, 2011.

[2] E. Eddy, E. J. Scheme, and S. Bateman, "A framework and call to action for the future development of EMG-based input in HCI," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, (New York, NY, USA), Association for Computing Machinery, 2023.

[3] R. Antonius and H. Tjahyadi, "Electromyography gesture identification using cnn-rnn neural network for controlling quadcopters," *Journal of Physics: Conference Series*, vol. 1858, p. 012075, apr 2021.

[4] M.-S. Song, S.-G. Kang, K.-T. Lee, and J. Kim, "Wireless, skin-mountable emg sensor for human–machine interface application," *Micromachines*, vol. 10, no. 12, 2019.

[5] F. Haque, M. Nancel, and D. Vogel, "Myopoint: Pointing and clicking using forearm mounted electromyography and inertial motion sensors," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, (New York, NY, USA), p. 3653–3656, Association for Computing Machinery, 2015.

[6] C. Savur and F. Sahin, "American sign language recognition system by using surface emg signal," in *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 002872–002877, 2016.

[7] L. B. Jaloza, "Inside Facebook Reality Labs: Wrist-based interaction for the next computing platform," *Tech at Meta*, Mar. 2021.

[8] E. Campbell, A. Phinyomark, and E. Scheme, "Current trends and confounding factors in myoelectric control: Limb position and contraction intensity," *Sensors*, vol. 20, no. 6, 2020.

[9] U. Cote-Allard, G. Gagnon-Turcotte, A. Phinyomark, K. Glette, E. Scheme, F. Laviolette, and B. Gosselin, "A transferable adaptive domain adversarial neural network for virtual reality augmented emg-based gesture recognition," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 29, pp. 546–555, 2021.

[10] B. Lock, K. Englehart, and B. Hudgins, "Real-time myoelectric control in a virtual environment to relate usability vs. accuracy," Myoelectric Symposium, 2005.

[11] J. W. Sensinger, B. A. Lock, and T. A. Kuiken, "Adaptive pattern recognition of myoelectric signals: Exploration of conceptual framework and practical algorithms," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 17, no. 3, pp. 270–278, 2009.

[12] H. Dantas, V. J. Mathews, and D. Warren, "Semi-supervised adaptive learning for decoding movement intent from electromyograms," in *2019 27th European Signal Processing Conference (EUSIPCO)*, pp. 1–5, IEEE, 2019.

[13] P. M. Pilarski, M. R. Dawson, T. Degris, F. Fahimi, J. P. Carey, and R. S. Sutton, "Online human training of a myoelectric prosthesis controller via actor-critic reinforcement learning," in *2011 IEEE international conference on rehabilitation robotics*, pp. 1–7, IEEE, 2011.

[14] J. L. Nawfel, K. B. Englehart, and E. J. Scheme, "A multi-variate approach to predicting myoelectric control usability," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 29, pp. 1312–1327, 2021.

[15] L. Hargrove, L. Miller, K. Turner, and T. Kuiken, "Control within a virtual environment is correlated to functional outcomes when using a physical prosthesis," *Journal of neuroengineering and rehabilitation*, vol. 15, no. 1, pp. 1–7, 2018.

[16] C. L. Chicoine, A. M. Simon, and L. J. Hargrove, "Prosthesis-guided training of pattern recognition-controlled myoelectric prosthesis," in *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 1876–1879, IEEE, 2012.

[17] H. Zhang, Y. Zhao, F. Yao, L. Xu, P. Shang, and G. Li, "An adaptation strategy of using lda classifier for emg pattern recognition," pp. 4267–4270, 2013.

[18] R. B. Woodward and L. J. Hargrove, "Adapting myoelectric control in real-time using a virtual environment," *Journal of NeuroEngineering and Rehabilitation*, vol. 16, p. 11, Jan. 2019.

[19] U. Côté-Allard, G. Gagnon-Turcotte, A. Phinyomark, K. Glette, E. J. Scheme, F. Laviolette, and B. Gosselin, "Unsupervised domain adversarial

self-calibration for electromyography-based gesture recognition," *IEEE Access*, vol. 8, pp. 177941–177955, 2020.

[20] M. A. Oskoei and H. Hu, "Adaptive myoelectric control applied to video game," *Biomedical Signal Processing and Control*, vol. 18, pp. 153–160, 2015.

[21] U. Côté-Allard, E. Campbell, A. Phinyomark, F. Laviolette, B. Gosselin, and E. Scheme, "Interpreting deep learning features for myoelectric control: A comparison with handcrafted features," *Frontiers in bioengineering and biotechnology*, vol. 8, p. 158, 2020.

[22] E. Campbell, A. Phinyomark, and E. Scheme, "Deep cross-user models reduce the training burden in myoelectric control," *Frontiers in Neuroscience*, vol. 15, p. 657958, 2021.

[23] K. Szymaniak, A. Krasoulis, and K. Nazarpour, "Recalibration of myoelectric control with active learning," *Frontiers in Neurorobotics*, vol. 16, p. 277, 2022.

[24] W. Wu, K. R. Saul, and H. Huang, "Using reinforcement learning to estimate human joint moments from electromyography or joint kinematics: An alternative solution to musculoskeletal-based biomechanics," *Journal of Biomechanical Engineering*, vol. 143, no. 4, p. 044502, 2021.

[25] J. P. Vásconez, L. I. Barona López, Á. L. Valdivieso Caraguay, and M. E. Benalcázar, "Hand gesture recognition using emg-imu signals and deep q-networks," *Sensors*, vol. 22, no. 24, p. 9613, 2022.

[26] A. L. Edwards, A. Kearney, M. R. Dawson, R. S. Sutton, and P. M. Pilarski, "Temporal-difference learning to assist human decision making during the control of an artificial limb," *arXiv preprint arXiv:1309.4714*, 2013.

[27] W. B. Knox and P. Stone, "Interactively shaping agents via human reinforcement: The tamer framework," in *Proceedings of the fifth international conference on Knowledge capture*, pp. 9–16, 2009.

[28] J. P. Hanna, S. Niekum, and P. Stone, "Importance sampling in reinforcement learning with an estimated behavior policy," *Machine Learning*, vol. 110, no. 6, pp. 1267–1317, 2021.

[29] G. D. Abowd, A. K. Dey, P. J. Brown, N. Davies, M. Smith, and P. Steggles, "Towards a better understanding of context and context-awareness," in *Handheld and Ubiquitous Computing* (H.-W. Gellersen, ed.), (Berlin, Heidelberg), pp. 304–307, Springer Berlin Heidelberg, 1999.

[30] J.-y. Hong, E.-h. Suh, and S.-J. Kim, "Context-aware systems: A literature review and classification," *Expert Systems with applications*, vol. 36, no. 4, pp. 8509–8522, 2009.

[31] B. Schilit, N. Adams, and R. Want, "Context-aware computing applications," in *1994 First Workshop on Mobile Computing Systems and Applications*, pp. 85–90, 1994.

[32] A. Schmidt, M. Beigl, and H.-W. Gellersen, "There is more to context than location," *Computers & Graphics*, vol. 23, no. 6, pp. 893–901, 1999.

[33] S. Kulkarni and S. F. Rodd, "Context aware recommendation systems: A review of the state of the art techniques," *Computer Science Review*, vol. 37, p. 100255, 2020.

[34] O. Buyukkokten, J. Cho, H. Garcia-Molina, L. Gravano, and N. Shivakumar, "Exploiting geographical location information of web pages," in *ACM SIGMOD Workshop on The Web and Databases, WebDB 1999, Philadelphia, Pennsylva-*

nia, USA, June 3-4, 1999. Informal Proceedings (S. Cluet and T. Milo, eds.), pp. 91–96, INRIA, 1999.

[35] S. T. P. Raghu, D. MacIsaac, and E. Scheme, "Decision-change Informed Rejection Improves Robustness in Pattern Recognition-based Myoelectric Control," 4 2023.

[36] E. J. Scheme, B. S. Hudgins, and K. B. Englehart, "Confidence-based rejection for improved pattern recognition myoelectric control," *IEEE Transactions on Biomedical Engineering*, vol. 60, no. 6, pp. 1563–1570, 2013.

[37] J. W. Sensinger, B. A. Lock, and T. A. Kuiken, "Adaptive pattern recognition of myoelectric signals: Exploration of conceptual framework and practical algorithms," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 17, no. 3, pp. 270–278, 2009.

[38] G. K. Patel, J. M. Hahne, C. Castellini, D. Farina, and S. Dosen, "Context-dependent adaptation improves robustness of myoelectric control for upper-limb prostheses," *Journal of Neural Engineering*, vol. 14, p. 056016, sep 2017.

[39] M. A. Trout, T. C. Hansen, C. D. Olsen, D. J. Warren, J. L. Segil, and G. J A, "Shared control decreases the physical and cognitive demands of maintaining a secure grip," in *Proceedngs of the Myoelectric Controls and Upper Limb Prosthetics Symposium*, 2022.

[40] G. Ghazaei, A. Alameer, P. Degenaar, G. Morgan, and K. Nazarpour, "Deep learning-based artificial vision for grasp classification in myoelectric hands," *Journal of neural engineering*, vol. 14, no. 3, p. 036025, 2017.

[41] K. Z. Zhuang, N. Sommer, V. Mendez, S. Aryan, E. Formento, E. D'Anna, F. Artoni, F. Petrini, G. Granata, G. Cannaviello, *et al.*, "Shared human–robot

proportional control of a dexterous myoelectric prosthesis," *Nature Machine Intelligence*, vol. 1, no. 9, pp. 400–411, 2019.

[42] J. W. Sensinger and S. Dosen, "A review of sensory feedback in upper-limb prostheses from the perspective of human motor control," *Frontiers in Neuroscience*, vol. 14, p. 345, 2020.

[43] B. Hudgins, P. Parker, and R. Scott, "A new strategy for multifunction myoelectric control," *IEEE Transactions on Biomedical Engineering*, vol. 40, no. 1, pp. 82–94, 1993.

[44] I. Al-Tahir, J. Sensinger, and E. Scheme, "A better framework for the assessment of performance and stability of co-adaptive myoelectric systems," in *2022 International Conference on Rehabilitation Robotics (ICORR)*, pp. 1–5, 2022.

[45] E. Eddy, E. Campbell, A. Phinyomark, S. Bateman, and E. Scheme, "Libemg: An open source library to facilitate the exploration of myoelectric control," *IEEE Access*, pp. 1–1, 2023.

[46] E. Campbell, A. Phinyomark, and E. Scheme, "Linear discriminant analysis with bayesian risk parameters for myoelectric control," in *2019 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pp. 1–5, IEEE, 2019.

[47] I. S. MacKenzie, "Fitts' law as a research and design tool in human-computer interaction," *Human–Computer Interaction*, vol. 7, no. 1, pp. 91–139, 1992.

[48] E. J. Scheme and K. B. Englehart, "Validation of a selective ensemble-based classification scheme for myoelectric control using a three-dimensional fitts' law test," *IEEE transactions on neural systems and rehabilitation engineering*, vol. 21, no. 4, pp. 616–623, 2012.

[49] E. Eddy, E. Campbell, S. Bateman, and E. Scheme, "Leveraging task-specific context to improve unsupervised adaptation for myoelectric control," in *2023 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. NA–NA, IEEE, 2023.

[50] E. Scheme and K. Englehart, "A comparison of classification based confidence metrics for use in the design of myoelectric control systems," in *2015 37th annual international conference of the IEEE engineering in medicine and biology society (EMBC)*, pp. 7278–7283, IEEE, 2015.

[51] L. Wang, X. Li, Z. Chen, Z. Sun, and J. Xue, "Electrode shift fast adaptive correction for improving myoelectric control interface performance," *IEEE Sensors Journal*, pp. 1–1, 2023.

[52] E. Scheme, B. Lock, L. Hargrove, W. Hill, U. Kuruganti, and K. Englehart, "Motion normalized proportional control for improved pattern recognition-based myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 22, no. 1, pp. 149–157, 2014.

[53] I. Rodríguez-Fdez, A. Canosa, M. Mucientes, and A. Bugarín, "STAC: a web platform for the comparison of algorithms using statistical tests," in *Proceedings of the 2015 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, 2015.

[54] H. Finner, "On a monotonicity problem in step-down multiple test procedures," *Journal of the American statistical association*, vol. 88, no. 423, pp. 920–923, 1993.

[55] R. M. Hinson, J. Berman, W. Filer, D. Kamper, X. Hu, and H. Huang, "Offline evaluation matters: Investigation of the influence of offline performance on real-time operation of electromyography-based neural-machine interfaces,"

IEEE Transactions on Neural Systems and Rehabilitation Engineering, vol. 31, pp. 680–689, 2022.

[56] N. Jiang, I. Vujaklija, H. Rehbaum, B. Graimann, and D. Farina, "Is accurate mapping of emg signals on kinematics needed for precise online myoelectric control?," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 22, no. 3, pp. 549–558, 2013.

[57] A. Krasoulis, S. Vijayakumar, and K. Nazarpour, "Effect of user practice on prosthetic finger control with an intuitive myoelectric decoder," *Frontiers in neuroscience*, vol. 13, p. 891, 2019.

[58] F. Chamberland, É. Buteau, S. Tam, E. Campbell, A. Mortazavi, E. Scheme, P. Fortier, M. Boukadoum, A. Campeau-Lecours, and B. Gosselin, "Novel wearable hd-emg sensor with shift-robust gesture recognition using deep learning," *IEEE Transactions on Biomedical Circuits and Systems*, 2023.

[59] S. Yoneda and A. Furui, "Bayesian approach for adaptive emg pattern classification via semi-supervised sequential learning,"

[60] C. Igual, J. Igual, J. M. Hahne, and L. C. Parra, "Adaptive auto-regressive proportional myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 27, no. 2, pp. 314–322, 2019.

[61] A. Gijsberts, R. Bohra, D. Sierra González, A. Werner, M. Nowak, B. Caputo, M. A. Roa, and C. Castellini, "Stable myoelectric control of a hand prosthesis using non-linear incremental learning," *Frontiers in neurorobotics*, vol. 8, p. 8, 2014.

Chapter 4

Screen Guided Training Does Not Capture Goal-Oriented Behaviours: Learning Myoelectric Control Mappings From Scratch Using Context-Informed Incremental Learning

Evan Campbell, Ethan Eddy, Xavier Isabel, Scott Bateman, Benoit Gosselin, Ulysse Côté-Allard, Erik Scheme

4.1 Foreword

This chapter presents the next incarnation of classification-based context-informed incremental learning, using neural networks. The linear discriminant analysis classifier explored in Chapter 3 required hard labels, which forced a selection among goal-directed options for generating a pseudo-label. In contrast, neural networks allow soft labels, which enable the use of a pseudo-label created by equally weighting among goal-directed options. As a separate contribution outside of CIIL, the incremental learning of proportional control thresholds and the use of a nonlinear transfer function were also explored. Finally, this experiment explores incremental learning from an uninitialized model. My contribution to this chapter spans ideation, experimental design, data collection, data analysis, manuscript preparation, and revision corrections.

4.2 Reference

This work was peer-reviewed and published under ©2024 IEEE. Reprinted, with permission, from:

E. Campbell, E. Eddy, X. Isabel, S. Bateman, B. Gosselin, U. Côté-Allard, and E. Scheme, *Screen Guided Training Does Not Capture Goal-Oriented Behaviours: Learning Myoelectric Control Mappings From Scratch Using Context Informed Incremental Learning,* Transactions on Neural Systems and Rehabilitation Engineering, vol. 33, no. 1, p. 332, 2024.

4.3 Abstract

Human-machine interfaces based on myoelectric signals typically use screen-guided training (SGT) for model calibration, but this approach fails to capture realistic user behaviors. This study evaluates a user-in-the-loop context-informed incremental learning (CIIL) framework, comparing SGT, SGT followed by CIIL adaptation (SGT-A), and a novel zero-shot adaptation (ZS-A) CIIL approach that begins adapting with no prior training. Sixteen participants completed a Fitts' Law targeting task using these control schemes, with performance measured via online throughput and offline classification accuracy. Despite lower offline accuracy, the ZS-A model achieved the highest online throughput (1.47 ± 0.46 bits/s), significantly outperforming the SGT baseline (1.15 ± 0.37 bits/s) and reached competitive performance within 200 seconds. To further enhance control performance, a novel adaptive sigmoid-based proportional control mapping was introduced, dynamically adjusting control signals to allow precise control near neutral positions and rapid movements at higher activation levels, better aligning with natural user behaviors. These findings demonstrate that CIIL can surpass traditional SGT methods in online performance and emphasize the value of real-time user-in-the-loop data for developing adaptable and intuitive myoelectric interfaces, with implications for prosthetics, rehabilitation, and telerobotics.

4.4 Introduction

Pattern recognition-based myoelectric control, which uses electrical signals from muscles to derive control inputs for a device or system, has long been used in upper-limb powered prostheses [1]. Recently, there has been renewed interest in applying this technology to broader human-computer interaction (HCI) use cases such as

mixed reality [2], rehabilitation [3], and general ubiquitous computing [4, 5]. This resurgence is partially driven by academic [6] and industry breakthroughs [2] in cross-user modelling, which are beginning to show the ability to enable "plug-and-play" capabilities for discrete (i.e., dynamic) gesture recognition. These emerging cross-user models are a substantial deviation from the historical perspective that myoelectric control necessitates user-specific models. Conversely, despite an abundance of cross-subject studies [7, 8, 9], the same level of benefits have not yet been demonstrated in online continuous myoelectric control [10] due to a number of additional challenges. Furthermore, even when cross-subject models demonstrate promising results in lab experiments with non-amputees, these approaches may still face significant limitations when applied to amputees where varied residual limb anatomy poses additional challenges.

The difficulties of continuous control, compared to the aforementioned discrete systems, stem from two main factors. First, the highly non-stationary properties of the electromyography (EMG) signal [11] make generalization across users and confounding factors quite difficult when decisions must be made over shorter time windows (in the order of milliseconds) [12, 13]. Correspondingly, these controllers often need to be re-trained or additional data must be accumulated to account for variations in signals. Second, unlike discrete control, where a complete gesture is performed before an action is triggered, continuous control involves ongoing user modulation based on real-time feedback (i.e., moving a cursor across a screen or guiding a prosthesis into a specific position or grasp) [14]. Historically, training these models has required screen-guided training (SGT) protocols where users follow along with a set of prompted gestures on a screen in front of them or as guided by a prosthesis [15]. This structured follow-along approach, however, does not necessarily capture realistic user-in-the-loop patterns and behaviours, particularly when describing the short non-stationary windows used for continuous control. Consequently, even cross-user

model strategies typically rely on some degree of end-user calibration, and recent work has shown that these models benefit greatly from additional user-in-the-loop data [16, 17, 18].

Online end-user calibration techniques for continuous myoelectric control generally fall into two categories based on how they determine pseudo-labels to assign to user-in-the-loop data. The first category, *classifier-dependent calibration*, uses the classifier's own prediction outputs derived from the data as pseudo-labels for adaptation. This approach often employs some form of selective filtering so as to collect only highly-confident [19] or representative [20] samples, to avoid model deterioration due to incorrect pseudo-labels. Classifier-dependent calibration works well under known conditions, such as when adequate training data is supplied and when confounding factors like limb position effects are avoided. Conversely, it may struggle to recover when the model is perturbed, for example, if the contact between the skin and electrodes changes or if user behaviours drastically change [16]. The second category, *environment-dependent calibration*, assumes the presence of an oracle (i.e. an external source of knowledge) that can verify the suitability of actions in specific situations. Reinforcement learning exemplifies this approach, using a reward signal from the environment to influence future predictions for similar inputs [21]. Context-informed incremental learning (CIIL) operates similarly, but accelerates training by using both a reward signal and potential pseudo-labels based on assumed environmental goals (e.g. context) [16, 17]. In this situation, "context" refers to environmental cues or known user task goals that help evaluate the suitability of model predictions [22, 23]. For example, during a target acquisition task, context refers to information such as whether the user is moving toward the target or performing precise adjustments near it. A recent work has demonstrated the effectiveness of using a form of context, eye-gaze-derived and direct environmental-suitability data, to determine the optimal cursor direction in a target acquisition

task, employing a CIIL-like framework to improve 2D myoelectric regression models [24]. Other CIIL work has shown that these approaches can produce significantly higher throughput than classifier-dependent calibration, even in the case of large covariate shifts [16] or in the case of limited training data [17].

This work expands on the past environmental-dependent calibration literature to investigate whether the CIIL framework can be used to train a new model completely from scratch in a reasonable amount of time.Toward this goal, this work makes four distinct contributions:

1. Whereas most studies aim to maximize the offline classification accuracy of a model, this work shows that improving online usability can actually decrease offline accuracy due to a divergence between SGT-collected and real-world data.

2. This work expands upon past CIIL work by using neural networks, where hard pseudo-labels are no longer required. By introducing "mixed" soft targets from environmental feedback rather than one-hot targets, usability increased much faster during adaptation.

3. To counteract the noise in weakly-supervised contexts, we propose novel pseudo-label relabeling methods, including label spreading and a Gaussian Mixture Model (GMM). These techniques refine pseudo-labels by leveraging patterns in the collected data, improving the stability of adaptation even under highly ambiguous context.

4. Traditional proportional control mappings often fail to align with real-world user behavior, leading to inefficiencies in both speed and precision. For the first time, we introduce a nonlinear, adaptive sigmoid-based control mapping, which dynamically adjusts to user-specific contraction patterns. This approach

supports dual control modes (fast, coarse movements and slow, precise adjustments), directly improving throughput and target acquisition.

4.5 Methods

4.5.1 Participants

As approved by the University of New Brunswick Research Ethics Board (on file as REB 2022-122), 16 intact-limbed participants (age: 29±8, male: 11, female: 5) varying from novice to expert myoelectric users took part in this online experiment. No inducements or incentives were provided, and informed consent was obtained from all participants prior to beginning. The experiment lasted approximately 1 hour.

4.5.2 Hardware and Data Processing

The hardware used for this experiment (see Figure 4.1) was the commercially available *OYMotion gForcePro+* EMG armband[1]. This 8-channel armband samples EMG at 1000 Hz with 12-bit resolution. Moreover, this device has built-in filtering for common sources of noise (e.g., high-frequency noise and motion artifact). All data streamed from the device were split into overlapping windows of 200 ms, with a 50 ms step size. To increase the information density of each window, the wavelet energy (WENG) feature was extracted [25]. This feature was chosen as it was found in pilot studies to perform well compared to other common features. Finally, all data collection, feature extraction, and classification were done using LibEMG, an open-source Python library for creating and evaluating myoelectric interfaces [26].

[1]http://www.oymotion.com/en/product32/149

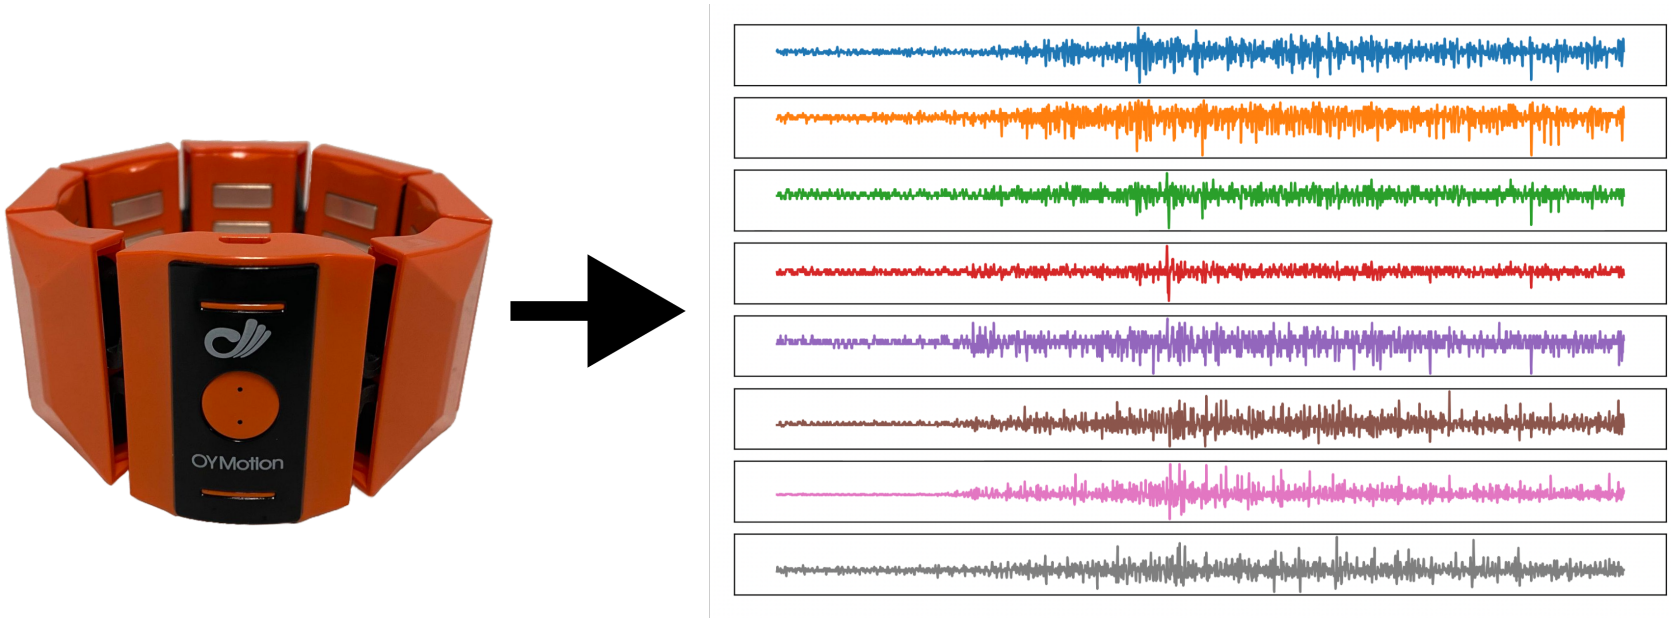


Figure 4.1: The 8-channel, 1000 Hz, *OYMotion gForcePro+* EMG armband that was used in this study and the raw EMG from a 3-second recording.

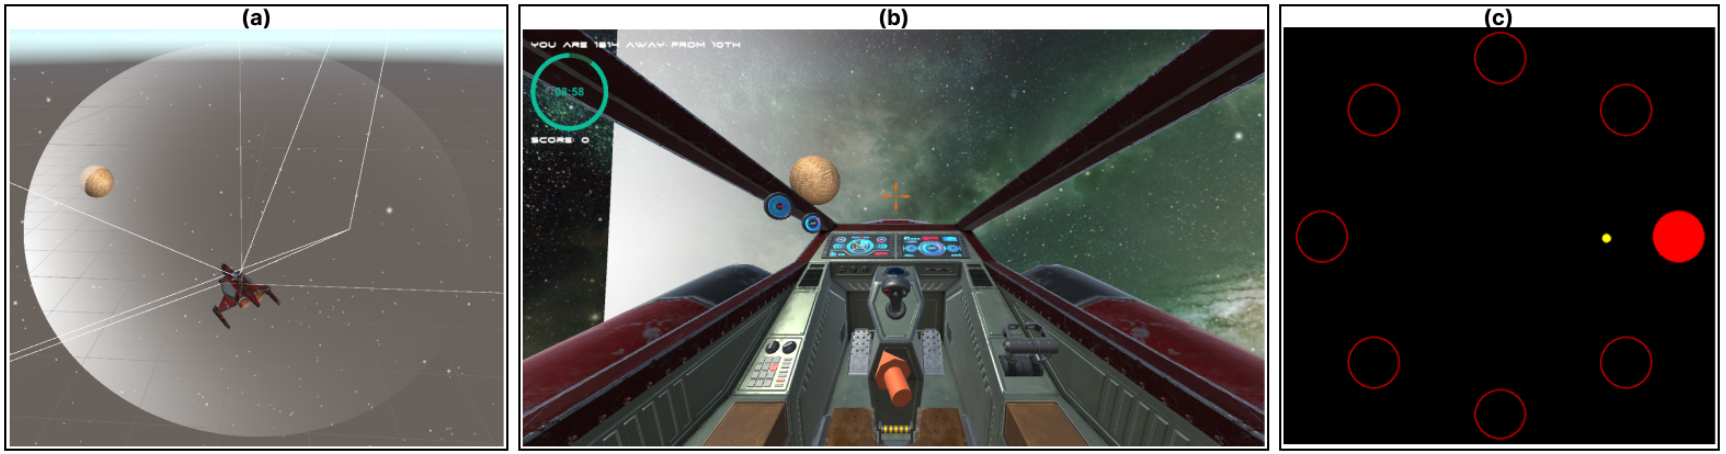


Figure 4.2: **(a)** The Unity-built scene used for the game. The white hemisphere represented the playable region which was centred around the last meteor with 170 degree range on both axes. The ship did not move forward or backward; all movements only changed the ship orientation.**(b)** The cockpit view from the Myo-Shoot 2 game. The arrow in the center indicates the direction of the planet. **(c)** The ISO 9241-9 Fitts' law testing environment.

4.5.3 Experimental Environments

4.5.3.1 Screen Guided Training

After donning the EMG armband, approximately 1/3 the length proximal to the elbow, participants underwent a screen-guided training (SGT) session, as is typical in most EMG studies. During each of the three pipelines (see Section 4.5.4), 5 repetitions of 3 seconds were acquired for each of the five control gestures: no motion, wrist extension, wrist flexion, hand open, and hand close. Participants were instructed to perform ramp contractions, whereby they were coached to gradually increase their contraction intensity to their maximum comfortable contraction level. During this collection, participants were seated in a chair, and no instructions were provided regarding the positioning of the limb so as not to confound the experiment. Most of the participants ended up moving their limb throughout the experiment.

4.5.3.2 Context Available Game Environment

An extension of the environment used in [16], Myo-Shoot 2 (see Figure 4.2A/B) was a simple game in which participants were placed inside the cockpit of a space-ship and were tasked with saving the earth from incoming meteors. The mapping from control gestures to spaceship movement was no motion (do nothing), wrist extension (rotate right), wrist flexion (rotate left), hand open (rotate up), and hand close (rotate down). To successfully destroy a meteor, participants had to navigate to the target using the control gestures and hover within its bounds (i.e., hold the crosshair inside the meteor's diameter) for 1.5 seconds. After successfully destroying the meteor, a new target was randomly generated within the game scene (see Figure 4.2A). This order of generation was kept consistent across participants and algorithms. During the game, *context suitability* could be inferred by whether the user was accomplishing

the game's goals. Positive context was when the user moved toward a target or remained still inside its bounds. In contrast, negative context was when the user moved away from the target or was within it and moved away from its center. Due to the nature of the environment and the sequential nature of classification-based control, most of the model decisions could be associated with two potential *context labels* (i.e., the components that reduce the diagonal distance to the target, such as *up* or *right* for a target placed at 2 o'clock) [16, 17].

4.5.3.3 Fitts' Law Environment

An ISO 9241-9 inspired Fitts' Law test (see Figure 4.2C) was used to evaluate the different algorithms [27]. The same control mapping defined above was used to control the cursor in this 2D environment. Like the Myo-Shoot 2 game, participants had to move the cursor over the top of the different targets and dwell within their bounds for one second. Participants acquired all eight targets twice, leading to 16 evaluation trials for each algorithm. No adaptation occurred during the evaluation, meaning the model and proportional control mappings did not change.

4.5.4 Pipelines

Throughout the online experiment, three pipelines were used by the participants: (1) a model initially trained with screen guided training data that remained constant during the user-in-the-loop phase (SGT), (2) a model initially trained with screen guided training data that underwent online adaptation via CIIL (SGT-A), and (3) a zero-shot model that started from a randomly initialized model that underwent online adaptation via CIIL (ZS-A).

All pipelines used a *neural network* topology, composed of a small multi-layer percep-

tron consisting of 4 linear layers of 32, 16, 10, and 5 neurons, respectively. Each of the layers, except the last, was proceeded by batch normalization, rectified linear unit activation function, and dropout (p=0.2). The final layer output was transformed via a softmax function, and the logits were treated as motion class probabilities.

4.5.4.1 Initial Training

An initial training was performed for two of three pipelines (SGT, SGT-A) using the screen guided training introduced in 4.5.4. The data were prepared following the processing introduced in 4.5.2. All 5 reps were used as training data, formed into batches of 64 windows. The model was optimized over 150 epochs using the Adam optimizer with a learning rate of 0.002 and a mean-squared error criterion applied to one-hot labels.

4.5.4.2 Online Adaptation

Throughout the game session, adaptation was performed for 2 of 3 pipelines (SGT-A, ZS-A). The adaptation consisted of an incremental learning approach via CIIL, with additional components added in this work for the first time; (1) relabeling of weakly-supervised pseudo-labels via label spreading, (2) separation of active classes and no motion via a Gaussian mixture model (GMM), and (3) adaptive proportional control. Pseudo-code for the online adaptation can be found in Algorithm 1.

