

# Application of real-time machine learning to myoelectric prosthesis control: A case series in adaptive switching

Ann L. Edwards<sup>1,2</sup>, Michael R. Dawson, MSc<sup>3</sup>, Jacqueline S. Hebert, MD, FRCPC<sup>2,3</sup>, Craig Sherstan<sup>2,4</sup>, Richard S. Sutton, PhD<sup>4</sup>, K. Ming Chan, MD, FRCPC<sup>2,3</sup>, and Patrick M. Pilarski, PhD<sup>2,\*</sup>

<sup>1</sup>Faculty of Rehabilitation Medicine, University of Alberta, Edmonton, AB, Canada

<sup>2</sup>Division of Physical Medicine & Rehabilitation, University of Alberta, Edmonton, AB, Canada

<sup>3</sup>Glenrose Rehabilitation Hospital, Alberta Health Services, Edmonton, AB, Canada

<sup>4</sup>Department of Computing Science, University of Alberta, Edmonton, AB, Canada

\* *Corresponding Author:* Patrick M. Pilarski, Ph.D., Division of Physical Medicine and Rehabilitation, Department of Medicine, 5-005 Katz Group Centre for Pharmacy and Health Research, University of Alberta, Edmonton, Alberta, Canada, T6G 2E1; patrick.pilarski@ualberta.ca; +1 (780) 248-5636

## Abstract

**Background:** Myoelectric prostheses currently used by amputees can be difficult to control. Machine learning, and in particular learned predictions about user intent, could help to reduce the time and cognitive load required by amputees while operating their prosthetic device.

**Objectives:** The goal of this study was to compare two switching-based methods of controlling a myoelectric arm: non-adaptive (or conventional) control and adaptive control (involving real-time prediction learning).

**Study Design:** Case series study.

**Methods:** We compared non-adaptive and adaptive control in two different experiments. In the first, one amputee and one non-amputee subject controlled a robotic arm to perform a simple task; in the second, three able-bodied subjects controlled a robotic arm to perform a more complex task. For both tasks, we calculated the mean time and total number of switches between robotic arm functions over three trials.

**Results:** Adaptive control significantly decreased the number of switches and total switching time for both tasks compared with the conventional control method.

**Conclusion:** Real-time prediction learning was successfully used to improve the control interface of a myoelectric robotic arm during uninterrupted use by an amputee subject and able-bodied subjects.

**Clinical Relevance:** Adaptive control using real-time prediction learning has the potential to help decrease both the time and the cognitive load required by amputees in real-world functional situations when using myoelectric prostheses.

## Background

Myoelectrically controlled prostheses use electromyographic (EMG) signals generated by muscle activation and detected by surface electrodes.<sup>1</sup> In traditional direct control, one pair of electrodes over each muscle site is used to drive electromechanical actuators that move each prosthetic joint. In this manner, each pair of antagonistic muscle sites directly controls one motion of the prosthesis, and various methods of switching can be used as needed to control additional motions of the prosthesis.<sup>1-4</sup>

State-of-the-art myoelectric hands have over a dozen possible grip patterns that can be manually selected by the user. A prosthetic arm with many available joint movement and grasp options can present a problem, since there exist more degrees of freedom than the available number of control signals from the human user.<sup>1,5-7</sup> One technique to address this problem is for the user to switch between all available joints or grip patterns in a predesigned, optimized order. As another option, the prosthetist may selectively reduce the number of available control options, such that the user only has access to a small subset of the devices available functions. Both of these options require trade-offs between switching effort and device functionality.

While switching between functions is used in clinical settings to extend prosthesis functionality, it can be laborious<sup>1</sup>. Switched or gated control is slow and non-intuitive, requiring additional time and sustained cognitive effort by the user.<sup>1,5</sup> Non-intuitive control represents one of the main reasons amputees stop using their myoelectric prostheses.<sup>1-3</sup> These limitations have been a driving force for the development of more advanced control paradigms such as pattern recognition, which increase the number of degrees of freedom that can be intuitively controlled, but are typically still limited to sequential control.<sup>1,3,8</sup> However, as device functionality increases and control becomes more challenging,<sup>7</sup> an alternate solution may be for prostheses to begin to as-

**This paper is a preprint of the following article:** A. L. Edwards, M. R. Dawson, J. S. Hebert, C. Sherstan, R. S. Sutton, K. Ming Chan, and P. M. Pilarski, "Application of real-time machine learning to myoelectric prosthesis control: A case series in adaptive switching," *Prosthetics and Orthotics International*, first published on September 30, 2015 as doi:10.1177/0309364615605373

sume more autonomy in interpreting and executing a users intended movements.<sup>8</sup> This could reduce the burden on the user to consciously and independently control every individual joint motion of the prosthesis.

