An Investigation of Transition-Informed Classifier Adaptation for Myoelectric Control

by

Julia Meneley

BSE.EE – University of New Brunswick, 2020

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

Master of Science in Engineering

In the Graduate Academic Unit of Electrical and Computer Engineering

Supervisors:	Dawn MacIsaac, Ph.D., Electrical and Computer Engineering Erik Scheme, Ph.D., Electrical and Computer Engineering
Examining Board:	Kevin Englehart, Ph.D., Electrical and Computer Engineering, Chair Shivam Saxena, Ph.D., Electrical and Computer Engineering Wei Song, PhD., Computer Science

This thesis is accepted by the
Dean of Graduate Studies

THE UNIVERSITY OF NEW BRUNSWICK

December, 2023

Abstract

Myoelectric prostheses use pattern recognition of surface electromyography (SEMG) to interpret a user's intent. Over time, changes in the SEMG worsen the usability of these prostheses, requiring cumbersome retraining. Adaptive learning, although able to update the classifier, suffers from mislabelling errors during unsupervised use. This study aimed to overcome this by investigating the impact of transitions between classes, often associated with elevated misclassification, on the adaptation process. Several adaptation techniques, some based on explicitly avoiding transitions and others based on leveraging awareness of transitions to improve decision stream labelling, were explored. Finally, these transition-informed adaptation techniques were tested on two datasets that included sequences of transitions between known classes. Results suggest that an awareness of transience in the SEMG can inform the data selection process and improve the labelling of unsupervised data for adaptation. A resulting LC-SSL technique yielded significant (p¡0.05) improvement to several offline classifier performance metrics.

Acknowledgements

My sincere thanks go to my supervisors, Dawn MacIsaac and Erik Scheme, for their continued support and guidance. I am grateful for the enthusiasm and encouragement to continually attempt new solutions to this elusive problem. I would also like to thank Shriram Raghu who was instrumental in the completion of this work. You not only allowed access to your data set, but were also continually available for both academic discussions and walks in Odell.

I also acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC).

Finally, my family deserves endless gratitude for putting up with me throughout this process. There have been innumerable dramatic phone calls and late-night ice cream trips. Your support was essential - I could not have done this without you.

Table of Contents

Abstract ii

Acknowledgments iii

Table of Contents iv

List of Tables vii

List of Figures xi

Abbreviations xiii

1 Introduction 1

1.1 Objective 2

1.2 Outline 3

2 Background and Related Work 5

2.1 Pattern Recognition based Myoelectric Control 5

2.2 Adaptive Learning in Myoelectric Control 6

2.3 Transitions in EMG Data 10

2.3.1 Continuous Transition Datasets 10

2.3.1.1 Raghu Dataset 11

2.3.1.2 3DC Dataset 12

2.3.2 Existing Transition Detection Algorithms 13

2.3.2.1 Decision Stream Segmentation 16

2.3.2.2 RuLSIF . . . 17
2.3.2.3 Bottom-Up Clustering . . . 18
2.3.3 Performance Metrics . . . 20

3 Evaluating the Effect of Transition-Informed Adaptation of Myoelectric Classifiers 22

3.1 Methods . . . 23
3.1.1 Data Set Selection . . . 23
3.1.2 Adaptation Schemes . . . 23
3.1.3 Analysis . . . 28
3.2 Results . . . 29
3.3 Discussion . . . 31

4 SEMG-Based Transience Detection and Labelling 34

4.1 Methods . . . 35
4.1.1 Selection of Data Sets . . . 35
4.1.2 Transition Detection Algorithms . . . 36
4.1.2.1 Evaluation of Performance . . . 37
4.1.2.2 Parameter Selection . . . 38
4.1.3 Data Selection for Adaptation . . . 40
4.1.3.1 Parameter Selection . . . 41
4.1.4 Transition-Informed Adaptation Labels . . . 43
4.2 Results . . . 44
4.2.1 Transition Detection Algorithms . . . 44
4.2.2 Transition-Informed Data Selection for Adaptation . . . 45
4.2.3 Transition-Informed Adaptation Labels . . . 45
4.3 Discussion . . . 48

5 SEMG-Based Transition-Informed Adaptation 54

5.1 Methods 55
5.1.1 Data Set Selection 55
5.1.2 Adaptation 56
5.1.3 Performance Evaluation 57
5.2 Results 58
5.2.1 Supervised Adaptation Results 59
5.2.2 Unsupervised Adaptation Results 59
5.3 Discussion 62

6 Conclusions and Future Work **66**

Bibliography **79**

A Covariate Shift Adaptation of Linear & Quadratic Discriminant Analysis Classifiers **80**

B SEM-Based Transition-Informed Adaptation of the 3DC Data Set **82**

Vita

List of Tables

2.1 Summary of steady state and transition classifier performance metrics used by Raghu et al. [45] 21

3.1 Nomenclature used to describe adaptation of LDA classifier 27
3.2 Effect of including the beginning of transitions during adaptation on classifier metrics when LMC-assigned adaptation labels are used . . . 30
3.3 Effect of including the end of transitions during adaptation on classifier metrics when LMC-assigned adaptation labels are used 31

4.1 Parameter values when not under test 39
4.2 Results of parameter grid-search for DSS selection technique to optimize myoelectric classifier performance. 39
4.3 Results of parameter grid-search for RuLSIF selection technique to optimize myoelectric classifier performance. 39
4.4 Results of parameter grid-search for BUC selection technique to optimize myoelectric classifier performance. 40
4.5 Weights Used to Combine Performance Metrics of Myoelectric Classifier 41
4.6 Results of parameter grid-search for high-confidence and low-confidence selection techniques to optimize myoelectric classifier performance. . . 42
4.7 Results of parameter grid-search for DSS selection technique to optimize myoelectric classifier performance. 42
4.8 Results of parameter grid-search for RuLSIF selection technique to optimize myoelectric classifier performance. 42

4.9 Results of parameter grid-search for BUC selection technique to optimize myoelectric classifier performance. 43
4.10 Transition detection performance on Raghu data set using parameters optimized for transition detection 44
4.11 Transition detection performance on 3DC data set using parameters optimized for transition detection 45
4.12 Transition detection performance on Raghu data set using parameters optimized for myoelectric classifier performance 48
4.13 Transition detection performance on 3DC data set using parameters optimized for myoelectric classifier performance 48
4.14 Steady state label accuracy and the number of frames identified as steady state by LMC which were not adapted to due to failing to meet the SSL threshold requirement. Results included for both no SSL and SSL with a threshold of 25% applied to the Raghu data set. 49
4.15 Steady state label accuracy and number of frames identified as steady state by LMC which were not adapted to due to failing to meet the steady state labelling threshold requirement. Results included for both no SSL and SSL with a threshold of 25% applied to the 3DC data set. 49

5.1 Classifier metrics resulting from SAL adaptation after 6 trials of initial ramp training and an adaptation weight of $W_{init} = 1.0$. LMC-assigned labels with SSL (th=25%). 59
5.2 Classifier metrics resulting from SAL adaptation after 6 trials of initial ramp training and an adaptation weight of $W_{init} = 0.5$. LMC-assigned labels with SSL (th=25%). 59

5.3 Classifier metrics resulting from SAL adaptation after 6 trials of initial ramp training and an adaptation weight of $W_{init} = 2.0$. LMC-assigned labels with SSL (th=25%). 60
5.4 Classifier metrics resulting from SAL adaptation after 1 trial of initial ramp training and an adaptation weight of $W_{init} = 1.0$. LMC-assigned labels with SSL (th=25%). 60
5.5 Classifier metrics resulting from SAL adaptation after 1 trial of initial ramp training and an adaptation weight of $W_{init} = 0.5$. LMC-assigned labels with SSL (th=25%). 60
5.6 Classifier metrics resulting from UAL adaptation after 6 trials of initial ramp training and an adaptation weight of $W_{init} = 1.0$. Classifier-assigned labels with SSL (th=25%). 61
5.7 Classifier metrics resulting from UAL adaptation after 6 trials of initial ramp training and an adaptation weight of $W_{init} = 0.5$. Classifier-assigned labels with SSL (th=25%). 61
5.8 Classifier metrics resulting from UAL adaptation after 6 trials of initial ramp training and an adaptation weight of $W_{init} = 2.0$. Classifier-assigned labels with SSL (th=25%). 61
5.9 Classifier metrics resulting from UAL adaptation after 1 trial of initial ramp training and an adaptation weight of $W_{init} = 1.0$. Classifier-assigned labels with SSL (th=25%). 62
5.10 Classifier metrics resulting from UAL adaptation after 1/2 trial of initial ramp training and an adaptation weight of $W_{init} = 1.0$. Classifier-assigned labels with SSL (th=25%). 62

B.1 Classifier metrics following adaptation on 3DC data set with relative importance of 100% $W_{init} = 1.00$. LMC-assigned labels with SSL (th=25%). 83

B.2 Classifier metrics following adaptation on 3DC data set with relative importance of 50% $W_{init} = 0.50$. LMC-assigned labels with SSL (th=25%). 83
B.3 Classifier metrics following adaptation on 3DC data set with relative importance of 200% $W_{init} = 2.00$. LMC-assigned labels with SSL (th=25%). 83
B.4 Classifier metrics following adaptation 3DC data set with relative importance of 100% $W_{init} = 1.00$. Classifier-assigned labels with SSL (th=25%). 83
B.5 Classifier metrics following adaptation on 3DC data set with relative importance of 50% $W_{init} = 0.50$. Classifier-assigned labels with SSL (th=25%). 84
B.6 Classifier metrics following adaptation on 3DC data set with relative importance of 200% $W_{init} = 2.00$. Classifier-assigned labels with SSL (th=25%). 84

List of Figures

2.1 Example of SEMG transience within both steady state and transitions. First trial of subject 1 of the 3DC data set. 14

3.1 Illustration of nomenclature used during LMC/Prompt-based assigning of adaptation labels. 25

3.2 Illustration of adaptation weights assigned using Equations 3.4 and 3.7 to adjust cumulative weight ratios. 27

3.3 The steady state active error (left) and transition duration (right) metrics over continuous trials as adaptation occurs. TR-Inclusive adaptation adapts during steady state and first 20% of transitions with 25% of cumulative adaptation occurring during transitions ($\gamma_1 = 20\%$ and R = 25%). No Adaptation, Naive Adaptation (i.e. adapting to all frames using classifier decision), and TR-Avoidant Adaptation (i.e. adapting only during steady-state regions) are included for comparison. 30

4.1 Detected transitions for P10 of Raghu data set. (1-top) Leap-identified transitions, (2) TDA-identified transitions when transition detection is optimized, (3) frames avoided during adaptation by ASTs when optimized for myoelectric classifier performance, and (4-bottom) MSR feature of SEMG signal 46

4.2 Detected transitions for P0 of 3DC data set. (1-top) Leap-identified transitions, (2) TDA-identified transitions when transition detection is optimized, (3) frames avoided during adaptation by ASTs when optimized for myoelectric classifier performance, and (4-bottom) MSR feature of SEMG signal 47

List of Symbols and Abbreviations

EMG	Electromyography
SEMG	Surface electromyography
LMC	Leap motion controller
LDA	Linear discriminant analysis
SAL	Supervised adaptive learning
UAL	Unsupervised adaptive learning
CPD	Change-point detection
DSS	Decision stream segmentation
RuLSIF	Relative unconstrained Least-Squares Importance Fitting
BUC	Bottom-Up Clustering
HC	High confidence
LC	Low confidence
AST	Adaptation selection technique
TDA	Transition detection algorithm
SSL	Steady state labelling

Chapter 1

Introduction

Myoelectric control, which uses electromyography (EMG) signals generated by the contraction of muscles to interpret a user's intent, has been widely used for decades to control powered prostheses [9]. Despite its popularity, surveys have found that myoelectric-controlled upper limb prostheses are abandoned by 35% and 23% of pediatric and adult users, respectively [7]. Along with reducing device weight and cost, and improving durability, [6], improving the robustness of classification-based myoelectric control remains one of the most important unmet clinical and research goals. Incorporating data from a variety of usage conditions into the training of classification-based myoelectric controllers is known to improve their robustness to those conditions during regular use [39]. However, changes in perspiration, electrode placement, limb position, force, and user behaviours can all affect the resulting SEMG signal [4, 31, 16, 37, 1], and collecting an initial training set that accounts for such wide-ranging variables isn't feasible. Consequently, adaptive learning has been proposed as a technique to adjust to these changes as they occur [27, 44, 52, 54].

In myoelectric control, classifiers interpret user intention by assigning segments of SEMG signal to existing motion classes (e.g. wrist flexion). Adaptive learning uses new data to update existing classifiers and is implemented in two main forms: super-

vised (when the intended motion class associated with a segment of data is known) and unsupervised (when the intended motion class is unknown). In supervised adaptive learning (SAL), the user engages in a prescribed data collection protocol when a decrease in the performance of their device is noted, by following a series of motion class prompts [50][10]. These prompts are then treated as labels for the associated SEMG data, and used to conduct SAL. Alternatively, unsupervised adaptive learning (UAL) updates the classifier continuously using data collected during regular prosthesis use. Although more appealing, in the absence of prompts, unsupervised classifier adaptation must rely on inferring intended motion classes, which isn't always accurate[44]. Correspondingly, misclassification during UAL introduces the risk of compounding classifier errors and worsening the performance of the prosthesis over time. In response, most UAL myoelectric control studies have been selective about which data are used to adapt; for instance, [44, 53, 18] report on using data with highly confident classification decisions to adapt. However, as users transition between trained motion classes during regular use of their prosthesis, they must travel through poorly defined gesture spaces, leading to erratic classifier responses and misclassifications with poor estimates of confidence [54, 3]. Despite this, the role of transitions has yet to be explicitly explored in the performance of unsupervised adaptation. The purpose of this work was to conduct such an exploration.

1.1 Objective

The main objective of this work was to leverage the consideration of transitions to improve unsupervised adaptation of myoelectric control. As an important component of their treatment, the automatic identification of transition regions was also explored. The following three specific aims delineate the focus of this work.

- The first aim was to investigate the effect of transitions on adaptation and

to determine if and how they should be included to benefit adaptation. This was explored using a data set that included continuous transitions between all combinations of motion classes. Both SEMG and kinematic data (e.g. hand position, velocity etc.), which served as pseudo-ground truth labels for gestures and transition boundaries, were available. How best to include transitions during SEMG adaption (e.g. fully/partially included, fully removed, etc) was explored, assuming ground truth knowledge of when transitions occurred.

- The second aim was to explore the potential of identifying transitions between motion classes from the SEMG itself using statistical and temporally-informed techniques. Identifying these regions was necessary for the real-world applicability of their specific treatment during unsupervised adaptation.

- The third aim was to integrate the developed transition detection and adaptation techniques and to evaluate their performance as part of a transition-informed unsupervised adaptation scheme.

1.2 Outline

Chapter 2 covers an overview of pattern recognition in myoelectric control, an introduction to adaptive learning, an explanation of terms used to describe the temporal nature of SEMG signals and a description of datasets which include controlled transitions between movement classes. A survey of algorithms which may be used to identify transitions in myoelectric control is also included. The effect of including transitions in adaptation is explored in Chapter 3, where available kinematic data is used to both determine intended motion classes and to identify transitions between motion classes. Transition detection and labelling techniques are evaluated in Chapter 4 with the intention of replacing the kinematic data in Chapter 3 with information derived solely from SEMG data. Chapter 5 combines the schemes of

transition inclusion identified in Chapter 3 with SEMG-based alternatives to kinematic data identified in Chapter 4 to determine the efficacy of EMG-based transition detection to improve both supervised and unsupervised adaptive learning in myoelectric control. Finally, Chapter 6 discusses the contributions, limitations, and future work informed by this research.

Chapter 2

Background and Related Work

2.1 Pattern Recognition based Myoelectric Control

Myoelectric control of prosthetic devices relies on the interpretation of user intent via EMG signals as a control input. This is often accomplished using pattern recognition to classify EMG data by recognizing similarities between elicited data and previous examples belonging to the same motion class (e.g. flexion, extension, etc.). This process generally consists of EMG acquisition, pre-processing, feature extraction, classification, and post-processing.

EMG acquisition consists of collecting EMG signals from the user of the system. Pre-processing includes filtering and signal segmentation; the collected EMG signal is typically filtered to remove both noise and artifacts. The signal is then segmented into overlapping frames from which features are extracted. These features are supplied to the classifier which labels the user intention as one of a set of predefined motion classes. Following classification, various post-processing techniques can be applied to improve system performance. Majority vote [19] and rejection [43] are two commonly implemented schemes in myoelectric control.

The classification approach relies on the assumption that the features corresponding to a class will remain consistent over time. Unfortunately, many factors such as perspiration, electrode placement, limb position, force variation, and user behaviour are all known to affect the resulting SEMG signal [4, 31, 16, 37, 1]. For this reason, adaptation has been introduced to allow the classifier to gradually adjust to the changing EMG features.

2.2 Adaptive Learning in Myoelectric Control

Classifier adaptation is the process by which new data are incorporated to adjust an existing classifier. The linear discriminant analysis (LDA) classifier, which assumes that all classes follow a normal distribution and share a common covariance, offers an intuitive example of this process. The LDA classifier decision, CD[n], is selected to maximize the probability that the SEMG features in frame n, x[n], originated from a given class c. This is described in Equation 2.1.

$$CD[n] = \arg\max_{c} P(c|x[n]) \tag{2.1}$$

To adapt an LDA classifier, the mean and covariance of the classes are simply adjusted using new data in order to influence $P(c|x[n])$, as described in Appendix A. This method of adaptation is called covariate shift adaptation [49].

Adaptation is generally separated into two categories: supervised and unsupervised. Supervised adaptation assumes that the user intention is known during the collection of adaptation data. This is common in myoelectric control but requires that the user follow an inconvenient and focused protocol of prescribed motions to update the model. This form of training can be screen-guided (SGT), where the user follows prompts displayed on a screen, or prosthesis-guided (PGT), where users follow preprogrammed prosthesis movements [10]. Although PGT improves convenience,

the user is still required to interrupt device use and follow the prompts provided to them. Additionally, both forms of SAL imply that the control performance has degraded sufficiently for the user to recognize a need to update the classifier. As a result, although SAL has been adopted clinically with SEMG pattern-recognition control schemes [50, 48], it remains cumbersome and reactive.

In comparison, unsupervised adaptation aims to proactively update the learned models without the need for supervised labels, relying instead on inferring the intention from the incoming data. In this case, incoming frames of data are assigned an adaptation label based on the current classifier decision. This inherently assumes that the current classifier is correct. However, the need for adaptation arises when the data are changing or the model has become outdated or insufficient in some way, meaning that it is likely to be making errors. Any such classification errors are then amplified as the data are fed back into the classifier for adaptation with incorrect labels. This severely limits the feasibility of long-term UAL in clinical applications, creating a cascade effect when models begin to degrade.

The application of some form of selection criteria to determine which data should be included in adaptation, here called *adaptation selection*, has been proposed as a potential solution. Established adaptation selection techniques (AST) have largely relied on isolated decision indicators, such as the current classifier's confidence in its adaptation label, or simple decision-stream error checking, such as excluding rapid, physiologically impossible, changes in classifier outputs [44, 36].

Sensinger et al. [44] reported a significant reduction in classifier error when adapting, in both supervised and unsupervised manners, to frames with high classifier confidence. However, while high-confidence adaptation is useful in reducing the amount of mislabelling error included during adaptation, it also restricts adaptation to only the regions in feature space which are already well-defined. Regions which would benefit from additional data are automatically excluded due to inherently ex-

periencing lower classifier confidence. Furthermore, some classifiers can yield erratic and erroneously high confidences during transitions, because they are largely undefined by current SGT and PGT practice. Consequently, if an adaptation selection technique could be developed that avoids mislabelling errors and allows adaptation to occur when classifier confidence is low, then additional improvements may be achieved.

