

ADAPTIVE SWITCHING IN PRACTICE: IMPROVING MYOELECTRIC PROSTHESIS PERFORMANCE THROUGH REINFORCEMENT LEARNING

Ann L. Edwards^{1,2}, Michael R. Dawson³, Jacqueline S. Hebert^{2,3},
Richard S. Sutton¹, K. Ming Chan², Patrick M. Pilarski^{1,2}

¹*Department of Computing Science, University of Alberta, Edmonton, AB, Canada;*

²*Division of Physical Medicine and Rehabilitation, University of Alberta, Edmonton, AB, Canada;*

³*Glenrose Rehabilitation Hospital, Alberta Health Services, Edmonton, AB, Canada.*

INTRODUCTION

Myoelectrically controlled prostheses are a class of assistive device that use electrical signals generated by muscle activation. These electromyographic (EMG) signals are used to control one or more electromechanical actuators that move prosthetic joints. Myoelectric control signals are typically measured with electrodes on the surface of the skin, with one pair of electrodes over each muscle site. In this manner, each muscle site directly controls one motion of the prosthesis, and various methods of switching can be used as needed to control additional motions of the prosthesis [1] [2] [3].

Some state-of-the-art myoelectric hands currently used by amputees have over a dozen possible grip patterns that can be manually selected by the user. Despite increasing possible control options, a robotic arm with so many available motions presents a problem, since there exist more degrees of freedom than there are available control signals from the human user [1] [4] [5]. One solution to circumvent this problem is for the user to switch between all available joints or grip patterns in a pre-designed, optimized order. As another option, the amputee and their prosthetist may selectively reduce the number of available control options (i.e., the amputee will have access to and switch between only a small subset of the device's available functions during regular use). Both of these options require trade-offs between switching effort and device functionality.

While switching between functions continues to be used in clinical settings to extend prosthesis functionality, it can be laborious. Switched or gated control is considered to be slow and non-intuitive, requiring both time and sustained cognitive effort on the part of the user [1] [4]. Non-intuitive control in fact represents one of the main reasons amputees stop using their myoelectric prostheses [1] [2] [3]. These limitations have been a driving force for more advanced control paradigms such as pattern recognition [1] [3]. However, as functionality increases and control becomes more challenging, one acknowledged solution is for prostheses to begin to assume more autonomy in interpreting and executing a user's intended movements.

Previous work by our group has therefore examined ways to streamline and optimize prosthetic control interfaces such as the switching system indicated above, potentially increasing the number of available and accessible modes or functions through the use of machine intelligence [5] [6] [7] [8]. In particular, our prior work showed how predictions about sensorimotor signals, such as signals pertaining to arm movements, could be learned and maintained using a reinforcement learning technique known as *General Value Functions* (GVFs). GVFs are temporally extended predictions about a signal of interest that have been applied to building up real-time anticipatory knowledge in relation to human-machine interactions [5] [6] [7]. We have shown in experiments using reinforcement learning offline (prior to use in prosthesis control) that GVFs may offer a way to help streamline control interfaces with robotic arms. In particular, we demonstrated the use of GVFs and reinforcement learning to predict which joint of a robotic arm an amputee user intends to actuate next, and proposed the idea of an adaptive or situation-specific switching list [6]. A natural extension of this work would be to apply predictions to actual human interaction with artificial limbs with the intent of improving control. Applying predictions to human machine interaction is consistent with the knowledge that, similar to GVFs, the human brain makes motor predictions of its own, using both knowledge of context and immediate sensory input [9].

In the current paper we extend our prior studies to provide preliminary evidence that our method of adaptive switching does in fact provide benefit during actual use of a robotic arm by an amputee. This work is the first simple demonstration of the use of prediction learning in real time to improve the control of a prosthetic device during its use by an amputee subject. Predictions are learned and used in real time by the control system to reduce the burden of control on the user, making it easier and faster to switch to the user's intended next joint or function.

