Online Learning for Proactive Obstacle Avoidance with Powered Transfemoral Prostheses

Max Gordon[1], Nitish Thatte[2], and Hartmut Geyer[2]

Abstract— **Avoiding obstacles poses a significant challenge for amputees using mechanically-passive transfemoral prosthetic limbs due to their lack of direct knee control. In contrast, powered prostheses can potentially improve obstacle avoidance via their ability to add energy to the system. In past work, researchers have proposed stumble recovery systems for powered prosthetic limbs that provide assistance in the event of a trip. However, these systems only aid recovery after an obstacle has disrupted the user's gait and do not proactively help the amputee avoid obstacles. To address this problem, we designed an adaptive system that learns online to use kinematic data from the prosthetic limb to detect the user's obstacle avoidance intent in early swing. When the system detects an obstacle, it alters the planned swing trajectory to help avoid trips. Additionally, the system uses a regression model to predict the required knee flexion angle for the trip response. We validated the system by comparing obstacle avoidance success rates with and without the obstacle avoidance system. For a non-amputee subject wearing the prosthesis through an adapter, the trip avoidance system improved the obstacle negotiation success rate from 37% to 89%, while an amputee subject improved his success rate from 35% to 71% when compared to utilizing minimum jerk trajectories for the knee and ankle joints.**

I. INTRODUCTION

Avoiding obstacles on the ground is a necessity for maintaining safety while performing a variety of locomotion tasks. This behavior requires anticipation of an obstacle and active leg control strategies to avoid it [1]. Transfemoral amputees, however, have a compromised ability to negotiate obstacles, as shown in Figure 1, as current prosthesis technology relies on mechanically passive knees that necessitate significant compensation at the hip in order to replicate able-bodied trip recovery strategies [2]. Compromised ability to avoid and recover from trips may contribute to the large number of falls that leg amputees suffer. For instance, 58% of unilateral amputees reported a fall within a year [3]. Moreover, the fear of falling can cause amputees to avoid activity, leading to further deterioration of their physical condition [4].

An increasing availability of powered prostheses in research labs provides the opportunity to study active obstacle avoidance strategies in prosthetics, although so far only a limited number of studies exist on this topic. These studies focus on detecting and classifying the correct response strategy after the amputee has tripped. For example, Lawson et al. [5] developed a classifier that uses fast Fourier transform and the root mean square of accelerometer data as features to classify stumbles and recovery strategies, respectively. Zhang et al. [6] found that adding EMG signals from the residual limb to accelerometer data can help reduce false positives for stumble and strategy detection. Finally, Shirota et al. [7] identified the optimal sliding window lengths and increments for feature calculation for trip detection and strategy selection classifiers. While detecting and classifying trip recovery strategies after their occurrence is a necessary step towards obstacle avoidance, it does not provide a proactive prosthesis control strategy that prevents obstacle encounters in the first place.

Another major drawback of the previous studies is that they train and test the classifiers offline. However, a deployed trip classifier needs to function online and deal with temporal adaptation of the learner and amputee. The adaptation is required as the obstacle avoidance behavior triggered by a trip classifier alters the amputee's movements and, therefore, the data used to train the classifier. Consequently, trip classifiers trained offline may be ineffective due to the mismatch of training and testing data, a common problem faced in imitation and reinforcement learning [8].

Here we present the first pilot study that combines online learning and proactive control of a powered transfemoral prosthesis to implement obstacle avoidance in amputee locomotion. The obstacle avoidance system uses early-swing measurements of the residual limb angle, angular velocity, and linear acceleration to recognize in-process obstacle avoidance attempts. To address the online learning aspect of this system, we adapted a previously proposed algorithm for detecting gait modes [9]. We also changed the existing swing leg behavior of the prosthesis to facilitate obstacle avoidance. This change includes a regression to predict the appropriate degree of knee and ankle flexion given the user's previous obstacle response motions. Finally, we evaluated the system behavior in trials with both non-amputee and amputee subjects.