Incremental updates to the model were continuously performed throughout the game phase using the memory buffer containing user-in-the-loop data. The incremental learning was performed using the same hyper-parameters introduced in 4.5.4.1. The memory buffer began empty, but grew throughout the 10 minute game phase. Model updates were performed in 5 epoch intervals. After each interval, the new

Algorithm 1 Online Adaptation Procedure

Define

S_t: *EMG features at time step t*
C_t: *Predicted motion class at time step t*
PC_t: *Predicted proportional control at timestep t*
$E_{s,t}$: *Environment suitability at time step t*
$E_{o,t}$: *Environment options at time step t*
$E(C_t)$: *Environment perception from motion at time t*
P_t: *Pseudo-label at time step t*
$LS(M)$: *Update pseudo-labels in memory buffer based on neighboring points using label spreading*
$GMM(M)$: *Relabel no motion samples in memory buffer using Gaussian Mixture Model*
$APC(M)$ *Update class-specific proportional control thresholds using samples in memory buffer*
$APC(S_t)$: *Get proportional control value from EMG signal*
$NN(S_t)$: *Get predicted motion class of EMG signal from Neural Network*

Initialize

Initialize memory buffer M *as empty*
Label Spreading and GMM start at 2 minutes
Adaptive Proportional Control start at 30 seconds
Initialize NN *as pretrained via SGT or random*

while *closed-loop setting is active* **do**
- **Process** *Prediction*:
  - Check for new NN $C_t \leftarrow NN(S_t)$ $PC_t \leftarrow APC(S_t)$ Act in environment
- **Process** *Memory Update*:
  - $E_{s,t}$, $E_{o,t} \leftarrow E(C_t)$ **if** $E_{s,t}$ *is positive* **then**
    - $P_t \leftarrow C_t$
  - **else**
    - $P_t \leftarrow E_{o,t}$
  - $M.append(\langle S_t, C_t, E_{s,t}, E_{o,t}, P_t \rangle)$
- **Process** *Model Update*:
  - $M.fetch()$ **for** *epoch from 1 to 5* **do**
    - $NN.update(M)$
  - **if** *Label Spreading enabled* **then**
    - $M \leftarrow LS(M)$
  - **if** *GMM enabled* **then**
    - $M \leftarrow GMM(M)$
  - **if** *Adaptive Proportional Control enabled* **then**
    - $APC \leftarrow APC(M)$
  - $NN.save()$

weights were saved and overwrote the model the user was actively using, and the memory buffer was updated with the data that had been collected in the interim. Correspondingly, these update cycles were nearly instant at the beginning of the adaptation phase, but scaled to approximately 15 seconds with 10 minutes of data.

4.5.5 Context Informed Incremental Learning

A CIIL framework was used to assign the weakly-supervised labels that arose directly from the game. Although CIIL has been employed in previous works [16, 17], a number of improvements were explored in this study. As with past versions of CIIL, the inputs to the network during the user-in-the-loop setting were stored along with the model prediction, environment suitability, and environment context labels in a memory buffer. Context suitability was operationalized based on the alignment of user actions with task goals within the game environment. Positive context suitability was defined as movements or actions that reduced the diagonal distance between the cursor and the target, or when the cursor remained stable within the target bounds. Conversely, negative context suitability was assigned when movements increased the diagonal distance from the target or deviated from its center. These criteria, derived from real-time positional feedback within the game environment, ensured consistency across participants.

When environment suitability was positive, the model prediction was taken as the pseudo-label for that input. For negative suitability, past CIIL works would have assigned the most probable environment context label according to the model. While these past CIIL works were not entirely dependent on the classifier, and therefore could improve even when the classifier was poor, using the classifier to select between environment context labels was inefficient and could lead to slower recovery. In this study, soft pseudo-labels were introduced to address this limitation.

The adoption of a neural network architecture in this work allowed for the use of soft pseudo-labels composed of both options (i.e., the diagonal) for any input associated with a negative environment suitability. Consequently, in the event of negative suitability, pseudo-labels were typically initially established with a 0.5 output logit for both possible options. This encouraged the model to increase the likelihood of

both options.

4.5.5.1 Label Spreading Relabeling

Label spreading was introduced to gradually refine the uncertainty associated with these new soft labels as greater quantities of user-in-the-loop data was collected. This process constructs a spanning tree from a graph Laplacian representing the k-nearest neighbours affinity matrix to determine regions with high class certainty. Using these regions, unsupervised data then can be relabelled based largely on their proximity to supervised data [28]. Label spreading also allows for the supervised data to be relabelled if the kernel determines its label is overwhelmingly unlikely. The relabeling regularization factor for supervised data was set to 0.2.

Pseudo-labels arising from positive context suitability were treated as the "supervised" data, given it was quite certain the label was correct, and these labels were propagated onto the negative context data to gradually change them from equal weight distribution (e.g., 0.5 for both options) towards the most representative option. Since the efficacy of label spreading depends on well-represented classes in the spanning tree, it was beneficial to wait until the memory buffer had accumulated a sufficiently large number of samples before beginning label spreading. Pilot testing indicated that 2 minutes of gameplay provided adequate data diversity, so label spreading began at this time.

4.5.5.2 Active/Non-Active Separation via a GMM

To address the class imbalance inherent in the data collection process, a GMM was added to the adaptation procedure. This approach was motivated by the disproportionate representation of active motions in the overall distribution of classes. During

gameplay, users predominantly engaged in maneuvering the spaceship towards meteors, with comparatively little time spent in stationary positions. Consequently, the "no motion" class was at risk of being underrepresented or even disappearing from the model due to label spreading, despite its critical importance in achieving the task. Although users could still achieve stationary positions through proportional control, the potential loss of a distinct "no motion" class might have posed significant challenges to the system's responsiveness and accuracy.

The GMM addressed this issue by separating the proportional control inputs into two clusters, corresponding to active and no-motion classes. The lower-intensity cluster was periodically relabeled as no motion, ensuring that this class retained a robust representation throughout adaptation. Figure 4.3 illustrates the separability of no-motion data from active classes based on the proportional control. The GMM process was initiated after 2 minutes of gameplay to prevent early-stage overfitting from sparse data, complementing the label spreading procedure.

Together, these methods—initial context suitability gathering, label spreading, and GMM-based relabeling—were integrated into the CIIL framework to dynamically refine pseudo-labels and maintain robust class boundaries during adaptation. This iterative process ensured the system could adapt effectively to user behavior while mitigating the impact of noise or data imbalance.

4.5.5.3 Adaptive Proportional Control

As is typical for most continuous myoelectric control systems, proportional control was used to enable participants to control the speed of the cursor (i.e., lighter and harder contractions moved the cursor more slowly or quickly, respectively) [29, 30, 31]. Generally, this strategy can significantly improve the throughput of a myoelectric interface as it can (1) minimize the effect of misclassifications on lower

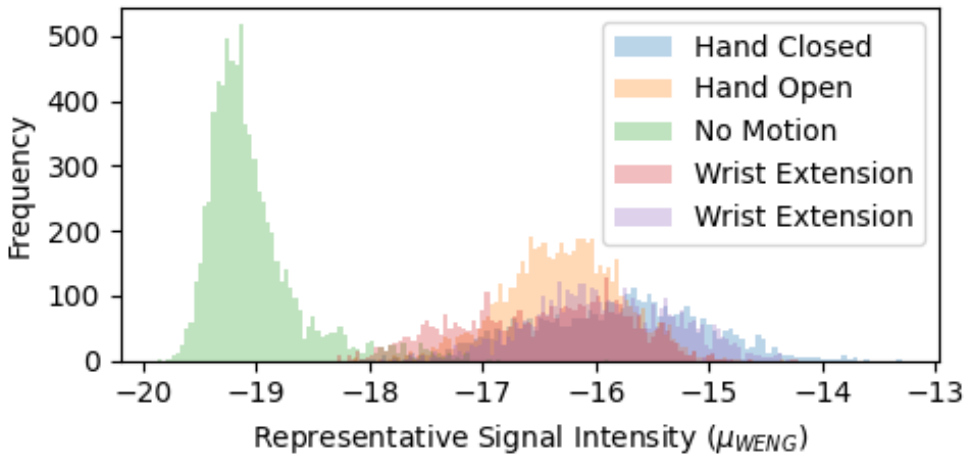


Figure 4.3: Histogram of the proportional control inputs from a pilot gameplay session colored by predicted motion class.

amplitude contractions (when errors are most likely) and (2) enable users to adapt their speed to more accurately achieve a desired target position. This mapping, determined based on previously recorded data (such as that acquired during SGT), has historically remained fixed as the user interacts with the system. Moreover, as described in [31], the transfer function, mapping the contraction intensity with the speed of the cursor, has typically been linear.

During pilot studies, two issues arose when relying on traditional linear proportional control mappings. First, as highlighted above, offline data does not represent online behaviours well. Correspondingly, the ramp contractions acquired during SGT, which are ultimately used to define the proportional control mappings for each class, are unlikely to be a good representation of the user's desired speed over the cursor. Because myoelectric control performance is coupled so tightly to proportional control [32], a static proportional control mapping could theoretically limit even the best adaptive classifier. Second, users found it difficult to control the speed precisely when combined with pattern-recognition based control. To enable fine control at slow speeds (needed to accurately acquire targets), this necessitated a lower overall

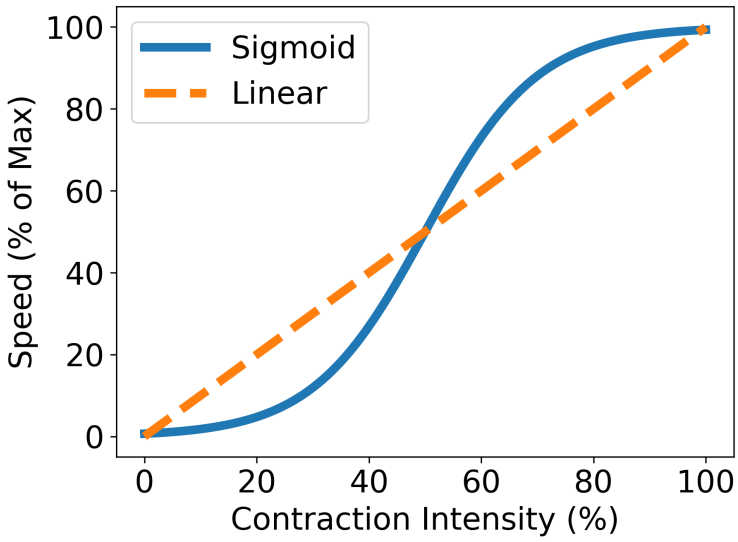


Figure 4.4: The mapping between the normalized contraction intensity and cursor speed (represented as a % of the maximum gain). Both the traditional linear mapping and the proposed sigmoid mapping are shown.

gain on the cursor, limiting the potential throughput.

Similarly, research in human-computer interaction has shown that for other pointing devices, like computer mice, linear transfer functions (i.e., constant gains) are not always ideal [33]. As highlighted by Casiez et al., this is because users have a tendency in pointing tasks to operate in two modes: (1) fast and imprecise (i.e., quickly moving to the approximate location of a target) and (2) slow and precise, (i.e., once close to the target precise movements are made to orient the cursor on top of the target) [34]. Such behaviours, which were noticeable during early pilots from this work, suggested the potential of adopting a sigmoid-based transfer functions.

Consequently, in this work, we overcome some of the limitations of previous proportional control mappings by using an adaptive sigmoid-based transfer function (see Equation 4.1). As determined through pilot studies, all data were first thresholded to find the 10% and 90% percentiles of the maximum WENG value (averaged across channels) for each class. These thresholds were then continuously adapted using

the new data acquired within the game environment, with the old thresholds updated by a factor of 0.25 for each adaptation cycle. The WENG for each window was then converted to a linear range between the lower and upper bounds, scaling them between 0 and 1 before passing through the sigmoid transfer function shown in Equation 4.1. The "squashing factor" in Equation 4.1 was set to 10, determined through pilot studies with experienced myoelectric control users. Adaptive proportional control was enabled 30 seconds into the gameplay phase. This was chosen to disentangle the effect of adaptive proportional control from CIIL, label spreading, and GMM procedures procedures.

$$\%\ \text{Max Speed}(\mu_{WENG}) = \frac{1}{1 + e^{-10(\mu_{WENG} - 5)}} \tag{4.1}$$

where μ_{WENG} is the normalized contraction intensity.

4.5.6 Evaluation Criteria

While many metrics exist to evaluate myoelectric interfaces, this work focuses primarily on the throughput (see Equations 4.2 and 4.3) enabled by the system. In this case, the throughput (measured in bits/s) describes the amount of control information that can flow through the system (i.e., how quickly users can acquire targets) [27]. Compared to offline metrics like classification accuracy and active error, which may not fully predict usability [35], throughput directly measures how well the user can use the control system to accomplish a specific task. As such, the goal of any adaptation should not be to improve classification accuracy but, instead, to improve the throughput.

$$ID = \log_2 (\frac{D}{W} + 1) \tag{4.2}$$

where ID is the index of difficulty, D is the distance between targets, and W is the width of the target.

$$\text{Throughput} = \frac{ID}{MT} \tag{4.3}$$

where MT is the time taken to acquire the target subtracted by the dwell time.

4.6 Results & Discussion

This section explores the results of this work, including an analysis of the offline and online performance of the adaptive and non-adaptive control schemes. Also highlighted is an evaluation of the adaptive sigmoid-based proportional control mapping. Statistical differences between conditions were computed using the Statistical Tests for Algorithms Comparison (STAC) [36]. For all comparisons, Friedman tests with a Finner correction were used.

4.6.1 Offline Metrics Only Tell a Small Part of the Story

As highlighted in Figure 4.5A/D, the average offline accuracy of the different approaches were: SGT (92.4% $\pm$ 5.2%), SGT-A-Pre (92.3% $\pm$ 3.2%), SGT-A-Post (85.8% $\pm$ 5.1%), and ZS-A-Post (83.1% $\pm$ 8.6%). As SGT and SGT-A-pre followed the same calibration procedure, the slight difference in accuracy observed were due to the model being trained with a different set of five repetitions acquired before each specific trial. Without additional insights, these results might suggest that the models without adaptation would yield the best performance if deployed in a real-world setting, significantly outperforming the accuracies of SGT-A-post and ZS-A ($p < 0.05$). These results, if taken out of context, may even appear to suggest that

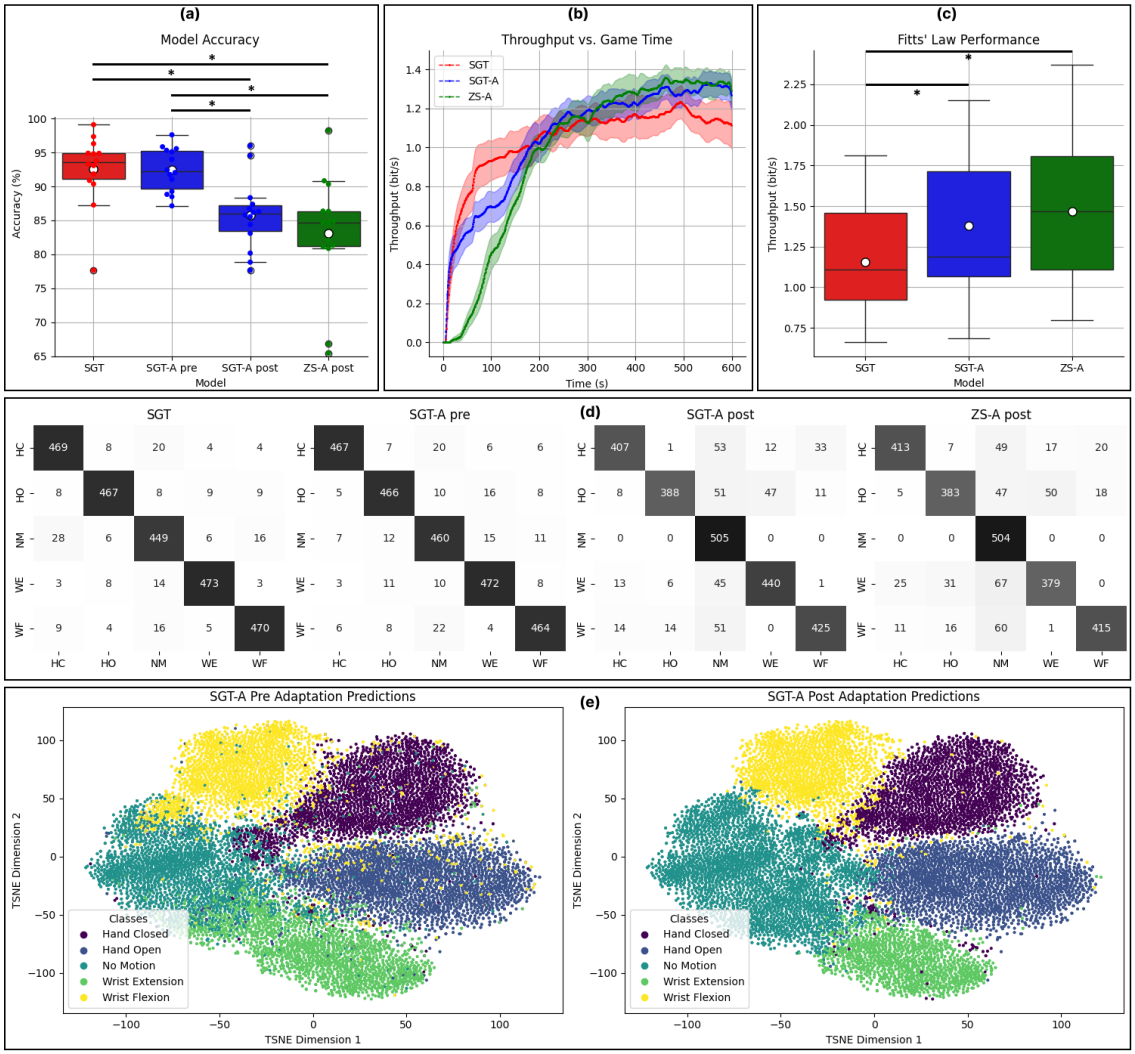


Figure 4.5: Classifier dependent results. **(a)** Classification accuracy across different models. **(b)** Throughput progression during gameplay for SGT, SGT-A, and ZS-A models. **(c)** Fitts' Law performance comparison between models. **(d)** Confusion matrices for SGT, SGT-A (pre and post), and ZS-A (post) models. **(e)** t-SNE visualization of SGT-A model predictions before (left) and after (right) adaptation on the gameplay data.

adaptive schemes might harm performance due to an average drop in accuracy of 6.5% after adaptation; however, the online throughput actually increased as adaptation took place. This contradicts some of the myoelectric control adaptation research suggesting that offline accuracy is a reasonable predictor of online performance [37], and other works that typically show improvements in classification accuracy after adaptation, even in fully unsupervised situations [38].

When looking beyond average classification accuracy alone, the confusion matrices shown in Figure 4.5D start to tell more of the story. In particular, it becomes clear that many misclassifications in SGT-A and ZS-A are between active classes and no motion. This pattern was likely caused by the relabeling of lower amplitude contractions during adaptation (where misclassifications are common) to the no motion class, acting as a sort of rejection. Such rejection can be beneficial as no onset detection was performed during screen guided training [39]. This relabeling ensured that uncertain or ambiguous inputs—common during transient behaviors—did not propagate as incorrect active class labels, which could destabilize control.

Figure 4.5E further highlights this phenomenon, showing that the no motion class effectively expanded into regions of active class space after adaptation. This was driven by the interplay of the GMM, which reclassified low-intensity contractions as no motion, and label spreading, which refined boundaries for the active classes. This dynamic adjustment process was critical for achieving stable and usable active class boundaries. Unlike screen-guided training, which operates on controlled and predictable data, the adaptation processes in SGT-A and ZS-A leveraged transient and noisy real-world inputs to develop a more robust class structure reflective of gameplay behavior.

Ultimately, the structural rejection of low-amplitude active class windows to the no motion class inevitably results in some misclassifications of active classes as no

motion. However, in systems that use proportional control within a user-in-the-loop setting, these misclassifications would have a negligible impact on throughput. This is because low-amplitude contractions contribute minimally to the overall movement of the system, making their misclassification less critical to task performance. While such rejection might appear detrimental in offline accuracy metrics, it ensures a cleaner decision boundary between movement and rest states, ultimately enhancing system stability. These findings highlight that for systems with proportional control, prioritizing robust no motion classification over perfect active class accuracy is a valid trade-off that supports smoother transitions and improved real-time usability.

4.6.2 Online Performance Tells the Full Story

First shown by Hargrove et al. in 2007, there is often a mismatch between classification accuracy and online performance for myoelectric control [40]. The online results corroborate this idea (see Figure 4.5B/C) as they tend to tell a much different story than the offline metrics highlighted above, with the ZS-A-Post algorithm having the highest throughput (1.47 ± 0.46), significantly outperforming ($p < 0.05$) the baseline SGT model (1.15 ± 0.37). Similarly, the SGT-A-Post algorithm significantly outperformed ($p < 0.05$) the baseline SGT model, having an average throughput of (1.38 ± 0.45). Although not significant, the ZS-A-Post algorithm outperformed the SGT-A-Post model by an average of 0.09 bits/s, highlighting a potential benefit of starting from scratch (perhaps reducing the risk of model biases learned from the SGT data). Also interesting to note is the evolution of the model throughputs during gameplay, as shown in Figure 4.5B. The SGT model generally started with the highest throughput but quickly plateaued (around 200s) as users grew accustomed to the controller. The SGT-A model also started with relatively high throughput, but due to its ability to learn and account for real-world user behaviours, it increased

throughout the entire trial and arguably still had not plateaued even after 600s of adaptation. In contrast, the ZS-A algorithm started with an extremely low throughput (due to its randomized initialization) but within 200s (around the same time as it took to conduct the original SGT session), it converged to a similar usability of the other two models. Moreover, after converging, the ZS-A algorithm continued to improve throughout the session, exceeding the performance of SGT-A at around 300s.

4.6.3 Fast or Slow With Very Little In-Between

A perhaps equally important contributor to the improved throughput of the adaptive pipelines was the incorporation of adaptive proportional control. Figure 4.6 highlights the dynamics of the model's proportional control thresholds and user behaviour over the course of the game.

At the beginning, the ZS-A model started with unusable classifier boundaries and proportional control thresholds corresponding to the dynamic range of the device. However, when the adaptive proportional control was enabled at 30 seconds, the thresholds of the active classes were rapidly learned. After 120 seconds, when the GMM had been activated, the classification boundaries of the no-motion class were greatly improved, and the proportional control thresholds narrowed for no-motion, thus better describing the noise floor. After 200 seconds, the proportional control boundaries did not change meaningfully; yet, the performance of ZS-A had only reached around 70% of its terminal value at this time (Figure 4.5B), indicating that user learning and classifier boundary refinement continued to yield benefits even after the adaptive proportional control had settled. Despite the change in proportional control thresholds, in Figure 4.6D the user was not required to produce harder or softer contractions throughout this process.

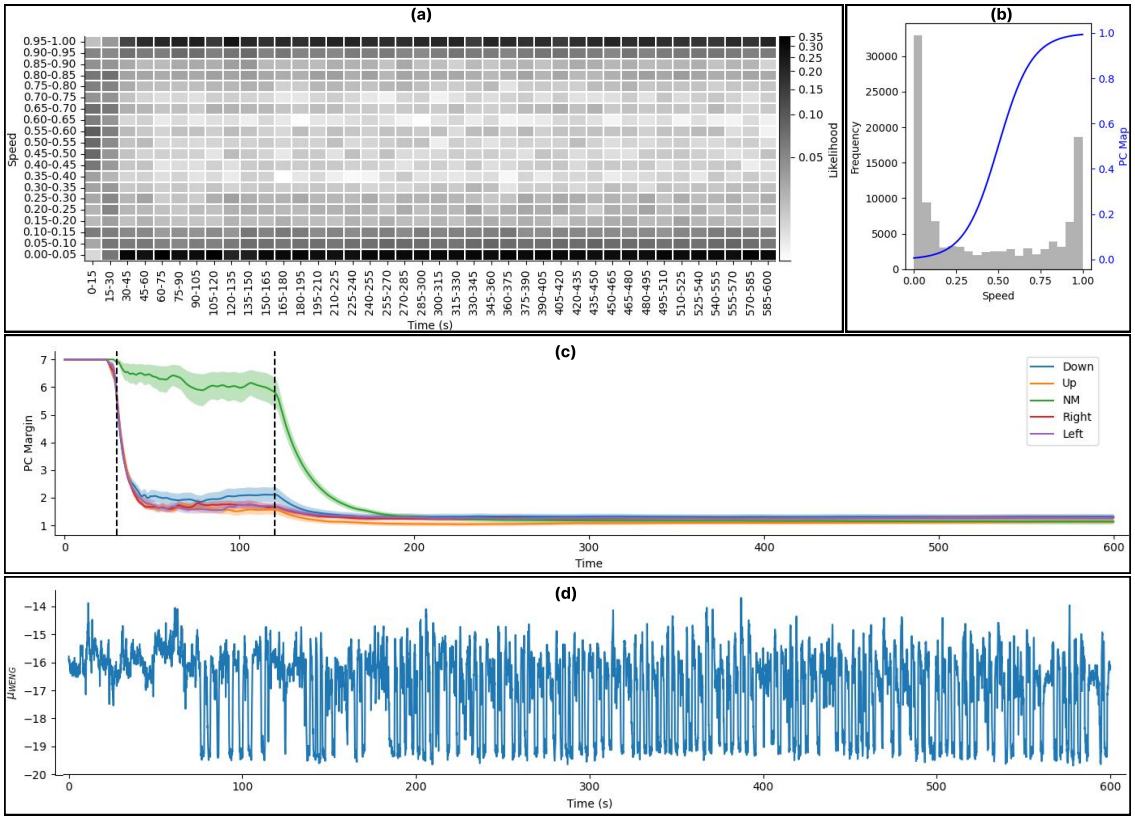


Figure 4.6: Analysis of proportional control in the ZS-A model. (a) Heatmap showing the distribution of proportional control speeds over 15-second game intervals. (b) Overall speed distribution throughout the game, with frequency histogram. (c) Class-specific proportional control threshold margins over time. Vertical dashed lines indicate activation of adaptive proportional control and GMM/label spreading at 30 and 120 seconds, respectively. (d) Time series of proportional control input for a single subject over the course of the game.

This study introduced a non-linear sigmoid mapping for proportional control, which aligns with the observed user behavior. As shown in Figure 5B, users predominantly employed two distinct control modes: fast but imprecise movements, and slow but precise adjustments - consistent with previous findings [34]. Notably, the middle range of speeds was mainly observed at the beginning of gameplay, before the proportional control boundaries were optimized. These two control modes directly correspond to the dual objectives in Fitts' Law: reaching the target area quickly (fast and imprecise) and then accurately positioning within it (slow and precise) [41]. This behavior mirrors natural limb movements, which, despite having smooth velocity profiles to minimize jerk, typically consist of a brief initial acceleration followed by a longer deceleration phase [42]. To further optimize this behavior, we propose future work should explore increasing the sigmoid function's "squashing factor" (currently set at 10 in Equation 4.1) to more distinctly separate these two control modes. This adjustment could particularly benefit control classes with smaller dynamic ranges, such as wrist pronation/supination. Further, we believe that, akin to the large body of transfer function work done for pointing tasks in human-computer interaction [33], a deeper exploration of the mappings between contraction intensity and speed in the context of pattern recognition-based myoelectric control could yield even further benefit.

4.6.4 Future Directions

This study improved on previous CIIL work [16] by using a multilayered perceptron instead of a linear discriminant analysis classifier. This enabled more efficient data use and future work should continue to explore more sophisticated architectures. For example, EMG is a highly dynamic time series that exhibits non-stationary properties. Traditionally, researchers treated neighboring windows as independent;

however, many recent studies have demonstrated benefits from models such as recurrent neural networks [43], long short term memory [44], self-attention [45], or transformers [46]. In these studies, the transitions between classes occupy unique regions in latent space [47], which was not the case when visualizing our feature space based on independent windows (Figure 4.5E). Combining this more descriptive latent space alongside CIIL could provide combined benefits.

The label spreading strategy in this work effectively consolidated ambiguous contexts over time; however, more advanced approaches could further improve performance. Transductive propagation networks, for example, dynamically construct task-specific graph structures in a learnable, end-to-end framework, enabling efficient label propagation without reliance on heuristics or pre-defined similarity metrics [48]. Additionally, while label spreading denoised pseudo-labels effectively, its reliance on an ever-growing memory bank is impractical for untethered embedded systems—where myoelectric control is deployed in practice. Potential solutions include decomposing data into relevant and obsolete components, as proposed by Yeung [49] which offers a potential solution for maintaining a fixed-size memory bank, as does Yamanoi's "survival of the fittest" approach [50], which retains only relevant data over a specified horizon in incremental myoelectric learning Implementing these methods could make CIIL more scalable and suitable for real-world deployment.

In this work, the adaptive pipelines yielded a higher throughput than has typically been observed in past studies. Although some previous works have achieved similar throughputs, like 1.37 bit/s in a 1D Fitts' Law environment [2], throughput in comparable higher dimensional tasks has typically been reported in the 0.4-0.75 bit/s range [51, 52, 53], including our past CIIL studies [16]. This increase arose in part from the improved model adaptation enabled by the novel CIIL implementation, but also from the change to a sigmoid transfer function, which allowed the overall gain

to be increased while maintaining a controllable low speed control mode. In future work, such as in the simultaneous activation of degrees of freedom through regression, this concept should be likewise considered to reduce drift at low amplitudes (quiescence [54]) and unintentional co-activation (action interference [55]).

Furthermore, the proposed weakly-supervised incremental learning framework was demonstrated here in a game environment as a training tool. The noted improvement of the SGT-A and ZS-A approaches suggests that this is worth continued exploration, and future work should explore the subsequent performance of CIIL-trained models in real-world use cases of myoelectric control (such as virtual reality, robotic teleoperation, or prosthesis control). The greater dexterity requirements of these challenging environments poses a unique opportunity to attribute regions of feature space through context with the richness that can be described by an oracle. Further, for virtual reality (VR), the immersive nature of the environment could benefit from integrating CIIL with multi-modal sensory inputs, such as visual or haptic feedback, to notify the user of positive or negative context suitability and improve transparency of the adaptation. Future work could address these by incorporating unsupervised domain adaptation techniques or transfer learning pipelines, enabling consistent performance with minimal recalibration.

Finally, CIIL could also be used to maintain and improve performance beyond initial training in general ubiquitous computing environments where context can be measured or inferred. Such situational context has been studied in shared control systems, where a vision-based object recognition system determines the motion class based on the nearest object and EMG is used only for onset detection [56, 57]. Although these systems boast high performance thanks to the use of state-of-the-art vision models, they significantly reduce the sense of user agency [58]. Alternatively, the vision labels could be used to inform CIIL, leaving the user in control. Future

work should focus on deriving situational context from environmental perception, providing a more convincing arena for CIIL to be viewed as a persistent incremental learning strategy beyond training and gamified applications. This direction will enable a broader range of real-world scenarios in which CIIL-informed myoelectric control could be deployed.

Bibliography

[1] E. Scheme and K. Englehart, "Electromyogram pattern recognition for control of powered upper-limb prostheses: state of the art and challenges for clinical use.," *Journal of Rehabilitation Research & Development*, vol. 48, no. 6, 2011.

[2] C. Labs, D. Sussillo, P. Kaifosh, and T. Reardon, "A generic noninvasive neuromotor interface for human-computer interaction," *bioRxiv*, 2024.

[3] D. J. Berger and A. d'Avella, "Myoelectric control and virtual reality to enhance motor rehabilitation after stroke," *Frontiers in Bioengineering and Biotechnology*, vol. 12, p. 1376000, 2024.

[4] M. Zheng, M. S. Crouch, and M. S. Eggleston, "Surface electromyography as a natural human–machine interface: a review," *IEEE Sensors Journal*, vol. 22, no. 10, pp. 9198–9214, 2022.

[5] E. Eddy, E. J. Scheme, and S. Bateman, "A framework and call to action for the future development of emg-based input in hci," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–23, 2023.

[6] E. Eddy, E. Campbell, S. Bateman, and E. Scheme, "Big data in myoelectric control: large multi-user models enable robust zero-shot emg-based discrete

gesture recognition," *Frontiers in Bioengineering and Biotechnology*, vol. 12, p. 1463377, 2024.

[7] X. Jiang, C. Ma, and K. Nazarpour, "One-shot random forest model calibration for hand gesture decoding," *Journal of Neural Engineering*, vol. 21, no. 1, p. 016006, 2024.

[8] E. Campbell, A. Phinyomark, and E. Scheme, "Deep cross-user models reduce the training burden in myoelectric control," *Frontiers in Neuroscience*, vol. 15, p. 657958, 2021.

[9] U. Côté-Allard, C. L. Fall, A. Campeau-Lecours, C. Gosselin, F. Laviolette, and B. Gosselin, "Transfer learning for semg hand gestures recognition using convolutional neural networks," in *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 1663–1668, 2017.

[10] K. Englehart and B. Hudgins, "A robust, real-time control scheme for multifunction myoelectric control," *IEEE Transactions on Biomedical Engineering*, vol. 50, no. 7, pp. 848–854, 2003.

[11] E. Shwedyk, R. Balasubramanian, and R. Scott, "A nonstationary model for the electromyogram," *IEEE Transactions on Biomedical Engineering*, no. 5, pp. 417–424, 1977.

[12] E. Campbell, A. Phinyomark, and E. Scheme, "Current trends and confounding factors in myoelectric control: Limb position and contraction intensity," *Sensors*, vol. 20, no. 6, 2020.

[13] E. Eddy, E. Campbell, S. Bateman, and E. Scheme, "Understanding the influence of confounding factors in myoelectric control for discrete gesture recognition," *Journal of Neural Engineering*, vol. 21, no. 3, p. 036015, 2024.

[14] C. Igual, A. Castillo, and J. Igual, "An interactive training model for myoelectric regression control based on human–machine cooperative performance," *Computers*, vol. 13, no. 1, p. 29, 2024.

[15] B. Lock, A. M. Simon, K. Stubblefield, and L. J. Hargrove, "Prosthesis-guided training for practical use of pattern recognition control of prostheses," in *Proceedings of the Myoelectric Controls Symposium (MEC)*, pp. 61–64, 2011.

