Progress Toward the Shared Control of a Prosthetic Arm

Ann L. Edwards; Michael R. Dawson; Jacqueline S. Hebert; Craig Sherstan; Richard S. Sutton; K. Ming Chan; Patrick M. Pilarski

Reinforcement Learning and Artificial Intelligence Laboratory, Dept. Computing Science & Dept. of Medicine, University of Alberta, Edmonton, Alberta, Canada

Abstract

State-of-the-art myoelectric arms typically have a greater number of functions than the possible number of control signals, requiring amputees to manually switch through a fixed list to select a desired function. Previous studies have demonstrated that reinforcement learning techniques, in particular General Value Functions (GVFs), can be applied to develop temporally extended predictions about signals related to prosthetic arm movement. Using GVFs, we can learn and update a list of possible prosthetic arm functions, termed adaptive switching. In this work, we demonstrate the real-time use of adaptive switching by subjects in a simple control task with a myoelectric arm. We also present results from subjects controlling a myoelectric arm in a more complex task, providing evidence for the scalability of the learning system. Our results suggest that adaptive switching can significantly decrease the amount of time and the number of switches required for the control of a robotic arm. We anticipate the future blending of human and machine decision making for the shared control of a robotic arm.

Background

Reinforcement learning (RL) methods in the form of General Value Functions (GVFs) have been shown to accurately capture anticipatory knowledge during human-robot interaction (Pilarski et al. 2013). GVFs represent temporally extended predictions about signals of interest; they can be learned in real time using standard RL methods during the use of an assistive robotic system—e.g., a powered artificial arm or leg.

Figure 1. Example of the control disparity facing prosthesis users

The Switching Problem: Powered artificial limb controllers use recorded muscle signals (electromyographic recordings, or EMG) to inform their control decisions. In more advanced prostheses there is often a disparity between the number of available EMG recording sites on an amputee's body and the number of controllable functions on the prosthesis (Figure 1). As such, an amputee can actuate only a small subset of a device's function at any given time. Switching between functions in a fixed or pre-determined order can help increase the number of useable functions, but can increase both the actuation time and cognitive effort.

Methods

Two myoelectrically controlled robot arms were used in separate tasks (Figure 2). The first robot arm was used by subjects to perform a simple task, and the second arm was used to perform a more complex task. In both tasks, electrodes were affixed to the skin of amputee and non-amputee subjects and used to measure EMG signals from muscles on the user's arm. These EMG signals were mapped to two control channels: one to actuate a robotic joint, and one to switch between joints, sequentially.

Two types of switching methods were tested:

- **Adaptive switching:** Joints of the arm were continually reordered in the switching list based on their predicted likelihood of being used next
- **Non-adaptive switching:** Joints were presented as a fixed list; the user must switch through joints one by one in order to select and use a joint