II. METHODS

A. Forward-Backward Classifier

In order to learn to classify trips online with minimal hand-labeling of data, we rely primarily on the forward-backwards classifier approach first proposed by Spanias et al. [9] for the purpose of classifying different modes of gait such as level ground walking, standing, and stair climbing. In their work, a *forward classifier* predicts the next step's gait mode using data in a window shortly before the transition. In parallel, a

*This work was supported by the National Science Foundation under Grant No. 1734559.

[1]Max Gordon mjgordo3@ncsu.edu is with the Department of Electrical and Computer Engineering, North Carolina State University, 890 Oval Dr, Raleigh, NC 27606, USA

[2]Nitish Thatte nitisht@cs.cmu.edu and Hartmut Geyer hgeyer@cs.cmu.edu are with the Robotics Institute, Carnegie Mellon University, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh PA 15213, USA

Fig. 1. a) Utilizing minimum jerk trajectories during swing does not allow for appropriate adaptation of swing trajectories to enable obstacle avoidance. b) Our adaptive system learns online to detect the presence of an obstacle from the amputee's late stance/early swing movements. Once detected, the controller modifies the trajectories of the knee and ankle to achieve improved obstacle clearance.

backward classifier labels completed steps with their correct gait mode in hindsight. Because the backwards classifier has access to features from the completed step, it can achieve accurate labels with a small amount of hand-annotated data. Once trained, the backwards classifier can provide labels for training the forward classifier, obviating the need for further hand-labeling of steps.

In our work, the forward classifier predicts, shortly after toe-off, if the upcoming swing will require obstacle avoidance or not. For this purpose, we use a linear support vector machine and features of the residual limb motion in the last 210 ms of stance and first 90 ms of the swing phase. Because user behavior changes over time in response to changes in prosthesis obstacle avoidance behavior, we retrain the forward support vector machine every ten steps using labels from the backward classifier.

The backward classifier is another linear support vector machine, trained once for each user, which uses features extracted from the entire swing phase to label a step as an avoidance attempt after the fact. To train the backwards classifier we hand label obstacle avoidance attempts and normal steps for roughly ten obstacles. Figure 2 provides an overview of this system.

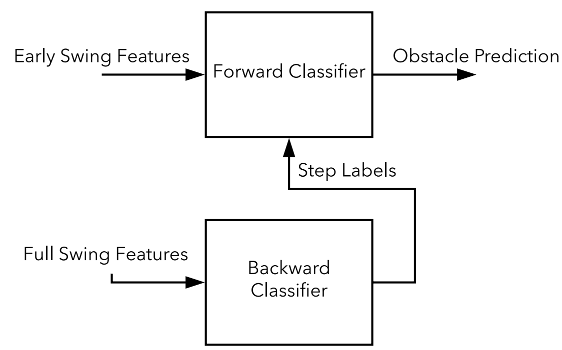

Fig. 2. Forward-Backward Classifier Overview. The backwards classifier uses features from the entire swing to provide training class labels to a forwards classifier. The forwards classifier uses features from late stance and early swing in order to predict if an upcoming swing will be an obstacle avoidance attempt.

B. Target Knee Angle Regression

A prosthesis user will not always encounter obstacles of the same height. As the obstacle avoidance response can be disruptive to the user, it is desirable to give the user control over the magnitude of the prosthesis response. We seek to achieve this functionality by using the normalized backward classifier score as a metric for the difficulty of avoiding an obstacle. We then implement a simple linear feedback law that assigns higher target flexion knee angles to obstacle avoidance attempts that are more difficult according to this metric. Figure 3 outlines this feedback mechanism, which has the form

$$\theta_{n+1}^{tgt} = \theta_n^{tgt} - k_{\text{decay}}(\theta_n^{tgt} - \theta_{\min}) + k_{\text{score}}\hat{\xi}, \tag{1}$$

