To appear: 2019 IEEE-RAS-EMBS International Conference on Rehabilitation Robotics (ICORR), 24-28 June, 2019, Toronto, Canada, 8 pages.

Exploring the Impact of Machine-Learned Predictions on Feedback from an Artificial Limb

Adam S. R. Parker, Ann L. Edwards, and Patrick M. Pilarski*

***Abstract*— Learning to get by without an arm or hand can be very challenging, and existing prostheses do not yet fill the needs of individuals with amputations. One promising solution is to improve the feedback from the device to the user. Towards this end, we present a simple machine learning interface to supplement the control of a robotic limb with feedback to the user about what the limb will be experiencing in the near future. A real-time prediction learner was implemented to predict impact-related electrical load experienced by a robot limb; the learning system's predictions were then communicated to the device's user to aid in their interactions with a workspace. We tested this system with five able-bodied subjects. Each subject manipulated the robot arm while receiving different forms of vibrotactile feedback regarding the arm's contact with its workspace. Our trials showed that using machine-learned predictions as a basis for feedback led to a statistically significant improvement in task performance when compared to purely reactive feedback from the device. Our study therefore contributes initial evidence that prediction learning and machine intelligence can benefit not just control, but also feedback from an artificial limb. We expect that a greater level of acceptance and ownership can be achieved if the prosthesis itself takes an active role in transmitting learned knowledge about its state and its situation of use.**

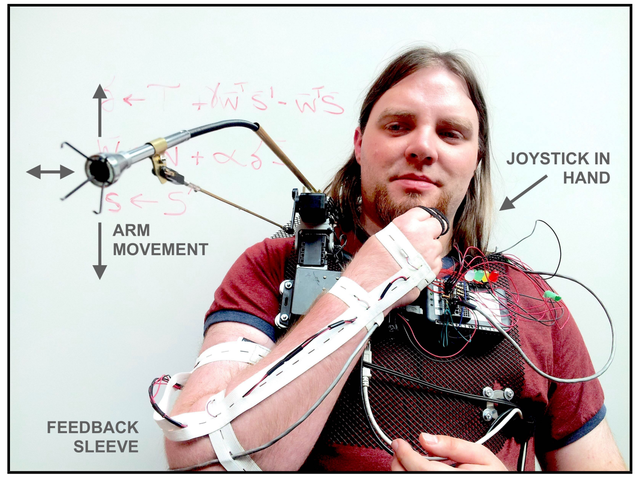

Fig. 1. Wearable robot limb system used in these experiments. The four degree-of-freedom arm is controlled by a joystick in the user's hand which sends signals to an ADC and then to a laptop, which in turn commands the servos. A vibrotactile feedback sleeve can provide feedback to the user.

I. INTRODUCTION

The loss of a limb, especially an upper limb, can have a significant impact on an individual. A person may be missing a limb from birth, or it could be the result of illness or injuries sustained over the course of one's life. Artificial limbs, also called prosthetic limbs, are often seen as a means of mitigating the absence of a biological limb. In all cases, but particularly when a limb is lost later in life, it can be very difficult to adapt to interacting with the world through a mechanical or electronic device [1]–[8]. There are many prostheses on the market that attempt to fill the needs of individuals with amputations, and many of these have tremendous potential to restore lost functionality and independence to the user; however, even the best prostheses currently available have limitations [1]–[3]. There are two major areas where current prostheses begin to show the strain of insufficient technology to properly support them. The first area is a lack of feedback [3], [4], [7], [8]—e.g., the sense of touch—and more important to this work, lack of proprioception when using a prosthesis [2], [7]. The second area is insufficient control [2]–[6]. Under most current techniques, the person who needs to control the limb has fewer control channels available to them than their device has functions [4]–[6]. This leads to some clever, but non-natural, control solutions such as routing some of the control channels to alternate locations on the user's body. A final challenge which results from the first two is acceptance of the prosthesis by the user [1]–[3]. Despite the great clinical potential of many modern prostheses, as a result of the first two limitations a prosthetic can be perceived by the user as insufficient or as a reminder of the functionality that they lost and that the device simply cannot restore [1]–[3]. Lack of acceptance is especially prominent in the newer myoelectric (EMG) prostheses, i.e., electrically driven robot limbs, versus the older mechanical types despite the increased potential that myoeletric prostheses have in overcoming the other challenges [2], [5].

Operating a device that interacts with the world is a learned motor function. As infants, we learn the way our limbs interact with our environment through general motion and play [9]–[11]. This develops the control channels and models required for us to use our bodies to sense and manipulate the world we live in [9]. This interaction involves two parts [9], [10]. The first is the internal forward copy of the action—in effect, knowledge that moving specific muscles will cause a motion which results in the desired sensory feedback. There is also a reverse copy that is processed at the same time. The reverse copy starts at the desired interaction

This research was undertaken, in part, thanks to funding from the Canada Research Chairs program, the Alberta Machine Intelligence Institute (Amii), Alberta Innovates, the Government of Alberta, the Natural Sciences and Engineering Research Council (NSERC), and the UofA Undergraduate Research Initiative.

