2020 8th IEEE International Conference on Biomedical
Robotics and Biomechatronics (BioRob)
New York, USA. Nov 29 - Dec 1, 2020

Coordinated Movement for Prosthesis Reference Trajectory Generation: Temporal Factors and Attention

Vijeth Rai [1], Abhishek Sharma [2], Pornthep Preechayasomboon [2], and Eric Rombokas [12]

Abstract— **Data-driven gait prediction can provide a reference trajectory for a wide variety of simple and complex movements captured in the training data. Coordinated Movement (CM) is a data-driven approach that maps movements of the body to movements of target joints, such as the ankle and knee. We have previously shown that the performance of CM for complex activities can be improved by adding more training data. In this paper we demonstrate that performance can also be improved by 1) including a history of the target joint angles as inputs to the model and 2) dynamic reallocation of the importance of the inputs over time using a neural network technique called Attention. These modifications are applicable when additional training data is limited. We also observe that Attention can follow important events in gait over time, adding interpretability to the system.**

I. INTRODUCTION

Powered prostheses are constantly improving, becoming lighter, more rugged, and more capable with each iteration [1, 2], but controlling the devices remains a significant challenge [3, 4]. Most powered prosthetic lower limbs have modes of operation corresponding to terrain types, such as flat-ground walking or stair ascent. Data from sensors measuring motion or forces, or body signals such as those from electromyography (EMG) can be used to estimate user intent. Threshold-based rules were used in earlier implementations to trigger transitions from one mode to another [5]. More recently, machine learning has become an integral part of this approach and has demonstrated improved performance and more robust detection of the current mode and transitions [6, 7]. These advancements in control have enabled people with limb-loss to experience increased ease in mobility through various everyday environments.

However, powered lower limb users still experience challenges performing difficult non-periodic movements [8]. Sports, getting in and out of cars or restaurant booths, obstacle avoidance, and uneven terrain are examples of challenging activities of daily living that do not cleanly correspond to commonly-used control modes.

In [9, 10] we showed that a data-driven method inspired by the continuous and coordinated movement of the human body can address a range of challenging movements. Our Coordinated Movement (CM) controller is a neural network model, which uses the history of movements of the rest of the body to estimate the kinematics of a target joint omitted from the input set of joints. For example, in the case of a trans-tibial prosthetic application this target joint would be the ankle joint. The body motion of the person with the amputation would serve as the inputs and the predictions generated by the network could be used for the control of the prosthesis.

Our CM approach produces high-quality reference trajectories that resemble the movements chosen by able-bodied individuals in the training set for activities such as ambulation of flat ground and stairs [10], and more complex movements such as weaving around cones, side stepping and backwards walking [9].

Previously we observed that the type of activity and the complexity of the movements involved can impact performance. For example, certain sections of the obstacle crossing activity, which contained punctuated and infrequent movements, were not well anticipated [9]. One way to improve performance is to increase the training set size to include more variation of movements [10]. In this study, we present two alternative modifications of our CM network: 1) augmentation of the input data to the network and 2) modification of the network architecture.

A. Data augmentation: Including the history of target joints

In the previous implementations of CM, estimates for target joints (e.g. right ankle angles) were generated using current and previous movements of the *intact joints* as inputs, but not the previous movements of the target joints. In other words, the output of the network was a function of body coordination, but was not strongly constrained to be continuous over time.

We hypothesized that augmenting the inputs with the history of target joint angles provides context and would encourage continuity in the outputs. This approach has been widely used in the field of multivariate time series forecasting to predict temporal patterns of a time-series with dependencies on multiple variables [11, 12].

Such models predict the current value of a time series using past values of the same series as well as current and past values of an external (exogenous) driving series. For our application, the target joint (ankle) trajectory is the time series of interest and the trajectories of the other intact joints of the body are the exogenous driving time series.

Time-series data vary greatly based on the the underlying phenomenon they represent. Seasonal processes such as the weather patterns exhibit more stationarity compared to economic forecasting with less distinct trends. Consequently, models are often chosen specific to the domain. Several linear, non-linear [13], and hybrid [14] models have been used

[1] Department of Electrical Engineering, University of Washington Seattle,WA, 98102 USA e-mail: (rombokas@uw.edu).
[2] Department of Mechanical Engineering, University of Washington Seattle,WA, 98102 USA .

