IMPROVING USER-IN-THE-LOOP MYOELECTRIC CONTROL USING CONTEXT INFORMED INCREMENTAL LEARNING

Evan Campbell[1], Ethan Eddy[1], Ulysse Côté-Allard[2], and Erik Scheme[1]
[1]*Institute of Biomedical Engineering, University of New Brunswick, Canada*
[2]*Department of Technology Systems, University of Oslo, Norway*

ABSTRACT

Screen or prosthesis guided training is typically used to train pattern recognition-based myoelectric control by providing controlled calibration samples with known labels. When these models are used with a user-in-the-loop, however, the observed patterns are much more variable, resulting in poor model extrapolation to these conditions and thus poor usability. Incremental and reinforcement learning approaches can continue learning from user-in-the-loop settings, but are limited in their reliability due to the lack of supervised labels and increased training times. In this work, we propose context informed incremental learning (CIIL), which adapts by drawing contextual information from the control task, to solve these issues. We test our claims across two conditions: a short training data scenario and a simulated electrode shift scenario. With only one second of initial training data per class, CIIL achieves similar throughput as SGT in a Fitts' law-style usability test after only two minutes of adaptation (the same amount of time taken for SGT). In the harder electrode shift scenario, CIIL significantly outperformed the pre-shifted SGT model after 5 minutes of adaptation, offering a promising direction for future clinical validation of user-in-the-loop training.

INTRODUCTION

Screen guided (or prosthesis guided) training (SGT) has become the standard practice for pattern recognition-based myoelectric control, achieving reasonable classification accuracy with acceptable effort and time invested by the user. By directing users through a sequence of motions to collect a representative dataset of electromyography (EMG) signals corresponding with known motion labels, models can be trained via supervised learning. An ongoing challenge, however, is how to make this data collection more representative of the data produced while the user is actively controlling their device, called the closed-loop/user-in-the-loop/online control setting. Solutions like collecting "ramp" contractions and known confounding factors like variable limb position have improved the performance in the online setting [1]; however, SGT still has meaningful shortcomings because of behavioural variability that arises from the online control setting (compensatory motions, proportional control, and error correction) which ultimately degrades model transfer to the online setting [2].

While strategies that continue to learn with the user-in-the-loop have tried to bridge this gap and produce more representative intent recognition models, these approaches have their own challenges. Among these strategies, unsupervised adaptation and reinforcement learning have been successful in addressing behavioural variability. Unsupervised adaptation, or incremental learning, continues to update the model in a weakly supervised learning process with the assumption that the model's prediction can be taken as a substitute for true labels. Approaches generally filter these predictions to only reincorporate highly confident [3], or representative [4] predictions to improve the robustness of the adaptation. Nevertheless, unsupervised incremental learning is generally not viable when the model's output is not trustworthy, such as after electrode-shift [5], and thus is not a reliable strategy when adaptation is arguably needed the most. Reinforcement learning leverages an environmental reward signal that truthfully describes the appropriateness of its actions for the situation, irrespective of classifier performance. Unfortunately, reinforcement learning is sample inefficient, which consequently increases the training burden to an unaffordable amount ($\sim$10-30 minutes [6]) as a standalone strategy employed online. Correspondingly, the challenge of all incremental learning and reinforcement learning approaches for incorporating user-in-the-loop behaviours is that they fail to simultaneously address the need for guaranteed performance improvements while ensuring sample efficiency.

Our proposed approach, context informed incremental learning (CIIL), aims to achieve the positives of both approaches. By using an environmental reward-like signal through context (similar to reinforcement learning), CIIL can reliably improve performance.Further, by using supervised learning to update model weights, CIIL can rapidly improve performance while minimizing the training burden on users. Across two simulated scenarios, including a short training data initialization and an introduced electrode shift, the proposed CIIL approach is compared to unsupervised high-confidence adaptation – the established baseline adaptation approach [3]. The results corroborate that CIIL quickly improves the model by incorporating user-in-the-loop patterns, irrespective of starting performance, which may lead to more robust long-term models.

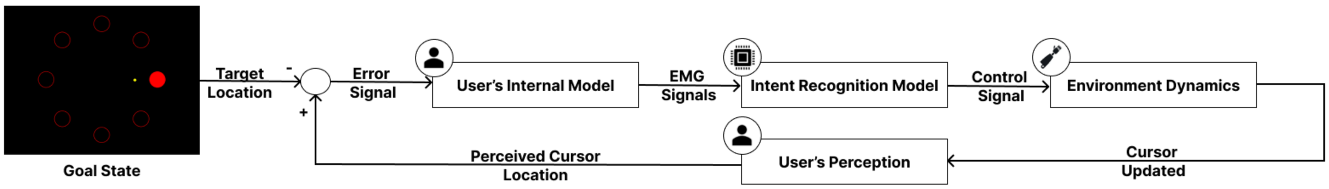

Fig. 1: Diagram of closed-loop myoelectric control. In this depiction, the user tries to direct a small yellow cursor to the red goal using myoelectric control. While the user is in-the-loop, they uses their internal model of the task and visual feedback to produce the necessary EMG to bring the cursor to the target. Consequently, produced EMG signals can be much more variable in this scenario compared to open-loop control that does not contain feedback.