$$\hat{\xi} = \frac{\xi - \xi_{10^{\text{th}}\ \text{percentile}}}{\xi_{90^{\text{th}}\ \text{percentile}} - \xi_{10^{\text{th}}\ \text{percentile}}}, \tag{2}$$

where θ_{tgt} is the current target angle for a given set of features, n is the current time step, k_{decay} is a gain that prevents continual target angle growth by decaying target angles towards $\theta_{\min}$, and k_{score} is a gain on the normalized class score, $\hat{\xi}$. The system shifts class scores, ξ, so that scores below the 10^{th} percentile of tripped step scores result in a reduction of the target knee angle. Furthermore, the system normalizes the scores by $\xi_{90^{\text{th}}\ \text{percentile}} - \xi_{10^{\text{th}}\ \text{percentile}}$ so that

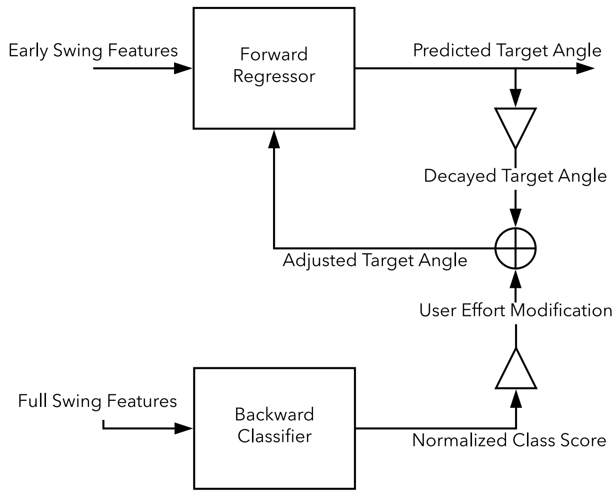

Fig. 3. Knee Angle Regression Feedback. In order to enable volitional control of the knee and ankle flexion angles to allow users to achieve greater flexion angles for larger obstacles, we implement a feedback system that uses the backwards classifier class score to quantify obstacle difficulty. After each step, the system increases the desired target angle for that step's forward features proportionally to the normalized back classifier score. We also decay the current desired target angle for those features to prevent continual growth of the target angle. The regression is retrained every ten avoidance attempts.

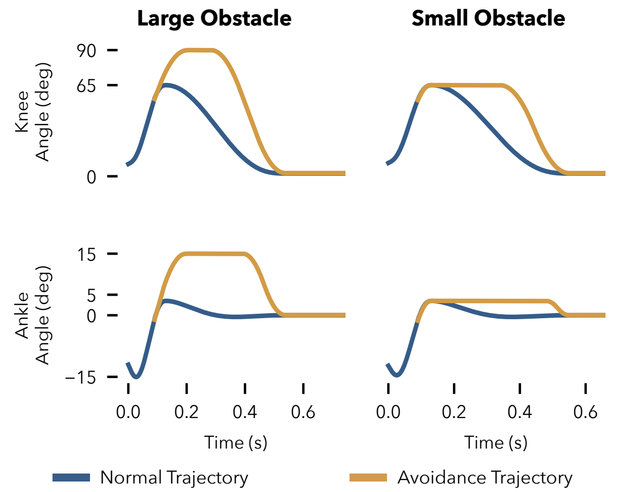

Fig. 4. Bang-bang obstacle avoidance trajectories (yellow) vs normal minimum jerk trajectories (blue) for the knee and ankle.

the gain k_{score} has a predictable effect across subjects whose score ranges vary.

The system fits the target knee angles with a linear support vector regression. Every time the trip avoidance triggers, it appends an additional target angle, specified by eq. (1), to a training data set. The system retrains the regression using this data set every ten trip-avoidance steps.

