

Gait Phase and Foot Trajectory Prediction During Dynamic Walking Using Gated Recurrent Units

Hayley Wisman

Thesis submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Master of Science

in

Computer Engineering

Creed F. Jones, Chair

Alan T. Asbeck, Co-chair

William T. Baumann

February 19, 2025

Blacksburg, Virginia

Keywords: Deep learning, gait phase, foot trajectory

Copyright 2025, Hayley Wisman

Gait Phase and Foot Trajectory Prediction During Dynamic Walking Using Gated Recurrent Units

Hayley Wisman

(ABSTRACT)

In the field of assistive robotics and exoskeletons, foot trajectory prediction has the potential to play a pivotal role in improving the functionality and user experience of worn devices. Rather than operating as a reactive system which only responds to user movement, a device which predicts future foot position can anticipate an action before it occurs, reducing latency and moving with the wearer for a more natural, uninhibited motion. While previous studies have focused on predicting continuous motion, they often overlook critical transitions between walking and standing, which are essential for natural locomotion. We propose in this study a foot trajectory prediction approach which leverages a recurrent deep learning architecture to make predictions based on sequential walking data. The first of the two machine learning models predicts the gait phase as a value between 0 and 1, while the second model leverages the gait phase prediction output to predict foot position in three dimensions. The models were trained and evaluated on IMU sensor data collected from three subjects instructed to walk on a treadmill at speeds varying from 0.5 mph to 1.5 mph. The resulting mean absolute error on gait phase percentage across subjects and velocity was 1.92%. For foot trajectory prediction, the cross-subject trained model achieved mean distance error of 2.85 ± 2.89 cm, 3.29 ± 2.82 cm, 4.15 ± 4.12 cm, 5.33 ± 5.46 cm, and 6.92 ± 6.56 cm with prediction horizons of 0.1s, 0.25s, 0.5s, 1s, and 2s, respectively.

Gait Phase and Foot Trajectory Prediction During Dynamic Walking Using Gated Recurrent Units

Hayley Wisman

(GENERAL AUDIENCE ABSTRACT)

Foot trajectory prediction, which is the process of estimating the movement path of a person's foot during a walking sequence, can be used in the advancement of assistive robotics such as exoskeletons. Knowing where the next step will end and the path it will take to get there can help align the device with the user's intended movements and make the experience more natural. Previous work has often been limited to predicting foot placement during periods of constant walking, but not across sequences of both walking and standing. These transitions are important to capture in order to effectively emulate the human walking pattern and make assistive walking devices more practical and comfortable for everyday use. The approach to this problem proposed in this paper uses two machine learning models trained on walking data gathered from IMUs placed near the feet of the user. The results indicate that this method could be effective for predicting foot trajectory during the use of a walking exoskeleton across walking and standing transitions.

Dedication

This study is dedicated to my family who have offered me unwavering support in all my life's endeavors and to the educators throughout my life who encouraged my dreams.

Acknowledgments

I would like to thank my advisor, Professor Alan Asbeck, for his support, guidance, and patience throughout my journey to completing this thesis. His expertise in assistive robotics was invaluable to this research, and his mentorship allowed me to grow academically, professionally, and personally. He offered me a research opportunity that aligns with my passions and challenged me, and despite the obstacles I faced along the way, he provided continued support and a positive outlook. I also owe thanks to Musharrat Mau of the Assistive Robotics Lab for performing data collection and developing the algorithms used to compute ground truth foot positions from the IMU data.

Furthermore, I extend my gratitude to Professor William Baumann and Professor Creed Jones for not only serving on my committee, but for their support throughout my academic journey. They exemplify the ideals of professorship through their roles as researchers and educators, and through their teachings encouraged me to strive for excellence.

Contents

List of Figures	ix
List of Tables	xiv
1 Introduction	1
1.1 Contributions	3
2 Review of Literature	5
2.1 Gait Phase Estimation	5
2.2 Foot Placement and Trajectory Prediction	6
2.3 Limitations of Prior Work	7
3 Methods	8
3.1 Exoskeleton	8
3.2 Data Processing	12
3.3 Gait Phase Prediction	14
3.3.1 Data Processing and Setup	14
3.3.2 Model Architecture	18
3.3.3 Post-Processing	19

3.4	Foot Trajectory Prediction	23
3.4.1	Data Processing and Setup	23
3.4.2	Model Architecture	23
3.4.3	Post-Processing	29
3.5	Lookback Window Sampling	29
4	Results	30
4.1	Metrics	30
4.2	Gait Phase Prediction	31
4.3	Foot Trajectory Prediction	40
4.4	Lookback Window	53
5	Discussion	54
6	Conclusions	57
	Bibliography	58
	Appendices	61
	Appendix A First Appendix	62
A.1	Additional Input Features	62
A.2	Supplemental Graphs	63

A.3 Early Foot Trajectory Prediction	64
--	----

List of Figures

3.1	Front and side views of the exoskeleton.	9
3.2	Rear view of the exoskeleton	10
3.3	Left foot sensor setup along with the degrees of freedom between the exoskeleton foot and human foot.	11
3.4	High-level system flowchart.	12
3.5	Lookback window linear vs. logarithmic algorithm sample distributions.	14
3.6	Foot Angular Velocity with markers indicating gait events.	16
3.7	Gait phase progression with markers indicating gait events.	17
3.8	Gait phase for the left foot in continuous form with markers for heel strikes.	18
3.9	Gait phase continuous representation without Savitzky-Golay filter applied.	21
3.10	Gait phase continuous representation after Savitzky-Golay filter application.	22
3.11	Side view ground truth foot position visualization with left foot (red) in early swing phase.	24
3.12	Side view ground truth foot position visualization with left foot (red) in late swing phase.	25
3.13	Isometric view ground truth foot position visualization with left foot (red) in late swing phase.	26

3.14	Isometric view ground truth foot position visualization with left foot (red) in late swing phase.	27
3.15	Ground truth foot positions during walking generated using IMU data.	28
4.1	Left foot histogram showing the MAE of gait phase predictions.	32
4.2	Right foot histogram showing the MAE of gait phase predictions.	33
4.3	Component predictions during a period of steady-state walking.	34
4.4	Percent predictions during a period of steady-state walking. The MAE taken over this segment was 2.05% for the left foot and 0.67% for the right foot.	35
4.5	Transitional gait phase component predictions at start of walking.	36
4.6	Transitional gait phase percent at start of walking.	37
4.7	Transitional gait phase component predictions at start of walking.	38
4.8	Transitional gait phase percent at start of walking.	39
4.9	Foot trajectory predictions and ground truth over approximately one foot step, predicted 0.10 seconds in advance.	41
4.10	Histogram showing the combined prediction error of the left and right foot positions 10 timesteps in advance.	42
4.11	Histogram showing the combined prediction error of the left and right foot positions 25 timesteps in advance.	43
4.12	Histogram showing the combined prediction error of the left and right foot positions 50 timesteps in advance.	44

4.13	Histogram showing the combined prediction error of the left and right foot positions 100 timesteps in advance.	45
4.14	Histogram showing the combined prediction error of the left and right foot positions 200 timesteps in advance.	46
4.15	Right foot position predictions in the x-direction (side-to-side) predicted 0.1 seconds into the future.	47
4.16	Right foot position predictions in the y-direction (forwards and backwards) predicted 0.1 seconds into the future.	47
4.17	Right foot position predictions in the z-direction (up and down) predicted 0.1 seconds into the future.	48
4.18	Right foot position predictions in the x-direction (side-to-side) predicted 0.5 seconds into the future.	48
4.19	Right foot position predictions in the y-direction (forwards and backwards) predicted 0.5 seconds into the future.	49
4.20	Right foot position predictions in the z-direction (up and down) predicted 0.5 seconds into the future.	49
4.21	Right foot position predictions in the x-direction (side-to-side) predicted 1 second into the future.	50
4.22	Right foot position predictions in the y-direction (forwards and backwards) predicted 1 second into the future.	50
4.23	Right foot position predictions in the z-direction (up and down) predicted 1 second into the future.	51

4.24	Right foot position predictions in the x-direction (side-to-side) predicted 2 second into the future.	51
4.25	Right foot position predictions in the y-direction (forwards and backwards) predicted 2 second into the future.	52
4.26	Right foot position predictions in the z-direction (up and down) predicted 2 second into the future.	52
A.1	Unfiltered gait phase percent predictions plotted alongside ground truth. . .	63
A.2	Gait phase percent predictions after Savitzky-Golay filtering plotted alongside ground truth.	64
A.3	Foot trajectory prediction error over increasing prediction horizons.	65
A.4	Histogram showing the combined prediction error of the left and right foot positions 10 timesteps in advance.	66
A.5	Histogram showing the combined prediction error of the left and right foot positions 25 timesteps in advance.	67
A.6	Histogram showing the combined prediction error of the left and right foot positions 50 timesteps in advance.	68
A.7	Histogram showing the combined prediction error of the left and right foot positions 100 timesteps in advance.	69
A.8	Histogram showing the combined prediction error of the left and right foot positions 200 timesteps in advance.	70
A.9	Right foot position predictions in the x-direction (side-to-side) predicted 0.1 seconds into the future.	71

A.10 Right foot position predictions in the y-direction (forwards and backwards) predicted 0.1 seconds into the future.	71
A.11 Right foot position predictions in the z-direction (up and down) predicted 0.1 seconds into the future.	72
A.12 Right foot position predictions in the x-direction (side-to-side) predicted 0.5 seconds into the future.	72
A.13 Right foot position predictions in the y-direction (forwards and backwards) predicted 0.5 seconds into the future.	73
A.14 Right foot position predictions in the z-direction (up and down) predicted 0.5 seconds into the future.	73
A.15 Right foot position predictions in the x-direction (side-to-side) predicted 1 second into the future.	74
A.16 Right foot position predictions in the y-direction (forwards and backwards) predicted 1 second into the future.	74
A.17 Right foot position predictions in the z-direction (up and down) predicted 1 second into the future.	75
A.18 Right foot position predictions in the x-direction (side-to-side) predicted 2 second into the future.	75
A.19 Right foot position predictions in the y-direction (forwards and backwards) predicted 2 second into the future.	76
A.20 Right foot position predictions in the z-direction (up and down) predicted 2 second into the future.	76

List of Tables

3.1 Hyperparameter Search Ranges	20
--	----

List of Abbreviations

AO Adaptive Oscillator

FSR Force Sensitive Resistor

GRF Ground Reaction Force

GRU Gated Recurrent Unit

IMU Inertial Measurement Unit

LSTM Long Short-Term Memory

MAD Median Absolute Deviation

MAE Mean Absolute Error

RMSE Root Mean Squared Error

RNN Recurrent Neural Network

TPE Tree-structured Parzen Estimator

Chapter 1

Introduction

Assistive robotics is a broad field encompassing technologies such as prostheses, exoskeletons, and other wearable devices where the common denominator is the goal of restoring or augmenting the physiological capabilities of the user. Specifically, lower body assistive devices may aim to help individuals with walking disabilities regain mobility or to help an able-bodied person bear weight that they normally cannot handle comfortably for long periods of time. Designing devices that emulate and cooperate with the intricacies of bipedal motion, however, is no easy task as some phases of the gait cycle are not strictly controlled but rather depend on the overall pattern of motion to maintain stability. Foot placement is of vital importance in this process, forming the basis of the walking pattern, and is used extensively to maintain a stable gait pattern and avoid falls [1]. Understanding and leveraging foot position and placement is, therefore, essential for advancing assistive walking technologies.

Foot placement prediction involves estimating in three dimensions where the next footstep in a sequence will fall based on prior movement patterns and gait phase. Foot trajectory extends this to include predictions for the full path of motion, offering more precise insights into the subject's movements through the gait cycle. This is a multi-faceted problem—foot placement is affected not only by intent, but by factors beyond the individual's control such as terrain characteristics, physical and cognitive health, and external events which may occur late in the swing phase, forcing the individual to suddenly change their intended foot placement.

Traditional approaches use simplified models which struggle to capture the complexities of the human gait [2] and more advanced approaches leveraging machine learning are often limited in practicality for various reasons including data requirements and their failure to account for transitions between walking and standing.

The sequential nature of walking data makes recurrent neural networks (RNNs) a suitable approach for foot trajectory prediction. By retaining information from previous inputs through their hidden states, RNNs can learn and model temporal dependencies within the data; thus, the model can use previous footsteps to predict where the next one will fall. Standard RNNs, however, have certain pitfalls such as a tendency towards vanishing and exploding gradients, which make it difficult for the model to learn long-term dependencies [3]. The gated recurrent unit, or GRU, is a type of RNN that was designed to solve the vanishing gradient problem in RNNs by introducing update and reset gates. These gates regulate how past information influences the current hidden state, allowing the model to better capture long-term dependencies. This capability makes GRUs particularly well-suited for foot trajectory prediction, where capturing long-term temporal dependencies is crucial for accurately estimating future foot positions [4].

Foot trajectory prediction has the potential to significantly improve the functionality and practicality of lower-body assistive technology, such as prostheses and exoskeletons, by allowing them to anticipate user movements in real-time. By predicting intent and adjusting proactively, these devices could offer a more natural walking experience and improve the comfort and stability of assistive technology for diverse groups of people such as those with motor impairments, workers who routinely perform heavy lifting, and people undergoing physical rehabilitation therapy, among others.

This body of research is focused on developing and evaluating deep learning models that are capable of accurately predicting human gait phase and foot trajectory over dynamic walking

conditions including changes in speed, as well as transitions between stopping and starting.

1.1 Contributions

This research presents several novel contributions to the field of gait phase and foot trajectory prediction, including the development of a model which is capable of predicting both walking and standing, and the transitions between the two. Further, the potential impact of two input feature lookback window algorithms is investigated, along with the integration of sensor data from IMUs placed on both the subject’s feet and the exoskeleton. Finally, future foot positions are predicted at each timestep in the walking sequence, forming a complete picture of the foot trajectory throughout the gait cycle and allowing the model to more closely emulate the subject’s motion for a smoother, more natural walking experience.

The current exoskeleton control system relies on an error-prone method that uses discrete gait events, such as minima and maxima in foot velocity, to control ground reaction force (GRF). Gait phase prediction, however, has the potential to improve this by allowing the control system to adapt the GRF during walking using continuous, real-time information about the gait cycle. Beyond the direct application of the predicted gait phase to control the GRF, it also has value as an input feature to the model predicting foot trajectory. Due to the delay present between when a command is sent and when the receiving motor on the exoskeleton reaches the desired position, it is imperative that the control system be able to anticipate future foot positions far enough in advance that the exoskeleton does not lag behind the user.

Unlike many previous studies, our approach does not have major additional equipment requirements to train. Rather than using expensive setups such as motion capture, gaze tracking, or cameras, the ground truth labels for both models in this study are generated

using data from sensors that are already on the exoskeleton. This not only decreases the cost to train the generalized model, but makes it more practical for training on new subject data outside of a lab environment.

Chapter 2

Review of Literature

The concepts of gait phase and foot placement or trajectory prediction are not new—various bodies of work have been produced exploring different methodologies from mechanical models [2] to more complex approaches leveraging machine learning. Despite this, opportunities for growth in robustness, accuracy, and real-world practicality remain plentiful.

2.1 Gait Phase Estimation

Kang et al. [5] proposed a sensor fusion-based approach to estimate gait phase using data collected from thigh and trunk-mounted IMUs, an encoder measuring hip angle, and force sensitive resistors (FSR) detecting heel strikes. The model architecture included a single layer fully connected neural network trained on 300 ms sliding windows of data. Three similar models were created and evaluated—one which was tested on a hold-out subject to evaluate generalization to new users; one trained on a single subject’s data and tested on hold-out data from the same subject; and a hybrid model that was pre-trained on other subjects, fine-tuned using a subset of the target subject’s data, and tested on the remaining hold-out data from that subject. The hybrid model achieved an average gait phase RMSE of $5.22 \pm 0.81\%$ during dynamic-speed walking over the range 0.6 m/s to 1.1 m/s.