Previous work has examined ways to streamline and optimize prosthetic control through the use of more robust pattern classification and regression techniques,<sup>1,3,8</sup> supervised and unsupervised adaptation,<sup>9,10</sup> and real-time machine intelligence.<sup>6,11-13</sup> In particular, prior work on increasing the number of accessible prosthetic functions demonstrated how predictions about sensorimotor signals pertaining to prosthetic joint movements could be learned and maintained using a technique from reinforcement learning known as *general value functions* (GVFs).<sup>14</sup> GVFs are temporally extended predictions about signals of interest, and have been applied to build up real-time anticipatory knowledge in relation to human-machine interactions.<sup>6,11,12</sup> It was shown in experiments that GVFs may provide a way to streamline control interfaces with robotic arms.<sup>6,11-13</sup> In particular, it was demonstrated that the use of GVFs to predict which joint of a robotic arm a user will actuate next could lead to an adaptive or situation-specific switching list, termed *adaptive switching*.<sup>11</sup> A natural extension of that work would be to apply predictions to actual human interaction with prosthetic limbs with the intent of streamlining control. Applying GVF predictions to human-machine interaction is consistent with the idea that the human brain makes forward motor predictions of its own, using knowledge of context and immediate sensory input.<sup>15-18</sup>

In the current paper, we extend prior studies to present evidence that adaptive switching does in fact provide benefit during the operation of a robotic arm by a prosthetic user. This work includes a simple demonstration of the use of prediction learning in real time to improve the control of a prosthetic device during use by an amputee subject and an able-bodied subject. We also include preliminary results from three able-bodied subjects performing a second, more complex task. In both cases predictions are learned and used in real time by the control system to reduce the burden of switching on the user, making it easier and faster to switch to the users intended next joint or function. The goal of this work is to demonstrate that adaptive switching, as a core application of machine learning, could have a direct effect on reducing the effort of amputee users operating complex multifunctional prosthetic devices.

## Methods

### Simple Task

In order to implement and assess adaptive switching, a transhumeral amputee subject and an able-bodied subject were recruited to perform a simple, semi-repetitive task using an experimental robotic arm. The amputee subject was a body-powered prosthetic user and had no experience using myoelectric control or using our experimental robotic arm. The able-bodied subject had previous experience controlling myoelectric devices. In the amputee user, electrodes were attached to the skin over the subjects wrist extensor muscle on the intact arm, which provided a signal for switching

between robotic joints. Separate sets of electrodes were attached to the biceps and triceps muscle of the residual limb. Those electrodes became the source of control signals for antagonistically moving selected joints of the robot arm. In the able-bodied subject, one set of electrodes was placed over the wrist extensor muscle on one arm, and separate sets of electrodes were placed over the wrist flexor and extensor muscles of the opposite arm. An 8-channel Bagnoli EMG system (Delsys Inc.) was used in the acquisition of EMG control signals from the experimental subjects at a frequency of 1 kHz. The subjects provided written informed consent to participate, and the trial was approved by the human research ethics board at our institution.

We used a custom-built research platform known as the Myoelectric Training Tool (MTT) in our experiments.<sup>19</sup> 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.<sup>19</sup> 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 1a shows the amputee subject using the MTT to perform a simple task. The MTT operated within a table-top task workspace centred on the robots axis of humeral rotation.

Each subject was given time to become familiar with the operation of the MTT. After familiarization, the subjects were presented with a specific task that involved a subset of the available joints (hand open/close, wrist flexion/extension, elbow flexion and extension, and humeral internal/external rotation). The task was chosen to be functionally comparable to tasks of daily living such as picking up a dish and placing it on a shelf. The instruction given to each subject in both the non-adaptive and adaptive trials was to manipulate the MTT to repetitively open and close the hand (i.e., as if grasping and releasing an imaginary object) on one side of the task workspace, perform humeral rotation to the opposite side of the workspace, repeatedly flex and extend the wrist joint (i.e., as if waving), and then perform humeral rotation back to the starting side of the workspace. Each trial involved repeating this sequence as many times as possible for 3 minutes.

Two types of trials were performed in order to test the predictive capabilities of our proposed control approach: trials using our adaptive switching algorithm (“adaptive trials”) were compared with conventional non-predictive switching methods (“non-adaptive trials”). Three 3-minute trials were done for each condition of non-adaptive and adaptive switching, and in both types of trials, an auditory cue was provided to the subject upon switching, naming the selected joint.

In non-adaptive trials, the subjects switched their myoelectric control between four joints in a fixed switching order: hand, wrist, elbow, and humeral rotation. This non-adaptive setup was selected to reflect how traditional direct myoelectric control is programmed in order to access all 4 joint motions with limited control sites. In contrast, in adaptive trials, the joints were continually reordered in the switching list in real time, based on their predicted likeli-