C. Feature Extraction

For the forwards and backwards classifiers, as well as the target knee angle regression, we use features of the thigh angle, angular velocity, and linear accelerations in a time window. Specifically, we use the mean, standard deviation, minimum value, and maximum value of each signal. For forward classification and regression the time window begins 210 ms before toe-off and ends 90 ms after toe-off, while for the backward classification we use a window consisting of the entire swing phase between toe-off and heel strike.

D. Trajectory Planning

To generate the knee and ankle motions for unperturbed swing, we follow the method proposed by Lenzi et al. [10] to generate and follow human-like minimum jerk trajectories that start at the toe-off state of each joint (angle, angular velocity, and angular acceleration), go to a target flexion state, and then extend to desired final angles at the estimated heel strike time. We estimate the swing period to be 65% of the stance period.

When the forward classifier triggers an obstacle avoidance attempt, we switch to bang-bang trajectories for the knee and ankle joints. These trajectories maximize foot clearance while respecting joint angle, velocity, and acceleration limits. The bang-bang trajectories achieve desired flexion angles as quickly as possible and then extend as late as possible such that they achieve extension before the predicted heel strike time. The trajectory planner uses the target knee angle regression to determine the appropriate peak angle for the knee trajectory, while the ankle trajectory's target flexion angle is a linear function of the knee's target angle. The knee trajectory's peak flexion angle is constrained to lie within 65 and 90 degrees while the peak ankle flexion is constrained within 5 and 15 degrees. Examples showing the minimum jerk swing trajectories and obstacle avoidance trajectories planned for large and small obstacles are given in fig. 4.

III. Experiments and Results

A. Experimental Protocol

We tested the ability of the proposed online learning system to accurately classify trips and normal swings, help subjects avoid tripping on obstacles, and modulate knee and ankle flexion appropriately for obstacles of different heights. To evaluate these aspects of system performance, we conducted experiments with a powered knee and ankle prosthesis shown in fig. 5 (previously described in Thatte et al. [11]).

Two subjects, one non-amputee with prior experience using this prosthesis, and one inexperienced amputee subject, performed walking trials with the obstacle avoidance system enabled. As subjects walked, an experimenter placed objects on the treadmill belt in front of each subject's prosthetic leg, necessitating an obstacle avoidance reaction. To obtain a baseline performance level for non-reactive prosthetic swing control, we also performed obstacle avoidance trials with the minimum jerk swing trajectories designed for undisturbed swing. Before the online trials, the backwards classifier was trained for the prosthesis user with 75 steps. The able bodied subject completed 446 total steps, with 53 box avoidance steps, while the amputee completed 222 total steps, with 40 box avoidance steps. The amputee subject performed trials in an ABBA order, where A is minimum jerk control and B is the reactive control, in order to average out potential

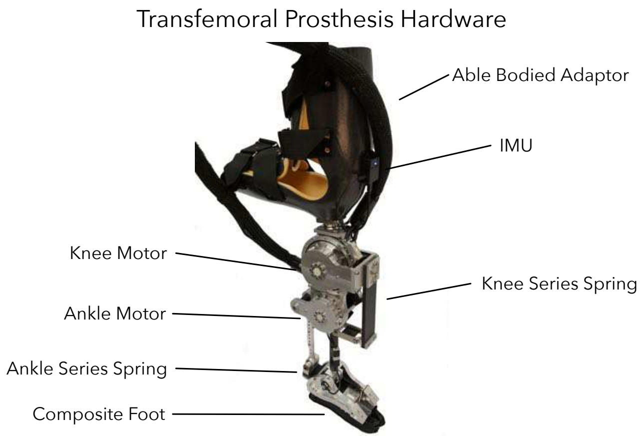

Fig. 5. Our powered transfemoral prosthesis prototype features series elastic actuators at both the knee and ankle joints for accurate torque control. We mount an IMU (3-Space Sensor, Yost Labs) to the thigh in order to measure hip angle and angular velocity and thigh linear accelerations. Able-bodied subjects wear the prosthesis via an L-shaped adapter (shown), whereas amputee subjects can attach the prosthesis to their personal socket via a standard pyramid adapter.

learning effects. The amputee subject also had an additional practice session the day prior to the box avoidance trials in which he acclimated to walking with the powered prosthesis without obstacles.