978-1-7281-5907-2/20/$31.00 ©2020 IEEE

— each with their own benefits and trade-offs. Linear models such as Auto-Regressive Moving Average (ARMA) are fast and interpret-able but often fail to capture complex, non-stationary, and nonlinear characteristics of the time-series. Non-linear models, such as Non-linear Auto Regressive with Exogenous (NARX) or artificial neural networks (ANN), are more complex, time consuming, and come at a higher computational cost, but often perform better for domains with complex dependencies [13]. We have previously shown that this applies to human movement data as well, with the non-linear recurrent neural network model significantly improving performance compared to linear regression model [10]. Here, we compare the performance of two recurrent neural networks against a simple linear ARIMAX regression model.

B. Attention models for temporal dynamics of gait

For time-series prediction, we expect performance to benefit with longer time-histories comprising more temporal information. Recurrent neural networks are particularly suited for this purpose[15, 16]. We observed this to be true [10] but the improvements plateaued after a gait time-history of about 160 ms. However, studies on human gait have shown that changes in movement related to execution of a complex task occur within 1-2 preceding steps [17, 18], suggesting that a time history of joint angles at least 500-1000 ms long should contain useful information for predicting current target joint angles.

Our previous results, showing that prediction performance saturated for shorter time histories than 500ms, suggests that the models we chose were unable to take advantage of the cues present during the long time histories. The immense data present in the longer histories could be beyond the capability of the relatively simple RNN structures we used to learn the relevant movement cues. A common approach in machine learning to overcome this sort of problem is to increase the depth, or the number of layers of the network [19]. However, this adds to the complexity, training time, required amount of training data, and further reduces the interpretability of the model. This difficulty in interpreting how the decisions are made, and what inputs contribute to the outputs, is a common critique of neural network and deep learning methods.

To increase the efficiency in learning long term dependencies and to introduce better interpretability of the outputs, we borrow the Attention technique [20] from the field of Natural Language Processing (NLP). Attention models are the state-of-the-art in speech-to-speech translation research [21]. They have demonstrated improved efficiency in learning long term dependencies in long input sentences by computing the correlation of targets with other parts of the sentence and selectively focusing on the most useful parts.

We hypothesize that this ability to dynamically place unequal "attention weights" on different parts of the input series will allow the system to emphasize the informative events in the time history and improve performance. Another advantage of Attention models is that the attention weights corresponding to different points in time allow for visualization of the relative importance of these inputs in generating target values.

II. METHODS

A. Participants

Ambulation data was collected for 10 healthy participants (5 males, median age of 25) with no amputation or other mobility impairments. The experiment was completed in a single session which lasted less than 2 hours. Recruitment and human subject protocols were performed in accordance with the University of Washington Institutional Review Board approval and each subject provided informed consent. De-identified data can be made available, via a data use agreement, upon request to the authors.

B. Experiments

The subject's anthropometric details were recorded and 17 wearable Xsens Awinda (Xsens Technologies B.V.) sensors were placed on locations according to [22], followed by a calibration procedure. The subjects then performed 10-15 minute trials of traversing through an obstacle course at a self selected speed, as described in [9]. The experiment was completed in a single session which lasted less than 2 hours. In total, 180 minutes of data were collected from all subjects combined.

Obstacle Course: The obstacle course was designed to require the participants to perform a large number of heading changes, discrete events of stepping over obstacles, and adjustments of foot placement. Participants were not instructed on the path they should take, but simply to "traverse the course to the end and come back, weaving around the cones on your way out, and stepping over the obstacles on your way back." The cones were arranged in a single line separated by 1.2 meters, and 3-cone "figure eights" [23] were placed at the ends. The other side of the course consisted of three rectangular obstacles of varying height. The dimensions of the obstacles were (L $\times$ W $\times$ H) 0.7 m $\times$ 1.7 m $\times$ 0.2 m, 0.6 m $\times$ 0.6 m $\times$ 0.1 m, and 0.5 m $\times$ 0.6 m $\times$ 0.5 m, and is also pictured in [9].