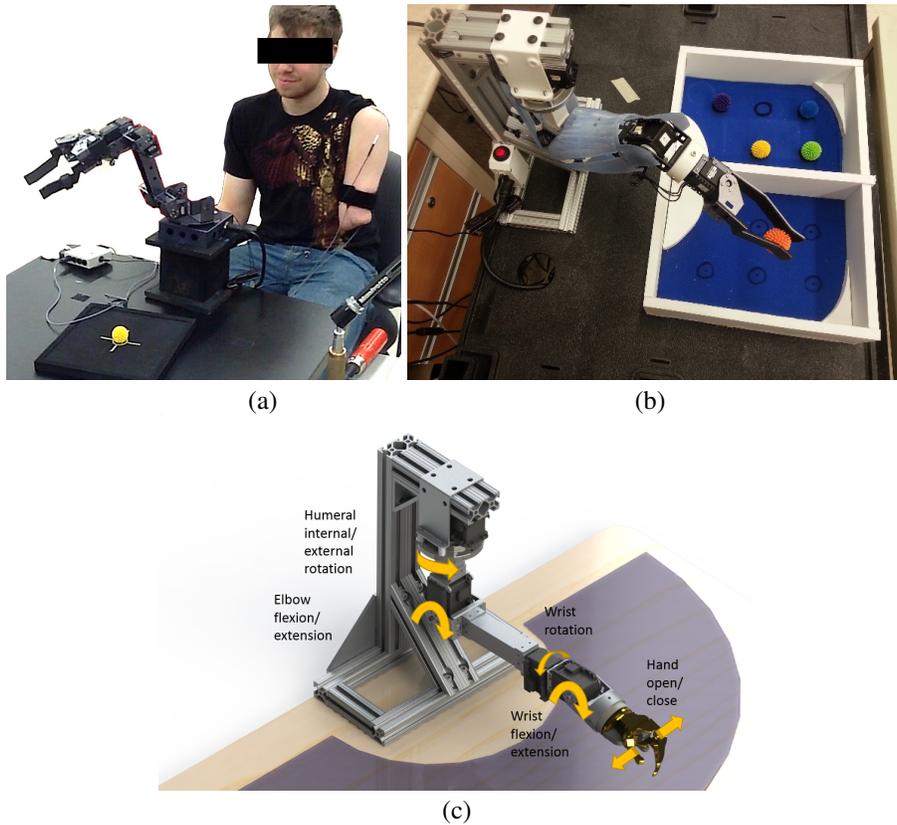


Figure 1: (a) Amputee participant performing simple tasks with the MTT arm using myoelectric control signals; (b) the Bento Arm, controlled by an able-bodied subject performing a modified box-and-blocks task; and (c) a schematic of the Bento arm (without outer casing) depicting the different motions available to the user. The MTT also operates as depicted in (c).

hood of being used next. This was accomplished throughout the course of the task through the use of GVs. GVs allow the control systems programming to continually learn and adapt with each successive control signal received from the user and subsequent joint motion of the robot arm—in essence, the system learns to present the appropriate switch suggestion via on-going observations of a users actions.

The procedure for using GVF predictions during adaptive trials was as described by Pilarski et al.<sup>11</sup> Learned GVs are able to represent predictions about a subjects situation-specific use of each joint in a myoelectric switching list. In the current work, and in contrast to prior demonstrations, GVF predictions were learned during *real-time* robotic arm use by the subjects and were continually ranked based on their relative magnitudes. 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.

To learn real-time predictions of the next active joint, we combined sensorimotor data from the robot with EMG data from the human user. Each of the AX-18 motors that make up the joints of the MTT produced a number of useful sen-

sory signals, including measures of angular position, angular velocity, load (current), temperature, and voltage. We provided joint angular position and angular velocity observations to the learning system as information about the current state. Features based on the current state of the arm enabled 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 such that GVs started each trial with no stored knowledge (predictions) about the user or the task in question.

### **Modified Box-and-Blocks Task**

In order to demonstrate transferability of this approach to a functional task commonly used in the clinical setting, we conducted a second experiment with three able-bodied subjects using a different robotic arm. A more complex task was designed to evaluate predictive learning, based on a modified box-and-blocks task. The traditional box-and-blocks task assesses a subjects manual dexterity by counting the number of blocks a subject can move from one side of a divided box to the other in a predetermined amount of time.<sup>20</sup> Subsequent studies have shown that simplifying the task to fewer block movements does not result in a loss of valid information on performance.<sup>21</sup> For our experiments, we therefore used a modified task that measured the amount of time re-

quired for a subject to move five rubber balls from one side of a box to the other.<sup>19</sup>

For this task, each subject controlled an anthropometric robotic arm called the Bento Arm (Figure 1b). The Bento Arm was designed with the MX series of actuators (Robotis Inc.), which are more powerful and robust than the AX-18 actuators.<sup>22</sup> The Bento Arm, similar to the AX-18 smart robotic arm, can be controlled by switching between its five degrees of freedom: humeral internal/external rotation, elbow flexion/extension, wrist rotation, wrist flexion/extension, and gripper open/close.

Two of the three able-bodied subjects had no previous experience controlling a myoelectrically driven robotic arm. Myoelectric signals were acquired from each subject through an EMG setup identical to the Bagnoli setup used on the simple task. The Bento Arm was controlled using the Robot Operating System (ROS) in a multi-computer configuration, with a central computer handling the direct communication with the arm and recording the EMG. A second computer recorded data over TCP/IP communication using ROSs data logging functionality. Visualization and management of EMG and robotic arm parameters was managed on a third computer.

After a period of familiarization, subjects were asked to control the Bento Arm, alternating between adaptive and non-adaptive switching trials (for a total of three adaptive and three non-adaptive trials). Each trial consisted of five iterations of moving all five balls from one side of the divider to the other. Throughout the trials, position and velocity signals from each of the motors in addition to torque from the gripper motor were used by the GVF learning system to build up expectations about future switching decisions and to continually rank the Bento Arms joints in real time (i.e., re-order the switching list presented to the user).

