Determining the Time until Muscle Fatigue using Temporally Extended Prediction Learning

Patrick M. Pilarski, Liping Qi, Martin Ferguson-Pell, and Simon Grange

***Abstract*—Anticipating the time remaining until muscle fatigue during the real-time use of an assistive rehabilitation device promises to improve the health and mobility of patients. In this work, we present a data-driven machine learning approach for assessing and reporting on situation-specific muscle fatigue in a point-of-care setting. The time until the onset of fatigue is predicted from interactions between a user and their wheelchair, as captured through myoelectric signals and pushrim data from a SmartWheel recording system. Our findings indicate that a real-time learning approach is able to accurately forecast the time remaining until a subject reaches fatigue-related activity thresholds. Our results also suggest the potential for generalizing these personalized anticipations between different patients. The approach presented in this work therefore promises to allow both a user and their device's control system to observe endurance-related future effects of current motor control choices; ongoing feedback of this kind may prove to be a valuable tool for improving patient mobility outside of the clinical environment.**

I. Introduction

Undetected muscle fatigue can cause injury—often irreversible—to the users of assistive rehabilitation technologies. The careful regulation of muscle contractions and fatigue is important for unifying mind, body, and machine in myoelectric prostheses, assistive robotics, and functional electrical stimulation systems; the onset of fatigue can impact the sustained efficacy of these devices and the health of the patients using them [1], [2]. The consequences of fatigue occurring during wheelchair propulsion are also of particular concern—propulsion is a demanding activity involving repetitive loading of the upper extremities through a precarious range of motion [3]. Great interest has therefore been placed in the development of systems that can predict when and how muscles will fatigue [4]–[7]. Assessing the level of muscle fatigue during a patient's use of assistive technologies promises to help avoid or reduce the risk of injury, inform how occupational therapists teach movement patterns that minimize the risk of long-term motion-related injuries, and provide valuable state information for device control systems (e.g., stimulation controllers [1]).

Despite the potential for improved care for users, the development of an autonomous system for predicting fatigue remains under-explored in the literature [4]. Open problems in this deployed setting include predicting the impact current patient movement patterns will have on the onset time of fatigue, estimating the non-linear time-endurance relationships of muscles in near real time, and anticipating the rate and timing at which different muscles will fatigue (as related to muscle imbalance in long-term device use). Portability (the ability to be affixed to a patient or their device) and persistence (the ability to maintain accuracy over time) are also critical for the long-term deployed use of devices that predict muscle fatigue in the users of assistive technologies.

Progress has been made toward the goal of a versatile deployed system for fatigue prediction, and several groups have demonstrated the detection of fatigue from myoelectric activity and/or kinematic recordings. Of note, Al-Mulla et al. have demonstrated a wearable device to classify fatigued, near-fatigued, and non-fatigued muscle states in near real time (e.g. during isometric weight lifting) [4]. Work by Frey-Law et al. and Gonzalez-Izal et al. further addresses the prediction of muscle endurance time during isometric activity [5], and the prediction of fatigue-related changes in muscle force generation [6]. A recent study by Qi et al. also showed that during wheelchair use certain recorded EMG variables are sensitive to fatigue, specifically EMG amplitude (increases with fatigue) and mean power frequency (decreases with fatigue) [7]. However, systems still require the ability to efficiently obtain continuous-valued endurance forecasts in non-isometric use domains. In addition, it remains important to demonstrate a portable system with the ability to update personalized fatigue predictions during the ongoing activity of individual patients in near real time, for example during daily use. Robustly adapting fatigue prediction systems to individual users and their changing environments and use patterns is an important unsolved challenge.

In this preliminary study, we demonstrate a computationally efficient, real-time machine learning approach to address these open problems. Specifically, we demonstrate a system for learning and predicting the time until muscle fatigue from a user's current activity and motion choices (Fig. 1). Our approach forms the basis for a system which could guide users' choices and optimize control system actions, maintaining an optimal fatigue state and avoiding unnecessary strain on muscles so as to prevent injury.

