Towards Prediction-Based Prosthetic Control

Pilarski PM[1], Dawson MR[2], Degris T[3], Carey JP[4], Chan KM[5], Hebert JS[2], Sutton RS[1]

[1] Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada

[2] Glenrose Rehabilitation Hospital, Edmonton, Alberta, Canada

[3] INRIA, Bordeaux, France

[4] Department of Mechanical Engineering, University of Alberta, Edmonton, Alberta, Canada

[5] Centre for Neuroscience, University of Alberta, Edmonton, Alberta, Canada

Abstract

Predictions and anticipations form a basis for pattern-recognition-based control systems, including those used in next-generation prostheses and assistive devices. In this work, we outline how new methods in real-time prediction learning provide an approach to one of the principal open problems in multi-function myoelectric control—robust, ongoing, amputee-specific adaptation. Techniques from reinforcement learning and general value functions are applied to learn and update predictions during continuing human interactions with multi-function robotic appendages. Tests were performed with a targeted motor and sensory reinnervation amputee subject and non-amputee participants. Results show that this online prediction learning approach is able to accurately anticipate a diverse set of human and robot signals by more than two seconds, and at different levels of temporal abstraction. These methods allow predictions to be adapted during real-time use, which is an important step toward the intuitive control of powered prostheses and other assistive devices.

Keywords: real-time machine learning, adaptive control, myoelectric control, prosthetics, assistive medical devices.

Introduction

Simultaneous myoelectric control of multiple joints remains a challenging unsolved problem [1]. In particular, conventional myoelectric control is faced with the challenge of scaling up to the expanded sensing and actuation capabilities of new multi-function prosthetic devices. Even with advanced function switching, conventional control approaches are only able to make use of a fraction of the movements available to next-generation devices like the Johns Hopkins Modular Prosthetic Limb. This problem becomes more pronounced with more proximal levels of amputation. While patients with higher levels of limb loss require more assistive technology to replace lost function, they also have fewer discrete sources (i.e., remaining muscles) from which to acquire control information for their prosthetic devices [1,2].

These challenges for myoelectric prosthesis users are being addressed by improvements in surgical techniques, new sensor and actuator technology, and advanced control and pattern-recognition paradigms. Targeted motor and sensory reinnervation surgery is opening new ground for intuitive control and feedback [3]. At the same time, physical platforms for sensation and actuation are being improved through new biomedical device technologies and work on high-degree-of-freedom limb systems.

An important avenue for improving myoelectric control has been the use of machine learning and pattern recognition techniques [2,4,5]. The state-of-the-art for myoelectric pattern recognition relies on sampling a number of training examples in the form of recorded signals, identifying relevant features within these signals, and then classifying these features into a set of control commands [2,4,5]. This approach has been largely implemented in an offline context, meaning that systems are calibrated and then not changed significantly during regular (non-calibration) use by an amputee; demonstrated offline methods include support vector machines, linear discriminant analysis, artificial neural networks, and principal component analysis using time and frequency domain EMG information.

Work to date has made it clear that effective myoelectric control must take into account real-time changes to the control environment, patient physiology, and the prosthetic hardware [6]. However, a robust, unsupervised approach to online adaptation has yet to be demonstrated [2].

In this article we outline the use of online, real-time prediction learning as one foundation for truly adaptive myoelectric devices. An extended presentation of these methods and results will be provided in the future via a full-length article (under review) [7].

Materials and Methods

Real-time Machine Learning

Reinforcement learning (RL) is one form of machine learning that has demonstrated the ability to learn in an ongoing, incremental prediction and control setting [8]. An RL system uses interactions with its environments to build up expectations about future events. Specifically, it learns to estimate the value of a one-dimensional feedback signal termed *reward*.

Recent work has provided a straightforward way to use RL for acquiring expectations pertaining to non-reward signals and observations [9]. Specifically, *general value functions* (GVFs) have been proposed as a way of asking and answering temporally extended questions about future sensorimotor experience [9]. Predictive questions can be defined for different time scales, and may take into account different methods for weighting the importance of future observations. Expectations comprising a GVF are acquired using standard RL techniques; this means that learning can occur in an incremental, online fashion, with constant demands in terms of both memory and computation. For the present work, GVFs were implemented as described by Sutton et al. [9] and Pilarski et al. [10]; anticipations were learned in an incremental fashion during online operation. This was done using temporal difference (TD) learning, a standard technique from RL [8].