METHODS

A. Hardware, Gestures, and Data Processing

The commercial Myo Armband was used to collect EMG data at 200Hz from 8 channels. Wrist flexion/extension and hand open/close gestures, plus a rest class, were mapped to left/right, up/down, and no movement of the cursor. To recognize the five gestures, all EMG data were first split into windows of 200ms with 100ms increments. Next, Hudgins' time domain features (Mean Absolute Value, Slope Sign Changes, Zero Crossings, and Waveform Length) were extracted from each window. All data processing was done using LibEMG, an open source Python library for myoelectric control [7]. Finally, all features were passed through an adaptive Linear Discriminant Analysis (aLDA) classifier [8].

B. Incremental Learning Strategies

The purpose of incremental learning is to gather behaviours that arise during the closed-loop setting. In this setting, subjects vary their contractions to move the cursor to a target position. The patterns produced are dependent on several factors, including the user's internal model and task-specific behaviours (Figure 1), and consequently are more variable than signals observed during SGT. Within this study, we evaluated an established incremental learning approach, Unsupervised High Confidence (UHC) compared to our proposed approach Context Informed Incremental Learning (CIIL).

Unsupervised High-Confidence (UHC): As proposed by Sensinger et al. [3], UHC adaptation involves updating a classifier based on predictions whose probability are above a predefined threshold.

Context Informed Incremental Learning (CIIL): Our proposed method makes use of a novel source of information returned from the environment: *context*. Context contains *suitability* information related to whether the actions taken at this time step were aligned with goal-seeking behaviour or not (binary value), and can be instrumented within many scenarios, such as the target acquisition environment shown in Figure 2. The control actions that are viable *options* for goal-seeking behaviour are also returned at each time step, which realistically can be inferred even from complex environments like prosthesis control using detection of proximal objects to curate likely classes of motion. Context is then used to inform the selection of labels that were suitable for the environment given the task, reinforcing positive actions. For unsuitable actions, wherein the classifier prediction disagrees with all viable options, the most viable option is used to relabel the sample for adaptation. This promotes the prediction of more viable alternatives the next time similar EMG patterns appear, adapting the model behaviour in this uncertain region.

C. Experimental Overview

Both experiments followed the same experimental structure, outlined as follows, with minor modifications outlined in the individual experimental sections. Both experiments were approved by the University of New Brunswick Research Ethics Board and are on file as REB 2022-122. Further details can be found in [9].

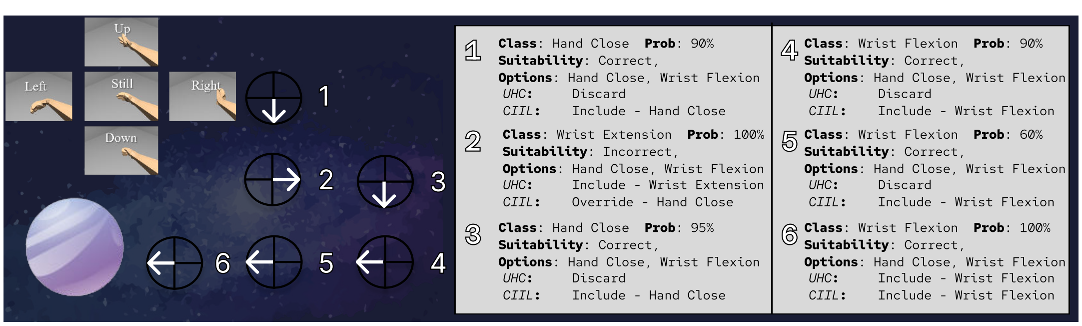

Fig. 2: Diagram of how adaptation strategies filter or correct pseudo-labels from within the gamified environment. The gesture mapping to cursor movement map is given in the top left. Crosshairs and arrows indicate cursor position at a given timestep and intent recognized. A table is given that provides information available to the incremental learning approaches, and how the sample is ultimately used in future model updates.