In an effort to solve the problem of predicting gait phase during non-steady state walking, Wang et al. [6] used a dual adaptive oscillator (AO) approach, implementing a high-level

and low-level AO. The high-level AO learned three parameters—phases, amplitudes, and intrinsic frequency—from analyzing gait events within the last stride. These parameters were then used in place of the low-level AO’s initial values, giving it insight into the subject’s gait from the beginning rather than starting with generic estimates. The low-level adaptive oscillator then updated its parameters based on real-time limb angle input to estimate gait phase. This method demonstrated improvement over conventional, single-layer AO implementations, particularly in reducing the number of transition strides, which were defined as the strides required to reduce the phase reset error to less than 5% of the gait cycle. The dual gait phase AO solution achieved 2.63 ± 0.50 transition strides during interval walking, whereas the conventional AO required 4.00 ± 0.88 .

2.2 Foot Placement and Trajectory Prediction

In [7], orientation data gathered from an IMU mounted to the heel of the subject was used to predict foot placement in the early swing phase with a probability distribution model, achieving mean absolute difference error of 13.51 cm in cross-subject testing across walking velocities ranging from 0.5 m/s to 1.4 m/s. This demonstrates the possibility of using motion features to predict foot placement far in advance of what had been done previously, but leaves room for improvement in terms of accuracy. [8] achieved mean distance error (MDE) on foot placement predictions during steady-state walking of 6.99 cm, 3.32 cm, and 3.22 cm using pre-trained, base, and fine-tuned GRU models taking as input gait phase and IMU data. While the previous studies focused on level-ground walking on a treadmill, foot placement was predicted over complex terrain in [9] using a fusion of environmental context and foot motion data with RMSE 8.19 ± 1.20 cm.

Tanghe et al. [10] used a probabilistic principal component analysis model to predict gait

events and joint trajectories during walking. Joint angles, velocity, acceleration, and jerk were predicted with high accuracy for the hip, knee, and ankle using treadmill walking data collected from motion capture cameras and ground reaction forces under the feet. This data was gathered at walking speeds ranging from 2 - 5 km/h (about 0.56 - 1.39 m/s) with and without the exoskeleton. Error for each joint was calculated for predictions at 10%, 50%, and 100% of a future gait cycle, all achieving a mean less than 1%, although standard deviations ranged from 4.4% to 28.5%. One study by Dey and Schilling [11] predicted angular foot positions in the sagittal plane 10 samples (at 100 Hz) into the future using a convolutional neural network taking IMU input with mean NRMSE of 6.0%.

2.3 Limitations of Prior Work

Much of the research done on this topic previously focuses solely on steady-state walking—that is, predictions are made only on periods of constant walking with no pauses in between or at constant speed, although [12] performance activity classification which includes standing. The lack of real-time gait phase estimation over both walking and standing limits the practical application of the techniques since everyday locomotion is much less consistent and includes periods of walking and standing, as well as acceleration and deceleration. The earliest in the gait cycle predictions can be made also limits some approaches, such as [7] and [13], which are bound to the latter 50% of the gait cycle. Further, a great deal of the literature on gait phase and foot placement prediction use complex setups for motion capture and gaze tracking, making the implementations less mobile [10], [14]. Lastly, the majority of studies predict the final placement of the foot at the end of a step, whereas our approach makes continuous future predictions of foot position in three dimensions.


Chapter 3

Methods

3.1 Exoskeleton

The design of the lower body exoskeleton used in this study (Figures 3.1 and 3.2) includes a waist belt and back module held to the user’s torso. Two exoskeleton legs attach to the torso module, running down next to the user’s legs. These include an “exoskeleton foot” at the bottom that contacts the ground directly, and a linkage connecting the exoskeleton leg to the user’s foot. There are unpowered joints at the back and side between the torso module and exoskeleton legs, so that the user can move their hip (with the exoskeleton) freely in any direction. The exoskeleton legs consist of two prismatic joints that can extend and contract. These are powered, and push down on the ground. When they are in contact with the ground, the exoskeleton harness pushes up on the wearer’s waist and thighs and can support up to 100 lbs of their bodyweight.

As shown in Figure 3.3, the human foot and exoskeleton foot together are able to rotate in abduction/adduction due to Axis 1. The human foot can move in inversion/eversion due to Axis 2 between the human foot and exoskeleton foot. The human foot can lift up and down and tilt forward/back (plantarflexion/dorsiflexion) due to the joints at Axes 3 and 4. IMUs on the exoskeleton foot (IMU 1) and the human foot heel (IMU2) measure accelerations at and the orientation of these locations. An encoder is used to calculate the vertical displacement of the human foot from the exoskeleton foot in conjunction with the



(a) Front view of the exoskeleton



(b) Side view of the exoskeleton.

Figure 3.1: Front and side views of the exoskeleton.



Figure 3.2: Rear view of the exoskeleton




Figure 3.3: Left foot sensor setup along with the degrees of freedom between the exoskeleton foot and human foot.

link between Axis 3 and Axis 4.

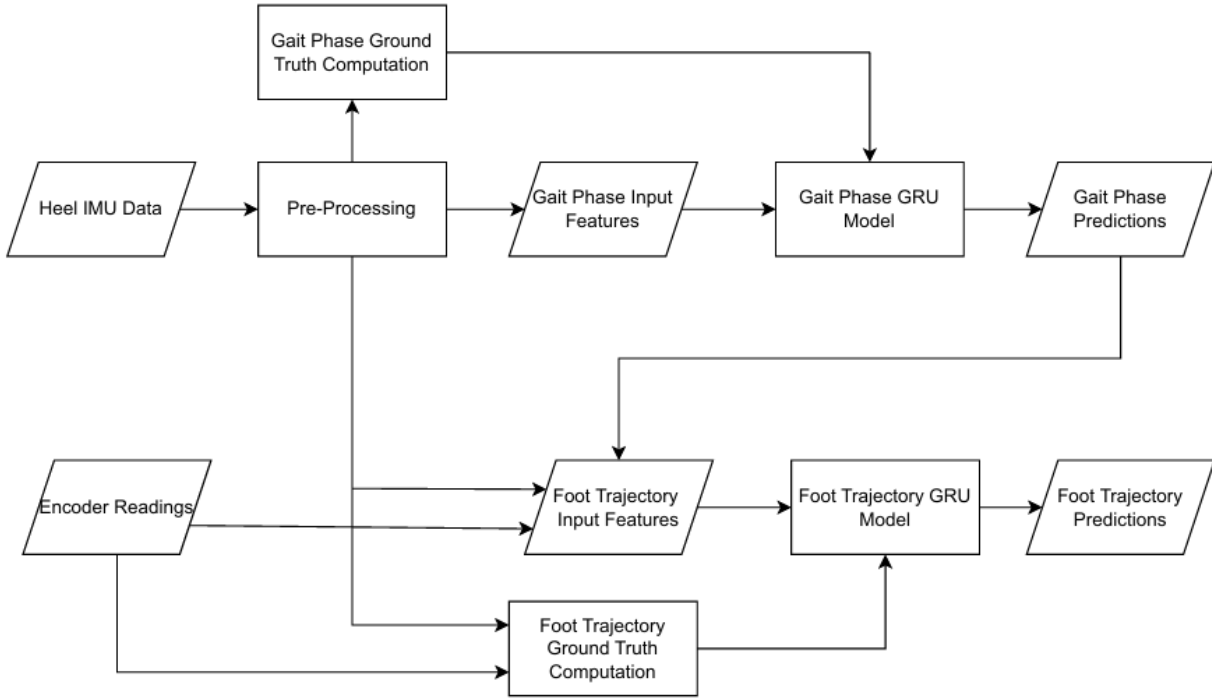


Figure 3.4: High-level system flowchart.

3.2 Data Processing

The data used in this study was collected from sensors worn by $N=3$ female subjects (ages 23.0 ± 6.1 years, heights 163.4 ± 2.9 cm) in an experiment approved by the Virginia Tech Institutional Review Board (IRB# 21-165). The sensors included IMUs placed on the back of the subject's heels and on the exoskeleton feet, as well as encoders which measured foot rotation relative to the exoskeleton, as shown in Figure 3.3.

Participants were instructed to walk on a treadmill at speeds varying from 0.5 mph to 1.5 mph (approximately 0.22 m/s - 0.67 m/s) while wearing the exoskeleton. In between these periods of walking, the treadmill was slowed to a stop and then restarted. Data was

collected on periods of constant walking speed as well as on those which accelerated and decelerated during walking. Before constructing the input feature vectors, multiple pre-processing techniques were applied to prepare the input data for training the models. First, rows in the sensor data CSV files which did not contain valid data were corrected by forward filling them with a copy of the data from the previous timestep. Since it is relative timing that is most important in gait phase prediction, the time column in each file was also scaled to start at zero to maintain consistency across datasets and avoid inadvertently biasing the model with large starting values, and outliers in the time data were replaced with the interpolation of the two surrounding values.

The quaternions output by the IMU sensors were used to compute the foot angles in the x, y, and z axes to reduce dimensionality and make the data more interpretable for the model. The angles were subsequently filtered using a 4th order low-pass Butterworth filter with a cutoff frequency of 10 Hz to attenuate high-frequency noise. Since most walking motions occur at lower frequencies, the 10 Hz cutoff was adequate to reduce noise while preserving important signal characteristics. The velocity and acceleration data, on the other hand, were band-pass filtered with a low cutoff of 0.01 Hz and a high cutoff of 5 Hz, thereby mitigating drift and reducing high-frequency noise. The phase response of the Butterworth filters introduced a time shift of about 0.05 seconds, which was acceptable for our purposes, but could be reduced by decreasing the order of the filter, widening the passband, or using a filter with a more mild frequency response. Lastly, since the ranges of the input features vary significantly, they were also standardized to zero mean and unit variance to minimize bias and improve convergence.

In order to capture the temporal dependencies of the data, the input feature vectors were augmented to include a lookback window of up to 125 preceding samples to the current timestep, collected at a rate of 100 Hz, by combining the previous timesteps along the

vertical axis. Windows created at the beginning or end of a dataset were zero-padded to be of uniform length. Two sample selection methods were tested and evaluated based on model accuracy, the first being linear sampling and the second being logarithmic. The linear sampling method took samples at a fixed step size across the entire lookback history, whereas the logarithmic method prioritized samples nearer to the current timestep.




Figure 3.5: Lookback window linear vs. logarithmic algorithm sample distributions.

3.3 Gait Phase Prediction

3.3.1 Data Processing and Setup

The input feature vector for the gait phase model consisted of 19 features—this data, used for training and testing the model, was gathered from two IMUs mounted to the subject’s feet. These provided time, velocity, acceleration, and orientation data at 100 Hz. In addition to the time and raw velocity and acceleration data, the input feature vector also contained foot angles in the x, y, and z directions which were derived from the quaternions output by the IMUs. The lookback window for this model contained 100 samples taken linearly over a 3

second window preceding the target timestep. The model saw little to no improvement with more samples in the lookback window, and taking these samples over a 3 second history saw the best performance. The ground truth gait phase labels were generated using a heel-strike detection algorithm which identifies heel-strikes as the first local minimum of foot angular velocity after a local maximum. Gait phase percentage for each foot was then computed as a linear progression from 0 to 1 between two successive heel-strikes.

The ground truth gait phase labels were determined using the foot IMU sensor readings by detecting consecutive peaks in angular velocity, where heel-strikes correspond to the first minimum after a maximum and toe-offs occur at the second minimum [15]. To eliminate erroneous detections, heel-strikes were retained only if followed closely by a corresponding toe-off or if they marked the final heel-strike in a walking sequence. The gait phase percentage was then computed as the linear progression from 0 to 1 between two consecutive heel-strikes with the same foot. Additionally, a pause threshold of 5 seconds was used to define separate walking sequences—that is, if the time elapsed after a heel-strike exceeded the pause threshold without encountering another heel-strike, the gait phase was reset to zero as this was considered a pause in walking. The threshold value of 5 seconds was chosen to accommodate the full range of walking speeds evaluated in this study.

Although this method provides a useful approximation of the gait phase, the abrupt transition between 100% at the end of one step to 0% at the start of the next introduces a discontinuity that may hinder the model’s ability to learn the smooth, cyclic nature of gait. To address this, the gait phase was transformed into a continuous 2D representation using sine and cosine components, eliminating the artificial linearity of the labels that could pose challenges to the model [5], [16], [8].

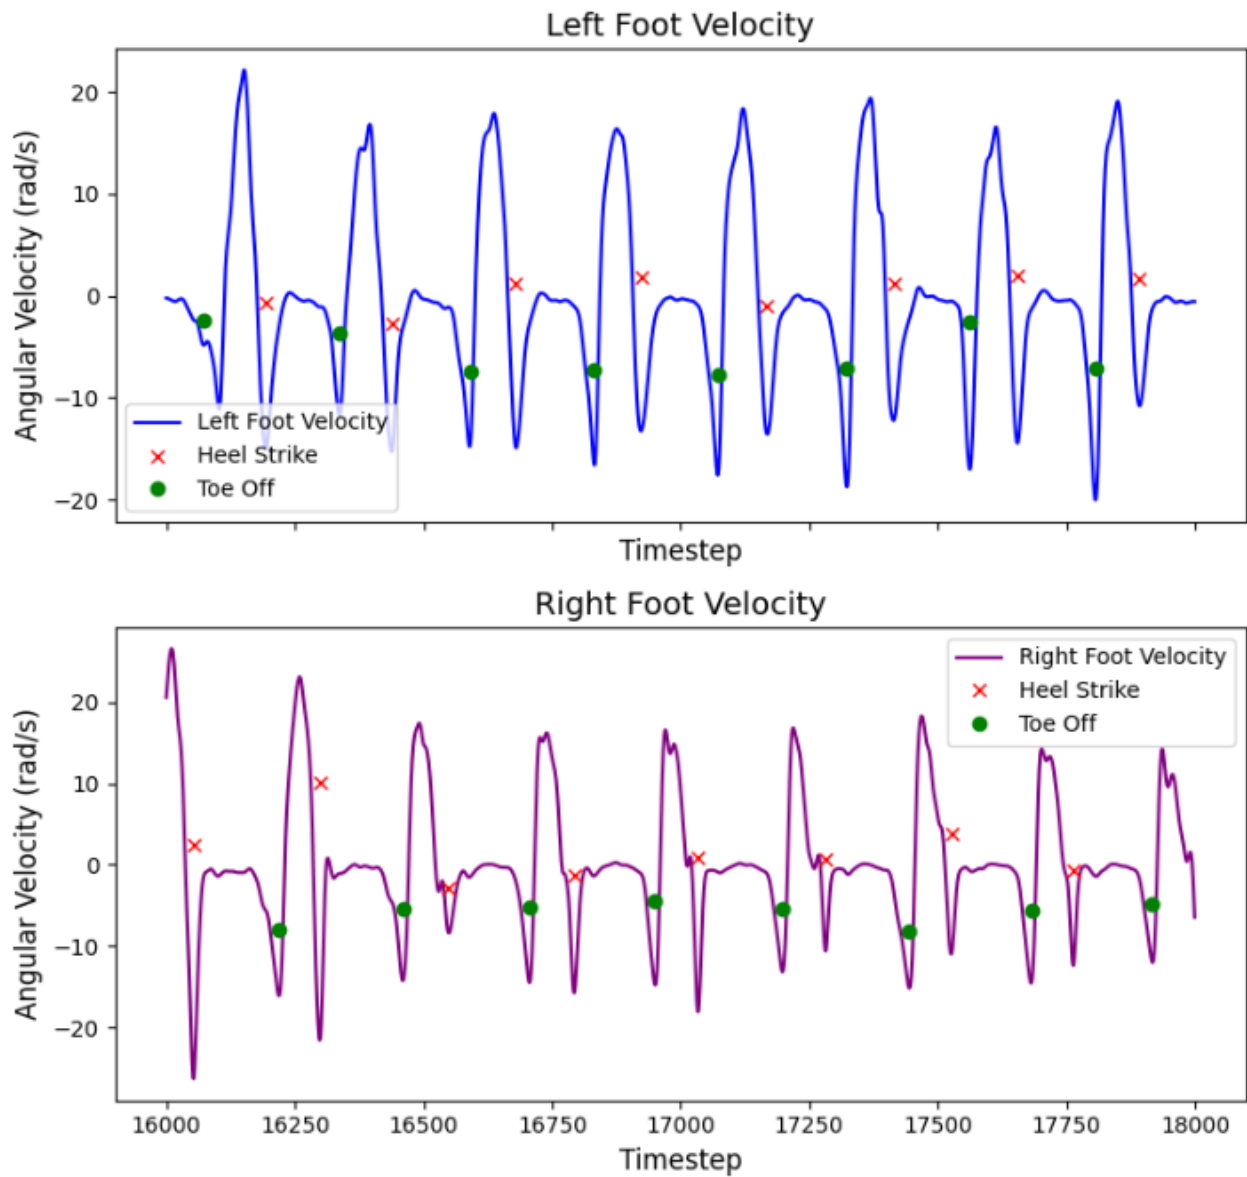


Figure 3.6: Foot Angular Velocity with markers indicating gait events.