C. Instrumentation

We collected locomotion data using the Xsens Awinda suit [22], consisting of 17 body-worn sensors placed at key locations. Each sensor has a tri-axial gyroscope, accelerometer, magnetometer, and barometer. Xsens MVN Analyze software integrates these individual sensors and renders a full-body avatar. After a system specified calibration, the software provides position and joint kinematics in a 3D environment. Although other data such as limb-segment position, orientation, acceleration are available, we used only joint angles for this study. All angles are in Euler representation of the joint angle vector (x, y, z) in degrees, calculated using the Euler sequence ZXY following the International Society of Biomechanics standard joint angle coordinate system [24]. Data, sampled at 60 Hz, from a total of 22 joints in 3 anatomical planes (sagittal, frontal, transverse) were captured

Fig. 1. Xsens Motion Wearable Motion Capture system consisting of 17 IMU sensors placed at key locations on the body

for each trial, which results in 66 total possible features for our machine learning methods.

D. Data Processing

Sensor data was visually inspected to see any aberrant errors. During data collection, some sensors might get displaced from their original calibrated location. If this was detected during the experiment, sensor placement was corrected followed by recalibration and reinitialization of the suit.

The Xsens software features a real-time engine that processes raw sensor data for each frame, fits the human body model to estimate anthropomorphic joint and segment data. A post processing engine includes information from the past, present, and future of each timestep to get an optimal estimate of the position and orientation of each segment. This High-Definition or 'HD' processing increases the joint data quality by extracting more information from larger time windows and modeling for skin artifacts. Given that most human movement data information is contained within 15Hz [25], we sub-sampled the data to 30 Hz, which reduces the training time without loss of performance.

E. Machine Learning Model and Architecture

In [9] we show that a variant of Recurrent Neural Networks, the Long Short-Term Memory (LSTM), using a short history of gait movements provided the best prediction of the right ankle joint. In this study, we analyzed performance on our dataset using the same network as well as a slightly modified version to also include the history of target joints as inputs. We also analyse the performance of the attention based networks with and without target ankle joint history.

Given a time series trajectory of M intact joints $x \in \Re^{M \times T-1}$, and a history of target joint angles $y_{hist} \in \Re^{T-1}$, we employ the following network models to estimate current target joint values $\hat{y_T}$ at time instant T.

- LSTM : Single layer LSTM network that maps intact joint trajectories to current target joint angles.

$$\hat{y_T} = f(x) \tag{1}$$

where f is the LSTM network.

- LSTM+H : Single layer LSTM network with the target joint angle history included in the input features. The most straightforward way to implement this is to simply concatenate the target time series history with the intact joint history.

$$\hat{y_T} = f([x : y_{hist}]) \tag{2}$$

where $[x : y_{hist}] \in \Re^{M+1 \times T-1}$ is a concatenated matrix of intact joints and target joint history.

- ATTN: Attention network [26] consists of a single layer LSTM encoder followed by a single layer LSTM decoder. Using only intact joints trajectory as input, this model predicts current ankle value at time T as:

$$\hat{y_T} = f(x) \tag{3}$$

where f is the Attention based encoder-decoder network.

- ATTN+H: Attention network that includes the target ankle joint history as one of the inputs along with the intact joint history. However, the implementation relies on concatenation with a context vector that is computed based on temporal attention weights. For more details, see [26]

$$y_T = f(x, y_{hist}) \tag{4}$$

where f is the Attention based encoder-decoder network.

The neural networks were implemented on the PyTorch framework [27]. All experiments were run on a system with an Intel Core i7-6850K 3.6 GHz 6-Core Desktop Processor $\times$ 12 core CPU, 4 $\times$ GeForce GTX 1080 Ti 11GB VRAM.

In a real-time prosthetic controller, the trained network would be applied to predict joint trajectories for a user whose movements would not have been captured in the training dataset. To analyze performance for this use case, our evaluations were performed on a random test subject who was omitted from the training set cohort. A random trial from this test subject used as the validation set, and different trial from same test subject was used as the test set.

a) Data Normalization and Reshaping: Each of the joint angles exhibits a different Range of Motion (ROM). In order to prevent high-ROM joints from dominating predictions, it is common practice to normalize all features (generally 0 to 1). We normalized all joint angles for every trial and saved the average scaling factor of the training samples for de-normalizing the predicted joint angles.