The outlier detection approach proposed by Ding et al. implemented a hybrid classifier that used multiple one-class Support Vector Data Descriptions (SVDD) to identify frames that were outliers from the existing classes. Considering the entire duration of an active class, if a certain proportion of frames were found to be outliers, the LDA classifier responsible for prosthesis control was adapted[18]. This technique relies on identifying segments of SEMG belonging to one active class but the study provided no practical way to do this beyond requiring a mandatory rest period between contractions. Therefore, a more advanced method of segmenting continuous data is required for this technique to be implemented for UAL during device use.

The synergy-based data selection developed by Yeung et al. was used for UAL during a modified Fitts' law test, ensuring that adaptation occurred while subjects transitioned between motion classes [53]. However, SAL significantly out-performed the UAL performance for able-bodied subjects suggesting the need for continued improvement. Moreover, the technique relies on the simultaneous activation of opposing synergies (e.g. wrist flexion/extension, pronation/supination) to inform adaptation and may not translate well to classifiers that include classes that aren't paired with an opposing synergy such as multiple different hand grips or rest. The study also included only two degrees of freedom (i.e. four active classes).

In comparison, both Gu et al. and Huang et al. focused on adaptation selection to improve the limitations associated with the adaptation of support vector classifiers (SVC). By appending online data to the initial training set, the complexity of

the SVC increased exponentially along with computation cost. Gu et al. aimed to simplify the adaptation process by adapting to a representative sample set instead of all online data [21]. This was accomplished by first clustering the potential adaptation samples and then selecting samples from each cluster, randomly, to include during adaptation. A new classifier was then trained using the appended data set. It was found that this selection technique matched or improved upon the classification accuracy when all available online data was used to train new classifiers. The authors suggest that including all online data during training may result in worse performance as erroneous adaptation labels would be included.

Huang et al. implemented a representative-particle (RP) based adaptation which also relies on restricting the number of samples used to train the classifier [27]. Similarly to Gu et al., an initial SVC is trained on RPs randomly selected from the clusters identified in the original data set. However, during adaptation, the position of each RP is then modified using online data which is both close to, and shares the same label as, the RP being updated. A new SVC is then trained on the updated RPs - not an appended data set as with Gu et al. This RP-based classifier was able to maintain performance during UAL and improve it during SAL. In comparison, including all samples during UAL resulted in a significant degradation of performance.

While several ASTs show promising results, none directly address how transitions between movement classes should be managed during adaptation. Hakonen et al. [22] explained that classifier models that are trained to detect static classes are undefined for transitions between active classes. The importance of transitions is reinforced by the fact that the majority of misclassifications occur near or during regions of transition between classes [54, 3]. Despite this, the majority of reported studies on adaptation selection are implemented with data that didn't contain transitions between active classes. In addition to the previously discussed limitations, this may

cause the reported results to be overly confident in their ability to generalize to clinical use [27, 21, 55, 18, 44]. This is because they may be unable to identify the mislabelled frames that occur during, and in proximity to, regions of transition. Adaptation that considers transitions in the selection of appropriate data to be for adaptation will be referred to as transition-informed adaptation throughout the remainder of this work.

2.3 Transitions in EMG Data

To facilitate the study of transition-informed adaptation, it is important to differentiate between transitions and transient EMG behaviour. In this work, the term transition is used solely to represent the period of time during which the subject moves from one known motion class to another. Comparatively, any region in which an intended motion class is sustained is referred to as a steady state region. Transitions differ from transient SEMG behaviour, which describes any short-term temporal changes observed in SEMG signals. While transient EMG behavior is expected during transitions because of the dynamics associated with moving between motion classes, it can also be present in steady states, through small modulations of muscle synergies or variations in contraction intensity. In the remainder of this work, transient-EMG behaviour will be collectively termed transience.

2.3.1 Continuous Transition Datasets

While it is essential to study the impact of proposed classification strategies with users in the loop controlling a device, offline evaluation in myoelectric control research is often beneficial because removing user feedback allows for the rapid evaluation and comparison of multiple strategies and the re-use of pre-existing data sets of EMG signals for training and testing. Eventually, online studies are necessary to

determine the translation of offline gains to usability, but offline evaluations provide an important means for exploratory studies. As a result, publicly available EMG data stores for offline evaluation are growing (e.g. NinaPro datasets). Nevertheless, there are very few data sets available that include the explicit collection of continuous contractions; ones that contain transitions between active classes of movement along with secondary data (e.g. kinematic data) necessary to identify the bounds of each transition. Fortunately, two such sets were available for this work as described below.

2.3.1.1 Raghu Dataset

The Raghu Dataset, first used in [38], consists of 6 channels of SEMG collected from 43 subjects using Trigno electrodes [17] equally spaced around the circumference of the forearm of the dominant limb. This data set, sampled at 2 KHz and filtered between 20 Hz and 450 Hz. Kinematic data were also collected with a Leap Motion Controller (LMC) infrared sensor positioned below the hand and synchronized to SEMG features using timestamps.

Each subject engaged in 6 ramp trials and 8 continuous trials involving the wrist and hand. Ramp trial contractions started in the neutral class and gradually increased in contraction intensity to each of the 7 motion classes (including no movement) and held for 3 seconds. Subjects were guided through the trials using a screen-guided training approach. During continuous trials, prompts displayed to the participants changed every 3 seconds without pause - meaning the participants transitioned between classes to perform the newly prompted gesture. All 42 possible transitions, stemming from the 7 included motion classes, were captured in random order for each trial. Subjects were instructed to elicit moderate contraction intensities, but no feedback was provided to subjects.

All data were collected during the same session, electrodes were not adjusted between trials, and throughout the session, the forearm position was maintained

through the use of a supportive rest. These conditions helped to a) limit the amount of steady state transient SEMG behaviour in the data and b) ensure that the LMC was able to reliably identify the regions in which transitions occur. However, this also limited the amount of domain shift that occurred - suggesting that the potential benefit of adaptation may be limited.

2.3.1.2 3DC Dataset

The Long-Term 3DC Dataset has been made publicly available by Côté-Allard et al. [14]. This dataset consists of 10 channels of SEMG collected from 20 subjects using the 3DC Armband [15]. Data were filtered between 20 Hz and 495 Hz. Kinematic data were also collected with an LMC infrared sensor attached to a virtual reality (VR) headset and synchronized to SEMG features using timestamps.

The experimental protocol consisted of one training and two evaluation trials for each day of the study. Training trials involved the subjects mimicking gestures shown to them via VR. The subject ramped into each of 11 gestures, always starting from neutral, and held the contraction for 5 seconds. This was repeated 4 times per trial. During evaluation trials, a virtual reality environment (VRE) was used to display randomized gesture, limb position, and contraction intensity targets to the subjects every 5 seconds. Positional feedback was provided to the subjects through the use of the LMC to represent the limb in the VRE and mean absolute value (MAV) was used to provide intensity feedback.

A main feature of this data set is that it includes collections across 3 to 4 days per user, over a span of 14 to 21 days. In addition to electrode shift, variations in the user intentions, limb position, and contraction intensity likely result in domain shift which presents a greater opportunity for adaptation to both improve and sustain performance. Combined with the presence of feedback (even if it wasn't SEMG-based), these factors introduced substantial variability in the SEMG data of this

dataset.

Through visual inspection, it is evident that the 3DC data set often experiences transient SEMG behaviour during what was supposed to be steady state. This may have been caused by a variety of factors, including:

- reduced efficacy of LMC to identify transitions in non-neutral limb positions
- positional and intensity feedback leading to subjects performing minor corrections (not large enough in magnitude to be identified as a transition)
- an increase in mental load on subjects leading to performing transitions in a piece-wise manner

Figure 2.1 provides an example of transient SEMG behaviour that occurs during LMC-identified steady states. A time-series of the displayed prompt with four EMG features (maximum fractal length, Willison amplitude, mean-square-root, and the second L-moment) is included with the transitions detected in kinematic data provided by LMC superimposed. From the SEMG features it is clear that a substantial amount of SEMG transient behaviour is occurring during supposed steady states.

2.3.2 Existing Transition Detection Algorithms

Several techniques are already used to detect changes in SEMG signals. However, these mostly focus on locating the onset and offset of muscle contractions within the time series, assuming that transitions only occur between rest and active classes. As a result, these techniques tend to rely on applying amplitude thresholds to an aggregate of the amplitude measured across all channels. This is sufficient for onset and offset detection but often fails for transitions between two active classes, where there is not always an appreciable change in total EMG amplitude (some channel amplitudes may get larger, but others may get smaller). Therefore, thresholds would be required to be applied to each channel individually, and then somehow combined.

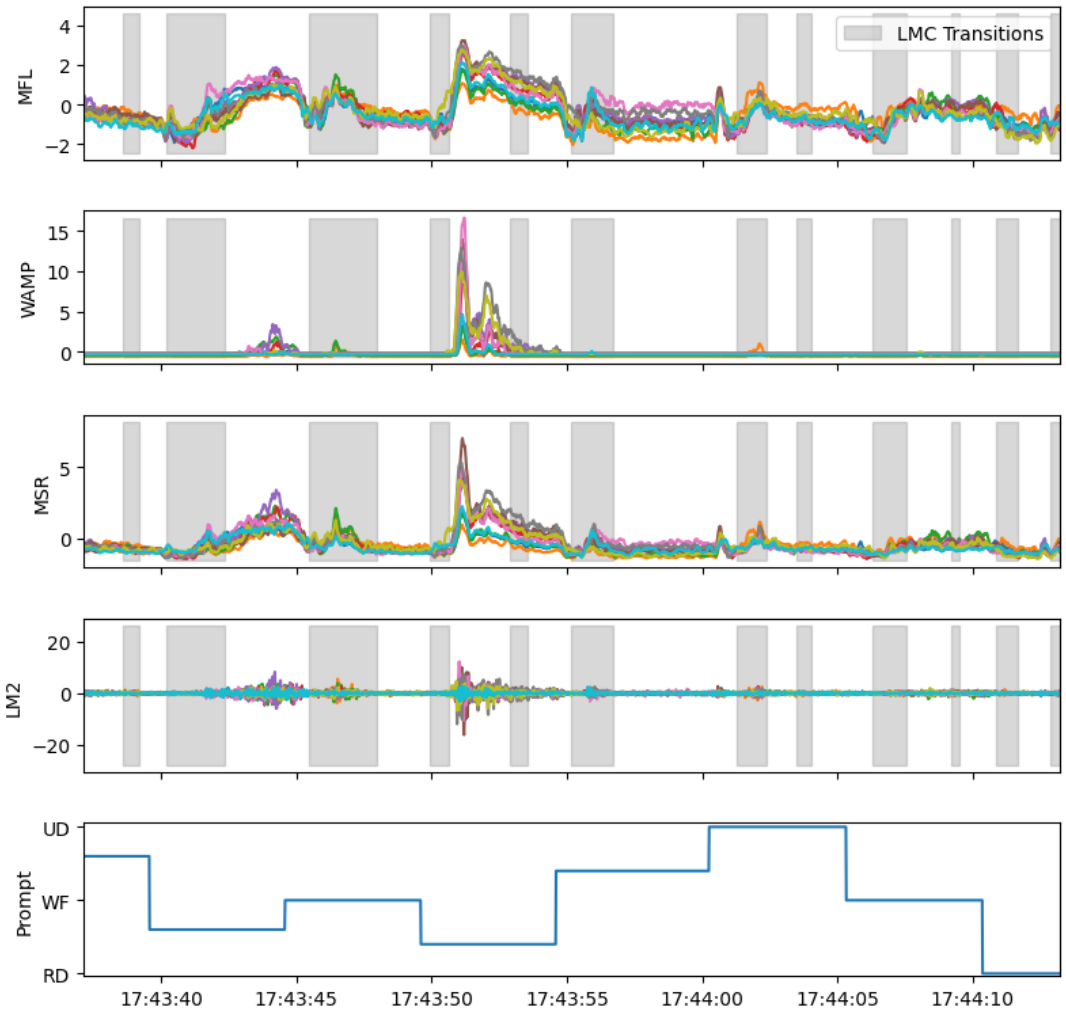

Figure 2.1: Example of SEMG transience within both steady state and transitions. First trial of subject 1 of the 3DC data set.

Additionally, the change in each channel during a transition depends on the transition being performed, meaning thresholds may need to be class-to-class specific. These obstacles suggest that simple thresholding may not be appropriate to detect transitions between active classes.

Many anomaly and change-point detection (CPD) algorithms have been developed to identify changes in time series signals. Supervised CPD algorithms encompass any algorithm that relies on classifier decisions [13, 2].Aminikhanghahi and Cook further separate supervised CPD algorithms into multi-class classifiers, binary classifiers, and virtual classifiers [2]. Multi-class CPD relies on changes in classifier decisions to identify changes in the underlying time series signal. In comparison, binary classifiers have been trained to explicitly classify between transitions and non-transitions [13, 2]. However, the number of different transitions, sparsity of training examples, and the variation in ways transitions in the myoelectric control problem occur result in prohibitively large within-class variation for the 'transition' class for real-world applicability.

Unsupervised CPD techniques are more varied in their approaches but generally rely on evaluating the likelihood that subsequent segments of the signal stem from the same distribution as those they follow [2, 5, 32, 29]. Two traditional approaches are density ratio comparison and subspace modelling. Density ratio comparison utilizes the probability density of two consecutive segments to calculate the likelihood that they originate from the same distribution [5][32]. Subspace modelling represents a time-series as belonging to states. If the state parameters from two consecutive segments are sufficiently different, then the underlying state is assumed to have changed and therefore a change-point in the time-series is identified [29]. These are both examples of parametric techniques, where an assumption about the time-series distribution or model must be made. Alternatively, kernels can be used to estimate the distributions non-parametrically [20] and clustering techniques can

identify change-points as resulting in different clusters [46].

2.3.2.1 Decision Stream Segmentation

A multi-class classifier CPD algorithm can be adopted for transition detection in SEMG signals. Change points in the underlying signal are identified by the classifier changing its decision from one class to another. When applied to SEMG signals, each movement should have a corresponding class in the multi-class classifier. Change points, $\delta[x]$, are identified as described in Equation 2.2, where d[x] describes the decision of the classifier at for frame x.

$$\delta[x] = \begin{cases} 0 & d[x] = d[x+1] \\ 1 & d[x] \neq d[x+1] \end{cases} \tag{2.2}$$

When applied to EMG transition detection, instability in the assigned classifier labels or transient SEMG behaviour not originating from a transition could erroneously be identified as transitions. Therefore, a windowing technique may be applied to identify how many decision changes have occurred within a given amount of time. If a sufficient number of change points are incorporated within the window, then the central frame of the window is identified as belonging to a region of transition. These relationships are defined in Equation 2.3, where the score (s) is dependent on the length of the window being considered (L), the changes identified in classifier decision (δ), and the windowing function selected (w).

$$s[x] = \sum_{n=-L/2}^{L/2-1} \delta[x+n] * w[n] \tag{2.3}$$

If this score is above a set threshold, then it is considered to be a transition, as described in Equation 2.4 where TR[x] is a binary signal where a value of 1 indicates that the EMG signal at time x belongs to a region of transition.

$$TR[x] = \begin{cases} 0 & s[x] \leq threshold \\ 1 & threshold < s[x] \end{cases} \tag{2.4}$$

2.3.2.2 RuLSIF

Relative unconstrained Least-Squares Importance Fitting (RuLSIF), a form of density ratio comparison, can also be adopted to detect transitions in SEMG signals. It is implemented in the manner described by Liu et al. [32] which is summarized below.

Given the SEMG feature signal y[x] of dimension d, a subsequence of length k can be described as shown in Equation 2.5, and a set of n subsequences can be defined as shown in Equation 2.6.

$$Y[x] := [y[x]^T, y[x+1]^T, ..., y[x+k-1]^T]^T \tag{2.5}$$

$$\varphi[x] := [Y[x], Y[x+1], ..., Y[x+n-1]] \tag{2.6}$$

A symmetric Pearson divergence is used to measure the similarity between two distributions, P and Q. It is defined in Equation 2.8.

$$PE_{sym}(P||Q) = PE(P||Q) + PE(Q||P) \tag{2.7}$$

$$PE_{sym}(P||Q) = \frac{1}{2}\int q(\mathbf{Y})(\frac{p(\mathbf{Y})}{q(\mathbf{Y})} - 1)^2 d\mathbf{Y} + \frac{1}{2}\int p(\mathbf{Y})(\frac{q(\mathbf{Y})}{p(\mathbf{Y})} - 1)^2 d\mathbf{Y} \tag{2.8}$$

However, both p($\mathbf{Y}$) and q($\mathbf{Y}$) are unknown so they would first need to be estimated, which is difficult in real-time analysis of dynamic time-series.

As an alternative, Liu et al. calculate the divergence using an α-regularized

estimation of the ratio between the distributions as described in Equation 2.9[32].

$$\widehat{PE_\alpha} = -\frac{\alpha}{2n}\sum_{i=1}^{n}\hat{g}(\mathbf{Y}_i)^2 - \frac{1-\alpha}{2n}\sum_{i=1}^{n}\hat{g}(\mathbf{Y}'_j)^2 + \frac{1}{n}\sum_{i=1}^{n}\hat{g}(\mathbf{Y}_i) - \frac{1}{2} \tag{2.9}$$

The density ratio estimate is defined by a Gaussian kernel model where θ_k describes the model parameters associated with kernel k.

$$\hat{g}(\mathbf{Y}) = \sum_{k=1}^{n}\theta_k \times exp\left(-\frac{\|\mathbf{Y}-\mathbf{Y}'\|^2}{2\sigma^2}\right) \tag{2.10}$$

When applied to SEMG transition detection, equation 2.9 is used to estimate the similarity between two sets of SEMG features $\varphi[x]$ and $\varphi[x+\ell]$ which are separated by a lag of ℓ frames. This divergence is assigned to the frame at the midpoint of all data included in the divergence calculation (i.e. $x + \frac{\ell}{2} + \frac{n}{2}$ rounded to the nearest integer). Distribution estimation was performed with 100 kernels with randomized means and sharing the same width. This width is assigned via a grid search of the resulting PE divergence.

Once the Pearson divergence is calculated using RuLSIF, transitions are determined to have occurred where this divergence is above a specified threshold as described in Equation 2.11 where $\widehat{PE_\alpha}[x]$ is the estimate of the Pearson divergence, associated with frame x, and t is the threshold.

$$TR[x] = \begin{cases} 0 & \widehat{PE_\alpha}[x] \leq t \\ 1 & t < \widehat{PE_\alpha}[x] \end{cases} \tag{2.11}$$

2.3.2.3 Bottom-Up Clustering

A bottom-up cluster (BUC) algorithm, as described in [30], can also be implemented for change point detection in the context of SEMG transition detection. It first splits the available signal into segments with a length of 2 frames. Sequential segments are

then iteratively combined in order of discrepancy between neighbouring segments. The discrepancy, $d(y_{a,t}, y_{t,b})$, where $y_{a,t}$ is the signal belonging to the first segment, $y_{t,b}$ is the signal belonging to the second segment, and $y_{a,b}$ is the signal resulting from the combination of both segments, is defined in Equation 2.12.

$$d(y_{a,t}, y_{t,b}) = c(y_{a,b}) - c(y_{a,t}) - c(y_{t,b}) \tag{2.12}$$

This equation relies on the cost function $c(.)$, which can be selected from a variety of applicable functions[47]. A commonly used cost function in change point detection is defined in Equation 2.13, where $\bar{y}_{a,b}$ is the mean of the signal $y_{a,b}$.

$$c(y_{a,b}) = \sum_{t=a+1}^{b} \|y_t - \bar{y}_{a,b}\|^2 \tag{2.13}$$

The sequential segments with the least discrepancy (i.e. the two segments are the most similar) are combined and the ranking is then readjusted with the newly formed segments. This continues until a stopping criterion is met. This technique is intended for use in detecting step changes in signals while transitions in EMG are generally gradual changes. Therefore, an additional stage can be implemented which identifies regions of transitions based on how long the standard BUC technique determines the signal to be maintained. During regions of transition, these segments will be much shorter than in regions of steady state. These relationships are defined in Equation 2.14 where s is the score, L is the window length, CP is the change-points determined by BUC, and w is the window function. A threshold is then applied to this score, as in Equation 2.4 to identify regions of transition.

$$s[x] = \sum_{n=-L/2}^{L/2-1} CP[x+n] * w[n] \tag{2.14}$$

2.3.3 Performance Metrics

To evaluate the benefit of transition-informed unsupervised adaptation, a technique must be selected to assess the performance of myoelectric control. Online tests that involve the manipulation of either physical or virtual objects are the most representative of the use case of prosthesis control. A common test is the clothespin task where the user is required to move clothespins from a vertical bar to a horizontal bar [8]. The box-and-block test is also commonly used to evaluate the dexterity of the upper limb. It involves the user moving blocks from one side of a partition to another [26]. Sensory feedback, such as tactile and force, have also been incorporated into such tasks and is known to allow users better control of the prostheses [34].