B. Results

Tables I and II show the overall classification accuracies, sensitivities, and specificities for the forward and backwards classifiers for the able-bodied and amputee subjects respectively. The forward and backwards classifiers for both subjects achieve high specificity (the number of normal steps classified correctly) and accuracy (> 95%). The sensitivity, the percentage of true trips classified correctly, of the classifiers for both subjects is substantially lower than the specificity or accuracy. For the forward classifier, we see that because the model is trained online, the sensitivity improves from the first half of the trial to the second half, which explains some of the low overall sensitivity.

TABLE I
CLASSIFIER PERFORMANCE, ABLE-BODIED[1]
TOTAL STEPS: 446, AVOIDANCE ATTEMPTS: 53

Controller	Classification Accuracy	Sensitivity	Specificity
Forward, 1st Half	96% ⋆	73% ⋆⋆⋆	99%
Forward, 2nd Half	99%	93%	99%
Forward Overall	98%	85% ⋆⋆⋆	99%
Backward	99%	100%	99%

TABLE II
CLASSIFIER PERFORMANCE, AMPUTEE[1]
TOTAL STEPS: 222, AVOIDANCE ATTEMPTS: 40

Controller	Classification Accuracy	Sensitivity	Specificity
Forward, 1st Half	95%	80%	98%
Forward, 2nd Half	96%	85%	98%
Forward Overall	95%	83% ⋆	98%
Backward	98%	90%	99%

Importantly, the ability of the forward classifier to correctly trigger the bang-bang obstacle avoidance trajectories improves obstacle avoidance success rates as shown in table III. Both subjects were able to avoid significantly more obstacles with the obstacle avoidance controller than with the minimum jerk trajectory controller.

TABLE III
OBSTACLE AVOIDANCE SUCCESS RATES[1]

Controller	Able-Bodied Success Rate	Amputee Success Rate
Minimum Jerk	37% ⋆⋆⋆	35% ⋆⋆⋆
Adaptive Bang-Bang	89%	71%

We also compared our online learning approach for obstacle avoidance to an offline approach similar to that taken by Lawson et al. [5], Zhang et al. [6], and Shirota et al. [2]. To do this, we trained a classifier offline using the first half of the amputee subject's bang-bang control data and tested it on the second half of the data. Table IV shows that the classifier trained offline has trouble generalizing to the second half of the data, as it performs significantly worse than the online-trained model in terms of accuracy and sensitivity.

TABLE IV
ONLINE AND OFFLINE FORWARD CLASSIFIER PERFORMANCE, AMPUTEE[1]

Classifier	Classification Accuracy	Sensitivity	Specificity
Offline	89% ⋆	39% ⋆⋆⋆	100%
Online	95%	83%	98%

Finally, we examined the ability of the knee angle regression to choose a target knee angle that is appropriate for the object size. The feedback law proposed in eq. (1) assumes we can use the backwards classifier score as a metric of obstacle difficulty. For the able-bodied subject, this assumption seems warranted, as there is a strong relationship between the obstacle height and the classifier score (fig. 6a, $R^2 = 0.50$). However, for the amputee subject, who was less experienced with walking with the powered prosthesis, this relationship is less clear (fig. 6B, $R^2 = 0.22$).

As shown in fig. 6c&d, our system is able to ensure that high classification score steps, associated with high user effort, obtain larger target flexion angles. This relationship led to noisy volitional control of the knee flexion angle for the able-bodied subject (fig. 6e) as evidenced by the linear relationship between knee angle and obstacle height ($R^2 = 0.31$). However, for the amputee subject, there is no clear relationship between the obstacle height and knee flexion angle (fig. 6f, $R^2 = 0.10$).