Figure 2. Experimental platforms

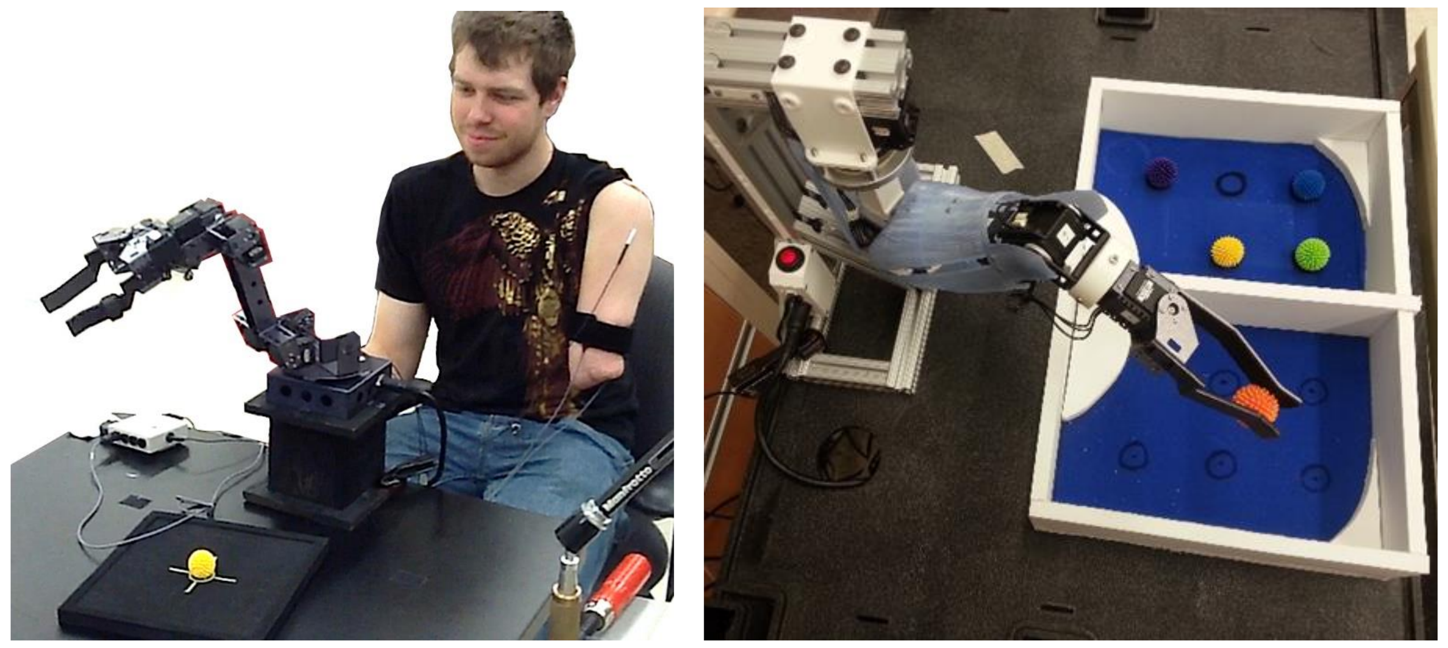

Simple Task: One amputee and one non-amputee subject used the myoelectric interface in Figure 2 (left) to perform a repetitive task with the robot for 3 minutes. This task involved rotating the shoulder joint back and forth, waving the wrist joint, and opening and closing the hand.

Box and Blocks Task: Three non-amputee subjects used the Bento arm, shown in Figure 2 (right), to perform a complex task that involved repeatedly moving 5 balls from one side of a divided box to the other (5 iterations total).

Predictions regarding signals of interest were acquired through online GVF-based reinforcement learning:

- TD-learning
- Eligibility traces
- Tile-coding function approximation

The following parameters were used in the state representation:

- Motor positions
- Motor velocities
- Motor torque (gripper)

Predicted signals of interest:

- The user's manual prompt to initiate switching between joints
- The motion (activity) of each user-driven joint

Results

Following a period of learning, adaptive switching significantly reduced the time and number of switches required to successfully complete a task.

Figure 3. Number of switches required per switching event during a 3 minute trial of simple task (single amputee subject)

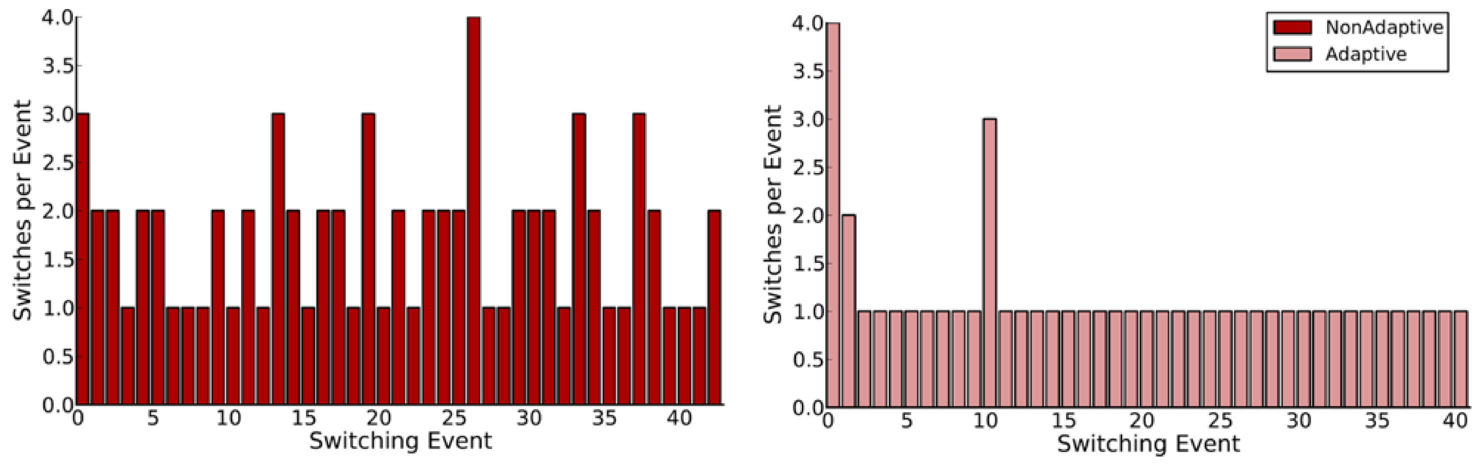

With adaptive switching, after an initial period of learning by the system, typically only one switch was required by the user.