## Results

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 amputee subject. Each switching event began when the user triggered a joint switch, and ended when the user initiated movement of any of the MTT joints. Therefore, all switches made while shifting control to a new joint were counted as a single switching event. In adaptive switching mode, the system very quickly adjusted to choose the correct joint for each part of the task with only one exception, as compared to the non-adaptive mode which consistently required additional switches through the trials.

Figure 3 shows that with the simple task performed by both subjects, the average amount of time (measured in seconds) dedicated to switching and the total number of switches required to complete the task were significantly less for the adaptive trials. Figure 4 compares each iteration of the non-adaptive and adaptive switching trials in the modified box-and-blocks task, averaged over three data sets, where each iteration involved moving all five of the balls from one side of the box to the other. Figure 4a illustrates the mean total time spent by each subject completing each iteration of the task, and Figure 4b illustrates the mean total num-

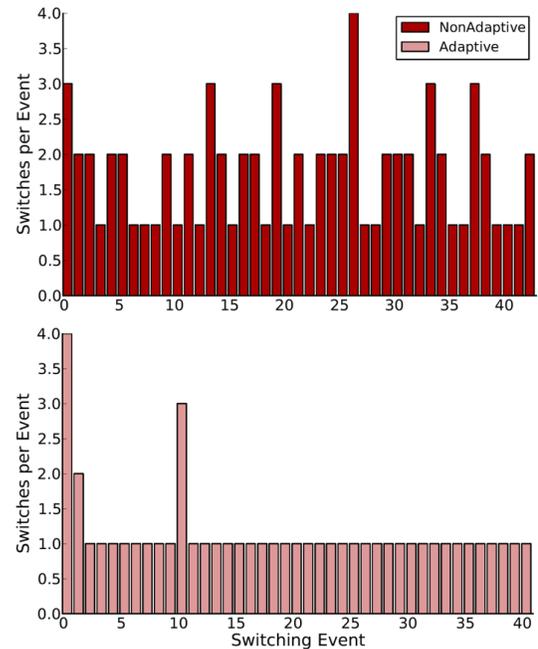


Figure 2: Number of voluntary switches initiated by the amputee subject per switching event over the course of a single 3 min trial. Shown for both non-adaptive (top) and adaptive control (bottom) approaches. The figure is also representative of switching by the non-amputee subject.

ber of switches required for each subject per iteration of the task. The dotted line in Figure 4b represents the minimum number of switches required to complete one iteration of the non-adaptive trial (i.e., the best-case performance that a subject can obtain using the static switching list for this task); the dashed line depicts the optimal number of switches (i.e., the number of switches required if the system predicted each joint with 100% accuracy and switched perfectly with no errors). To perform a single iteration of the box-and-blocks task using the non-adaptive switching list for this robot, a minimum of 75 switches would be required. As illustrated by Figure 4b, after the first or second iteration (in which the system is still learning), adaptive switching required fewer than 75 switches to complete each iteration. Furthermore, adaptive switching came close to meeting the optimal target of 35 switches. For the modified box-and-blocks task, 35 switches represents the perfect situation where the user made no errors and was given the correct joint every time they switched, with only one manual prompt per switching event. For all subjects, both experienced and inexperienced, in the fifth and final iteration of the modified box-and-blocks task adaptive switching demonstrated significant improvements in terms of both the number of switches required to complete the task iteration and the task completion time.

Figure 5 shows the mean time spent switching and the total number of switches per trial for adaptive and non-adaptive switching during the box-and-blocks task. For the experienced myoelectric user (Subject 1), the mean comple-