Fitts' law-style tests are a target acquisition task that are widely adopted in HCI and adopted for use in myoelectric control. They measure the trade-off between the speed and accuracy achieved in target acquisition when using myoelectric control to maneuver the cursor [41, 42, 51]. The Fitts' law, clothespin task, and box-and-block test are examples of online assessments which require the user to be present and perform a test for each individual scheme that is to be tested. As this work is an exploration of possible techniques to be used, the number of schemes compared is prohibitively high for these methods of evaluation.

As an alternative, offline metrics can be used as a proxy to evaluate the performance of myoelectric classifiers. One of the most prevalent of these is the classification accuracy i.e. the percentage of correct classifier decisions. Accuracy can be calculated offline - meaning that a classifier can be trained on pre-recorded training data and then evaluated on pre-recorded evaluation data. This allows for rapid comparison of various classifiers without the need for additional data collection.

Unfortunately, classification accuracy alone is known to correlate poorly with the usability of a resultant classifier [24, 35]. Nawfel et al. performed a study to determine the correlation between several offline and online metrics [35]. Accuracy

Metric Name	Abv.	Description
Total Error Rate	TER	The percentage of steady state frames assigned an incorrect label.
Active Error Rate	AER	The percentage of steady state frames assigned an incorrect, but active, label.
SS Instability	INS_SS	The number of times, in a steady-state region, that class decisions change. This is normalized by the total number of frames in that region.
Offset Delay	OFF	The delay, in milliseconds, between the prompt change and the change in classifier decision.
Onset Delay	ON	The delay, in milliseconds, between the prompt change and the classifier reaching the new steady state.
TR Duration	DUR	The duration, in milliseconds, of the transition between two class steady states.
TR Instability	INS_TR	The number of times, in a transition region, that class decisions change. This is normalized by the total number of frames in that region.
Tertiary Class Error	TCE	The percentage of frames during a transition labelled as an active class that is neither the originating nor terminating class.
Percent No Movement	PNM	The percentage of frames during a transition labelled as no movement/rest.

Table 2.1: Summary of steady state and transition classifier performance metrics used by Raghu et al. [45]

is also only useful during the steady state portions of the signal where an assumed label can be assigned. Here, the temporal performance of the classifiers, especially in and around regions of transitions, was of particular interest as this work involves incorporating temporal information (i.e. location of transitions) during adaptation. For these reasons, the framework proposed by Raghu et al. for use with continuous transition SEMG was adopted [45]. This framework consists of 3 commonly used steady state metrics and 6 novel transition metrics as delineated in Table 2.1. Of these transition metrics, the first three focus on the timing and duration of the transition while the last three reflect the classifier decisions during transition.:

Chapter 3

Evaluating the Effect of Transition-Informed Adaptation of Myoelectric Classifiers

As described in Chapter 2, steady state regions in EMG data are the regions during which an intended motion class is maintained. In comparison, transitions describe regions during which there is movement from one motion class to another. These transitions are often mislabeled during classification, in part because they are ill-defined during conventional SGT or PGT. Previous works have correspondingly noted an improvement in usability when myoelectric classifiers are trained on ramp data that includes transitions from rest to each active class [24, 28, 40]. It is reasonable to question, therefore, if the inclusion of portions of transitions between active classes during adaptation may better inform classifier boundaries and enhance the responsiveness and performance of the classifier.

Since there has been little work reported about transitions from one active class to another, it is currently unknown if these active-to-active transitions may prove to be beneficial. Therefore, a central aim of this study was to determine if includ-

ing regions of transitions during adaptation benefits the performance of the resulting classifier. This was accomplished by either avoiding or including, in various amounts, transition data during adaptation and comparing the offline performance of the resulting LDA classifiers.

3.1 Methods

3.1.1 Data Set Selection

To enable the exploration of transition-informed adaptation, the accurate identification of transitions in this study was essential. For this reason, the Raghu dataset, introduced and described in Section 2.3.1.1, was adopted for the aim of the work. It was chosen because data collection was conducted with the limb constrained to one position and with the same target force, producing clear and repeatable gestures that were reliably captured by the LMC. This ensured the reliable identification of transitions during the continuous transition trials which involved participants performing all 42 possible transitions resulting from the 7 included classes.

In keeping with the work of Raghu et al. [45], the decision stream was segmented into window lengths of 160 ms with 16 ms overlap. From these windows, the maximum fractal length (MFL), Wilson's amplitude (WAMP), mean-square root (MSR), and 2 nd L-moment features (LM2) were extracted [45]. The LMC data used to identify transition bounds were synchronized with the SEMG frames via timestamps.

3.1.2 Adaptation Schemes

An LDA classifier was trained and adapted using covariate shift adaptation, as described in Appendix A. The LDA classifier was trained on all available SEMG ramp data and adapted offline using the continuous trials. The adaptation of the LDA classifier occurred in 10-second windows. Performance was continuously evaluated

on the adaptation window that followed the window being processed (i.e. the next 10-second window where adaptation hasn't occurred yet).

To initially decouple the interaction between detecting transitions and how these transitions were included during adaptation, the LMC-identified transitions were used in this study. A combination of the LMC-identified transitions and knowledge of the displayed prompts were used to assign labels to the EMG data. This allowed for a best-case scenario, where the benefit of transitions could be evaluated under the assumption that the frames were appropriately labelled, as if during supervised adaptation. Even when a change in prompt occurred, the frames remained labelled as the previous class until a transition was detected by the LMC to account for delays in user response. The data were then divided as defined in Equation 3.3, where N_{pc} is the frame at which the prompt is changed, $N_{TR,start}$ and $N_{TR,end}$ are the LMC-identified start and end frames of transitions, $P[n]$ is the prompt being displayed during frame n, and $L[n]$ is the adaptation label assigned to frame n.

The transition split factors γ_1 and γ_2 allow for a percentage of the total transition, from the beginning and end respectively, to be assigned labels and included during adaptation. For this work, both γ_1 and γ_2 were set to 20%, 40%, 60%, and 80% while the other remains at 0%. The relationship between these transition split factors and the assigned adaptation labels are described in Equations 3.1 to 3.3. An illustration of the nomenclature used during this process is included in Figure 3.1.

$$N_1 = N_{TR,start} + \gamma_1 * (N_{TR,end} - N_{TR,start}) \tag{3.1}$$

$$N_2 = N_{TR,end} - \gamma_2 * (N_{TR,end} - N_{TR,start}) \tag{3.2}$$

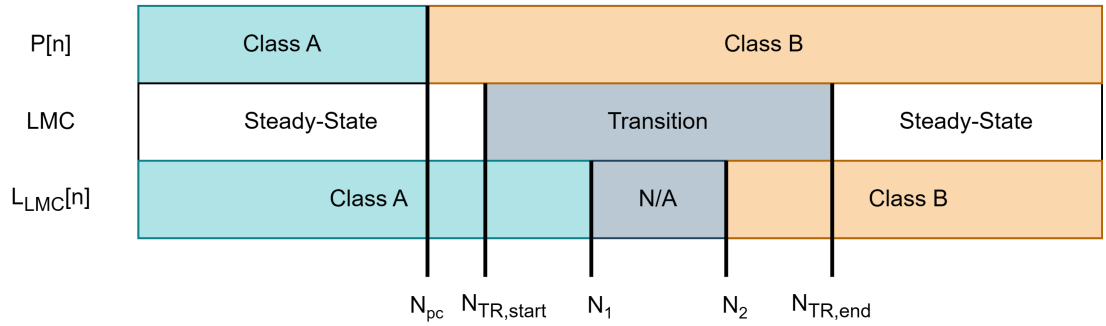

Figure 3.1: Illustration of nomenclature used during LMC/Prompt-based assigning of adaptation labels.

$$L_{LMC}[n] = \begin{cases} P[n] & n \leq N_{pc} \\ P[N_{pc} - 1] & N_{pc} < n \leq N_1 \\ nan & N_1 < n \leq N_2 \\ P[n] & n > N_2 \end{cases} \tag{3.3}$$

Adaptation weights were also manipulated to explore the effect of including varying amounts of transitions. The value $W_\alpha[n]$, defined in Equation 3.4, was used to assign the initial importance of adaptation frames, W_{init}, to both the regions of steady state and the regions of transition that were included during adaptation. In this study, an initial value of 1.0 was used, meaning that each frame of adaptation data had the same impact on the resulting classifier as one frame of the initial training data. Adjustments were then made to these initial weights to allow for a desired cumulative adaptation ratio between steady state and transition (e.g. 25% of adaptation occurs during transitions while 75% of adaptation occurs during steady states). These adjustments are defined in Equations 3.5 and 3.6 where the adaptation window spans from n_0 to n_1.

$$W_\alpha[n] = \begin{cases} W_{init} & n \leq N_1 \\ 0 & N_1 < n \leq N_2 \\ W_{init} & n > N_2 \end{cases} \tag{3.4}$$

$$\Delta_{SS} = \frac{\sum_{n=n_0}^{n_1} W_\alpha[n]}{\sum_{n=n_0}^{n_1} (1 - TR[n]) * W_\alpha[n]} \tag{3.5}$$

$$\Delta_{TR} = \frac{\sum_{n=n_0}^{n_1} W_\alpha[n]}{\sum_{n=n_0}^{n_1} (TR[n]) * W_\alpha[n]} \tag{3.6}$$

The adaptation weights adjusted for cumulative adaptation, $W_{R-adjust}[n]$, were then assigned using Equation 3.7 where R is the percentage of adaptation that is desired to occur during transitions, and $TR[n]$ is a binary signal equal to 1 when the LMC determines a transition is occurring and is otherwise 0. The value of R was set to 25%, 50%, and 75% to identify the effect that the proportion of transitions during adaptation has on the performance of the resulting classifier. The effect of Equations 3.4 and 3.7 are illustrated in Figure 3.2.

$$W_{R-adjust}[n] = \begin{cases} (1 - R) * \Delta_{SS} * W_\alpha[n] & TR[n] = 0 \\ R * \Delta_{TR} * W_\alpha[n] & TR[n] = 1 \end{cases} \tag{3.7}$$

A summary of the nomenclature used to describe the various adaptation schemes is included in Table 3.1

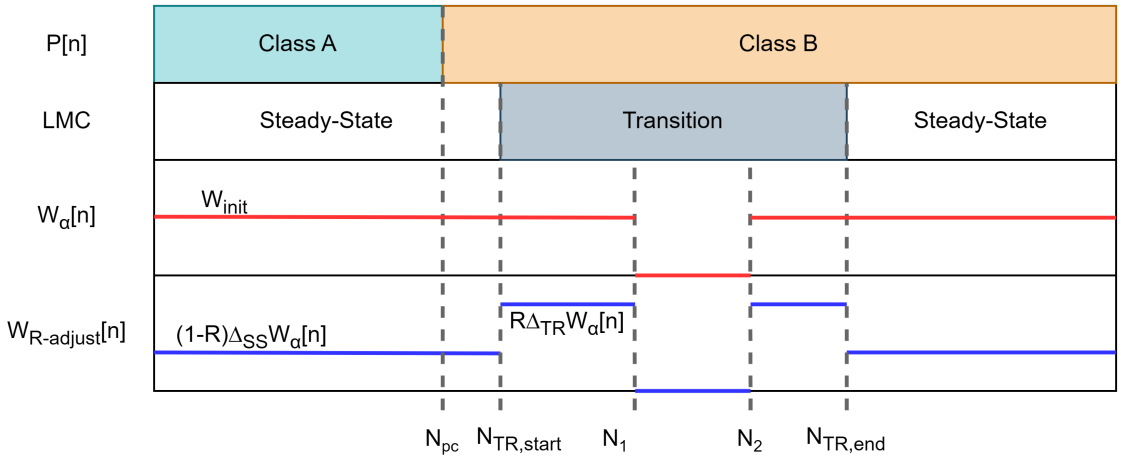

Figure 3.2: Illustration of adaptation weights assigned using Equations 3.4 and 3.7 to adjust cumulative weight ratios.

Parameter Name	Symbol	Description
Adaptation Weights	W[n]	The adaptation weight associated with frame n of SEMG data.
Adaptation Label	L[n]	The adaptation label associated with frame n of SEMG data.
Initial Importance	W_{init}	The initial importance of adaptation frames, relative to the importance of training frames, before modification to account for controlling cumulative adaptation amounts are performed. Used, in Equation 3.4, to determine the adaptation weight.
Transition Split Factors	γ_1 and γ_2	The percentage of a transition, from the beginning and end respectively, to be included during adaptation.
Cumulative Adaptation During Transition	R	The amount of adaptation that will occur during regions of transition included using γ_1 and γ_2 (i.e. 25% means 25% of adaptation occurs during transition and 75% during steady state).

Table 3.1: Nomenclature used to describe adaptation of LDA classifier

As a baseline comparison for the effect of the various adaptation techniques, the performance of the resulting classifiers was compared with those resulting from:

- **No Adaptation**: The initial training remained static without any adaptation updates, $W[n] = 0$.
- **Naive adaptation**: Adaptation to all frames of continuous trials with equal weight, $W[n] = W_{init}$, and using classifier-decisions as adaptation labels, $L[n] = CD[n]$. See Equation 2.1 of Chapter 2 for a definition of $CD[n]$.
- **TR-Avoidant Adaptation**: Adaptation only during LMC-identified steady state, $W[n] = W_{\alpha}[n]$, using LMC-assigned adaptation labels $L[n] = L_{LMC}[n]$.

3.1.3 Analysis

As previously mentioned, the adapted classifier was continuously evaluated using the window of data that followed the window being processed (i.e. 10 seconds of data that has not yet been adapted to). The metrics detailed in Section 2.3.3 were used to compare the various methods of transition-informed adaptation explored in this chapter.

- steady state Metrics: TER, AER, INS_SS
- Transition Metrics: OFF, ON, DUR, INS_TR, TCE, PNM

A two-sided paired t-test was performed to compare each metric of each adaptation strategy to that obtained by TR-Avoidant Adaptation. In addition, a Cohen D test is performed to evaluate the effect size associated with each comparison. A Cohen D result of 0.2 to 0.5 is interpreted as a small effect, 0.5 to 0.8 as a medium effect, and >0.8 as a large effect [12].

3.2 Results

Figure 3.3 shows the steady state active error and transition duration as adaptation progresses when using the transition-inclusive adaptation scheme described in Section 3.1.2 ($\gamma_1 = 20\%$ and R = 25%). It is noteworthy that all forms of adaptation outperform the classifier when no adaptation occurs and both LMC-informed transition-including and transition-avoidant adaptation perform better than naive EMG decision-based adaptation. Although the avoidance of transitions does appear to offer some benefit, the impact of their exclusion on AER is minimal compared to the relative benefit of using the LMC-defined labels rather than EMG-based labels derived from the classifier decision stream itself, as applied in naive adaptation. Also, the inclusion of transitions during adaptation produced a noticeable improvement in the transition duration. Nevertheless, the same relative impact of using LMC-defined labels (as opposed to EMG-based labels) held for transition duration performance.

Table 3.2 and 3.3 summarize the classifier performance for the transition-including scheme described in 3.1.2. The colours indicate the comparison of each classifier metric to what would be achieved when following TR-Avoidant Adaptation. These colours are intended only as a visual aid and do not indicate statistical significance or effect size. Statistical significance ($\alpha < 0.05$) is denoted by an asterisk (*) while effect size is designated by carets (ˆ- small effect, ˆˆ- medium effect, ˆˆˆ- large effect). When the amount of adaptation occurring during transitions was low, the inclusion of regions of transition during adaptation had a net benefit to classifier metrics. However, when the amount of adaptation occurring during transitions increased, the steady state metrics experienced a noticeable decline.

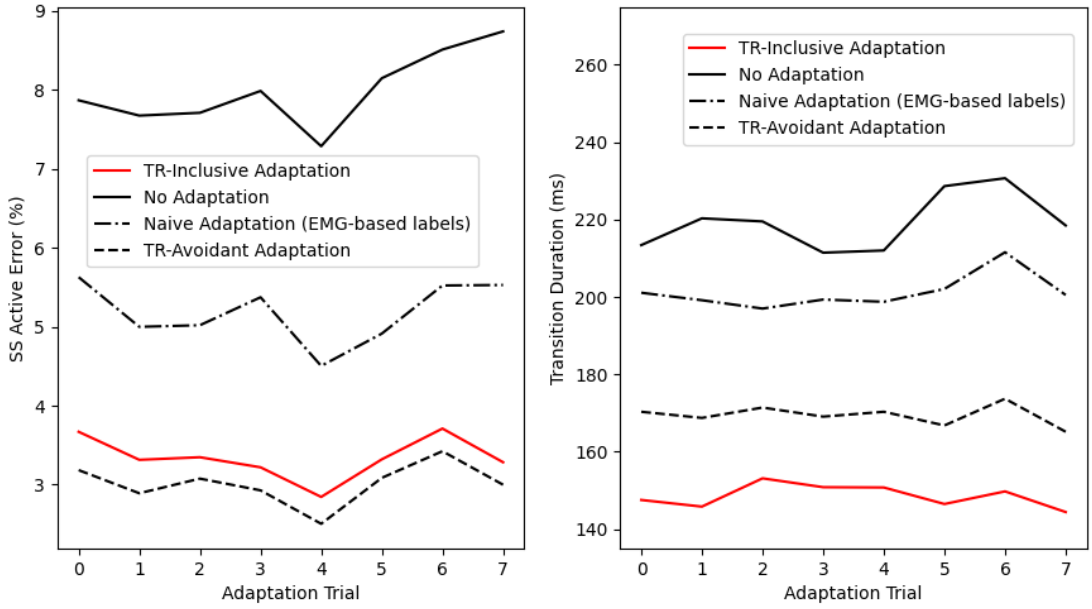

Figure 3.3: The steady state active error (left) and transition duration (right) metrics over continuous trials as adaptation occurs. TR-Inclusive adaptation adapts during steady state and first 20% of transitions with 25% of cumulative adaptation occurring during transitions ($\gamma_1 = 20\%$ and $R = 25\%$). No Adaptation, Naive Adaptation (i.e. adapting to all frames using classifier decision), and TR-Avoidant Adaptation (i.e. adapting only during steady-state regions) are included for comparison.