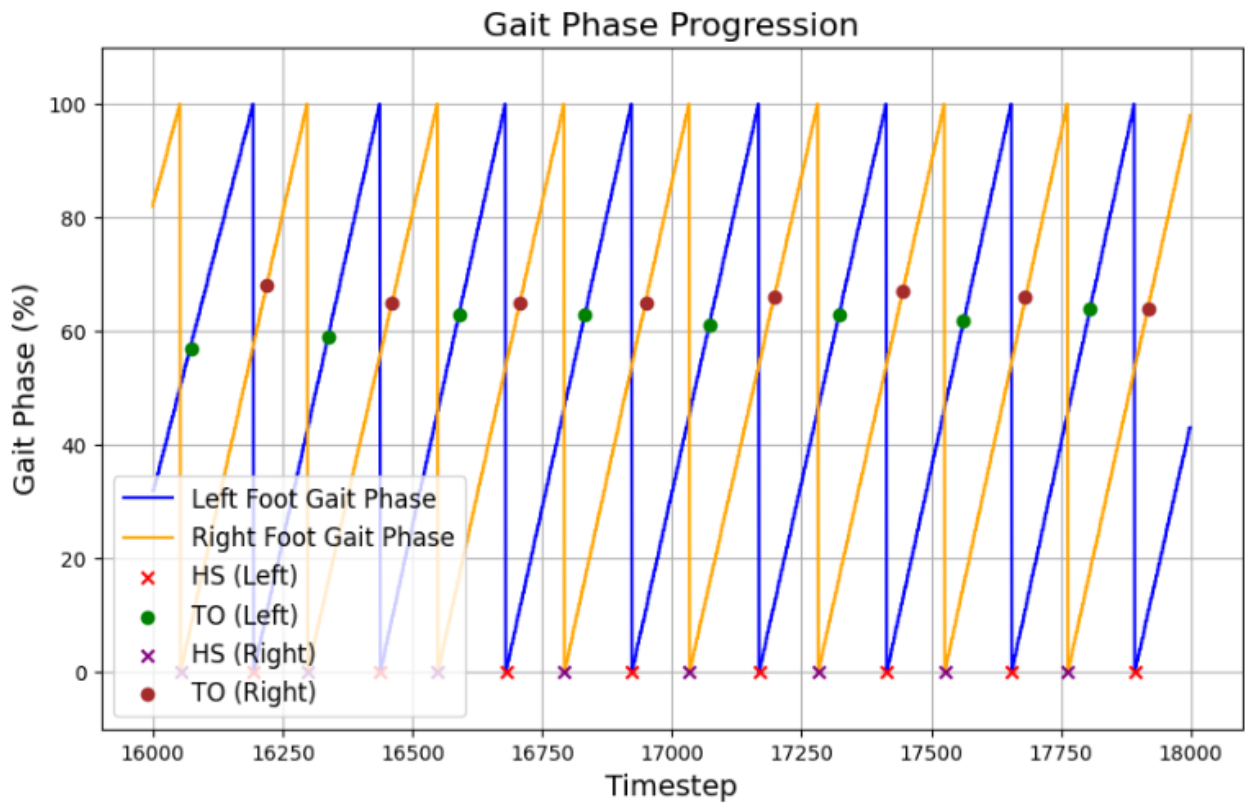


Figure 3.7: Gait phase progression with markers indicating gait events.

$$x_1 = \sin(GP \times 2\pi) \quad (3.1)$$

$$x_2 = \cos(GP \times 2\pi) \quad (3.2)$$

where GP is the gait phase percentage scaled to the range $[0, 1]$.

Thus, the gait phase segment shown in Figure 3.3 is represented in continuous form by the sine and cosine waves plotted in Figure 3.4 below.

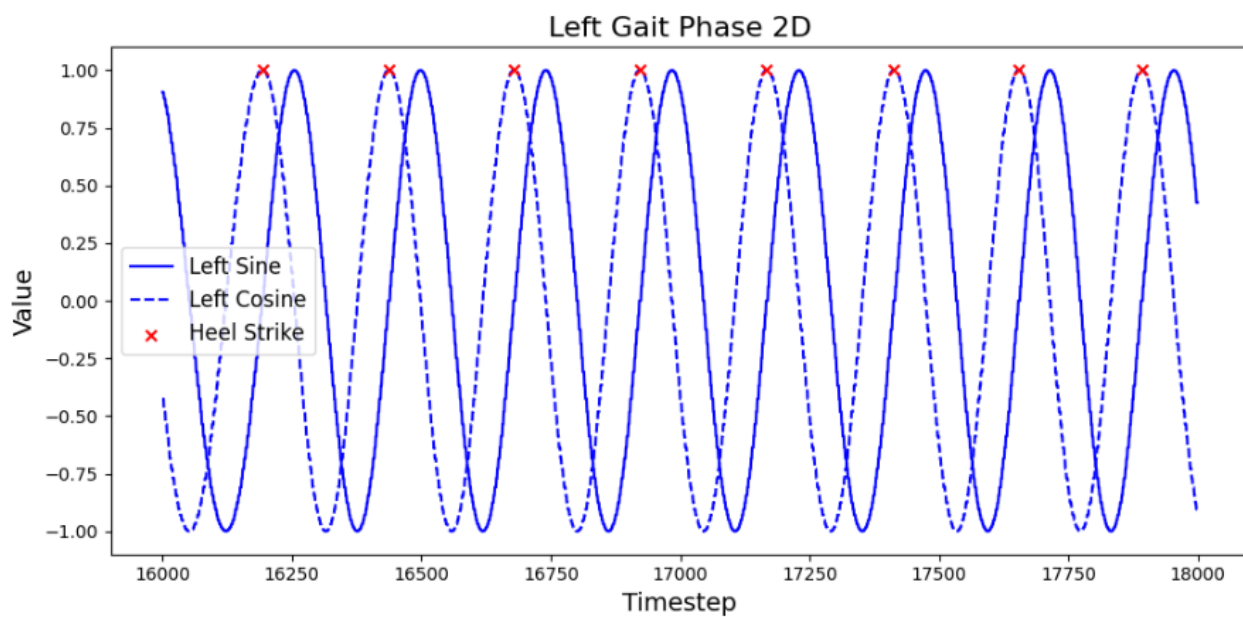


Figure 3.8: Gait phase for the left foot in continuous form with markers for heel strikes.

3.3.2 Model Architecture

The gait phase prediction model architecture consists of two gated recurrent unit (GRU) layers with 256 nodes to process the time series data, followed by a fully connected layer with the same hidden state size, all implemented using the PyTorch library [17]. The tanh

activation function,

$$\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}} \quad (3.3)$$

is applied to the model output to constrain the predictions to the range $[-1, 1]$, as this is the range of the gait phase sine and cosine components produced by the 2D transform [18]. The model was trained using the L1 loss function, which measures the mean absolute error (MAE) of the predictions, making it robust to noisy data and providing an easily interpretable performance metric.

$$MAE = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (3.4)$$

The model was trained for up to 25 epochs on each data file. To prevent the model from overfitting to the training data, early stopping was implemented, which would halt training on a dataset when the validation loss failed to improve for a specified number of epochs.

Additionally, the AdamW algorithm [19] was used for optimization, with a learning rate of $1e-5$, weight decay $1e-6$, and batch size 16. The hyperparameters used to train the model were chosen based on the results of a hyperparameter search using the Tree-structured Parzen Estimator sampler by Optuna [20], [21]. The hyperparameters explored in this study included the hidden size, number of GRU layers, learning rate, batch size, and weight decay.

3.3.3 Post-Processing

The predictions output by the model were post-processed with the application of a Savitzky-Golay filter and subsequent transformation back into percentages from the sine and cosine

Table 3.1: Hyperparameter Search Ranges

Hyperparameter	Range Tested
Learning Rate	$[10^{-5}, 10^{-3}]$
Hidden Size	$\{64, 128, 256, 512\}$
Batch Size	$\{16, 32, 64, 128\}$
Number of Layers	$\{1, 2, 3\}$
Weight Decay	$[10^{-6}, 10^{-4}]$

components used for training. The Savitzky-Golay filter was chosen for its ability to smooth results while preserving the important features of the data by fitting a polynomial to the data points within a sliding window, making it a suitable choice given the sinusoidal nature of the component gait phase predictions [22]. The smoothing filter, with window size 99 and polynomial order 2, resulted in minor improvements to the accuracy of the final results. It is important to note that this filter introduces a delay as a function of window length which may not be suitable for real-time operation. The resulting performance improvements, typically under 1%, were small enough that removing or reducing the size of the filter window would be of no severe detriment.

After filtering, the inverse 2D transform was applied to convert the sine and cosine components back into percentages for interpretability and input to the foot trajectory prediction model [8]:

$$GP = \frac{\text{atan2}(-x_1, -x_2)}{2\pi} + 0.5 \quad (3.5)$$

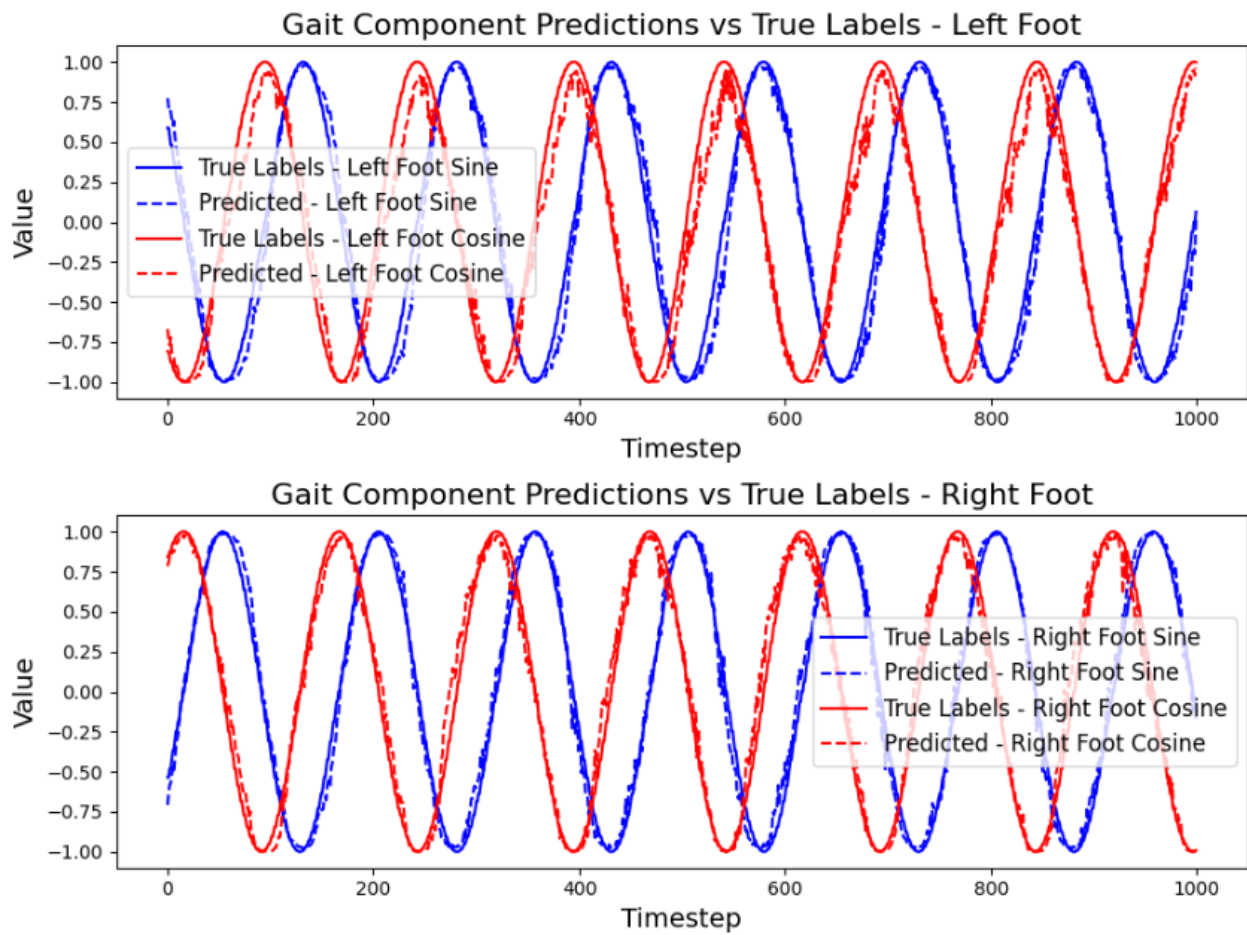


Figure 3.9: Gait phase continuous representation without Savitzky-Golay filter applied.