IV. DISCUSSION

We developed an online learning system to help users of powered transfemoral prostheses avoid obstacles. Our system uses information from an inertial measurement unit during the late stance to early swing period to classify the upcoming swing as either normal or a trip avoidance

[1]$\star \implies p < 0.05$, $\star\star \implies p < 0.01$, $\star\star\star \implies p < 0.001$, Chi-squared test

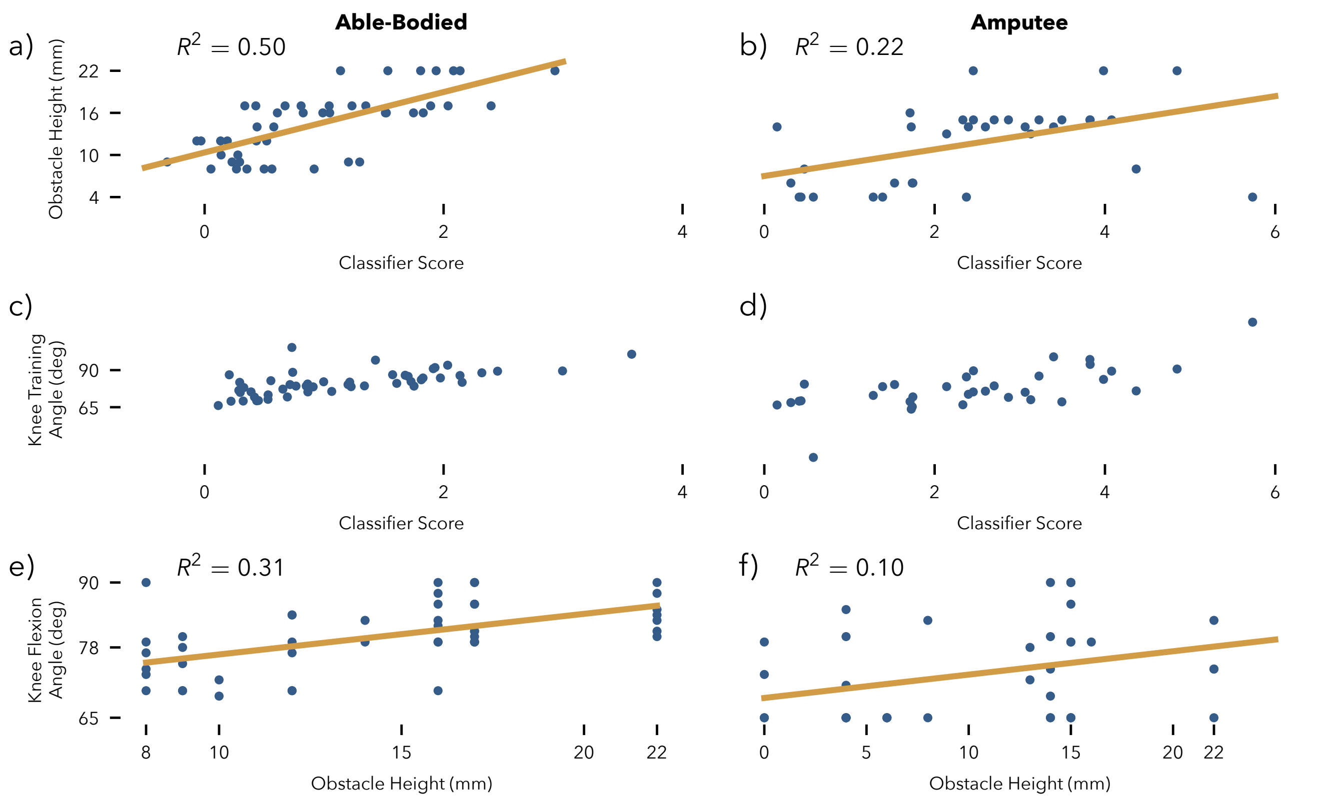

Fig. 6. Obstacle height vs backwards classifier score for (a) the able-bodied and (b) amputee subjects. The system uses the backwards classifier score as a metric for obstacle avoidance difficulty. This score is used in a feedback loop that forms the training set for the flexion target angle regression (c-d). With this feedback system, the able bodied user (e) is able to achieve a degree of volitional control over flexion angle as evidenced by the linear relationship between knee flexion angle and obstacle height ($R^2 = 0.31$). However, the amputee (f) was not able to achieve meaningful control over the flexion of the prosthesis ($R^2 = 0.10$), possibly due to the decreased experience level of this subject.

attempt. Unlike previous work on obstacle negotiation for transfemoral prostheses [5–7], our system learns online on an actual transfemoral prostheses. We compared the classification performance of our online system with a hypothetical offline system using online trials to provide testing and training data for offline analysis. This comparison showed that the online learning system provided an improvement in sensitivity and accuracy to obstacle avoidance attempts. Both an experienced, able-bodied subject and an inexperienced, amputee subject were able to improve their obstacle avoidance success rates significantly. However, only the experienced, able-bodied subject was able to achieve some level of volitional control of the prosthesis flexion as a function of obstacle height.

There are several reasons why the amputee subject may not have been able to achieve volitional control of prosthesis flexion. First, the amputee had far less experience using the prosthesis than the able-bodied subject. Consequently, even though both subjects were informed that trying harder to lift the leg over bigger obstacles would likely lead to greater flexion once the prosthesis learns, it is likely that only the first subject was able to incorporate and implement this information. The amputee may have concentrated on more rudimentary aspects of gait, as evidenced by his use of the handrails to walk, whereas the able-bodied subject