End of TR	% Adapt on TR	AER	TER	INS_SS	ON (ms)	OFF (ms)	DUR (ms)	INS_TR	TCE	PNM
20%	25	3.1%*	3.7%*	1.5%	464.32*	298.37*^	165.94*	7.8%	37.4%*	18.8%*
	50	3.5%*^	4.3%*^	1.6%*	463.68*	293.45*^	170.24	7.8%	36.3%*^	20.6%*^
	75	4.4%*^^^	5.5%*^^^	1.9%*^^	465.30*	280.56*^^	184.74*^	7.9%	35.6%*^	22.0%*^
40%	25	3.1%*	3.8%*	1.5%	457.30*^	297.20*^	160.10*^	7.9%	37.5%*	18.1%*
	50	3.6%*^	4.4%*^	1.7%*^	452.42*^	290.78*^	161.64*	7.7%	36.6%*^	19.4%*^
	75	4.5%*^^^	5.6%*^^^	2.0%*^^	451.70*^	277.74*^^	173.96	8.0%	35.9%*^	20.4%*^
60%	25	3.2%*	3.7%*	1.6%	451.31*^	295.27*^	156.04*^	7.9%	37.8%*	17.1%
	50	3.7%*^	4.4%*^	1.7%*^	439.10*^^	284.05*^^	155.05*^	8.0%	36.7%*^	17.7%
	75	4.7%*^^^	5.6%*^^^	2.0%*^^^	432.59*^^	269.64*^^^	162.94	7.9%	35.9%*^	18.4%
80%	25	3.2%*	3.8%*	1.6%*	444.51*^	291.58*^	152.93*^	8.0%	38.0%*	16.1%
	50	3.7%*^	4.4%*^	1.7%*^	428.91*^^	279.19*^^	149.72*^	8.0%	37.0%*^	16.3%
	75	4.8%*^^^	5.6%*^^^	2.1%*^^^	419.53*^^^	265.23*^^^	154.30*^	8.1%^	36.0%*^	16.8%
TR-Avoidant		3.0%	3.6%	1.5%	476.74	307.31	169.43	7.8%	39.4%	17.0%
No Adaptation		8.0%*^^^	9.7%*^^^	2.9%*^^^	502.95*^	283.64*^^	219.31*^^^	8.1%*^	40.1%	17.5%
Naïve Adaptation		5.2%*^^^	6.1%*^^^	1.8%*^	494.78*^	293.60*^	201.18*^^	7.4%*^	41.3%*	16.3%

Table 3.2: Effect of including the beginning of transitions during adaptation on classifier metrics when LMC-assigned adaptation labels are used

Start of TR	% Adapt on TR	AER	TER	INS_SS	ON (ms)	OFF (ms)	DUR (ms)	INS_TR	TCE	PNM
20%	25	3.3%*^	3.9%*	1.7%*^	444.59*^	296.02*^	148.58*^	7.6%	38.7%	12.9%*^
	50	3.9%*^^	4.5%*^^	2.0%*^^^	434.98*^^	295.01*^	139.97*^^	7.6%^	38.2%	11.2%*^^
	75	5.2%*^^^	5.8%*^^^	2.5%*^^^	433.06*^^	295.09*^	137.96*^^	7.6%^	37.7%	9.8%*^^
40%	25	3.4%*^	3.9%*^	1.7%*^	436.39*^^	291.70*^	144.69*^^	7.6%^	38.2%	12.2%*^
	50	4.1%*^^	4.7%*^^	2.1%*^^^	423.49*^^	289.48*^	134.00*^^^	7.6%^	37.3%^	10.3%*^^
	75	5.6%*^^^	6.3%*^^^	2.7%*^^^	416.70*^^^	288.02*^	128.68*^^^	7.3%*^^	36.1%*^	8.6%*^^
60%	25	3.4%*^	4.0%*^	1.8%*^	437.66*^^	292.99*^	144.67*^^	7.7%	38.4%	12.2%*^
	50	4.2%*^^^	4.8%*^^	2.1%*^^^	418.77*^^^	286.74*^	132.03*^^^	7.6%^	37.1%*^	10.0%*^^
	75	5.9%*^^^	6.7%*^^^	2.8%*^^^	409.89*^^^	284.63*^^	125.25*^^^	7.4%*^	35.3%*^	8.5%*^^
80%	25	3.4%*^	4.0%*^	1.8%*^	440.02*^^	294.11*^	145.91*^^	7.7%	38.6%	12.3%*^
	50	4.2%*^^^	4.8%*^^	2.1%*^^^	420.04*^^^	286.59*^	133.45*^^^	7.5%*^	37.1%*^	10.2%*^^
	75	5.9%*^^^	6.7%*^^^	2.9%*^^^	406.06*^^^	280.69*^^	125.37*^^^	7.4%*^	34.8%*^	8.7%*^^
TR-Avoidant		3.0%	3.6%	1.5%	476.74	307.31	169.43	7.8%	39.4%	17.0%
No Adaptation		8.0%*^^	9.7%*^^	2.9%*^^^	502.95*^	283.64*^^	219.31*^^^	8.1%*^	40.1%	17.5%
Naïve Adaptation		5.2%*^^^	6.1%*^^^	1.8%*^	494.78*^	293.60*^	201.18*^^	7.4%*^	41.3%*	16.3%

Table 3.3: Effect of including the end of transitions during adaptation on classifier metrics when LMC-assigned adaptation labels are used

3.3 Discussion

Tables 3.2 and 3.3 indicate that the inclusion of transitions during adaptation leads to mixed results, depending on the classifier metric. Generally, though, the inclusion of regions of transition during adaptation tends to degrade steady state metrics and improve transition metrics in the resulting classifier. The effect size of these changes is also proportional to the value of R (i.e. the amount of adaptation that occurs during transition regions). One exception to the transition metric trends is delineated in Table 3.2 which shows that including the beginning of transitions during adaptation leads to a worsened percent no movement metric. One possible cause for this is the shortening of transition duration as this likely leaves the difficult-to-classify frames remaining. Therefore, the transitions become shorter but more difficult to classify. In support of this notion, in Table 3.3, the transition duration metric experiences no improvements with an effect size larger than 0.2 while the PNM metric is either maintained or improved.

As there is a trade-off between steady state and transition classifier metrics, there is no definitive recommendation for transition inclusion which can be made - even when appropriate adaptation labels are available. Once a relationship between the classifier metrics and the resulting classifier usability is established, then

those metrics which most affect the usability should be prioritized. However, since this relationship is as yet unknown, it may be prudent to avoid transitions during adaptation. This is in part due to the unknown percentage of time in which users of prostheses spend in transition. Here, the R-value was artificially manipulated to reflect environments in which 25% to 75% of adaptation occurs during transitions. However, this value is unknown for clinical applications. Should users spend less than 25% of time performing transitions between classes, then any amount of the transition can be included during adaptation without worry of decreasing performance. However, as users spend more time performing transitions, the inclusion of transitions during adaptation may result in proportional decreases in steady state metrics. In the Raghu data set, an average of 13.3% of frames are identified as transitions by the LMC. However, this was collected under a controlled protocol and is therefore not representative of behaviour which may be seen in clinical settings. Additionally, it is still unknown if labels can be appropriately assigned during regions of transition. By definition, the SEMG signal during these periods does not correspond with any of the individual motion classes. This study used kinematic data from the LMC, in combination with information about the prompt being displayed, to assign adaptation labels. However, as it is unknown if these static labels are truly valid in a transition, it is prudent to avoid transitions during adaptation.

The apparent trade-off between steady state and transition classifier metrics also motivates the consideration of a hybrid myoelectric classifier consisting of two sub-classifiers. If one sub-classifier implements transition-avoidant adaptation while another implements transition-including adaptation, then incorporating the assigned labels from both of these classifiers may result in improvements to both steady state and transition metrics. In such cases, the transition-avoidant adapted classifier should maintain control during steady state behaviour but defer to the classifier which underwent transition-including adaptation when it is determined that a tran-

sition is occurring.

Ultimately, the results of this study motivate further exploration of the benefits of transition-informed adaptation. When using the LMC for context, including transition data during adaptation yielded an apparent tradeoff between improving offline transition metrics but worsening steady state metrics. However, for real-world applications, any transition-informed adaptation would require a robust transition (or possibly transience) detection system based solely on SEMG. It is also not known how any differences between the labels inferred from SEMG alone and those of the LMC might impact the overall adaptation process. Consequently, as a first step, Chapter 4 explores a variety of transience-informed data selection techniques for adaptation.

Chapter 4

SEMG-Based Transience Detection and Labelling

The work detailed in Chapter 3 indicated that knowledge of transience and transitions could be beneficial, whether to inform adaptation (as is the main focus of this work) or as part of some other hybrid control approach. In the previous chapter, the detection of transitions was accomplished with an external kinematic sensor (i.e. the LMC). Because this isn't viable in a real-world deployment of myoelectric control, the primary purpose of this chapter was to explore various transition detection algorithms (TDA), using only the SEMG data. These techniques were, therefore, largely reliant upon the detection of transience in the SEMG data to indicate the occurrence of transitions.

In addition to the identification of transitions, the work detailed in Chapter 3 also leveraged external sensor data to support the accurate assignment of adaptation labels (i.e. using a combination of the LMC-identified transitions and knowledge of the prompt being displayed to provide the labels). Without the LMC-based transition information or other knowledge of the displayed prompts, both the timing of transitions and the adaptation labels must be assigned solely based on the SEMG

signal. In UAL, this is most commonly accomplished using the current classifier decision as the label for subsequent adaptation. However, the identification of transitions allows for an opportunity to incorporate additional context into the assignment of adaptation labels. When the user is maintaining a steady state contraction, it directly follows that the user's intention should not be changing. Therefore, all frames between two sequential transitions could be assigned the same adaptation label, referred to here as steady state labelling (SSL). Consequently, a secondary purpose of this study was to investigate the potential of leveraging transience and transition information to improve the SEMG-based assignment of adaptation labels.

4.1 Methods

4.1.1 Selection of Data Sets

This study focused on the ability to identify transitions using only SEMG data and to further incorporate this information in the assignment of adaptation labels. Unlike the LMC sensor data, SEMG signals are highly stochastic and may contain a variety of transient behaviours. To explore these behaviours and their impact on transition detection, both the Raghu and 3DC datasets were included in this study. As in Chapter 3, the SEMG data (from both datasets) were segmented into 160 ms long frames with 16 ms overlap, and the fractal length (MFL), Wilson's amplitude (WAMP), mean-square root (MSR), and 2 nd L-moment features were extracted from each frame. LMC data were available from both data sets and used in this portion of the study to identify transitions as pseudo ground-truth for comparison with SEMG-identified transitions.

The Raghu dataset, described in Section 2.3.1.1, was collected as subjects continuously transitioned between 7 motion classes while maintaining constant limb position and contraction intensity. Because the data collection was well controlled,

there is little transient behaviour outside of transitions in this data set. In comparison, the Long-Term 3DC data, introduced in Section 2.3.1.2 was collected over the span of several weeks. When performing evaluation trials, VR prompts with randomized limb position, contraction intensity, and gesture were displayed every 5 seconds. Subjects received both positional feedback from the LMC and contraction intensity feedback from the MAV of the SEMG signal, with the intention of ensuring compliance with the prompt. However, this also inadvertently divided subject focus and incited them to continually refine their contractions. As a result, the 3DC data set has substantially more transient SEMG behaviours outside of transitions.

4.1.2 Transition Detection Algorithms

The first goal of this work was to determine if the TDAs proposed in Section 2.3.2 were capable of accurately identifying transitions within the SEMG signal. As a reminder, these included:

- **Decision Stream Segmentation (DSS)**

  If the stream of classifier-assigned labels for a sequence of consecutive SEMG frames changes more than a set threshold number of times, the frame in the center of this sequence of frames is identified as belonging to a transition.

- **Relative unconstrained Least-Squares Importance Fitting (RuLSIF)**

  The RulSIF algorithm provides a measure of similarity between two sequences of frames of SEMG features. If these two sequences are sufficiently different, then the SEMG frame directly between these sequences is identified as belonging to a transition.

- **Bottom-up Clustering (BUC)**

  The BUC algorithm groups similar segments of the stream of SEMG feature frames together. If there are too many segment changes within a given sequence

of frames, then the frame in the center of this sequence is identified as belonging to a transition.

A Tukey window was used for both the DSS and BUC windowing function [25]. The RuLSIF algorithm was implemented in part through the use of code made available by Makiyama et al. [33] and the BUC algorithm through the use of code made available by Truong et al. [47].

Because DSS is a classifier-dependent approach, a classifier must first be chosen. This classifier is, importantly, involved in both the initial transition detection and, subsequently, the selection of appropriate frames for inclusion in adaptation. To reflect the intended use of these TDAs, an adaptive LDA classifier, undergoing adaptation in 10-second segments, was used. For a given segment, transitions were first identified via the TDA. For the purposes of this portion of the work, the frames to be included during adaptation were then selected in a transition-avoidant manner. The next 10-second window was then processed based on the updated classifier.

4.1.2.1 Evaluation of Performance

The ability to identify transitions was measured on a frame-by-frame basis by comparing transitions identified by the TDAs with those identified by the LMC. Because the data sets were imbalanced between transitions and steady state, the F-β score was used. It is a combined metric of precision and recall, where precision (p) represents the rate at which frames identified as belonging to transitions are correct and recall (r) is the rate at which transition frames were identified as such [11, 23]. It is defined in Equation 4.1.

$$F_\beta = (1 + \beta^2) \times \frac{p \times r}{(\beta^2 \times p) + r} \tag{4.1}$$

β is a weighting factor that can be used to prioritize precision or recall. In this application, however, it is yet unknown which is more important: that no transitions are missed or that no steady states are erroneously labelled as transitions. Consequently, common β-values of 0.5, 1, and 2 were reported [11]. Accuracy, recall, and precision are also reported for completeness.

The recall is defined in Equation 4.2, where the true positive (TP) is the number of frames where both TDA and LMC identify the frame as a transition and the false negative (FN) the number of LMC-identified transition frames erroneously labelled as steady state.

$$r = \frac{TP}{TP + FN} \tag{4.2}$$

Similarly, the false positive (FP) is the number of LMC-identified steady state frames erroneously labelled as transitions and the true negative (TN) is the number of LMC-identified steady state frames correctly identified as such. The precision (p) and accuracy (a) can then be calculated using Equations 4.3 and 4.4 respectively.

$$p = \frac{TP}{TP + FP} \tag{4.3}$$

$$a = \frac{TP + TN}{TP + FN + FP + TN} \tag{4.4}$$

4.1.2.2 Parameter Selection

The selection of values for all TDA hyper-parameters was accomplished through a grid-search in which all but one parameter value was held constant. Table 4.1 defines the parameter values used when not under selection. A subset of the subjects from each data set (4 of 22, and 10 of 43 subjects, from the 3DC and Raghu data sets respectively) were randomly selected for use during the grid search selection of these parameters.

Table 4.1: Parameter values when not under test

Technique	Parameter Name	Value
DSS	Window Length	500 ms
	Threshold	0.3
RuLSIF	Window Lag	Combined (2,4, and 6)
	Alpha	0.1
	Threshold	0.05
BUC	Window Length	300 ms
	Average Segment Length	3000 ms
	Threshold	0.60

As described in Section 4.1.2.1, the ability of the TDAs to identify transitions can be measured by the F-β score. For the purpose of parameter selection, it was assumed that recall is prioritized over sensitivity (i.e. it was more important to correctly identify transitions than to avoid incorrectly identifying steady state frames as belonging to a transition). Therefore, an F-2 score was used as the metric being optimized during parameter selection.

The parameters under selection, the potential values, and the selected values are summarized in Table 4.2 to 4.4.

Table 4.2: Results of parameter grid-search for DSS selection technique to optimize myoelectric classifier performance.

Parameter Name	Threshold	Window Length
Tested Values	0.0:0.05:1.0	200:100:800 ms
Raghu	0.05	800
3DC	0	800

Table 4.3: Results of parameter grid-search for RuLSIF selection technique to optimize myoelectric classifier performance.

Parameter Name	Threshold	Window Lag	Alpha
Tested Values	0.0:0.05:1.0	2, 4, 6, combined	0.1:0.2:0.9
Raghu	0	2	0.7
3DC	0	2	0.7

Table 4.4: Results of parameter grid-search for BUC selection technique to optimize myoelectric classifier performance.

Parameter Name	Threshold	Window Length	Mean Segment Length
Tested Values	0.0:0.05:1.0	250:250:1000 ms	1000:1000:4000 ms
Raghu	0.05	250	3000
3DC	0.05	250	3000

4.1.3 Data Selection for Adaptation

The transition detection algorithms were evaluated based on their ability to simply detect transitions. However, since the intended use for them is to inform ASTs, their ability to appropriately select frames for inclusion in transition-avoidant adaptation was also evaluated. To do this, TDA parameters were also chosen to optimize the performance of the myoelectric classifier resulting from adaptation. This evaluation also made use of the F-β score.

The performance of the TDA approaches were also compared to two prevalent ASTs implemented in myoelectric adaptive learning: High Confidence (HC) and Low Confidence (LC) selection. In HC selection, adaptation occurs only when the confidence associated with the classifier is above a set threshold. Conversely, LC selection results in adaptation occurring only when the classifier confidence remains below a set threshold.

Several of these ASTs are necessarily classifier-dependent, meaning that a classifier is involved in the selection of appropriate frames for inclusion in adaptation. For the DSS-based AST, this is because the classifier is involved in transition detection. For both LC and HC, the classifier is directly responsible for selecting frames based on its decision confidence. For this reason, an LDA classifier, which undergoes adaptation as described in Section 4.1.2, was used.

	Metric Name	Weight
Steady state Metrics	Total Error Rate	2.5%
	Active Error Rate	30%
	Instability	15%
Transition Metrics	Transition Onset Delay	2.5%
	Transition Offset Delay	2.5%
	Transition Duration	2.5%
	Transition Instability	15%
	Tertiary Class Error	15%
	Percent No Movement	15%

Table 4.5: Weights Used to Combine Performance Metrics of Myoelectric Classifier

4.1.3.1 Parameter Selection

As the goal of the ASTs was to inform adaptation, the hyper-parameters associated with each AST were selected to optimize classifier performance after adaptation. To determine the performance of the adapted classifier, any of the steady state and transition performance metrics outlined by Raghu et al. could be utilized [45]. However, because these each capture a different aspect of performance, a weighted metric combination (WMC) of them, with the weights of each metric as included in 4.5, were used. As the absolute importance of each metric to the usability of myoelectric control is unknown, the weights were assigned using an intuitive understanding of their effect. The total and active error rates both provide similar information, with the active error being considered more important as it doesn't consider errors to no movement (which are deemed to be less hurtful) [40]. The duration of transitions, as well as the delay in transition offset and onset, were considered to be less important than transition instability, tertiary class error, and percent no movement, for similar reasons - the first 3 metrics do not require the user to make active corrections, only account for delays in the control scheme. The important steady state metrics were given the same cumulative weight as the important transition metrics with the remaining weight being equally divided among the remaining metrics.

Because the various metrics were of different scales (i.e. error rates span from 0 to

1 while transition timing is generally hundreds of milliseconds), they were normalized before being combined. A z-score normalization technique was used as described in Equation 4.5, where the metric of subject p, M[p], is normalized by subtracting the mean of, and then dividing by the standard deviation of, the metric across all subjects.

$$M_{norm}[p] = \frac{M[p] - \frac{1}{P}\sum_{p=0}^{P} M[p]}{\sqrt{\frac{1}{P} * \sum_{p=0}^{P}(M[p] - \frac{1}{P}\sum_{p=0}^{P} M[p])^2}} \tag{4.5}$$

The parameters under evaluation, the potential values, and the selected values are summarized in Table 4.6 to 4.9 for both the Raghu and 3DC data sets.

Table 4.6: Results of parameter grid-search for high-confidence and low-confidence selection techniques to optimize myoelectric classifier performance.

Parameter Name	HC - Threshold	LC - Threshold
Tested Values	0.0:0.05:1.0	0.0:0.05:1.0
Raghu	0.20	0.95
3DC	0.25	0.95

Table 4.7: Results of parameter grid-search for DSS selection technique to optimize myoelectric classifier performance.

Parameter Name	Threshold	Window Length
Tested Values	0.0:0.05:1.0	200:100:800 ms
Raghu	0.40	300
3DC	0.40	600

Table 4.8: Results of parameter grid-search for RuLSIF selection technique to optimize myoelectric classifier performance.

Parameter Name	Threshold	Window Lag	Alpha
Tested Values	0.0:0.05:1.0	2, 4, 6, combined	0.1:0.2:0.9
Raghu	0.15	6	0.7
3DC	0.15	combined	0.9

Table 4.9: Results of parameter grid-search for BUC selection technique to optimize myoelectric classifier performance.