* Please direct correspondence to: asparker@ualberta.ca; pilarski@ualberta.ca; Department of Medicine, University of Alberta, Edmonton, AB, Canada, T6G 2E1.

with the environment and links the required muscle action to it. In order to skillfully interact with the environment, both the forward and reverse models must be present [9], [10].

Artificial intelligence offers a promising solution to the control problems encountered by the users of electromechanical prostheses [12]. Offline machine learning in the form of pattern recognition is for the first time seeing use in commercial prostheses, and is considered to be the state-of-the-art in controlling multiple prosthetic joints [4], [5]. Real-time machine learning has also recently been used to ease the control burden on a user by learning joint activation sequences as a limb is being used [13], [14]; as one example, predictions about a user's control choices have been learned so as to minimize the number of switches between joints, and consequently the time required to perform a task [13].

A lack of feedback is frequently responsible for abandonment of prosthetic devices, especially upper-limb prostheses [21]. Feedback is an important aspect of control, and how to provide feedback from upper-limb prostheses to individuals with amputations is an active area of research [7] [20]. There are many modalities and means of feedback that are being explored currently. Some examples are substitution, where a signal that is not meant to imitate the lost physiological system is used, and modality matched, where an attempt is made to imitate the physiological sensations [20]. Providing feedback in these ways has been shown on subjects without amputation to improve performance on grasping tasks, as outlined by Schofield et al. (2014) [20]

The primary contribution of the present work is to suggest that machine intelligence can be used to enhance not just control—the focus of most prosthesis-related machine intelligence research to date—but also *feedback from a prosthesis*. This feedback was part of a user's intact biological system, and contained information used in operation of their natural limb. In the case of a prosthetic limb, motor awareness and forecasting are now at least partly encoded in the hardware of the prosthesis rather than in a user's biology. Therefore, we may need to provide assistance to the natural system in interfacing with its electronic components. We suggest that machine intelligence can be used to take the internal state of the assistive device and interpret it in ways the biological system cannot do naturally; the results of this interpretation can be communicated to the user in a variety of ways to improve their control over the device. Thus, using machine intelligence, we can help create a forward prediction of an action electrically and communicate it to the user, similar to the operation of the intact biological system.

This work therefore contributes a preliminary exploration of the application of machine-learned predictions, expanding upon the work started by Parker et al. [22]. A simple system for communicating machine-learned predictions is used to assist a user in refining their own forward model of motor actions while using a prosthetic limb analog. Specifically, temporal-difference learning is used to generate a prediction about the electrical load the servos of a human controlled robot arm will experience as they near a potentially dangerous collision with objects in the user's environment. This prediction is communicated to the user through a vibration motor. In this way, we emulate the forward predictive model present in a biological limb's motor function. We expect that, similar to the way that the biological operation of a limb is dependent on its forward copy, the addition of an electronic/computational equivalent during human-robot collaboration will yield control improvements over purely reactive feedback. In this study, the amount of load experienced by a servo over the course of an experimental run when the user receives this predictive feedback is compared to the same user receiving the same indication when the servo is actively experiencing high load (reactive feedback).

II. METHODS

A. Robot and Experimental Platform

The experimental platform used in this work was a custom-designed robotic arm called the ExArm (Fig.1), which was wearable by individuals without amputation. The arm was designed to model the gross motor functionality of joints in a human arm. It had four controllable actuators: shoulder, elbow, wrist flexion, and hand open/close (AX-12/18+ Dynamixel servo motors). Subjects used a 2-axis thumb joystick (SparkFun) to control the motion of the ExArm's joints, and pressing the joystick could change the active joint. The joystick was connected to an ADC (DI-149 data acquisition starter kit, DATAQ Instruments), which digitized the 3.3 V signal modified by the user's control of the joystick. The resulting output signal was sent via USB to a computer, which interpreted the signals and sent commands to the robot's servos. The control software only utilized information from a single axis of the joystick for motion, as well as the joystick button press to indicate a joint switch, to emulate EMG control of a prosthetic limb. The velocity of motion was fixed for all participants in all trials; speed of arm motion was a constant value.

AX-12/18+ servos used in the design of the ExArm provided several useful output signals, including their angular position, angular velocity, motor temperature, voltage, and load. To communicate feedback about these sensors to the user, we designed a custom sleeve embedded with four vibration motors (termed *tactors*) similar to those used in a cellphone or pager. With the sleeve donned, one tactor each was located over the user's shoulder, elbow, wrist, and hand, as shown in Figs. 1 and 2. The platform therefore emulated the capacity for actuation in many commmon prosthetic devices while adding vibrotactile feedback.