did not need to use the hand rails. Moreover, the amputee's socket may have provided less control over the prosthesis than did the intact subject's able-bodied adapter (shown in fig. 5). Finally, we noted that the relationship between joint flexion and obstacle height tended to oscillate over the course of our trials. This may imply that the gains we used for the target knee angle regression (eq. (1)) were too high.

Before settling on the specifics of the obstacle avoidance system presented here, we also tested other options for its components. For example, we also evaluated incorporating EMG signals from the non-prosthetic limb in our obstacle avoidance classifier. Previous research showing that able-bodied subjects utilize stance leg musculature to help raise the hip during obstacle avoidance motivated this choice of EMG placement [1]. However, as was found by Spanias et al. [9], using EMG data along with mechanical data in the forwards-backwards online learning algorithm did not seem to improve classification accuracy, which is already high. This lack of improvement may also result from a significant delay in our wireless EMG sensors (Delsys Trigno). It is possible that a low-latency wired EMG sensor would be able to improve classification performance or the performance of the target angle regression.

We also tried using imitation learning techniques to model able-bodied strategies for stepping over obstacles. Specifically, we employed maximum margin inverse optimal control [12] to learn, offline, cost functions for the knee that explained obstacle avoidance trajectories. However, when used online, the generated trajectories tended to diverge and produce unexpected results because the initial toe-off state of the prosthesis did not match those in the able-bodied data set. For the obstacle avoidance classifier, we correct this sort of offline-online mismatch via the backwards classifier that provides labels to train the forwards classifier online. It is less clear how to update trajectories in hindsight as we never see the obstacle. For this reason, we used bang-bang trajectories during obstacle avoidance, which maximize the time the joints remain flexed.

In the future, we plan to overcome this issue by incorporating a laser distance sensor into the prosthesis. This sensor should allow precise measurement of the ground and obstacle shape during the initial part of swing as the hip moves forward. We plan to then use this information to explicitly plan knee and ankle trajectories that will avoid the obstacle and the floor until the appropriate touch down time.