[16] E. Eddy, E. Campbell, S. Bateman, and E. Scheme, "Leveraging task-specific context to improve unsupervised adaptation for myoelectric control," in *2023 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 4661–4666, 2023.

[17] E. Campbell, E. Eddy, S. Bateman, U. Côté-Allard, and E. Scheme, "Context-informed incremental learning improves both the performance and resilience of myoelectric control," *Journal of NeuroEngineering and Rehabilitation*, vol. 21, p. 70, May 2024.

[18] R. B. Woodward and L. J. Hargrove, "Adapting myoelectric control in real-time using a virtual environment," *Journal of neuroengineering and rehabilitation*, vol. 16, pp. 1–12, 2019.

[19] J. W. Sensinger, B. A. Lock, and T. A. Kuiken, "Adaptive pattern recognition of myoelectric signals: Exploration of conceptual framework and practical algorithms," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 17, no. 3, pp. 270–278, 2009.

[20] K. Szymaniak, A. Krasoulis, and K. Nazarpour, "Recalibration of myoelectric control with active learning," *Frontiers in Neurorobotics*, vol. 16, p. 1061201, 2022.

[21] P. M. Pilarski, M. R. Dawson, T. Degris, F. Fahimi, J. P. Carey, and R. S. Sutton, "Online human training of a myoelectric prosthesis controller via actor-critic reinforcement learning," in *2011 IEEE international conference on rehabilitation robotics*, pp. 1–7, IEEE, 2011.

[22] G. D. Abowd, A. K. Dey, P. J. Brown, N. Davies, M. Smith, and P. Steggles, "Towards a better understanding of context and context-awareness," in *Handheld and Ubiquitous Computing: First International Symposium, HUC'99 Karlsruhe, Germany, September 27–29, 1999 Proceedings 1*, pp. 304–307, Springer, 1999.

[23] J.-y. Hong, E.-h. Suh, and S.-J. Kim, "Context-aware systems: A literature review and classification," *Expert Systems with applications*, vol. 36, no. 4, pp. 8509–8522, 2009.

[24] A. H. Chou, M. Madduri, S. J. Li, J. Isa, A. Christensen, F. L. Hutchison, S. A. Burden, and A. L. Orsborn, "Using eye gaze to train an adaptive myoelectric interface," *bioRxiv*, pp. 2024–04, 2024.

[25] R. N. Khushaba, S. Kodagoda, S. Lal, and G. Dissanayake, "Driver drowsiness classification using fuzzy wavelet-packet-based feature-extraction algorithm," *IEEE Transactions on Biomedical Engineering*, vol. 58, no. 1, pp. 121–131, 2011.

[26] E. Eddy, E. Campbell, A. Phinyomark, S. Bateman, and E. Scheme, "Libemg: An open source library to facilitate the exploration of myoelectric control," *IEEE Access*, vol. 11, pp. 87380–87397, 2023.

[27] I. S. MacKenzie, "Fitts' law," *The wiley handbook of human computer interaction*, vol. 1, pp. 347–370, 2018.

[28] X. Zhu and Z. Ghahramani, "Learning from labeled and unlabeled data with label propagation," *ProQuest number: information to all users*, 2002.

[29] A. Fougner, Ø. Stavdahl, P. J. Kyberd, Y. G. Losier, and P. A. Parker, "Control of upper limb prostheses: Terminology and proportional myoelectric control—a review," *IEEE Transactions on neural systems and rehabilitation engineering*, vol. 20, no. 5, pp. 663–677, 2012.

[30] A. E. Olsson, N. Malešević, A. Björkman, and C. Antfolk, "Learning regularized representations of categorically labelled surface emg enables simultaneous and proportional myoelectric control," *Journal of NeuroEngineering and Rehabilitation*, vol. 18, pp. 1–19, 2021.

[31] E. Scheme, B. Lock, L. Hargrove, W. Hill, U. Kuruganti, and K. Englehart, "Motion normalized proportional control for improved pattern recognition-based myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 22, no. 1, pp. 149–157, 2014.

[32] R. Meredith, E. Eddy, S. Bateman, and E. Scheme, "Comparing online wrist and forearm emg-based control using a rhythm game-inspired evaluation environment," *Journal of Neural Engineering*, 2024.

[33] G. Casiez and N. Roussel, "No more bricolage! methods and tools to characterize, replicate and compare pointing transfer functions," in *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, UIST '11, (New York, NY, USA), p. 603–614, Association for Computing Machinery, 2011.

[34] G. Casiez, N. Roussel, and D. Vogel, "1 € filter: a simple speed-based low-pass filter for noisy input in interactive systems," in *Proceedings of the SIGCHI*

Conference on Human Factors in Computing Systems, CHI '12, (New York, NY, USA), p. 2527–2530, Association for Computing Machinery, 2012.

[35] J. L. Nawfel, K. B. Englehart, and E. J. Scheme, "A multi-variate approach to predicting myoelectric control usability," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 29, pp. 1312–1327, 2021.

[36] I. Rodríguez-Fdez, A. Canosa, M. Mucientes, and A. Bugarín, "STAC: a web platform for the comparison of algorithms using statistical tests," in *Proceedings of the 2015 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, 2015.

[37] R. M. Hinson, J. Berman, W. Filer, D. Kamper, X. Hu, and H. Huang, "Offline evaluation matters: investigation of the influence of offline performance on real-time operation of electromyography-based neural-machine interfaces," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 31, pp. 680–689, 2022.

[38] J. W. Sensinger, B. A. Lock, and T. A. Kuiken, "Adaptive pattern recognition of myoelectric signals: exploration of conceptual framework and practical algorithms," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 17, no. 3, pp. 270–278, 2009.

[39] S. Solnik, P. DeVita, P. Rider, B. Long, and T. Hortobágyi, "Teager–kaiser operator improves the accuracy of emg onset detection independent of signal-to-noise ratio," *Acta of bioengineering and biomechanics/Wroclaw University of Technology*, vol. 10, no. 2, p. 65, 2008.

[40] L. Hargrove, Y. Losier, B. Lock, K. Englehart, and B. Hudgins, "A real-time pattern recognition based myoelectric control usability study implemented in

a virtual environment," in *2007 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 4842–4845, 2007.

[41] R. Plamondon and A. M. Alimi, "Speed/accuracy trade-offs in target-directed movements," *Behavioral and brain sciences*, vol. 20, no. 2, pp. 279–303, 1997.

[42] E. V. Todrov, *Studies of goal directed movements*. PhD thesis, Massachusetts Institute of Technology, 1998.

[43] P. Xia, J. Hu, and Y. Peng, "Emg-based estimation of limb movement using deep learning with recurrent convolutional neural networks," *Artificial organs*, vol. 42, no. 5, pp. E67–E77, 2018.

[44] M. Jabbari, R. N. Khushaba, and K. Nazarpour, "Emg-based hand gesture classification with long short-term memory deep recurrent neural networks," in *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pp. 3302–3305, IEEE, 2020.

[45] M. Jabbari, R. N. Khushaba, and K. Nazarpour, "Combined dynamic time warping and spatiotemporal attention for myoelectric control," in *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pp. 5940–5943, IEEE, 2021.

[46] M. Montazerin, E. Rahimian, F. Naderkhani, S. F. Atashzar, S. Yanushkevich, and A. Mohammadi, "Transformer-based hand gesture recognition from instantaneous to fused neural decomposition of high-density emg signals," *Scientific reports*, vol. 13, no. 1, p. 11000, 2023.

[47] S. T. P. Raghu, D. MacIsaac, and E. Scheme, "Enabling myoelectric control training using continuous data through self-supervised representation learning," in *Myoelectric Controls Symposium*, 2024.

[48] Y. Liu, J. Lee, M. Park, S. Kim, E. Yang, S. J. Hwang, and Y. Yang, "Learning to propagate labels: Transductive propagation network for few-shot learning," *arXiv preprint arXiv:1805.10002*, 2018.

[49] D. Yeung, D. Farina, and I. Vujaklija, "Directional forgetting for stable co-adaptation in myoelectric control," *Sensors*, vol. 19, no. 9, p. 2203, 2019.

[50] Y. Yamanoi, S. Togo, Y. Jiang, and H. Yokoi, "Learning data correction for myoelectric hand based on "survival of the fittest"," *Cyborg and Bionic Systems*, 2021.

[51] E. N. Kamavuako, E. J. Scheme, and K. B. Englehart, "On the usability of intramuscular emg for prosthetic control: A fitts' law approach," *Journal of Electromyography and Kinesiology*, vol. 24, no. 5, pp. 770–777, 2014.

[52] A. Belyea, K. Englehart, and E. Scheme, "Fmg versus emg: A comparison of usability for real-time pattern recognition based control," *IEEE Transactions on Biomedical Engineering*, vol. 66, no. 11, pp. 3098–3104, 2019.

[53] A. Waris, I. Mendez, K. Englehart, W. Jensen, and E. N. Kamavuako, "On the robustness of real-time myoelectric control investigations: A multiday fitts' law approach," *Journal of neural engineering*, vol. 16, no. 2, p. 026003, 2019.

[54] C. Morrell, E. Campbell, and E. Scheme, "Exploring user compliance in the training of regression-based myoelectric control," in *Myoelectric Controls Symposium*, 2024.

[55] M. Nowak, T. Eiband, E. R. Ramírez, and C. Castellini, "Action interference in simultaneous and proportional myocontrol: comparing force-and electromyography," *Journal of neural engineering*, vol. 17, no. 2, p. 026011, 2020.

[56] J. DeGol, A. Akhtar, B. Manja, and T. Bretl, "Automatic grasp selection using a camera in a hand prosthesis," in *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 431–434, IEEE, 2016.

[57] D. Xiong, X. Fu, D. Zhang, Y. Chu, Y. Zhao, and X. Zhao, "Robotic telemanipulation with emg-driven strategy-assisted shared control method," *Science China Technological Sciences*, 09 2024.

[58] J. W. Sensinger and S. Dosen, "A review of sensory feedback in upper-limb prostheses from the perspective of human motor control," *Frontiers in neuroscience*, vol. 14, p. 345, 2020.

Chapter 5

Context-Informed Incremental Learning Improves Throughput and Reduces Drift in Regression-Based Myoelectric Control

Christian Morrell, **Evan Campbell**, Ethan Eddy, Erik Scheme

5.1 Foreword

This chapter presents the translation of context-informed incremental learning from the classification domain to the regression domain. I led the formulation and implementation of regression-based CIIL, and shared in the initial piloting, before my role transitioned to a supervisory capacity for the undertaking of the experiment

and the remainder of a collaborator's MSc degree.

5.2 Reference

This work was peer-reviewed and published under ©2025 IEEE. Reprinted, with permission, from:

Morrell, C., Campbell, E., Eddy, E., Scheme, E., *Context-Informed Incremental Learning Improves Throughput and Reduces Drift in Regression-Based Myoelectric Control,* IEEE Transactions on Neural Systems and Rehabilitation Engineering, Vol 33, No 1, May 2025, pp 1841, DOI: 10.1109/TNSRE.2025.3567245

5.3 Abstract

Despite decades of research, commercially available powered myoelectric prostheses continue to use sequential, classification-based control. While regression-based approaches can improve the dexterity offered through simultaneous, independent, and proportional control, current training protocols lack consistency across studies and fail to capture realistic user behaviours, resulting in robustness issues. To address these challenges, this work employs context-informed incremental learning (CIIL) in an unconstrained, velocity-based environment for regression-based myoelectric control. Two new adaptive models, one inspired by previous works (O-CIIL) and one modified to factor in user compliance and behaviours (T-CIIL), were compared with two models trained using traditional screen-guided training. Sixteen participants completed an online Fitts' Law target acquisition task. Both adaptive approaches significantly outperformed ($p < 0.05$) the non-adaptive models across a variety of metrics. Additionally, T-CIIL outperformed O-CIIL in alleviating drift and action interference, key issues that have plagued existing regression-based myoelectric control systems. These findings are supported by two novel metrics, namely *action interference* and *simultaneity gain*, which show that adding simultaneity often increases instability in the form of undesired and uncontrollable simultaneous motions. These findings demonstrate the viability of CIIL in an unconstrained, velocity-controlled environment for regression-based myoelectric control, and highlight the importance of capturing user behaviours when training regression-based myoelectric control systems. Source code is available on GitHub.

5.4 Introduction

Myoelectric control uses electromyogram (EMG) signals to operate an input device, allowing a wide array of interactions in powered prosthetics and human-machine interaction [1, 2]. Despite decades of research and success that have propelled control schemes from traditional dual-site control to state-of-the-art classification-based pattern recognition, the functional capabilities of current prostheses remain somewhat limited [3]. One contributing factor is the need for sophisticated coordination between the hand and wrist for advanced tasks, yet even the most advanced prostheses typically rely on sequential control schemes [4, 3]. For improved dexterity and function, simultaneous control of multiple degrees of freedom (DoFs) has been desired by prosthesis users since at least 1996 [5]. Although this feature has yet to reach the commercial market, research has shown that it can improve performance and naturalness when completing functional tasks in laboratory settings [6, 7, 8, 9].

A variety of methods exist for achieving simultaneous control [10, 11, 12, 7], but regression-based control uniquely allows simultaneous and independent proportional control of multiple DoFs [13]. Despite these advantages, much of the myoelectric control field—and commercially available devices—still rely on sequential, classification-based control [14, 15]. Arguably, this is due to concerns about robustness, partially caused by the difficulty associated with training regression-based myoelectric control systems [16, 17]. In contrast to classification-based control schemes that rely on categorical motion classes (e.g., hand open or wrist pronation), regression-based control requires collecting labelled data along each continuous DoF with high granularity. Various supervised training protocols have been used to synchronize these labelled data with EMG signals, including force transducers [18, 19], limb kinematics tracking [20, 21], mirror training (i.e., contra-lateral) [22, 23], and visual prompts [17, 16, 24, 25, 26]. While force and position mappings are viable for able-bodied

participants [18], they are not feasible for users with upper limb differences, motivating the use of visual prompts (e.g., a rising and falling bar plot that indicates the strength of the contraction) [17, 27, 24, 28]. However, unlike classification-based tasks, visual prompting in regression not only assumes that the users perform the intended contraction, but also that they closely track changes in the DoF over the entire progression.

Given the likelihood of deviations in compliance and timing, the assumption that the user accurately tracks the contraction strength throughout the data collection process may not always be valid, potentially causing discrepancies between the elicited contractions and the ground truth labels used to train the model [28, 16]. This is concerning, as preliminary findings indicate that such compliance issues degrade online performance [16, 23]. Likewise, changes in the data acquisition protocol—such as the visual prompting style or the amount of steady-state and transience used during a repetition—can drastically impact online performance, further complicating best practices for initial calibration[16, 29, 30]. Whether or not simultaneous motions should be included in the training set also affects performance. For example, some studies suggest that simultaneous motions are unnecessary in the initial calibration [31, 20, 13], while others recommend their inclusion [32, 33], and some offline results even demonstrate simultaneous motions in the initial calibration actually harm performance of isolated movements [26]. Altogether, the many uncertainties surrounding regression-based myoelectric control system calibration and the compliance issues that ultimately arise likely contribute to its limited adoption. This leads to several unsolved issues including poor performance during periods of *quiescence* (i.e., unintended cursor or prosthesis drift when the user wants to stop) [17, 16, 26] and high *action interference* (i.e., unintended activation of other DoFs when the user tries to perform a single DoF action) [17, 27, 26, 34, 35].

To overcome these challenges, this work leverages a recently proposed approach, *context-informed incremental learning* (CIIL) [36], to acquire realistic user-in-the-loop data as users interact with the system. We compare four different training protocols, two traditional non-adaptive approaches and two CIIL-based adaptive approaches, to demonstrate three primary contributions. First, we demonstrate, for the first time, the viability of CIIL in an unconstrained, velocity-based regression environment and show that adaptive models significantly improve performance in all facets of myoelectric control measured in this work. Next, we show that previous implementations of incremental learning in regression-based myoelectric control make naive assumptions about user behaviours, and that modifying the adaptation process to improve robustness to these behaviours mitigates robustness issues in downstream models. Lastly, we employ multiple metrics, including two newly proposed - namely *action interference* and *simultaneity gain* - to demonstrate that naively adding simultaneity often results in instability in the form of undesired and uncontrollable simultaneous motions. Overall, this work provides insights into how the training protocol affects performance, simultaneity, and user experience, highlighting the importance of considering user behaviour when training myoelectric control systems.

5.5 Background

5.5.1 Adaptive Learning

To avoid the noted challenges of screen-guided training (SGT) for regression and to capture more realistic control behaviours, researchers have begun collecting training data for regression-based systems in real-time co-adaptive environments (where the user and algorithm are adapting at the same time) [13, 37, 38, 39, 40, 41, 42, 43]. For

example, many have employed the approach proposed by Hahne et al. [13], where a linear regressor is adapted in real-time by using the target *position* as the ground truth label [39, 38]. Despite its efficacy, the approach used by Hahne et al. [13] can be sensitive to hyperparameter selection, prompting others to explore the use of unsupervised adaptation in position control [42, 41]. Building on these adaptation strategies, Chou et al. [43] have shown that co-adaptive learning is viable in a *velocity*-controlled environment both when using a goal target and when using eye gaze coordinates as ground truth labels, demonstrating potential uses in real-world systems. While these co-adaptive approaches have been shown to outperform SGT in regression-based myoelectric control for persons with and without limb difference [13, 39], previous implementations have been limited by two main drawbacks: 1) the use of position control and 2) naive assumptions about user compliance.

First, previous uses of co-adaptation for regression-based myoelectric control often use position-based control due to the challenges of acquiring the proportional control magnitude and assessing simultaneity in a velocity-based environment [13, 40]. Modern powered prostheses, however, use velocity control. This remarkably different paradigm (wherein EMG patterns are mapped to the device's velocity and not its position) limits the potential to capture real use-case behaviours during position-based training. Of the approaches that have been evaluated in a velocity-controlled environment [38, 43], users were shown a pre-defined path and told to mimic the velocity of the cursor based on real-time feedback. While these approaches are a step in the right direction, they do not give the user agency in how they accomplish a task. For example, in [43], it was assumed that the velocity of the cursor should be proportional to the cursor's distance from the target. While this is a fair assumption, previous work suggests that users act primarily at two speeds: (1) fast and imprecise and (2) slow and precise [36]. Correspondingly, the constraints placed on learning in these environments limit the goal-oriented behaviours that they can

capture.

The second limitation is that previous works assume that the user is either on (for position-based), or going directly to (for velocity-based), the target [13, 43]. Due to the complexities of regression-based control, users often do not travel in the optimal direction, invalidating this oversimplified assumption. For example, Smith et al. found that users often use 1-DoF movements for finer movements near the target [25]. In this case, assuming that the user is always travelling along the optimal path would relabel 1-DoF movements as multi-DoF movements, providing the model with incorrect pseudo-labels and likely contributing to action interference (the unwanted activation of additional DoFs). We, therefore, argue that previous attempts at co-adaptation in regression environments have been built upon naive assumptions about the user and may not address some of the key issues hindering regression-based myoelectric control.

5.5.2 Context Informed Incremental Learning

Shifting the paradigm away from conventional methods of the historically used screen-guided or prosthesis-guided training, CIIL is a training approach that has previously shown promise for classification tasks [36, 44, 45]. By relying solely on the context provided by the environment (i.e., whether the user is accomplishing the general goal of the training environment, such as moving towards a target along any axis), CIIL enables flexible and sample-efficient user-in-the-loop adaptation. For example, it has been shown that CIIL can adapt a model with a limited amount of training data (1 second per class) [45], in the case of a catastrophic electrode shift [44], and can even learn from scratch and outperform traditional SGT approaches [36]. Moreover, this approach is suitable for simple statistical models like Linear Discriminant Analysis and more complex deep learning architectures.

CIIL leverages every model prediction and the contextual suitability vector for that action. For example, in the case of a target acquisition task, the action might move the cursor to the right, and the possible contextual suitability vector is ['Up' 'Right', 'Nothing'] since the target is up and to the right. Alternatively, the suitability vector might be ['Wrist Pronation', 'Hand Close', 'Nothing'] when approaching an object like a tennis ball with a prosthesis. If the action taken by the controller aligns with the suitability context vector, this is considered to be a positive context match, and the model's decision should be reinforced. In contrast, if they do not align, the context match is negative, and the model's parameters should be updated to reflect this error. If there are multiple suitable choices, soft labels can be used to update the model to eliminate issues associated with guessing [36]. Correspondingly, the only assumption of this approach is that user behaviours generally align with the goals of the tasks (e.g., they are trying to acquire the targets). Although similar to reinforcement learning [40], CIIL is much more sample efficient because the context extracted from the environment limits its need for exploration. Moreover, compared to previous adaptive approaches, there is no direct assumption of velocity or direction. While this approach has never been employed in a regression environment, it is hypothesized that it will be suitable for overcoming many of the limitations of current training approaches.

5.6 Methods

5.6.1 Participants

Sixteen participants (ages 22-32, 8 male, 8 female) took part in this experiment as approved by the University of New Brunswick Research Ethics Board (REB 2022-122). Informed consent was obtained from all participants in accordance with the

Declaration of Helsinki, and no incentives were provided. The experiment lasted approximately one hour.

5.6.2 Signal Processing

LibEMG, an open-source Python library for processing EMG, was used to facilitate data streaming, the creation of visual prompts, the presentation of the online validation environment, and real-time myoelectric control [46]. A 2^{nd} order highpass filter ($f_c = 20$ Hz) was applied to reduce DC noise and motion artifact, and a 3^{rd} order notch filter ($f_c = 60$ Hz) was used to reduce power line interference. The Thalmic Labs Myo Armband, a wearable 8-channel cuff with a sampling frequency of 200 Hz, was placed on the participant's dominant forearm. EMG signals were segmented into 150 ms windows with an increment of 40 ms. Wavelet energy (WENG) features were then extracted from each window and passed to the model for real-time prediction [47].

5.6.3 Data Acquisition

As shown in Fig. 5.1, participants sat in an upright position facing a computer display screen. The height of the chair was adjustable to provide a comfortable setup for all participants. The Thalmic Labs Myo Armband was placed on the participant's dominant forearm. Participants were free to place their arm wherever they felt was most comfortable as long as nothing pressed against the electrodes.

All SGT-based data used to initialize or fully train the various models were collected using mimic training, a standard protocol used throughout the field for collecting training data for regression-based myoelectric control systems [23, 16, 24, 26, 27]. Briefly, participants were asked to follow (*mimic*) a moving cursor in real-time and



Figure 5.1: Illustration of data acquisition setup. EMG electrodes were placed on the participant's dominant forearm. Participants sat in front of a computer screen that displayed relevant information for each task.

the values at each frame were used as ground truth labels. Similar to other works [31, 20, 16], a point moved along a Cartesian plane, where the x-axis represented one DoF (hand open/close) and the y-axis represented the other DoF (wrist rotation). Participants were asked to map the strength of their contraction to the position of the point in real time as it moved in multiple directions. For isolated (single DoF) motions, a 50-second repetition of SGT was composed of ten 2-second oscillations along the x-axis for the first DoF (hand open/hand closed), followed by 5 seconds of rest, followed by ten 2-second oscillations along the y-axis for the second DoF (wrist pronation/wrist supination), followed by another 5 seconds of rest. For simultaneous motions (both DoFs active), participants followed a similar pattern but the cursor's path was rotated 45° to prompt both DoFs simultaneously.

5.6.4 Models

Four distinct models were investigated to compare the effects of including combined motions in the calibration data, adapting with user-in-the-loop data, and mitigating assumptions about user compliance. All models used the same underlying neural network architecture but were trained using different protocols. Each model was composed of a multi-layer perceptron with layers of 64, 32, 10, and 2 neurons. Batch normalization, a rectified linear unit activation function, and dropout ($p = 0.2$) followed each linear layer except the last. Training for all models used the Adam optimizer, L1 loss, and a batch size of 64 windows. All models were provided with 300 seconds of training data, acquired from SGT and/or CIIL adaptation, followed by 300 seconds of online validation in a Fitts' Law environment. For non-adaptive models, the training data consisted of 300 seconds of SGT-acquired data. For the adaptive models, training data comprised 50 seconds (i.e., 1 repetition) of isolated SGT data followed by 250 seconds of CIIL-based adaptation. In this way, participants spent the same amount of time collecting data for both the adaptive and non-adaptive approaches, and each was given the same amount of data for model training. For each model, participants underwent a training phase, followed by a 5-minute validation phase, and then completed a survey before proceeding to the next model. To minimize bias in survey responses and user behaviours, participants were not informed of the model they were using at any given time and were only told that the study compared adaptive and non-adaptive approaches. The order of the models was balanced across participants using a Latin square design. Each model is described in detail below.

5.6.4.1 W-SGT

A conventional non-adaptive model trained for 150 epochs using 300 seconds of SGT data collected from isolated, within-DoF motions (thus *W-SGT*).

5.6.4.2 C-SGT

A conventional non-adaptive model trained for 150 epochs using SGT data, with 150 seconds collected from isolated motions and 150 seconds from simultaneous, combined-DoF motions (thus *C-SGT*).

5.6.4.3 O-CIIL

An adaptive model seeded with 50 seconds of SGT data followed by 250 seconds of user-in-the-loop co-adaptation. The model used as a starting point for adaptation was trained for 150 epochs using only isolated motion SGT data. For each prediction made by the model throughout the adaptation phase, pseudo-labels were created based on the context of the environment and added to a memory buffer. Once new data were added to the buffer, the model was trained for 5 epochs with all the data currently available (i.e., SGT data + pseudo-labels collected during adaptation). Data were continuously being added to the memory buffer at all times, so any data that were collected while the model was adapting were used once the model was ready for the next round of adaptation. Pseudo-labels were created following an *oracle*-based approach (thus *O-CIIL*), so the user was assumed to be going directly to the target at all times during the adaptation phase. When the cursor was inside the target bounds, the corresponding pseudo-label was set to 0 for both DoFs (i.e., rest) to reinforce that the user wanted to dwell in the target. All predictions that occurred when the cursor was outside of the target were treated as shown in Fig.

5.2, where the corresponding pseudo-label points in the optimal direction from the cursor to the target. The magnitude of the pseudo-label vector was determined based on the distance between the cursor and the target at the time of the prediction, a novel concept introduced in this work. The relationship between target distance and proportional control magnitude (see Equation 5.1) was determined based on piloting.

$$p = \min(1, \sqrt{\frac{d - TR}{IR}}) \tag{5.1}$$

where d is the distance between the cursor and the target, TR is the target radius, IR is the ISO Fitts' radius (i.e., the radius of the circle that defines the target locations), and p is the corresponding proportional control magnitude of the pseudo-label vector. Equation 5.1 was formulated to normalize based on the distance between the targets with a maximum value of 1.

5.6.4.4 T-CIIL

An adaptive model seeded with 50 seconds of SGT data followed by 250 seconds of user-in-the-loop co-adaptation. The model used as a starting point for adaptation was trained for 150 epochs that only used isolated motion SGT data. For each adaptation, the model was trained for 5 epochs. As with O-CIIL, T-CIIL used the mapping from target distance to proportional control shown in Equation 5.1. Instead of assuming that the user always took the optimal path to the target, however, T-CIIL employed a tolerant approach (thus *T-CIIL*) by using the context of the task to modify the pseudo-labels to mitigate known issues associated with user compliance (i.e., the model is aware that user behaviours may deviate from the optimal action and can *tolerate* these variations). As shown in Fig. 5.3A, pseudo-label

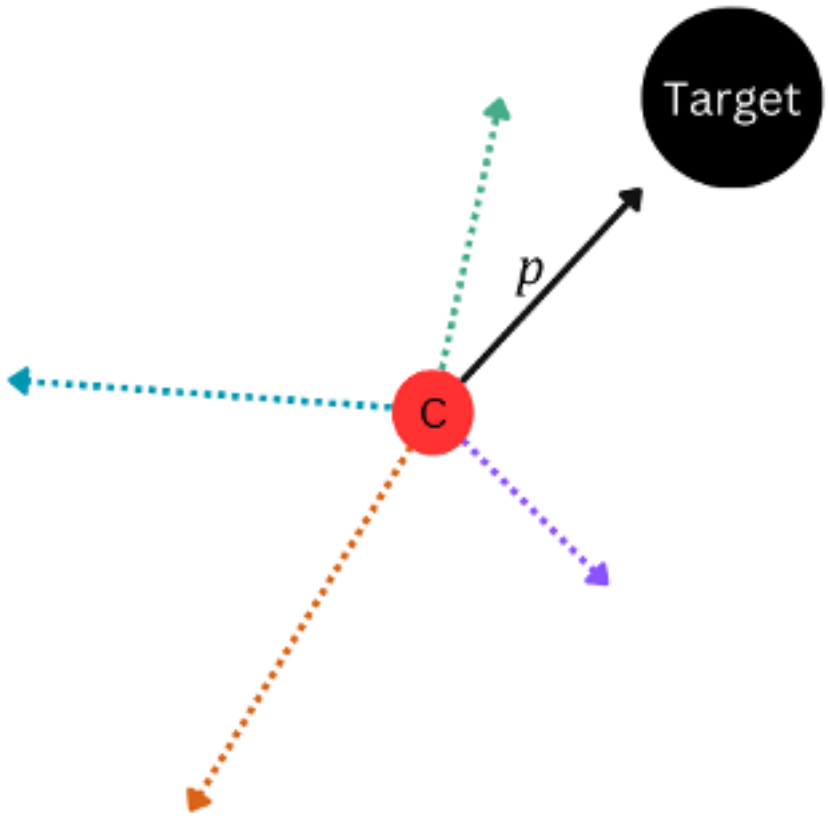


Figure 5.2: Illustration of pseudo-label creation for O-CIIL. Predictions are represented by dotted arrows and the corresponding pseudo-label is represented as a solid black arrow. The length of the vector, p, is determined using Equation 5.1.

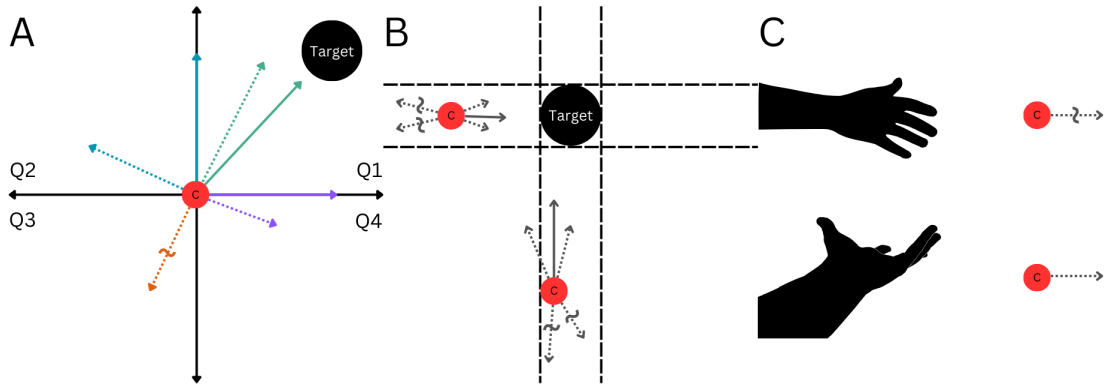


Figure 5.3: Illustration of pseudo-label creation for T-CIIL. Predictions are represented by dotted arrows and corresponding pseudo-labels are represented by solid arrows in the same colour. The red dot, labelled "C," represents the current cursor position, and the black circle, labelled "Target," shows the current goal target. Arrows marked with the "∼" symbol represent predictions that were considered poor for adaptation and have been ignored. A) Example of main pseudo-label decision flow. Predictions in the correct quadrant (green) are relabelled to point in the optimal direction to the target. Predictions for which one component is correct (blue and purple) are relabelled to point in the direction of the correct component. Predictions where neither component is in the correct direction (orange) are ignored. The magnitude of the corresponding pseudo-label vector is determined based on the distance from the cursor to the target. B) Illustration of target runways. In the case where the cursor is in line with the target (i.e., they can reach the target by moving in a single component), only the component that requires movement to reach the target is reinforced. C) Visualization of activation threshold. For all decisions, pseudo-labels are only created for adaptation if the model input is above the activation threshold (determined based on no motion data from calibration).

creation was first informed by the direction of the prediction. For each prediction, the cursor was assumed to be at the origin of an axis and the prediction vector. The subsequent pseudo-label was then crafted based on the quadrant of the prediction. If the prediction fell in the quadrant opposite of the target (quadrant 3, in this case), no pseudo-label was generated, and the prediction was excluded from adaptation (though the cursor still moved in that direction). If one of the components was correct (quadrants 2 or 4 in this case), the resultant pseudo-label pointed in the direction of only the correct component (i.e., on-axis). If both components were in the correct direction (quadrant 1 in this case), the pseudo-label vector pointed in the optimal direction to the target (similar to O-CIIL). This decision flow was used to create pseudo-labels in most cases, but there were two cases where it did not apply. First, as shown in Fig. 5.3B, if the cursor was in line with a target in either component, the pseudo-label pointed solely to the correct component that would move the cursor to the target. Second, as shown in Fig. 5.3C, an activation threshold was employed to improve robustness to periods of no activation (e.g., after a target was acquired or if a user was slow to respond) being labelled as high-magnitude predictions. Data were then only used for adaptation if the level of EMG activity (i.e., the mean wavelet energy across the frequency bands) was above this activation threshold. To avoid spurious labelling of low-intensity contractions, this decision was computed as a majority vote of the 5 most recent windows. The activation threshold (see Equation 5.2) and the number of decisions in the majority vote were determined empirically through pilot experiments.

$$at = \mu_{WENG}^{NM} + 0.2\sigma_{WENG}^{NM} \tag{5.2}$$

where μ_{WENG}^{NM} is the mean of the wavelet energy of all no-motion samples (i.e., both components have a label of 0) in the calibration data, σ_{WENG}^{NM} is the standard

deviation of the no motion samples in the calibration data, and at is the activation threshold.