II. Methods

The pre-clinical study data supporting our proposed approach was derived from subjects using a wheelchair in the controlled environment of a laboratory with an ergomoeter to measure the performance of the user. Kinetic and myoelectric data (Fig. 2) from fourteen wheelchair study participants was

PMP is with the Department of Computing Science, University of Alberta, Edmonton, AB T6G 2E8, Canada. LQ, MFP, and SG are with the Faculty of Rehabilitation Medicine, University of Alberta, AB, Canada. The authors gratefully acknowledge support from the Alberta Innovates Centre for Machine Learning, the Alberta Paraplegic Foundation, and Alberta Innovates – Technology Futures.

Please direct correspondence to: pilarski@ualberta.ca

To appear: *Proc. of the 18th International Functional Electrical Stimulation Society Conference (IFESS)*, Donostia-San Sebastián, Spain, June 6–8, 2013.

collected by Qi *et al.* as described in detail in their work [7], and summarized in the text that follows.

14 able-bodied participants (7 males, 7 females, age: 30±4 years, weight: 65±12 Kg) volunteered to participate in the study. They were all informed about the procedure and singed consent in accordance with the procedures approved by the University of Alberta Ethics Committee. Participants were advised not to perform any exercise 48 h before measurements.

Surface electromyography (sEMG) activity of upper extremity muscles was recorded using parallel-bar EMG Sensors (DE-3.1 double differential sensor, 1mm in diameter and separated by 10 mm, Bagnoli, Delsys Inc., Boston, MA, USA). sEMG signals were detected on eight muscles: anterior deltoid (AD), middle deltoid (MD), and posterior deltoid (PD), pectoralis major (PM), upper trapezius (UT), biceps brachii (BB), and triceps brachii (TB) on the right shoulder after prior removal of the hair and cleaning with alcohol swipes. The EMG signals were sampled at 2000Hz.

A test wheelchair (Quickie GP, Sunrise Medical, Longmont, CO, USA) was aligned and secured over the rollers of an ergometer, which connected to a monitor placed in front of the participant to provide visual speed feedback. The forces applied to the pushrim of the wheel were measured using a SMARTWheel (Three Rivers Inc., LLC, Mesa, AZ, USA). The SMARTWheel was placed on the right side of the test wheelchair. Pushrim kinetic data were collected at 240 Hz. Kinetic and EMG recordings were synchronized.

Participants were given several minutes to get used to propelling the wheelchair and to establish a comfortable propulsion technique. During testing, participants performed multiple trials of wheelchair propulsion at two different speed settings; for the present work we look specifically at fast propulsion data, wherein the participant would continue pushing at a speed of 1.6m/s until they felt the activity hard to maintain. This speed was faster than the normal adult walking speed and normal wheelchair propulsion speeds (i.e., 1.0–1.3m/s) to present a challenging and strenuous situation for the participants. Each participant continued pushing at the speed set-point until their Ratings of Perceived Exertion (RPE) measure met or exceeded a score of 15, considered to mark the highest level of effort prior to exhaustion in this study [7].

A. Prediction Learning

The key insight developed in our work is that long-term, temporally abstracted forecasts of fatigue can be learned during the use of an assistive device (in our present case, from either pre-recorded or real-time wheelchair usage data) and used to provide feedback on potential injury or fatigue-related stresses. As shown in our prior work, predictive machine learning can occur in an incremental fashion, making it suitable for ongoing use and patient-by-patient adaptation [8].

In the present work, prediction learning was performed using the framework of general value functions (GVFs) [9], [10], as described for prosthetic signal anticipation in Pilarski *et al.* 2013 [8]; predictions were learned using computationally efficient temporal-difference learning algorithms. The learning target of interest was the time remaining until some terminal