METHODS

In order to implement and assess adaptive switching, three subjects—two transhumeral amputee subjects and one able-bodied subject—were recruited to perform a simple,

semi-repetitive task using an experimental robotic arm. Because of the similarity between the data sets, only one representative data set is presented. The subject was a body-powered prosthetic user and had no experience using myoelectric control or using our specific robotic arm. We attached surface electrodes to the skin over his wrist extensor muscle on the intact arm, which provided control signals for switching between robot joints. Separate sets of electrodes were also attached to the biceps and triceps muscle of his residual limb. Those electrodes became the source of control signals for flexing and extending selected joints of the robot arm. An 8-channel Bagnoli EMG system (Delsys, Inc.) was used in the acquisition of EMG control signals from the experimental subject, at a frequency of 1 kHz. The subject gave informed consent to participate and the trial was approved by the human research ethics board at the University of Alberta.

We used a custom-built robot arm known as the Myoelectric Training Tool (MTT) in our experiments [10]. The MTT includes an AX-18 smart robotic arm (Crustcrawler, Inc.) that has five degrees of freedom and can be controlled via EMG signals by both amputees and able-bodied subjects. In addition, it can be used as a training tool for amputees preparing to use a myoelectrically controlled prosthetic arm, as it was designed to be functionally similar to commercial prostheses. Figure 1 shows the amputee subject using the MTT to perform a simple task.



Figure 1: Amputee participant performing simple tasks with the robot arm using myoelectric control signals.

The subject was given time to become familiar with the MTT. After familiarization, the subject was presented with a

specific task that involved a subset of the available joints (specifically hand open/close, wrist flexion/extension, and shoulder rotation). The task was chosen to be functionally comparable to other tasks of daily living—for instance, picking up a dish and placing it on a shelf. The instruction given to the subject in both the non-adaptive and adaptive trials was to manipulate the MTT to grasp an imaginary object on one side of the shoulder space, rotate the shoulder to the opposite side, wave with the wrist joint, and rotate the shoulder back to the other side. Each trial involved repeating this task for a total of 3 minutes.

Two types of trials were performed in order to test the predictive capabilities of our design compared with conventional switching methods. In the non-adaptive trial, the subject switched their myoelectric control between four joints in a fixed switching order: hand, wrist, elbow, and shoulder. In contrast, in the adaptive trial, the joints were continuously reordered in the switching list based on their likelihood of being used next. This was done in an ongoing fashion throughout the course of the task through the use of GVs. Three 3-minute trials were done each for non-adaptive and adaptive switching.

As described in Pilarski et al. (2012), GVs represented predictions about the subject’s situation-specific use of each joint in the switching list [6]. These predictions were learned during robot arm use by the subject and continuously ranked based on their relative magnitudes. In the current work, with adaptive switching turned on, the system learned to predict the intended joint for the given task in advance of the switch signal from the user. When a switch signal was received by the system, the highest-ranked joint in the adaptive switching list became the active joint, with the remaining joints filling in the new switching list in decreasing order of prediction strength. All GV learning was implemented as per Pilarski et al. [6].

In order to build up real-time predictions about the intended active joint, we combined ongoing sensorimotor data from the robotic arm with EMG data from the human user. Each of the AX-18 motors that make up the joints of the MTT relayed a number of useful sensorimotor outputs, including angular position, angular velocity, load (current), temperature, and voltage. We used a select number of these motor observations as features, or information about the current state, in the learning system. The included observations were angular position and angular velocity of each joint. Features based on the current state of the arm enable the system to build up expectations about future switching decisions made by the user. The machine learning system was re-initialized at the beginning of each trial—GVs started with no stored knowledge (predictions) about the user or the task in question.

RESULTS AND DISCUSSION

Figure 2 compares the number of switches required per event for non-adaptive switching (top) with the number of switches required during adaptive switching (bottom) for the subject. Each switching event was considered to begin when the user triggered a joint switch, and end when the user initiated movement of any of the MTT joints. Therefore, all switches made while shifting control to a new joint are counted as a single switching event. As shown in Figure 2, there was a significant difference between non-adaptive switching and adaptive switching. With adaptive switching enabled, after an initial period of learning by the system (i.e. the first several switching events), typically only one switch was required by the user to select the most appropriate joint.