Parameter Name	Threshold	Window Length	Mean Segment Length
Tested Values	0.0:0.05:1.0	250:250:1000 ms	1000:1000:4000 ms
Raghu	0.95	1000 ms	1000 ms
3DC	0.85	250 ms	3000 ms

4.1.4 Transition-Informed Adaptation Labels

In addition to selecting which data to include in adaption, which labels to apply during adaptation was explored. To investigate transition-informed adaptation labelling, Gaussian fusion was used to suggest a label which could be applied to the entire region between two transitions. Equation 4.6 defines this process where L_{SS} is the label assigned to the steady state region, $N_{SS,start}$ and $N_{SS,end}$ are respectively the start and end time of the steady state region, and $P(c|x[n])$ is the probability of the frame belonging to class c given the SEMG features, x[n].

$$L_{SS} = \arg\max_{c} \sum_{n=N_{SS,start}}^{N_{SS,end}} P(c|x[n]) \tag{4.6}$$

Adaptation labels and weights were then set using Equation 4.7 and 4.8 respectively. Note that this scheme avoids transitions during adaptation, although it could be modified to include regions of the transition as described in the previous chapter. The scheme also adapts conservatively to steady state regions using a confidence metric. Equation 4.9 describes the confidence that the steady state label applies to the entire steady state region $C_{L_{SS}}$. When the confidence is below a set threshold, C_{th}, no adaptation occurs over these regions. The primary motivation for this technique was to avoid the case of an entire steady state being labelled incorrectly if a transition was not detected. In such a scenario, two or more steady states would be detected as belonging to the same steady state with the SSL technique attempting to assign one label to all frames in this region. A preliminary empirical study resulted

in the selection of 25% as the confidence threshold.

$$L_{SS-Label}[n] = \begin{cases} L_{SS} & N_{SS,start} < n \leq N_{SS,end} \\ Nan & otherwise \end{cases} \tag{4.7}$$

$$W_{SS-Label}[n] = \begin{cases} W_{init} & N_{SS,start} < n \leq N_{SS,end} \; and \; C_{L_{SS}} > C_{th} \\ 0 & otherwise \end{cases} \tag{4.8}$$

$$C_{L_{SS}} = \frac{\max \sum_{n=N_{SS,start}}^{N_{SS,end}} P(c|x[n])}{N_{SS,end} - N_{SS,start}} \tag{4.9}$$

To evaluate the efficacy of this transition-informed steady state labelling (SSL) scheme, each frame's assigned label was compared to the LMC-assigned label to measure accuracy. The number of LMC-identified steady state frames that the SLL rejected was also reported.

4.2 Results

4.2.1 Transition Detection Algorithms

Table 4.10 and 4.11 include the accuracy, recall, precision, and F-β score, averaged across subjects, for the Raghu and 3DC data sets respectively. These results correspond to the case where parameters were chosen to optimize the F-2 score (i.e. a measure of transition detection).

Table 4.10: Transition detection performance on Raghu data set using parameters optimized for transition detection

Technique	Accuracy	Sensitivity	Precision	F-0.5 Score	F-1 Score	F-2 Score
DSS	0.761	0.954	0.362	0.412	0.520	0.711
RuLSIF	0.133	1.000	0.133	0.161	0.235	0.433
BUC	0.889	0.378	0.650	0.566	0.476	0.412

Table 4.11: Transition detection performance on 3DC data set using parameters optimized for transition detection

Technique	Accuracy	Sensitivity	Precision	F-0.5 Score	F-1 Score	F-2 Score
DSS	0.349	1.000	0.349	0.401	0.517	0.728
RuLSIF	0.340	1.000	0.340	0.391	0.507	0.719
BUC	0.647	0.085	0.409	0.231	0.140	0.100

Figure 4.1 provides an example of the transitions detected by each TDA for the Raghu data set. Both the LMC-identified transitions and the MSR of the SEMG signal are included for comparison. Additionally, the frames excluded by the AST are indicated for the same adaptation window to compare the transition detection of the TDA to the frames selected for inclusion by the ASTs. Figure 4.2 provides a similar example for the 3DC data set.

4.2.2 Transition-Informed Data Selection for Adaptation

The previous section focused on the transition detection performance alone. Although this is important, the motivation of this work was to use this to improve adaptation. Consequently, the results in Tables 4.12 and 4.13 show the transition detection performance when optimized for myoelectric classifier performance. These include the accuracy, recall, precision, and F-β score, averaged across subjects, for the Raghu and 3DC data sets respectively. It is worth noting that, although HC and LC are not explicitly transition-informed, the results of Tables 4.12 and 4.13 suggest that they are reflecting transient behaviours. This is likely tied to the fact that myoelectric classifier confidence is known to be higher during steady-state portions and lower, or more erratic, during transitions [45].

4.2.3 Transition-Informed Adaptation Labels

The results of applying SSL with a 25% threshold, as a method of transition-informed adaptation labelling, are reported in Table 4.14 and 4.15 for the Raghu data set and

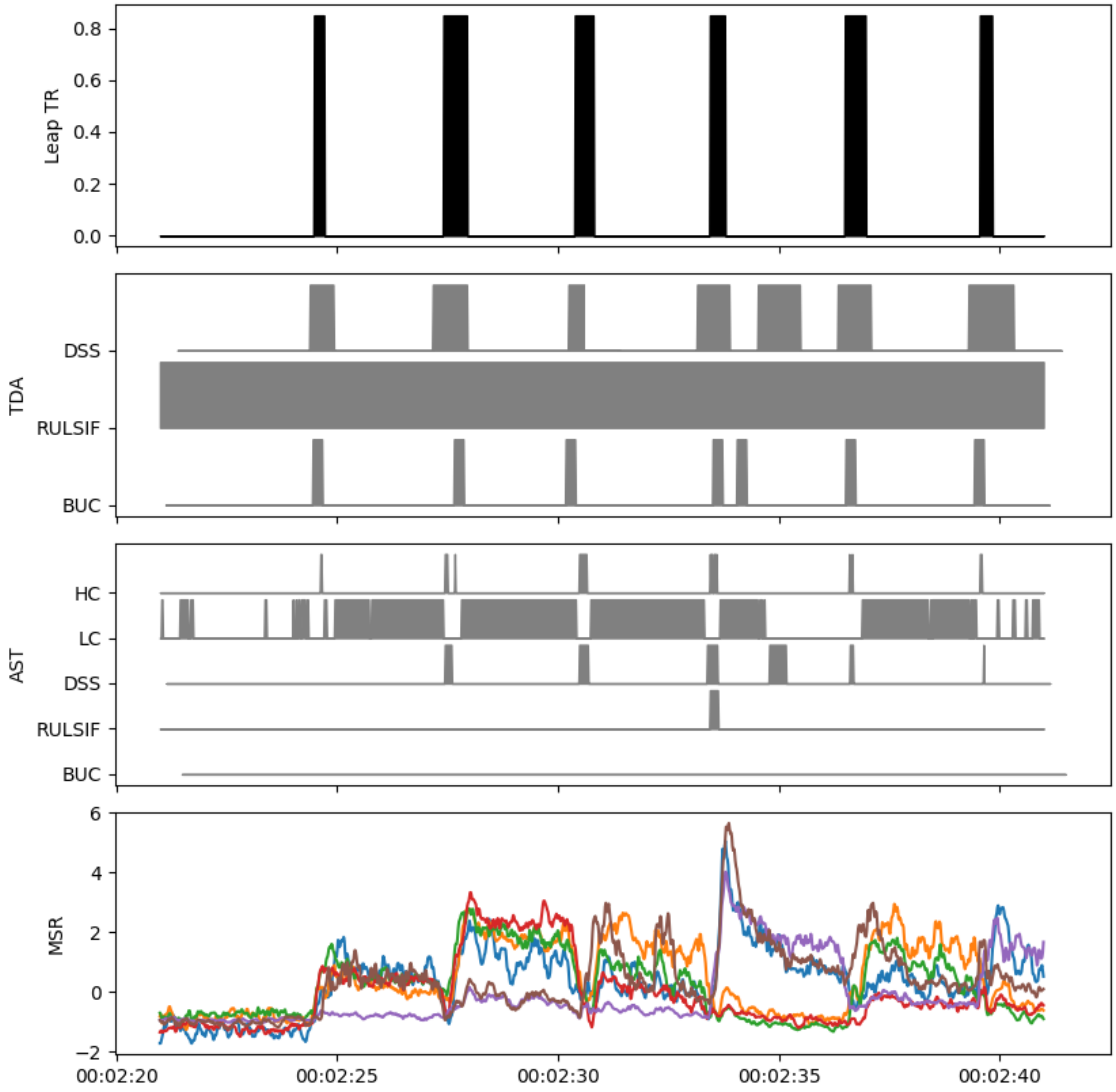

Figure 4.1: Detected transitions for P10 of Raghu data set. (1-top) Leap-identified transitions, (2) TDA-identified transitions when transition detection is optimized, (3) frames avoided during adaptation by ASTs when optimized for myoelectric classifier performance, and (4-bottom) MSR feature of SEMG signal

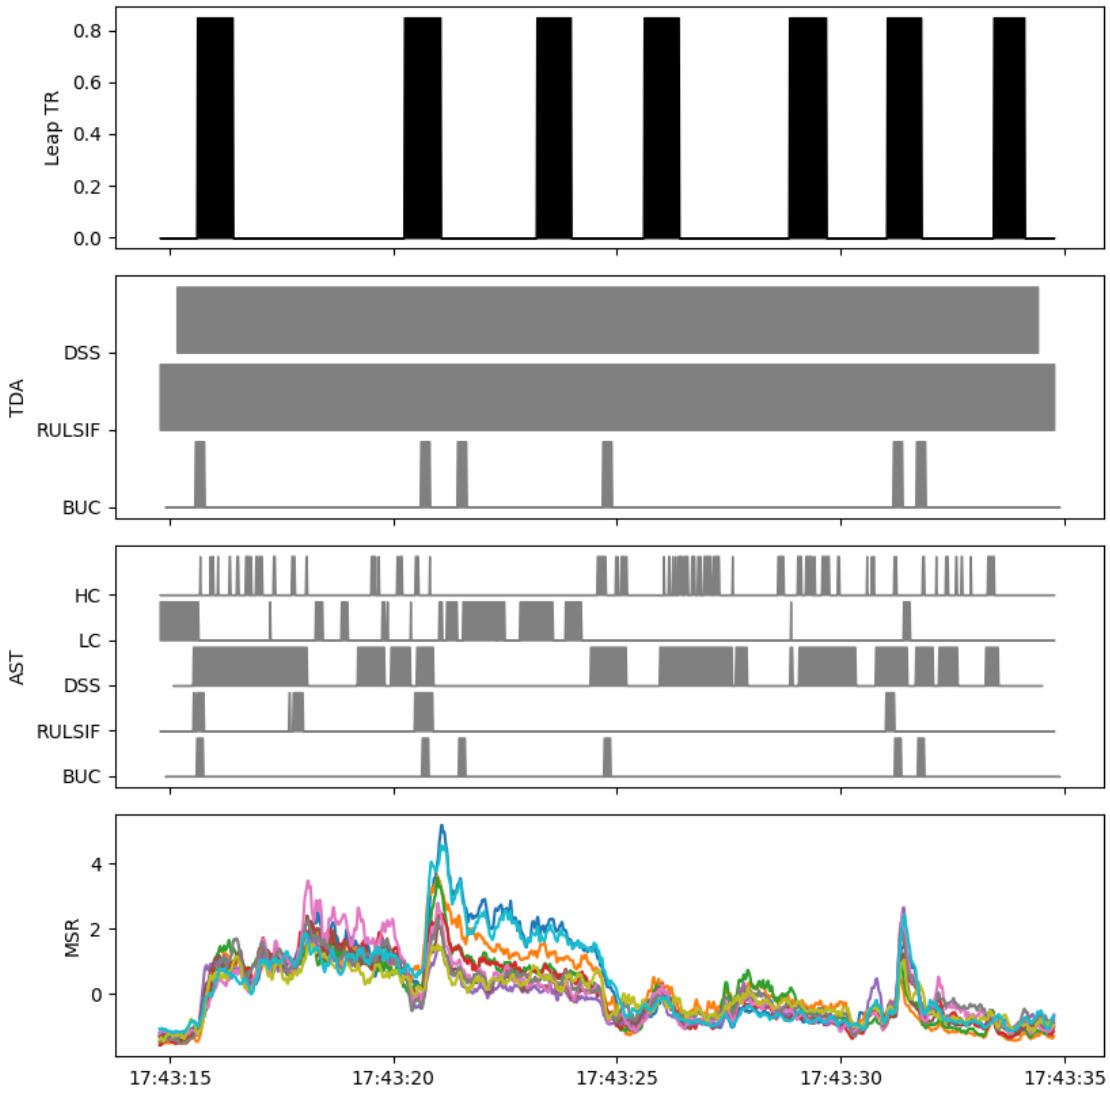

Figure 4.2: Detected transitions for P0 of 3DC data set. (1-top) Leap-identified transitions, (2) TDA-identified transitions when transition detection is optimized, (3) frames avoided during adaptation by ASTs when optimized for myoelectric classifier performance, and (4-bottom) MSR feature of SEMG signal

Table 4.12: Transition detection performance on Raghu data set using parameters optimized for myoelectric classifier performance

Technique	Accuracy	Recall	Precision	F-0.5 Score	F-1 Score	F-2 Score
HC	0.867	0.078	0.533	0.242	0.135	0.094
LC	0.509	0.167	0.056	0.064	0.082	0.117
DSS	0.860	0.112	0.441	0.267	0.174	0.131
RULSIF	0.872	0.133	0.624	0.332	0.211	0.156
BUC	0.867	0.010	0.758	0.046	0.019	0.012

Table 4.13: Transition detection performance on 3DC data set using parameters optimized for myoelectric classifier performance

Technique	Accuracy	Recall	Precision	F-0.5 Score	F-1 Score	F-2 Score
HC	0.632	0.091	0.346	0.213	0.141	0.106
LC	0.586	0.212	0.330	0.291	0.252	0.226
DSS	0.619	0.144	0.374	0.278	0.205	0.163
RULSIF	0.659	0.012	0.444	0.053	0.023	0.015
BUC	0.659	0.004	0.439	0.019	0.008	0.005

3DC data set, respectively. Both the average accuracy of assigned labels during LMC-assigned steady states and the number of LMC-identified steady state frames for which labels were rejected are reported. In both data sets, the LC technique rejected the most SS frames, and the HC and BUC techniques rejected the fewest. Implementing the SSL improved label accuracy for the more complex 3DC data set, but reduced that of the simpler Raghu data set.

4.3 Discussion

The F-β score, which is a harmonic mean of precision and recall, was used in this work as a measure of each TDA's ability to detect transitions. Comparisons of these values are included in Tables 4.10 and 4.11 for the Raghu and 3DC data sets respectively. The accuracy, recall, and precision are also reported for completeness. An F-0.5 score considers the precision to be twice as important as the recall. In such cases, the BUC is the preferred algorithm for the Raghu data set. However, when the recall is considered as, or more, important than precision, DSS is preferred.

Technique		SS Label Accuracy	Rejected LMC-SS Frames
Naïve		92.4%	0
No SSL	HC	92.8%	0
	LC	86.9%	0
	DSS	93.5%	0
	RuLSIF	92.5%	0
	BUC	92.4%	0
25% SSL	HC	83.7%	83
	LC	86.2%	589
	DSS	82.0%	241
	RuLSIF	84.2%	173
	BUC	91.2%	28

Table 4.14: Steady state label accuracy and the number of frames identified as steady state by LMC which were not adapted to due to failing to meet the SSL threshold requirement. Results included for both no SSL and SSL with a threshold of 25% applied to the Raghu data set.

Technique		SS Label Accuracy	Rejected LMC-SS Frames
Naïve		39.3%	0
No SSL	HC	40.9%	0
	LC	34.1%	0
	DSS	41.8%	0
	RuLSIF	39.3%	0
	BUC	39.4%	0
25% SSL	HC	44.1%	531
	LC	40.8%	6818
	DSS	45.0%	3081
	RuLSIF	40.5%	2151
	BUC	39.5%	516

Table 4.15: Steady state label accuracy and number of frames identified as steady state by LMC which were not adapted to due to failing to meet the steady state labelling threshold requirement. Results included for both no SSL and SSL with a threshold of 25% applied to the 3DC data set.

This is because the DSS algorithm has a near-perfect recall of 0.954 even though its precision is worse than BUC. The RuLSIF algorithm, despite its relatively good F-2 score, failed to appropriately identify transitions as it defaults to identifying all frames as transitions. This is shown in Figure 4.1 and reflected in the recall of 1 in Table 4.10.

This failing is caused, in part, by the use of β=2 for the F-β score used in the selection of parameter values. When $\beta > 1$, the recall is prioritized over the precision. In Table 4.3, for both the Raghu and 3DC data set, a threshold of 0.00 was selected for the RuLSIF algorithm. This indicates that all frames would be above this threshold and therefore identified as transitions. While this would result in perfect recall, as seen in Tables 4.10 and 4.11, if implemented in a transition-avoidant AST this would result in no adaptation at all. Although different beta-values could be explored, the underlying reason behind this collapse in the grid search is that the TDA approaches appear to be unable to differentiate between the stochastic nature of the SEMG, transience in the steady state, and true transitions. A combination of the poorly selected parameters and the insensitivity of TDAs for the 3DC data set indicates that it may be more difficult to detect transitions in this data than in the Raghu data set. As a result, the grid-search was unable to identify parameters that would result in good sensitivity while also maintaining adequate precision. This aligns with the expectation that, due to the increased complexity of transitions being performed and the presence of transience during the purportedly steady state, the identification of transitions in the 3DC data set would be more difficult.

Although selecting an F-β score with a lower β-value may change the chosen operating point for TDA, the inherent sensitivity remains poor. However, as the goal of these TDAs was ultimately to inform the selection of data to improve classifier adaptation, parameter optimization for ASTs was then explored based on classifier performance, as detailed in Section 4.1.3. This was performed for all ASTs, both

the TDA transition-avoidant (i.e. DSS, RULSIF, and BUC) and the traditional techniques (i.e. HC and LC). A comparison of these values is included in Tables 4.12 and 4.13 for the Raghu and 3DC data sets respectively. The accuracy, recall, and precision were again reported for completeness.

In Table 4.12, a comparison of the F-β scores achieved by each AST indicates that, regardless of β value, the RuLSIF AST achieved the best results for the Raghu data set. In comparison, Table 4.13 indicates that RuLSIF performed poorly on the 3DC data set. Alternatively, both the LC and DSS ASTs perform the best for the 3DC data set yet achieve middling performance on the Raghu dataset. As no TDA unilaterally outperformed the others, all evaluated ASTs were carried forward in the subsequent work.

While no single AST stood out, an understanding of their behaviour can be achieved through comparison of these values to those in Table 4.10 and 4.11. For the Raghu results, the precision of classifier-optimized AST improved in comparison to the precision of transition detection (i.e. when the F-2 score was optimized). The recall, however, was reduced substantially. This means that more LMC-identified transition frames were being included during adaptation. However, those frames which were avoided during adaptation were more likely to occur during a transition. One likely cause for this behaviour is that the ASTs intentionally included transient behaviour to improve classifier performance. As previously discussed, transient SEMG behaviour occurs both within and outside of transitions. The steady state transient behaviour may prove beneficial to the performance of a myoelectric classifier in the same way that training a classifier on ramp data improves the performance in comparison to training on solely static data. The transience of the SEMG signal may provide a more complete understanding of feature space and the boundaries between classes. Therefore, the ASTs may have included transient SEMG behaviour at the expense of allowing more transitions to be included during adaptation.

These trends do not apply to the LC AST as it avoided adaptation during high-confidence frames which resulted in many steady state frames not being included during adaptation. The precision, and therefore the F-β score, greatly suffered as reported in Table 4.12. In comparison, the HC AST resulted in precision roughly on par with the other TDA-based ASTs. This indicates that low-confidence classifier decisions have a tendency to occur during transitions. However, the HC AST also has low recall which suggests that the majority of LMC-identified transition regions actually maintain a reasonable level of confidence.