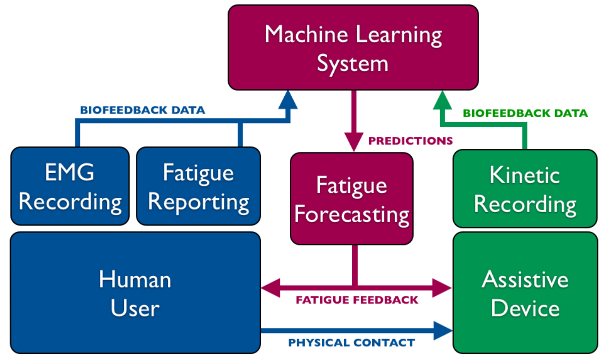

Fig. 1. A system to anticipate dangerous levels of muscle fatigue and provide user with temporally extended feedback about future exertion. Biofeedback signals are used to by a machine learning system to predict muscle fatigue; situation-dependent fatigue forecasting (feedback) can then be transmitted to the user and their device during real-time use and training to improve skill acquisition and help reduce the risk of injury during out-of-clinic use.

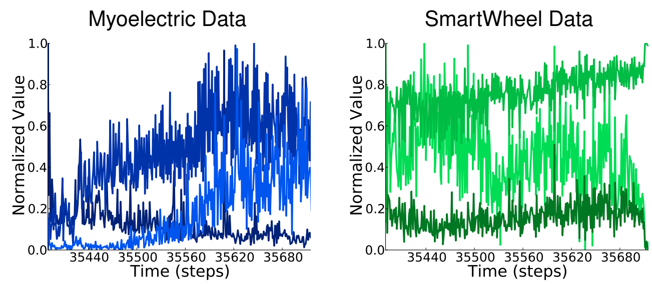

Fig. 2. Examples of recorded sEMG data (left) and SmartWheel data (right) during a single trial by one of the study's participants. As reported by Qi et al., visible changes occur as fatigue sets in [7]. Shown myoelectric data includes the mean absolute sEMG values for BB, TB, and PM muscle groups, while shown SmartWheel data includes force and force moment signals.

event, denoted *time-to-fatigue* (TTF), in this case indicated by a RPE report of 15 or more. As shown by the robot examples presented by Sutton *et al.* [9], time until a fixed event can be incrementally estimated by updating a set of learning weights $\mathbf{w}$ with a constant per-time-step value. For our work, the signal presented to the learning system at each step was Δt, where Δt was the length of a learning time step (in minutes). Prediction updates were subject to a continuation probability γ that was state conditional, so as to remain 1.0 during pre-fatigue motion and becoming 0.0 at the moment of the fatigue event of interest (RPE≥15). Replacing eligibility traces $\mathbf{e}$ were also severed when $\gamma = 0.0$. The full learning approach was implemented as follows:

initialize: $\mathbf{w}, \mathbf{e}, \mathbf{x}$
loop
    observe $\mathbf{s}$
    $\mathbf{x}' \leftarrow \mathsf{approx}(\mathbf{s})$
    if RPE ≥ 15 **then**
        $\gamma = 0.0$
        $\mathbf{e} = 0.0 * \mathbf{e}$
    else
        $\gamma = 1.0$
    $\delta \leftarrow \Delta t + \gamma \mathbf{w}^T \mathbf{x}' - \mathbf{w}^T \mathbf{x}$
    $\mathbf{e} \leftarrow \mathsf{min}(\lambda \mathbf{e} + \mathbf{x}, 1)$
    $\mathbf{w} \leftarrow \mathbf{w} + \alpha \delta \mathbf{e}$
    $\mathbf{x} \leftarrow \mathbf{x}'$

Tile-coding function approximation, denoted approx($\mathbf{s}$), was used to map the input signal space $\mathbf{s}$ into a binary state vector $\mathbf{x}$ suitable for learning, as described in prior work [8]; approx($\mathbf{s}$) was a 2-way encoding with 12x12 grids, 8 overlapping tilings, and single active bias unit. The system was initialized with $\alpha = 0.022, \lambda = 0.99$ and all weight vector elements set to 0. Recorded data was presented to the learning system in training batches, with testing occurring on independent testing data not previously visible to the learning system ($\alpha = 0$ during testing). Two cases were explored: 1) training on one example set for a given subject and testing on a second data set from the same subject, and 2) training on data from multiple subjects and testing on new data from one of the subjects. The input space of signals to the learning system included velocity, force, and momentum data from the SmartWheel system, and all seven sEMG signals. Input signals were normalized according to their observed maximum ranges during operation and averaged into 1s steps prior to presentation to the learning system. To assess the error of the system's anticipations, the true, *post hoc* time-to-fatigue values were compared to the learned predictions at each point during the training and testing runs. The learning time to process a single recorded time step was $\leq$1ms during offline training.