5.6.5 Online Validation Task

Each model was evaluated using an ISO 9241-9 inspired Fitts' Law validation task shown in Fig. 6.7 [48]. Pilot work found that mapping wrist rotation to vertical movement on a Cartesian plane was unintuitive for many participants, so the target-reaching task in this work was modified to map wrist rotation to the rotation of a cursor in a Polar plane. The direction of rotation was modified based on the handedness of the user to provide the most intuitive mapping (i.e., pronation rotated the cursor counter-clockwise for right-handed participants and vice-versa for left-handed participants). Hand open/close was mapped to the radius of the aperture, where the aperture was a circle with the origin centered at the bottom of the screen. The underlying cursor locations were recorded in a Cartesian plane and then rendered to the screen in the Polar plane to ensure the calculation of downstream usability metrics was unchanged. To acquire a target, the user had to move the cursor into the current goal target and dwell within the target for 2 consecutive seconds. Once the cursor was in the target, the cursor colour changed to blue and a 2-second timer was shown in the top left of the screen. If the user did not acquire the target within 10 seconds, the trial was considered a failure and the next trial began with a new target. Participants were informed that they could move simultaneously and were encouraged to do so, but simultaneous motions were not required to acquire the target.

Eight online usability metrics were used to evaluate models in this work to analyze multiple components of myoelectric control. A trial was considered successful if it was acquired and dwelled in for 2 seconds within a 10-second timeout. To account

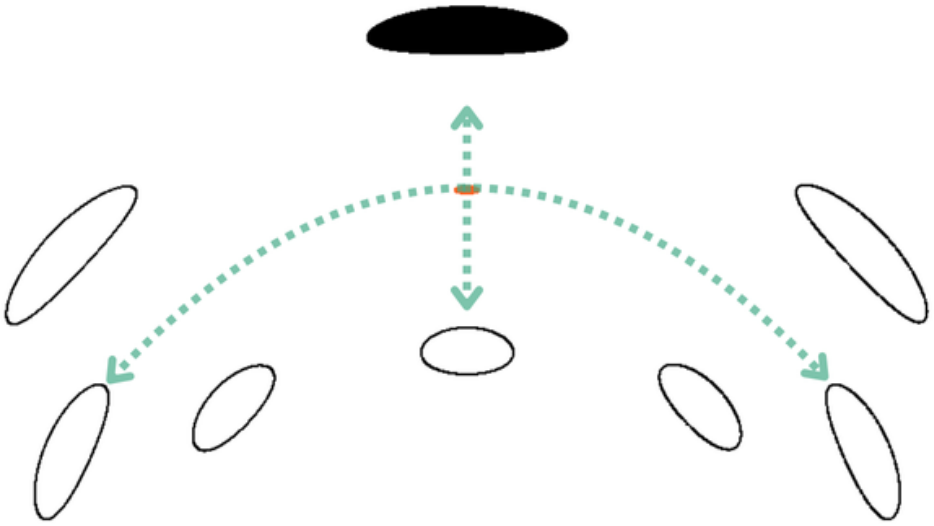


Figure 5.4: Online ISO 9241-9 inspired Fitts' Law task [48]. Targets are drawn in black and the cursor is drawn in orange. The current goal target is filled in, while potential target locations are outlined. Green arrows annotate the two DoFs, but this was not shown to the user during the experiment. As shown by the vertical arrow, hand open/close increases and decreases aperture size, respectively. effectively moving the cursor in the vertical direction. As shown by the curved arrow, rotating the wrist rotates the cursor in the corresponding rotational direction, where the direction of the mapping was dependent on the handedness of the participant.

for the user learning, the first 20 trials (during which the average performance across users began to plateau) were considered warm-up trials and excluded from the calculation of validation metrics. Validation metrics were computed using the remaining trials. Trials that timed out were still included in the calculation of performance metrics, and were optimistically assigned a movement time of 10 seconds (i.e., equal to the timeout). Each usability metric is described below.

5.6.5.1 Throughput

Throughput describes the amount of control information that can flow through the system [48]. It is a measure of the difficulty of the task (see Equation 5.3) and the speed at which it was completed (see Equation 6.2).

$$ID = \log_2(\frac{D}{W} + 1) \tag{5.3}$$

where ID is the index of difficulty, D is the distance from the cursor starting position to the center of the target, and W is the width of the target.

$$Throughput = \frac{ID}{MT} \tag{5.4}$$

where MT is the time taken to acquire the target. The dwell time is subtracted from the movement time for successful trials.

5.6.5.2 Path Efficiency

Path efficiency compares the path taken to acquire the target to the most optimal path, offering insight into the control of simultaneous motions [49]. It is a measure of the distance travelled relative to the minimum distance required to reach the

target (see Equation 5.5).

$$Path\ Efficiency = \frac{FP}{DT} \times 100\% \tag{5.5}$$

where FP is the optimal (i.e., fastest) path to the target from the cursor's starting position and DT is the total distance travelled during the trial.

5.6.5.3 Overshoots

Overshoots are a count of the number of times the cursor enters the target and then exits before it is finally acquired [49]. Overshoots help quantify the ability to stop and stay at rest when desired.

5.6.5.4 Completion Rate

The completion rate measures how many targets were successfully acquired before the timeout relative to the total number of trials (see Equation 5.6).

$$Completion\ Rate = \frac{ST}{T} \times 100\% \tag{5.6}$$

where ST is the number of successful trials and T is the total number of trials.

5.6.5.5 Number of Trials

The number of trials (successful or unsuccessful) reached was able to be computed as a metric in this work because the validation trial was based on a time limit for each model instead of a set number of trials.

5.6.5.6 Action Interference

A novel metric was used to quantify the action interference exhibited by the controller (see Equation 5.7). Action interference occurs when the participant attempts to move in isolation (on activating one DoF), but an unintended activation occurs in a different DoF, causing undesired simultaneous motion [35]. It was assumed that the user wanted to move in isolation when they were aligned with the target in one DoF, so action interference was calculated as the total distance travelled in the DoF associated with an axis that the cursor was already aligned in (i.e., if the cursor was below the target, movement to the left or right increased action interference). This metric describes the distance travelled in an unintended DoF, so a lower value indicates better performance. Trials where the cursor did not align with the target at any point were excluded when calculating action interference.

$$Action\ Interference = \frac{1}{S_x} \sum_{i=1}^{S_x} |p_x(i)| \tag{5.7}$$

where S_x is the number of samples where the cursor is in line with a target in a single component (i.e., only a single DoF activation is needed) and p_x represents the predictions associated with the unwanted component (e.g., in a Cartesian space, the unwanted component would be movement in the x-axis if the cursor was already directly below the target).

5.6.5.7 Drift

Drift occurs when the user wants to be at rest, but the cursor continues to move. It was assumed that the participant wanted to be at rest when in the target, so drift was quantified as the distance travelled while in the target (see Equation 5.8). Trials where the cursor did not enter the target at any point were excluded when

calculating drift.

$$Drift = \frac{1}{S_a} \sum_{i=1}^{S_a} \|p_a(i)\|_1 \tag{5.8}$$

where S_a is the number of samples where the target is acquired (i.e., the cursor is inside the target), p_a represents the predictions made while the cursor was in the target throughout the entire trial, and $\|\cdot\|_1$ indicates the L1 norm.

5.6.5.8 Simultaneity Gain

The amount of time spent moving simultaneously (see Equation 5.9) can capture how much simultaneity was used, but not whether this simultaneity was useful or intentional. Simultaneity gain, a novel metric proposed in this work, was devised to measure how well the participant could control the additional simultaneity provided by the model (see Equation 6.4). Models that provided higher simultaneity but similar path efficiency would suggest that the additional simultaneity could not be leveraged fully or did not add value, so the resulting simultaneity gain is lower. Trials that did not have any simultaneous predictions were excluded when calculating simultaneity gain.

$$\%\ Simultaneity = \frac{S_s}{S} \tag{5.9}$$

where S_s is the number of samples where both predictions are greater than 0 and S is the total number of samples.

$$Simultaneity\ Gain = \frac{Path\ Efficiency}{\%\ Simultaneity} \tag{5.10}$$

5.6.6 Statistics

All data were first checked for normality using the Shapiro-Wilk test. For those metrics that exhibited a normal distribution (throughput and trials complete), a repeated-measures ANOVA was conducted to determine if there were significant differences ($p < 0.05$) between models. For metrics that showed significant effects in the ANOVA, pairwise comparisons were performed using Bonferroni correction to identify which specific models differed significantly. For metrics that showed non-normal distributions (path efficiency, overshoots, completion rate, action interference, drift, and simultaneity gain), a Friedman test was conducted to determine if there were significant differences between models. Follow-up pairwise Nemenyi tests were performed to identify which models differed significantly. All stats were computed in R.

5.7 Results

5.7.1 Adaptive Models Outperform SGT Models

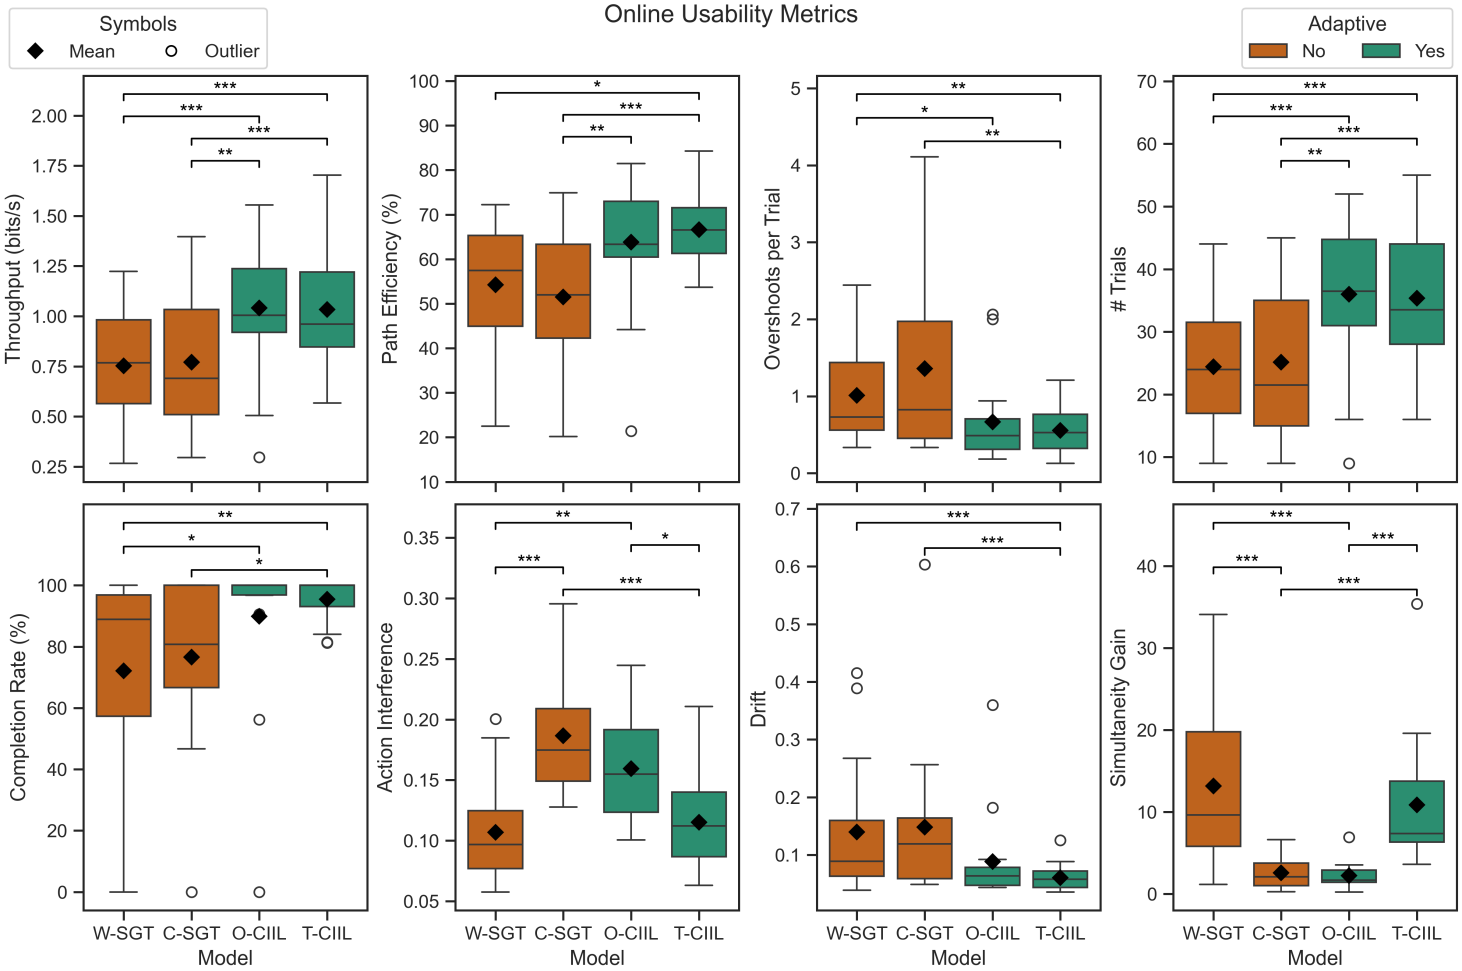


Figure 5.5: Summary of online usability metrics over the validation period. The first 20 trials were excluded as they were deemed 'user learning' trials. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

The primary goal of this work was to demonstrate the viability of co-adaptive models in an unconstrained, velocity-based regression environment and contrast them against non-adaptive protocols used throughout the field. Despite using equal amounts of training data, the usability metrics shown in Fig. 5.5 demonstrate that the adaptive CIIL models outperformed both non-adaptive SGT models on average for most metrics. O-CIIL significantly ($p < 0.05$) outperformed W-SGT in throughput, overshoots, trials, and completion rate, and C-SGT in throughput, path efficiency, and trials. Similarly, T-CIIL significantly ($p < 0.05$) outperformed W-SGT in all metrics except action interference and simultaneity gain, and was significantly better than C-SGT in all metrics ($p < 0.05$).

In contrast to the other usability metrics, W-SGT performed the best in terms of action interference, significantly improving action interference and simultaneity gain compared to C-SGT and O-CIIL ($p < 0.01$). Similarly, T-CIIL significantly outperformed C-SGT and O-CIIL for action interference and simultaneity gain ($p < 0.05$). Furthermore, T-CIIL was roughly equivalent to W-SGT on average for both metrics.

The loss curves shown in Fig. 5.6 confirm that the L1 loss monotonically decreased while training with SGT data and began to plateau, but that the loss increased again when incorporating new behavioural data using the adaptive models. Despite this, Fig. 5.7 shows that throughput steadily increased over the adaptation phase. Fig. 5.7 also shows that the throughput over the validation period exhibited a small learning effect that was similar across all models.

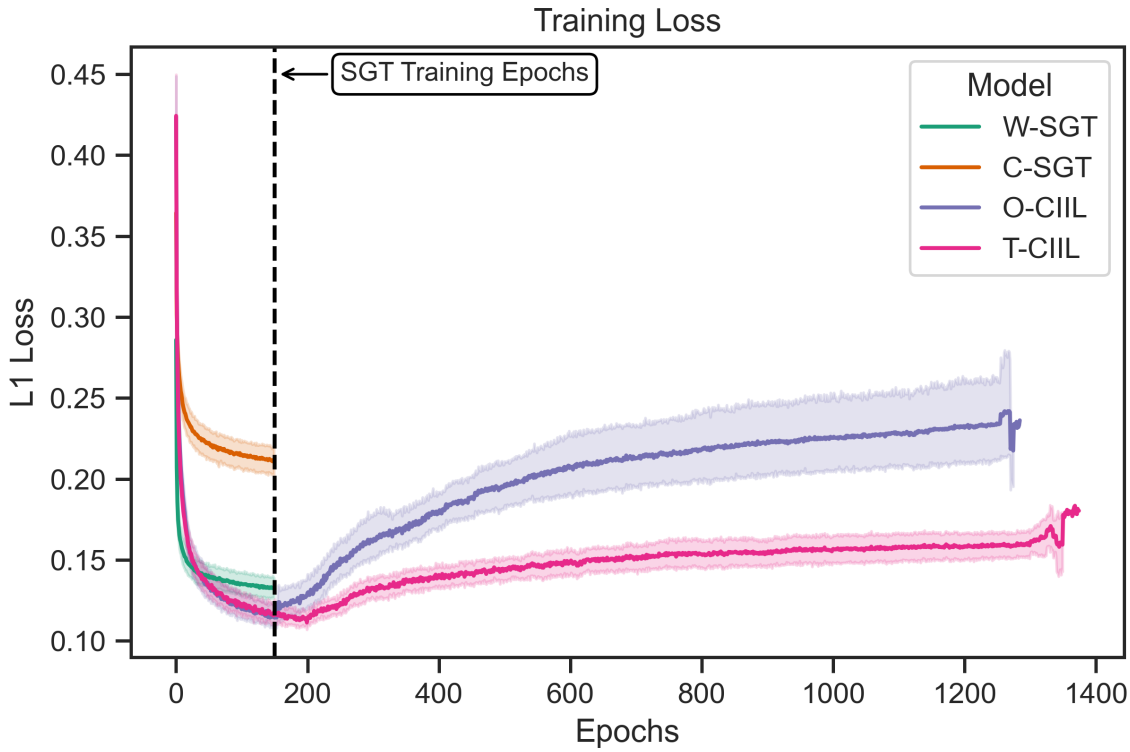


Figure 5.6: The L1 loss for each model. The mean across participants is shown as a solid line, and 95% confidence intervals are shown as shaded regions. Loss after the dotted line represents the loss during the adaptation period. The sudden changes in the mean and confidence intervals at the end of the loss curves for the adaptive models are simply a byproduct of the different number of epochs per participant during the adaptation period.

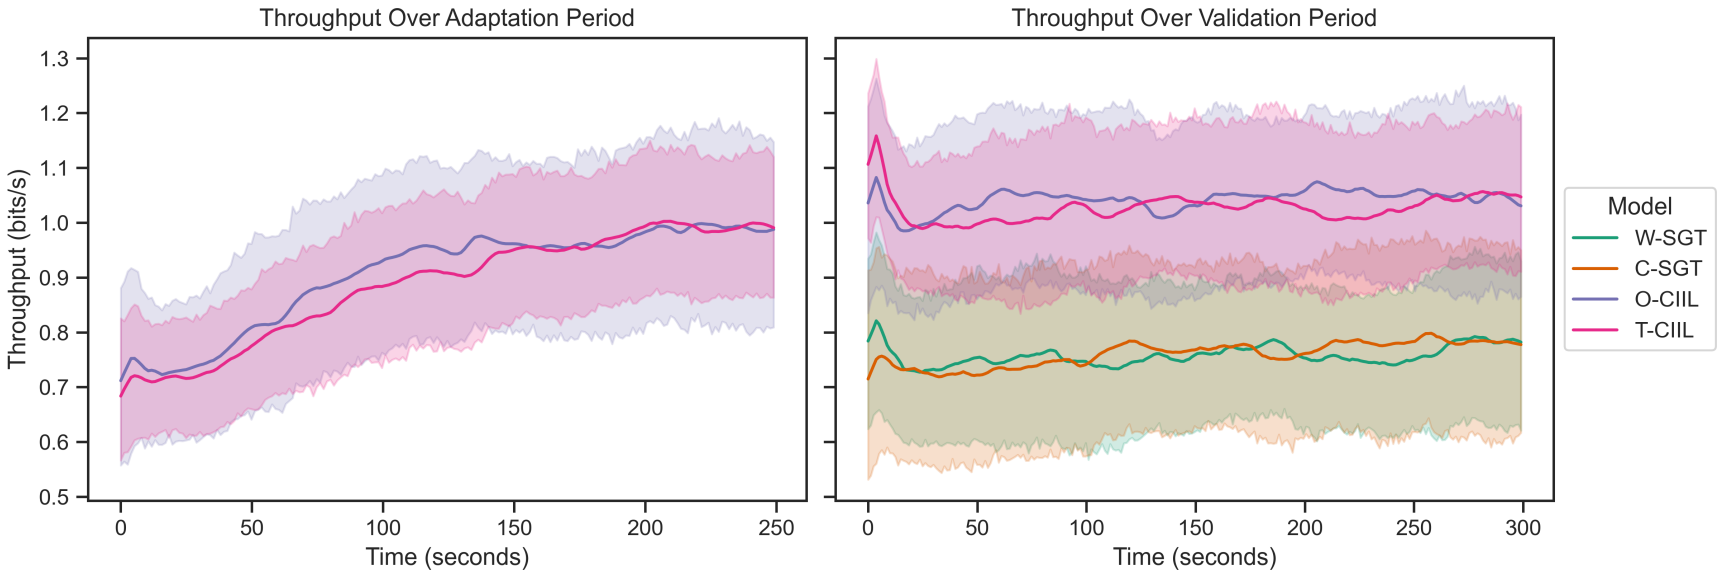


Figure 5.7: A rolling average (N = 60 seconds) of throughput over time during the adaptation (left) and validation periods (right). The mean across participants is shown as a solid line, and the 95% confidence intervals are shown as shaded regions.

5.7.2 The Effect of User Compliance Adjustments in an Adaptive Setting

The proposed novel method for capturing goal-oriented behaviours in a co-adaptive setting leveraged context from the environment while acknowledging known issues with user compliance to improve the creation of pseudo-labels for adaptation. The upward trend in throughput over time shown in Fig. 5.7 confirms that the user and model co-adapted effectively throughout the adaptation period for both adaptive approaches, but each adaptive approach offered subtle differences. For example, Fig. 5.6 shows that the loss throughout the adaptation period increased noticeably more for O-CIIL than T-CIIL.

The differences observed during the adaptation phase continued into the validation phase, leading to distinct performance trade-offs. To visualize the simultaneity exhibited by participants for each model, Fig. 5.8 presents a two-dimensional histogram (i.e., heatmap) of their outputs across all participants. Predictions at the center of the heatmap correspond to rest, those along the axes of a DoF represent isolated motions, and all other predictions indicate simultaneous motions. The % simultaneity for all predictions made by each model was then computed using Equation 5.9. As shown in Fig. 5.5 and 5.8, the modifications made to T-CIIL improved path efficiency, overshoots, completion rate, action interference, drift, and simultaneity gain, but not throughput or number of trials relative to O-CIIL. Among the online metrics, a significant difference between the two adaptive approaches was found only for action interference ($p < 0.05$).

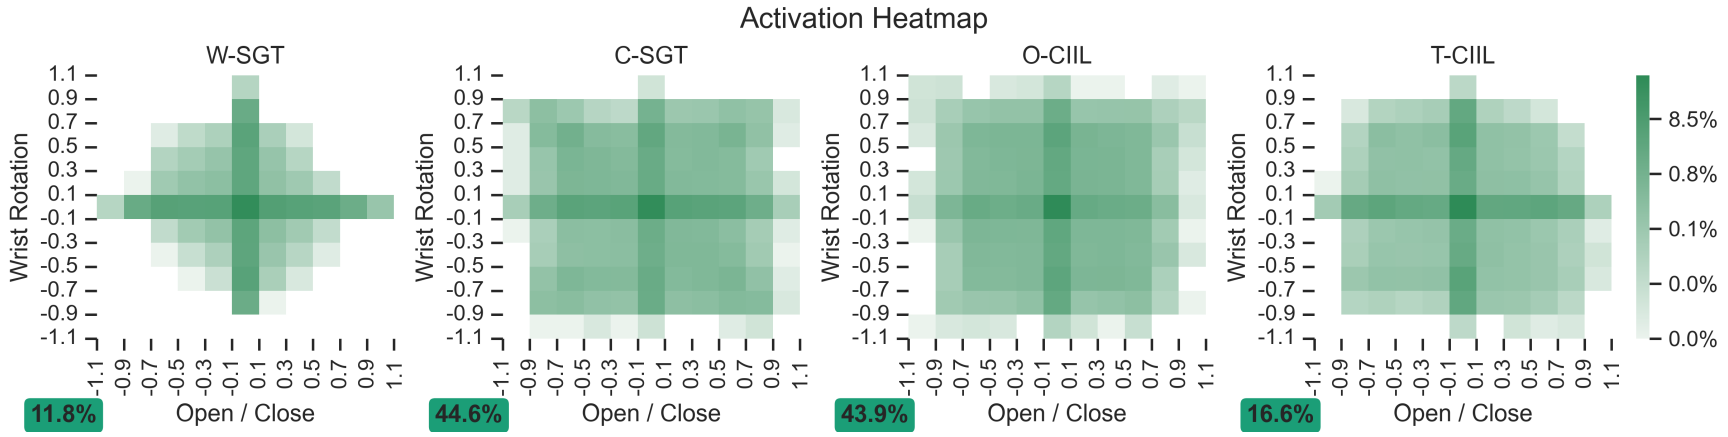


Figure 5.8: Visualization of the use of simultaneity per model based on model predictions during the validation period. The heatmaps show the frequency of each activation level for both DoFs at once. The total % simultaneity (see Equation 5.9) for each model is annotated in a text box below the corresponding heatmap. A colour bar, using a logarithmic scale, maps colours to the percentage of total predictions made by each model.

5.7.3 The Cost of Simultaneity

Non-adaptive SGT models with and without combined motions in the calibration data were tested in this work to confirm the impact of adding simultaneous training data. Participants used far more simultaneity with C-SGT based on Fig. 5.8, but this simultaneity was not always beneficial. As demonstrated by Fig. 5.5, C-SGT had a higher throughput than W-SGT, but the worsened performance in overshoots, action interference, drift, and simultaneity gain suggests that it came at a cost of higher instability. As reflected in the significant difference in simultaneity gain ($p < 0.001$), Fig. 5.5 shows that C-SGT had worse path efficiency than W-SGT on average despite the marginal increase in simultaneity used. The same can be said for the adaptive models, where T-CIIL had less than half the % simultaneity of O-CIIL but improved path efficiency and significantly better simultaneity gain ($p < 0.001$).

5.7.4 Participants Preferred Adaptive Models

In addition to the objective performance results, subjective results were captured in the form of survey questions. Fig. 5.9 showcases the opinions expressed by participants for each model based on multiple aspects of control. Generally, participants preferred adaptive models over non-adaptive models based on survey results, but the difference between adaptive models was not as apparent. T-CIIL was rated higher than or equal to O-CIIL for questions involving control, but T-CIIL was rated lower regarding the control of the cursor's speed. Additionally, it was found that participants generally found the actual training protocol of the adaptive approaches more engaging.



Figure 5.9: Results of the survey posed after each model was used. The distribution of answers for each question are shown.

5.8 Discussion

The results of this work highlight some key challenges associated with regression-based myoelectric control, particularly regarding how the training protocol can drastically affect usability. Previous uses of co-adaptation [13, 39, 40] have shown promising performance compared to models trained using SGT, but have made assumptions about user behaviours and compliance and often employ position-based control instead of velocity-based control. First and foremost, this work shows that co-adaptation can be achieved in an unconstrained, velocity-based environment and that models trained using this approach are significantly better than those trained using SGT. Both adaptive CIIL approaches outperformed the SGT approaches in

all metrics except for two, showcasing that carefully planned co-adaptive approaches can overcome many of the challenges associated with regression-based myoelectric control. The only metrics where adaptive models did not outperform non-adaptive models were action interference and simultaneity gain, for which W-SGT had the best values, significantly better than C-SGT and O-CIIL. This result may be misleading, though, because W-SGT predicted far fewer simultaneous motions in general than the other models. Despite this, T-CIIL was roughly equivalent to W-SGT on average for both metrics while significantly improving path efficiency, highlighting the benefit of appropriately leveraging context from both the environment and the user behaviours when designing user-in-the-loop systems.

Despite differences in training protocols, ensuring a fair comparison across all models was a key goal of the experimental design. Each model received 300 seconds of training data and the loss curves shown in Fig. 5.6 indicate that training began to plateau by 150 epochs for the SGT models. Therefore, we argue that all models were evaluated fairly in terms of data, training, and user burden. Although training loss decreased for all models during initial training, it increased for both adaptive models during the adaptation phase. In contrast, Fig. 5.7 shows that throughput steadily increased throughout the adaptation phase, confirming that the incremental learning employed in this work resulted in better models. This is consistent with previous work [36] implementing CIIL, where it was found that optimizing for online performance can degrade offline performance. In addition to data considerations (e.g. the inclusion of more variable training data during the adaptation process), it is possible that some of the benefits seen with the adaptive models could be associated with user learning. While this is a valid consideration, Fig. 5.7 demonstrates that all models showed a similar increase in throughput over the validation phase, indicating that the large benefits seen during adaptation — and, by extension, validation — are primarily driven by incremental learning rather than pure user learning. Overall,

the results of this work suggest that the benefits of incremental learning far outweigh those of increasing SGT data or simply allowing the user more time to learn the system.

Despite both proving beneficial, the results of this work demonstrate clear differences between the two adaptive approaches. The average throughputs between the adaptive approaches were very similar, but the other usability metrics show the benefits of more appropriately considering user compliance in T-CIIL. T-CIIL outperformed O-CIIL in path efficiency, overshoots, completion rate, action interference, drift, and simultaneity gain. Furthermore, T-CIIL significantly improved many of these metrics compared to the non-adaptive SGT models, whereas O-CIIL did not always. Specifically, it's worth noting that the action interference was significantly better for T-CIIL than O-CIIL, confirming the notion discussed in the introduction that assuming that the user always goes directly to the target contributes to action interference. Fig. 5.8 shows that O-CIIL had far greater % simultaneity than T-CIIL, but T-CIIL had better path efficiency and simultaneity gain, furthering the point that O-CIIL provides at least some degree of uncontrollable simultaneity. Smith et al. [25] have highlighted the importance of isolated motions for simultaneous control systems, specifically for finer movements. T-CIIL has improved on these problem areas compared to O-CIIL, confirming that the tweaks made to the pseudo-label decision flow are beneficial. Fig. 5.6 corroborates this by showing that the loss increased more throughout the adaptation phase for O-CIIL than T-CIIL, suggesting that the labels provided by O-CIIL, which assumes you're always going to the target, have a higher entropy than those provided by T-CIIL. Along with contributing to action interference, the incorrect labels produced by O-CIIL during inactive periods mislabel these periods as high-magnitude predictions, explaining why O-CIIL has higher drift than T-CIIL. Despite these limitations, both adaptive models reduce drift compared to the non-adaptive models, with T-CIIL showing a

significant reduction. Overall, the improved action interference, simultaneity gain, and drift of T-CIIL shows that the modifications made to T-CIIL helped improve the control of both isolated and simultaneous motions at high and low amplitudes, demonstrating the benefit of leveraging task-specific context when training user-in-the-loop systems.

Although T-CIIL raised both the upper and lower bounds of performance, its average throughput was slightly lower than that of O-CIIL. Most other metrics clearly suggest that T-CIIL should have yielded better control, suggesting that this was influenced by some other factor. Interestingly, once a strong machine learning model has been established, throughput is largely determined by cursor speed, whereas the other usability metrics reward movement in the correct direction over speed [27, 36]. The survey results shown in Fig. 5.9 also suggest that participants often preferred the control of T-CIIL over O-CIIL, but the speed of the cursor was better for O-CIIL. Based on these two results, it is likely that T-CIIL suffered from an issue with the speed mapping. Some parameters for T-CIIL, such as the mapping from cursor distance to proportional control magnitude (see Equation 5.1) and handling the activation threshold when creating pseudo-labels, were determined based on brief piloting. These parameters are almost certainly not ideal for all users though. For example, if the activation threshold was too high for some participants, T-CIIL may not collect data until the participant is at 30% contraction strength, which could harm the speed mapping as it misses out on the low end of proportional control. The low end of proportional control has been shown to be crucial to online performance [36, 17], so any issue with the activation threshold could drastically impact usability. It is difficult to know the root cause of this without further investigation, though, so future work will explore how to determine a speed mapping or activation threshold that is robust for all users, such as through an adaptive approach similar to one employed in a previous implementation of CIIL [36]. Despite these limitations, par-

ticipants clearly preferred the adaptive models over non-adaptive models, both in terms of the control and the training process. The superior usability metrics explain why users preferred the control of the adaptive approaches, while their preference for the adaptive training protocol was likely due to the more engaging game environment during adaptation compared to SGT. It is worth noting, however, that these preferences were expressed in a participant survey, exposing them to potential bias. Although participants were blinded as much as possible, the nature of the different training approaches—such as the inclusion of combined motions or the presence of adaptation—made it apparent which model was being used in some cases. However, while the difference between adaptive and non-adaptive approaches may have been noticeable, participants were not provided with any additional information to distinguish between the approaches. Additionally, these subjective findings aligned with the objective usability metrics, reinforcing the conclusions drawn from these results. Overall, the improved performance of the adaptive CIIL models demonstrates the benefits of incremental learning for myoelectric control, and the differences between the two adaptive approaches highlight the impact user compliance issues can have on performance.