B. Experimental Procedure

Five subjects were asked to participate in experiments with the ExArm, and gave informed consent in accordance with the study's institutional review board approval. Each user wore the sleeve containing the vibration tactors and controlled the back-and-forth motion of the robotic arm's shoulder joint using the thumb joystick. The other joints of the arm as well as the joint switching functionality were not used for this experiment to restrict the motion of the arm to a single path. The ExArm was affixed to a stationary

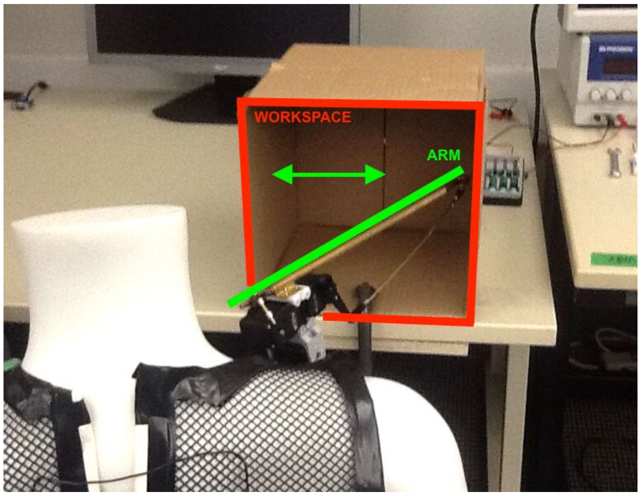

Fig. 2. The experimental setup: a confined workspace (red), the robotic arm (green), and as in Fig. 1, an experimental subject with attached vibrotactile feedback sleeve (seated to the left of the workspace, not shown).

mannequin as shown in Fig. 2 to ensure each experiment began with the robotic arm at a constant position and to mitigate the effect of a user's trunk movement. Thus, for this initial work, the position and movement of the user was unrelated to the outcome of the experiment. The workspace was a subspace of the shoulder joint's total range of motion, bounded by a 27 cm square box that was fastened in place. Prior to each experiment, the end effector was centered with respect to the workspace, perpendicular to the rear wall of the box and equidistant from the left and right walls. Each subject was asked to perform four separate five-minute tasks, structured as follows:

1) Training Task: The first task was designed to provide users with practice controlling the ExArm. For this training task, the user was asked to move the arm repetitively from one side of the box to the other using the joystick, pausing briefly ($\leq$1 second) upon reaching the center of the workspace. At the left and right walls of the box, the user was tasked with pushing the robotic arm against the wall until the arm was fully flexed, causing a temporary increase in the load reported by the servo motor. The user's shoulder vibration tactor was programmed to vibrate at a load threshold of 650 out of a maximum reportable load reading of 1024. This vibration communicated to the user that the load had exceeded the maximum threshold considered safe for the robotic arm, and the arm should be moved away from the wall. In addition to providing each user with practice manipulating the arm, this task produced the source data for prediction learning (described below).

2) No-Feedback Task: Each subject performed the second task without any knowledge of the position of the arm within the workspace other than its starting location. In order to establish a baseline with no visual, auditory, or tactile feedback, subjects were given a blindfold and listened to music through earphones throughout the task. The volume of the music was increased to a comfortable level at which they could not hear the arm tapping the walls of the box. During this task, vibratory feedback about load was also turned off. The instruction given to the user for this task and those that follow was to avoid excessive load on the servos by not colliding with the barriers too harshly while approaching the left and right walls closely in an alternating fashion.

3) Reactive-Feedback Task: The next task was identical to the no-feedback task; participants were blind-folded and sound isolated and asked to navigate from wall to wall without stressing the servos with collisions. For this trial the participant was provided with reactive vibration feedback when the current load experienced by the robot arm's shoulder servo reached a threshold of more than 420, determined experimentally. The maximum value of the load recorded during the trials was 827.87, which means the threshold was 50.7% of the maximum experienced. Thus, the tactor triggered every time the user hit a wall but not during travel in between. This task provided an indication of the effectiveness of having reactive tactile feedback only, and specifically examined how well the user could approach each wall without incurring a forceful impact when feedback was delivered at the moment the arm first contacted the wall.

4) Predictive-Feedback Task: For the final task given to participants, users were again blindfolded and sound-isolated, and given the same task as the previous two trials. In this case, they were provided with tactile feedback from predictions of the electrical load on the robot's arm servo motor. Predictions were provided by a real-time machine learning system trained while the participant was performing task 1. This prediction learning system is described in the following section. When the load prediction rose above 900, determined experimentally, the shoulder tactor was programmed to vibrate. The maximum prediction during the trial was 3857.5, which means the threshold was 23.3% of the maximum prediction value. This task was designed to determine how communicating the learned prediction of load changed the user's ability to approach the wall without incurring a forceful impact.

All load and prediction thresholds used were determined from the analysis of data prior to experiments. We determined the noise level of the load signal while traversing the workspace and set the thresholds so they would not trigger during travel. The prediction threshold and learning parameters were also set so as not to signal an impending high load event too early in travel.