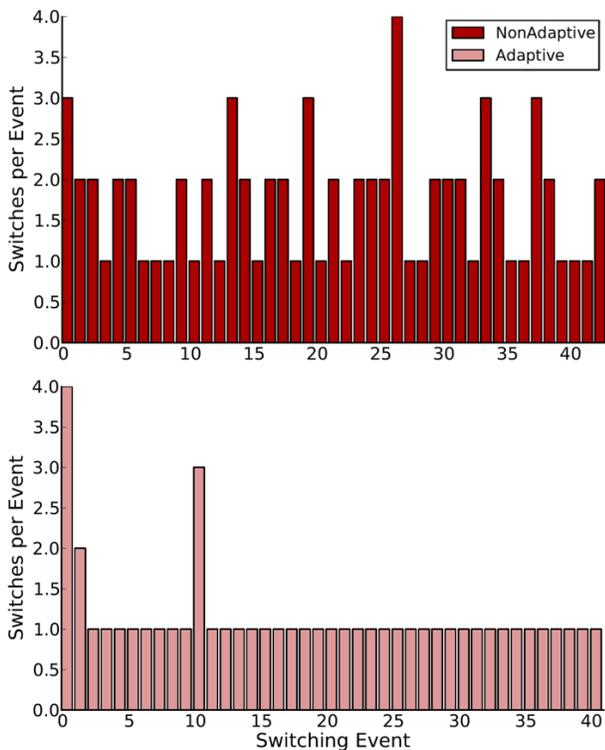


Figure 2: Number of voluntary switches initiated by the amputee subject per switching event over the course of a single 3 min trial. Shown for both non-adaptive (top) and adaptive control (bottom) approaches.

The decrease in the number of switches is also reflected in Figures 3 and 4. Figure 3 shows the average amount of time (measured in seconds) dedicated to switching, calculated over the three non-adaptive trials and the three adaptive trials. Adaptive switching showed a large decrease in time spent switching compared with non-adaptive. Thus, for each 3-minute trial, the subject saved an average of about 20 seconds when adaptive switching was enabled. Figure 4 is the total number of switches averaged over three

trials. The decrease in the amount of time spent switching is also illustrated in the decrease in the total number of switches per trial. Furthermore, the median time per switching event was consistently more than 1 second for all non-adaptive trials, and consistently under 1 second for adaptive trials. Not only was the median time per event lower, but, in some trials the total number of switching events completed in a task was also greater when adaptive switching was enabled.

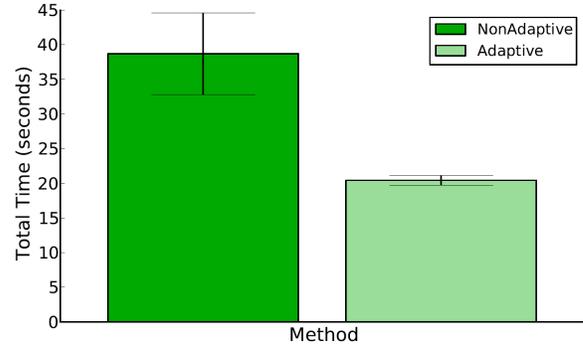


Figure 3: Average time the amputee subject spent switching per trial when using non-adaptive and adaptive switching (left and right, respectively, average over 3 trials).

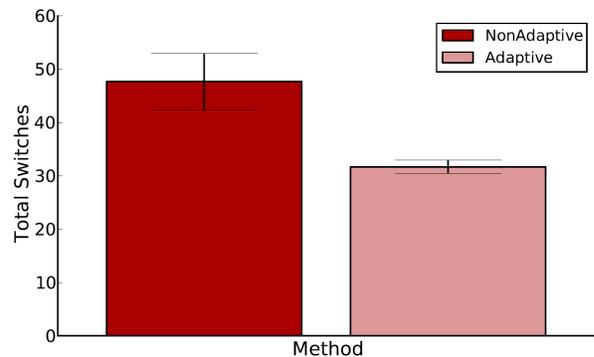


Figure 4: Average number of switches made by the amputee subject per trial when using non-adaptive and adaptive switching (left and right, respectively, average over 3 trials).

These results suggest there are efficiencies with adaptive switching, and agree with our expectations regarding the simple task presented to the subject: there were clear regions of the task space that corresponded to the use of specific joints. For this task, it would have been possible to hand-code several different switching lists in response to the different positions of the shoulder actuator. The simplicity allowed us to easily verify the correctness of the adaptive switching options proposed by the learning system. However, a key observation from the present work

is that situation-specific switching orders do not need to be hand-coded; our system learns situational delineations as the robotic arm is being used, and without prior information about the user or their task. Furthermore, we have observed that as the task changes or becomes more complex (and thus increasingly harder to engineer) the learning system can scale up naturally and easily without the need for manual tuning.