It was noted empirically that the benefits of T-CIIL were greater than O-CIIL for novice users, likely due to their increased propensity for compliance issues. This bodes well for the ability of T-CIIL to generalize to persons with limb differences who may not have the same level of proprioception and/or who may calibrate more frequently with reduced attention. Despite this promising result, the results of this work must be extended to include users with a limb difference to verify the efficacy of CIIL with the target population. Similarly, CIIL has yet to be employed on a physical device, which offers new confounds that cannot be captured in a virtual environment without a socket. While training on a physical device likely provides the most context-rich and relevant data for the end task, recent work[50] has shown that

operating in virtual environments can translate to skill transfer on-device. Along with additional confounds, employing adaptive approaches in a physical environment produces new experimental challenges, such as collecting context from the environment to create pseudo-labels. Chou et al.[43] have demonstrated that models trained using real-world context can produce similar performance to models trained using context from the virtual environment, motivating future work employing CIIL on a physical device. Furthermore, previous work with CIIL[44] has demonstrated its ability to recover from common confounds associated with on-device use, such as electrode shift. To validate the clinical viability of this approach, future work should verify CIIL both with users with a limb difference and on-device. Another key clinical factor that remains unexplored in this work is generalization across sessions, as many users of myoelectric prostheses must recalibrate regularly. Given that aggregating previous data has been shown to improve robustness across sessions [51, 52], the diverse and representative data provided by CIIL may offer similar benefits. A longitudinal study could confirm this by comparing user-in-the-loop data with SGT data when generalizing across sessions, incorporating data aggregation, and applying transfer learning.

This work offers important insights into the use of simultaneity for regression-based myoelectric control. Truly independent, simultaneous, proportional control is arguably the greatest advantage of regression-based control schemes over classification, but this improved functionality can often become a burden due to undesired activations of other DoFs (i.e., action interference) [17, 27, 26, 34, 35]. One approach to improve the control of simultaneous motions has been to add simultaneous motions to the calibration data, but this has not always shown to be a benefit [26, 33]. For example, Smith et al. [26] suggested based on offline results that adding in combined DoFs would harm performance for isolated motions, and the results in this work corroborate this through online experiments based on the significantly in-

creased action interference of C-SGT compared to W-SGT. The results shown in Fig. 5.5 demonstrate that adding combined motions to the training set slightly improved throughput and completion rate, but also provided a less robust model based on the inferior path efficiency, overshoots, action interference, drift, and simultaneity gain. In particular, the action interference and simultaneity gain exhibited the largest differences when adding combined motions to the calibration data. These metrics exhibited largely the same pattern, with W-SGT and T-CIIL significantly outperforming C-SGT and O-CIIL. Generally, adding combined motions to the training set increased the simultaneity used by each participant, but the subsequent degradation of action interference and simultaneity gain indicates that this simultaneity was often undesired and uncontrollable. Interestingly, simultaneity gain is slightly worse for O-CIIL than C-SGT despite O-CIIL having a better path efficiency and Fig. 5.8 suggesting that the two models have similar simultaneity.

The results of this work also suggest that simply adding simultaneous motions to the calibration data may not provide a more robust controller, and, in this case, often did more harm than good. This is in contrast to the findings of Paskett et al. [33], who found that adding simultaneous motions to the calibration data helped overall performance, but only for certain models. Consequently, future work could continue to explore the role of the model in simultaneity and adaptation.

Bibliography

[1] E. Scheme and K. Englehart, "Electromyogram pattern recognition for control of powered upper-limb prostheses: state of the art and challenges for clinical use.," *Journal of Rehabilitation Research & Development*, vol. 48, no. 6, 2011.

[2] E. Eddy, E. J. Scheme, and S. Bateman, "A framework and call to action for

the future development of emg-based input in hci," in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–23, 2023.

[3] H. Huang, L. J. Hargrove, M. Ortiz-Catalan, and J. W. Sensinger, "Integrating upper-limb prostheses with the human body: Technology advances, readiness, and roles in human–prosthesis interaction," *Annual Review of Biomedical Engineering*, vol. 26, 2024.

[4] A. D. Roche, H. Rehbaum, D. Farina, and O. C. Aszmann, "Prosthetic myoelectric control strategies: a clinical perspective," *Current Surgery Reports*, vol. 2, pp. 1–11, 2014.

[5] D. J. Atkins, D. C. Y. Heard, and W. H. Donovan, "Epidemiologic Overview of Individuals with Upper-Limb Loss and Their Reported Research Priorities," *JPO: Journal of Prosthetics and Orthotics*, vol. 8, no. 1, p. 2, 1996.

[6] J. M. Hahne, M. A. Schweisfurth, M. Koppe, and D. Farina, "Simultaneous control of multiple functions of bionic hand prostheses: Performance and robustness in end users," *Science Robotics*, vol. 3, June 2018. Publisher: American Association for the Advancement of Science.

[7] S. Amsuess, I. Vujaklija, P. Goebel, A. D. Roche, B. Graimann, O. C. Aszmann, and D. Farina, "Context-Dependent Upper Limb Prosthesis Control for Natural and Robust Use," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 24, pp. 744–753, July 2016. Conference Name: IEEE Transactions on Neural Systems and Rehabilitation Engineering.

[8] O. C. Aszmann, I. Vujaklija, A. D. Roche, S. Salminger, M. Herceg, A. Sturma, L. A. Hruby, A. Pittermann, C. Hofer, S. Amsuess, *et al.*, "Elective amputation and bionic substitution restore functional hand use after critical soft tissue injuries," *Scientific Reports*, vol. 6, no. 1, p. 34960, 2016.

[9] H. E. Williams, A. W. Shehata, K. Y. Cheng, J. S. Hebert, and P. M. Pilarski, "Myoelectric prosthesis control using recurrent convolutional neural network regression mitigates the limb position effect," *bioRxiv*, pp. 2024–02, 2024.

[10] L. Schmalfuss, J. Hahne, D. Farina, M. Hewitt, A. Kogut, W. Doneit, M. Reischl, R. Rupp, and D. Liebetanz, "A hybrid auricular control system: direct, simultaneous, and proportional myoelectric control of two degrees of freedom in prosthetic hands*," *Journal of Neural Engineering*, vol. 15, p. 056028, Aug. 2018. Publisher: IOP Publishing.

[11] A. J. Young, L. H. Smith, E. J. Rouse, and L. J. Hargrove, "A new hierarchical approach for simultaneous control of multi-joint powered prostheses," in *2012 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, pp. 514–520, June 2012. ISSN: 2155-1782.

[12] A. J. Young, L. H. Smith, E. J. Rouse, and L. J. Hargrove, "A comparison of the real-time controllability of pattern recognition to conventional myoelectric control for discrete and simultaneous movements," *Journal of NeuroEngineering and Rehabilitation*, vol. 11, p. 5, Jan. 2014.

[13] J. M. Hahne, S. Dähne, H. J. Hwang, K. R. Müller, and L. C. Parra, "Concurrent Adaptation of Human and Machine Improves Simultaneous and Proportional Myoelectric Control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 23, pp. 618–627, jul 2015.

[14] A. Phinyomark, E. Campbell, and E. Scheme, "Surface Electromyography (EMG) Signal Processing, Classification, and Practical Considerations," in *Biomedical Signal Processing: Advances in Theory, Algorithms and Applications* (G. Naik, ed.), pp. 3–29, Singapore: Springer, 2020.

[15] E. Campbell, A. Phinyomark, and E. Scheme, "Deep Cross-User Models Reduce the Training Burden in Myoelectric Control," *Frontiers in Neuroscience*, vol. 15, May 2021. Publisher: Frontiers.

[16] C. Morrell, E. Campbell, and E. Scheme, "Exploring user compliance in the training of regression-based myoelectric control," *Myoelectric Controls Symposium*, Aug. 2024.

[17] A. Ameri, E. N. Kamavuako, E. J. Scheme, K. B. Englehart, and P. A. Parker, "Support vector regression for improved real-time, simultaneous myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 22, pp. 1198–1209, nov 2014.

[18] A. Ameri, E. J. Scheme, E. N. Kamavuako, K. B. Englehart, and P. A. Parker, "Real-time, simultaneous myoelectric control using force and position-based training paradigms," *IEEE Transactions on Biomedical Engineering*, vol. 61, pp. 279–287, 2 2014.

[19] G. Hajian, E. Campbell, M. Ansari, E. Morin, A. Etemad, K. Englehart, and E. Scheme, "Generalizing upper limb force modeling with transfer learning: A multimodal approach using EMG and IMU for new users and conditions," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2024.

[20] J. M. Hahne, F. Bießmann, N. Jiang, H. Rehbaum, D. Farina, F. C. Meinecke, K. R. Muller, and L. C. Parra, "Linear and nonlinear regression techniques for simultaneous and proportional myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 22, no. 2, pp. 269–279, 2014.

[21] S. Salter, R. Warren, C. Schlager, A. Spurr, S. Han, R. Bhasin, Y. Cai, P. Walkington, A. Bolarinwa, R. Wang, *et al.*, "emg2pose: A large and di-

verse benchmark for surface electromyographic hand pose estimation," *arXiv preprint arXiv:2412.02725*, 2024.

[22] T. N. Tully, C. J. Thomson, G. A. Clark, and J. A. George, "User-Specific Mirror Training Can Improve Myoelectric Prosthesis Control," *Myoelectric Controls Symposium*, Aug. 2022.

[23] T. N. Tully, C. J. Thomson, G. A. Clark, and J. A. George, "Validity and impact of methods for collecting training data for myoelectric prosthetic control algorithms," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2024.

[24] L. H. Smith, T. A. Kuiken, and L. J. Hargrove, "Use of probabilistic weights to enhance linear regression myoelectric control," *Journal of Neural Engineering*, vol. 12, p. 066030, nov 2015.

[25] L. H. Smith, T. A. Kuiken, and L. J. Hargrove, "Myoelectric Control System and Task-Specific Characteristics Affect Voluntary Use of Simultaneous Control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 24, pp. 109–116, Jan. 2016. Publisher: Institute of Electrical and Electronics Engineers Inc.

[26] L. H. Smith, T. A. Kuiken, and L. J. Hargrove, "Evaluation of linear regression simultaneous myoelectric control using intramuscular EMG," *IEEE Transactions on Biomedical Engineering*, vol. 63, pp. 737–746, Apr. 2016. Publisher: IEEE Computer Society.

[27] A. Ameri, M. A. Akhaee, E. Scheme, and K. Englehart, "Regression convolutional neural network for improved simultaneous emg control," *Journal of Neural Engineering*, vol. 16, p. 036015, 4 2019.

[28] A. T. Wang, C. D. Olsen, W. C. Hamrick, and J. A. George, "Correcting Temporal Inaccuracies in Labeled Training Data for Electromyographic Control Algorithms," *IEEE International Conference on Rehabilitation Robotics*, 2023. Publisher: IEEE Computer Society ISBN: 9798350342758.

[29] A. Hagengruber, U. Leipscher, B. M. Eskofier, and J. Vogel, "A New Labeling Approach for Proportional Electromyographic Control," *Sensors 2022, Vol. 22, Page 1368*, vol. 22, p. 1368, Feb. 2022. Publisher: Multidisciplinary Digital Publishing Institute.

[30] J. A. George, M. R. Brinton, C. C. Duncan, D. T. Hutchinson, and G. A. Clark, "Improved Training Paradigms and Motor-decode Algorithms: Results from Intact Individuals and a Recent Transradial Amputee with Prior Complex Regional Pain Syndrome," *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, vol. 2018-July, pp. 3782–3787, Oct. 2018. Publisher: Institute of Electrical and Electronics Engineers Inc. ISBN: 9781538636466.

[31] J. M. Hahne, H. Rehbaum, F. Biessmann, F. C. Meinecke, K. R. Muller, N. Jiang, D. Farina, and L. C. Parra, "Simultaneous and proportional control of 2D wrist movements with myoelectric signals," *IEEE International Workshop on Machine Learning for Signal Processing, MLSP*, 2012.

[32] J. Nielsen, S. Holmgaard, Ning Jiang, K. Englehart, D. Farina, and P. Parker, "Enhanced EMG signal processing for simultaneous and proportional myoelectric control," in *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, (Minneapolis, MN), pp. 4335–4338, IEEE, Sept. 2009.

[33] M. D. Paskett, M. R. Brinton, T. C. Hansen, J. A. George, T. S. Davis, C. C. Duncan, and G. A. Clark, "Activities of daily living with bionic arm improved by combination training and latching filter in prosthesis control comparison," *Journal of NeuroEngineering and Rehabilitation*, vol. 18, p. 45, Feb. 2021.

[34] J. Li, Z. Zhu, W. J. Boyd, C. Martinez-Luna, C. Dai, H. Wang, H. Wang, X. Huang, T. R. Farrell, and E. A. Clancy, "Virtual regression-based myoelectric hand-wrist prosthesis control and electrode site selection using no force feedback," *Biomedical Signal Processing and Control*, vol. 82, p. 104602, Apr. 2023.

[35] M. Nowak, T. Eiband, E. R. Ramírez, and C. Castellini, "Action interference in simultaneous and proportional myocontrol: comparing force- and electromyography," *Journal of Neural Engineering*, vol. 17, p. 026011, Mar. 2020. Publisher: IOP Publishing.

[36] E. Campbell, E. Eddy, X. Isabel, S. Bateman, B. Gosselin, U. Côté-Allard, and E. Scheme, "Screen guided training does not capture goal-oriented behaviours: Learning myoelectric control mappings from scratch using context informed incremental learning," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, pp. 1–1, 2024.

[37] C. Igual, J. Igual, J. M. Hahne, and L. C. Parra, "Adaptive auto-regressive proportional myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 27, no. 2, pp. 314–322, 2019.

[38] C. Igual, A. Castillo, and J. Igual, "An Interactive Training Model for Myoelectric Regression Control Based on Human–Machine Cooperative Performance," *Computers*, vol. 13, p. 29, Jan. 2024. Publisher: Multidisciplinary Digital Publishing Institute.

[39] C. Igual and J. Igual, "Simultaneous three-degrees-of-freedom prosthetic control based on linear regression and closed-loop training protocol," *Sensors*, vol. 24, no. 10, p. 3101, 2024.

[40] J. L. Betthauser, R. Greene, A. Dhawan, J. T. Krall, C. L. Hunt, G. Levay, R. R. Kaliki, M. S. Fifer, S. Sikdar, and N. V. Thakor, "Online Adaptation for Myographic Control of Natural Dexterous Hand and Finger Movements," Dec. 2024. arXiv:2412.17991 [cs].

[41] D. Yeung, D. Farina, and I. Vujaklija, "Directional Forgetting for Stable Co-Adaptation in Myoelectric Control," *Sensors*, vol. 19, p. 2203, Jan. 2019. Number: 9 Publisher: Multidisciplinary Digital Publishing Institute.

[42] D. Yeung, I. M. Guerra, I. Barner-Rasmussen, E. Siponen, D. Farina, and I. Vujaklija, "Co-Adaptive Control of Bionic Limbs via Unsupervised Adaptation of Muscle Synergies," *IEEE Transactions on Biomedical Engineering*, vol. 69, pp. 2581–2592, Aug. 2022. Conference Name: IEEE Transactions on Biomedical Engineering.

[43] A. H. Chou, M. Madduri, S. J. Li, J. Isa, A. Christensen, F. L. Hutchison, S. A. Burden, and A. L. Orsborn, "Using eye gaze to train an adaptive myoelectric interface," *bioRxiv*, pp. 2024–04, 2024.

[44] E. Campbell, E. Eddy, S. Bateman, U. Côté-Allard, and E. Scheme, "Context-informed incremental learning improves both the performance and resilience of myoelectric control," *Journal of NeuroEngineering and Rehabilitation*, vol. 21, no. 1, p. 70, 2024.

[45] E. Eddy, E. Campbell, S. Bateman, and E. Scheme, "Leveraging task-specific context to improve unsupervised adaptation for myoelectric control," in *2023*

IEEE International Conference on Systems, Man, and Cybernetics (SMC), pp. 4661–4666, 2023.

[46] E. Eddy, E. Campbell, A. Phinyomark, S. Bateman, and E. Scheme, "LibEMG: An Open Source Library to Facilitate the Exploration of Myoelectric Control," *IEEE Access*, vol. 11, pp. 87380–87397, 2023.

[47] R. N. Khushaba, S. Kodagoda, S. Lal, and G. Dissanayake, "Driver drowsiness classification using fuzzy wavelet-packet-based feature-extraction algorithm," *IEEE Transactions on Biomedical Engineering*, vol. 58, no. 1, pp. 121–131, 2010.

[48] I. S. MacKenzie, "Fitts' law," *The Wiley Handbook of Human Computer Interaction*, vol. 1, pp. 347–370, 2018.

[49] E. J. Scheme and K. B. Englehart, "Validation of a selective ensemble-based classification scheme for myoelectric control using a three-dimensional fitts' law test," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 21, no. 4, pp. 616–623, 2012.

[50] B. Maas, J. Tchimino, B. van Dijk, A. Murgia, C. K. van der Sluis, and R. M. Bongers, "Skill transfer of upper limb prosthesis control after training in a virtual reality environment," *Technology and Disability*, p. 10554181241311685, 2025.

[51] E. Eddy, E. Campbell, S. Bateman, and E. Scheme, "Understanding the influence of confounding factors in myoelectric control for discrete gesture recognition," *Journal of Neural Engineering*, vol. 21, no. 3, p. 036015, 2024.

[52] C. J. Thomson, T. N. Tully, E. S. Stone, C. B. Morrell, E. J. Scheme, D. J. Warren, D. T. Hutchinson, G. A. Clark, and J. A. George, "Enhancing neuro-

prosthesis calibration: the advantage of integrating prior training over exclusive use of new data," *Journal of Neural Engineering*, vol. 21, no. 6, p. 066020, 2024.

Chapter 6

Transformer-Based Context-Informed Incremental Learning with Soft-DTW Alignment Unlocks Fast and Precise Regression-Based Myoelectric Control

Evan Campbell, Ethan Eddy, Erik Scheme

6.1 Foreword

This chapter presents the next progression of regression-based context-informed incremental learning towards adopting temporal neural networks, and loss functions

that provide additional robustness to pseudo-label inaccuracies. This chapter simultaneously provides benefit to context-informed incremental learning and illuminates an unreported practical limitation of temporal models when deployed in closed-loop settings using standard calibration procedures. My contribution to this chapter spans ideation, experimental design, data collection, data analysis, manuscript preparation, and revision corrections.

6.2 Reference

Under Review with the IEEE Transactions on Neural Systems and Rehabilitation Engineering as:

Campbell, E., Eddy, E., Scheme, E., *Transformer-Based Context-Informed Incremental Learning with Soft-DTW Alignment Unlocks Fast and Precise Regression-Based Myoelectric Control*, IEEE Transactions on Neural Systems and Rehabilitation Engineering, *Under Review*

6.3 Abstract

Regression-based myoelectric interfaces hold the promise of enabling intuitive proportional and simultaneous control but remain limited by calibration sensitivity, unpredictable dynamics, and inconsistent user behaviours. Temporal neural architectures have the potential to substantially improve these controllers by capturing the temporal structure of user behaviours, provided they are trained using dynamics that are sufficiently representative of closed-loop use. Context-informed incremental learning (CIIL) offers a mechanism for acquiring such data online; however, its reliance on environment-derived pseudo-labels makes it vulnerable to temporal deviations between assumed and true user intent. This study introduces T-sDTW-CIIL, a transformer-based incremental learning framework that integrates temporal modelling, closed-loop learning, and soft dynamic time warping (sDTW) to enable tolerant label alignment. Twelve participants completed an adaptive regression-based cursor-control task using four pipelines: static and CIIL variants of both MLP and transformer models. T-sDTW-CIIL achieved significantly higher success rates, throughputs, efficiencies, and simultaneity gains when evaluated in a high precision ISO-Fitts' environment. T-sDTW-CIIL achieved throughputs of 2.0×, 2.4×, and 3.7× those of an MLP trained using conventional screen-guided training when acquiring large, medium, and small targets, respectively. Perhaps more importantly, it maintained success rates of 98.4% for small targets, whereas the static MLP degraded to only 23.4% success. T-sDTW-CIIL-based adaptation also reduced overall contraction intensities by ∼10%. These results demonstrate the powerful combination of temporal learning with context-informed co-adaptation. T-sDTW-CIIL overcomes key limitations of existing regression-based myoelectric controllers, enabling robust, low-intensity human–computer interaction.

6.4 Introduction

Surface electromyography (sEMG) has a long history as a neural interface for prosthesis control [1] and is becoming increasingly attractive as a general-purpose input for human-computer interaction [2, 3]. Early myoelectric control strategies relied on direct control, where the envelope of the EMG signal was mapped proportionally to an output [4]. Subsequent advances introduced pattern classification, enabling users to select among a larger set of physiologically appropriate motor classes [5, 6]. While proportionality was later incorporated into classification schemes by scaling outputs according to signal energy within each class, these methods remain limited to producing one motion class at a time without adopting specialized architectures [7]. To enable more natural movements, characterized by simultaneous and proportional control across multiple degrees of freedom, regression-based solutions have emerged as a promising alternative [8].

6.4.1 Shortcomings of Regression

While regression models have been shown to improve performance and usability in lab studies, they are still largely limited by their susceptibility to action interference (i.e., unintended activation of an extraneous degree of freedom [9, 10]) and drift (i.e., where the absence of an accessible deadband leads to unstable outputs [11]). These characteristics, which may lead to user frustration and reduce the perception of agency, arise, at least in part, from the open-loop screen-guided calibration procedures used to obtain the continuous-valued labels required to train regression models [12]. During these sessions, users often struggle to comply perfectly with prompts, exhibiting behaviours such as "all-or-nothing" activations or anticipatory responses, struggling to match both the timing (temporal errors) and intensity (spatial errors)

of the prescribed training targets [13, 11]. These misalignments, along with an absence of representative closed-loop control behaviours, degrade the quality of the training data and, consequently, the model's performance [11, 14, 15].

6.4.2 Incremental Learning for Closed-Loop Behaviours

Environment-dependent incremental learning, in which pseudo-labels are derived autonomously from the task environment during user-in-the-loop interactions, is emerging as a compelling strategy to overcome these calibration-related issues by establishing a framework around the user's natural goal-seeking behaviours [16, 17, 18, 19]. A compelling example is context-informed incremental learning (CIIL) [13, 20]. The closed-loop pseudo-label generation strategy of CIIL decreases the prevalence of spatial errors by inferring reasonable approximations for what the user should be doing at any point during a goal-seeking task. Further, by capturing training data directly from user-in-the-loop interactions, CIIL mitigates the out-of-distribution problem inherent in screen-guided calibration, aligning the distribution of data seen by the model during training with the behavioural patterns that actually arise during closed-loop control.

When applied to regression-based myoelectric control, CIIL has been shown to increase throughput by approximately 30%, while also reducing action interference by 30% and drift by 25% compared to traditionally calibrated models, despite equivalent training durations [21]. Users even consistently reported a subjective preference for incremental learning over conventional screen-guided calibration. However, different CIIL strategies have so far appeared to optimize distinct performance axes: one strategy produced superior precision (i.e., minimizing action interference and drift), while another achieved greater speed (maximizing throughput). Furthermore, although drift was improved, it was not eliminated, and stable control still

depended on maintaining a conservative maximum speed. In more demanding situations, where both high speed and high precision are required, existing approaches would likely fall short [22].

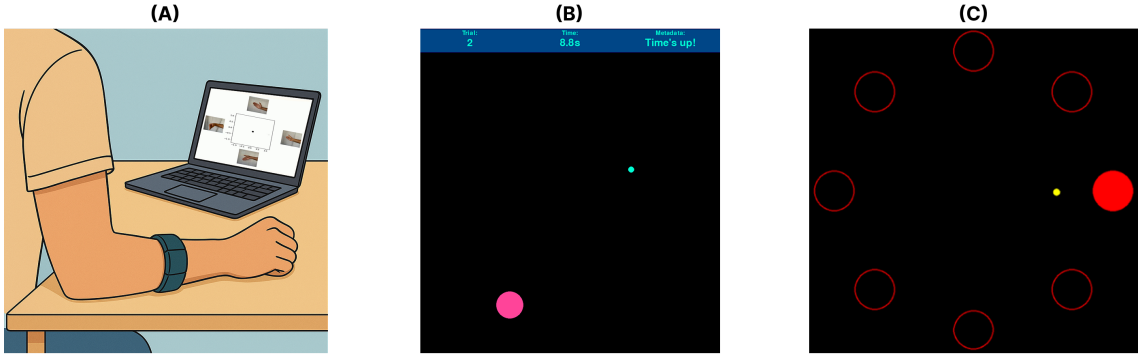


Figure 6.1: (A) Experimental setup and open-loop, screen-guided training interface with Cartesian animation displayed. (B) Closed-loop target acquisition environment used for user and model adaptation. (C) ISO 9241-9 Fitts' Law evaluation environment for assessing control performance. The largest of three target sizes is shown in this example.

6.4.3 Temporal Models for Temporal Dynamics

CIIL promotes the temporal dynamics that naturally arise during closed-loop control; however, previous versions have used static classifiers and deep networks that don't inherently model these dynamics. Consequently, it remains unclear whether they can be appropriately learned and integrated within the pseudo-labelling framework of CIIL, or if doing so could resolve the underlying precision–speed limitations of regression-based myoelectric control. Temporal neural architectures, such as recurrent neural networks (RNNs), gated recurrent units (GRUs), and transformers, are well-suited to capturing these temporal dependencies that underlie real-world motor behaviour [23]. However, this descriptive capacity introduces an important challenge: temporal models require a broader range of motion variability during calibration to adequately represent diverse temporal regimes. For example, a high-capacity temporal model trained with constant velocity oscillations may inadvertently learn to act as a narrow bandpass filter tuned to that frequency [24].

This limitation of temporal models, under-reported in the myoelectric control literature, can be profound when they are exposed to limited temporal variability during calibration. When trained using homogeneous temporal regimes, these models may appear accurate during offline analyses, but produce delayed or uncanny outputs online, reflecting epistemic uncertainty when faced with out-of-distribution inputs.

Nevertheless, numerous offline studies have reported substantial gains in intent recognition when temporal architectures are adequately exposed to diverse temporal dynamics [25]. Conversely, few studies have examined their behaviour during closed-loop control, and, to our knowledge, no incremental learning framework has explored whether temporal models can effectively assimilate the rich, variable temporal structure generated during user-in-the-loop adaptation. This gap motivates the integration of temporal architectures into environment-dependent incremental

learning approaches, but requires consideration of the assumptions of these approaches. For example, CIIL assumes a spatial approach strategy in goal-seeking behaviours, but has benefited significantly from incorporating tolerance to known variations in user behaviours [21]. It remains unclear, however, what impact temporal misalignment from assumed pseudo-labels may have when explicitly trying to learn temporal dynamics.

6.4.4 Soft Dynamic Time Warping for Pseudo-label Alignment

Soft dynamic time warping (sDTW) [26] provides a potential avenue for mitigating such temporal errors by enabling flexible alignment between intended and performed trajectories. Recent myoelectric work has even shown that sDTW can effectively compensate for temporal misalignments in an open-loop setting [27]. If similar benefits can be shown to facilitate temporal learning during a closed-loop setting, such as CIIL, the three approaches may offer complementary strengths. CIIL could resolve spatial inconsistencies arising from behavioural mismatch during goal-seeking activities, temporal models have the capacity to capture the dynamics arising in closed-loop control, while sDTW could compensate for variability in temporal execution. Integrating these three approaches effectively may therefore address both major dimensions of calibration error and out-of-distribution inputs.

6.4.5 Proposed Solution

This present work aims to advance the state-of-the-art in CIIL by capitalizing on the natural temporal variability that arises in user-in-the-loop settings. We introduce a novel framework, T-sDTW-CIIL, which combines a transformer-based architec-

ture with context-informed incremental learning (CIIL) and a soft dynamic time warping (sDTW) loss function. This design leverages the representational power of temporal models while simultaneously exploiting the behavioural diversity inherent to incremental learning. Moreover, the use of sDTW provides tolerance to deviations from ideal goal-seeking behaviour—such as reaction-time delays between trials or transient overshoots—relaxing many of CIIL's previous limitations. To evaluate the efficacy of this approach, we compare T-sDTW-CIIL against three baselines: (1) a non-temporal, non-incremental model (MLP-Static); (2) a temporal, non-incremental model (Transformer-Static); and (3) a non-temporal, incremental learning model (prior state of the art in CIIL: MLP-CIIL). The results show that, across three different precision constraints, T-sDTW-CIIL achieved the highest throughput and precision while enabling lower-effort, more stable control, demonstrating that temporal adaptation with differentiable alignment can overcome long-standing trade-offs in speed and precision.

6.5 Methods

6.5.1 Data Acquitision

Twelve participants (8 male, 4 female; age: mean 28.0 $\pm$ 5.8) were recruited for this study. All experimental procedures were conducted at the University of New Brunswick, with each participant completing a single session of approximately two hours. Surface EMG data were recorded using the OyMotion gForcePro+ armband, positioned on the right forearm at approximately two-thirds of the length proximal to the elbow. The device comprises eight bipolar recording channels and was configured to use 8-bit resolution and a 1000 Hz sampling rate. All participants provided written informed consent before participation, and the study protocol was approved

by the Research Ethics Board of the University of New Brunswick (REB 2022-122).

6.5.2 Experimental Design

The study employed a repeated-measures design in which each participant completed a common protocol under all four intent-recognition pipelines described in Section 6.5.3. To mitigate learning or order effects, the sequence of pipelines was balanced using a Latin Square arrangement. Each pipeline was divided into three phases: (1) a screen-guided training calibration phase, (2) a user-learning or co-adaptive phase, and (3) a final testing phase.

6.5.2.1 Screen-Guided Training

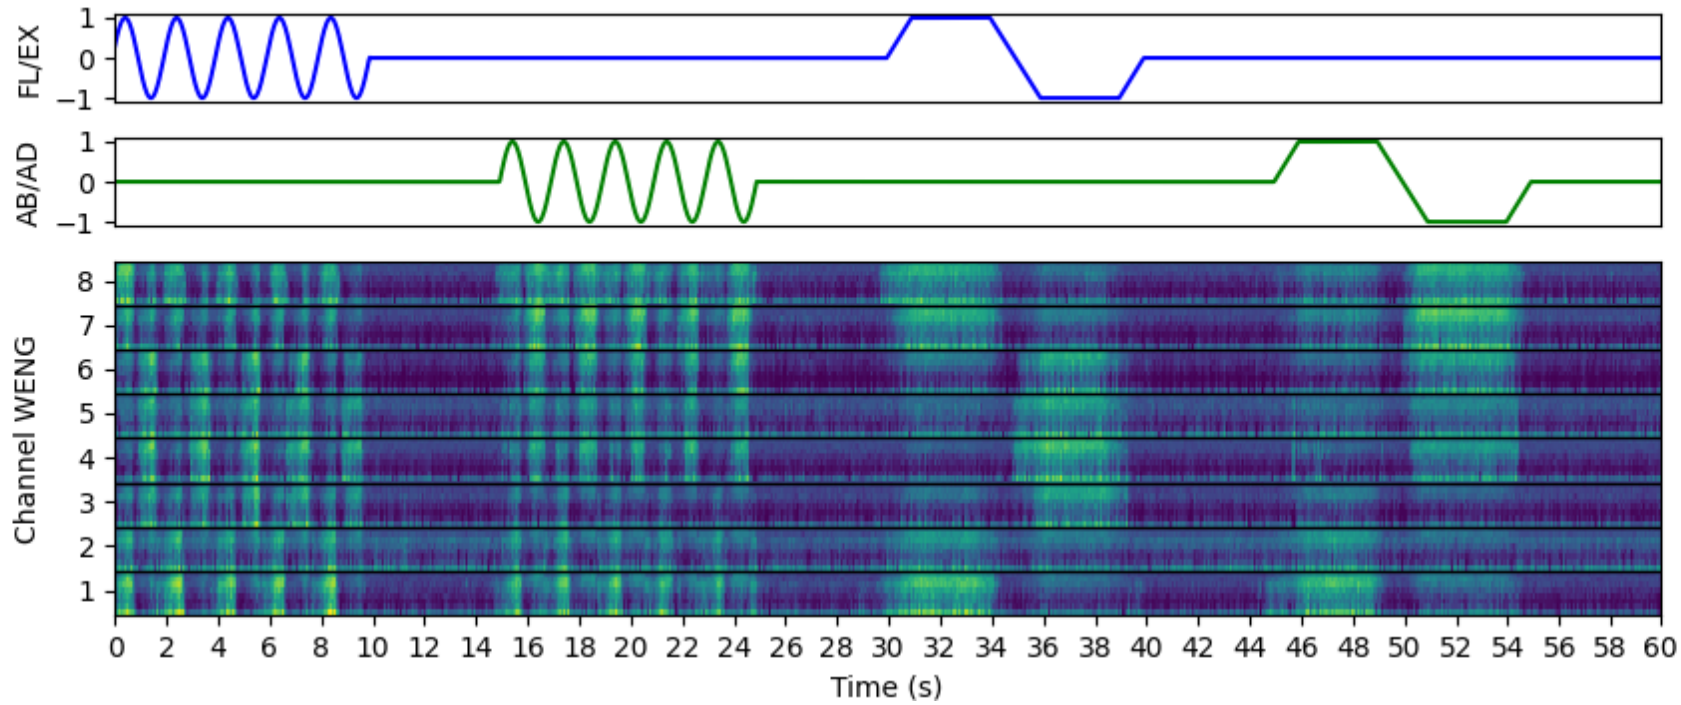


Figure 6.2: Visualization of one representative repetition of screen-guided training. The top two time series represent the two-degree-of-freedom (2-DoF; wrist flexion/extension and wrist abduction/adduction) choreography presented via the Cartesian animation. The heat maps (bottom) represent the corresponding EMG activation, as shown by the WENG feature obtained over this 60-second period across 8 channels.