C. Machine Intelligence and Prediction Learning

The main component of this study is an incremental prediction learner to generate expectations about future impact given learned knowledge about the user's previous motion choices, their outcomes, and the current state of the robot arm. To make predictions about the world, intelligent systems require sensory inputs. These inputs can then be divided into discrete states for increased or decreased resolution. The shoulder joint of the ExArm has a rotation range of 300°. In our protocol, we used the servo encoders value to determine the position of the shoulder joint as a sensory input, divided into 32 distinct states (termed *bins*). These states were motion-dependent; as such, each of the 32 states was further expanded into three: one set of 32 position bins used to represent the state when the servo is moving clockwise, a second set to represent the position while the servo is moving counter-clockwise, and a third set that represent the position when the servo is not moving. The immediate state of the arm was noted in a feature vector (denoted x, of length 96) as a single active bit indicating the current position and direction; this feature vector also contained a single active baseline unit. A weight vector of corresponding length, denoted w, was used to store the learned predictions about the interactions between the robot arm and the walls of the workspace.

The weight vector w was learned from data using standard techniques from temporal-difference learning and recent generalized value function methods, as outlined for the prosthetic setting in Pilarski et al. [12] and more generally in Modayil et al. [15]. Weights w were updated on each time step according to the temporal difference between the instantaneous load being reported by the servo (denoted τ) and predictions about the immediate and next load readings (the inner products $w_t^\top x_t$ and $\gamma w_t^\top x_{t+1}$, respectively, where γ is the timescale or level of temporal abstraction for the prediction of interest). The update to the weight vector on each timestep t was done according to:

$$w_{t+1} = w_t + \alpha(\tau_{t+1} + \gamma w_t^\top x_{t+1} - w_t^\top x_t)x_t,$$

where α represents a step-size (learning rate, set to $\alpha = 0.1$ in these experiments). The temporal abstraction for predicting the load signal of interest was set to $\gamma = 0.92$; this means the prediction learner was acquiring knowledge about the exponentially discounted expectation of the electrical load experienced by the robot's shoulder servo motor over the next $\sim$12 time steps, or 0.6 seconds; the system learned and operated within a control cycle of roughly 20 Hz (50 ms time steps). This knowledge could then be retrieved and used in predictive feedback by reporting the prediction as the inner product $w_t^\top x_t$. As noted above, in the predictive feedback task, vibratory feedback to the user was triggered when the prediction's value exceeded a fixed threshold, indicating an impending collision with the walls of the workspace.

Learning was only enabled during the training task, such that the system acquired and updated user-specific predictions about servo motor load while each subject was performing their first task. Learning weights were then frozen (i.e., $\alpha = 0$) during all remaining tasks, including the predictive feedback task. Learning could in principle continue during all tasks; however, for clear assessment of the principles of interest, our experimental protocol featured defined training and testing periods.

III. RESULTS

When compared to the case where purely reactive feedback was given to the user, giving learned predictive information as feedback was found to reduce the load experienced by the shoulder actuator of the robot limb. One way repeated measures ANOVA was used to analyze the difference in load when the feedback system was triggered differently. As

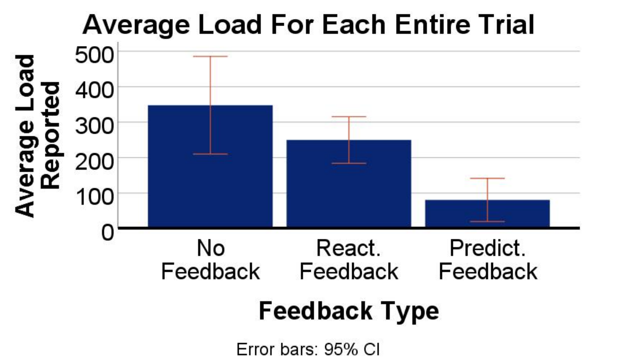

Fig. 3. Key finding: the use of predictive feedback reduced the load (a measure of impact intensity) experienced by the system during use. The load shown is averaged across all five participants and the entire duration of each trial.

shown in Fig. 3, the average load across all participants (N = 5) for each entire trial was significantly less with predictive feedback than with reactive feedback. For these comparisons, Maulchy's W indicated sphericity could be assumed (0.602). The uncorrected F statistic was F(2, 8) = 16.385, p = 0.001. The difference is specifically between the predictive feedback case and the other two (no feedback to predictive feedback p = 0.031, reactive feedback to predictive feedback p = 0.038). No significance was found between the no feedback and reactive feedback cases.

There was a notable increase in visits to more central positions when using predictive feedback. Figure 4(a–c) shows the frequency of visits to each position as seen by the system (bin). Results shown are the average for each trail across all five participants, where again a one way repeated measures ANOVA was used for statistical analysis. Maulchy's W indicated sphericity could not be assumed ($<$ 0.050). The Greenhouse-Geisser epsilon was $<$ 0.75 (0.081), so this correction was used. The corrected F statistic was F(2, 9) = 4.994, which is above the critical value of 4.26, and significance was detected (p = 0.030). The N = 5 potentially interferes with some bins being significant, as the variance is high. Despite this however, significance is noted in bins 15–19. Significance was not found on the extreme end bins, 13 and 22, due to the high variance. When using reactive feedback, subjects were observed to contact both walls of the workspace with approximately even frequency as can be seen in Fig. 4(b), with the robot arm deflecting noticeably on both sides due to the contact. When predictive feedback was provided to the user, the robot arm was also observed to approach the two sides of the workspace symmetrically, but with much less or no visible deflection to the arm upon contact. The figures show increased visits in central regions of the workspace under predictive feedback compared to the other feedback modes.