III. Results

The results in Figs. 3–5 show the ability to predict in advance the amount of time remaining until a user exceeds a self-reported fatigue threshold. In other words, the learning system was able to answer the following question: "Given how you are moving, how long until you exceed an acceptable level of muscle fatigue?" As shown by the similarity between predicted TTF and true computed TTF on both training and independent testing set data, the predicted TTF correctly approximated the true amount of time until self-reported fatigue (RPE $\geq$15). These initial results suggest that the learning system was able to form appropriate non-linear relationships between immediate features in the input data stream and the time until a future terminal event.

Performance differences were observed when learning systems were presented with different input signal subsets and different sources of training data. Figure 3ab presents the performance over the training and testing data following iterated offline learning when the learning system received a full set of 14 sEMG and SmartWheel data signals. Figure 3cd presents a similar comparison for a learning system trained only on SmartWheel data (i.e., no sEMG information was available to the machine learner). The learner without sEMG demonstrated more variability in its prediction accuracy after 100 learning iterations, as seen qualitatively in Fig. 3 and quantitatively via a comparison of post-learning median error values over training and testing data (Fig. 5). The results in Fig. 3 represent observed learning performance when training and testing data were both derived from the same subject. Figure 4 contrasts these single-subject results with results for a learning system trained on data from two subjects (100 learning iterations per subject, interleaved sequencing) and then tested on new data from each of the subjects.

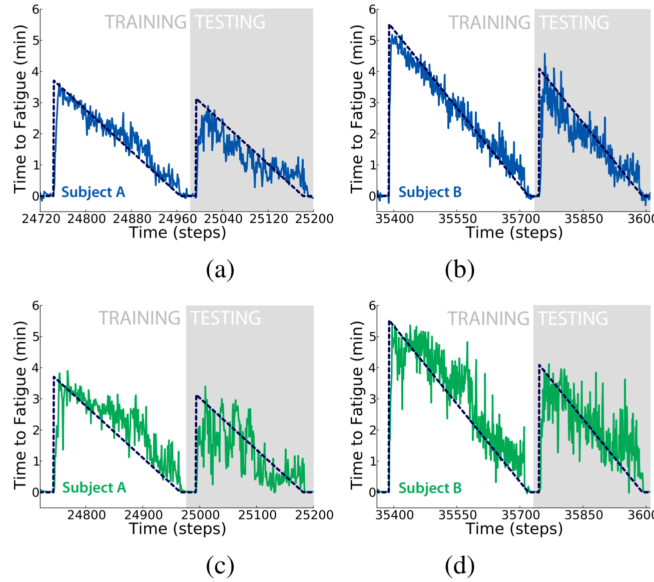

Fig. 3. Example Time-to-Fatigue (TTF) results on training and testing data for two different subjects (a/c and b/d respectively) when the learning system was trained using sEMG and SmartWheel signals (a and b) and only SmartWheel signals (c and d). Offline-learning involved 100 passes through the training data; the learner was then presented with a single non-learning pass through previously-unseen testing data from the same subject. Predicted TTF (solid line) and true TTF (dotted line) are shown for both training and testing data.

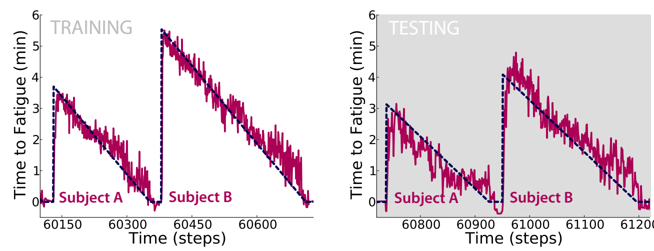

Fig. 4. Example Time-to-Fatigue (TTF) results on training and testing data when the learning system was trained using sEMG and SmartWheel signals from multiple subjects. Offline-learning involved 100 interleaved passes through the training data from each of the two subjects (200 total iterations); the learner was then presented with a single non-learning pass through new testing data from each subject. Predicted TTF (solid line) and true TTF (dotted line) are shown for both training and testing data.