Experimental Setup

We examined the ability of a GVF-based learning system to predict and anticipate human and robot signals during online interactions between able-bodied (non-amputee) and amputee subjects and a robotic device. These two domains are shown in Figs. 1 and 2. Specifically, we studied the anticipation user EMG signals, the angular position of user-controlled elbow and wrist joints of a robot limb, and the grip force detected by a robot's prehensor. The robotic platform used in able-bodied subject experiments was a Nao T14 robot torso (Aldebaran Robotics, France; Fig. 1), while the robot in the amputee trial was the Myoelectric Training Tool (MTT)—a system designed to help new amputees prepare for powered prosthesis use [11]. The MTT includes a robotic arm with five degrees of freedom that mimic the functionality of commercial myoelectric prostheses. These platforms provide a flexible interface to send actuator position, velocity, and stiffness commands to the robot, and receive real-time sensorimotor feedback for use by a controller or learning system. EMG signals used in device control and learning

were obtained using a Bagnoli-8 (DS-B03) EMG System with DE-3.1 Double Differential Detection EMG sensors (Delsys, Boston, USA). Electrode locations were selected according to EMG signal strength and differentiation on the arms of amputee and non-amputee subjects.


Fig. 1: Equipment for able-bodied subject trials, including an Aldebaran Nao T14 robotic platform, laptop computer, analog to digital converter, and Bagnoli-8 EMG recording equipment.


Fig. 2: Experimental setup for amputee trials, including the Myoelectric Training Tool (MTT) with robot arm, tactor, control computer, and EMG system.

Procedures

To generate a rich stream of sensorimotor data in an online, interactive setting, participants worked with the robotic platforms to complete a series of actuation tasks, including reaching and grasping motion, and (for able-bodied subjects) randomized elbow and wrist actuation tasks. Informed consent was acquired as per the study's ethics authorization from the University of Alberta Health Research Ethics Board.

Subjects for the non-amputee experiments included healthy participants with no cognitive or motor impairments. The subject for the amputee trial was a twenty year old male injured in a work accident sixteen months prior to the study, with a resultant left transhumeral amputation. Six months prior to the trial the subject underwent surgical revision of his limb, involving Targeted Muscle Reinnervation (TMR), as well as Targeted Sensory Reinnervation (TSR). Reinnervation resulted in a widely distributed discrete representation of digital sensation on the upper arm. To supplement MTT actuation via conventional myoelectric control, force information was communicated to the participant via a tactor (micro servomotor) placed over the reinnervated skin of his upper arm.

Results

Accurate Anticipation of Diverse Signals

These experiments showed that a GVF-based prediction learning system was able to successfully anticipate events initiated by both able-bodied and amputee subjects—i.e., joint angle changes, EMG signals, and effector force fluctuations. Accurate predictions were observed after training on only 16 minutes of data using real-time sampling and learning. Iterative training was found to further improve the accuracy of these predictions.

Figure 3 shows a representative example of grip force prediction after three passes through the training data. Here the normalized predictions at two time scales (red lines) are compared to the measured force sensor readings (grey line); prediction signals were found to anticipate changes to the actual measured signal by ~2–3 seconds.


Fig. 3: Example of force signal prediction on amputee testing data, shown after three learning iterations for two time scales (solid and dotted red lines). Normalized predictions (red lines) precede the measured force activity (wide grey lines) by 2–3 seconds.


Fig. 4: Example of myoelectric signal prediction during able-bodied subject trials, after ~10 min of learning. Normalized predictions (red line) precede the observed signal activity (wide grey lines) by 0.5–2.0s.

The system was also able to learn accurate predictions for time-averaged myoelectric signals (Fig. 4). Normalized EMG predictions (red line) rise visibly in advance of the actual myoelectric events (grey line); changes to the myoelectric signal are anticipated up to 1500 milliseconds before change actually occurs. Similar anticipatory capabilities were observed for robot elbow, wrist, and hand actuator signals in both amputee and non-amputee experimental domains.

Discussion

These results illustrate the potential utility for an online prediction approach within the myoelectric setting. Using GVFs, it was possible to learn accurate temporally abstracted predictions about a human's interaction with a robotic device. Few application-specific changes and tuning operations were needed to shift between the different experiments and signal types presented in this work. It would be straightforward to apply this approach to other assistive robotic domains.