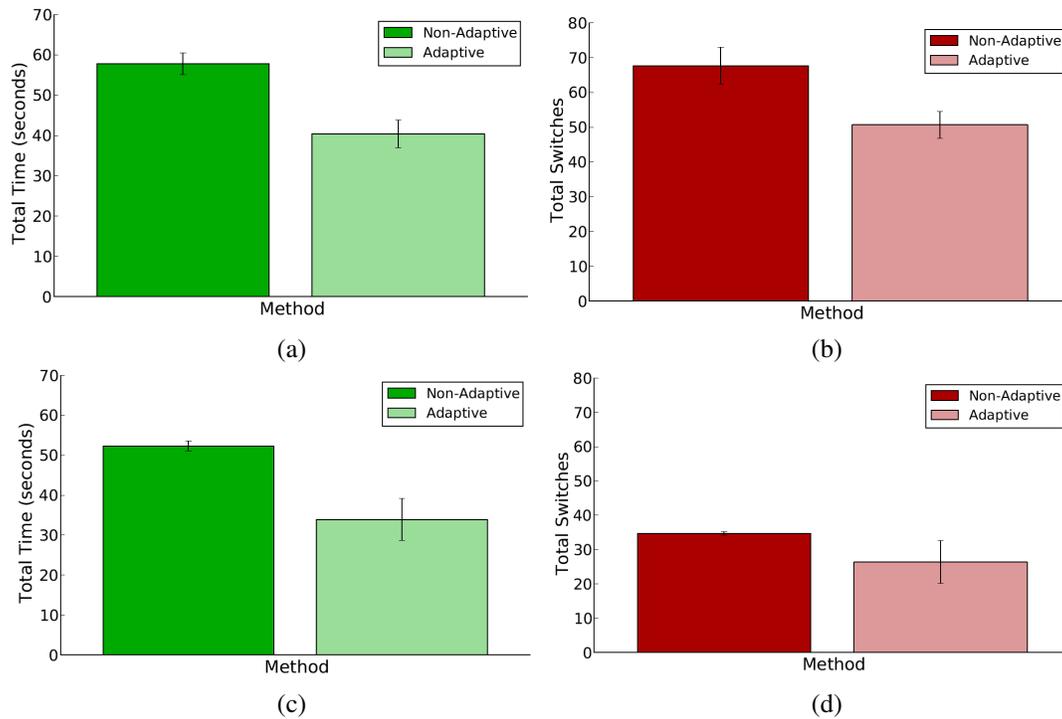


Figure 3: (a) Average time the amputee subject spent switching per trial when using non-adaptive and adaptive switching (left and right, respectively, average over 3 trials); (b) average number of switches made by the amputee subject per trial when using non-adaptive and adaptive switching (left and right, respectively, average over 3 trials); (c) average time the able-bodied subject spent switching per trial when using non-adaptive and adaptive switching (left and right, respectively, average over 3 trials); (d) average number of switches made by the able-bodied subject per trial when using non-adaptive and adaptive switching (left and right, respectively, average over 3 trials).

tion times for the adaptive and non-adaptive trials of the box-and-blocks task were  $11.9 \pm 1.2$  min and  $13.4 \pm 0.8$  min, respectively, whereas the mean times spent switching were  $4.9 \pm 0.4$  min and  $6.5 \pm 0.1$  min, respectively (Figure 5a). The mean total number of switches for this subject was  $337 \pm 15$  for the adaptive trials and  $463 \pm 22$  for the non-adaptive trials (Figure 5b). The non-experienced myoelectric users spent  $12.6 \pm 1.1$  min and  $14.2 \pm 1.2$  min completing the adaptive trials; they spent  $13.7 \pm 2.1$  min and  $14.3 \pm 0.6$  min completing the non-adaptive trials. They also spent  $9.8 \pm 1.2$  min and  $10.5 \pm 1.1$  min switching during adaptive trials and  $11.0 \pm 3.2$  min and  $10.5 \pm 0.7$  min switching during non-adaptive trials (Figure 5a). The mean total number of switches for adaptive trials were  $359 \pm 29$  and  $369 \pm 30$ ; for non-adaptive trials, the mean number of switches required were  $528 \pm 10$  and  $503 \pm 17$  (Figure 5b).

## Discussion

The goal of this work was to directly compare adaptive to non-adaptive switching to determine if there is potential utility of adaptive switching in reducing the effort and burden of controlling a robotic arm with more degrees of freedom available than direct control sites. There was a significant difference between non-adaptive switching and adap-

tive switching during the simple task. With adaptive switching enabled, after a brief initial period of learning by the system (i.e., the first several switching events), typically only one switch was required by the user to move control to the most appropriate joint. Adaptive switching also produced a large decrease in time spent switching compared with non-adaptive switching. For each 3-minute trial with the MTT, each subject saved an average of about 20 seconds when adaptive switching was enabled (11% of the total task time). This could have potential implications on prosthetic users performing more complex tasks requiring multiple joint modal switching by reducing the total amount of time and burden to affect the desired joint control. If joint switching could be made more intuitive (i.e., if the prosthesis reliably selected the correct joint at the right time), it might encourage prosthetic users to utilize additional joint control motions more often, rather than to deliberately restrict control options.

Adaptive switching was also able to reduce time and switching effort in the more complex task. In comparing the decrease in time and number of switches (Figure 4) between the first and fifth task iterations, for all subjects there was only a minor decrease during non-adaptive switching, which is most likely due to improved performance with experience.