For each pipeline, a conventional screen-guided training (SGT) session was conducted to establish an initial model for two-degree-of-freedom regression, wrist flexion/extension and wrist adduction/abduction, which were mapped to horizontal (left/right) and vertical (up/down) cursor movements. The choreography was presented visually using the Cartesian animation interface provided by LibEMG [28, 29] (Figure 6.1A)–where the cursor's position along the axes reflected the user's subjective force level of the DoFs (velocity-based control). The training choreography comprised only isolated motions within these degrees of freedom. Each trial contained two distinct temporal behaviours. First, a rhythmic template was collected by prompting participants to complete five oscillations between the extremes of a degree of freedom at 0.5 Hz, followed by a 5-second rest period. Second, a slower ramp template was collected to increase temporal variability: participants executed a one-second ramp to the extreme of the degree of freedom, held the contraction for three seconds, and then ramped down over one second, after which the sequence was mirrored for the opposite extreme and concluded with 5 seconds of rest. Each trial generated approximately one minute of data, and five trials were recorded per session to provide sufficient calibration for the initial model [21]. The labels and feature representation of the EMG signals were shown for one repetition of SGT in Figure 6.2.

6.5.2.2 User-/Co-Adaptation Phase

Following model calibration (Section 6.5.3.2), participants entered a user-learning/co-adaptive phase in which they controlled a cursor within a target acquisition environment (LibEMG AdaptiveFitts; Figure 6.1B). Each participant completed 80 target acquisition trials. On each trial, a circular target with a 45-pixel radius appeared at a random location within a 1000 $\times$ 1000-pixel workspace. Participants were given

10 seconds to move the cursor to the target and were required to maintain the cursor within the target for 2 seconds to complete the trial. This phase facilitated adaptation in two possible ways: in static pipelines, it enabled user learning of the control interface; in pipelines supporting incremental updates, it enabled co-adaptation, where both the user and model could adapt simultaneously.

In this environment, pseudo-labels representing the ideal goal-seeking behaviour were generated continuously using context-informed incremental learning (CIIL). The central premise of CIIL is to exploit environmental context—specifically, the relationship between the current task state and the model's ongoing predictions—to derive pseudo-labels that promote adaptation while minimizing the risk of introducing action interference. The pseudo-label generation process for regression tasks is described in full detail in [21]. Briefly, each component of the model's predicted output is compared to the ideal direction of motion required to achieve the task objective. When all components share the same sign, the pseudo-label adopts the direction of the ideal motion. If some components differ in sign, those components are zeroed out, providing an allowance for users moving towards the target in one DoF at a time. When all components disagree in sign, the sample at that time point is given a placeholder value (NaN), as no reliable inference can be made regarding the user's true intent given available context. The magnitude of each pseudo-label is then scaled according to the cursor's distance from the goal position, following the function defined in Equation 6.1. This mapping follows a similar motivation as its conception in [21]; however, it has been modified to have a baseline speed of 0.1. When the cursor was located within the target, the same logic was applied if it continued moving towards the target center; otherwise, a zero vector was assigned. These feature-pseudo-label pairs were continuously aggregated into a buffer managed by the LibEMG MemoryManager [29], which serves data to another LibEMG utility that ultimately performs incremental learning steps for the neural network.

$$\text{SCALE} = min(1, 0.1 + \frac{0.9}{1 + e^{-10(\frac{\text{distance}-\text{target size}}{300} - 0.5)}} \tag{6.1}$$

6.5.2.3 Evaluation Phase

In the final phase, performance metrics were collected for each pipeline to enable comparison. The evaluation followed the ISO 9241-9 layout [30], in which eight circular targets were arranged evenly along the circumference of a circle with radius 275 pixels (Figure 6.1C). As in the user-/co-learning phase, participants were required to dwell within each target for 2 seconds to complete the trial successfully. Each participant acquired 16 targets (two rounds of 8) at three different target sizes (45, 30, and 15 pixels) to assess model performance across varying indexes of difficulty (representing varying precision demands). Performance was quantified for each trial using four metrics: success rate, throughput, path efficiency, and simultaneity gain [21].

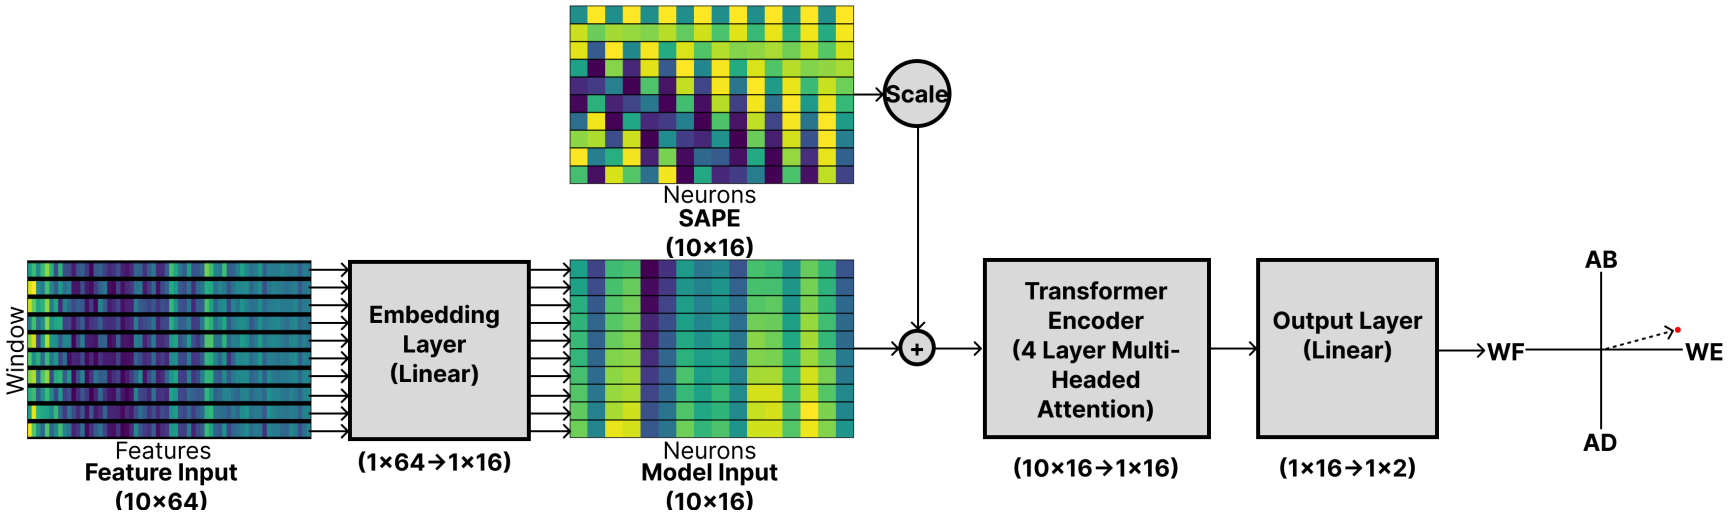


Figure 6.3: Visualization of the transformer pipeline. A horizon of 10 consecutive feature windows was independently embedded through a linear layer to 16 dimensions. Relative positioning information was injected over the horizon using SAPE. Subsequently, a 4-layer multi-headed attention network extracted temporal dynamics, which ultimately was processed by a single linear layer to render a 2-DoF prediction.

Throughput was computed to quantify a general information transfer rate based on Fitts' law (Equation 6.2), where D, W, T, and t_{dwell} are the optimal distance to complete the trial, the width of the target, the trial time upon completion, and the required dwell time, respectively. The dwell time was only subtracted for successful trials.

$$\text{throughput} = \frac{log_2(\frac{D}{W} + 1)}{T - t_{dwell}} \tag{6.2}$$

Path efficiency was computed to quantify a general accuracy metric within the user-in-the-loop phase (Equation 6.3), where c_t is the cursor position at time t.

$$\text{efficiency} = 100\% \times \frac{D}{\sum_{i=1}^{N-1} ||c_{i+1} - c_i||} \tag{6.3}$$

Simultaneity gain quantifies the controllability of a multi-DoF controller with respect to action interference, where % simultaneity is the percentage of predictions that had both degrees of freedom elicited more than some ϵ (1e-3 in this study).

$$\text{simultaneity gain} = \frac{\text{efficiency}}{\%\ \text{simultaneity}} \tag{6.4}$$

6.5.3 Intent Recognition Pipelines

The experiment evaluated four intent-recognition pipelines designed to isolate the contributions of temporal modelling and incremental adaptation. Together, these pipelines formed an ablation-style comparison, ranging from traditional non-adaptive approaches to the proposed framework. Pipeline 1 (MLP-Static) was a multilayer perceptron (MLP) trained offline without user-in-the-loop adaptation, representing a conventional SGT-trained myoelectric controller. Pipeline 2 (T-Static) employed

a transformer architecture, also trained offline, to assess the impact of introducing temporal modelling without adaptation. Pipeline 3 (MLP-CIIL) combined an MLP with context-informed incremental learning, reflecting prior work demonstrating the benefits of user-in-the-loop adaptation for regression-based myoelectric control [21]. Finally, Pipeline 4 (T-sDTW-CIIL) extended this framework by pairing a transformer with CIIL trained using an sDTW loss, thereby leveraging both temporal modelling and robust adaptation to advance the state of the art.

6.5.3.1 Model Architectures

Both neural architectures operated on features derived from the raw EMG signals through an overlapping windowing process. Windows of 100 ms were extracted with a 40 ms increment between consecutive segments. Within each window, the wavelet energy feature (WENG [31]) was computed, yielding eight frequency-banded energy measurements per channel across the eight bipolar EMG channels. This process produced a 64-dimensional representation of the EMG signal for each window. All features were subsequently normalized using a standard scaler fit to the screen-guided training data.

MLP Architecture: The MLP processed each windowed input independently, without explicit modelling of temporal dependencies. The architecture consisted of four sequential blocks, followed by a final output layer that mapped to two continuous output dimensions. Each block comprised a fully connected layer with 32 neurons, batch normalization, a rectified linear unit (ReLU) activation, and dropout ($p = 0.1$). This architecture served as the non-temporal baseline pipeline against which the temporal model was compared.

Transformer Architecture: The transformer encoder [32], shown in Figure 6.3, operated on sequences comprising the 10 most recent feature windows, corresponding

to a horizon of ∼0.5 seconds. At each time step, the 64-dimensional feature vector was projected into a 16-dimensional embedding through a fully connected layer. The resulting sequence of embeddings (10×16) was then augmented with sinusoidal absolute positional embeddings (SAPE) to provide relative temporal context [32]. This sequence was processed by a transformer encoder composed of four stacked layers, each containing a multi-head self-attention mechanism (four heads) and a feed-forward sub-layer with 16 hidden units. The output representation from the final time step was passed through a fully connected layer with two output neurons to generate the predicted degrees of freedom.

6.5.3.2 Initial Calibration

For the initial calibration of each model, data from all five repetitions of the SGT protocol were combined to form a single training dataset. To preserve the full extent of user-specific calibration data, no separate validation subset was withheld. Instead, the number of training epochs was fixed based on pilot analyses conducted with a larger dataset, in which validation loss trends were used to identify the point of diminishing returns. This analysis indicated that 125 epochs provided a stable balance between convergence and overfitting for both neural architectures. During calibration, the network parameters were optimized using the Adam optimizer (learning rate = 0.001) to minimize the mean-squared error between the predicted and true trajectory labels.

6.5.3.3 Incremental Learning

During the user-/co-learning phase, the MLP-CIIL and T-sDTW-CIIL models were incrementally updated using the LibEMG AdaptationManager [29], which interfaces directly with the continuously growing memory buffer maintained by the aforemen-

tioned MemoryManager. Upon initialization, the AdaptationManager was seeded with the SGT data to ensure stable early updates. Thereafter, it executed a repeating loop in which it first checked whether new closed-loop samples were available from the MemoryManager. When new data were detected, these feature–pseudo-label pairs were appended to the working buffer and used for a single-epoch update of the network weights. Both adaptive pipelines employed the same optimizer and learning rate as the initial calibration phase (Adam, learning rate = 0.001), though their loss functions differed. After each single-epoch update, the newly adapted model was automatically re-deployed within the intent-recognition pipeline, allowing the system to evolve continuously as the user interacted with the environment.

The MLP-CIIL model employed a mean-squared-error loss between the predicted 2-DoF output and the corresponding CIIL-generated pseudo-labels contained in the memory buffer. Samples whose pseudo-labels contained undefined values (NaN), corresponding to cases in which all model-predicted direction components conflicted entirely with the ideal direction of motion, were excluded from weight updates to prevent erroneous adaptation.

In contrast, the T-sDTW-CIIL model employed the soft dynamic time warping (sDTW) loss function, which extends classical dynamic time warping (DTW) to be fully differentiable (a requirement for backpropagation) [26]. Given a predicted trajectory $\hat{Y} = [\hat{y_1}, ...\hat{y_T}]$ and the pseudo-label trajectory from CIIL $Y^* = [y_1^*, ...y_S^*]$, where T and S were chosen to be 128 (reflecting an $\sim$ 5.5s time frame), a pairwise cost matrix is first computed using:

$$D_{ij} = ||\hat{y}_i - y_j^*||_2^2, \tag{6.5}$$

where any pseudo-labels with NaN values were given a distance of 100. A dynamic-

programming matrix $R \in \mathbb{R}^{(T+1)\times(S+1)}$ is then filled initially with

$$R_{0,0} = 0, \quad R_{i,0} = R_{0,j} = \infty, \tag{6.6}$$

then for each i, j

$$R_{ij} = D_{ij} + \text{softmin}_\gamma(R_{i-1,j}, R_{i,j-1}, R_{i-1,j-1}), \tag{6.7}$$

where the soft minimum operation replaces the non-differentiable min with a smooth log-sum-exponential form:

$$\text{softmin}_\gamma(a, b, c) = -\gamma \log(e^{-a/\gamma} + e^{-b/\gamma} + e^{-c/\gamma}). \tag{6.8}$$

This recursion aggregates all possible alignment paths between $\hat{Y}$ and Y^* with exponentially weighted contributions, controlled by the smoothing parameter γ. As γ approaches 0, the function converges to the hard minimum used in the standard DTW; as γ increases, alignments become softer and gradients smoother. This experiment used $\gamma = 0.1$, which is the recommended default value for sDTW. An example of the R matrix is given in Figure 6.4 between a simulated pathological "all or nothing" behaviour and a portion of the SGT protocol.

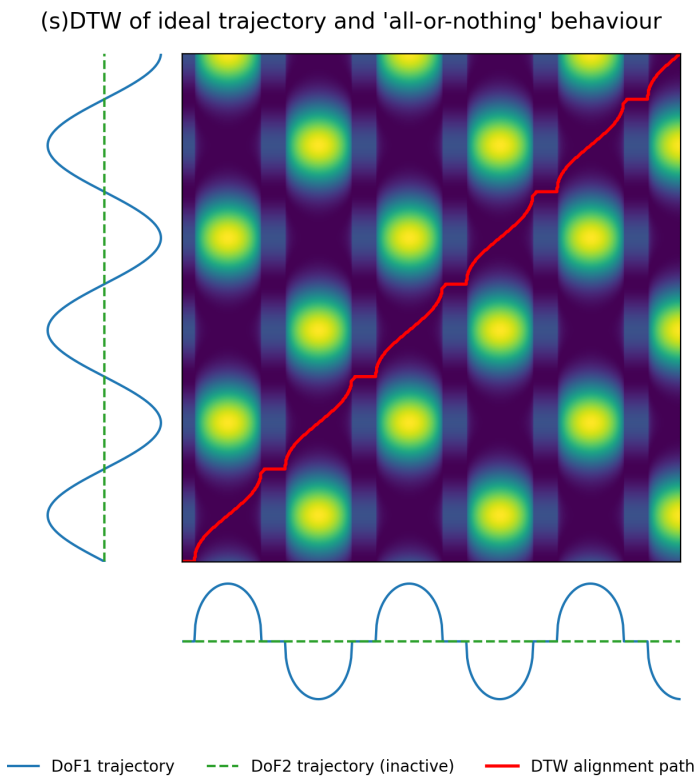


Figure 6.4: Visualization of R matrix of sDTW loss function. Blue and green lines represent the two degrees of freedom of the control space. An ideal label trajectory (obtained via screen-guided training) is shown on the vertical axis, and emulated predictions from a participant eliciting "all-or-nothing" behaviours (where they struggle with the continuous transition between directions of the DOF) are shown on the horizontal axis.

The sDTW loss for one sequence pair is then the value at $\{i = T, j = S\}$. This scalar is differentiable with respect to every entry in D, and its gradients are computed through a reverse dynamic-programming pass that distributes partial derivatives backwards along all alignment paths with exponentially decaying weights. In the context of incremental learning, this loss function allows the transformer to align its predicted and pseudo-labelled motion trajectories flexibly, tolerating temporal shifts and variable execution speeds while maintaining smooth, differentiable updates suitable for real-time adaptation.

6.5.4 Statistical Analyses

The four intent-recognition pipelines were compared using throughput, path efficiency, success rate, and simultaneity gain across the three target sizes of the testing phase. Performance values were aggregated to yield one data point per participant per target size for each intent recognition pipeline. After verifying performance metrics were normally distributed, a two-way repeated-measures analysis of variance (ANOVA) was conducted with pipeline (four levels) and target size (three levels) as within-subject factors, and subject as the repeated measure for each metric. This design tested both main effects and their interaction while accounting for inter-participant variability. When sphericity was violated, the Huynh–Feldt corrected p-value is reported. If significant, post hoc pairwise comparisons were conducted for the main effects of model type and target size, with Bonferroni adjustment for multiple testing. Interactions were examined by comparing model types within each target size and target sizes within each model type.

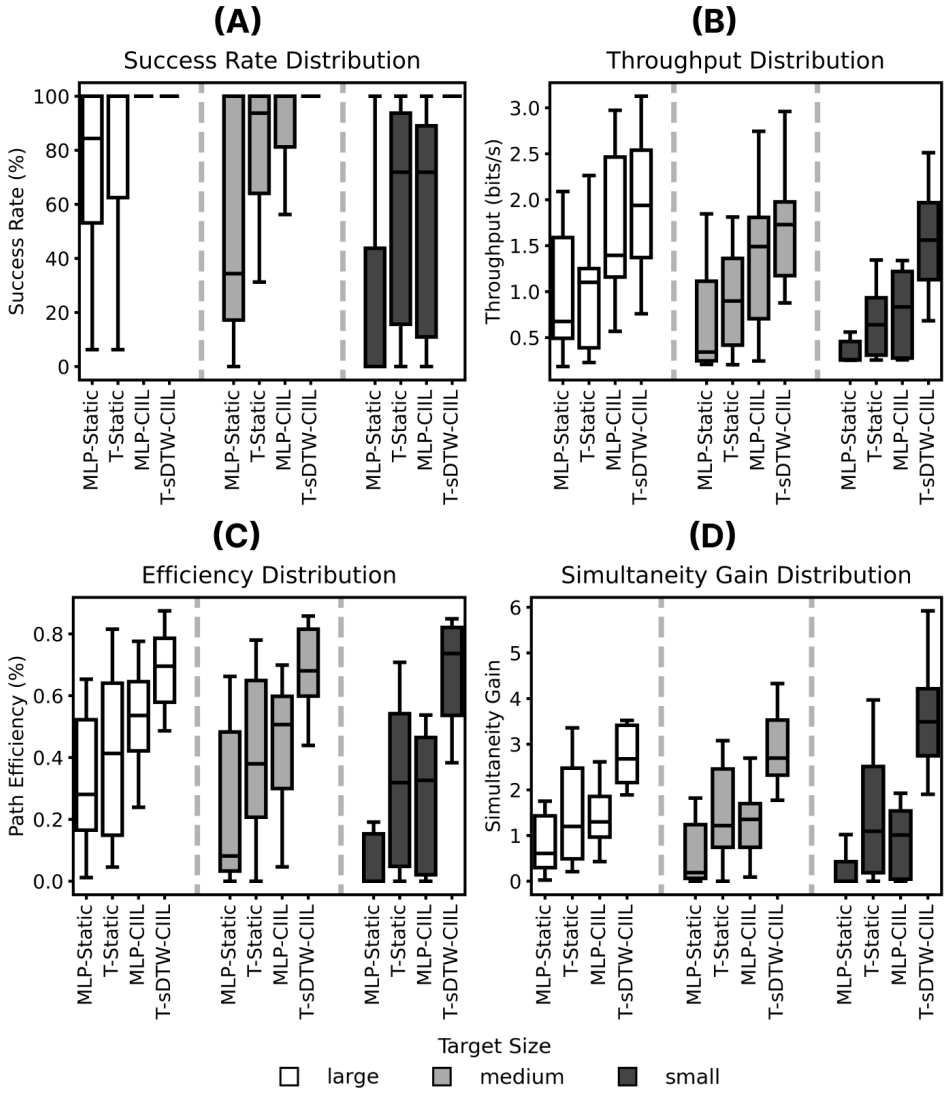


Figure 6.5: Success rate, throughput, efficiency, and simultaneity gain outcome measures from the evaluation phase presented as box plots. Each sample represents the average metric value obtained across the 16 target acquisitions at each target size per subject.

6.6 Results

The performance metrics across all subjects are shown in Figure 6.5. A series of repeated-measures ANOVAs was conducted on four dependent variables: success rate (Table 6.1), throughput (Table 6.2), efficiency (Table 6.3), and simultaneity gain (Table 6.4). Two within-subject factors were examined: model and target size.

6.6.1 Success Rate

All trials in which users successfully acquired the target within the 10-second time limit were classified as successful; all others were considered unsuccessful. The mean success rates across the four models are shown in Table 6.1. A repeated-measured ANOVA (Table 6.1) revealed a significant main effect of model ($p < .001$), target size ($p < .001$), and a significant model $\times$ target size interaction ($p < .001$).

Follow-up pairwise comparisons confirmed that the success rate differed significantly across models and target sizes. On average, across target sizes, T-sDTW-CIIL (99.1%) achieved a significantly higher success rate than MLP-Static (49.1%, $p = .0015$) and MLP-CIIL (80.6%, $p = .043$), while the difference from T-Static (68.9%) did not reach significance ($p = .088$). This nonsignificant difference was driven largely by the high between-subject variability of T-Static, which resulted in substantially wider confidence intervals despite a large difference in means. MLP-CIIL also outperformed MLP-Static ($p = .034$), but did not differ significantly from T-Static ($p = 1.000$).

Across target sizes, success rate decreased systematically with increasing task difficulty (large $\rightarrow$ medium $\rightarrow$ small; all $p < .002$). Model $\times$ target size interactions revealed that the performance of T-sDTW-CIIL persisted even for the smallest targets, where its mean success rate (98.4%) exceeded MLP-CIIL (57.8%, $p = .021$), T-Static (55.2%, $p = .020$), and MLP-Static (23.4%, $p < .001$). In contrast, the other three models all showed sharp declines in performance as target size decreased ($p < .03$ for pairwise large–small comparisons).

Table 6.1: Success Rates Rate (%) Across Different Target Sizes (top) with Repeated Measures ANOVA Results (bottom).

Model	Small (%)	Medium (%)	Large (%)
MLP-Static	23.4	49.5	74.5
MLP-CIIL	57.8	85.4	98.4
T-Static	55.2	74.0	77.6
T-sDTW-CIIL	**98.4**	**99.5**	**99.5**

Effect	DFn	DFd	F	p	η^2
(Intercept)	1	11	192.79	$< .001$	0.871
Model	3	33	10.25	$< .001$	0.286
Target Size	2	22	30.36	$< .001$	0.147
Model × Target Size	6	66	6.82	$< .001$	0.071

6.6.2 Throughput

The mean throughput across the four models is shown in Table 6.2. A repeated-measures ANOVA (Table 6.2) revealed a significant main effect of model ($p < .001$), target size ($p < .001$), and a significant model × target size interaction ($p < .001$).

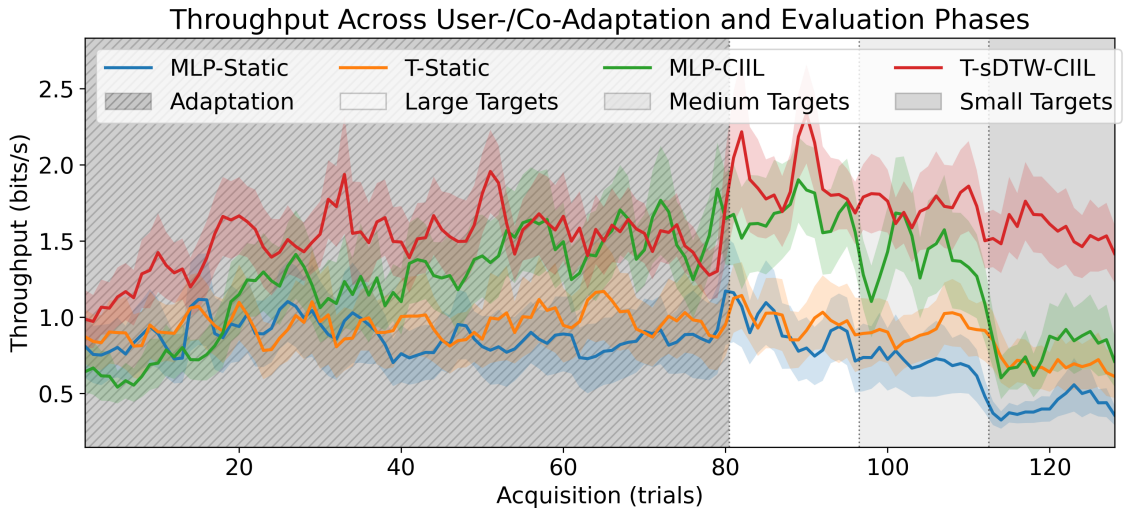


Figure 6.6: Throughput over User-/Co-Adaptation and Evaluation Phases across all subjects. The model's throughput was represented via the average throughput for each trial across subjects, where the shaded region represents the standard error.

Follow-up pairwise comparisons showed clear differences in throughput across models. When averaged over target sizes, T-sDTW-CIIL achieved the highest through-

put (1.75 bits/s), significantly outperforming MLP-Static ($p < .001$), T-Static ($p < .001$), and MLP-CIIL ($p = .016$). MLP-CIIL yielded significantly higher throughput than MLP-Static ($p = .018$) and T-Static ($p = .012$), while MLP-Static and T-Static did not differ ($p = .892$). Throughput also decreased systematically with increasing task difficulty (large $\rightarrow$ medium $\rightarrow$ small; all $p < .03$).

The model $\times$ target size interaction further revealed that T-sDTW-CIIL maintained its advantage at every precision level. For large targets, T-sDTW-CIIL outperformed MLP-Static and T-Static (both $p < .001$), but not MLP-CIIL ($p = 0.874$). This pattern held for medium targets (both MLP-Static and T-Static: $p < .001$; MLP-CIIL: $p = 0.249$) and was even more pronounced for small targets, where T-sDTW-CIIL exceeded MLP-CIIL ($p \ll .001$), T-Static ($p \ll .001$), and MLP-Static ($p \ll .001$). Although MLP-CIIL showed moderate improvements over MLP-Static across all target sizes ($p < .045$), its performance dropped sharply for the smallest targets, whereas T-sDTW-CIIL retained high throughput across all conditions.

Table 6.2: Throughput (bits/s) Across Different Target Sizes (top) with Repeated Measures ANOVA Results (bottom).

Model	Small (bit/s)	Medium (bit/s)	Large (bit/s)
MLP-Static	0.42	0.70	0.98
MLP-CIIL	0.77	1.38	1.67
T-Static	0.69	0.91	0.99
T-sDTW-CIIL	**1.57**	**1.73**	**1.94**

Effect	DFn	DFd	F	p	η^2
(Intercept)	1	11	55.13	$< .001$	0.787
Model	3	33	32.31	$< .001$	0.315
Target Size	2	22	25.32	$< .001$	0.119
Model $\times$ Target Size	6	66	6.64	$< .001$	0.027

6.6.3 Path Efficiency

The path efficiency metric quantified how closely the cursor trajectory followed the ideal path to the target. The mean efficiencies across the four models are shown in Table 6.3. A repeated-measures ANOVA (Table 6.3) revealed significant main effects of model ($p < .001$) and target size ($p < .001$), as well as a significant model $\times$ target size interaction ($p < .001$).

Follow-up pairwise comparisons indicated clear differences in path efficiency across models. Averaged over target sizes, T-sDTW-CIIL (67.9%) achieved substantially higher efficiency than MLP-CIIL (41.1%, $p < .001$), T-Static (37.3%, $p < .001$), and MLP-Static (21.8%, $p < .001$). MLP-CIIL also outperformed MLP-Static ($p = .035$), whereas its mean efficiency did not differ significantly from T-Static ($p = 1.000$), and MLP-Static and T-Static likewise did not differ ($p = .199$). Across target sizes, efficiency decreased systematically with increasing task difficulty (large $\rightarrow$ medium $\rightarrow$ small; all $p < .007$).

The model $\times$ target size interaction showed that T-sDTW-CIIL again maintained a clear advantage at every precision level. For large targets, its efficiency exceeded that of MLP-Static, T-Static, and MLP-CIIL (all $p < .001$), with a similar pattern for medium targets (all $p < .003$) and small targets (all $p < .001$). Within-model comparisons further revealed that efficiency for T-sDTW-CIIL remained stable across target sizes (no significant differences), whereas all three comparison models showed marked declines from large to small targets ($p < .01$ for large–small comparisons).

6.6.4 Simultaneity Gain

Simultaneity gain highlights the tradeoff between accessible simultaneity and action interference. The mean simultaneity gain across the four models is shown in Table

Table 6.3: Efficiency (%) Across Different Target Sizes (top)
with Repeated Measures ANOVA Results (bottom).

Model	Small (%)		Medium (%)		Large (%)	
MLP-Static	10.5		22.8		32.2	
MLP-CIIL	27.0		44.1		52.1	
T-Static	30.3		40.9		40.8	
T-sDTW-CIIL	**67.0**		**67.9**		**68.7**	
Effect	DFn	DFd	F	p	η^2	
(Intercept)	1	11	65.54	$< .001$	0.807	
Model	3	33	30.19	$< .001$	0.394	
Target Size	2	22	53.36	$< .001$	0.083	
Model × Target Size	6	66	9.28	$< .001$	0.035	

6.4. A repeated-measured ANOVA (Table 6.2) revealed a significant main effect of model ($p < .001$) and significant model × target size interaction ($p = .003$).

Follow-up pairwise comparisons showed that simultaneity gain differed across models. Averaged over target sizes, T-sDTW-CIIL achieved the highest simultaneity gain (3.43), significantly exceeding MLP-CIIL (1.18, $p < .001$), T-Static (1.45, $p < .001$), and MLP-Static (0.55, $p < .001$). MLP-CIIL also outperformed MLP-Static ($p = .034$), whereas its mean simultaneity gain did not differ significantly from T-Static ($p = 1.000$). T-Static further showed higher simultaneity gain than MLP-Static ($p = .046$). In contrast, there was no significant main effect of target size (all $p > .68$), indicating that, on average, simultaneity gain was relatively stable across large, medium, and small targets.

The significant model × target size interaction reflected the fact that target size influenced some models. At each target size, T-sDTW-CIIL produced substantially higher simultaneity gain than all other models (all $p < .006$), with the largest advantages observed for the smallest targets. Both MLP-CIIL and T-Static tended to achieve higher simultaneity gain than MLP-Static, particularly for large and small targets ($p < .048$), but remained far below T-sDTW-CIIL. Overall, T-sDTW-CIIL was uniquely able to combine high simultaneity with maintained path quality,

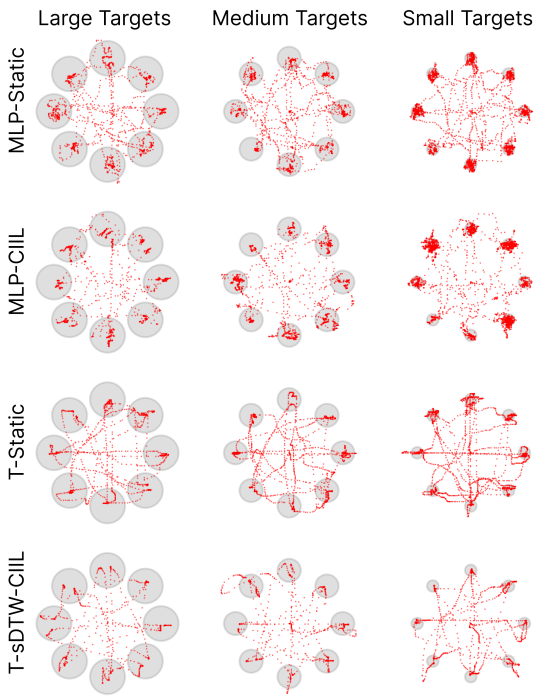


Figure 6.7: Paths traversed during the evaluation phase for one representative subject for all four models and three difficulties. Longer distances between neighbouring points represent higher speeds.

yielding the most favourable simultaneity gain profile across all task difficulties.

Table 6.4: Simultaneity Gain Across Different Target Sizes (top) with Repeated Measures ANOVA Results (bottom).