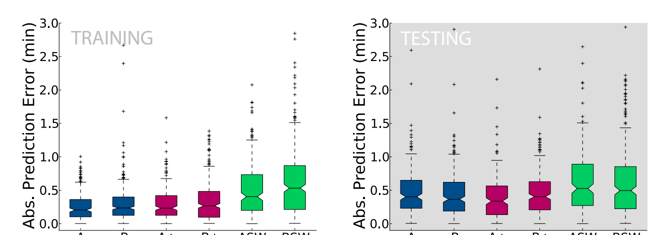

Fig. 5. Quartile analysis of absolute error values over a single post-learning pass through training and testing trial data; performance is shown by the median values (notch centres), their 95% confidence bounds (width of notches), and location of the first and third quartiles (box edges). Comparisons shown for two subjects when the learner used both sEMG and SmartWheel data from a single subject (A and B), sEMG and SmartWheel data from both subjects (A+ and B+) and single-subject SmartWheel data (ASW and BSW).

As shown in Figs. 4 and 5, results for single-subject and multi-subject learning on the training data were similar in both qualitative and quantitative terms, with slightly less variability on training data in the single-subject learning case. However, performance on testing data suggests that multi-subject training may allow better generalization to new data (Fig. 5). This potential for cross-subject transfer is promising, but requires further testing to determine statistical significance.

IV. Discussion

The machine learning approach presented above is data-driven and computationally inexpensive, making it suitable for real-time deployment during clinical or daily use. The ability to predict time-to-fatigue provides a basis for delivering real-time rehabilitation feedback to patients, potentially increasing the effectiveness of in-clinic training [11] and reducing the risk of injury during daily life. Specifically, our results suggest a viable way to give real-time predictive feedback directly to devices and their users regarding the effect their choices and motor policies will have on future fatigue. We expect the anticipatory methods presented in this work will extend well to other related biofeedback scenarios, including direct metabolic recordings and invasive (*in vivo*) biofeedback signals. One clear area of application is the TTF-based adaptation of multi-electrode stimulation patterns to maximally engage muscles during functional electrical stimulation (e.g., modulating the sequential activation of spatially distributed stimulators [2]).

The generality of the proposed methods also suggests that the current system will extend to predicting fatigue events based on non-RPE measures that can be extracted directly from time-space features in the input data. While this study used self-perceived effort as a measure of exhaustion and muscle fatigue, and would be suitable for ongoing use during daily wheelchair propulsion, it is also feasible to derive true TTF values directly and objectively from biofeedback data; a direct measure of fatigue is fully compatible with our proposed prediction learning approach. The benefits of one approach over the other is a useful area for future study, as is the observed potential to anticipate time until fatigue directly from kinetic (e.g., SmartWheel) recordings during daily use without the need for cumbersome myoelectric recording equipment.

Another important feature of the proposed approach is that it allows for the predictive relationships developed through offline (a priori) training of a machine learning system to be extended with new online data without extensive re-training. As new examples are presented to the learning system it is capable of incrementally updating its predictions; an incremental learning approach as suggested here is beneficial for ongoing improvement and adaptation during daily use, and also for rapidly adapting a pre-calibrated system (for instance, one trained offline on data from multiple subjects) for personalized use by a new individual.

The preliminary study presented in this work leaves a number of areas open for future investigation. We are currently engaged in a comprehensive evaluation of the accuracy of learned predictors for all 14 subjects, and at fast and slow wheelchair propulsion speeds. We are also cross-validating the ability of the system to adapt a predictor trained on multiple subjects to data from a new, previously unseen subject. Finally, to assess the potential impact of real-time predictive feedback on occupational training and chronic injury prevention, it will be important to deploy future extensions of our approach in continuing trials with human users and their assistive devices.