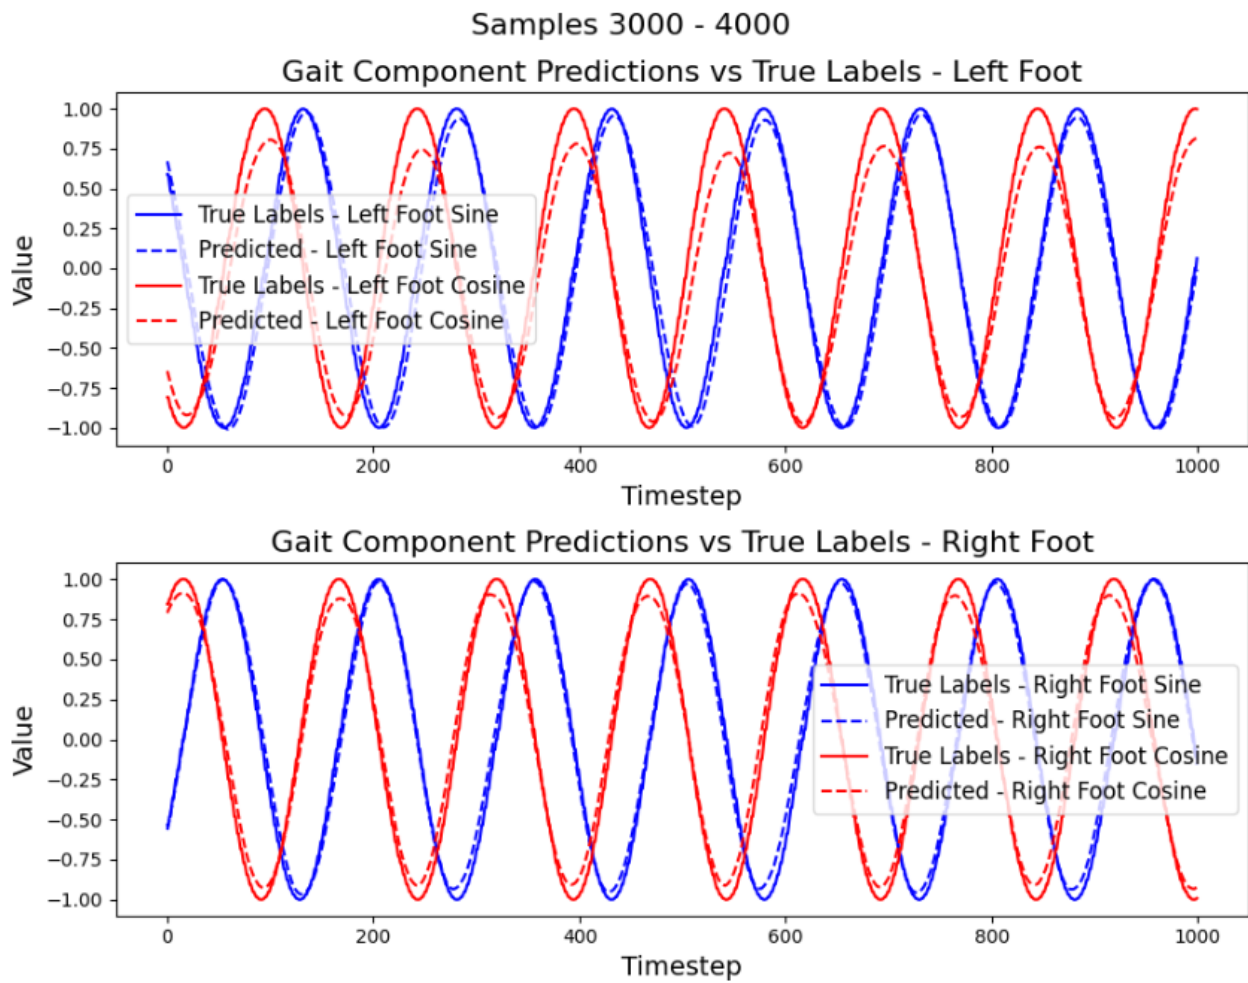


Figure 3.10: Gait phase continuous representation after Savitzky-Golay filter application.

3.4 Foot Trajectory Prediction

3.4.1 Data Processing and Setup

The input feature vector for the foot trajectory model consisted of 29 features, many of which overlapped with those used to train the gait phase model. It contained time, velocity, acceleration, and angle data from the foot IMUs, the gait phase for both feet as predicted by the model, encoder positions providing the rotation of the foot relative to the exoskeleton, and past foot positions in three dimensions, all collected at 100 Hz. The foot trajectory model’s lookback window was longer than that of the gait phase model at 125 samples in length, taken linearly over a 4 second history, as this yielded the best performance after training. The ground truth foot positions were determined using the angles of the foot and exoskeleton, the readings from encoders mounted near the heel, and geometric parameters of the exoskeleton. The exoskeleton components were transformed to a global reference frame defined by the hip bar, and positions over time were derived using rotations and translations. This resulted in an estimate of the exoskeleton’s pose, as seen in Figures 3.11-3.14. An example of the ground truth foot positions resulting from the exoskeleton model is shown in Figure 3.15.

3.4.2 Model Architecture

The foot trajectory prediction model uses 2 GRU layers and one linear layer, all with hidden state size of 256. Like the gait phase model, this network was also trained with early stopping enabled, with the L1 loss function, a batch size of 16, and optimized using the AdamW algorithm with learning rate 1e-5 and weight decay 1e-6. These hyperparameters were determined through the use of the Optuna TPE sampler and experimentation.

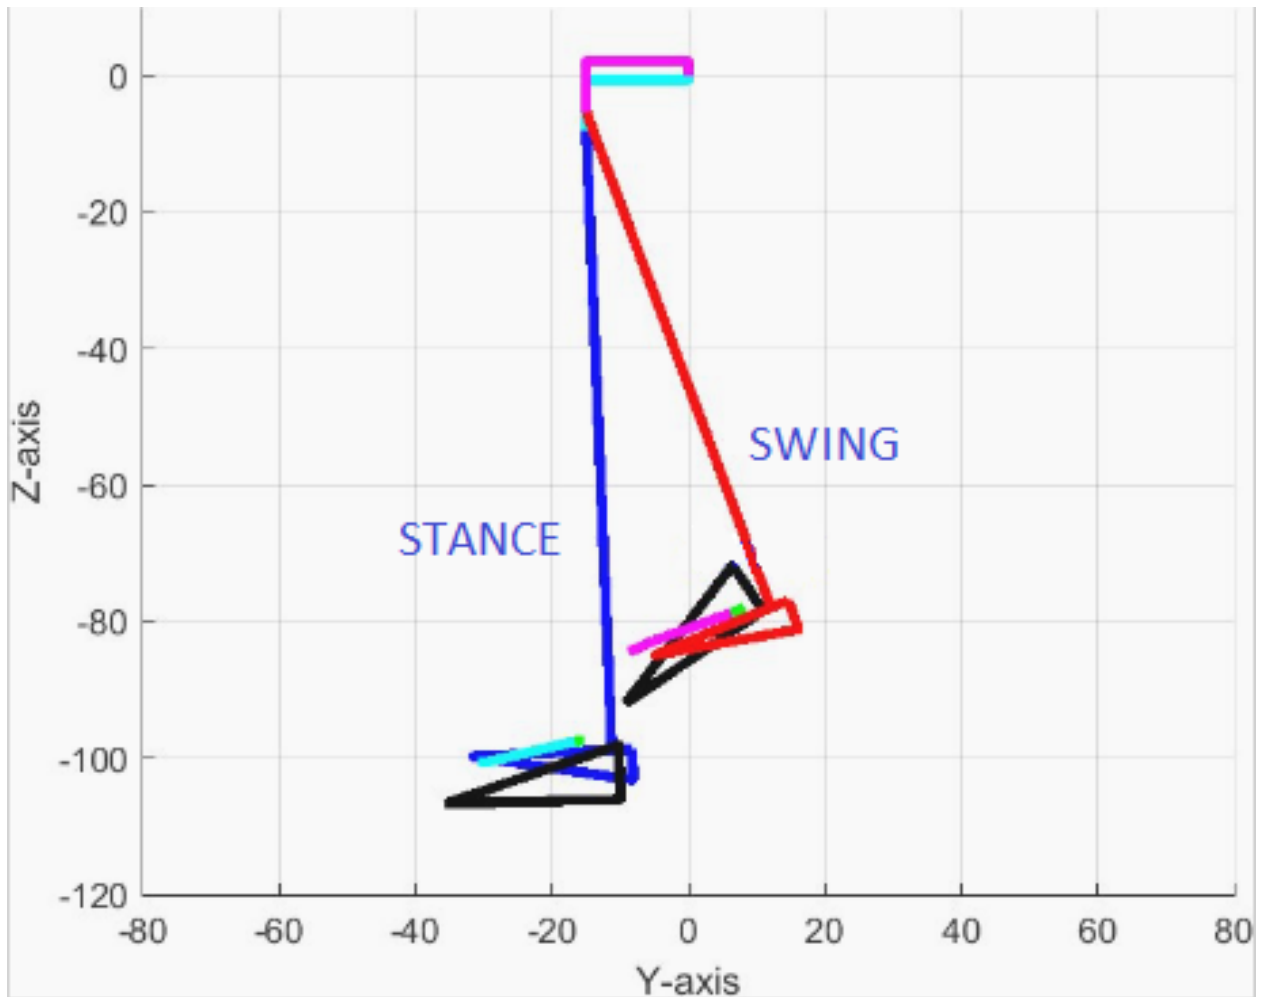


Figure 3.11: Side view ground truth foot position visualization with left foot (red) in early swing phase.

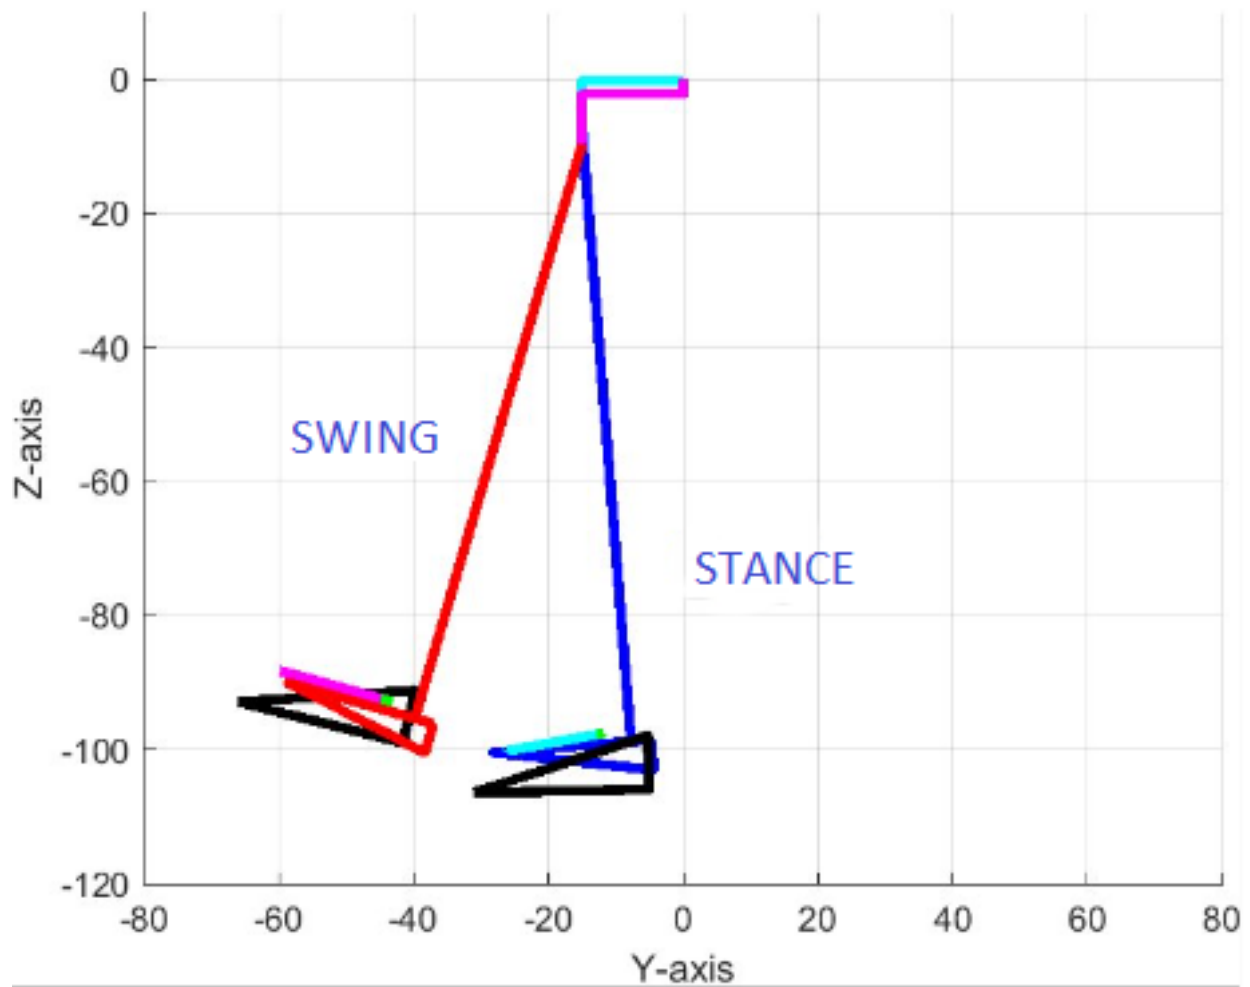


Figure 3.12: Side view ground truth foot position visualization with left foot (red) in late swing phase.

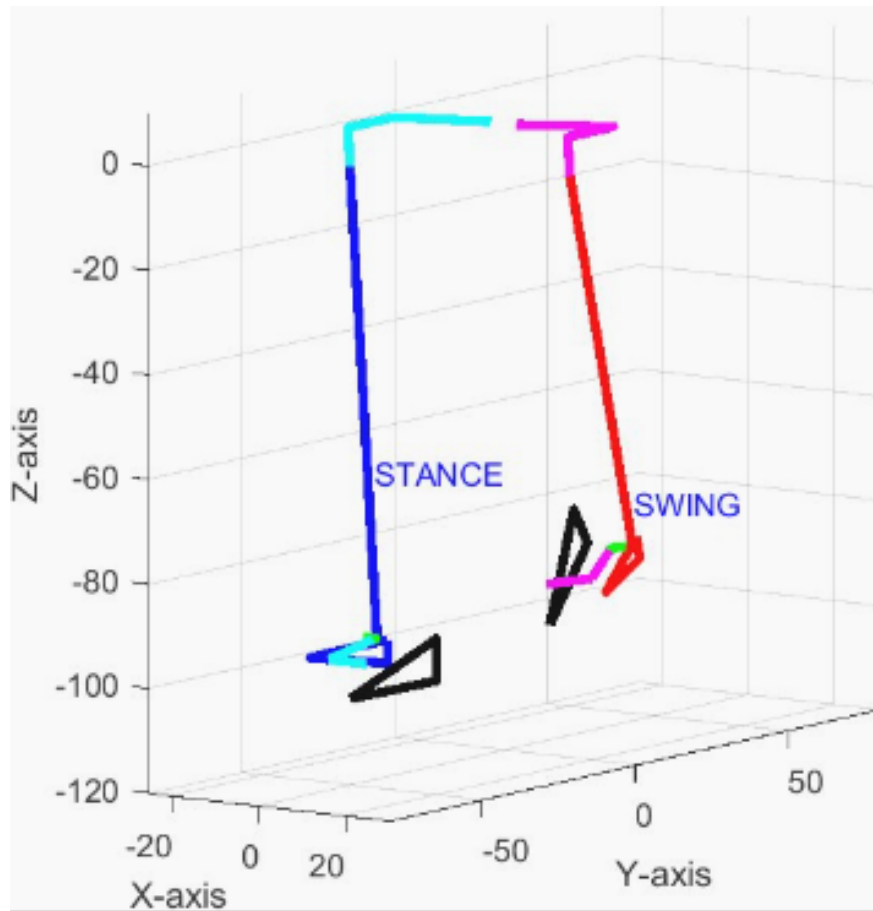


Figure 3.13: Isometric view ground truth foot position visualization with left foot (red) in late swing phase.