The edges of the workspace show less average load for the predictive feedback case than the other feedback cases examined here. The relationship between the feedback type, position, and load can be seen in Fig. 4(d–f). The load is shown as an average across all 5 participants, using the visits per bin to average the load in each bin individually. As with previous analyses, one way repeated measures ANOVA was used to determine any statistical significance. Again, Maulchy's W indicated sphericity could not be assumed (< 0.050), but the Greenhouse-Geisser epsilon was again < 0.75 (0.088) so was again used to correct the results. The Greenhouse-Geisser correction returns an F statistic of $F(2,10) = 6.805$ compared to a critical value of 4.10, with $p = 0.010$. Examined on a bin by bin basis, significance was found in bins 13, 20, and 21. Bin 22 was not found have statistically different results between the feedback types. The figures show greater load on the extreme ends of travel for the no feedback and reactive feedback cases. This is the region where collisions with the barrier of the workspace would occur. The load that is incurred in this region is not seen in the predictive feedback case. The visitation frequencies in Fig. 4(a–c) appear to coincide with the lower load experienced by the system in the end region.

IV. DISCUSSION

Feedback is an important aspect of skilled control. As noted above, we defined the control of the robotic device to be successful and skilled if the load experienced by the device while moving near the border of the work area is low—the task objective given to our subjects during testing was to closely approach but not impact the walls of the workspace. With different forms of feedback or different settings, we expected a subject might never get near the wall (overly sensitive predictions, thresholds, or too much temporal extension), that they might do so with high variability (as when operating with minimal feedback), or that they might impact the wall consistently but forcefully if feedback comes too late (e.g., with overly delayed or reactive feedback). Our observations support these expectations.

When the machine-learned predictions about collisions were used in providing feedback to the user, the user was able to reduce the overall strain on the system. Figure 3 demonstrates the effect that different types of feedback had on skilled control of the robotic device. In the no feedback case, the load experienced by the device was large and variable (the maximum load that can be reported by the servos is 1024). The variance improved in the reactive feedback case, but while the overall load decreased, it was not enough improvement for a study with N=5 to find statistical significance. No matter how sensitive the threshold is to initiate the reactive feedback, the user must still perceive the feedback and act in an appropriate, timely way; load is inevitable since it is already occurring. The strain on the system can potentially be reduced via fast human reaction time—subject-specific reaction time is one possible source of the variance in Fig. 3.

The similarity in the means of the no feedback and reactive feedback cases has interesting implications. It seems that if we only cared to limit the load experienced by the system during operation that there is little reason to use reactive feedback, which would be a typical first solution, over no feedback. This has major implications as to the importance of feedback in the operation of prosthetic devices. The significance of the predictive feedback case is a little surprising for the small sample size. Increasing the sample size may begin to differentiate the no feedback and reactive feedback cases, but that there is significance in the predictive feedback case with such a small sample size suggests it's power. *The simple machine learning agent is capable of learning something that when communicated to the human user improved their performance according to at least one outcome measure.*

The source of the overall reduced load can be seen in how participants moved the arm differently using the different feedback sources. A more detailed indication of the motion of the device is illustrated in Fig. 4(a–c). In particular, the figures highlight differences in the feedback modes in the area of travel between the borders of the workspace (between bins 14 and 21). Specifically in bins 15 to 18 the bin-by-bin cumulative visits are shown to be higher for predictive feedback than the same measures using reactive or no feedback. Despite this, all three modes have similar grand means, as would be expected for the constant fixed velocity motion the participants used. The higher frequency of visits in the central region is an indication of successful operation, as it shows that the device spent a greater portion of the moving time during the trial transitioning from border to border rather than impacting the walls, under which condition the servo can move a small amount while the arm flexes. This observation is further supported by the differences that can be seen between the types of feedback at and beyond bins 13 and 21—the borders of the workspace. With predictive feedback, the user moved into the borders of the workspace less frequently (Fig. 4c)). The impact this had on the load can be seen in the bin by bin comparisons of Fig. 4(d–f). As a result of the system visiting the border cases less frequently the system experienced significantly less (or no) load in bins 13 and 21–22 (Fig. 4(c–f)), indicating less time spent under impact conditions or flexing of the physical device. The similarities in average load reported by the servo for the central bins, bins 15 to 18 is expected as these bins are where the arm would be moving steadily with no perturbations to cause changes in the load experienced. This also aligns with the area more frequently visited with predictive feedback. There is some discontinuity in the load and position values for bins 20-22 when compared to bins 13 and 14. This may be the result of the discretization of the encoder positions not being related to the physical workspace—the difference in load reported on the left and right sides may be because the physical workspace barrier fell in between two bins, causing the load from contact with the wall to divide between two locations as perceived by the system.