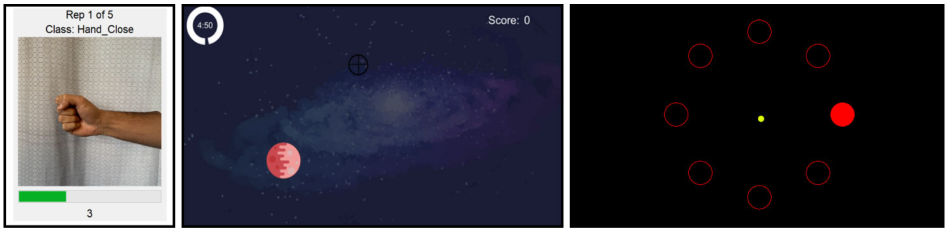

Fig. 3: (Left) Shows the SGT environment. (Middle) Shows Myo Shoot — the gameplay environment where adaptation took place. Users controlled the black cross-hair, and the planet represented the target. (Right) Shows the adapted ISO 9241-9 Fitt's law test. The circles represent targets, and the smaller yellow icon represents the cursor.

1) *Armband Placement:* The Myo Armband was placed 2/3 up the participant's right forearm, proximal to the elbow.
2) *Initial Model Training:* The initial model was trained. For experiment scenario one, one second of training data for each class was acquired. For experiment scenario two, 5 reps of 3 seconds were recorded for each gesture using SGT (see Figure 3, Left).
3) *Initial Model Evaluation:* The initial model was evaluated as per *Step 5 - Online Assessment*.
4) *Adaptive Gameplay:* Participants played the Myo-Shoot game (Figure 3, Middle), whereby real-time adaptation occurred in ten-second batches. The goal during this game was to acquire as many targets (planets) as possible. To successfully acquire a target, users had to hover the cross-hair within the planet's bounds and hover on top of it for three seconds. After successfully acquiring a planet, another was randomly re-generated at another position on the screen. Models were updated during this phase based on the given adaptation strategy. The order of adaptation strategies was randomized across participants.
5) *Online Assessment:* Using the adapted model from step 4, participants completed an ISO 0241-9 inspired Fitt's law test (Figure 3, Right), whereby they had to hover a cursor over a set of targets and hover within their bounds for three seconds. Once acquired, the user was prompted to move the cursor to a different target. This continued until the participant acquired all eight targets in the ring. During this evaluation, no adaptation occurred.
6) *Offline Assessment:* To enable offline testing of the adapted model with representative data, two additional three-second repetitions of each gesture were acquired.

Experimental Scenario 1: Limited Training Data Eleven participants took part in this initial study on the use of a limited amount of training data (1 second per class). The adaptive gameplay phase lasted for 2 minutes. A constant velocity was employed, meaning the cursor moved at a fixed speed, regardless of contraction intensity. Finally, due to the high confidence profile of LDA classifiers, a UHC threshold of 100% was leveraged. Even at this seemingly high threshold, approximately 50% of decisions were used to adapt the model [10].

Experimental Scenario 2: Electrode Shift In a subsequent study, 21 individuals evaluated the adaptation strategies after a severe 45° electrode shift. After training the initial model via SGT, participants were instructed to rotate the Myo Armband clockwise by one electrode. The adaptive gameplay phase lasted 5 minutes. Due to the decreased confidence profile of the aLDA classifier, a lower confidence threshold of 99% was selected for UHC adaptation. Additionally, a proportional control scheme was adopted, meaning harder contractions resulted in faster cursor speeds meaning subjects could generally achieve higher throughput [1].

RESULTS

The results across the online assessment and offline assessment of both experiments were compiled into Table I. In the limited training data scenario, all incremental learning approaches led to improved performance across all metrics when compared to the model seeded with only one second of calibration data. CIIL performed comparably to a full SGT protocol for online metrics, despite only beginning with one second of calibration data. In the same setting, UHC performed marginally worse than CIIL, with a 12.5% relative difference in throughput. In the electrode shift scenario, the offline accuracy of the model dropped from 87.8% to 48.3%, demonstrating the severity of the degradation introduced by the electrode shift. After a 5 minute adaptation phase, CIIL was able to significantly outperform the UHC approach, which was unable to recover to a usable state. CIIL also significantly outperformed the pre-shift SGT, tested before the shift was introduced, and the unusable post-shift SGT.

DISCUSSION

In the limited training data scenario, a single second of calibration data per class was sufficient for both UHC and CIIL adaptation to positively influence the throughput of the model. Past work has evaluated UHC under favourable conditions, where model performance is extremely high when adaptation begins; however, this study shows that UHC can improve performance for a distinguishable gesture set when classification accuracy begins as low as 65.7%. Regardless, CIIL was

TABLE I: Summary of performance for the two experimental scenarios. Bold values indicate the best value of the metric across the approaches within an experiment. Dashes denote the fact that the SGT Post-Shift model was unusable and thus no acquisitions were completed. An asterisk (*) indicates a significant difference with CIIL determined through a Friedman test with Finner posthoc correction.