b) Rolling Time Window: During training, LSTMs backpropagate errors a specific number of time steps back. The length of this history parameter, known as the sequence length, affects the time scale that the LSTM cell state reasons about. Choosing a longer sequence length increases the number of parameters that need to be trained, increasing computational load and requiring more training data. Choosing a shorter sequence length increases the difficulty of learning time dependencies in the data. In practice, choosing a sequence length appropriate to the inherent temporal

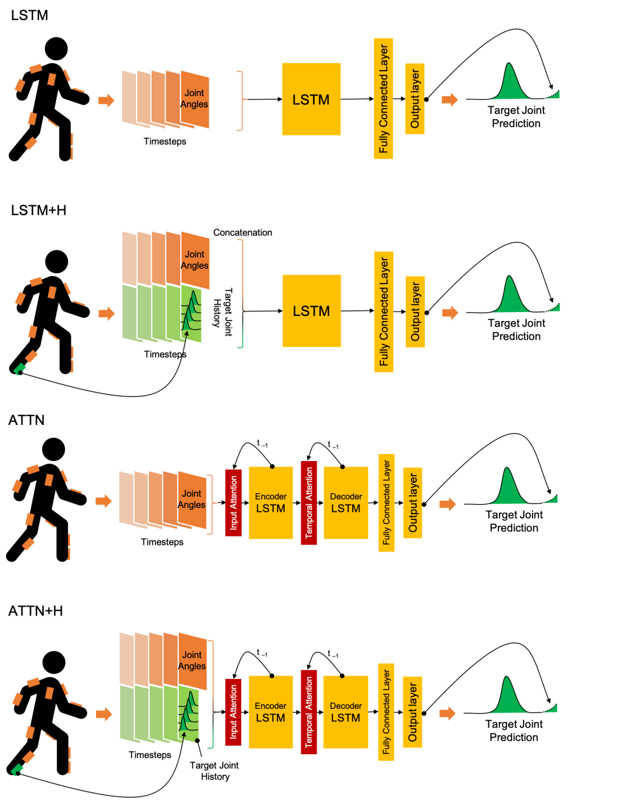

Fig. 2. System diagram. A time history of full-body joint angles (orange) are concatenated with a history of target joint angles (green) in the +H networks. The neural networks, two of which includes Attention, are trained to predict the current target joint angle.

Hyperparameter	Range/Values	Optimal
Learning Rate	$[10^{-5}:10^{-2}]$	5×10^{-4}
Batch Size	[50,100,300]	100
Number of Epochs	[20,50,100]	60
Input Sequence Length	[1,2,5,10,15,20,30]	10
Number of LSTM Layers	[1,2,4]	2
Number of Hidden Units	[4,16,32,64,128]	128
Regularization Rate	[0.5,0.05]	0.05
Standard Deviation of Random Noise	[0.01:0.02]	0.015

TABLE I

HYPERPARAMETER VALUES TESTED FOR OPTIMAL PERFORMANCE OF THE LSTM NETWORK ON THE OBSTACLE COURSE DATASET.

dynamics of the problem greatly simplifies training and performance of the network [28]. Training input samples were prepared as a overlapping rolling window of time series data of desired sequence length. The optimal sequence length was a hyperparamter we tuned for.

c) Loss Function and Neural Network Hyperparameter Optimization: We used the mean squared error (MSE) between the predicted and measured joint angle as loss function to be optimized. This is common metric used for regression tasks in machine learning. Apart from the sequence length, the network also has several hyperparameters that need to be optimized for different application domains.

Hyperparameter Optimization: A combination of random and grid search was applied to optimize LSTM network hyperparameters. Each minibatch was shuffled and random Gaussian noise was added to each sample to reduce overfitting.

Optimized hyperparameters included batch size, number of epochs, number of layers (L), number of units in each layer (HU), the standard deviation of the injected noise, the regularization parameter for L2 loss (λ), and learning rate.

Every 5 epochs, the performance of the model was evaluated on a validation set. The best performing model was saved and used to generate predictions and metrics on a test set. 30 trials were evaluated for each parameter set and the average RMSE was recorded. The optimal parameter value selection was based not just on the absolute best performance but also considering the overhead in time and computation needed to reach that performance. The range of parameter values tested is shown in Table I. The optimal hyperparameter set was used to compare and evaluate performance.

d) Denormalization: To report results in their original scale, all predictions were denormalized using avarage minimum and maximum scaling factors extracted from the training set. This is common practice in machine learning as test set scaling factors are not known a priori.