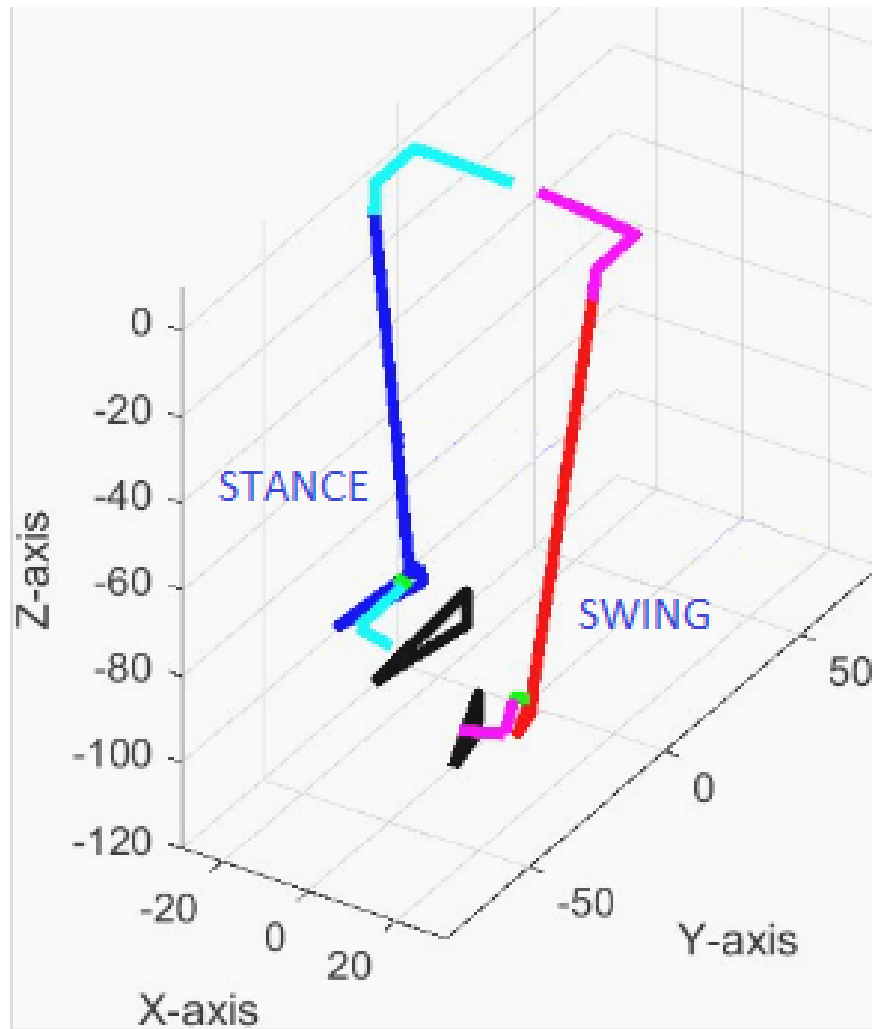


Figure 3.14: Isometric view ground truth foot position visualization with left foot (red) in late swing phase.




Figure 3.15: Ground truth foot positions during walking generated using IMU data.

3.4.3 Post-Processing

A Savitzky-Golay filter was used to smooth the foot position predictions. In contrast to the filter applied to the gait phase predictions, this filter used a smaller window of length 15 and polynomial order 2 to accommodate the more irregular patterns in the output, preserving short-term variations while still reducing noise.

3.5 Lookback Window Sampling

The input feature lookback window was constructed using two different algorithms for sample selection. The first, the linear sampling method, constructed the lookback window by selecting the desired number of samples from a given history length using a fixed step size between each sample. For instance, choosing 100 samples from a history length of 2 seconds would require a step size of 2; thus, the algorithm would add every other sample to the lookback window. The logarithmic lookback window, on the other hand, prioritized a greater concentration of points nearer to the current timestep in order to emphasize recent data. The selected values were spaced logarithmically between the current timestep and the previous points up to the given history length. The rationale for testing the logarithmic sampling algorithm was that recent data would be more strongly correlated with future gait events and movements; however, having sparse samples from more distant history would still expose long-term trends.

Chapter 4

Results

4.1 Metrics

The dataset consisted of 21 files of $32,430 \pm 26,711$ timesteps, with a range from 3,754 to 105,648, containing walking data collected from sensors mounted to the exoskeleton and the subject, although only 20 were used to train the foot trajectory model due to missing data in one file. To evaluate the gait phase model’s performance, we used a reconstructed MAE metric which subtracts from 1 any errors in the percent predictions exceeding 0.5:

$$e_i = |y_i - \hat{y}_i| \tag{4.1}$$

$$e'_i = \begin{cases} e_i, & \text{if } e_i \leq 0.5 \\ 1 - e_i, & \text{if } e_i > 0.5 \end{cases} \tag{4.2}$$

$$MAE_{\text{recon}} = \frac{1}{n} \sum_{i=1}^n e'_i \tag{4.3}$$

This adjustment was used to account for wrapping effects caused by converting periodic target values to percentages. Examination of the predictions in comparison to the targets showed that any false reduction of error not due to wrapping would be minimal as most

genuine errors were significantly smaller than the 0.5 threshold used to determine wrapping errors.

4.2 Gait Phase Prediction

We employed a repeated hold-out method to evaluate model performance on unseen data. The model was trained on 80% of the collected data, validated on an additional 10%, and tested on the remaining 10% that the model did not see previously. This process was repeated 5 times with test and validation files selected at random. No test files were repeated, and the model was tested on data from each subject at least once over the full range of walking speeds, with periods of both constant-speed and dynamic-speed walking. A mean error of $1.92\% \pm 4.35\%$ was obtained using this hold-out technique to evaluate the model on a diverse sample of the dataset, ensuring the model was tested on data from each subject and at varying walking speeds to avoid bias. The high standard deviation relative to the mean is due to outliers in the long tail seen in the histograms in Figures 4.1 and 4.2. The median absolute deviation (MAD) of the same data was calculated as 0.35%, given by the formula:

$$\text{MAD} = \text{median}(|x_i - \text{median}(x)|) \quad (4.4)$$

The model also exhibited the ability to adapt quickly to changes in walking, taking few or no transition strides to closely estimate the gait at the start or end of a walking cycle. In all cases, the first true stride was estimated within 5% or less of the ground truth, but the model often predicted small oscillations just prior to the first step. These frequently corresponded to real patterns in the input data which were simply not classified as full steps in the ground truth generation.

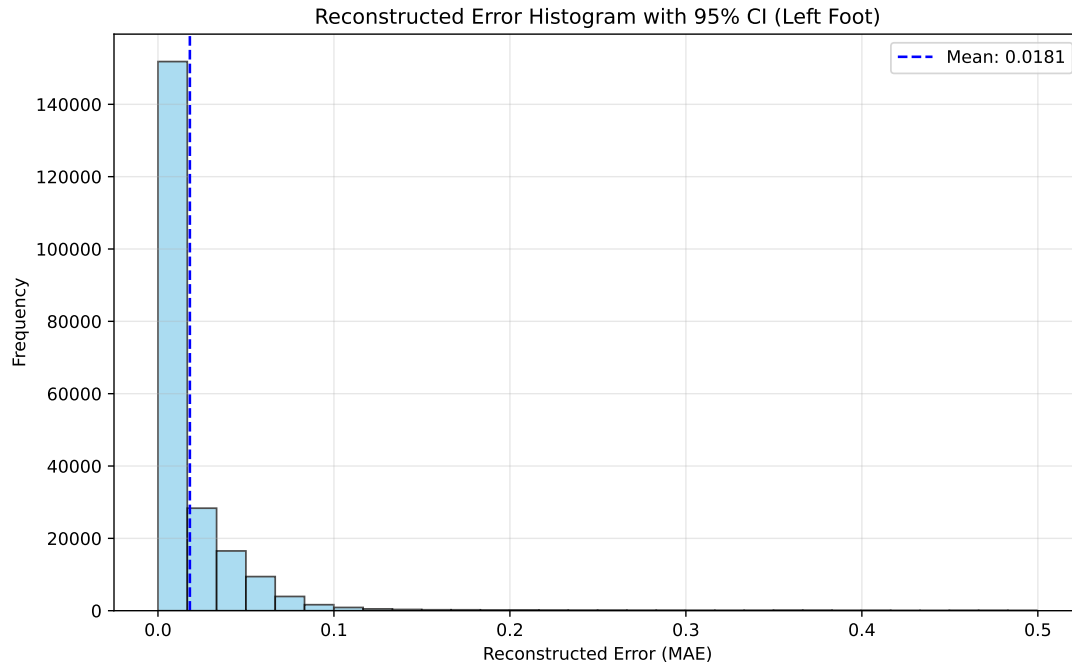


Figure 4.1: Left foot histogram showing the MAE of gait phase predictions.

The transition shown in Figures 4.5 and 4.6 was predicted with 1.38% error over the first three strides. Some transitions exhibited higher error due to imperfect target values and variations caused by the aforementioned input feature patterns—such as the case in Figures 4.7 and 4.8.

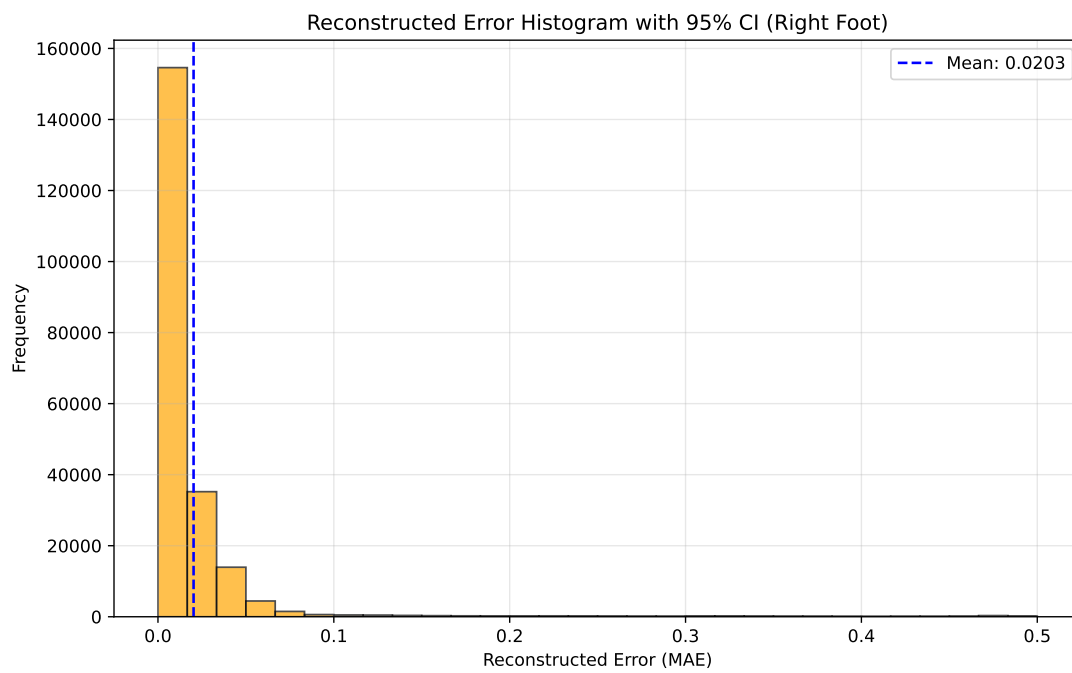


Figure 4.2: Right foot histogram showing the MAE of gait phase predictions.

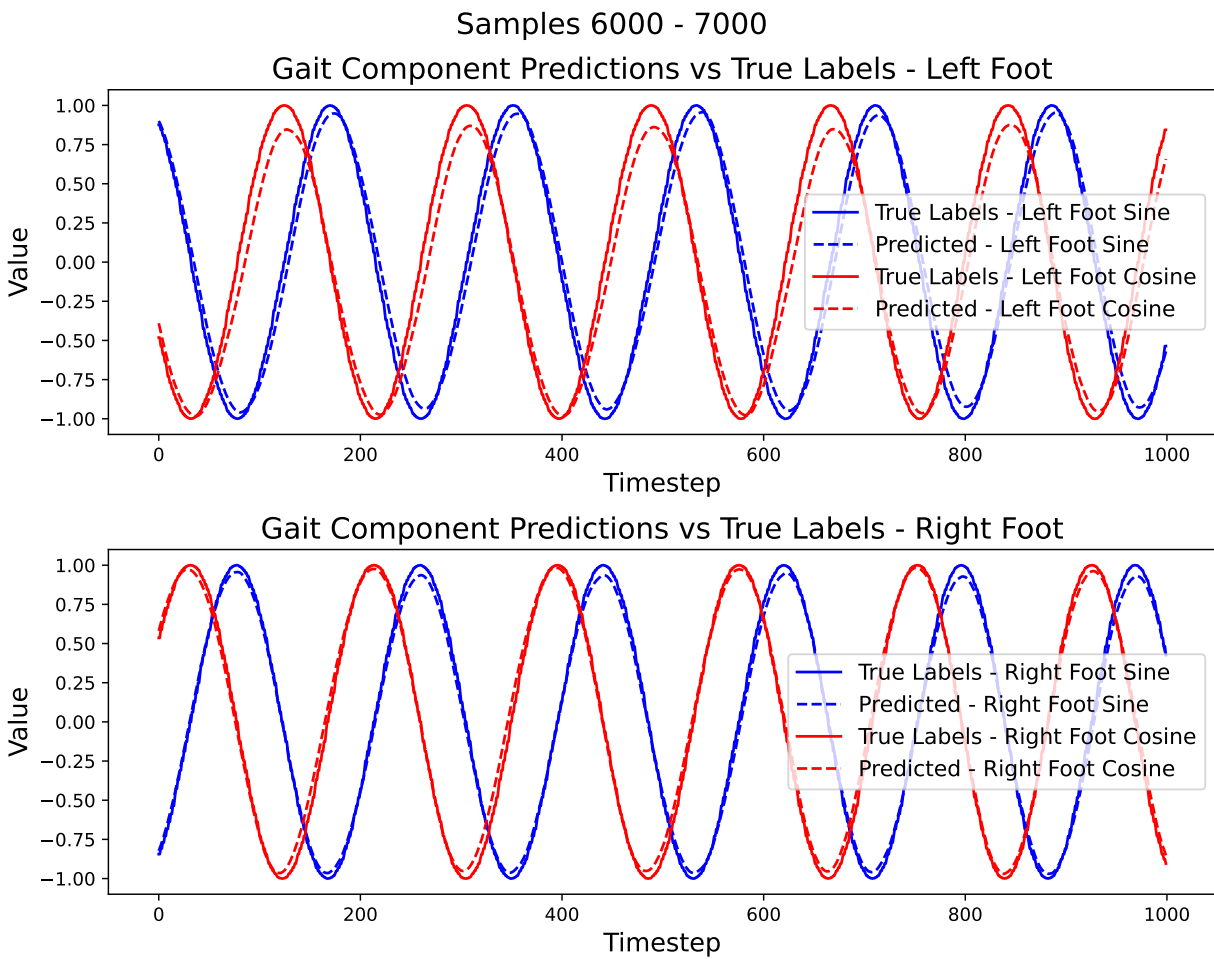


Figure 4.3: Component predictions during a period of steady-state walking.