CONCLUSION

The primary contribution of this paper is a concrete demonstration of adaptive switching in an applied setting. This study is the first time that real-time prediction learning has been used to improve the control interface of a robotic device during un-interrupted use by an amputee subject. Our experiments with an amputee subject showed that for simple tasks, enabling adaptive switching on a robotic arm significantly decreased the time spent switching. This is consistent with and extends previous studies using pre-recorded (non-real-time) data that indicated the potential merit of adaptive switching.

We believe that adaptive switching would help to decrease the cognitive load required by amputees during more complex tasks and real-world functional situations using a wearable prosthesis. In particular, in our future work we will study the use of adaptive switching in tasks with multiple solution pathways—i.e., situations where many possible (and user specific) movement sequences could be used to achieve the task's objective.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge support from the Alberta Innovates Centre for Machine Learning (AICML), Alberta Innovates – Technology Futures (AITF), the National Science and Engineering Research Council (NSERC), and the Glenrose Rehabilitation Hospital Foundation. The study was approved by the human research ethics board at the University of Alberta.

REFERENCES

- [1] E. Scheme and K. B. Englehart, "Electromyogram pattern recognition for control of powered upper-limb prostheses: State of the art and challenges for clinical use," *The Journal of Rehabilitation Research and Development*, vol. 48, no. 6, pp. 643-660, 2011.
- [2] B. Peerdeman, H. D Boere, W. R H in 't Veld, H. Hermens, S. Stramigioli, H. Rietman, P. Veltink and S. Misra, "Myoelectric forearm prostheses: State of the art from a user-centered perspective," *The Journal of Rehabilitation Research and Development*, vol. 48, no. 6, pp. 719-738, 2011.
- [3] S. Micera, J. Carpaneto and S. Raspopovic, "Control of hand prostheses using peripheral information," *IEEE Rev. Biomed. Eng.*, vol. 3, pp. 48-68, 2010.
- [4] T. W. Williams, "Guest editorial: Progress on stabilizing and controlling powered upper-limb prostheses," *J. Rehabil. Res. Dev.*, vol. 48, no. 6, pp. ix-xix, 2011.
- [5] P. M. Pilarski, M. R. Dawson, T. Degris, J. P. Carey, K. M. Chan, J. S. Hebert and R. S. Sutton, "Adaptive artificial limbs: a real-time approach to prediction and anticipation," *IEEE Robotics & Automation Magazine*, vol. 20, pp. 53-64, 2013.
- [6] P. M. Pilarski, M. R. Dawson, T. Degris, J. P. Carey and R. S. Sutton, "Dynamic switching and real-time machine learning for improved human control of assistive biomedical robots," in *Proc. 4th IEEE RAS & EMBS Int. Conf. Biomedical Robotics and Biomechanics (BioRob)*, Roma, Italy, 2012.
- [7] 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," in *Proc. 1st Multidisciplinary Conf. on Reinforcement Learning and Decision Making*, Princeton, NJ, 2013.
- [8] P. M. Pilarski, T. B. Dick and R. S. Sutton, "Real-time Prediction Learning for the Simultaneous Actuation of Multiple Prosthetic Joints," in *Proc. of the 2013 IEEE International Conference on Rehabilitation Robotics (ICORR)*, Seattle, USA, 2013.
- [9] R. S. Sutton and A. G. Barto, Time-derivative models of pavlovian reinforcement. In *Learning and Computational Neuroscience: Foundations of Adaptive Networks*, M. Gabriel and J. Moore, Eds., MIT Press, 1990, pp. 497-537.
- [10] M. R. Dawson, F. Fahimi and J. P. Carey, "The development of a myoelectric training tool for above-elbow amputees," *Open Biomed. Eng. J.*, vol. 6, pp. 5-15, 2012.
- [11] R. S. Sutton, J. Modayil, M. Delp, T. Degris, P. M. Pilarski, A. White and D. Precup, "Horde: a scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction," in *Proc. 10th Int. Conf. Autonomous Agents and Multiagent Systems (AAMAS)*, Taipei, Taiwan, 2011.