III. ANALYSIS

We use Root Mean Squared Error (RMSE) as our outcome measure to assess performance.

A. Inclusion of target joint history:

We assess the change in performance with inclusion of the target joint's history for two recurrent neural network models with and without the Attention mechanism, totalling 4 models: LSTM, LSTM+H, ATTN and ATTN+H. An ARIMAX model is used as baseline comparison for linear models.

B. Input sequence duration

We assess the benefit of longer time history for the two recurrent models. Both models (LSTM+H, ATTN+H) factor the target joint history as an input and hence the input sequence here refers to trajectories of both, intact joints as well as the target joint. The input sequence duration is varied from 160 ms to 1000 ms.

C. Temporal Dynamics of Attention

A significant gait event would be 'attended' to by the network and will be reflected as a higher attention weight corresponding to that instant in time. To show the dynamics of the attention weights progress through time, we use a 2D heat map. Input at each timestep contained an gait history of 1 second. Inputs were fed as overlapping rolling time widows with step size 30 ms. (See Section II-E.b). Hence an event would propagate diagonally along the next 30 timesteps.

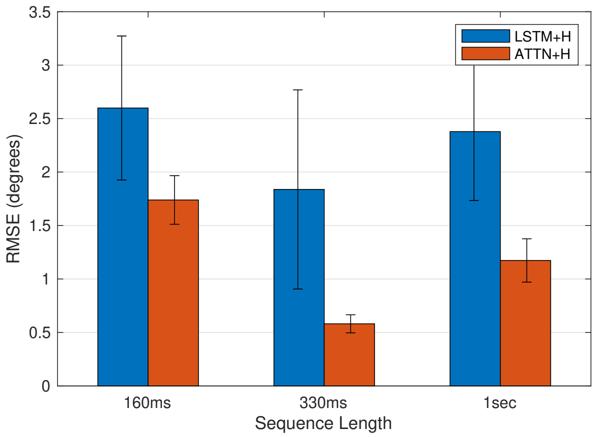

Fig. 3. Bar plot comparing the RMSE of ankle angle predictions in the sagittal plane using input sequences of lengths 160 ms, 330 ms, and 1 second for the LSTM+H and ATTN+H models. Performance improved for both models with longer time history but not beyond 330 ms

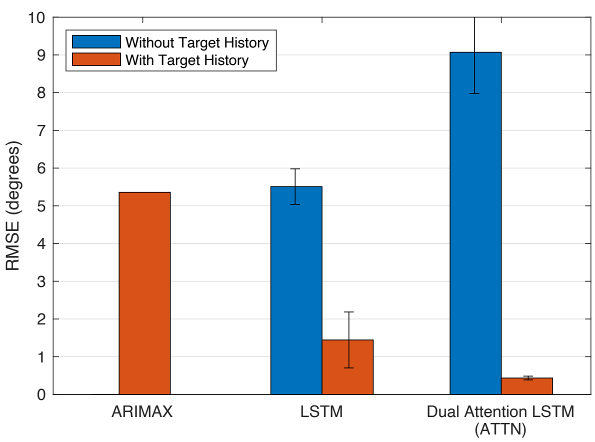

Fig. 4. Bar plot comparing the RMSE of ankle angle predictions in the sagittal plane between different models. Adding ankle angle history to both LSTM-based models show a significant reduction of error.

IV. RESULTS

A. Including the history of target joints

Fig. 4 shows that factoring the target joint's past movements can drastically improve performance for estimating current position. Of the 3 models that do consider target joint history, the linear ARIMAX model had the highest RMS error and the attention based model had the least. Interestingly, the attention based model without target joint history had the worst performance amongst all models.

Fig. 5 compares predictions of the model from our previous study [9] with the attention based model (ATTN+H).

B. Input sequence duration

Including a longer time history generally improved performance up to a certain point. This improvement was more pronounced for models with attention than without, indicating better emphasis on important aspects of data. But the trend of diminishing return was applicable to both networks with error increasing with very long time history.