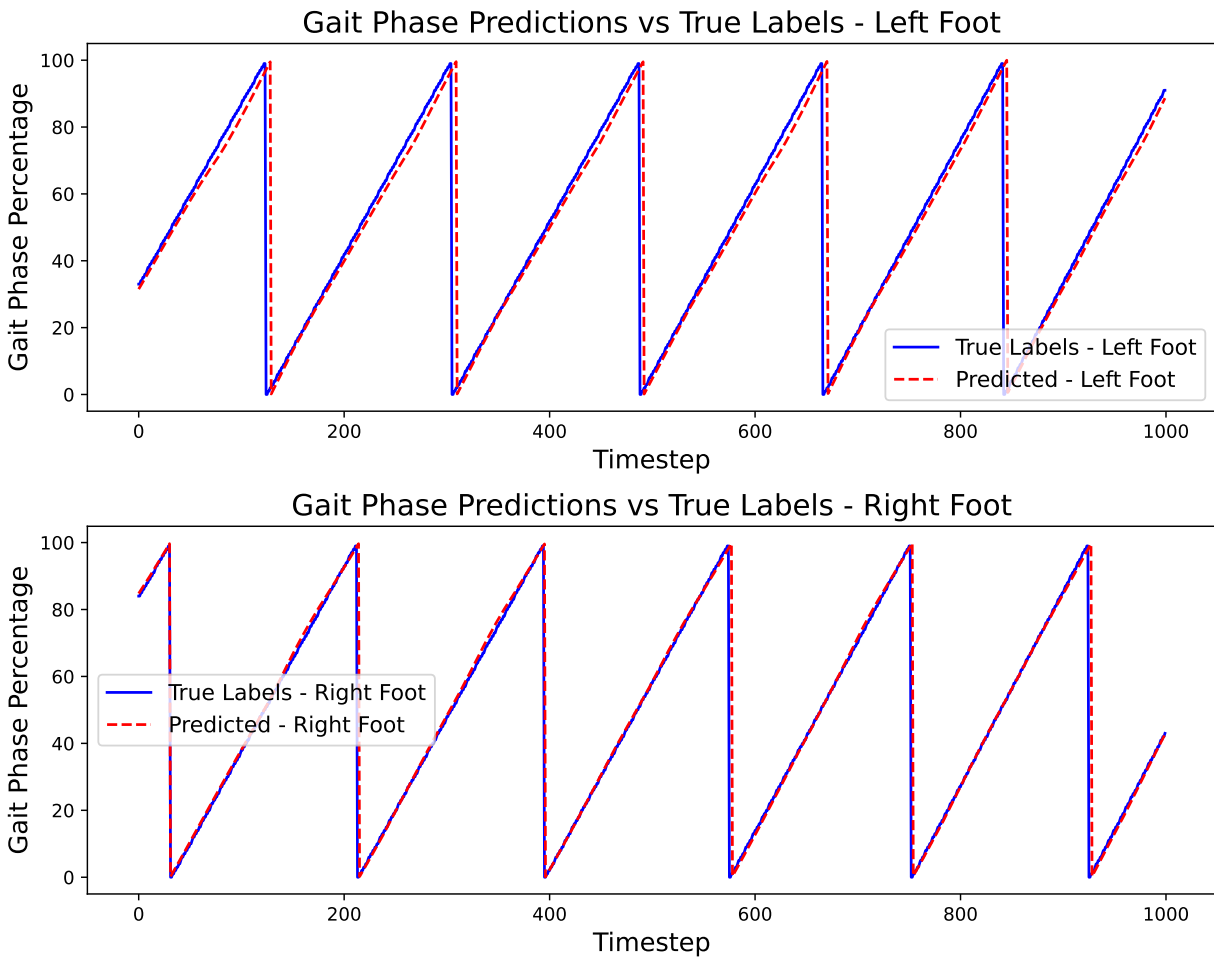


Figure 4.4: Percent predictions during a period of steady-state walking. The MAE taken over this segment was 2.05% for the left foot and 0.67% for the right foot.




Figure 4.5: Transitional gait phase component predictions at start of walking.

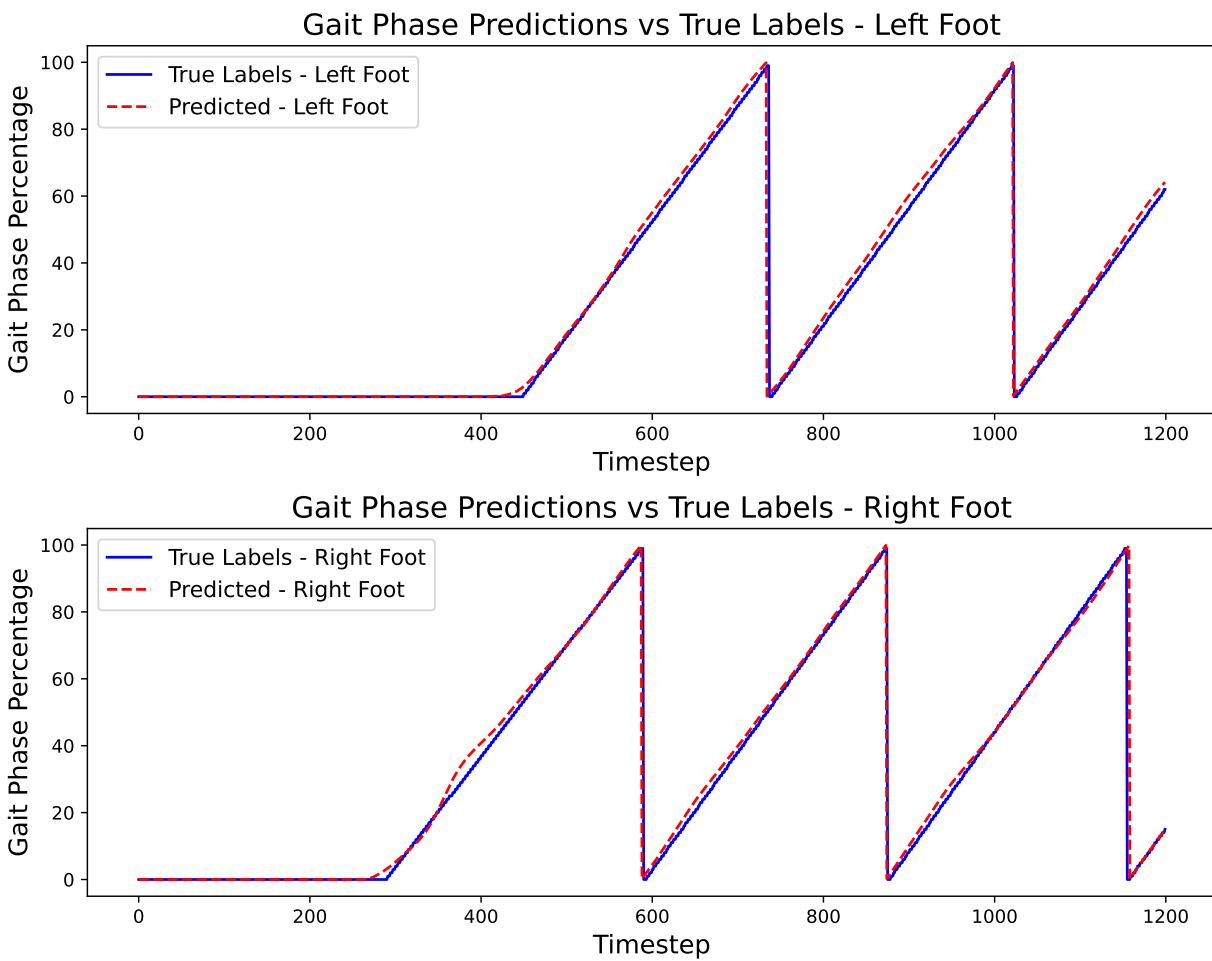


Figure 4.6: Transitional gait phase percent at start of walking.

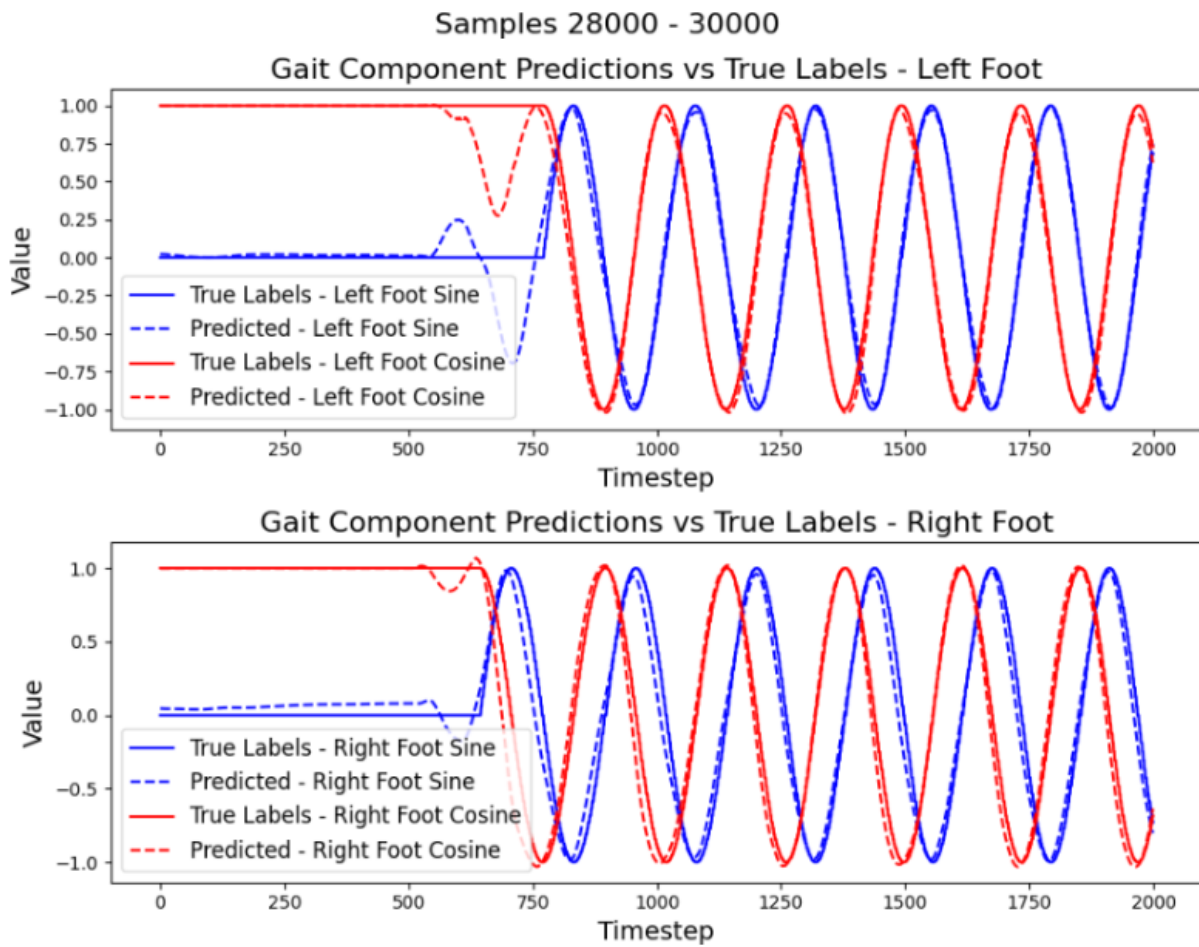


Figure 4.7: Transitional gait phase component predictions at start of walking.

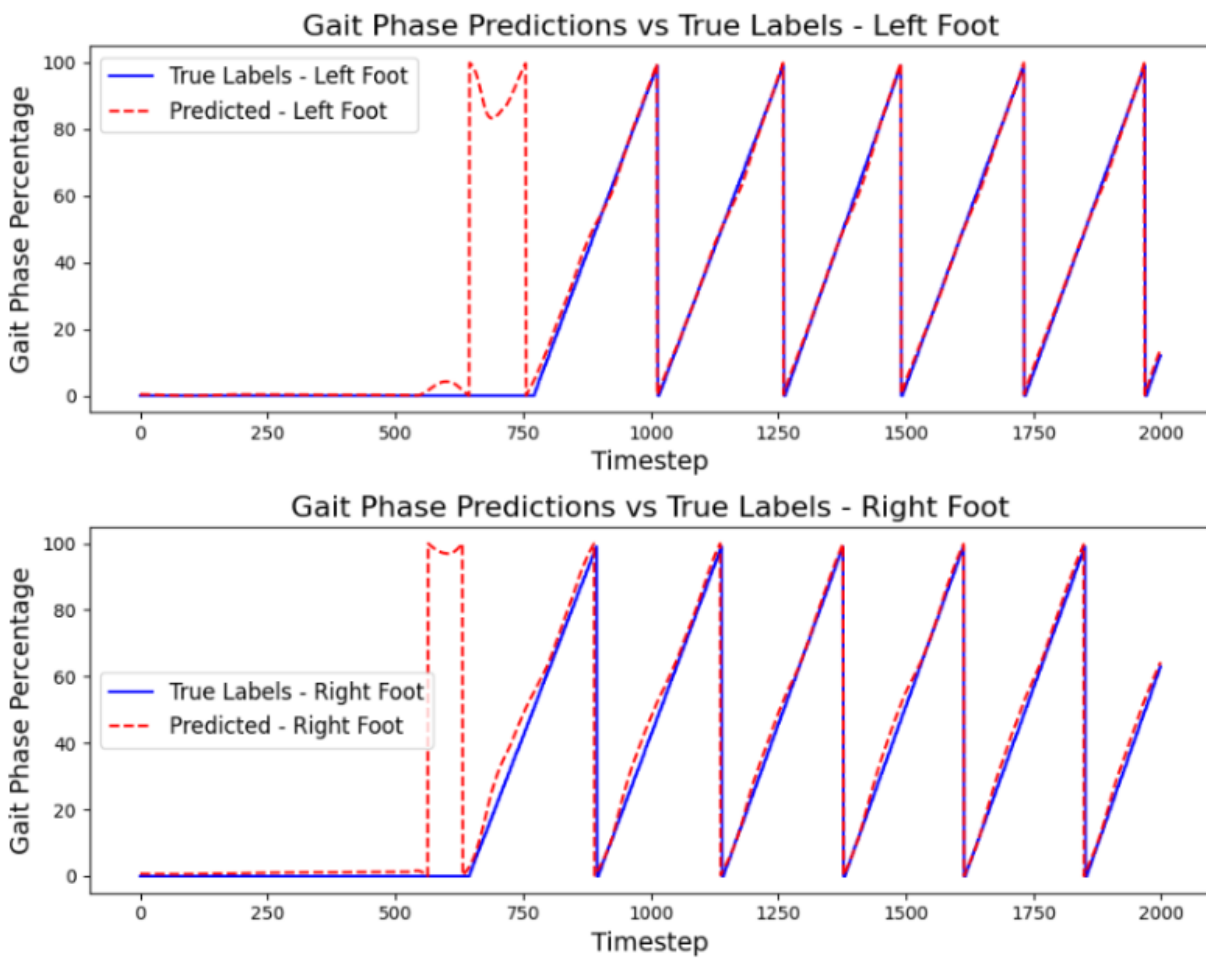


Figure 4.8: Transitional gait phase percent at start of walking.

4.3 Foot Trajectory Prediction

The foot trajectory prediction model was evaluated in the same way as the gait phase model using a repeated pair hold-out strategy. To accommodate the time delay of the exoskeleton motors, we evaluated the model’s performance for predictions from 0.10s to 2s into the future. The results show a gradual decline in accuracy the farther into the future the model tries to predict, which aligns with the results obtained in [23], although we tested a range of more distant future timesteps. The 5-fold hold-out method yielded a mean absolute error of 2.85 ± 2.89 cm predicting 10 timesteps in advance, 3.29 ± 2.82 cm predicting time $t+25$, 4.15 ± 4.12 cm at $t+50$, 5.33 ± 5.46 cm at $t+100$, and 6.92 ± 6.56 cm at $t+200$ in cross-subject cross-velocity testing. Additionally, the model demonstrated improved accuracy over all but the 200 sample prediction horizon when compared to a model trained without previous foot positions included in the input feature vector, shown in Appendix A.3.