The behaviour of the LC AST is most easily understood through consideration of Figures 4.1 and 4.2. These figures provide a comparison of the LMC-identified transitions and the frames of data included during adaptation by the ASTs for both the Raghu and 3DC data set respectively. The MSR of the SEMG signal is also included as a naive indication of SEMG transience. In Figure 4.1 the LC AST, which avoids adaptation during regions of high classifier confidence, often avoided adaptation to steady state frames.

While the HC AST reliably avoids transitions, the duration of these avoidances is notably shorter than the LMC-identified transition. As the F-β score is calculated using a frame-by-frame decision, this performance is seen as slightly worse than DSS. This is because while the DSS TDA identifies fewer transitions, the transitions which are detected occur for longer. Therefore, alternative metrics which may be used to evaluate the TDA's ability to detect transitions may be valuable in achieving a more conclusive comparison between the TDAs.

These TDAs were also used to implement steady state labelling with results included in Table 4.14 and 4.15. As expected, when no SSL was applied, the LC adaptation selection technique resulted in the lowest label accuracy. This is because it adapted only when the classifier's confidence was low - leading to higher rates of mislabelling. The application of the SSL technique to the Raghu data set, which

achieved remarkably low labelling error due to imposed restraints, actually worsens the labelling accuracy. However, when applied to the 3DC data set, which achieves comparatively poor performance, the application of the SSL technique improves or maintains the labelling accuracy for all ASTs. For both data sets, the LC-SSL scheme rejects the most frames from adaptation. This is a logical result as the spans selected by the LC technique have lower confidence and are, therefore, less likely to be able to meet the 25% confidence threshold that one label may be applied to the entire span.

Chapter 5

SEMG-Based Transition-Informed Adaptation

The final aim of this work was to evaluate the SEMG-based transition-informed adaptation techniques developed in Chapter 4 to determine the potential of transition-informed adaptation for myoelectric control. The developed TDAs were used as data selection techniques that labelled and selected frames from the continuous transition SEMG stream for use in both SAL and UAL. The SAL approach was compared to a classifier adapted to all incoming data using available prompt information. The UAL approaches were compared to no adaptation case and a naively adapted classifier (adapting to all data based on SEMG classification outputs). The susceptibility of these adaptation schemes to changes in the importance of frames during adaptation and the quality of initial classifier training were also explored.

5.1 Methods

5.1.1 Data Set Selection

This study extended the work in Chapter 4 by applying the developed TDAs in transition-avoidant ASTs. Both the Raghu and 3DC datasets were included in this study. As in Chapter 4, the SEMG data (from both datasets) were segmented into 160 ms long frames with 16 ms overlap, and the fractal length (MFL), Wilson's amplitude (WAMP), mean-square root (MSR), and 2 nd L-moment features were extracted from each frame. LMC data were available from both data sets and used in this portion of the study to assign adaptation labels in the SAL adaptation schemes.

The Raghu dataset, described in Section 2.3.1.1, was collected as subjects continuously transitioned between 7 motion classes while maintaining constant limb position and contraction intensity. Because the data collection was well controlled, there is little transient behaviour outside of transitions in this data set and the TDAs were able to accurately identify transitions. The relatively good performance of the classifier, however, limits the potential improvement of adaptation.

In comparison, the Long-Term 3DC data, introduced in Section 2.3.1.2 was collected over the span of several weeks. When performing evaluation trials, VR prompts with randomized limb position, contraction intensity, and gesture were displayed every 5 seconds. Subjects received both positional feedback from the LMC and contraction intensity feedback from the MAV of the SEMG signal, with the intention of ensuring compliance with the prompt. The division of subject focus and continual adjustments in contraction, however, led to increased SEMG transience in regions of steady state. As a result, the TDAs in Chapter 4 struggled to adequately identify transitions. Despite this, the poor initial classifier performance allows an opportunity for adaptation to provide further benefits.

5.1.2 Adaptation

A transition-informed adaptation, as described in Section 4.1.4 was used to adapt an LDA classifier. Based on the results of the studies detailed in Chapters 3 and 4, a transition-avoidant scheme was adopted with the SSL technique applied with a 25% confidence threshold. This corresponds to the adaptation weights, $W[n]$, being set as $W_{SS-Label}[n]$, as described in Equation 4.8, and the adaptation labels, $L[n]$, being set as $L_{SS-Label}[n]$, as described in Equation 4.7. The value of W_{init} was set to multiple values as described later in this section. Adaptation occurred in 10-second windows, where transitions were identified, labels were assigned, and then covariate shift adaptation was used to update the classifier.

The performance using each of the ASTs introduced in Chapter 4 was compared against the following adaptation schemes to gauge effectiveness:

- **Naive**: Adapt to all frames of available data using classifier-assigned labels (L[n] = CD[n]).
- **No Adaptation**: The initial training remained static without any adaptation updates.
- **Prompt**: Adapt to all frames of available data using the displayed prompts as adaptation labels (L[n] = P[n]).

As the prompt adaptation scheme mimics an SAL strategy, it was compared to transition-informed adaptations that used the LMC-assigned adaptation label. Besides the TDAs, two other experiment parameters were manipulated in this study to determine their influence on performance, as described below.

Quality of Initial Training: Before adaptation occurs, the initial LDA classifier was first trained using available ramp data. To evaluate the sensitivity of the TDA-informed adaptation schemes to the quality of this initial training, the amount of training data used was controlled. Given the differences in the Raghu and 3DC data

sets, and the higher performance ceiling of the Raghu set (and thus larger dynamic range of performances), it was used for this evaluation. Three starting points were compared: A baseline that included all 6 trials of ramp data during initial training, a single trial initial training that used 1 randomly selected trial, and a restricted case where only half of a randomly chosen trial was used. When using half of one trial, half of the frames belonging to each class were randomly selected for inclusion.

Adaption Rate When using AST schemes that choose which frames to use for adaptation, frames that are excluded can simply be assigned adaptation weights of zero - meaning no adaptation occurs. However, when frames are selected for adaptation, the associated weights may be scaled to control their relative importance as compared to the initial training. This is controlled by the W_init parameter. In this work, these weightings were chosen so that when the importance was 100%, the effect of one frame of adaptation was equal to that of one frame of training data. When the importance was 50%, the effect of one frame of adaptation was half that of one frame of training data. It is worth noting that the importance does not describe the ratio between the cumulative effect of all adaptation frames to that of all initial training (i.e. an importance of 50% does not mean that the effect of all adaptation which occurred is equal to that of all training). It simply describes the importance of one frame of adaptation data relative to one frame of training data. Modifying this value allows for an exploration of the effect of changing the ratio of adaptation to the initial training without having to reduce the amount of data included in the initial training (which is known to reduce classifier performance). Values of 50%, 100%, and 200% were compared in this study.

5.1.3 Performance Evaluation

As in Chapters 3 and 4, the classifier undergoing adaptation was continuously evaluated using the adaptation window that followed the one being processed. This was

to simulate a real-world causal implementation wherein a classifier is adapted on the fly, and subsequent frames are classified using the updated model before then being used to adapt it again. The steady state and transition metrics used by Raghu et. al. [45] and detailed in Section 2.3.3 were used to compare the quality of the myoelectric classifiers resulting from the various adaptation strategies:

- Steady State Metrics: TER, AER, INS_SS
- Transition Metrics: OFF, ON, DUR, INS_TR, TCE, PNM

A two-sided paired t-test was performed to compare each metric for the different adaptation strategies to that obtained by naive adaptation. In addition, a Cohen D test was performed to evaluate the effect size associated with each comparison. As is common, a Cohen D result of 0.2 to 0.5 was interpreted as a small effect, 0.5 to 0.8 as a medium effect, and >0.8 as a large effect [12].

5.2 Results

The results of transition-informed SAL applied to the Raghu data set are included in Tables 5.1 to 5.5. The results for UAL on this data set are included in Tables 5.6 to 5.10. The colouring was assigned based on the standard deviation-normalized difference between each metric and the naive adaptation case to aid in visualization but does not signify statistical significance or effect size. Statistical significance ($\alpha < 0.05$) is instead denoted by an asterisk (*) with effect size designated by carets (ˆ- small effect, ˆˆ- medium effect, ˆˆˆ- large effect). The results for the 3DC data set showed similar trends as those in the Raghu data set so, for conciseness, they are included in Appendix B.

5.2.1 Supervised Adaptation Results

Table 5.1 includes the results for SAL adaptation when all training data were included and adaptation frames were assigned the same importance as training frames (i.e. $W_{init} = 1.0$). All data selection techniques led to significant improvements in all SS metrics, although only LC and LC-SSL yielded a noticeable (though small) effect size. This trend continued in Tables 5.2 to 5.5, which show various combinations of SAL using 50%, 100%, and 200% adaptation weightings, and 6, 1, and 0.5 training trials. Despite these different adaptation rates and initial amounts of training data, the LC-SSL outperformed the other data selection techniques.

	Prompt	No Adapt	HC	HC-SSL	LC	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	3.1%	8.0%*^^^	2.9%*	2.9%*	2.5%*^	2.5%*^	3.0%*	2.9%*	2.9%*
TER	3.7%	9.7%*^^^	3.4%*	3.4%*	3.0%*^	3.0%*^	3.5%*	3.4%*	3.4%*^
INS_SS	1.6%	2.9%*^^^	1.5%*	1.5%*	1.4%*^	1.4%*^	1.5%*	1.5%*^	1.5%*
ON (ms)	451.85	502.95*^^	453.63	453.63	432.91*^	432.91*^	458.51*	456.21*	451.88
OFF (ms)	290.09	283.64	296.94*	296.94*	291.20	291.20	299.74*^	298.33*	296.87*
DUR (ms)	161.76	219.31*^^^	156.70*	156.70*	141.72*^	141.72*^	158.77*	157.88*	155.01*
INS_TR	8.0%	8.1%	7.9%	7.9%	8.0%	8.0%	8.0%	8.0%	8.0%
TCE	39.5%	40.1%	39.3%	39.3%	38.8%	38.8%	39.7%	39.4%	39.2%
PNM	15.9%	17.5%	15.4%*	15.4%*	14.1%*^	14.1%*^	15.4%*	15.7%	15.3%*

Table 5.1: Classifier metrics resulting from SAL adaptation after 6 trials of initial ramp training and an adaptation weight of $W_{init} = 1.0$. LMC-assigned labels with SSL (th=25%).

	Prompt	No Adapt	HC	HC-SSL	LC	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	3.4%	8.0%*^^^	3.1%*	3.1%*	2.7%*^	2.7%*^	3.2%*	3.1%*^	3.1%*^
TER	4.0%	9.7%*^^^	3.7%*^	3.7%*^	3.3%*^	3.3%*^	3.8%*	3.7%*^	3.6%*^
INS_SS	1.7%	2.9%*^^^	1.6%*	1.6%*	1.6%*^	1.6%*^	1.6%	1.6%*	1.6%*
ON (ms)	456.96	502.95*^^	459.64	459.64	443.80*	443.80*	462.39*	461.21*	456.71
OFF (ms)	292.00	283.64	299.18*	299.18*	295.17	295.17	300.21*	299.93*	298.17*
DUR (ms)	164.97	219.31*^^^	160.47*	160.47*	148.63*^	148.63*^	162.19*	161.28*	158.53*
INS_TR	8.0%	8.1%	8.0%	8.0%	8.1%	8.1%	8.0%	7.9%	7.9%
TCE	39.6%	40.1%	39.3%	39.3%	39.2%	39.2%	39.4%	39.3%	39.0%*
PNM	16.0%	17.5%*	15.8%	15.8%	14.6%*	14.6%*	15.8%	16.2%	15.8%

Table 5.2: Classifier metrics resulting from SAL adaptation after 6 trials of initial ramp training and an adaptation weight of $W_{init} = 0.5$. LMC-assigned labels with SSL (th=25%).

5.2.2 Unsupervised Adaptation Results

Table 5.6 includes the results for UAL adaptation when all ramp trials were included during training and adaptation frames were given an equal weighting to those from

	Prompt	No Adapt	HC	HC-SSL	LC	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	3.0%	8.0%*^^^	2.9%*	2.9%*	2.4%*^	2.4%*^	2.9%*	2.8%*	2.8%*
TER	3.6%	9.7%*^^^	3.4%*	3.4%*	2.9%*^	2.9%*^	3.4%*	3.3%*	3.3%*^
INS_SS	1.6%	2.9%*^^^	1.5%*	1.5%*	1.4%*^	1.4%*^	1.5%*	1.5%*	1.5%*
ON (ms)	448.06	502.95*^^^	449.70	449.70	426.34*^	426.34*^	454.82*	453.18*	447.92
OFF (ms)	288.55	283.64	295.40*	295.40*	288.94	288.94	297.98*^	296.83*	294.81*
DUR (ms)	159.51	219.31*^^^	154.30*	154.30*	137.40*^^	137.40*^^	156.84*	156.35*	153.11*
INS_TR	7.9%	8.1%^	7.9%	7.9%	7.9%	7.9%	8.0%	7.9%	7.9%
TCE	39.2%	40.1%	39.3%	39.3%	38.3%	38.3%	39.6%	39.5%	39.3%
PNM	15.8%	17.5%^	15.0%*	15.0%*	14.0%*^	14.0%*^	15.1%*	15.4%	14.9%*

Table 5.3: Classifier metrics resulting from SAL adaptation after 6 trials of initial ramp training and an adaptation weight of $W_{init} = 2.0$. LMC-assigned labels with SSL (th=25%).

	Prompt	No Adapt	HC	HC-SSL	LC	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	3.0%	12.6%*^^^	2.8%*	2.8%*	2.5%*^	2.5%*^	2.9%*	2.8%*	2.8%*
TER	3.6%	15.2%*^^^	3.3%*	3.3%*	3.0%*^	3.0%*^	3.4%*	3.3%*	3.3%*^
INS_SS	1.6%	3.9%*^^^	1.5%*	1.5%*	1.6%	1.6%	1.5%*	1.5%*	1.5%*^
ON (ms)	446.46	528.66*^^^	448.84	448.84	425.29*^	425.29*^	452.14*	452.23*	447.03
OFF (ms)	288.13	272.71^	295.46*	295.46*	287.69	287.69	296.53*	297.04*^	294.93*
DUR (ms)	158.32	255.96*^^^	153.38*	153.38*	137.61*^^	137.61*^^	155.61*	155.19*	152.10*
INS_TR	8.0%	8.0%	8.0%	8.0%	7.9%	7.9%	8.0%	8.0%	8.0%
TCE	39.2%	40.5%	39.5%	39.5%	38.8%	38.8%	39.7%*	39.7%	39.4%
PNM	15.7%	16.1%	14.8%*	14.8%*	13.1%*^	13.1%*^	14.7%*	15.1%*	14.7%*

Table 5.4: Classifier metrics resulting from SAL adaptation after 1 trial of initial ramp training and an adaptation weight of $W_{init} = 1.0$. LMC-assigned labels with SSL (th=25%).

	Prompt	No Adapt	HC	HC-SSL	LC	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	3.0%	18.4%*^^^	2.8%*	2.8%*	2.7%*^	2.7%*^	2.9%*	2.8%*	2.8%*
TER	3.6%	20.8%*^^^	3.3%*	3.3%*	3.1%*^	3.1%*^	3.4%*	3.3%*	3.3%*^
INS_SS	1.6%	5.1%*^^^	1.6%	1.6%	1.6%	1.6%	1.5%*	1.5%*	1.5%
ON (ms)	445.75	527.24*^^^	449.45*	449.45*	424.70*^	424.70*^	452.00*	451.66*	446.17
OFF (ms)	287.44	252.56*^^^	296.67*^	296.67*^	285.63	285.63	297.50*^	297.25*^	295.05*
DUR (ms)	158.31	274.68*^^^	152.78*	152.78*	139.07*^	139.07*^	154.50*	154.40*	151.12*
INS_TR	8.0%	8.0%	8.0%	8.0%	7.9%	7.9%	8.0%	8.0%	8.0%
TCE	39.4%	43.2%*^^	39.8%	39.8%	40.0%	40.0%	40.1%*	39.9%*	39.6%
PNM	15.5%	10.9%*^^	14.5%*	14.5%*	11.4%*^^	11.4%*^^	14.3%*	14.8%*	14.5%*

Table 5.5: Classifier metrics resulting from SAL adaptation after 1 trial of initial ramp training and an adaptation weight of $W_{init} = 0.5$. LMC-assigned labels with SSL (th=25%).

training (i.e. $W_{init} = 1.0$). The results for the HC and LC techniques are included for both when the SSL was applied with a threshold of 25%, and when it wasn't. Both HC-SSL and LC-SSL yielded improved metrics compared to their non-SSL counterparts, suggesting that the temporal context afforded by SSL improved the labelling for adaption. While HC-SSL and LC-SSL generally outperformed naive adaptation, the DSS, RuLSIF, and BUC TDA-based approaches yielded few improvements. Of those metrics which did significantly improve, only the transition duration for DSS

achieved an appreciable effect size. This trend again continued through Tables 5.7 to 5.10 which show various combinations of UAL using 50%, 100%, and 200% adaptation weightings, and 6, 1, and 0.5 training trials.

	Naive	No Adapt	HC	HC-SSL	LC	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	5.2%	8.0%*^^	5.2%	4.8%	4.9%	4.6%*	5.5%	5.5%	5.1%
TER	6.1%	9.7%*^^^	6.1%	5.6%	5.8%*	5.4%*	6.4%	6.5%	6.1%
INS_SS	1.7%	2.9%*^^^	1.8%*	1.7%	1.7%	1.6%	1.9%*^	1.8%	1.8%
ON (ms)	496.16	502.95	495.17	477.80*^	484.30*	461.56*^	482.43*	486.60*	494.49
OFF (ms)	293.18	283.64*	293.10	290.07	292.59	284.35*	291.10	290.13	289.96
DUR (ms)	202.98	219.31*^	202.07	187.73*^	191.71*	177.21*^	191.33*^	196.47*	204.52
INS_TR	7.3%	8.1%*^^	7.3%	7.5%	7.4%	7.5%^	7.6%*^	7.6%*^	7.3%
TCE	41.7%	40.1%	41.6%	39.9%*	41.9%	39.4%*^	40.5%	41.0%	40.4%*
PNM	15.9%	17.5%*	15.8%	15.3%	14.6%*	14.1%*	15.3%	15.3%	17.0%*

Table 5.6: Classifier metrics resulting from UAL adaptation after 6 trials of initial ramp training and an adaptation weight of $W_{init} = 1.0$. Classifier-assigned labels with SSL (th=25%).

	Naive	No Adapt	HC	HC-SSL	LC	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	5.2%	8.0%*^^^	5.2%	5.0%	5.1%	4.8%	5.3%	5.3%	5.1%
TER	6.1%	9.7%*^^^	6.1%	5.9%	6.0%	5.7%	6.2%	6.3%	6.0%
INS_SS	1.8%	2.9%*^^^	1.8%	1.9%*	1.8%	1.7%	1.9%*^	1.8%	1.8%
ON (ms)	494.78	502.95	493.95	477.20*^	489.92	465.82*^	486.00*	487.55*	494.83
OFF (ms)	293.60	283.64*	293.62	290.32	294.00	284.14*	290.38	290.29	292.09
DUR (ms)	201.18	219.31*^	200.32	186.89*^	195.93*	181.68*^	195.61*	197.27	202.74
INS_TR	7.4%	8.1%*^^	7.4%	7.7%*^	7.6%	7.5%	7.7%*^	7.6%*	7.4%
TCE	41.3%	40.1%	41.2%	40.0%*	41.9%	39.5%*	41.0%	40.8%	40.4%*
PNM	16.3%	17.5%	16.3%	15.5%*	15.2%*	14.5%*	15.4%	15.7%	17.2%*

Table 5.7: Classifier metrics resulting from UAL adaptation after 6 trials of initial ramp training and an adaptation weight of $W_{init} = 0.5$. Classifier-assigned labels with SSL (th=25%).