Our online prediction learning approach is also well suited to practical real-time implementation. For both subject types, the system learned to interpret signals from structured and unstructured human interactions with a robotic platform. Learning was performed under both the online case (learning during ongoing experience) and offline case (learning from recorded data); no application-specific changes were needed to address these two domains, and all learning updates and predictions were made under real-time computation constraints and with realistic learning times (<30min). Offline and online predictions also share the same incremental learning framework. This observation further reinforces the idea that online prediction learning may provide a basis for rapidly adapting and calibrating factory-made devices for use by specific individuals.

We have shown in related work how extended predictions about a user's control behaviour and motor objectives can be used to prioritize control options for the user [10], or to directly co-ordinate the simultaneous movement of multiple joints [12]. The use of real-time prediction learning to enhance an amputee's control interactions is the subject of ongoing studies by our group, and preliminary results indicate tangible benefit in terms of both effort and cognitive load.

Conclusions

This work outlines an online approach to acquiring and continuously updating a set of predictions and anticipations about a human user and their assistive biomedical device. For both amputee and able-bodied subjects it was possible for a real-time machine learning system to provide advance knowledge about signals such as grip force and myoelectric activity. Our system demonstrated the ability to learn accurate predictions in both amputee and non-amputee settings, and could do so without the need to significantly alter the learning framework for each specific domain or signal of interest. This approach provides a fundamental tool to address one major unsolved problem for prosthetic devices: amputee-specific adaptation during ongoing use. Future work will examine the prediction of signals during complex real-world activities, and evaluate prediction-based control adaptation techniques for multi-joint robot actuation by amputees.

References

[1] T. W. Williams, "Guest Editorial: Progress on stabilizing and controlling powered upper-limb prostheses," *J. Rehab. Res. Dev.,* 48(6):ix–xix, 2011.

[2] 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," *J. Rehab. Res. Dev.,* 48(6):643–660, 2011.

[3] Todd A. Kuiken, Guanglin Li, Blair A. Lock, Robert D. Lipschutz, Laura, A. Miller, Kathy A. Stubblefield, Kevin B. Englehart, "Targeted muscle reinnervation for real-time myoelectric control of multifunction artificial arms," *JAMA: J. American Medical Association,* 301(6):619-628, 2009.

[4] P. Parker, K. B. Englehart, and B. Hudgins, "Myoelectric signal processing for control of powered limb prostheses," *J. Electromyography Kinesiol.,* 16(6):541–548, 2006.

[5] M. A. Oskoei and H. Hu, "Myoelectric control systems—a survey," *Biomed. Signal Process. Control*, 2(4):275–294, 2007.

[6] J. Sensinger, B. Lock, and T. Kuiken, "Adaptive pattern recognition of myoelectric signals: exploration of conceptual framework and practical algorithms," *IEEE T. Neural Syst. Rehab. Eng.*, 17(3):270–278, 2009.

[7] P.M. Pilarski, T. Degris, M.R. Dawson, K.M. Chan, J.S. Hebert, J.P. Carey, and R.S. Sutton, "Prediction and Anticipation for Adaptive Artificial Limbs," 9 pages (*submitted*) March 2012.

[8] R. S. Sutton and A. G. Barto, *Reinforcement learning: an introduction.* MIT Press Cambridge, Massachusetts; London, England, 1998.

[9] R. S. Sutton, J. Modayil, M. Delp, T. Degris, P. M. Pilarski, A. White, D. Precup, "Horde: a scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction," in *Proc. of 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, May 2-6, 2011, Taipei, Taiwan, pp. 761–768, 2011.

[10] 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," *Proc. of the 4th IEEE International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, June 24-28, Roma, Italy, 7 pages, 2012.

[11] M. R. Dawson, F. Fahimi, J. P. Carey, "The development of a myoelectric training tool for above-elbow amputees," *The Open Biomedical Engineering Journal*, 6:5–15, 2012.

[12] P. M. Pilarski, M. R. Dawson, T. Degris, F. Fahimi, J. P. Carey,and R. S. Sutton, "Online human training of a myoelectric prosthesis controller via actor-critic reinforcement learning," in *Proc. 2011 IEEE International Conference on Rehabilitation Robotics*, pp. 134–140, 2011.

Acknowledgements

Author's Address

Patrick M. Pilarski
pilarski@ualberta.ca