The lack of notable difference between the no feedback and reactive feedback cases is also seen in Fig. 4. Despite having no awareness of the robot arm during the no feedback trail, participants still navigated the full space between the boundaries of the workspace, although with much greater variance. For this low sample size, the way participants moved under the no feedback and reactive feedback cases can

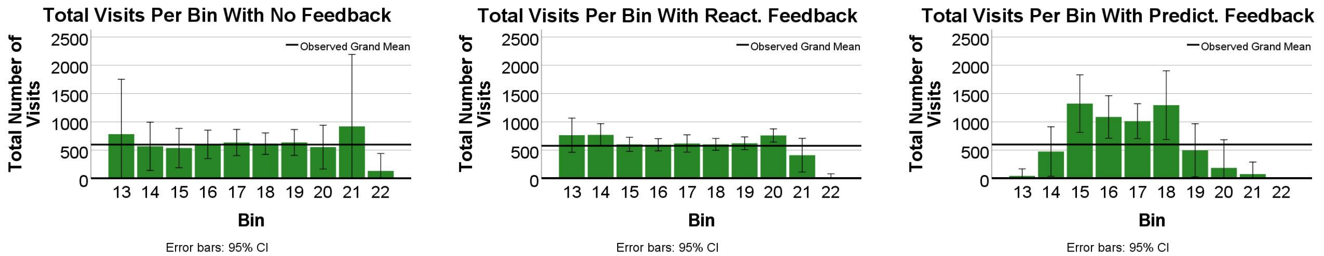

(a) (b) (c)

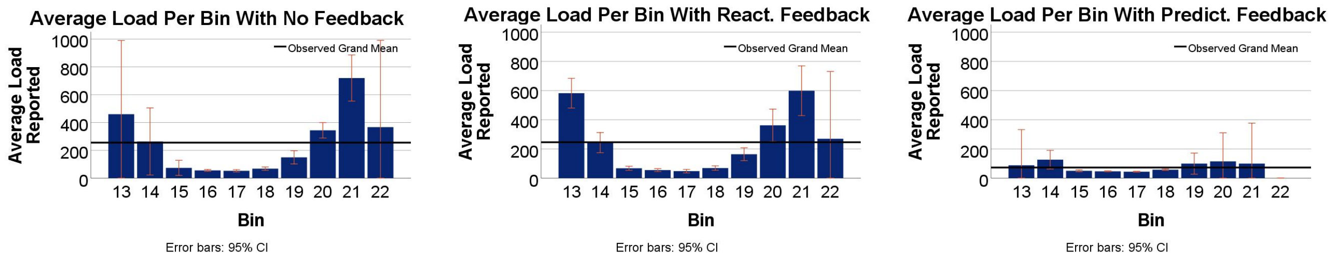

(d) (e) (f)

Fig. 4. Aggregate results for all five subjects showing (a–c) the frequency of visiting any given servo motor positional bin and (d–f) the average load in each bin as reported by the servo, using the frequency of visits to average.

be seen to be similar. It took the addition of machine-learned predictions to minimize the load *and* improve participants ability to travel inside the workspace. This reinforces the importance good feedback solutions for device users; machine-learned predictions, we suggest, are one such solution.

A. Tuning, Training, and Adaptability

With predictive feedback and the settings described above, we qualitatively observed subjects stopping the robot arm's motion such that it made only light, unloaded contact with the wall. This level of contact could be modified by varying the learning parameters of the artificial intelligence, and parameters could be adjusted in a number of ways to achieve a number of outcomes. There is no one "correct" setting for sensitivity; instead, there are a number of possibilities for how the device can assist the user in achieving their objectives. Learning parameters could be tuned to provide feedback behaviour that duplicates that of the reactive feedback case. The converse is not true—reactive feedback is not capable of providing preemptive feedback about future events. Also, when using predictive feedback, we observed that the threshold for indicating an impending load could be made less sensitive than the equivalent reactive load without incurring false positives. The reactive threshold was set to the lowest value that would not be triggered by the normal noise of motion. As the learned predictions are mathematical expectations conditioned on servo position, they are not affected by spurious load variance or noise due to motion, as would be a purely reactive approach, and allowed us to set a much lower threshold for triggering feedback in the predictive case.

For clarity of assessment, the artificial intelligence system in this preliminary study only acquired and updated its predictive knowledge during a defined training period. In any machine learning setting with a fixed training period, variability in training can noticeably affect learning system performance, but should not affect fixed or reactive approaches. Differences to the training of the learning system or a slight shift in the experimental setup may have resulted in an earlier feedback prompt to this user in terms of absolute servo-motor position on one side of the workspace. Omissions during training or changes to the domain of use may be corrected or updated through the use of continuing or ongoing machine learning. This has been suggested in prior work [12], and is a natural way to robustly extend the present study. As learning is already done in a per-time-step, incremental way during training, there are no technical or algorithmic barriers to continuing the learning of feedback-related predictions during operational use. Specifically, off policy algorithms would allow the system to learn in real time during the trials. The issue with the current learning approach is that when the system is successful in avoiding collisions with the workspace, it "forgets" the workspace is there. Off-policy learning offers a solution to this issue, and will be tested in a subsequent study. While many offline or batch prediction learning methods could potentially be used to generate expectations for use in feedback (e.g., the work of Pulliam et al. [16]), the continuing and computationally inexpensive nature of our chosen learning approach makes it well suited for use in a prosthetic environment [12]. Our prediction learning approach is suitable for subject-specific, task-specific learning with no requirement for a priori domain knowledge; it is also well suited for adapting to ongoing changes in a task or a user's behaviour during persistent, real-time use.