	Naive	No Adapt	HC	HC-SSL	LC	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	5.4%	8.0%*^^	5.4%	5.3%	5.0%*	4.9%*	5.4%	6.9%*^	5.5%
TER	6.4%	9.7%*^^	6.3%	6.1%	5.9%*	5.7%*	6.3%	7.9%*^	6.5%
INS_SS	1.7%	2.9%*^^^	1.7%	1.7%	1.8%	1.7%	1.9%*^	1.8%	1.8%
ON (ms)	495.99	502.95	493.22*	477.54*^	483.09*	456.66*^^	478.73*^	478.97*^	493.83
OFF (ms)	288.95	283.64	287.61	288.47	293.80	281.89	288.94	279.32*	287.07
DUR (ms)	207.04	219.31*^	205.61*	189.07*^	189.29*^	174.77*^^	189.79*^	199.65*	206.76
INS_TR	7.1%	8.1%*^^^	7.1%	7.5%*^	7.4%*^	7.6%*^	7.6%*^	7.4%	7.2%
TCE	42.0%	40.1%*	41.9%	39.7%*^	41.6%	38.7%*^	40.4%	40.7%	40.4%*
PNM	15.4%	17.5%*	15.3%	15.3%	14.4%	13.9%*	15.1%	14.5%	16.8%*

Table 5.8: Classifier metrics resulting from UAL adaptation after 6 trials of initial ramp training and an adaptation weight of $W_{init} = 2.0$. Classifier-assigned labels with SSL (th=25%).

	Naive	No Adapt	HC	HC-SSL	LC	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	8.9%	12.6%*^^	8.8%	8.4%	8.5%	9.7%	9.0%	10.6%*^	9.0%
TER	10.1%	15.2%*^^	10.0%	10.4%	9.7%	11.8%^	10.9%	12.7%*^	10.9%
INS_SS	1.9%	3.9%*^^^	1.9%*	1.9%	2.1%*^	2.0%	2.1%*^	2.0%	2.1%*^
ON (ms)	505.26	528.66*^	502.94	478.25*^	501.39	461.60*^^	483.73*^	481.99*^	504.11
OFF (ms)	274.94	272.71	275.14	273.68	285.44*	258.92*^	274.85	263.70	274.62
DUR (ms)	230.33	255.96*^	227.80*	204.56*^	215.94*^	202.68*^	208.87*^	218.29^	229.49
INS_TR	6.7%	8.0%*^^^	6.7%	7.0%^	7.1%*^	6.6%	6.9%	6.7%	6.9%^
TCE	42.2%	40.5%	42.2%	38.1%*^	41.3%	37.4%*^	38.9%*^	39.7%*^	40.6%*
PNM	13.5%	16.1%*^	13.5%	14.7%	13.6%	12.5%	13.8%	12.9%	14.9%*

Table 5.9: Classifier metrics resulting from UAL adaptation after 1 trial of initial ramp training and an adaptation weight of $W_{init} = 1.0$. Classifier-assigned labels with SSL (th=25%).

	Naive	No Adapt	HC	HC-SSL	LC	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	12.7%	18.4%*^^	12.8%	13.5%	12.4%	12.6%	13.5%	15.3%*^	15.5%*^
TER	14.1%	20.8%*^^	14.1%	15.2%	13.4%	14.7%	15.1%	18.8%*^^	18.4%*^^
INS_SS	2.2%	5.1%*^^^	2.2%	2.2%	2.5%*^	2.3%	2.3%	2.1%	2.5%*^
ON (ms)	497.10	527.24*^	493.53	463.51*^	511.34*	472.90*^	465.36*^	461.79*^	489.86
OFF (ms)	252.75	252.56	253.55	250.62	269.97*^	242.46	247.85	242.70	236.09*^
DUR (ms)	244.35	274.68*^	239.98	212.89*^	241.37	230.44^	217.52*^	219.09*^	253.78
INS_TR	6.3%	8.0%*^^^	6.3%	6.5%	6.9%*^	6.4%	6.4%	6.2%	6.2%
TCE	43.0%	43.2%	42.6%*	37.8%*^^	44.7%^	39.7%*^	38.5%*^^	35.8%*^^^	39.2%*^^
PNM	10.6%	10.9%	10.6%	11.8%	9.2%	10.2%	11.5%	13.0%*^	12.1%

Table 5.10: Classifier metrics resulting from UAL adaptation after 1/2 trial of initial ramp training and an adaptation weight of $W_{init} = 1.0$. Classifier-assigned labels with SSL (th=25%).

5.3 Discussion

Table 5.1, which shows the baseline results when using an adaptation weighting of $W_{init} = 1.0$ and all 6 trials of ramp data were included during initial training for the Raghu data set, indicates that incorporating some form of adaptation selection is beneficial in SAL. Of those techniques considered, LC-SSL resulted in the best performance. It was resilient to changes in adaptation rate and the quality of initial training, as shown in Tables 5.2 to 5.5. Table 5.1 also includes a comparison of both HC-SSL and LC-SSL with their non-SSL counterparts that shows that there is no appreciable change in performance when implemented in a supervised environment. This is likely because the LMC-assigned adaptation labels are already consistent throughout the LMC-identified steady states and, therefore, the regions likely to be identified as steady states by the ASTs. The only labels that are likely to be modified by the SSL technique would be those near the edges of LMC-identified transitions.

When applied to the Raghu dataset, as shown in Table 5.1, all adaptation selection techniques for SAL resulted in general improvements to the myoelectric classifier metrics. While ON, OFF, and PNM are all shown to experience significant worsening, the effect size is almost always either small or negligible. Although all of the evaluated data selection techniques resulted in improved classifier performance, the LC-SSL technique performed best, yielding a significant improvement to all steady state metrics as well as ON and DUR.

In contrast to these findings, when SAL was applied to the 3DC data set, several techniques worsened classifier performance. This is shown in Tables B.1 to B.3 in Appendix B, where both HC-SSL and DSS consistently result in significant worsening of nearly all metrics. This degradation, even when using SAL, suggests that the transient behaviours observed in the D3C data set (due to the provision of feedback) may have introduced conflicting information between the SEMG data and LMC labels. This could be due to differences between force-related EMG activation and the LMC-derived hand kinematics, especially for low levels of contraction. Despite these differences, the LC-SSL technique again resulted in significant improvements to AER, TER, DUR, and PNM with a small effect size with the only significantly declining metric being OFF with a negligible effect size. As a result, it is recommended that the LC-SSL technique be used in SAL applications as it appears to provide improved performance generalized at least over 2 data sets.

While comparing the HC-SSL and LC-SSL to their non-SSL counterparts resulted in no appreciable difference when applied in SAL, this is not true for UAL. As indicated in Table 5.6, the inclusion of the SSL technique led to all metrics but INS_TR being either maintained or improved for both HC and LC. When applied to the 3DC dataset, included in Table B.4, HC-SSL yielded relatively minor improvements in comparison to LC-SSL which mitigated the decline in ON, OFF, INS_TR, and TCE metrics seen when solely applying LC. Therefore, it is recommended that

the SSL technique be incorporated during both HC and LC unsupervised adaptation. Furthermore, LC should not be implemented without SSL as it may result in a significant decline in several classifier metrics. This is in agreement with Sensinger et al., who found that the high mislabelling error associated with traditional LC prohibited it from implementation in UAL [44].

A comparison between the various adaptation selection techniques in Table 5.6 indicates that both HC-SSL and LC-SSL outperform all other techniques for the Raghu data set - with LC-SSL slightly outperforming HC-SSL. In Tables 5.8 and 5.7 this trend is maintained as the importance of adaptation is varied. However, Table 5.9 and 5.10 indicate a slight preference for HC-SSL when the classifier's initial training is restricted to fewer samples. This could suggest that the confidence profiles of classifiers trained with too few samples aren't sufficiently robust to support the 25% SSL approach when applied to LC. HC-SSL and LC-SSL also continued to result in improved classifier performance when utilizing the 3DC data set, as indicated in Table B.4 found in Appendix B.

Interestingly, despite the hypothesized advantage of leveraging transition information to inform adaptation, the TDAs largely failed to improve classifier performance. The DSS technique yielded significant improvements to only TER and INS_SS, and RuLSIF and BUC resulted in very few significant improvements in comparison to the naive adaptation case. In comparison, the LC-SSL, which only includes low-confidence data during adaptation, experienced the best results. This is contrary to the expected results as low-confidence decisions are more likely to contain transient behaviour and therefore occur during transitions. A likely cause for these results is the addition of the SSL technique which rejects adaptation during regions where one adaptation label cannot be assigned with sufficient confidence. When combined together, the LC selection often incorporates portions of transience in adaptation but the SSL technique prevents transition regions where labels cannot

be appropriately assigned. Therefore, LC-SSL predominantly focused adaptation on learning from transience that occurred within steady state regions. In future work, the benefits of adapting to appropriately labelled transient SEMG data may be reproduced, or even improved upon, by inverting the transition detection algorithms. That is, the TDAs could be used to identify and include SEMG transience during adaptation. Then the SSL technique could be implemented to both assign appropriate labels and prevent adaptation when no label can be selected.

With the results from both data sets considered, it is therefore recommended that LC-SSL be used in most applications as it provides slightly better performance. However, if the amount of initial training data supplied to the classifier is limited, then HC-SSL should be implemented as it is not as susceptible to these changes. The moderate performance improvement afforded by SSL when combined with these approaches suggests that leveraging temporal context – or other forms of context – to improve the assignment of adaptation labels in UAL may be a viable path worth further exploration.

Chapter 6

Conclusions and Future Work

The overarching aim of this work was to explore whether transition-informed adaptation may benefit the performance of myoelectric control. Because very little work has explored the inclusion of transition data in the training and adaptation of myoelectric classifiers, this was explored within the context of two data sets comprised of collections of SEMG and LMC position data as participants transitioned continuously between target classes. These data were first used to investigate the effect of including various portions of the transition regions during adaptation. This study showed that when supervised adaptation labels are known, as determined by the LMC, including some sections of the transitions can improve the responsiveness (i.e. ON, OFF, and DUR metrics) of the resulting classifier. However, as more adaptation occurs during these transitions, the steady-state control metrics may degrade.

The initial exploration leveraged the presumed-to-be ground truth LMC kinematic data to both segment and assign adaptation labels. It was not yet clear whether, a) the segmentation could be replicated in a more realistic scenario using only the SEMG data stream, or b) whether the results might hold due to possible differences in classifier-assigned labels. Consequently, several transition detection algorithms, informed solely by the SEMG data, were evaluated. While these TDAs

performed adequately on the Raghu data set, the overall lack of sensitivity and specificity in separating transitions from the transient and stochastic nature of SEMG prevented the selection of appropriate algorithm hyper-parameters. The TDA parameters were therefore re-computed to optimize the performance of the classifiers adapted in a transition-avoidant scheme informed by the TDAs. Because this meant that the TDA algorithms were optimized to improve the end performance, their ability to detect transitions was reduced. Nevertheless, this enabled more meaningful adaptation selection to occur in both data sets, suggesting that the TDAs are capable of identifying transience but not transitions. In addition, the intentional inclusion of some transience, at the expense of including some transitions, during adaptation is beneficial to the performance of the myoelectric classifier. This agrees with previous work that has shown that including some variations in training through the use of ramp contractions improves overall performance.

In addition to TDAs, a temporally aware labelling technique (again solely reliant on SEMG data) was also developed. Instead of explicitly incorporating or excluding transitions, this technique used knowledge of transience to identify sequences of steady-state frames. Knowing that these frames likely shared a common user intent, this steady-state labelling approach resulted in improved accuracy in the assignment of adaptation labels for the 3DC data set.

Finally, the developed adaptation data selection techniques were evaluated within a full adaptation framework across two continuous transition data sets. Both supervised and unsupervised approaches were compared, along with variations in the adaptation rates and amounts of initial training data. It was generally determined that data selection techniques often led to marginal but significant improvements across various performance metrics. Although the TDA-based approaches failed to yield meaningful improvements, the implementation of the temporally- and transience-informed SSL approach improved some performance metrics when applied to con-

ventional HC and LC. In particular, combining the SSL approach with LC adaptation achieved the best performance. This suggests that leveraging additional context - in this case, the knowledge of presumably shared labels across a collection of frames - could improve the labelling of the decision stream for UAL.

One major limitation of this work is the unknown relationship between the classifier metrics and the usability of the corresponding prosthesis. This issue reoccurs throughout all of the individual studies, as it continues to do within the broader field. When identifying how transitions should be incorporated during adaptation, the trade-off between the steady state metrics and the responsiveness of the classifier was not fully understood. Preference was given to the steady-state metrics resulting in the exclusion of transitions from adaptation when using the TDA approaches. This was due in part to the importance of accuracy during sustained intentional contractions, but also because labelling transitions is problematic, as they are undefined (between target classes). In Chapter 4, however, the selection of parameter values to optimize for classifier performance required a single value to describe classifier performance. A weighted combination of the included metrics was therefore used, with the weights assigned to each metric based on an intuitive, but subjective, understanding of their importance to classifier usability. Finally, in the comparison of ASTs in Chapter 5, similar to the results in Chapter 3, it was difficult to provide a decisive recommendation as the true trade-off between metrics is unknown. Future work should, therefore, focus on bridging the gap between the offline classifier metrics and the usability of the corresponding prosthesis. Offline metrics were used in this study to allow for multiple adaptation schemes to be compared. This culminated in the recommendation of Chapter 5 to use an LC-SSL AST in both SAL and UAL adaptation. A usability study, such as a Fitt's law test, is feasible for a single, or small number, of adaptation schemes (along with a baseline comparison) and would verify that LC-SSL improves classifier usability. This closed-loop usabil-

ity study would also confirm the stability of co-adaptation between the user and the device. A usability study, however, would not be feasible in the optimization of AST parameters in Chapter 4. Therefore, a link between the offline classifier performance metrics and the corresponding usability is still suggested.

Another limitation of this work was the implementation of TDAs for transition detection. While this work did evaluate the ability of TDAs to detect transitions, the aim was the implementation of these TDAs in choosing frames for inclusion in transition-informed adaptation. Therefore, most of the work focused on TDAs with parameters optimized for classifier performance. If the proposed TDAs are intended for different purposes in other work, alternate methods for parameter selection should be explored.

The characteristics of transitions and transient behaviour varied widely between the two data sets included in this work. It is highly likely, therefore, that they would continue to vary for additional data collection protocols. This is especially true for both SEMG data collected during activities of daily life (ADL) and SEMG data collected from amputee subjects. Data collected during the completion of ADLs may contain varying amounts of transitions, and likely high degrees of low-level transience. The data sets used in this work prompted subjects to perform transitions at specific intervals. In the first study of this work, it was found that the ratio between the amount of adaptation occurring during steady state and transitions affects the resulting classifier performance. As this work focused on transition-avoidant adaptation, this was not a concern. However, if a transition-inclusive adaptation scheme is to be implemented, the optimal ratio should be explored.

In comparison to the data sets used in this study, which were collected from able-bodied subjects, amputee subjects are known to be less able to maintain and reproduce SEMG patterns. This, in combination with socket movement during changes in limb position, may result in additional transience occurring during steady state

regions. This will affect both the ability of TDAs to detect transitions and also the potential benefits of implementing transience-informed adaptation. Yeung et al. found that SAL was less beneficial for amputee subjects as the inability to maintain and reproduce contractions results in mislabelling errors even during supervised adaptation [53]. The ASTs included in this study were shown to improve SAL for able-bodied subjects by avoiding adaptation during SEMG transience. Future work should include a study of the result of including these ASTs in SAL for amputee subjects. The inability to maintain consistent contractions is likely to result in transience in the SEMG signal which may be detected and avoided. In addition, non-classifier-dependent ASTs may also be used to better select initial training data.

In summary, this work explored various transition-informed approaches to improve the adaptation of myoelectric control. Throughout the work, it was found that knowledge of transience in general, may also provide additional benefits. The impact of selecting subsets of data for adaptation when a ground-truth label was known was first explored. Subsequently, the ability to detect these subsets directly from the SEMG decision stream was explored. Finally, the ability to detect transitions, and conversely, to detect transience as a way of improving steady-state labelling, was explored in terms of its impact on classifier performance after adaptation. Consequently, the main contributions of this work are as follows:

1. This work produced evidence that suggests that transition information could be useful in improving adaptation, whether through explicit recognition of transitions or through awareness of transience in the signal.

2. Three different SEMG-based transition detection algorithms and a new transition-informed steady-state labelling technique were developed and evaluated. When using TDA parameters optimized for adaptation selection, fewer transitions are identified than by the TDAs when optimized for transition detection. This is likely to allow for adaptation to steady state transience at the risk of also in-

cluding transitions. The implementation of the steady-state labelling technique resulted in improved adaptation label accuracy but the confidence threshold needs to be set appropriately to ensure an appropriate amount of adaptation can occur.

3. Seven adaptation schemes, including those based on transition detection and improved steady-state labelling, were evaluated. The results suggest that the additional context afforded by SSL, as derived from an awareness of transitions and transience, improved the viability of LC-SSL for both UAL and SAL adaption.

While many questions remain unanswered, these contributions provide a strong starting point for future work to explore transition- and transience-informed adaptation.

Bibliography

[1] Mohammad Abdoli-Eramaki, Caroline Damecour, John Christenson, and Joan Stevenson, *The effect of perspiration on the semg amplitude and power spectrum*, Journal of Electromyography and Kinesiology **22** (2012), no. 6, 908–913.

[2] Samaneh Aminikhanghahi and Diane J. Cook, *A survey of methods for time series change point detection*, Knowledge and Information Systems **51** (2017), 339–367.

[3] Sebastian Amsuss, Peter M. Goebel, Ning Jiang, Bernhard Graimann, Liliana Paredes, and Dario Farina, *Self-correcting pattern recognition system of surface emg signals for upper limb prosthesis control*, IEEE Transactions on Biomedical Engineering **61** (2014), 1167–1176.

[4] Mojisola Grace Asogbon, Oluwarotimi Williams Samuel, Yanjuan Geng, Olugbenga Oluwagbemi, Ji Ning, Shixiong Chen, Naik Ganesh, Pang Feng, and Guanglin Li, *Towards resolving the co-existing impacts of multiple dynamic factors on the performance of EMG-pattern recognition based prostheses*, Computer Methods and Programs in Biomedicine **184** (2020), 105278.

[5] Michele Basseville, Igor V Nikiforov, et al., *Detection of abrupt changes: theory and application*, vol. 104, prentice Hall Englewood Cliffs, 1993.

[6] Elaine Biddiss, Dorcas Beaton, and Tom Chau, *Consumer design priorities for upper limb prosthetics*, Disability and Rehabilitation: Assistive Technology **2** (2007), no. 6, 346–357.

[7] Elaine Biddiss and Tom Chau, *Upper limb prosthesis use and abandonment: A survey of the last 25 years*, Prosthetics and Orthotics International **31** (2007), no. 3, 236–257.

[8] Alexander Boschmann, Dorothee Neuhaus, Sarah Vogt, Christian Kaltschmidt, Marco Platzner, and Strahinja Dosen, *Immersive augmented reality system for the training of pattern classification control with a myoelectric prosthesis*, Journal of NeuroEngineering and Rehabilitation **18** (2021).

[9] Alix Chadwell, Laurence Kenney, Sibylle Thies, Adam Galpin, and John Head, *The Reality of Myoelectric Prostheses: Understanding What Makes These Devices Difficult for Some Users to Control*, Frontiers in Neurorobotics **10** (2016), no. August, 1–21.

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

[11] Nancy Chinchor, *Muc-4 evaluation metrics*, Proceedings of the 4th Conference on Message Understanding (USA), MUC4 '92, Association for Computational Linguistics, 1992, p. 22–29.

[12] Jacob Cohen, *Statistical power analysis for the behavioral sciences*, Lawrence Erlbaum Associates, 1988.

[13] Diane J. Cook and Narayanan C. Krishnan, *Activity learning: Discovering, recognizing, and predicting human behavior from sensor data*, 2 2015, pp. 78–92.