Model	Accuracy (%)	Active Error (%)	Instability (%)	Overshoots (#)	Efficiency (%)	Throughput (bit/s)
Experiment 1: Limited Training Data						
SGT	**91.3 ± 8.6**	**6.2 ± 5.7**	**5.0 ± 3.8**	4.5 ± 3.7	68.5 ± 6.8	**0.41 ± 0.06**
1 Second	65.7 ± 14.2 *	32.6 ± 12.3 *	19.9 ± 6.0 *	13.6 ± 19.4	42.6 ± 18.8 *	0.24 ± 0.10 *
UHC	79.1 ± 12.8	18.2 ± 12.0 *	12.1 ± 7.3 *	6.1 ± 7.5	60.6 ± 14.1	0.35 ± 0.08
CIIL	85.3 ± 11.0	10.5 ± 8.0	8.1 ± 6.2	**3.6 ± 3.9**	**70.1 ± 8.2**	0.40 ± 0.06
Experiment 2: Electrode Shift						
SGT Pre-Shift	**87.8 ± 7.6**	**10.0 ± 6.2**	7.8 ± 4.6	7.9 ± 7.9	61.6 ± 13.6	0.50 ± 0.18 *
SGT Post-Shift	48.3 ± 15.5 *	61.5 ± 18.4 *	11.4 ± 5.3 *	–	–	–
UHC	45.2 ± 17.8 *	65.1 ± 22.5 *	**6.5 ± 4.5**	9.6 ± 12.4 *	11.6 ± 21.7 *	0.09 ± 0.18 *
CIIL	79.5 ± 17.2	21.0 ± 20.4	7.3 ± 4.3	**5.0 ± 7.0**	**66.0 ± 15.8**	**0.61 ± 0.21**

shown to improve throughput more than UHC in the limited training data setting, where both models were provided the same one second calibration and two minute adaptation phase. This outcome suggests that CIIL can quickly reduce epistemic errors through adaptation, and with a higher ceiling than UHC, even when beginning with a correctly trained model.

In the second experiment, both adaptation approaches began with an abrupt and intentional concept shift that degraded classifier accuracy from 87.8% to 48.4%. With this lower starting accuracy, UHC was unable to recover to a usable state post-adaptation. Further, despite filtering predictions using high confidence, UHC marginally decreased accuracy and active error compared to the initialization; indicating the lack of reliability of current adaptation approaches to address large concept shifts. Conversely, CIIL was shown to reliably improve performance irrespective of classifier performance and even outperformed the pre-shift SGT model's throughput. This outcome suggests that CIIL can reliably improve the model regardless of initial classifier performance using context from the environment. Correspondingly, CIIL warrants future work exploring the numerous other confounding factors that introduce concept shifts, such as limb position [11].

This study was an initial validation of CIIL, with future work needed to make the approach viable for clinical populations. If used as a training environment, the proposed CIIL approaches are viable in their current state. However, if leveraged during real-time use, future work on how to gather context within a prosthesis use-case, and how frequently it can be incorporated, is required. The benefits of using CIIL would need to merit the cost of creating advanced prostheses with the ability to perceive the environment to recognize nearby objects and grasps suitable for the situation. Once fully instrumented, it could be possible to rely on the environmental context alone for semi-autonomous gesture elicitation, however, this would decrease the user's agency over their device [12]. Consequently, the proposed CIIL approach may offer a favorable tradeoff for improved control. Ultimately, CIIL enables the incorporation of user-in-the-loop training data with context-based pseudo-labels. The ability to integrate user behaviours, confounding factors, and changes over time based on situational context is promising, and warrants further research.

References

[1] E. Scheme, B. Lock, L. Hargrove, W. Hill, U. Kuruganti, and K. Englehart, "Motion normalized proportional control for improved pattern recognition-based myoelectric control," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 22, no. 1, pp. 149–157, 2013.

[2] L. Hargrove, Y. Losier, B. Lock, K. Englehart, and B. Hudgins, "A real-time pattern recognition based myoelectric control usability study implemented in a virtual environment," in *2007 29th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 2007, pp. 4842–4845.

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

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

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

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

[7] E. Eddy, E. Campbell, A. Phinyomark, S. Bateman, and E. Scheme, "LibEMG: An open source library to facilitate the exploration of myoelectric control," *IEEE Access*, vol. 11, pp. 87 380–87 397, 2023.

[8] H. Zhang, Y. Zhao, F. Yao, L. Xu, P. Shang, and G. Li, "An adaptation strategy using LDA classifier for EMG pattern recognition," in *2013 Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2013, pp. 4267–4270.

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

[10] E. Scheme and K. Englehart, "A comparison of classification based confidence metrics for use in the design of myoelectric control systems," in *2015 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, 2015, pp. 7278–7283.

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

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