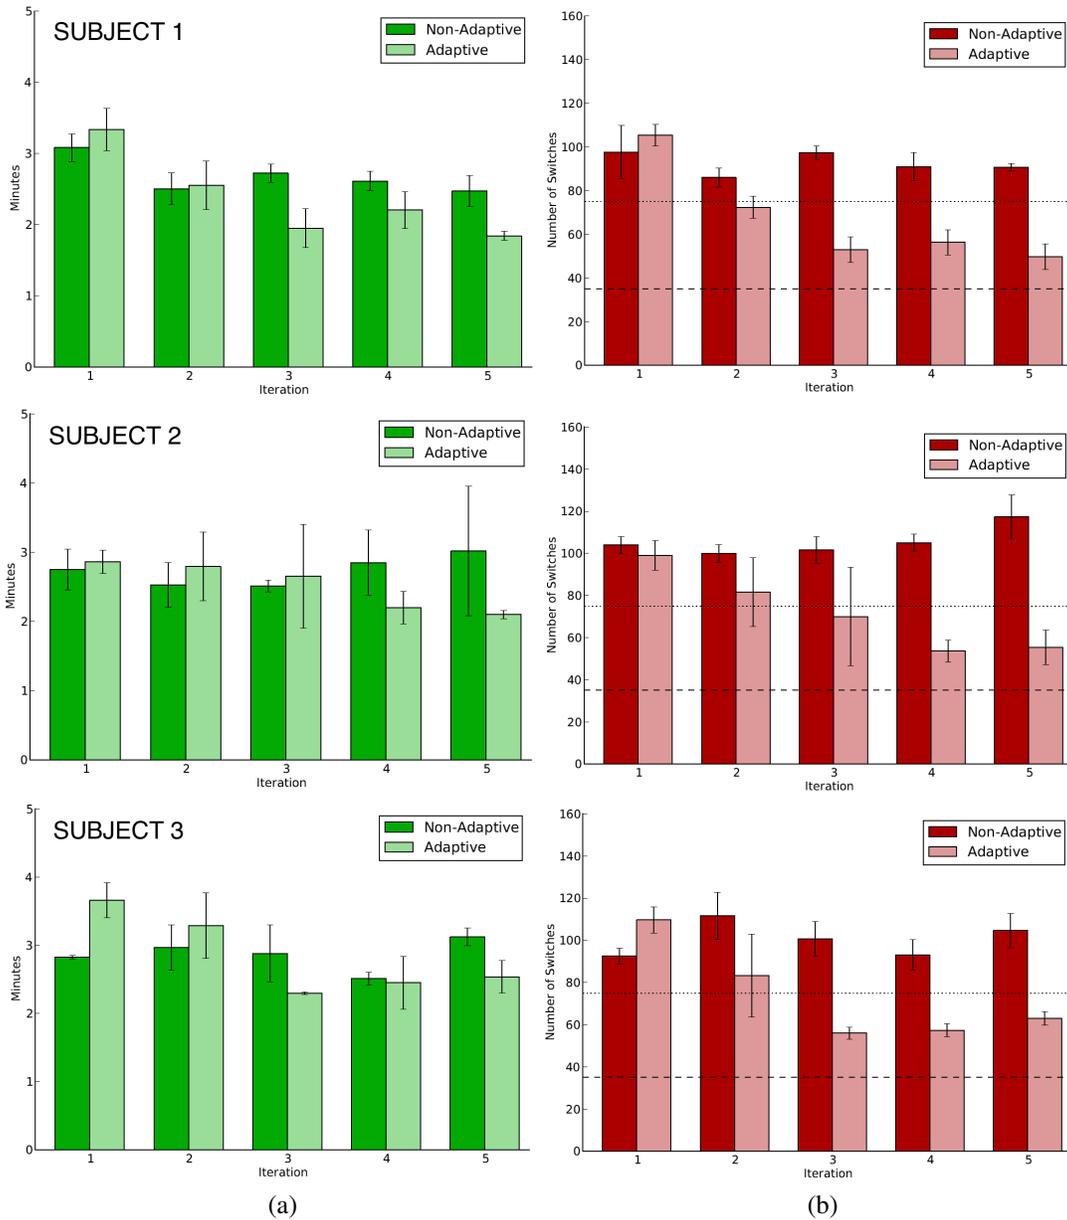


Figure 4: Comparison between adaptive and non-adaptive switching averaged over 3 trials. Total time (a) and number of switches (b) of the modified box-and-blocks task are shown by iteration of the task. Dotted line shows the minimum number of switches possible using a static switching list. Dashed line shows the best possible switching performance possible on this task with no user errors and perfect switching (a single switch) at each switching event.

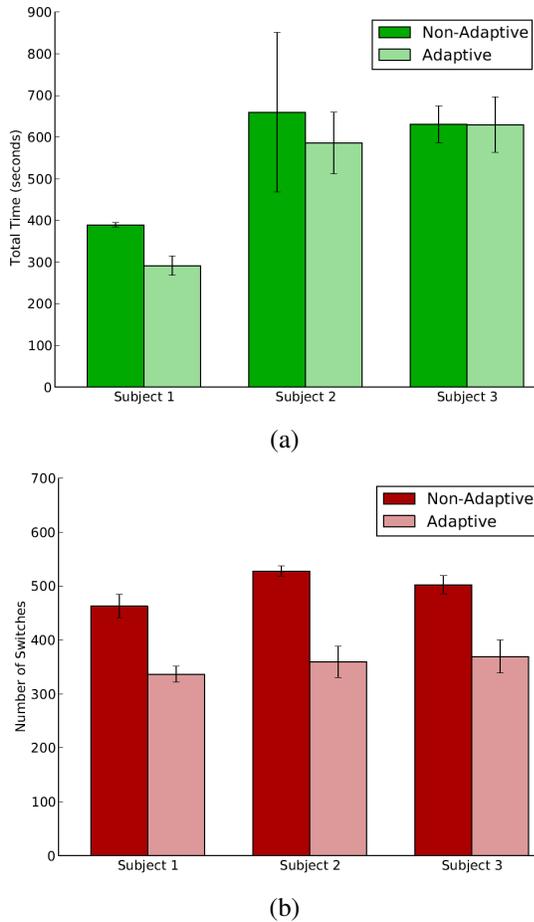


Figure 5: (a) Average time the able-bodied subjects spent switching per trial when using non-adaptive and adaptive switching (average over 3 trials); (b) average number of switches made by the able-bodied subjects per trial when using non-adaptive and adaptive switching (average over 3 trials).

However, the decrease in time and switches was more rapid and greater for the adaptive switching. By the final iteration of the task, the average number of switches made during adaptive switching decreased to approximately half that of non-adaptive switching. Consequently, by the final iteration, adaptive switching saved the experienced subject more than 30 seconds, or 20% of the total time.