There are several other limitations of the current study we should address in future work as well. First, we only tested the algorithm with two subjects. More subjects of varying skill levels are necessary to determine how applicable the system is to a broader population. Additionally, a likely reason why the forward classifier's sensitivity was relatively low, was that there were many more normal steps than obstacle avoidance attempts in the training data set. This may cause the SVM loss function's minimum to focus more heavily on classifying normal steps correctly. Deploying this system on a commercial prosthesis, for which trips are more rare, would exacerbate this issue. Therefore, future development should investigate how to train a classifier given heavily unbalanced class frequencies.

ACKNOWLEDGMENT

The authors would like to thank Professor Goeran Fiedler of the University of Pittsburgh for subject recruitment and help running the experiment. We would also like to thank Timothy Kyung and Nandagopal Srinivasan for their help conducting experiments as well.

REFERENCES

[1] A. E. Patla and S. D. Prentice, "The role of active forces and intersegmental dynamics in the control of limb trajectory over obstacles during locomotion in humans," *Experimental Brain Research*, vol. 106, no. 3, pp. 499–504, 1995.

[2] C. Shirota, A. M. Simon, and T. A. Kuiken, "Transfemoral amputee recovery strategies following trips to their sound and prosthesis sides throughout swing phase," *Journal of neuroengineering and rehabilitation*, vol. 12, no. 1, p. 79, 2015.

[3] J. Kulkarni, S. Wright, C. Toole, J. Morris, and R. Hirons, "Falls in patients with lower limb amputations: prevalence and contributing factors," *Physiotherapy*, vol. 2, no. 82, pp. 130–136, 1996.

[4] W. C. Miller, M. Speechley, and B. Deathe, "The prevalence and risk factors of falling and fear of falling among lower extremity amputees," *Archives of physical medicine and rehabilitation*, vol. 82, no. 8, pp. 1031–1037, 2001.

[5] B. E. Lawson, H. A. Varol, F. Sup, and M. Goldfarb, "Stumble detection and classification for an intelligent transfemoral prosthesis," in *Engineering in medicine and biology society (EMBC), 2010 annual international conference of the IEEE*. IEEE, 2010, pp. 511–514.

[6] F. Zhang, S. E. D'Andrea, M. J. Nunnery, S. M. Kay, and H. Huang, "Towards design of a stumble detection system for artificial legs," *IEEE transactions on neural systems and rehabilitation engineering*, vol. 19, no. 5, pp. 567–577, 2011.

[7] C. Shirota, A. M. Simon, and T. A. Kuiken, "Recovery strategy identification throughout swing phase using kinematic data from the tripped leg," in *Engineering in Medicine and Biology Society (EMBC), 2014 36th Annual International Conference of the IEEE*. IEEE, 2014, pp. 6199–6202.

[8] S. Ross, G. Gordon, and D. Bagnell, "A reduction of imitation learning and structured prediction to no-regret online learning," in *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 2011, pp. 627–635.

[9] J. Spanias, A. Simon, S. Finucane, E. Perreault, and L. Hargrove, "Online adaptive neural control of a robotic lower limb prosthesis," *Journal of neural engineering*, vol. 15, no. 1, p. 016015, 2018.

[10] T. Lenzi, L. Hargrove, and J. Sensinger, "Speed-adaptation mechanism: Robotic prostheses can actively regulate joint torque," *IEEE Robotics & Automation Magazine*, vol. 21, no. 4, pp. 94–107, 2014.

[11] N. Thatte, H. Duan, and H. Geyer, "A method for online optimization of lower limb assistive devices with high dimensional parameter spaces," in *Robotics and Automation (ICRA), 2018 IEEE International Conference on*. IEEE, 2018.

[12] N. D. Ratliff, J. A. Bagnell, and M. A. Zinkevich, "(online) subgradient methods for structured prediction," *Robotics Institute*, p. 55, 2007.