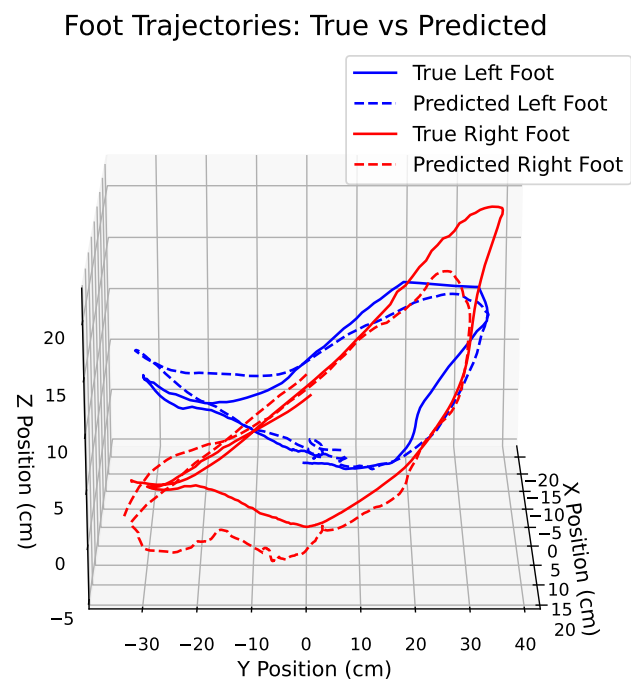


Figure 4.9: Foot trajectory predictions and ground truth over approximately one foot step, predicted 0.10 seconds in advance.




Figure 4.10: Histogram showing the combined prediction error of the left and right foot positions 10 timesteps in advance.




Figure 4.11: Histogram showing the combined prediction error of the left and right foot positions 25 timesteps in advance.




Figure 4.12: Histogram showing the combined prediction error of the left and right foot positions 50 timesteps in advance.

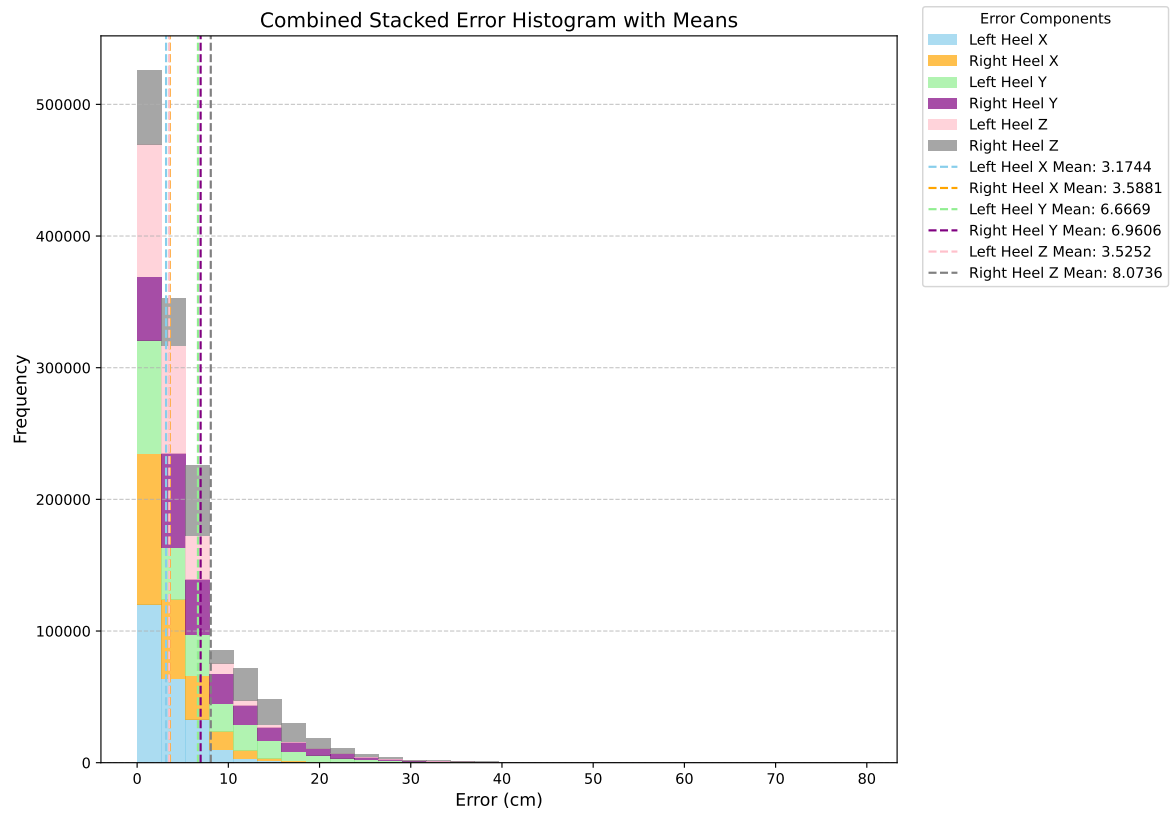


Figure 4.13: Histogram showing the combined prediction error of the left and right foot positions 100 timesteps in advance.

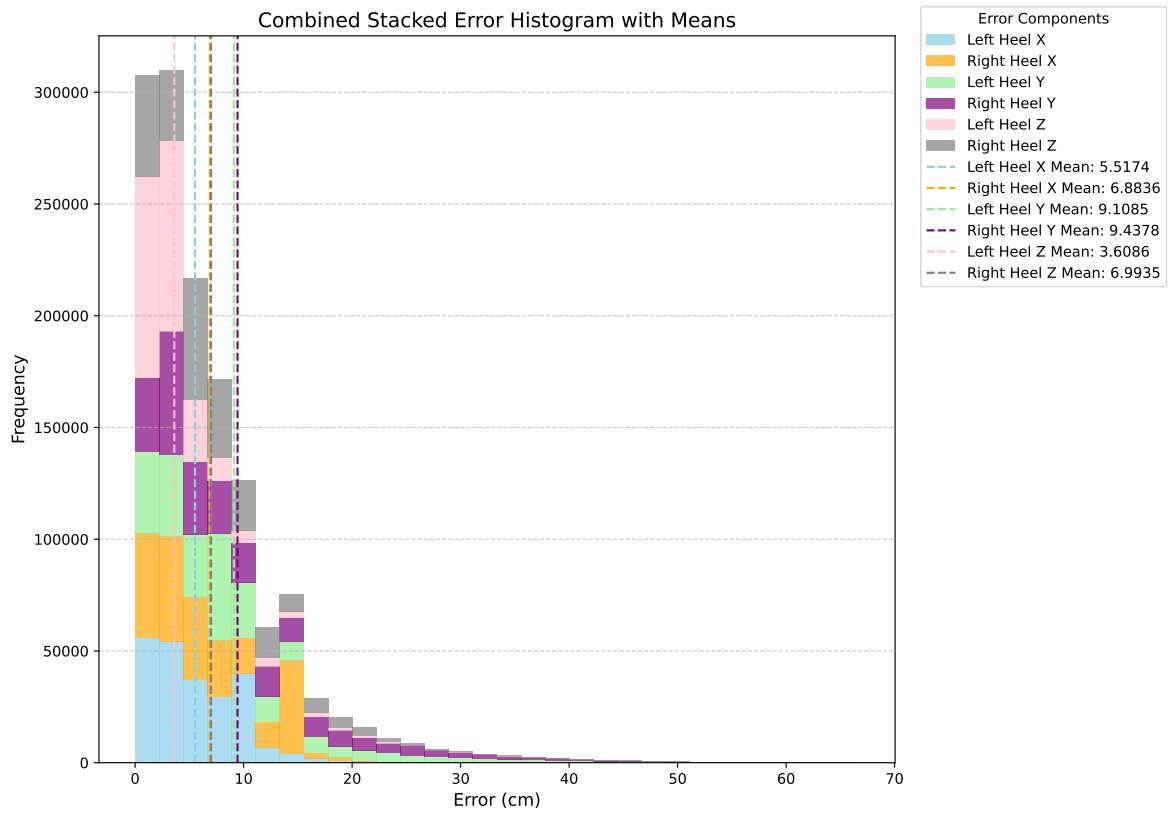


Figure 4.14: Histogram showing the combined prediction error of the left and right foot positions 200 timesteps in advance.

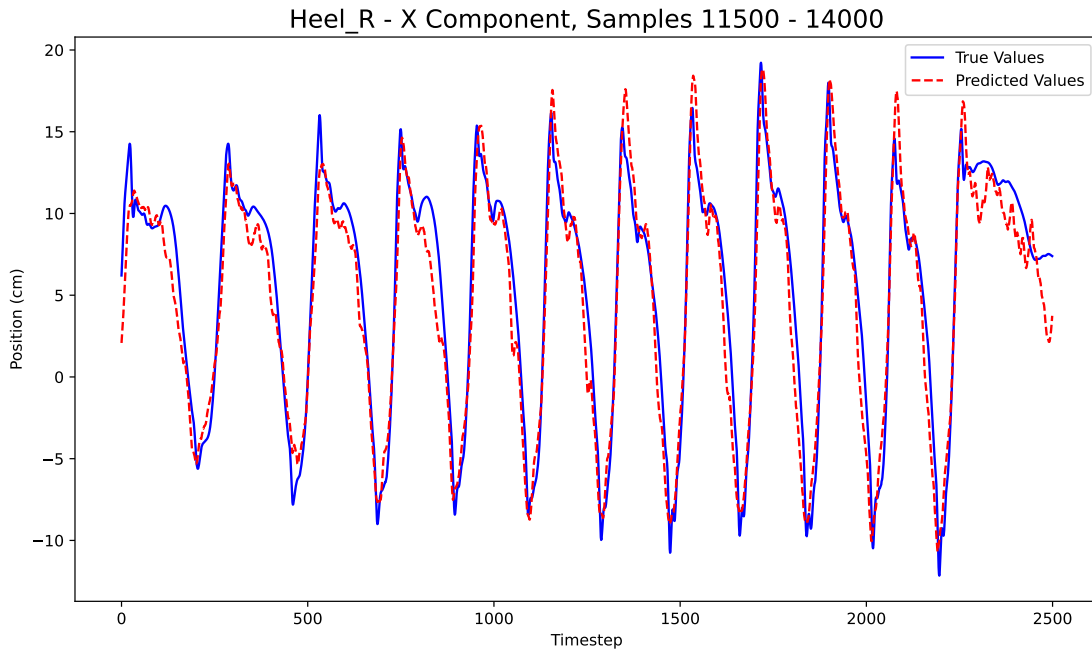


Figure 4.15: Right foot position predictions in the x-direction (side-to-side) predicted 0.1 seconds into the future.




Figure 4.16: Right foot position predictions in the y-direction (forwards and backwards) predicted 0.1 seconds into the future.




Figure 4.17: Right foot position predictions in the z-direction (up and down) predicted 0.1 seconds into the future.

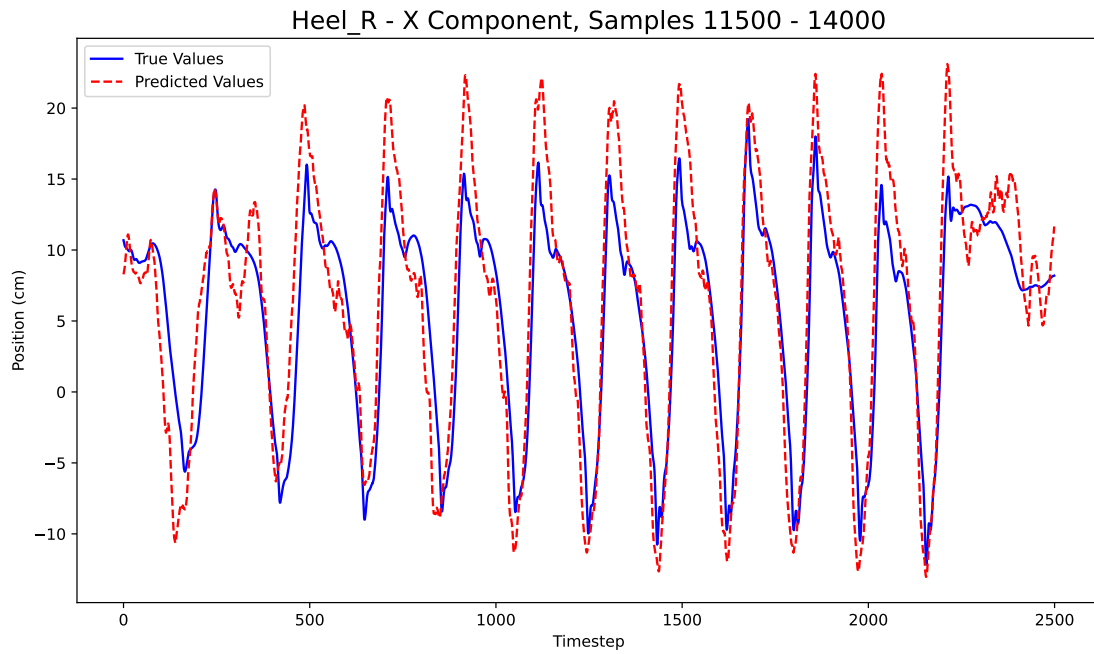


Figure 4.18: Right foot position predictions in the x-direction (side-to-side) predicted 0.5 seconds into the future.




Figure 4.19: Right foot position predictions in the y -direction (forwards and backwards) predicted 0.5 seconds into the future.




Figure 4.20: Right foot position predictions in the z -direction (up and down) predicted 0.5 seconds into the future.




Figure 4.21: Right foot position predictions in the x-direction (side-to-side) predicted 1 second into the future.




Figure 4.22: Right foot position predictions in the y-direction (forwards and backwards) predicted 1 second into the future.




Figure 4.23: Right foot position predictions in the z-direction (up and down) predicted 1 second into the future.




Figure 4.24: Right foot position predictions in the x-direction (side-to-side) predicted 2 second into the future.




Figure 4.25: Right foot position predictions in the y-direction (forwards and backwards) predicted 2 second into the future.




Figure 4.26: Right foot position predictions in the z-direction (up and down) predicted 2 second into the future.

4.4 Lookback Window

We found that the effects of window size and lookback history length differed across datasets. This was likely due in part to variation in walking speed, since slower speeds, which result in longer gait cycles, benefit more from a longer lookback window that can capture a greater portion of the gait cycle. In most cases, performance differences between the linear and logarithmic sampling methods were not significant and often inconsistent, although the logarithmic method performed slightly better when the lookback history was long relative to the number of samples being selected. Even so, the average difference was a fraction of 1% and thus not considered significant. Due to the computational simplicity and negligible performance differences, the linear lookback method was chosen to train the final models.

Chapter 5

Discussion

The goal of this study was to investigate the use of a deep learning-based approach for gait phase estimation and foot trajectory prediction in the control of a lower-body exoskeleton. Our implementation leverages two GRU networks; the first model uses IMU data collected from the exoskeleton during subject walking to estimate gait phase as a percentage while the second includes in its input features IMU data, encoder readings, and the gait phase predictions output by the first model to predict future foot trajectory in three dimensions. Our results demonstrate that the proposed models can accurately predict gait phase and foot trajectory across dynamic walking conditions, which could enhance exoskeleton control by providing predictions in real-time, mitigating the effects of communication delays onboard the exoskeleton that would otherwise cause lag.