The potential impact that user training may have on the functionality of adaptive switching is suggested by the relationship between the per-iteration data in Figure 4 and the aggregate data in Figure 5. In those datasets, it is clearly evident that the use of adaptive switching reduced the number of switching interactions needed by the user to well below the level possible with non-adaptive switching (Figures 4b and 5b). Indeed, by the third iteration of the task, experienced and non-experienced users averaged less switches per iteration than the 75-switch best-case performance of the non-adaptive approach. If the time to perform a switch and

resume motion was constant for all subjects, the switching data would suggest significant time savings for both experienced and non-experienced users (as in the simple task, Figure 3). However, for non-experienced myoelectric users, the difference in the total time spent switching between the adaptive and non-adaptive methods was less (Figure 5a, Subjects 2 and 3). Non-experienced users were observed to hesitate more following each adaptive switching event and before moving the selected prosthetic joint, primarily in earlier task iterations (as seen in the completion times for earlier iterations as compared to later iterations; Figure 4a). One subject indicated that during adaptive switching he paused slightly after each switch to determine the correctness of the choice made by the adaptive switching system. Conversely, the experienced myoelectric user reported greater trust in the choice made by the adaptive switching system, and demonstrated less total time and per-iteration time to complete the task (Figures 4a and 5a, Subject 1). Based on both these quantitative and qualitative observations, we believe that with more repetitions and greater experience with the adaptive switching, our non-experienced users would demonstrate time savings similar to those of the experienced user. More detailed studies are needed to determine the roles that trust, experience, and training play in the successful use of adaptive switching and related control adaptation paradigms.

The results from both the simple and more complex tasks suggest there are efficiencies with adaptive switching, and they agree with our expectations inferred from the simple task presented to the subject: there were clear regions of the task space that corresponded to the use of specific joints. For the simple 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 of the task design 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 the task (thus implementing a form of adaptation, or ongoing self-calibration, which has been pointed out to be of great clinical interest as it removes the need for regular recalibration by clinical staff<sup>1,9</sup>). Furthermore, we have observed that as the task changes or becomes more complex, as in the case of the modified box-and-blocks task, the learning system can scale up naturally and easily without the need for manual tuning.

It is important to note that while our results closely align with expectations from prior work,<sup>11,12</sup> the present experiments are single case studies which limits the generalizability of the conclusions. In addition, transferability to other tasks will need to be explored in order to assess relevance of this approach to the multitude of tasks required of prosthetic users in daily life. One strength of our proposed approach is that it is able to optimize prosthetic control for a repeated pattern of prosthetic movements, without that pattern being explicitly specified by a user or clinician. The adaptive switching system begins with simple determinis-

tic switching, as would be familiar to myoelectric prosthesis users, and gradually optimizes control as regularities are observed in the movements made by the user. Because of the way that information is presented to the adaptive switching system (*function approximation*, c.f., Pilarski et al.<sup>11</sup>), the system is able to leverage similarities and generalities in the movements being performed—i.e., it is able to form generalizations that are applicable to multiple tasks, such as “after reaching forward the user usually grabs an object”. While these generalities allow some translation of learned patterns to novel tasks, it is important to note that the system will still need experience with a new class of tasks before being able to form a reasonable set of adaptive switching suggestions. For daily-life usability, we suggest that adaptive switching may also need to be engaged and disengaged based on the control systems confidence in its own predictions; in other words, it may be more intuitive for users if prosthesis control defaults to simple deterministic switching until the system has enough experience to appropriately adapt the switching list in a given setting. How best to engage and disengage adaptive switching in this way, and what function approximation methods will allow the best generalization between real-world tasks remain open questions. Future work will also need to identify the most effective way to communicate the systems switching suggestions back to the user—auditory (as used here), visual, tactile, and vibratory signals are all valid possibilities for feedback.

Ideally these experiments will be repeated in multiple subjects with a range of experience in myoelectric control, and on subjects using donned prostheses during daily-life tasks. However, the current case study is a necessary and encouraging first demonstration of the practical applicability of our approach with both able-bodied and amputee subjects.

## Conclusion

The primary contribution of this paper is a concrete demonstration of adaptive switching in an applied setting of robotic arm control, in both able-bodied subjects and an amputee subject. This study is the first time that real-time prediction learning has been used to improve the control interface of a prosthetic device during uninterrupted use. Our experiments with an amputee subject showed that for simple tasks, enabling adaptive switching on a robotic arm significantly decreased the time spent switching between functions. This extends previous work using pre-recorded (non-real-time) data that indicated the potential merit of adaptive switching. These results have been extended to a box-and-blocks task that is more representative of complex everyday tasks. We believe that adaptive switching can help to decrease the time and cognitive load required by amputees during complex tasks and real-world functional situations involving wearable prostheses. In future work we will study the use of adaptive switching in shared-control functional tasks with prosthetic users, wherein switching control is further delegated to a control system to reduce the cognitive burden on the user.

## Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

## Ethical approval

Study conducted as approved by the Human Research Ethics Board at the University of Alberta.