B. Feedback Modalities

As noted above, much work is being done to restore missing feedback to prosthesis users [7], [8], [20]. Focus has been placed on restoring touch, including sensations such as pressure, texture, temperature, and even pain. A large body of this research has explored feedback using sensory substitution, wherein one sensation is replaced with another different sensation that the user must be trained to skillfully interpret; use of this approach is largely due to the physiological constraints of prosthetic human-machine interaction [7]. Modality-matched feedback is also receiving growing attention; in matched feedback, sensations are restored either invasively or non-invasively to the natural or proxy locations that convey sensations of the lost or damaged biological system as closely as possible [7], [8], [20].

Our present study can be thought of as a form of substitution feedback—predictions about the electrical current drawn by the device during operation (perhaps thought of as the device's "pain" or motor fatigue) are communicated to the user via a vibratory buzzing sensation in order to prompt the user to take action to prevent it. This buzzing is not a natural sensation, and it is not communicated at an equivalent natural location on the user's body. What separates this choice from the usual form of sensory substitution is that fact that the information being transmitted from the user to the device is not a biological sensation—it is specific to the internal hardware of the device and encodes a prediction about future changes to that hardware. While communicating these anticipations is helpful to the successful operation of the device, it is not a natural thing for the user to feel; as with most substitution feedback, it takes training to interpret such a sensation (as noted in Hebert et al. [8]). This training need was perhaps minimized for our participants because of the precedent in modern society to interpret the vibration of personal device as a prompt to act (e.g., cellphone vibration in response to a new text message).

However, our work should not be thought of solely in terms of sensory substitution. Our study is intended to be a small window into a larger area for research: the use of machine intelligence as a method for filtering, selecting, and *communicating* salient information about the internal state of a complex device. This communication can be thought of as a form of *transparency*, as used by Thomaz and Breazeal [17]. Communication of such non-biological knowledge to the device's user—e.g., prompts regarding a device's internal state, decisions, and anticipatory knowledge—promises to streamline human-machine interaction in many domains, and should be equally suited to feedback via both sensory substitution and modality-matched percepts.

C. Future Work

The results presented in this work are preliminary, and there is much room for further study in this area. The incremental learning algorithm used in this experiment was effective but monolithic. If a control-learning system were used in conjunction with the present prediction-learning algorithm, it may be possible for a device to adapt the timing and magnitude of its feedback to better suit its domain of use. For instance, the feedback threshold or level of temporal abstraction γ could be tuned on the basis of reward-like signals of approval or disapproval delivered by the user, using techniques from related work on the human training of machine learners [17]–[19]; predictive load information could be communicated at distances from the collision which have been learned to be appropriate for a specific user and their task preferences. Exploration could also be done into how effective the predictive feedback is when it is learned in real time while the user is doing the task, rather than freezing learning during the trials. Further, as artificial intelligence use in artificial limbs becomes more prevalent, finding ways of communicating the actions that the system has learned to the user, rather than solely a predefined environmental signal, may help allow more control to pass to the prosthetic—the case of shared control and sliding-scale autonomy. Transparent communication between the operator and their device could be the keystone which allows an intelligent prosthetic and a human user to co-operate, combine processing power, and more effectively restore lost function.

V. CONCLUSIONS

Feedback is important to prosthetic limb control. While machine intelligence has been used to improve the interpretation of control signals given to a limb from the user, its use in modulating feedback is often overlooked. This article contributed a look at the potential value of predictions and machine learning in feedback to close the loop between a human and their artificial limb. To our knowledge, this is the first study investigating the use of real-time prediction learning in the feedback path of a human controlled robotic limb, and suggested the potential value of continuing this line of exploration.

When compared to strictly communicating momentary electrical load to the user, communicating a machine-learned forecast of the same load was found to decrease the load experienced by a robotic limb as a result of impacts with a workspace, and to increase the ability of our subjects to navigate the limb despite the absence of all other feedback. The increase in precision in terms of both position and load for the predictive feedback case over the no feedback case was dramatic, especially given the low subject pool. Additionally, the improvement in load minimization over purely reactive feedback was significant. Though preliminary, these results promise two related outcomes for the user of a prosthetic limb. First, we expect that increased communication from the device about its internal state and setting of use may allow the user more personalized and more trustworthy options for control. Over the long term, predictive feedback could therefore lead to greater acceptance and assimilation of the device as part of the user. Further, by creating a computational predictive forward copy of an action and communicating it to the user, operating an assistive device may become more precise. These expectations remain to be verified during the use of predictive feedback in real-life functional tasks.