Model	Small	Medium	Large
MLP-Static	0.277	0.581	0.784
MLP-CIIL	0.871	1.26	1.42
T-Static	1.39	1.45	1.51
T-sDTW-CIIL	**3.96**	**3.05**	**3.28**

Effect	DFn	DFd	F	p	η^2
(Intercept)	1	11	44.78	$< .001$	0.691
Model	3	33	38.59	$< .001$	0.487
Target Size	2	22	0.96	.399	0.004
Model $\times$ Target Size	6	66	3.81	.003	0.045

6.7 Discussion

6.7.1 Speed-Precision Trade-off

Across the evaluation phase, T-sDTW-CIIL consistently outperformed the three other pipelines on measures that jointly represent speed, accuracy, precision, and control stability. Qualitatively, trajectory plots during the evaluation phase (Figure 6.7) show that T-sDTW-CIIL produced target-directed paths that were both rapid and well-contained within the target boundaries, whereas non-temporal or non-adaptive baselines more often overshot and required corrective motions. The time-series analysis of throughput over the closed-loop phases of the protocol (Figure 6.6) shows that T-sDTW-CIIL sustained its performance even as precision demands increased, while MLP-CIIL faltered at the smaller targets–consistent with precision-limited corrections and longer dwells to successfully acquire the targets. This is evidenced by the greater concentration of points within the small targets of Figure 6.7 for the MLP-based approaches. Together, these patterns indicate that temporal modelling plus soft alignment during online updates enables the controller to retain speed under tighter precision constraints, largely addressing the speed-precision limitations that were demonstrated in past CIIL variants.

6.7.2 CIIL with Temporal Insights Enable Lower Contraction Intensity

The intensity of contraction required for robust myoelectric control has long been a subject of debate, with lower intensities being favourable for minimizing fatigue and naturalness, but typically being associated with lower accuracy [33, 34]. Both in clinically-focused prosthesis control and emerging applications in consumer-grade

human-computer interaction, low-effort operation is desirable. Participants reported that T-sDTW-CIIL encouraged more subtle, less forceful contractions over the course of co-adaptation relative to other models. This is corroborated by the EMG statistics, with the distributions of windowed mean absolute value (MAV) shifting downwards for T-sDTW-CIIL over the course of adaptation, corresponding to a $\sim$10-percent reduction in MAV across the cohort (Figure 6.8). This outcome aligns with the method's design from a theoretical perspective: sDTW tolerates small temporal misalignments in pseudo-labels, allowing the model to learn stable mappings without users "spiking" their activations to stay proportional to calibration templates obtained through screen-guided training. Practically, these lower-intensity control inputs would likely reduce fatigue and can broaden the applications where myoelectric control feels natural (e.g., prolonged desktop, mobile, or AR/VR use), and the maintained throughput on small targets suggests that neither speed nor precision was traded away to achieve this improvement in comfort.

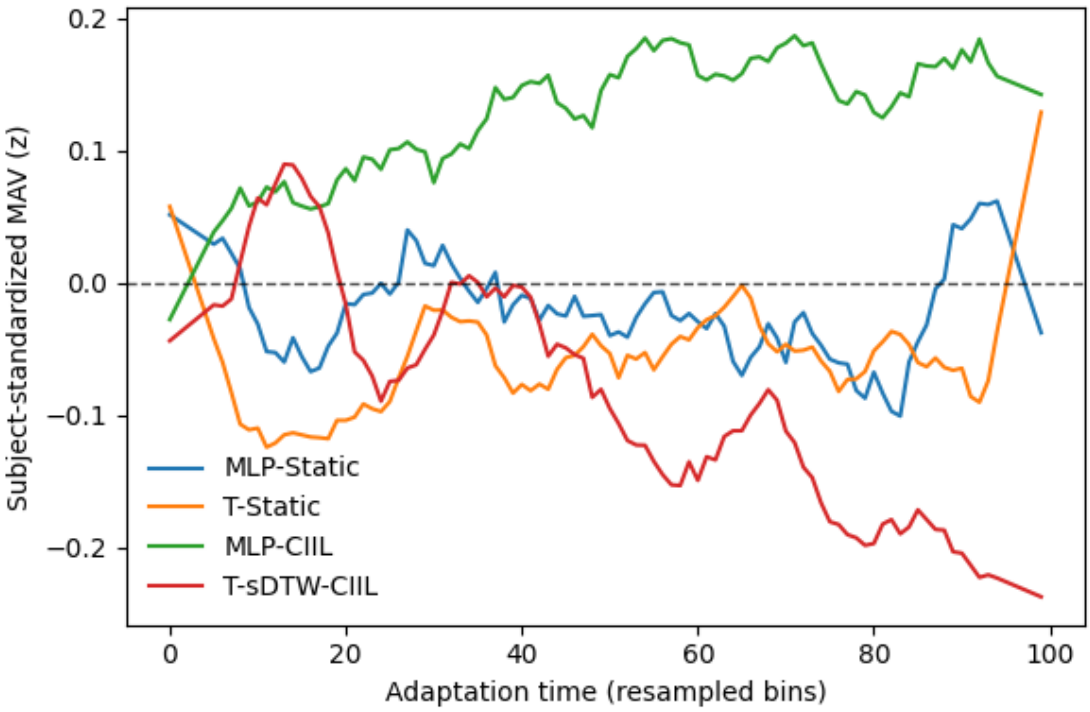


Figure 6.8: Trajectories of mean absolute values (MAV) from windows observed within the user-/co-adaptation phase of four models. MAVs were transformed into subject-specific z-scores with statistics computed from the data pooled from all models. User-/Co-adaptation phase time was not uniform between models or subjects, so each pass through was normalized to 100 temporal bins. A horizontal dashed line at 0 depicts the mean intensity of the subjects across all models. A black line shows the average MAV for time segments across the adaptation phase.

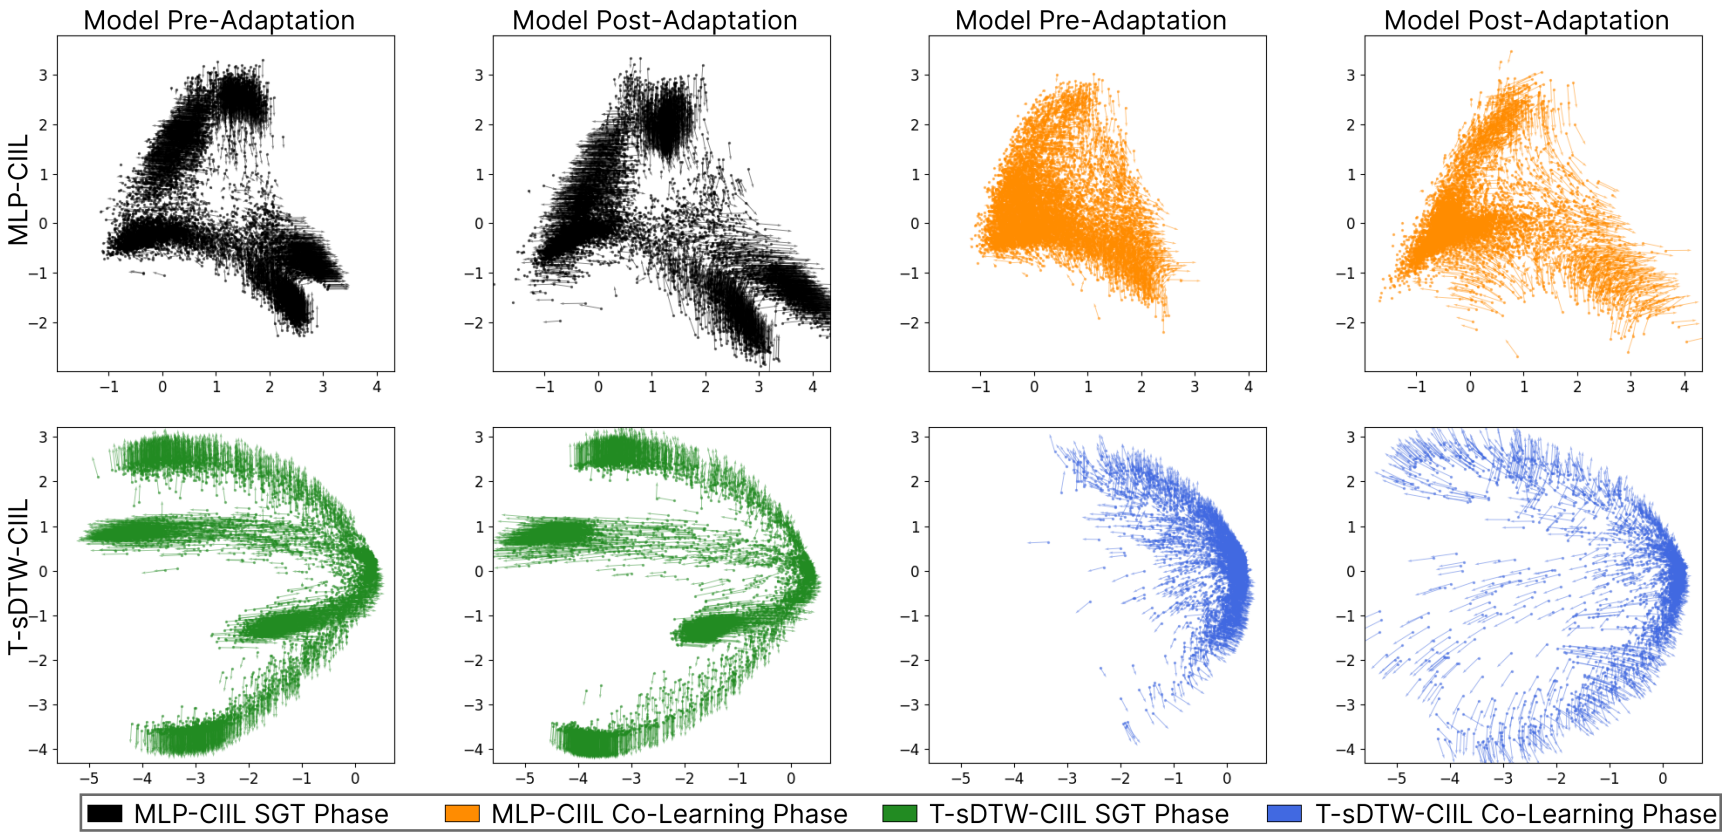


Figure 6.9: Quiver plots visualizing the latent space of MLP-CIIL and T-sDTW-CIIL models from the same representative subject. Each latent space was taken from the penultimate layer output of the corresponding model. The direction and length of the arrows depict the prediction of the model for that input. Plots that are the same colour represent the same data processed by different models (i.e, column 2 represents looking back on SGT data after the model was adapted, and column 3 represents looking ahead to co-learning phase data without having performed adaptation).

6.7.3 Closed-Loop Performance of Transformers

The superiority of the transformer when trained with CIIL compared to its static counterpart provides an important insight into the behaviour of temporal models when trained for myoelectric control that has not been previously reported. While transformers have shown remarkable performance in offline EMG analyses [35, 36, 37], the literature has lacked evidence of this performance being corroborated in online, user-in-the-loop settings. When trained with conventional short, single-session calibration protocols, the data contain only a narrow range of temporal regimes and do not expose the model to the user's natural temporal variability. As a result, conventional transformer controllers can tend to exhibit unstable or latency-laden outputs (feeling similar to a large window majority vote), reflecting epistemic uncertainty when faced with input data that falls outside the narrow calibration distribution. In this experiment, although the initial calibration of the transformer models resulted in remarkable R^2 values compared to the MLP models ($\sim$0.97 vs. $\sim$0.89, respectively), the closed-loop performance of T-Static was similar to MLP-Static (Figure 6.6 and pairwise comparison of $p = .892$ for throughput across difficulties). In contrast, the closed-loop co-adaptation used here exposed the transformer to the full spectrum of user-specific and task-relevant timing variations encountered during real use. Context-informed pseudo-labelling provides alignment-consistent supervisory signals, while the sDTW objective ensures that the model learns to associate temporally shifted but semantically equivalent behaviours. This combination effectively transforms the short initial calibration into an ongoing curriculum, allowing the transformer to specialize online without overfitting to limited temporal patterns. Consequently, the T-sDTW-CIIL model—supported by closed-loop adaptation that incorporates the typically out-of-distribution data that transformers struggle with achieved the highest throughput and precision of all pipelines,

demonstrating that temporal architectures can reach their potential when coupled with context-aware, differentiable alignment during incremental learning.

6.7.4 Model Latent Analysis

The latent-space quiver plots shown in Figure 6.9 provide additional insight into how the different pipelines encode and transform EMG activity into control commands. In these plots, each vector indicates the model's predicted control direction given an EMG feature embedding, revealing the topology of the model's decision space. The non-temporal MLP-based models exhibited fragmented, disjoint, orthogonal vector fields, with small regions of conflicting directionality and discontinuities near the decision boundaries—patterns that correspond to the erratic cursor oscillations and action interference that were observed. Introducing CIIL (MLP-CIIL) reduced some of this fragmentation by constraining updates through environment-aware pseudo-labels, resulting in a smoother transition between local regions, but residual discontinuities remained, particularly in regions corresponding to transitional muscle activations or situations where inputs for similar pseudo-labels were less homogeneous, like when inside the target. In contrast, the latent field produced by the transformer trained under the sDTW objective (T-sDTW-CIIL) exhibited smooth, continuous flow patterns that curved coherently around the origin of the latent manifold, mirroring the natural coupling between wrist flexion–extension and radial–ulnar deviation, which rely on different muscle synergies that can co-occur. Rather than forming disjoint clusters, the predicted vectors varied gradually across the latent space, suggesting that temporally adjacent EMG patterns were mapped to incrementally shifting outputs. This organization implies that the model learned a low-curvature control manifold, where small changes in muscle activation translate to proportionally small changes in control output—a hallmark of stable, high-fidelity

continuous control [38]. The smoothness of this latent flow also indicates that the transformer's temporal self-attention and the differentiable alignment introduced by sDTW together fostered representations that respect the inherent temporal and biomechanical continuity of movement.

6.7.5 Future Work

While the present findings demonstrate clear advantages of integrating temporal modelling and sequence alignment strategies into incremental myoelectric control, several limitations warrant consideration.

Although the adaptation protocol was conducted during a real target-acquisition control task, the evaluation was limited to a laboratory environment with a fixed task structure. Generalization to more realistic scenarios, such as more unstructured tasks (e.g., smart home), or more nuanced tasks (e.g., reach-to-grab), remains to be verified. Other work has explored CIIL within simple virtual reality environments [39]; however, deployment in more challenging scenarios is warranted.

The calibration and adaptation sessions were of relatively short duration (5 minutes and ∼5-10 minutes, respectively) because of constraints on the single-session experiment. This constrains conclusions regarding long-term stability for prolonged adaptation periods and retention over separate sessions of use. Continuous co-adaptation over days or weeks may reveal additional dynamics in both user behaviour and model drift, potentially uncovering new challenges (e.g. management of large volumes of data) or presenting an even greater opportunity for improving reliability [40].

However, just as providing an increasing amount of homogeneous repetitions during screen guided training would provide diminishing returns, it should also be considered that longer incremental learning within a static environment (that does not

modulate its difficulty) could also provide diminishing returns. Within this study, after only 10 minutes of incremental learning, the throughput of the incremental learning models had begun to plateau. Consequently, to produce a higher-speed and more precise model, future work should incentive users to push these boundaries via an environment with embedded curriculum learning [41].

Although sDTW introduced substantial robustness to temporal misalignment, it incurs a higher computational cost than conventional pointwise losses like mean squared error. In this experiment, where a desktop computer was used for live model adaptation this computational cost was acceptable. sDTW and mean squared error took about 1-second and 0.2-seconds per epoch for each minute of data in the AdaptationManager, respectively; both strategies gave the impression that adaptation was continuous even when periods between model updates reached their maximum ($\sim$10 seconds). Nonetheless, performing these updates on low-compute, wearable microprocessors may prove challenging when the stored data accrues. There are approaches to minimize computational cost for sDTW, like Sakoe-Chiba [42], which minimizes the off-diagonal terms in the dynamic programming process; however, the consequences of limiting the alignment path and the effectiveness of this shortcut for myoelectric control remain untested. Further, data prioritization strategies within LibEMG's AdaptationManager could keep computational costs low through leveraging the streaming data just once, without storing, in a manner that prioritizes recent emergent characteristics [43].

Finally, this study offered an initial examination of transformer architectures, which remain largely unexplored in online myoelectric control. Future work should investigate alternative model configurations, including variations in network depth and the number of attention heads. Moreover, recent success with conformer architectures in myoelectric handwriting recognition [3] suggests that related transformer-based

models may hold further promise for adaptive, user-in-the-loop control.

6.8 Conclusion

This work introduced T-sDTW-CIIL, a novel framework that combines temporal neural networks, a temporal-error-tolerant loss function, and closed-loop environment-dependent incremental learning to advance continuous myoelectric control. To our knowledge, this is the first online, user-in-the-loop experiment to evaluate transformers for continuous control, illuminating the limitations of standard calibration procedures that lack temporal variability, and providing a solution that simultaneously optimizes precision and speed. Beyond improving the success rate, throughput, path efficiency, and simultaneity gain by large margins across varying precision requirements, the adaptation of temporal models led to a reduction in contraction intensity during closed-loop use. Finally, analysis of the latent space revealed that T-sDTW-CIIL developed a smooth, coherent manifold after adaptation, indicating that the model had effectively integrated the temporal dynamics observed into a stable, well-organized representation. Together, these findings highlight the benefit of temporal modelling and differentiable alignment in incremental learning pipelines and, more generally, myoelectric control.

Bibliography

[1] D. S. Childress, "Historical aspects of powered limb prostheses," *Clin Prosthet Orthot*, vol. 9, no. 1, pp. 2–13, 1985.

[2] E. Eddy, E. J. Scheme, and S. Bateman, "A framework and call to action for the future development of emg-based input in hci," in *Proceedings of the 2023*

CHI conference on human factors in computing systems, pp. 1–23, 2023.

[3] P. Kaifosh, T. R. Reardon, and C.-l. at Reality Labs, "A generic non-invasive neuromotor interface for human-computer interaction," *Nature*, Jul 2025.

[4] A. Kobrinski, S. Bolkhovitin, L. Voskoboinikova, D. Ioffe, E. Polyan, B. Popov, Y. L. Slavutski, A. Y. Sysin, and Y. S. Yakobson, "Problems of bioelectric control," *IFAC Proceedings Volumes*, vol. 1, no. 1, pp. 629–633, 1960.

[5] F. R. Finley and R. W. Wirta, "A correlational study of myopotential response and force of muscle contraction during varying activity demands," tech. rep., 1965.

[6] P. Herberts, C. Almström, R. Kadefors, and P. D. Lawrence, "Hand prosthesis control via myoelectric patterns," *Acta Orthopaedica Scandinavica*, vol. 44, no. 4-5, pp. 389–409, 1973.

[7] L. J. Hargrove, E. J. Scheme, K. B. Englehart, and B. S. Hudgins, "Multiple binary classifications via linear discriminant analysis for improved controllability of a powered prosthesis," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 18, no. 1, pp. 49–57, 2010.

[8] J. M. Hahne, F. Biessmann, N. Jiang, H. Rehbaum, D. Farina, F. C. Meinecke, K.-R. Müller, and L. C. Parra, "Linear and nonlinear regression techniques for simultaneous and proportional myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 22, no. 2, pp. 269–279, 2014.

[9] M. Nowak, T. Eiband, E. R. Ramírez, and C. Castellini, "Action interference in simultaneous and proportional myocontrol: Comparing force-and electromyography," *Journal of neural engineering*, vol. 17, no. 2, p. 026011, 2020.

[10] L. H. Smith, T. A. Kuiken, and L. J. Hargrove, "Evaluation of linear regression simultaneous myoelectric control using intramuscular emg," *IEEE Transactions on Biomedical Engineering*, vol. 63, no. 4, pp. 737–746, 2015.

[11] C. Morrell, E. Campbell, and E. Scheme, "Exploring user compliance in the training of regression-based myoelectric control," in *Myoelectric Controls Symposium*, vol. 33, 2024.

[12] C. Igual, A. Castillo, and J. Igual, "An interactive training model for myoelectric regression control based on human–machine cooperative performance," *Computers*, vol. 13, no. 1, p. 29, 2024.

[13] E. Campbell, E. Eddy, X. Isabel, S. Bateman, B. Gosselin, U. Cote-Allard, and E. Scheme, "Screen guided training does not capture goal-oriented behaviours: Learning myoelectric control mappings from scratch using context informed incremental learning," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2024.

[14] A. T. Wang, C. D. Olsen, W. C. Hamrick, and J. A. George, "Correcting Temporal Inaccuracies in Labeled Training Data for Electromyographic Control Algorithms," *IEEE International Conference on Rehabilitation Robotics*, 2023. Publisher: IEEE Computer Society ISBN: 9798350342758.

[15] T. N. Tully, C. J. Thomson, G. A. Clark, and J. A. George, "Validity and impact of methods for collecting training data for myoelectric prosthetic control algorithms," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2024.

[16] E. Campbell, F. Egle, M. Oßwald, U. Cote-Allard, P. M. Pilarski, N. Boccardo, R. Meattini, I. Vujaklija, L. Hargrove, M. Canepa, *et al.*, "(un) supervised (co) adaptation via incremental learning for myoelectric control: Motivation,

review, and future directions," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2025.

[17] A. H. Chou, M. Madduri, S. J. Li, J. Isa, A. Christensen, F. Hutchison, S. A. Burden, and A. L. Orsborn, "Using eye gaze to train an adaptive myoelectric interface," *bioRxiv*, pp. 2024–04, 2024.

[18] M. M. Madduri, M. Yamagami, S. J. Li, S. Burckhardt, S. A. Burden, and A. L. Orsborn, "Predicting and shaping human-machine interactions in closed-loop, co-adaptive neural interfaces," *bioRxiv*, pp. 2024–05, 2024.

[19] R. B. Woodward and L. J. Hargrove, "Adapting myoelectric control in real-time using a virtual environment," *Journal of neuroengineering and rehabilitation*, vol. 16, no. 1, p. 11, 2019.

[20] E. Campbell, E. Eddy, S. Bateman, U. Côté-Allard, and E. Scheme, "Context-informed incremental learning improves both the performance and resilience of myoelectric control," *Journal of NeuroEngineering and Rehabilitation*, vol. 21, no. 1, p. 70, 2024.

[21] C. Morrell, E. Campbell, E. Eddy, and E. Scheme, "Context-informed incremental learning improves throughput and reduces drift in regression-based myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2025.

[22] R. Meredith, E. Eddy, S. Bateman, and E. Scheme, "Comparing online wrist and forearm emg-based control using a rhythm game-inspired evaluation environment," *Journal of Neural Engineering*, vol. 21, no. 4, p. 046057, 2024.

[23] C. Lin, C. Zhao, J. Zhang, C. Chen, N. Jiang, D. Farina, and W. Guo, "Continuous estimation of hand kinematics from electromyographic signals based on

power-and time-efficient transformer deep learning network," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2024.

[24] K. Yi, J. Fei, Q. Zhang, H. He, S. Hao, D. Lian, and W. Fan, "Filternet: Harnessing frequency filters for time series forecasting," *Advances in Neural Information Processing Systems*, vol. 37, pp. 55115–55140, 2024.

[25] T. N. Tully, A. E. Nelson, and J. A. George, "Training movement velocity significantly affects the performance of myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 33, pp. 3784–3792, 2025.

[26] M. Cuturi and M. Blondel, "Soft-dtw: a differentiable loss function for time-series," in *International conference on machine learning*, pp. 894–903, PMLR, 2017.

[27] R. Meattini, A. Bernardini, G. Palli, and C. Melchiorri, "semg-based minimally supervised regression using soft-dtw neural networks for robot hand grasping control," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10144–10151, 2022.

[28] E. Eddy, E. Campbell, A. Phinyomark, S. Bateman, and E. Scheme, "Libemg: an open source library to facilitate the exploration of myoelectric control," *IEEE Access*, vol. 11, pp. 87380–87397, 2023.

[29] E. Eddy, E. Campbell, C. Morrell, and E. Scheme, "Libemg documentation." `https://libemg.github.io/libemg/`, 2025. GitHub repository; Branch compose_emgpredictor; Accessed: Oct. 19, 2025.

[30] I. S. MacKenzie, "Fitts' law," *The wiley handbook of human computer interaction*, vol. 1, pp. 347–370, 2018.

[31] R. N. Khushaba, S. Kodagoda, S. Lal, and G. Dissanayake, "Driver drowsiness classification using fuzzy wavelet-packet-based feature-extraction algorithm," *IEEE transactions on biomedical engineering*, vol. 58, no. 1, pp. 121–131, 2010.

[32] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.

[33] E. Scheme, B. Lock, L. Hargrove, W. Hill, U. Kuruganti, and K. Englehart, "Motion normalized proportional control for improved pattern recognition-based myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 22, no. 1, pp. 149–157, 2013.

[34] E. Campbell, A. Phinyomark, and E. Scheme, "Current trends and confounding factors in myoelectric control: Limb position and contraction intensity," *Sensors*, vol. 20, no. 6, p. 1613, 2020.

[35] X. Li, X. Zhang, L. Zhang, X. Chen, and P. Zhou, "A transformer-based multi-task learning framework for myoelectric pattern recognition supporting muscle force estimation," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 31, pp. 3255–3264, 2023.

[36] S. T. P. Raghu, H. E. Williams, and E. Scheme, "Efficient multi-positioned training for regression-based myoelectric control: Exploring transformer models with transfer learning," in *2025 International Conference On Rehabilitation Robotics (ICORR)*, pp. 1597–1603, IEEE, 2025.

[37] W. Zhang, T. Zhao, J. Zhang, and Y. Wang, "Lst-emg-net: Long short-term transformer feature fusion network for semg gesture recognition," *Frontiers in Neurorobotics*, vol. 17, p. 1127338, 2023.

[38] H. Li, Z. Xu, G. Taylor, C. Studer, and T. Goldstein, "Visualizing the loss landscape of neural nets," *Advances in neural information processing systems*, vol. 31, 2018.

[39] G. Gagné, A. Azad, T. Labbé, E. Campbell, X. Isabel, E. Scheme, U. Côté-Allard, and B. Gosselin, "Context informed incremental learning improves myoelectric control performance in virtual reality object manipulation tasks," *arXiv preprint arXiv:2505.06064*, 2025.

[40] A. Waris, I. Mendez, K. Englehart, W. Jensen, and E. N. Kamavuako, "On the robustness of real-time myoelectric control investigations: A multiday fitts' law approach," *Journal of neural engineering*, vol. 16, no. 2, p. 026003, 2019.

[41] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48, 2009.

[42] H. Sakoe and S. Chiba, "Dynamic programming algorithm optimization for spoken word recognition," *IEEE transactions on acoustics, speech, and signal processing*, vol. 26, no. 1, pp. 43–49, 2003.

[43] D. Yeung, D. Farina, and I. Vujaklija, "Directional forgetting for stable co-adaptation in myoelectric control," *Sensors*, vol. 19, no. 9, p. 2203, 2019.

Chapter 7

General Discussion

7.1 Addressing Thesis Objectives

Throughout the chapters of this thesis, the three objectives established in the introduction were addressed. These are briefly reviewed here, as follows.

7.1.1 Objective 1: To establish CIIL as a viable adaptation strategy across both favourable and adverse deployment regimes—those where unsupervised high-confidence adaptation is expected to succeed and those where it is expected to fail.

Chapter 3 provided the first experimental validation of CIIL, directly addressing this objective through two complementary closed-loop evaluations. In the minimal-data condition, representing a regime where unsupervised high-confidence adaptation is theoretically well-suited, the chapter demonstrates that CIIL not only matches but

surpasses existing unsupervised adaptation by improving throughput while avoiding the drift that often emerges when pseudo-labels are trusted prematurely. In the electrode-shift condition, representing a distributional mismatch that typically causes unsupervised high-confidence adaptation to collapse, CIIL recovers performance by reassigning pseudo-labels based on context-derived suitability rather than raw classifier confidence. Empirical results across 32 participants showed that CIIL consistently improved online performance and significantly outperformed unsupervised high-confidence adaptation in both a sparse calibration scenario and after a 45° shift. These findings substantiate CIIL's ability to leverage weak, task-derived supervisory signals to support rapid, stable, and resilient adaptation across both subtly and severely mismatched operational regimes.

7.1.2 Objective 2: To investigate the deployment of CIIL across a spectrum of model classes, including Bayesian classifiers, static neural networks, and temporal neural networks.

The thesis evaluates CIIL broadly across model families, demonstrating that the framework generalizes beyond a single modelling paradigm. Chapter 3 introduced CIIL within statistical classification models (linear discriminant analysis), showing its capacity to stabilize and reorganize decision boundaries under perturbations. Chapter 4 extended this evaluation to static neural networks (multilayer perceptrons), including novel proportional-control adaptations, where CIIL enables zero-shot learning from scratch and surpasses screen-guided training. Chapter 5 further tests neural network-reliant CIIL in regression-based architectures, demonstrating substantial improvements in throughput, reductions in control drift, and coherent

re-alignment of proportional outputs during user-/co-adaptation. Finally, Chapter 6 explored CIIL within transformer-based temporal architectures combined with differentiable temporal alignment (sDTW), and demonstrated that CIIL enables these models to incorporate temporal dynamics into a smoother, more stable latent manifold. Collectively, these chapters show that CIIL is not limited by model class and can be adopted with significant benefit across statistical models, static neural networks, and temporal neural networks.

7.1.3 Objective 3: To investigate CIIL in both classification-based and regression-based myoelectric control paradigms.

Classification-based CIIL was introduced and validated in Chapters 3 and 4 across target-acquisition tasks, zero-shot continuous control, and object manipulation in virtual reality. In all cases, CIIL improved controllability by leveraging task context to correct or reinforce class predictions during real-time use. Regression-based CIIL was then addressed in Chapters 5 and 6, where the framework was adapted to support vector-style pseudo-labels instead of probabilities. These chapters illustrate that context suitability can inform both the direction and magnitude of regression outputs, enabling throughput improvements and markedly reducing drift during prolonged closed-loop operation. Temporal variants of regression-based CIIL (Chapter 6) further demonstrate that context-informed updates can guide the organization of temporal feature embeddings, yielding coherent vector fields and smoother trajectories. Together, these results verify that CIIL is compatible with both discrete and continuous paradigms, providing a unifying mechanism for correcting model predictions based on the behavioural goals implicit in the task.

7.2 Further Contributions

Beyond the contributions that were achieved as a direct consequence of the three main objectives of the thesis, several other noteworthy contributions were made within this body of research.

7.2.1 Contribution 1: To establish an incremental learning strategy for nonlinear, adaptive proportional control within classification-based myoelectric control.

Alongside the advancements of CIIL in Chapter 4, the benefits of incremental learning for proportional control thresholds were also explored. This study leveraged the experience buffer, stored primarily for the training of the neural network, to re-establish the upper- and lower-thresholds of signal energy for each motion category on a periodic basis. Further, a sigmoid mapping function was used between these thresholds instead of the conventional linear mapping function to provide a more accessible low speed state, which allowed for a greater speed multiplier to be applied in this study to the output of the model without increasing overshoots excessively. Empirically, this mapping fit observed behaviours better than the linear transfer function, as observed signal energies were most prevalent at either extreme of the proportional control mapping rather than uniformly distributed throughout the range (Chapter 4 Figure 4.6 A, B). Unfortunately, the long session length of the study (and resulting fatigue) did not allow for all permutations of pipelines to be tested to fully disentangle the influence on throughput of CIIL versus adaptive proportional control.

7.2.2 Contribution 2: To examine CIIL in conventional target acquisition environments and extend its application to more complex interaction arenas.

This contribution was achieved through the progression of tasks spanning Chapters 3-6, as well as additional research conducted in collaboration with the University of Oslo and l'Université Laval [1]. CIIL was first tested in conventional Fitts' law target-acquisition environments, where its strengths in stability, drift recovery, and throughput enhancement are established. The work then advanced into increasingly complex interaction settings. The experiment from Chapter 4 evaluated CIIL in continuous control from scratch, where the task dynamics are less constrained, and the model must simultaneously learn both class boundaries and proportional control mappings. The temporal neural architectures explored in Chapter 6 further tested CIIL in environments with more dynamic behavioural sequences, showing that CIIL helps temporal models maintain coherent latent organization even in heterogeneous, prolonged user-interaction sessions. The experiment from the aforementioned collaboration deployed CIIL in a virtual-reality object-manipulation scenario involving life-like grasps, variable task demands, and richer motion repertoires, revealing improvements not only in task success and speed but also in subjective workload [1].

Collectively, these studies demonstrate that CIIL generalizes effectively across task complexity, interaction modality, and control objectives, supporting robust real-time adaptation in both structured and ecologically richer environments.

7.2.3 Contribution 3: To compile a comprehensive review of state-of-the-art incremental learning approaches for myoelectric control, with emphasis on how different strategies address distribution shift between calibration and deployment.

Chapter 2 delivered this review by situating incremental learning within the broader historical and technical landscape of myoelectric control. The work synthesizes decades of research, transitioning from direct control to pattern-recognition approaches and identifies the persistent challenge of distribution shift arising from non-stationary EMG signals and behavioural variability. It introduces a structured taxonomy covering dedicated on-demand recalibration, unsupervised incremental learning, predictor-dependent incremental learning, and environment-dependent incremental learning, mapping each to the specific types of distribution shift they are designed to mitigate. The review clarifies the limitations of purely supervised calibration—particularly its sensitivity to screen-guided data and static behavioural assumptions—and frames incremental learning as a necessary paradigm shift toward continuous, user-in-the-loop adaptation. Further categorizations of incremental learning goals (domain-, task-, and class-incremental) and emerging methods such as transfer learning, domain adaptation, and self-supervised regression deepen the analysis, establishing a comprehensive conceptual foundation for understanding how various strategies attempt to preserve robustness as conditions drift between calibration and deployment.

Although other review papers have touched on the out-of-distribution problem from the perspective of confounding factors [2, 3], or user-learning/internal model formation [4], it should be noted that this is the first review to focus explicitly on

model-learning in user-in-the-loop myoelectric control.