V. Conclusions

In this work we demonstrated real-time machine learning techniques with the potential to address the problem of fatigue prevention during the use of an assistive device, for example a wheelchair or an electrical stimulation system. Specifically, we presented a situation-sensitive method for estimating the time until muscle fatigue during ongoing wheelchair activity. Our proposed approach was able to accurately forecast the time to fatigue using instantaneous sEMG and/or SmartWheel signals, and could integrate training data from multiple subjects. This approach therefore forms the basis for a system to guide motor control choices during interactions between a human and their assistive machine, potentially reducing the risk of chronic injury and avoiding unnecessary strain on a patient's muscles and joints. We expect this approach will transfer well to other point-of-care rehabilitation domains where it is important to predict and deliver real-time fatigue information during the use of human-machine interfaces and assistive medical devices.

References

[1] J. d. R. Millán, R. Rupp, G. R. Müller-Putz, R. Murray-Smith, C. Giugliemma, M. Tangermann, C. Vidaurre, F. Cincotti, A. Kübler, R. Leeb, C. Neuper, K.-R. Müller, D. Mattia, "Combining brain-computer interfaces and assistive technologies: State-of-the-art and challenges," *Front. Neurosci.,* 4:161, 2010. doi:10.3389/fnins.2010.00161

[2] R. Nguyen, K. Masani, S. Micera, M. Morari, M. R. Popovic, "Spatially distributed sequential stimulation reduces fatigue in paralyzed triceps surae muscles: A case study," *Artif. Organs*, 35: 1174–1180, 2011.

[3] M. M. Rodgers, G. W. Gayle, S. F. Figoni, M. Kobayashi, J. Lieh, R. M. Glaser, "Biomechanics of wheelchair propulsion during fatigue," *Arch. Phys. Med. Rehabil.*, 75: 85–93, 1994.

[4] M. R. Al-Mulla, F. Sepulveda, M. Colley, "sEMG techniques to detect and predict localised muscle fatigue," in *EMG methods for evaluating muscle and nerve function*, Mark Schwartz (Ed.), InTech, 2012.

[5] L. A. Frey-Law, J. M. Looft, J. Heitsman, "A three-compartment muscle fatigue model accurately predicts joint-specific maximum endurance times for sustained isometric tasks," *J. Biomech.*, 45(10): 1803–1808, 2012.

[6] M. Gonzalez-Izal, D. Falla, M. Izquierdo, D. Farina, "Predicting force loss during dynamic fatiguing exercises from non-linear mapping of features of the surface electromyogram," *J. Neurosci. Methods*, 190(2): 271–278, 2010.

[7] L. Qi, J. Wakeling, S. Grange, M. Ferguson-Pell, "Changes in surface electromyography signals and kinetics associated with progression of fatigue at two speeds during wheelchair propulsion," *J. Rehabil. Res. Dev.*, 49(1): 23–34, 2012.

[8] P. M. Pilarski, M. R. Dawson, T. Degris, J. P. Carey, K. M. Chan, J. S. Hebert, R. S. Sutton, "Adaptive artificial limbs: A real-time approach to prediction and anticipation," *IEEE Robot. Autom. Mag.*, 20(1): 53–64, 2013.

[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. 10th Int. Conf. Autonomous Agents and Multiagent Systems (AAMAS)*, Taipei, Taiwan, 2011, pp. 761–768.

[10] J. Modayil, A. White, R. S. Sutton, "Multi-timescale nexting in a reinforcement learning robot," in *Proc. Int. Conf. Simulation of Adaptive Behaviour (SAB)*, Odense, Denmark, 2012, pp. 299–309.

[11] I. M. Rice, R. T. Pohlig, J. D. Gallagher, M. L. Boninger, "Handrim wheelchair propulsion training effect on overground propulsion using biomechanical real-time visual feedback," *Arch. Phys. Med. Rehabil.*, 94(2): 256–263, 2013.