## Funding

This work was supported by grants from the Alberta Innovates –Technology Futures (AITF), the Natural Sciences and Engineering Research Council of Canada (NSERC) (RGPIN-2015-03646), the Alberta Innovates Centre for Machine Learning (AICML), and the Glenrose Rehabilitation Hospital Foundation.

## References

1. Scheme E and Englehart KB. Electromyogram pattern recognition for control of powered upper-limb prostheses: State of the art and challenges for clinical use. *J Rehabil Res Dev* 2011; 48(6): 643–660.
2. Peerdeman B, Boere D, Witteveen H, Huis in 't Veld R, Hermens H, Stramigioli S, Rietman H, Veltink P and Misra S. Myoelectric forearm prostheses: State of the art from a user-centered perspective. *J Rehabil Res Dev* 2011; 48(6): 719–738.
3. Micera S, Carpaneto J and Raspopovic S. Control of hand prostheses using peripheral information. *IEEE Rev Biomed Eng* 2010; 3: 48–68.
4. Parker P, Englehart KB and Hudgins B. Myoelectric signal processing for control of powered limb prostheses. *J Electromyogr Kines* 2006; 16(6): 541–548.
5. Williams TW. Guest editorial: Progress on stabilizing and controlling powered upper-limb prostheses. *J Rehabil Res Dev* 2011; 48(6): ix–xix.
6. Pilarski PM, Dawson MR, Degris T, Carey JP, Chan KM, Hebert JS and Sutton RS. Adaptive artificial limbs: A real-time approach to prediction and anticipation. *IEEE Robot Autom Mag* 2013; 20(1): 53–64.
7. Resnik L, Meucci MR, Lieberman-Klinger S, Fantini C, Keltly DL, Disla R and Sasson N. Advanced upper limb prosthetic devices: Implications for upper limb prosthetic rehabilitation. *Arch Phys Med Rehabil* 2012; 93(4): 710–717.
8. Castellini C, Artemiadis P, Wininger M, Ajoudani A, Alimusaj M, Bicchi A, et al. Proceedings of the first workshop on peripheral machine interfaces: Going beyond traditional surface electromyography. *Front Neurorobot* 2014; 8: Article 22.
9. Sensinger J, Lock B and Kuiken T. Adaptive pattern recognition of myoelectric signals: Exploration of conceptual framework and practical algorithms. *IEEE Trans Neural Syst Rehab Eng* 2009; 17(3): 270–278.

10. Tommasi T, Orabona F, Castellini C and Caputo B. Improving control of dexterous hand prostheses using adaptive learning. *IEEE Trans Robot* 2013; 29(1), 207–219.
11. Pilarski PM, Dawson MR, Degrís T, Carey JP and Sutton RS. Dynamic switching and real-time machine learning for improved human control of assistive biomedical robots. In: *Proc 4th IEEE RAS EMBS Int Conf Biomed Robot Biomechatron*. Rome, Italy, 2012, pp. 296–302.
12. Edwards AL, Kearney A, Dawson MR, Sutton RS and Pilarski PM. Temporal-difference learning to assist human decision making during the control of an artificial limb. In: *Proc 1st Multidiscip Conf Reinf Learn Decis Making*. Princeton, NJ, 2013, arXiv:1309.4714 [cs.AI].
13. Pilarski PM, Dick TB and Sutton RS. Real-time prediction learning for the simultaneous actuation of multiple prosthetic joints. In: *Proc 2013 IEEE Int Conf Rehabil Robot*. Seattle, USA, 2013, pp. 1–8.
14. Sutton RS, Modayil J, Delp M, Degrís T, Pilarski PM, White A and Precup D. Horde: a scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In: *Proc 10th Int Conf Auton Agent Multi Agent Syst*. Taipei, Taiwan, 2011, pp. 761–768.
15. Sutton RS and Barto AG. Time-derivative models of pavlovian reinforcement. In: Gabriel M and Moore J (eds) *Learning and computational neuroscience: Foundations of adaptive networks*. Cambridge MA: MIT Press, 1990; pp. 497–537.
16. Wolpert DM, Ghahramani Z and Flanagan JR. Perspectives and problems in motor learning. *Trends Cogn Sci* 2001; 5(11): 487–494.
17. Redish AD. *The mind within the brain: How we make decisions and how those decisions go wrong*. New York: Oxford University Press; 2013.
18. Linden DJ. From molecules to memory in the cerebellum. *Science* 2003; 301(5640): 1682–1685.
19. Dawson MR, Fahimi F and Carey JP. The development of a myoelectric training tool for above-elbow amputees. *Open Biomed Eng J* 2012; 6: 5–15.
20. Mathiowetz V, Volland G, Kashman N and Weber K. Adult norms for the box and blocks test of manual dexterity. *Am J Occup Ther* 1985; 39(6): 386–391.
21. Hebert JS, Lewicke J, Williams TR and Vette AH. Normative data for a modified Box and Blocks test measuring upper limb function via motion capture. *J Rehabil Res Dev* 2014; 51(6): 919–932.
22. Dawson MR, Sherstan C, Carey JP, Hebert JS and Pilarski PM. Development of the Bento Arm: An improved robotic arm for myoelectric training and research. In: *Proc Myoelectric Control Symposium*. Fredericton, NB, 2014, pp. 60–64.