7.2.4 Contribution 4: To make incremental learning more accessible for the community by open-sourcing software into LibEMG.

Developing incremental learning systems requires substantial engineering and integration efforts to enable pseudo-real-time model adaptation. This added complexity can hinder their reproducibility and translation to the research community and industry despite the consistent performance gains reported in this thesis and in prior work. Through repeated implementation of incremental learning pipelines across multiple software environments (Chapters 3–4: Unity; Chapters 5–6: PyGame; [1]: Unity/OpenXR; ongoing work: MuJoCo, Vicon) and learning paradigms (Chapters 3–4: classification-based CIIL; Chapter 3: unsupervised high-confidence adaptation; Chapters 5–6: regression-based CIIL), the resulting software has evolved into a reasonably generalizable framework. Each iteration introduced challenges to contend with. At first, when relying on Bayesian classifiers, model updates were performed near instantaneously, and all processing could be handled sequentially. However, after neural networks were introduced in Chapter 4, sequentially performing all operations (model inference, model adaptation, and harvesting input-output pairs) was no longer feasible as an epoch of adaptation exceeded the time between windows. This motivated a multi-process structure with many interconnections to notify when a new pseudo-label was produced, when a new series of input-output pairs was ready to send to the adaptation thread, and when an updated model was ready to be swapped into the online prediction thread. To improve accessibility for the research community, this framework has been open-sourced within LibEMG [5]. This comprehensive Python library for myoelectric control, co-developed and released by the

candidate, supports and integrates existing datasets, input devices, algorithms, and outputs that many researchers already use and has seen rapid adoption in the field. LibEMG 1.0.0 was released in August 2024 and has been downloaded more than 11,000 times across 70 countries as of February 2026. Figure 7.1 briefly illustrates the various new components of the incremental learning software.

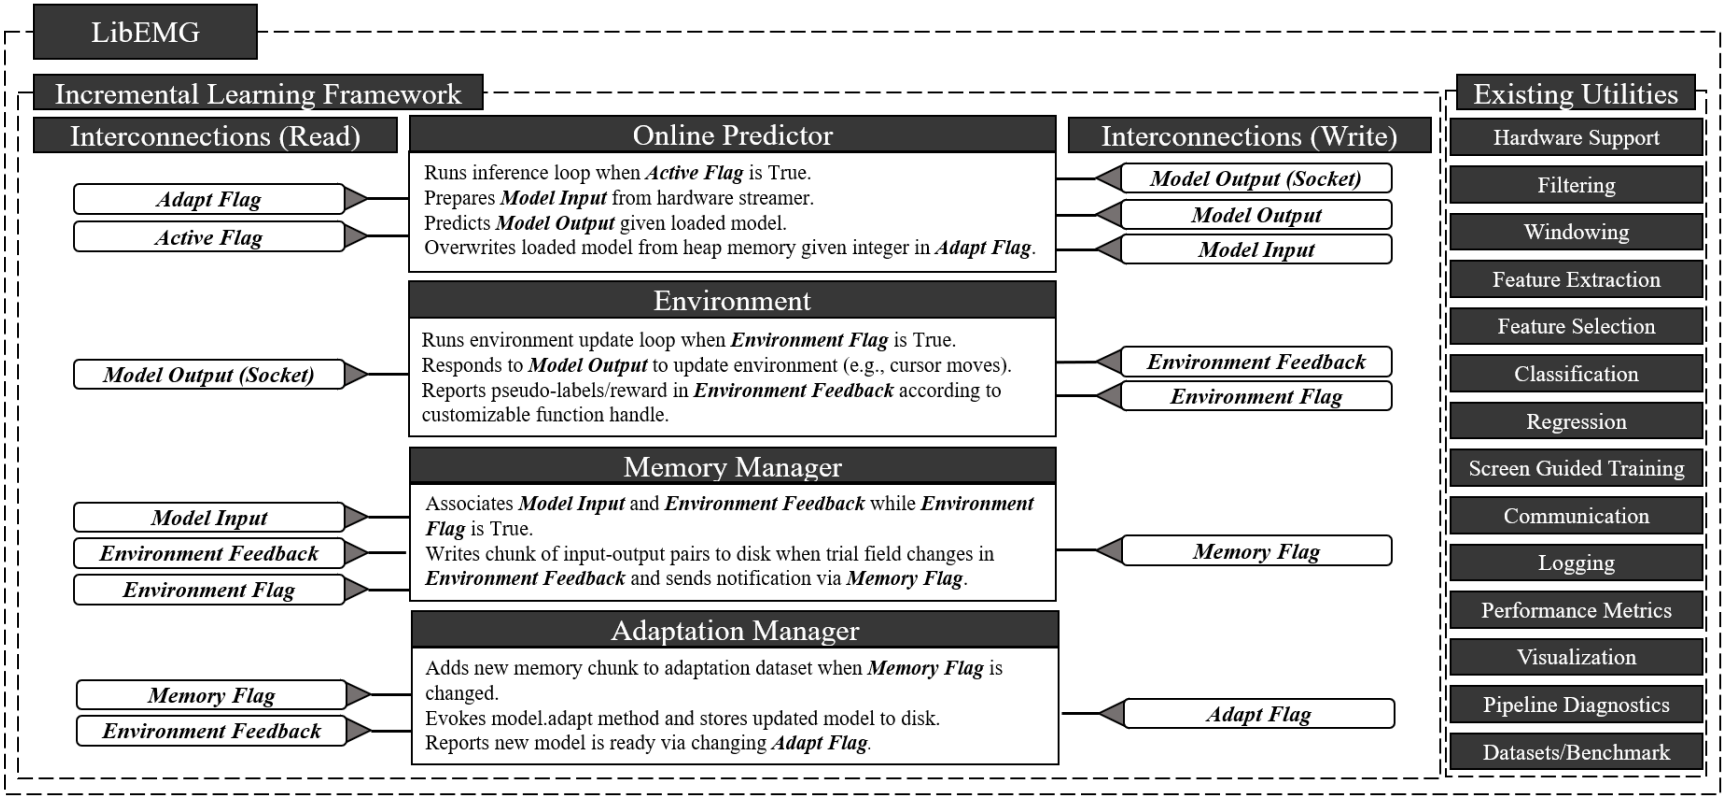


Figure 7.1: Overview of the incremental learning framework integrated into LibEMG. The system comprises four asynchronous processes: the Online Predictor, Environment, Memory Manager, and Adaptation Manager. The Online Predictor performs real-time inference and dynamically swaps to updated models when the Adapt Flag is triggered. The Environment consumes model outputs to drive task dynamics (e.g., cursor control) and returns Environment Feedback used for context-informed incremental learning or reinforcement learning. The Memory Manager associates model inputs with environment feedback and persists input–feedback pairs as adaptation-ready data. The Adaptation Manager incrementally updates the model using newly accumulated memory chunks, signals model availability to the Online Predictor, and coordinates data collection during adaptation.

7.3 Positioning CIIL Within Existing Context-Derived Adaptation Strategies

A brief comparison with prior context-derived adaptation strategies is valuable at this stage of the discussion. Although earlier work has demonstrated that task context can serve as a useful supervisory signal in user-in-the-loop adaptation, these approaches have been limited in scope and typically bound to specific control formulations or constrained task structures. In contrast, the findings presented throughout this dissertation show that CIIL provides a generalizable and principled mechanism for incorporating situational context into incremental learning. This section, therefore, highlights how CIIL extends and differentiates itself from existing methods, clarifying its unique contribution to the literature and reinforcing its role as a flexible, behaviour-aligned framework for real-time adaptive myoelectric control.

7.3.1 Regression-Based Context Leveraging

Several prior works have studied task-derived situational context, and it is important to position CIIL within this broader literature. Early evidence for using contextual information in incremental learning came from Hahne *et al.* [6], who adapted a position-based regression controller online during a target-acquisition task using the known target location as a supervisory signal. Chou *et al.* [7] recently validated this idea by demonstrating that target-derived (and eye-gaze–derived) pseudo-labels could reliably guide adaptation in a position-based environment. In these cases, the target location provides a unique or near-unique solution for the desired output, making pseudo-label assignment straightforward and, effectively, an extension of the initial calibration process.

However, the situation is fundamentally different for velocity-based control—the

paradigm used throughout Chapters 3–6. In velocity control, the model must infer the direction of motion at each timestep, and this direction is inherently non-unique. Users often pursue the target by decomposing the motion into sequential components along each axis, even when simultaneous multidimensional control is possible. If pseudo-labels were assigned using the same direct "ideal direction" approach used in position-based adaptation, the system would forcibly impose simultaneous movement even when the user intends sequential motion, thereby increasing the future likelihood of action interference.

CIIL addresses this mismatch by using situational context not to prescribe a single ideal action, but to define a set of context-eligible actions—supporting fully simultaneous, partially simultaneous, or fully sequential strategies. This flexibility makes CIIL uniquely compatible with velocity-based incremental learning, where user behaviour is more varied and existing task-derived pseudo-labelling strategies are insufficient.

7.3.2 Classification-Based Context Leveraging

A similar line of work exists in the classification-based myoelectric control literature, where task context was used to guide adaptation before the development of CIIL. Woodward *et al.* [8] demonstrated this in a serious-games environment designed to mitigate the limb-position effect. In their task, users were required to orient their limb toward a visual target while the controller classified the intended gesture (wrist flexion, wrist extension, or no motion). Whenever the controller's prediction did not correspond to the gesture indicated by the game's target, the system replaced the model's prediction with the correct label and used it as a pseudo-label for adaptation. This mechanism is conceptually similar to the N-CIIL strategy examined in Chapter 3, as both approaches exploit explicit task constraints to refine model predictions

over time. However, because Woodward's method depends on a task structure with a single correct gesture for each trial, it is inherently limited to scenarios with an unambiguous solution and does not generalize to tasks where users may employ multiple valid movement strategies.

Situational context has also begun to play a role in areas of myoelectric control beyond incremental learning, most notably in shared-control frameworks. These approaches operate on the premise that environmental perception can, in some cases, provide more reliable information about user intent than EMG alone. For example, Marković *et al.* [9] proposed systems in which external sensors infer the most likely grasp based on the surrounding scene, and the prosthesis executes that grasp whenever the user produces a generic contraction. While promising, early vision-based grasp-selection systems achieved recognition rates of only around 88% [10], and errors of this kind can meaningfully undermine user agency. More recently, Chamberland *et al.* [11] leveraged visual scene understanding to suppress implausible classes while still allowing the myoelectric controller to determine the final action. Their system, driven by a more capable vision model, achieved a 98% object identification rate and reduced misclassifications in myoelectric intent recognition by 60–100% across scenarios. As foundation models and scene-understanding systems continue to mature, both shared-control strategies and CIIL-like approaches that use situational context to guide or constrain myoelectric intent recognition will likely become increasingly effective and more viable for deployment in realistic, unconstrained environments.

7.4 Future Work and Limitations

7.4.1 Principled Analyses of Adaptation Dynamics

While this research demonstrated that CIIL is robust and versatile across architectures and tasks, future work should aim to characterize its adaptation behaviour more rigorously from a dynamical systems perspective. Existing analyses of co-adaptive myoelectric systems have typically relied on performance metrics such as throughput, path efficiency, or cursor error—metrics that, while informative, often fail to capture the nuanced temporal evolution of the user–model dyad. For instance, in Chapter 7, quiver plots provided general insights into the remapping that is established through CIIL, but this technique alone does not explain if input data still appears novel (high surprise), or if the incremental learning is currently bias-dominated (learning towards a minima) or variance-dominated (regularizing weights) [12]. Recent work by Al-Tahir *et al.* [13] introduced a compelling alternative: leveraging Poincaré maps—commonly used in bipedal locomotion studies—to visualize and quantify the periodic dynamics of co-adaptive control systems. Their analysis revealed that co-adaptive myoelectric interfaces, especially those involving continuous feedback and retraining, exhibit cyclic and sometimes chaotic behaviours that are poorly understood using traditional time-series plots. While the framework has shown empirical robustness, especially under distributional shift, there is limited understanding of how CIIL modulates user–model interactions over time. For example, such analyses might expose latent instability under aggressive pseudo-label reassignment or uncover patterns of stable convergence across repeated task exposures. These findings would complement the behavioural metrics already explored in the thesis and could inform adaptive parameter scheduling or more principled pseudo-label selection.

7.4.2 Psychological Aspects and Human Factors in Incremental Learning

Chapter 3 of this thesis included a NASA Task Load Index (NASA-TLX) to assess subjective workload during different adaptation strategies, but several psychologically grounded metrics remain underutilized in CIIL's evaluation. In particular, concepts from the neuroergonomics and HCI literature—such as sense of agency, ownership, and incorporation—are increasingly recognized as critical to understanding how users perceive and internalize adaptive systems [4]. While CIIL emphasizes behavioural goals inferred from task context, it is equally important to assess whether users feel in control of the adaptation process and whether the system supports or undermines their sense of autonomy.

Moreover, these metrics offer a pathway toward clinical and translational validation of adaptation frameworks. In rehabilitation or assistive device contexts, the degree to which users feel in control may be as important as measured performance. For example, a high-performing but disempowering system could reduce adherence or increase cognitive load [14]. Future iterations of CIIL could incorporate such metrics longitudinally to measure not only how performance evolves but also how the user–system relationship matures. Incorporation and agency scores may also serve as optimization targets during personalization, leading to CIIL variants that adapt not only for performance, but for human comfort and engagement.

7.4.3 Benchmarking CIIL Against Reinforcement Learning

The development of CIIL was heavily inspired by the limitations of unsupervised high-confidence adaptation and reinforcement learning in user-in-the-loop scenarios. While CIIL has already been benchmarked against confidence-based unsupervised

adaptation in Chapter 3, no direct comparisons have been made with reinforcement learning paradigms. This represents a key opportunity for future investigation, particularly given the growing interest in reinforcement learning-based control for prosthetics and assistive technologies.

Existing work in this space is limited but instructive. Freitag *et al.* [15] implemented a user-in-the-loop Q-learning controller in a simplified game-based environment (Guitar Hero), while Pilarski *et al.* [16] explored actor–critic methods for prosthesis control using sparse feedback signals. However, there is still a lack of consensus on which reinforcement learning strategies (on/off-policy, model-based/model-free, actor–critic, etc.) are best suited for classification- and regression-based myoelectric control.

Nevertheless, benchmarking CIIL directly against reinforcement learning under matched conditions—across tasks, architectures, and reward sparsity—could help illuminate where each approach excels or fails. Such comparative studies would also enable hybrid frameworks, where context-informed adaptation primes the model for reinforcement learning fine-tuning or vice versa. Ultimately, the goal is to understand whether CIIL's contextual pseudo-labelling offers comparable long-term benefit to the delayed but potentially more nuanced optimization offered by reinforcement learning.

7.4.4 Developing CIIL Strategies for Additional Myoelectric Control Arenas

While this work has validated CIIL in both classification-based and regression-based myoelectric control, several emerging areas of myoelectric research could also benefit from the framework. These include discrete control (e.g., event-based detection

of gestures such as snaps or swipes) [17], handwriting recognition (a hierarchical extension of discrete control involving character and word-level modelling) [18], and spike train decomposition (isolating individual motor unit activations from sEMG signals) [19]. First, in discrete control, task environments often contain sufficient structure to support context-derived pseudo-labels—for example, navigating a UI grid with sparse interconnections or interacting with widgets that permit only a limited set of valid gestures. These scenarios closely parallel the context-labelling mechanisms used throughout this thesis, making them well-suited for CIIL-based adaptation. Second, handwriting recognition systems similarly operate within constrained, structured domains. Lexicons provide a ground truth reference for valid outputs, enabling error detection when decoded words fall outside the allowable vocabulary. Pseudo-labels could be inferred by identifying lexicon entries with small Hamming distances to the observed output or via proximity in a learned latent embedding space. This form of error-aware correction aligns well with CIIL's use of context to guide adaptation. Finally, spike train decomposition—a process typically used to identify motor unit discharge timings from raw EMG—could serve as a learnable preprocessing block in a CIIL-enabled continuous control pipeline. During adaptation, this component could be fine-tuned alongside a downstream classifier or regressor. Once adaptation stabilizes, the decomposition module could revert to its canonical use, now better aligned with user-specific signal characteristics. In all these cases, the core CIIL principle—leveraging situational structure to guide online learning—remains applicable and valuable.

7.4.5 Translation to a Clinical Population

While interest in myoelectric control is growing within the broader human–computer interaction community, the strategies developed and evaluated in this thesis were

primarily motivated by the goal of improving intent recognition for powered upper-limb prostheses. Myoelectric interfaces are increasingly being used in clinical populations, including individuals with neuromuscular disorders or stroke-related motor impairments [20]. A key limitation of this work, however, is that all experiments were conducted exclusively with able-bodied participants. Future research should investigate the applicability and effectiveness of CIIL in these clinical populations, where signal variability, motor control strategies, and adaptation dynamics may differ substantially. Nonetheless, translation to these clinical populations may present the opportunity to overcome the systematically lower performance they experience with traditional calibration approaches [2], and harness gestures that feel more natural than those prescribed for normative populations (i.e., congenital amputees may struggle to produce gestures that their musculature doesn't support).

7.5 Conclusion

The studies described in this dissertation introduced Context-Informed Incremental Learning (CIIL) as a versatile and robust framework for real-time adaptation in myoelectric control. Through extensive user-in-the-loop experimentation, CIIL was shown to outperform conventional and confidence-based methods across a range of model types, control paradigms, and task complexities. By leveraging situational context to guide adaptation—even in the absence of explicit supervision—CIIL offers a scalable solution to the persistent challenge of distribution shift in user-in-the-loop systems. While demonstrated here using able-bodied participants, future work will be critical to validate its impact in clinical populations. Ultimately, CIIL represents a step toward more intuitive, resilient, and personalized myoelectric interfaces.

Bibliography

[1] G. Gagné, A. Azad, T. Labbé, E. Campbell, X. Isabel, E. Scheme, U. Côté-Allard, and B. Gosselin, "Context informed incremental learning improves myoelectric control performance in virtual reality object manipulation tasks," *arXiv preprint arXiv:2505.06064*, 2025.

[2] E. Scheme and K. Englehart, "Electromyogram pattern recognition for control of powered upper-limb prostheses: state of the art and challenges for clinical use.," *Journal of Rehabilitation Research & Development*, vol. 48, no. 6, 2011.

[3] E. Campbell, A. Phinyomark, and E. Scheme, "Current trends and confounding factors in myoelectric control: Limb position and contraction intensity," *Sensors*, vol. 20, no. 6, p. 1613, 2020.

[4] J. W. Sensinger and S. Dosen, "A review of sensory feedback in upper-limb prostheses from the perspective of human motor control," *Frontiers in neuroscience*, vol. 14, p. 345, 2020.

[5] E. Eddy, E. Campbell, A. Phinyomark, S. Bateman, and E. Scheme, "Libemg: An open source library to facilitate the exploration of myoelectric control," *IEEE Access*, vol. 11, pp. 87380–87397, 2023.

[6] J. M. Hahne, S. Dähne, H.-J. Hwang, K.-R. Müller, and L. C. Parra, "Concurrent adaptation of human and machine improves simultaneous and proportional myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 23, no. 4, pp. 618–627, 2015.

[7] A. H. Chou, M. Madduri, S. J. Li, J. Isa, A. Christensen, F. Hutchison, S. A. Burden, and A. L. Orsborn, "Using eye gaze to train an adaptive myoelectric interface," *bioRxiv*, pp. 2024–04, 2024.

[8] R. B. Woodward and L. J. Hargrove, "Adapting myoelectric control in real-time using a virtual environment," *Journal of neuroengineering and rehabilitation*, vol. 16, no. 1, p. 11, 2019.

[9] M. Markovic, S. Dosen, D. Popovic, B. Graimann, and D. Farina, "Sensor fusion and computer vision for context-aware control of a multi degree-of-freedom prosthesis," *Journal of neural engineering*, vol. 12, no. 6, p. 066022, 2015.

[10] G. Ghazaei, A. Alameer, P. Degenaar, G. Morgan, and K. Nazarpour, "Deep learning-based artificial vision for grasp classification in myoelectric hands," *Journal of neural engineering*, vol. 14, no. 3, p. 036025, 2017.

[11] F. Chamberland, T. Labbé, S. Tam, E. Scheme, and B. Gosselin, "Visual scene understanding for enhanced emg gesture recognition," in *2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 1–5, IEEE, 2024.

[12] R. Shwartz-Ziv and N. Tishby, "Opening the black box of deep neural networks via information," *arXiv preprint arXiv:1703.00810*, 2017.

[13] I. Al-Tahir, J. Sensinger, and E. Scheme, "A better framework for the assessment of performance and stability of co-adaptive myoelectric systems," in *2022 International Conference on Rehabilitation Robotics (ICORR)*, pp. 1–5, IEEE, 2022.

[14] E. Biddiss and T. Chau, "Upper-limb prosthetics: critical factors in device abandonment," *American journal of physical medicine & rehabilitation*, vol. 86, no. 12, pp. 977–987, 2007.

[15] K. Freitag, Y. Karayiannidis, J. Zbinden, and R. Laezza, "Fine-tuning myoelectric control through reinforcement learning in a game environment," *IEEE Transactions on Biomedical Engineering*, 2025.

[16] P. M. Pilarski, M. R. Dawson, T. Degris, F. Fahimi, J. P. Carey, and R. S. Sutton, "Online human training of a myoelectric prosthesis controller via actor-critic reinforcement learning," in *2011 IEEE international conference on rehabilitation robotics*, pp. 1–7, IEEE, 2011.

[17] E. Eddy, E. Campbell, S. Bateman, and E. Scheme, "Big data in myoelectric control: large multi-user models enable robust zero-shot emg-based discrete gesture recognition," *Frontiers in Bioengineering and Biotechnology*, vol. 12, p. 1463377, 2024.

[18] P. Kaifosh, T. R. Reardon, and C.-l. at Reality Labs, "A generic non-invasive neuromotor interface for human-computer interaction," *Nature*, Jul 2025.

[19] P. Capsi-Morales, D. Y. Barsakcioglu, M. G. Catalano, G. Grioli, A. Bicchi, and D. Farina, "Merging motoneuron and postural synergies in prosthetic hand design for natural bionic interfacing," *Science Robotics*, vol. 10, no. 98, p. eado9509, 2025.

[20] C. S. Klein, S. Li, X. Hu, and X. Li, "electromyography (emg) techniques for the assessment and rehabilitation of motor impairment following stroke," 2018.

Vita

Candidate's full name: Evan David Campbell

University attended (with dates and degrees obtained):

1. 2021 - Master's of Science in Engineering (Electrical), University of New Brunswick
2. 2018 - Bachelor's of Science in Engineering (Electrical), University of New Brunswick

Publications:

1. **Campbell, E.**, Egle, F., Oßwald, M., Côté-Allard, U., Pilarski, P., Boccardo, N., Meattini, R., Vujaklija, I., Hargrove, L., Canepa, M., Eddy, E., Del Vecchio, A., Castellini, C., Scheme, E., "(Un)supervised (Co)adaptation via Incremental Learning for Myoelectric Control: Motivation, Review, and Future Directions", IEEE Transactions on Neural Systems and Rehabilitation Engineering, Vol 33, No. 1, Aug 2025, pp. 3565, DOI: 10.1109/TNSRE.2025.3602397
2. Eddy, E., **Campbell, E.**, Morrell, C., Williams, H., Bateman, S., Scheme, E., "Raising the standard: an open source benchmarking platform and data repository to accelerate myoelectric control research", Machine Learning: Health, Vol 1, No 1, Jul 2025, pp 1, DOI: 10.1088/3049-477X/ade549
3. Morrell, C., **Campbell, E.**, Eddy, E., Scheme, E., "Context-Informed Incremental Learning Improves Throughput and Reduces Drift in Regression-Based Myoelectric Control", IEEE Transactions on Neural Systems and Rehabilitation Engineering, Vol 33, No 1, May 2025, pp 1841,
DOI: 10.1109/TNSRE.2025.3567245
4. Eddy, E., **Campbell, E.**, Bateman, S., Scheme, E., "EMG-based wake gestures eliminate false activations during out-of-set activities of daily living: an

online myoelectric control study", Journal of Neural Engineering, Vol 22, No 1, Jan 2025, pp. 1, DOI: 10.1088/1741-2552/ada4df

5. **Campbell, E.**, Eddy, E., Isabel, X., Bateman, S., Gosselin, B., Cote-Allard, U., Scheme, E., "Screen Guided Training Does Not Capture Goal-Oriented Behaviours: Learning Myoelectric Control Mappings From Scratch Using Context Informed Incremental Learning", IEEE Transactions on Neural Systems and Rehabilitation Engineering, Vol 33, No 1, Dec 2024, pp. 332, DOI: 10.1109/TNSRE.2024.3518059

6. Eddy, E., **Campbell, E.**, Bateman, S., Scheme, E., "Big data in myoelectric control: large multi-user models enable robust zero-shot EMG-based discrete gesture recognition", Frontiers in Bioengineering and Biotechnology, Vol 12, No 1, Sep 2024, pp. 1, DOI: 10.3389/fbioe.2024.1463377

7. Eddy, E., **Campbell, E.**, Côté-Allard, U., Bateman, S., Scheme, E., "Discrete Gesture Recognition Using Multimodal PPG, IMU, and Single-Channel EMG Recorded at the Wrist", IEEE Sensors Letters, Vol 8, No 9, Sep 2024, pp. 1, DOI: 10.1109/LSENS.2024.34472
40

8. Eddy, E., **Campbell, E.**, Bateman, S., Scheme, E., "Understanding the influence of confounding factors in myoelectric control for discrete gesture recognition", Journal of Neural Engineering, Vol 21, No 3, May 2024, pp. 1, DOI: 10.1088/1741-2552/ad4915

9. **Campbell, E.**, Eddy, E., Bateman, S., Côté-Allard, U., Scheme, E., "Context-informed incremental learning improves both the performance and resilience of myoelectric control", Journal of NeuroEngineering and Rehabilitation, Vol 21, No 1, May 2024, pp. 70, DOI: 10.1186/s12984-024-01355-4

10. Hajian, G., **Campbell, E.**, Ansari, M., Morin, E., Etemad, A., Englehart, K., Scheme, E., "Generalizing upper limb force modeling with transfer learning: a multimodal approach using emg and imu for new users and conditions", IEEE Transactions on Neural Systems and Rehabilitation Engineering, Vol 32, No 1, Jan 2024, pp. 391, DOI: 10.1109/TNSRE.2024.3351829

11. Chamberland, F., Buteau, E., Tam, S., **Campbell, E.**, Mortazavi, A., Scheme, E., Fortier, P., Boukadoum, M., Campeau-Lecours, A., Gosselin, B., "Novel wearable HD-EMG sensor with shift-robust gesture recognition using deep learning", IEEE Transactions on Biomedical Circuits and Systems, Vol 17, No 5, Sep 2023, pp. 968, DOI: 10.1109/TBCAS.2023.3314053

12. Eddy, E., **Campbell, E.**, Phinyomark, A., Bateman, S., Scheme, E., "LibEMG: an open source library to facilitate the exploration of myoelectric control", IEEE Access, Vol 11, No 1, Aug 2023, pp. 87380, DOI: 10.1109/ACCESS.2023.3304544

13. **Campbell, E.**, Phinyomark, A., Scheme, E., "Deep cross-user models reduce the training burden in myoelectric control", Frontiers in Neuroscience, Vol 15, No 1, May 2021, pp. 1, DOI:10.3389/fnins.2021.657958

14. **Campbell, E.**, Phinyomark, A., Scheme, E., "Current Trends and Confounding Factors in Myoelectric Control: Limb Position and Contraction Intensity", Sensors, Vol. 20, No. 6, Mar 2020, pp. 333, DOI: 10.3390/s20061613

15. Phinyomark, A., **Campbell, E.**, Scheme, E., "Surface Electromyography (EMG) Signal Processing, Classification, and Practical Considerations", Biomedical Signal Processing - Advances in Theory, Algorithms, and Applications, Springer, Singapore, ISBN 978-981-13-9096-8, pp. 3-29, 2020, DOI: 10.1007/978-981-13-9097-5_1

16. Côté-Allard, U., **Campbell, E.**, Phinyomark, A., Laviolette, F., Gosselin, B., Scheme, E. "Interpreting Deep Learning Features for Myoelectric Control: A Comparison with Handcrafted Features", Frontiers in Bioengineering and Biotechnology, Vol. 8, Mar 2020, DOI: 10.3389/fbioe.2020.00158

17. **Campbell, E.**, Phinyomark, A., Scheme, E., "Feature Extraction and Selection for Pain Recognition Using Peripheral Physiological Signals", Frontiers in Neuroscience, Vol. 13, No. 437, May 2019, DOI: 10.3389/fnins.2019.00437

Conference Presentations:

1. Gagné, G., Azad, A., Labbé, T., **Campbell, E.**, Isabel, X., Scheme, E., Côté-Allard, U., Gosselin, B., "Context Informed Incremental Learning Improves Myoelectric Control Performance in Virtual Reality Object Manipulation Tasks", 47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Copenhagen, Denmark, July 14-17, 2025, DOI: 10.48550/arXiv.2505.06064

2. Eddy, E., **Campbell, E.**, Bateman, S., Scheme, E., "Human-Machine Interaction Using Discrete Myoelectric Control: Contrastive Learning Reduces False Activations During Activities of Daily Living", 10th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob), Heidelberg, Germany, Sep 1-4, 2024, DOI: 10.1109/BioRob60516.2024.10719934

3. Morrell, C., **Campbell, E.**, Scheme, E., "Exploring user compliance in the training of regression-based myoelectric control", Myoelectric Controls Symposium, Fredericton, Canada, Aug 15, 2024, DOI: 10.57922/mec.2507

4. Chamberland, F., Isabel, X., **Campbell, E.**, Gagné, G., Gosselin, B., Scheme, E., Gagnon-Turcotte, G., Côté-Allard, U., "BioPoint: Single-site, Multi-sensor

Compound Gesture Recognition", Myoelectric Controls Symposium, Fredericton, Canada, Aug 15, 2024, DOI: 10.57922/mec.2505

5. **Campbell, E.**, Eddy, E., Côté-Allard, U., Scheme, E., "Improving User-in-the-Loop Myoelectric Control Using Context Informed Incremental Learning", Myoelectric Controls Symposium, Fredericton, Canada, Aug 15, 2024, DOI: 10.57922/mec.2498

6. Isabel, X., **Campbell, E.**, Olivier, L., Gagné, G., Scheme, E., Côté-Allard, U., Gagnon-Turcotte, G., "BioPoint: Enhancing Human-Computer Interaction through Single-Site, Multi-Sensor Gesture Recognition", 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Orlando, USA, July 15, 2024, DOI: 10.1109/EMBC53108.2024.10781656

7. Eddy, E., **Campbell, E.**, Bateman, S., Scheme, E., "Leveraging task-specific context to improve unsupervised adaptation for myoelectric control", IEEE international conference on systems, man, and cybernetics (SMC), Honolulu, USA, Oct 1, 2023, DOI: 10.1109/SMC53992.2023.10394393

8. **Campbell, E.**, Chang, J., Phinyomark, A., Scheme, E., "A Comparison Of Amputee And Able-Bodied Inter-Subject Variability In Myoelectric Control", Myoelectric Controls Symposium MEC '20, Fredericton, Canada, Online (due to COVID19), Aug 2020, DOI: arXiv:2003.03481

9. **Campbell, E.**, Phinyomark, A., Scheme, E., "Differences In Perspective On Inertial Measurement Unit Sensor Integration In Myoelectric Control", Myoelectric Controls Symposium MEC '20, Fredericton, Canada, Online (due to COVID19), Aug 2020, DOI: arXiv:2003.03424

10. **Campbell, E.**, Cameron, J., Scheme, E., "Feasibility of Data-driven EMG Signal Generation using a Deep Generative Model", 42nd Annual International Conference of the IEEE Engineering in Medicine and Biology Society, Montreal, Canada, July 20-24, 2020, DOI: 10.1109/EMBC44109.2020.9176072

11. **Campbell, E.**, Phinyomark, A., Scheme, E., "Linear Discriminant Analysis with Bayesian Risk Parameters for Myoelectric Control", 7th IEEE Global Conference on Signal and Information Processing, Ottawa, Canada, Nov 11-14, 2019, DOI: 10.1109/GlobalSIP45357.2019.8969237

12. **Campbell, E.**, Phinyomark, A., Al-Timemy, A., Khushaba, R., Petri, G., Scheme, E., "Differences in EMG Feature Space between Able-Bodied and Amputee Subjects for Myoelectric Control", IEEE Conference on Neural Engineering, San Francisco, USA, Mar 20-23, 2019, DOI: 10.1109/NER.2019.8717161

Workshops hosted:

1. Scheme, E., **Campbell, E.**, Eddy, E., Delsys, From Prosthetics to Human-Computer Interactions: Advances and Trends in Myoelectric Control, Exploring EMG-Driven Engineering Solutions with the Delsys API, July 16, 2025, https://delsys.com/event/webinar-emg-driven-engineering-solutions/

2. Egle, E., **Campbell, E.**, Oßwald, M., ICNR2024, WS10. (Co-)Adaptive (Un-)Supervised Learning for Myocontrol, November 20, 2024, https://2024.icneurorehab.org/workshops/

3. Scheme, E., **Campbell, E.**, Eddy, E., Morrell, C., MEC, LibEMG workshop: A New Open Source Library for Developing Myoelectric Interfaces and Conducting Myoelectric Control Research, Aug 11, 2024, https://www.unb.ca/ibme/mec/index.html