ACKNOWLEDGEMENTS

The authors thank Joseph Modayil, and Richard Sutton for insights and suggestions relating to this work.

REFERENCES

[1] L. Resnik, M. R. Meucci, S. Lieberman-Klinger, C. Fantini, D. L. Kelty, R. Disla, and N. Sasson, "Advanced upper limb prosthetic devices: implications for upper limb prosthetic rehabilitation," *Arch. Phys. Med. Rehabil.*, vol. 93, no. 4, pp. 710–717, 2012.

[2] T. W. Williams, "Guest Editorial: Progress on stabilizing and controlling powered upper-limb prostheses," *J. Rehab. Res. Dev.*, vol. 48, no. 6, pp. ix–xix, 2011.

[3] B. Peerdeman, D. Boere, H. Witteveen, R. Huis 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," *J. Rehab. Res. Dev.*, vol. 48, no. 6, pp. 719–738, 2011.

[4] S. Micera, J. Carpaneto, and S. Raspopovic, "Control of hand prostheses using peripheral information," *IEEE Rev. Biomed. Eng.*, vol. 3, pp. 48–68, 2010.

[5] 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.*, vol. 48, no. 6, pp. 643–660, 2011

[6] P. Parker, K. B. Englehart, and B. Hudgins, "Myoelectric signal processing for control of powered limb prostheses," *J. Electromyogr. Kines.*, vol. 16, no. 6, pp. 541–548, 2006.

[7] C. Antfolk, M. D'Alonzo, B. Rosén, G. Lundborg, F. Sebelius, and C. Cipriani, "Sensory feedback in upper limb prosthetics," *Expert Review of Medical Devices*, vol. 10, no. 1, pp. 45–54, 2013.

[8] J. S. Hebert, K. Elzinga, K. M. Chan, J. Olson, and M. Morhart, "Updates in targeted sensory reinnervation for upper limb amputation," *Curr. Surg. Rep.*, vol. 2, no. 3, art. 45, pp. 1–9, 2014.

[9] D. M. Wolpert, Z. Ghahramani, and J. R. Flanagan, "Perspectives and problems in motor learning," *Trends Cogn. Sci.*, vol. 5, no. 11, pp. 487–494, 2001.

[10] J. R. Flanagan, P. Vetter, R. S. Johansson, and D. M. Wolpert, "Prediction precedes control in motor learning," *Current Biology*, vol. 13, no. 2, pp. 146–150, 2003.

[11] J. Zacks, C. Kurby, M. Eisenberg, and N. Haroutunian, "Prediction error associated with the perceptual segmentation of naturalistic events," *J. Cogn. Neurosci.*, vol. 23, no. 12, pp. 4057–4066, 2011.

[12] 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 Robot. Autom. Mag.*, vol. 20, no. 1, pp. 53–64, 2013.

[13] 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 Biomechatronics (BioRob)*, Roma, Italy, 2012, pp. 296–302.

[14] P. M. Pilarski, T. B. Dick, and R. S. Sutton, "Real-time prediction learning for the simultaneous actuation of multiple prosthetic joints," *Proc. IEEE Int. Conf. on Rehabilitation Robotics (ICORR)*, Seattle, USA, June 24–26, 2013, pp. 1–8.

[15] J. Modayil, A. White, and R. S. Sutton, "Multi-timescale nexting in a reinforcement learning robot," *Adaptive Behavior*, vol. 22 no. 2, pp. 146–160, 2014.

[16] C. Pulliam, J. Lambrecht, and R. F. Kirsch, "Electromyogram-based neural network control of transhumeral prostheses," *J. Rehabil. Res. Dev.*, vol. 48, no. 6, pp. 739–754, 2011.

[17] A. L. Thomaz and C. Breazeal, "Teachable robots: Understanding human teaching behavior to build more effective robot learners," *Artif. Intell.*, vol. 172, no. 6, pp. 716–737, 2008.

[18] 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. IEEE Int. Conf. on Rehabilitation Robotics (ICORR)*, Zurich, Switzerland, 2011, pp. 134–140.

[19] W. B. Knox and P. Stone, "Learning non-myopically from human-generated reward," in *Proc. Int. Conf. on Intelligent User Interfaces (IUI)*, March 2013.

[20] J. S. Schofield, K. R. Evans, J. P. Carey, and J. S. Hebert, "Applications of sensory feedback in motorized upper extremity prosthesis: a review," in *Expert review of medical devices*, vol. 11, np. 5, pp. 499–511, 2014.

[21] E. A. Biddiss and T. T. Chao, "Upper limb prosthesis use and abandonment: A survey of the last 25 years," in *Prosthetics and Orthotics Int.*, vol. 3, no. 3, pp. 236–257, 2007.

[22] A. S. R. Parker, A. L. Edwards, and P. M. Pilarski, "Using learned predictions as feedback to improve control and communication with an artificial limb: Preliminary findings," in *arXiv:1408.1913 [cs.AI]*, 2014.