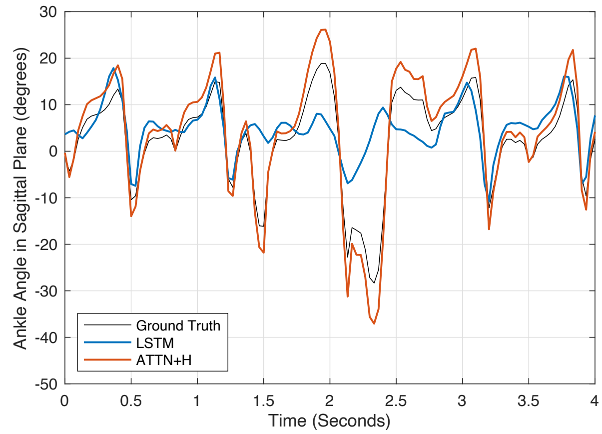

Fig. 5. Ankle angle predictions from our previously reported LSTM [9] and the ATTN+H models in the sagittal plane for the obstacle task.

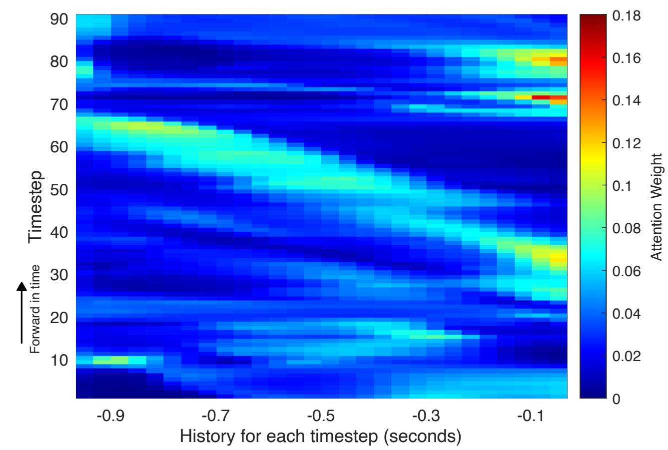

Fig. 6. Attention weights shown for a time window during a subject performing an obstacle crossing. The x-axis denotes the constant input history of 1 second for each timestep and the y-axis denotes progress of input timesteps. The Attention model placed larger weights on an event at the beginning of timestep 30. This event propagated along the history in subsequent timesteps and was maintained up to timestep 60, denoting high importance of that particular event.

C. Temporal Dynamics of Attention

Fig. 6 depicts an example of the evolving Attention weights. Each row depicts the Attention weights for a particular input timestep and each timestep contains about 1 sec of gait time history. The weights express the importance of information contained at that time. Columns on the right of the heatmap correspond to the most recent data. The procession of time corresponds to the positive direction on the y axis.

In this example, an important event occurs at approximately timestep 30 to 40. Attention places a high salience on the data generated then, depicted as "hotter" on the heatmap. As time continues, that event is maintained in high Attention weights, as can be seen by the sliding diagonal high values up to timestep 60. After that event exits the time window, attention is redirected to new movements at the beginning of the history window.

V. DISCUSSION

This line of research is motivated by the goal of relying on data to its maximum potential, and to hand-design as little as possible. Applying limited computational resources to the correct portions of that data is critical. The factors considered here have to do with time: continuity of predictions over time, and the appropriate time scale to pay attention to. In previous studies we have considered the role of different sensors, on the upper body or lower, or the contralateral limb. In the future we expect that Attention models can lend insight into the relative importance of input joints and sensors.

A. Inclusion of the time history of the target joints

This method produces dramatically superior results than those we have previously reported. For the challenging obstacle course data, we observed a drop from 6 degrees RMS error to about 1 degree RMS error with inclusion of target joint history. In practice, there could be complications that arise from including the target joint time history. At training time, the network learns to estimate the target joint angles of able-bodied individuals. However, at run time, the CM controller is used to generate reference trajectories for the prosthesis. The time history of target joint angles, then, is a time history of prosthetic joint angles. These data could differ in fundamental ways from those appearing in the training set, potentially causing unpredictable behavior. This issue will become more clear as we are in the process of conducting real usage tests using CM on a powered knee-ankle prosthesis.