[14] Ulysse Cote-Allard, Gabriel Gagnon-Turcotte, Angkoon Phinyomark, Kyrre Glette, Erik Scheme, Francois Laviolette, and Benoit Gosselin, *A transferable adaptive domain adversarial neural network for virtual reality augmented emg-based gesture recognition*, IEEE Transactions on Neural Systems and Rehabilitation Engineering **29** (2019), 546–555.

[15] Ulysse Côté-Allard, Gabriel Gagnon-Turcotte, François Laviolette, and Benoit Gosselin, *A low-cost, wireless, 3-d-printed custom armband for semg hand gesture recognition*, Sensors **19** (2019), no. 12.

[16] Carlo J. De Luca, *Surface Electromyography: Detection and Recording*, Tech. report, DelSys Incorporated, 2002.

[17] Delsys, *Discover beautiful emg. discover delsys.*, Jun 2023.

[18] Qichuan Ding, Xingang Zhao, Jianda Han, Chunguang Bu, and Chengdong Wu, *Adaptive hybrid classifier for myoelectric pattern recognition against the interferences of outlier motion, muscle fatigue, and electrode doffing*, IEEE Transactions on Neural Systems and Rehabilitation Engineering **27** (2019), 1071–1080.

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

[20] Arthur Gretton, Karsten Borgwardt, Malte J. Rasch, Bernhard Scholkopf, and Alexander J. Smola, *A kernel method for the two-sample problem*, 2008.

[21] Yikun Gu, Dapeng Yang, Qi Huang, Wei Yang, and Hong Liu, *Robust emg pattern recognition in the presence of confounding factors: features, classifiers and adaptive learning*, Expert Systems with Applications **96** (2018), 208–217.

[22] Maria Hakonen, Harri Piitulainen, and Arto Visala, *Current state of digital signal processing in myoelectric interfaces and related applications*, Biomedical Signal Processing and Control **18** (2015), 334–359.

[23] David J. Hand, *Assessing the performance of classification methods*, International Statistical Review **80** (2012), no. 3, 400–414.

[24] 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*, 2007 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2007, pp. 4842–4845.

[25] F.J. Harris, *On the use of windows for harmonic analysis with the discrete fourier transform*, Proceedings of the IEEE **66** (1978), no. 1, 51–83.

[26] Jacqueline S. Hebert and Justin Lewicke, *Case report of modified box and blocks test with motion capture to measure prosthetic function*, Journal of Rehabilitation Research and Development **49** (2012), 1163–1174.

[27] Qi Huang, Dapeng Yang, Li Jiang, Huajie Zhang, Hong Liu, and Kiyoshi Kotani, *A novel unsupervised adaptive learning method for long-term electromyography (emg) pattern recognition*, Sensors **17** (2017), no. 6.

[28] B. Hudgins, P. Parker, and R.N. Scott, *A new strategy for multifunction myoelectric control*, IEEE Transactions on Biomedical Engineering **40** (1993), no. 1, 82–94.

[29] Yoshinobu Kawahara, Takehisa Yairi, and Kazuo Machida, *Change-point detection in time-series data based on subspace identification*, Seventh IEEE International Conference on Data Mining (ICDM 2007), 2007, pp. 559–564.

[30] E. Keogh, S. Chu, D. Hart, and M. Pazzani, *An online algorithm for segmenting time series*, Proceedings 2001 IEEE International Conference on Data Mining, 2001, pp. 289–296.

[31] Rami N. Khushaba, Maen Takruri, Jaime Valls Miro, and Sarath Kodagoda, *Towards limb position invariant myoelectric pattern recognition using time-dependent spectral features*, Neural Networks **55** (2014), 42–58.

[32] Song Liu, Makoto Yamada, Nigel Collier, and Masashi Sugiyama, *Change-point detection in time-series data by relative density-ratio estimation*, Neural Networks **43** (2013), 72–83.

[33] Koji Makiyama, Ameya Daigavane, and Krzysztof Mierzejewski, *densratio 0.3.0*, 2022, Accessed July 17, 2023. https://pypi.org/project/densratio/.

[34] Marko Markovic, Meike A. Schweisfurth, Leonard F. Engels, Tashina Bentz, Daniela Wüstefeld, Dario Farina, and Strahinja Dosen, *The clinical relevance of advanced artificial feedback in the control of a multi-functional myoelectric prosthesis*, Journal of NeuroEngineering and Rehabilitation **15** (2018).

[35] Jena L. Nawfel, Kevin B. Englehart, and Erik J. Scheme, *A multi-variate approach to predicting myoelectric control usability*, IEEE Transactions on Neural Systems and Rehabilitation Engineering **29** (2021), 1312–1327.

[36] Mohammadreza Asghari Oskoei and Huosheng Hu, *Myoelectric based virtual joystick applied to electric powered wheelchair*, 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2008, pp. 2374–2379.

[37] Sebastien Racinais, *Hot ambient conditions shift the force / emg relationship*, 2013.

[38] Shriram T. P. Raghu, Dawn MacIsaac, and Erik Scheme, *Decision-change informed rejection improves robustness in pattern recognition-based myoelectric control pattern recognition-based myoelectric control*, IEEE Journal of Biomedical and Health Informatics **Epub ahead of print** (2023).

[39] E. Scheme, K. Biron, and K. Englehart, *Improving myoelectric pattern recognition positional robustness using advanced training protocols*, Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS (2011), 4828–4831.

[40] Erik Scheme and Kevin Englehart, *Training strategies for mitigating the effect of proportional control on classification in pattern recognition-based myoelectric control*, Journal of Prosthetics and Orthotics **25** (2013), 76–83.

[41] Erik J. Scheme and Kevin 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 **21** (2013), no. 4, 616–623.

[42] Erik J. Scheme and Kevin 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 **21** (2013), no. 4, 616–623.

[43] Erik J. Scheme, Kevin B. Englehart, and Bernard S. Hudgins, *Selective classification for improved robustness of myoelectric control under nonideal conditions*, IEEE Transactions on Biomedical Engineering **58** (2011), no. 6, 1698–1705.

[44] Jonathon W. Sensinger, Blair A. Lock, and Todd A. Kuiken, *Adaptive pattern recognition of myoelectric signals: Exploration of conceptual framework and practical algorithms*, IEEE Transactions on Neural Systems and Rehabilitation Engineering **17** (2009), no. 3, 270–278.

[45] Shriram Tallam Puranam Raghu, Dawn MacIsaac, and Erik Scheme, *Analyzing the impact of class transitions on the design of pattern recognition-based myoelectric control schemes*, Biomedical Signal Processing and Control **71** (2022), 103134.

[46] Charles Truong, Laurent Oudre, and Nicolas Vayatis, *Selective review of offline change point detection methods*, Signal Processing **167** (2020), 107299.

[47] Charles Truong, Laurent Oudre, and Nicolas Vayatis, *Selective review of offline change point detection methods*, Signal Processing **167** (2020), 107299.

[48] Marina M.-C. Vidovic, Han-Jeong Hwang, Sebastian Amsüss, Janne M. Hahne, Dario Farina, and Klaus-Robert Müller, *Improving the robustness of myoelectric pattern recognition for upper limb prostheses by covariate shift adaptation*, IEEE Transactions on Neural Systems and Rehabilitation Engineering **24** (2016), no. 9, 961–970.

[49] Marina M.-C. Vidovic, Liliana P. Paredes, Han-Jeong Hwang, Sebastian Amsu¨ss, Jaspar Pahl, Janne M. Hahne, Bernhard Graimann, Dario Farina, and Klaus-Robert Müller, *Covariate shift adaptation in emg pattern recognition for prosthetic device control*, 2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2014, pp. 4370–4373.

[50] Zachary Wright and Blair Lock, *Adaptive emg pattern recognition reduces frequency and improves quality of at-home prosthesis training for upper limb my-*

oelectric prosthesis wearers, Proceedings of the 2022 Myoelectric Controls and Upper Limb Prosthetics Symposium, 8 2022, pp. 17–20.

[51] Sophie M. Wurth and Levi J. Hargrove, *A real-time comparison between direct control, sequential pattern recognition control and simultaneous pattern recognition control using a fitts' law style assessment procedure*, Journal of Neuroengineering and Rehabilitation **11** (2014).

[52] Dapeng Yang, Yikun Gu, · Nitish, V Thakor, and Hong Liu, *Improving the functionality, robustness, and adaptability of myoelectric control for dexterous motion restoration*, Experimental Brain Research **237** (2019), 291–311.

[53] Dennis Yeung, Irene Mendez Guerra, Ian Barner-Rasmussen, Emilia Siponen, Dario Farina, and Ivan Vujaklija, *Co-adaptive control of bionic limbs via unsupervised adaptation of muscle synergies*, IEEE Transactions on Biomedical Engineering **69** (2022), 2581–2592.

[54] Bin Yu, Xu Zhang, Le Wu, Xiang Chen, Xun Chen, and Senior Member, *A novel postprocessing method for robust myoelectric pattern-recognition control through movement pattern transition detection*, IEEE Transaction On Human-Machine Systems **50** (2020).

[55] Yue Zhang, Zheng Wang, Zhuo Zhang, Yinfeng Fang, and Honghai Liu, *Comparison of online adaptive learning algorithms for myoelectric hand control*, 2016 9th International Conference on Human System Interactions (HSI), 2016, pp. 69–75.

Appendix A

Covariate Shift Adaptation of Linear & Quadratic Discriminant Analysis Classifiers

In these equations a new sample x_t, taken at time t, will affect both the mean $\mu_{t,c}$ and the covariance matrix $\Sigma_{t,c}$ associated with class c. $\Pi_{t,c}$ is a intermediary variable representing the auto-correlation matrix of class c at time t, used to calculate the covariance. Initial mean and covariance matrices, $\mu_{0,c}$ and $\Sigma_{0,c}$, are set using the available training data.

$$\mu_{t,c} = (1 - f_{t,c})\,\mu_{t-1,c} + f_{t,c}x_t \tag{A.1}$$

$$\Pi_{t,c} = (1 - f_{t,c})\,\Pi_{t-1,c} + f_{t,c}x_t x_t^T \tag{A.2}$$

$$\Sigma_{t,c} = \Pi_{t,c} - \mu_{t,c}\mu_{t,c}^T \tag{A.3}$$

The adaptation factor $f_{t,c}$ is assigned using the adaptation weight w_t and effective sample size $n_{t,c}$ as shown in Equation A.4. The adaptation weight, replacing the static value of 1 in the work of Anagnostopoulos et al., is used to assign im-

portance to the adaptation occurring at time t. As the importance is relative to the weight assigned to other frames, it is arbitrarily decided that all training frames were assigned weights of 1. Therefore, if the adaptation weight is larger than 1 the adaptation frame has a larger impact on the classifier than one frame of training data would have. When a weight of 0 is assigned, no adaptation occurs. The effective sample size, described in its recursive form in Equation A.5, is responsible for recording the summation of all previous adaptation weights. It also incorporates λ_t as an exponential forgetting factor, ranging from 0 to 1, where a value of 0 indicates no previous information is retained and a value of 1 indicates that all information is retained.

$$f_{t,c} = \frac{w_t}{n_{t,c}} \tag{A.4}$$

$$n_{t,c} = \lambda_{t-1} n_{t-1,c} + w_t \tag{A.5}$$

As Equations A.1 to A.5 are defined in a class-specific manner, as is required for a quadratic discriminant analysis (QDA) classifier, an additional stage is required to meet the assumption of similar covariances that is made by LDA classifiers. The effective sampling size of each class will control the contribution of each class-specific covariance matrix to the pooled covariance. However, the total effective sampling size, defined in Equation A.6, must first be calculated. The pooled covariance, used for the LDA classifier, is then calculated as a weighted combination of all class-specific covariance matrices using Equation A.7.

$$n_{t,pooled} = \sum_{c=1}^{C} n_{t,c} \tag{A.6}$$

$$\Sigma_{t,pooled} = \frac{1}{n_{t,pooled}} \sum_{c=1}^{C} n_{t,c} \Sigma_{t,c} \tag{A.7}$$

Appendix B

SEM-Based Transition-Informed Adaptation of the 3DC Data Set

The results of transition-informed UAL applied to the 3DC data set are included in Tables B.1 to B.6. As can be seen in Table B.1, the implementation of the SSL technique has very little effect in SAL. However, when applied in UAL, as in Table B.4, the benefits are clear for the LC adaptation scheme. For LC adaptation, SSL prevented the significant worsening of onset delay, offset delay, transition instability, and tertiary class error. It also significantly improved the steady state instability over the naive adaptation case.

Regardless of adaptation weight, SAL experienced significant worsening when HC, HC-SSL, or DSS was used for data selection. In opposition, LC, RuLSIF, and BUC generally resulted in significant improvements to the SS metrics while offset delay experienced slight worsening.

	Prompt	HC	HC-SSL	LC	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	25.2%	26.9%*^	26.9%*^	23.6%*^	23.6%*^	26.4%*	24.6%*	24.6%*
TER	27.4%	29.2%*^	29.2%*^	25.7%*^	25.7%*^	28.6%*	26.8%	26.8%*
INS_SS	7.3%	7.7%*^	7.7%*^	7.3%	7.3%	7.5%*	7.2%	7.2%
ON (ms)	1226.38	1289.93*^	1289.93*^	1230.67	1230.67	1262.21*	1245.52	1240.23
OFF (ms)	647.92	684.23*	684.23*	679.97*	679.97*	670.61*	672.28*	673.93*
DUR (ms)	578.47	605.69*	605.69*	550.71*^	550.71*^	591.60	573.24	566.30
INS_TR	11.7%	11.7%	11.7%	11.9%	11.9%	11.6%	11.6%	11.6%
TCE	58.8%	57.9%	57.9%	57.4%*	57.4%*	57.8%	57.8%*	57.7%*
PNM	6.6%	7.1%^	7.1%^	7.5%*^	7.5%*^	7.2%*^	7.1%*^	7.1%^

Table B.1: Classifier metrics following adaptation on 3DC data set with relative importance of 100% $W_{init} = 1.00$. LMC-assigned labels with SSL (th=25%).

	Prompt	HC-SSL	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	27.0%	28.5%*	25.6%*	28.1%*	26.4%*	26.3%*
TER	29.3%	30.8%*	27.8%*	30.6%*	28.6%*	28.6%*
INS_SS	7.6%	7.9%*	7.6%	7.8%*	7.5%*	7.5%
ON (ms)	1266.54	1302.98*	1259.70	1306.26*	1282.18	1282.86
OFF (ms)	665.25	680.86	678.59	683.39	680.87	682.23*
DUR (ms)	601.29	622.11*	581.10	622.87*	601.31	600.64
INS_TR	11.7%	11.8%	11.9%	11.6%	11.6%	11.6%
TCE	59.3%	58.1%	57.5%*^	58.1%*	57.9%*	58.2%*
PNM	6.9%	7.1%	7.3%	7.1%	7.1%	7.0%

Table B.2: Classifier metrics following adaptation on 3DC data set with relative importance of 50% $W_{init} = 0.50$. LMC-assigned labels with SSL (th=25%).

	Prompt	HC-SSL	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	24.5%	26.0%*^	22.5%*^	25.3%*	23.6%*	23.6%*
TER	26.5%	28.2%*^	24.5%*^	27.5%*	25.7%*	25.7%*
INS_SS	7.2%	7.6%*^	7.2%	7.3%	7.0%	7.0%*
ON (ms)	1180.87	1264.24*^	1220.18*	1234.65*^	1226.65*	1225.19*
OFF (ms)	638.54	685.59*^	686.61*^	671.69*	676.09*	674.52*
DUR (ms)	542.33	578.64*^	533.57	562.96*	550.55	550.67
INS_TR	11.6%	11.8%	12.2%*^	11.7%	11.6%	11.5%
TCE	58.1%	57.9%	57.4%	57.9%	57.3%	57.4%
PNM	6.5%	6.9%^	7.2%*^	7.1%^	7.0%^	7.0%^

Table B.3: Classifier metrics following adaptation on 3DC data set with relative importance of 200% $W_{init} = 2.00$. LMC-assigned labels with SSL (th=25%).

	HC	HC-SSL	LC	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL	Naive	No Adapt
AER	43.3%	42.2%	42.1%*	41.9%	42.3%	43.1%	43.4%	43.5%	45.8%*
TER	47.1%*	45.6%*	45.8%*	45.2%*	45.4%*	46.7%	47.2%	47.3%	49.4%*
INS_SS	7.6%	7.1%*^	7.9%	7.1%*^	7.1%*^	7.4%	7.5%	7.5%	9.6%*^^^
ON (ms)	1279.93*	1224.58	1313.56*	1248.49	1225.08	1272.43	1251.69	1258.62	1330.29*^
OFF (ms)	528.76	497.58	577.53*^	522.78	504.37	531.36	517.39	519.99	573.36*^
DUR (ms)	751.17*	727.00	736.03	725.71	720.71	741.07	734.31	738.62	756.93
INS_TR	8.0%	8.0%	8.9%*^	8.0%	7.8%	7.9%	7.9%	8.0%	9.9%*^^^
TCE	48.9%	48.9%	50.9%*	48.7%	48.6%	48.6%	48.3%	48.9%	51.3%*
PNM	5.4%	5.3%	5.7%	4.7%^	4.9%	5.4%	5.2%	5.4%	6.0%

Table B.4: Classifier metrics following adaptation 3DC data set with relative importance of 100% $W_{init} = 1.00$. Classifier-assigned labels with SSL (th=25%).

	Naive	No Adapt	HC-SSL	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	43.2%	45.8%*	42.0%*	42.3%	42.3%	42.9%	43.1%
TER	46.9%	49.4%*	45.4%*	45.4%*	45.5%*	46.4%	46.9%
INS_SS	7.9%	9.6%*^^^	7.3%*^	7.2%*^	7.4%*^	7.7%*	7.8%
ON (ms)	1283.78	1330.29	1245.48*	1232.87*	1249.14	1285.68	1278.96
OFF (ms)	543.31	573.36	513.37*	514.87*	520.69	542.43	532.43
DUR (ms)	740.48	756.93	732.11	718.00	728.44	743.25	746.53
INS_TR	8.3%	9.9%*^^	8.3%	8.3%	8.3%	8.3%	8.1%
TCE	49.8%	51.3%	49.8%	49.1%	49.2%	49.7%	49.7%
PNM	5.5%	6.0%	5.2%	5.0%	5.3%	5.3%	5.3%

Table B.5: Classifier metrics following adaptation on 3DC data set with relative importance of 50% $W_{init} = 0.50$. Classifier-assigned labels with SSL (th=25%).

	Naive	No Adapt	HC-SSL	LC-SSL	DSS-SSL	RULSIF-SSL	BUC-SSL
AER	44.3%	45.8%	42.9%*	42.4%	44.0%	43.9%	44.6%
TER	48.1%	49.4%	46.0%*	45.9%	47.0%	47.9%	48.4%
INS_SS	7.5%	9.6%*^^^	7.0%*^	7.0%*^	6.9%*^	7.3%	7.3%
ON (ms)	1241.19	1330.29*^	1207.40	1209.66	1174.25*^	1245.57	1222.16
OFF (ms)	517.87	573.36*^	481.14*	496.09	482.59	521.48	497.35
DUR (ms)	723.31	756.93^	726.26	713.58	691.65*	724.09	724.81
INS_TR	7.8%	9.9%*^^^	7.9%	7.8%	7.5%	7.8%	7.3%*
TCE	47.5%	51.3%*^	48.7%	46.8%	46.5%	47.3%	46.7%
PNM	5.2%	6.0%^	4.8%	5.2%	4.5%^	5.4%	4.9%

Table B.6: Classifier metrics following adaptation on 3DC data set with relative importance of 200% $W_{init} = 2.00$. Classifier-assigned labels with SSL (th=25%).

Vita

Julia Meneley
BScECE, University of New Brunswick 2020
Conference Presentations:

1. Meneley, J., MacIsaac, D., Scheme, E., "Segmentation of sEMG using time-series analysis," 24^{th} *Congress of the International Society of Electrophysiology and Kinesiology*, Quebec City, June 22-25, 2022