Figure 4. Mean time and number of switches per iteration of box and blocks task (single non-amputee subject)

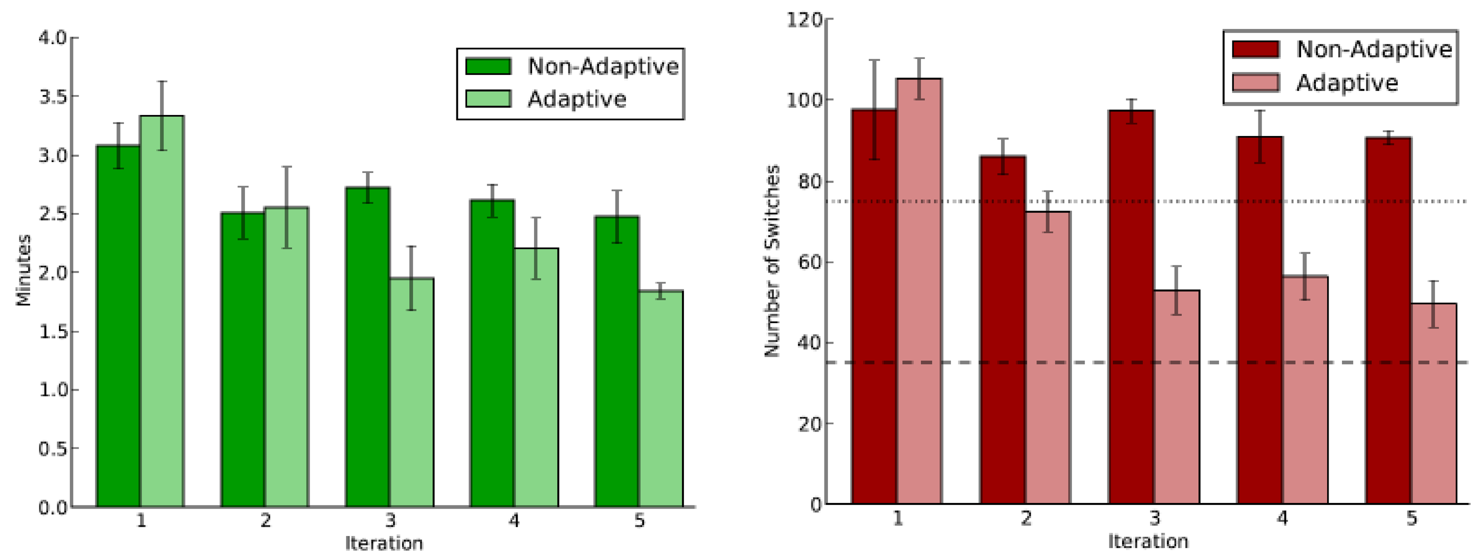

Adaptive switching required fewer than 75 switches (the minimum for non-adaptive) to complete each iteration, and approached 35 switches (the optimal number of switches, where the user made no mistakes and was given the correct joint each time they switched).

Conclusions

This work presents a concrete demonstration of adaptive switching in an applied setting. Real-time prediction learning was used for the first time to improve the control interface of a prosthetic device during uninterrupted use by both amputee and non-amputee subjects.

Based on our results, we believe that adaptive switching can help decrease the time and cognitive load required by amputees during complex tasks and real-world functional situations.

In future work we will study the use of adaptive switching in shared-control tasks wherein switching control itself may be further delegated to a control system to reduce the cognitive burden on the user.

The authors gratefully acknowledge support from the Alberta Innovates Centre for Machine Learning, Alberta Innovates – Technology Futures, the National Science and Engineering Research Council (NSERC), and the Glenrose Rehabilitation Hospital Foundation. The authors also thank Joseph Modayil, Thomas Degris, and Jason Carey for a number of helpful discussions leading up to this work. **Contact**: pilarski@ualberta.ca