B. Input sequence duration

Spatio-temporal behaviour patterns of human locomotion indicate that changes in movements needed for navigating a complex target occur within 1-2 steps prior [17, 18]. This would indicate that data comprising gait trajectory history of about 1 second would contain useful information for better joint angle prediction. However, we observed a diminishing return before that duration [10]. The useful cues from past movements present in the history could be lost in the volume of uninformative data. The improved performance of attention networks in this case indicates that efficient extraction of information could allow better use of the contained information. Compared to the single layer LSTM network, the attention model was more capable of taking advantage of additional information (Fig. 3). Even though the performance degraded slightly with very long input history (1 sec), it was still better than too short of a time history (160 ms).

C. Attention and temporal prioritization

Fig. 6 provides an example of how Attention, in addition to providing improved performance, can potentially provide insight. High Attention weightings correspond to data that are important for prediction. This provides an empirical measure of what data are actually being used by the network, which can help in hyperparameter design. For instance, in this study we observe high attention weights that occur even all the way at the "back" of the longest history we used (1 sec). This suggests that in some contexts, useful clues for predicting future movements appear a full second before those movements.

D. Limitations

The data used in this study to evaluate our novel approach is limited in size and the type of activities. We have previously shown that adding more data can improve network performance. This improvement was more pronounced with an increase in variability in data from including more subjects than with the addition of more data from the same subjects [10]. Although this analysis was performed on stair ambulation data, we expect this to apply to the obstacle avoidance data as well. This suggests that the small number of subjects (n=10) used is a limitation of this study. Moreover, the performance was evaluated on only one type of activity. However, the objective of this study was to show the relative improvement of prediction performance by augmenting the input data with the time history of the target prosthetic joint and the Attention-based network. We expect the results to hold with more subject data and for other activities as well.

The results we report here are limited to offline performance only and use RMS error as the outcome measure which is commonly used for regression tasks. It is unclear if the magnitude of the errors reported in this and our previous studies are within an acceptable threshold for a load-bearing prosthetic application. A more suitable outcome measure or a Minimal Clinically Important Difference (MCID) of error has not been established for gait and prosthetic control [29]. Live powered prosthesis experiments with human subjects would be needed to truly assess this control strategy.

VI. CONCLUSIONS

This study aims to address the challenges faced by powered lower limb users during unstructured activities such as side shuffling and weaving around obstacles. We have previously demonstrated that a data driven approach could be applicable to realize continuous control of powered prosthesis without explicit categorization of such movements. In this study, we improve performance at predicting target joint angles by including a time history of those target joint angles, and by using Attention models that modulate the importance of the data through time.

ACKNOWLEDGMENT

The authors would like to thank Melissa Calvert, Yunqian (Waterlily) Huang and Megan Bui for their assistance in data collection and processing.

REFERENCES

[1] Alejandro F Azocar et al. "Design and Characterization of an Open-Source Robotic Leg Prosthesis". In: *2018 7th IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob)*. IEEE. 2018, pp. 111–118.

[2] T. Lenzi et al. "Design, Development, and Validation of a Lightweight Nonbackdrivable Robotic Ankle Prosthesis". In: *IEEE/ASME Transactions on Mechatronics* 24.2 (2019), pp. 471–482. ISSN: 1941-014X. DOI: 10.1109/TMECH.2019.2892609.

[3] Michael R Tucker et al. "Control strategies for active lower extremity prosthetics and orthotics: a review". In: *Journal of neuroengineering and rehabilitation* 12.1 (2015), p. 1.

[4] Rene Fluit et al. "A comparison of control strategies in commercial and research knee prostheses". In: *IEEE Transactions on Biomedical Engineering* (2019).

[5] Huseyin Atakan Varol, Frank Sup, and Michael Goldfarb. "Multiclass real-time intent recognition of a powered lower limb prosthesis". In: *Biomedical Engineering, IEEE Transactions on* 57.3 (2010), pp. 542–551.

[6] JA Spanias et al. "Online adaptive neural control of a robotic lower limb prosthesis". In: *Journal of neural engineering* 15.1 (2018), p. 016015.

[7] Levi J Hargrove et al. "Intuitive control of a powered prosthetic leg during ambulation: a randomized clinical trial". In: *JAMA* 313.22 (2015), pp. 2244–2252.

[8] Brian J Hafner et al. "Characterizing mobility from the prosthetic limb user's perspective: Use of focus groups to guide development of the Prosthetic Limb Users Survey of Mobility". In: *Prosthetics and orthotics international* 40.5 (2016), pp. 582–590.

[9] Vijeth Rai and Eric Rombokas. "A Framework For Mode-Free Prosthetic Control For Unstructured Terrains". In: *2019 IEEE 16th International Conference on Rehabilitation Robotics (ICORR)*. IEEE. 2019, pp. 796–802.

[10] V. Rai, A. Sharma, and E. Rombokas. "Mode-free Control of Prosthetic Lower Limbs". In: *2019 International Symposium on Medical Robotics (ISMR)*. 2019, pp. 1–7. DOI: 10.1109/ISMR.2019.8710187.

[11] Richard Lewis and Gregory C Reinsel. "Prediction of multivariate time series by autoregressive model fitting". In: *Journal of multivariate analysis* 16.3 (1985), pp. 393–411.

[12] Kody M Powell et al. "Heating, cooling, and electrical load forecasting for a large-scale district energy system". In: *Energy* 74 (2014), pp. 877–885.

[13] Amaury Lendasse et al. "Non-linear financial time series forecasting-Application to the Bel 20 stock market index". In: *European Journal of Economic and Social Systems* 14.1 (2000), pp. 81–91.

[14] Hong Thom Pham, Bo-Suk Yang, et al. "A hybrid of nonlinear autoregressive model with exogenous input and autoregressive moving average model for long-term machine state forecasting". In: *Expert Systems with Applications* 37.4 (2010), pp. 3310–3317.

[15] Sepp Hochreiter and Jürgen Schmidhuber. "Long short-term memory". In: *Neural computation* 9.8 (1997), pp. 1735–1780.

[16] Andrej Karpathy. "The unreasonable effectiveness of recurrent neural networks". In: *Andrej Karpathy blog* 21 (2015), p. 23.

[17] Aftab E Patla and Joan N Vickers. "How far ahead do we look when required to step on specific locations in the travel path during locomotion?" In: *Experimental brain research* 148.1 (2003), pp. 133–138.

[18] Jonathan Samir Matthis, Jacob L Yates, and Mary M Hayhoe. "Gaze and the control of foot placement when walking in natural terrain". In: *Current Biology* 28.8 (2018), pp. 1224–1233.

[19] Tara N Sainath et al. "Deep convolutional neural networks for LVCSR". In: *2013 IEEE international conference on acoustics, speech and signal processing*. IEEE. 2013, pp. 8614–8618.

[20] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. "Neural machine translation by jointly learning to align and translate". In: *arXiv preprint arXiv:1409.0473* (2014).

[21] Ashish Vaswani et al. "Attention is all you need". In: *Advances in neural information processing systems*. 2017, pp. 5998–6008.

[22] *MTw Awinda - Products - Xsens 3D motion tracking*. https://www.xsens.com/products/mtw-awinda/. (Accessed on 10/29/2018).

[23] Rebecca J Hess et al. "Walking skill can be assessed in older adults: validity of the Figure-of-8 Walk Test". In: *Physical therapy* 90.1 (2010), pp. 89–99.

[24] Ge Wu et al. "ISB recommendation on definitions of joint coordinate system of various joints for the reporting of human joint motion—part I: ankle, hip, and spine". In: *Journal of biomechanics* 35.4 (2002), pp. 543–548.

[25] Ming Sun and JO Hill. "A method for measuring mechanical work and work efficiency during human activities". In: *Journal of biomechanics* 26.3 (1993), pp. 229–241.

[26] Yao Qin et al. "A dual-stage attention-based recurrent neural network for time series prediction". In: *arXiv preprint arXiv:1704.02971* (2017).

[27] Adam Paszke et al. "Automatic differentiation in PyTorch". In: *NIPS-W*. 2017.

[28] Nils Reimers and Iryna Gurevych. "Optimal hyperparameters for deep lstm-networks for sequence labeling tasks". In: *arXiv preprint arXiv:1707.06799* (2017).

[29] *Microprocessor-controlled Lower Limb Prostheses*. https://www.hca.wa.gov/assets/program/mclowerprostheticfinalreport[1].pdf. (Accessed on 03/30/2020).