The resulting models produce promising results for both gait phase estimation and foot trajectory prediction. Kang et al. [5] achieved average RMSE of $5.22 \pm 0.81\%$ predicting gait phase during constant walking at changing speeds in the range of 0.6 m/s to 1.1 m/s using IMUs and FSR sensors, and $4.83 \pm 0.62\%$ at constant speed, whereas our model was able to achieve a mean RMSE value of 3.96% and 1.92% MAE under dynamic walking conditions across speeds ranging from approximately 0.22 m/s to 0.67 m/s that contained transitions between starting and stopping. A lower error of $1.67 \pm 1.36\%$ was attained in [16] during slow speed steady-state walking (0.5 m/s) and $1.45 \pm 1.47\%$ at a higher speed of 1.5 m/s using an LSTM model. [8] managed 3.0% RMSE with their base trained model. Although this is

lower than our error measurements, our model still has merit in its ability to handle starting and stopping transitions with high accuracy. Additionally, our GRU approach achieved these results using a shorter sequence length of only 100 in comparison to Kang’s and Lee’s [16] 300. This can make the model less computationally intensive and reduce the hardware overhead required for training and inference.

Wang et al. [6] proposed an approach to address the problem of handling walking and standing transitions using a system composed of two adaptive oscillators which showed improvement over conventional AO solutions. They were successful in reducing the average required transition strides during interval walking of conventional AO from 4.00 ± 0.88 to 2.63 ± 0.50 with their specialized gait phase AO model, and during constant walking speed changes from 4.86 ± 0.88 to 3.00 ± 0.58 . Our deep learning approach is able to dynamically adapt to speed changes during walking and from standing to walking with 0 transition strides.

In [7], Chen found a mean absolute difference error of 13.51 cm across subjects and walking velocities using a probability distribution model with IMU data to predict foot placement in the early swing phase. In comparison, we calculated a mean absolute difference of 6.92 ± 6.56 cm in our least accurate model which predicted foot positions in three dimensions 2 seconds in advance, which, based on the walking speeds tested in our study, could mean predicting the next foot step from as early as the previous step. Lee achieved steady-state walking foot placement MDE of 3.32 cm with the base-trained model (i.e., without fine-tuning) within the first 56% of the gait phase and 4.59 ± 0.91 cm when prediction was made just 1% into the gait phase. This exceeds the accuracy of our model at comparable prediction horizons, but does not account for walking and standing transitions.

[11] found 6.0% NRMSE on predicted angular foot positions in the sagittal plane 10 samples (at 100 Hz) into the future. With the same prediction horizon, our model, which output predictions in all three dimensions, achieved RMSE over the same direction of motion of 5.62

cm. The average range of ground truth values in the forward and backward direction along the sagittal plane in our dataset was 72 cm, which would put 5.62 cm at roughly 7.81% of the step length. While this error is higher, Dey did not account for pauses in walking.

The results of this study demonstrate the feasibility of a GRU network for foot trajectory prediction in the control of a walking exoskeleton. The ability of our model to make accurate predictions early in the gait phase, and even prior to the predicted step, have significant implications for exoskeleton control. By anticipating foot motion in real-time, our approach makes it possible to reduce or even eliminate lag resulting from communication delays between exoskeleton components.

In contrast to the bulk of prior work, which primarily focuses on foot placement prediction during steady-state walking, our research presents findings on the prediction of foot positions over time not only across constant and variable walking speeds, but over transitions between walking and standing. It is critical to account for these transitional periods if walking exoskeletons are to be made practical for everyday use, as people frequently shift between walking and standing states during daily activities.

Future work should consider expanding our data collection to include more subjects for better generalization to different gaits, a wider range of walking speeds, and diverse ground conditions. Our training data was gathered only on treadmill walking in a straight line, but it would be useful and more practical to examine performance differences from overground walking that incorporates turns and obstacles. While the model can predict foot positions in advance, it is also important to evaluate inference time, especially on embedded hardware, to determine the feasibility of real-time deployment of the model on the exoskeleton. Additionally, the results presented in this paper were obtained by base training on data from all subjects rather than pre-training and fine-tuning on the target subject, but our model could likely be improved further with the use of fine-tuning.

Chapter 6

Conclusions

Exoskeletons and wearable technology are becoming ever more present and integrated into daily life as well as the medical field. They have the potential to make everyday tasks easier, improve physical rehabilitation outcomes, and restore mobility and bodily autonomy for people with disabilities. Towards this end, predicting human motion in advance of its initiation is crucial in order for assistive devices to operate naturally with the user. We proposed an approach leveraging two GRU models for gait phase and foot trajectory prediction which achieved, respectively, $1.92\% \pm 4.35\%$ and 4.21 ± 3.59 cm on foot positions with prediction horizon $t+10$, 4.75 ± 4.96 cm with prediction horizon $t+25$, 4.76 ± 4.40 cm at $t+50$, 5.63 ± 5.16 cm at $t+100$, and 6.60 ± 6.41 cm at $t+200$ in cross-subject cross-velocity testing.

We have demonstrated in this paper that a GRU-based approach can effectively anticipate gait phase and foot trajectory across dynamic walking conditions given data collected from IMU sensors attached to the exoskeleton. This work brings us closer to providing safer, more intelligent assistive exoskeletons that can support user needs at home, in the workforce, and on some of life's most challenging journeys.

Bibliography

- [1] S. M. Bruijn and J. H. van Dieën. Control of human gait stability through foot placement. *Journal of The Royal Society Interface*, 15(143), 2018.
- [2] Lei Zhang and Chenglong Fu. Predicting foot placement for balance through a simple model with swing leg dynamics. *Journal of Biomechanics*, 77:155–162, 2018.
- [3] N. Passalis A. Tsantekidis and A. Tefas. *Deep Learning for Robot Perception and Cognition*. Academic Press, 2022.
- [4] Rahul Dey and Fathi M. Salem. Gate-variants of gated recurrent unit (gru) neural networks. In *2017 IEEE 60th International Midwest Symposium on Circuits and Systems (MWSCAS)*, pages 1597–1600, 2017.
- [5] Inseung Kang, Pratik Kunapuli, and Aaron J. Young. Real-time neural network-based gait phase estimation using a robotic hip exoskeleton. *IEEE Transactions on Medical Robotics and Bionics*, 2(1):28–37, 2020.
- [6] Xiangyang Wang, Yue Ma, Chunjie Chen, Zhuo Wang, Sheng Guo, Kin Huat Low, and Xinyu Wu. Effective prediction of gait phase for assisted walking by means of gait-based adaptive oscillators. *IEEE Transactions on Automation Science and Engineering*, pages 1–11, 2024.
- [7] Xinxing Chen, Kuangen Zhang, Haiyuan Liu, Yuquan Leng, and Chenglong Fu. A probability distribution model-based approach for foot placement prediction in the early swing phase with a wearable imu sensor. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 29:2595–2604, 2021.

- [8] Sung-Wook Lee and Alan Asbeck. A deep learning-based approach for foot placement prediction. *IEEE Robotics and Automation Letters*, 8(8):4959–4966, 2023.
- [9] Jingfeng Xiong, Chuheng Chen, Yuanwen Zhang, Xinxing Chen, Yuepeng Qian, Yuquan Leng, and Chenglong Fu. A probability fusion approach for foot placement prediction in complex terrains. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 31:4591–4600, 2023.
- [10] Kevin Tanghe, Friedl De Groote, Dirk Lefeber, Joris De Schutter, and Erwin Aertbeliën. Gait trajectory and event prediction from state estimation for exoskeletons during gait. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 28(1):211–220, 2020.
- [11] Sharmita Dey and Arndt F. Schilling. A function approximator model for robust online foot angle trajectory prediction using a single imu sensor: Implication for controlling active prosthetic feet. *IEEE Transactions on Industrial Informatics*, 19(2):1467–1475, 2023.
- [12] S et al. Yu. Artificial neural network-based activities classification, gait phase estimation, and prediction. *Annals of biomedical engineering*, 51(7):1471–1484, 2023.
- [13] Kuangen Zhang, Haiyuan Liu, Zixuan Fan, Xinxing Chen, Yuquan Leng, Clarence W. de Silva, and Chenglong Fu. Foot placement prediction for assistive walking by fusing sequential 3d gaze and environmental context. *IEEE Robotics and Automation Letters*, 6(2):2509–2516, 2021.
- [14] Xinxing Chen, Zijian Liu, Jiale Zhu, Kuangen Zhang, Yuquan Leng, and Chenglong Fu. Comparison of machine learning regression algorithms for foot placement prediction. In *2021 27th International Conference on Mechatronics and Machine Vision in Practice (M2VIP)*, pages 169–174, 2021.

- [15] Joana Figueiredo, Paulo Félix, Luís Costa, Juan C. Moreno, and Cristina P. Santos. Gait event detection in controlled and real-life situations: Repeated measures from healthy subjects. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 26(10):1945–1956, 2018.
- [16] Jinwon Lee, Woolim Hong, and Pilwon Hur. Continuous gait phase estimation using lstm for robotic transfemoral prosthesis across walking speeds. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 29:1470–1477, 2021.
- [17] PyTorch. *PyTorch Documentation — PyTorch Master Documentation*, 2019.
- [18] PyTorch. *Tanh — PyTorch 1.13 Documentation*, 2022.
- [19] K Ding, N Xiao, and K.-C. Toh. Adam-family methods with decoupled weight decay in deep learning. 2023.
- [20] Optuna. *optuna.samplers.TPESampler — Optuna 3.6.1 Documentation*, 2023.
- [21] Bernd Bischl, Martin Binder, Michel Lang, Tobias Pielok, Jakob Richter, Stefan Coors, Janek Thomas, Theresa Ullmann, Marc Becker, Anne-Laure Boulesteix, Difan Deng, and Marius Lindauer. Hyperparameter optimization: Foundations, algorithms, best practices, and open challenges. *WIREs Data Mining and Knowledge Discovery*, 13(2), 2023.
- [22] S. R. Krishnan and C. S. Seelamantula. On the selection of optimum savitzky-golay filters. *IEEE Transactions on Signal Processing*, 61(2):380–391, January 2013.
- [23] Jennifer K Leestma, Courtney R Smith, Gregory S Sawicki, and Aaron J Young. A data-driven approach to estimate human center of mass state during perturbed locomotion using simulated wearable sensors. *Annals of biomedical engineering*, 52(8):2013–2023, 2024.

Appendices

Appendix A

First Appendix

A.1 Additional Input Features

A total of 37 different features were tested for inclusion in the input feature vector of the gait phase prediction model, including velocity, acceleration, and angle data from the IMUs mounted to the subject's feet and the exoskeleton feet. Ultimately, the data from the exoskeleton sensors was excluded due to feature redundancy which caused the model to see no performance improvement from their inclusion and, in some cases, performance degradation. This was confirmed via permutation feature importance computation, which showed small permutation importance values for the corresponding foot and exoskeleton features.

Similarly, the foot trajectory model was tested with variations of the input feature vector with 41 features. These included the foot and exoskeleton velocity, acceleration, and angles, as well as encoder readings for foot rotation and gait phase predictions. The best performance was achieved with the encoder data, gait phase predictions, and foot velocity, acceleration, and angles, forming a vector of 23 input features. The addition of exoskeleton foot angles had a negligible impact on prediction accuracy and were, therefore, excluded.

A.2 Supplemental Graphs




Figure A.1: Unfiltered gait phase percent predictions plotted alongside ground truth.

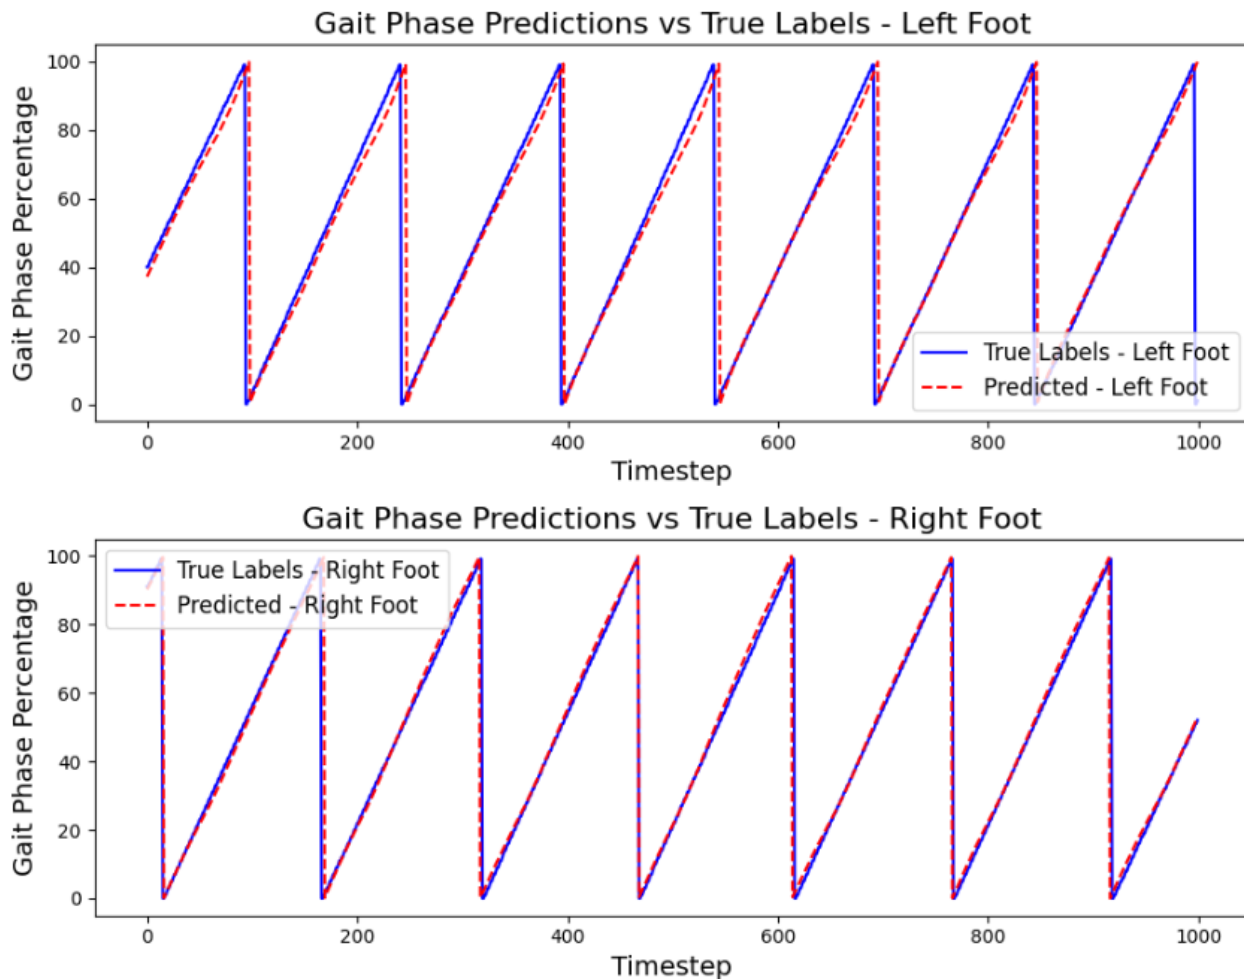


Figure A.2: Gait phase percent predictions after Savitzky-Golay filtering plotted alongside ground truth.

A.3 Early Foot Trajectory Prediction

An initial foot trajectory prediction model was implemented which predicted future foot positions using only IMU data, gait phase, and encoder readings—that is, the model did not have the advantage of knowing previous foot positions directly. This model, using the same repeated hold-out strategy detailed in the results section, achieved MAE of 4.21 ± 3.59 cm predicting 10 timesteps in advance, 4.75 ± 4.96 cm predicting time $t+25$, 4.76 ± 4.40 cm at

$t+50$, 5.63 ± 5.16 cm at $t+100$, and 6.60 ± 6.41 cm at $t+200$ in cross-subject cross-velocity testing.




Figure A.3: Foot trajectory prediction error over increasing prediction horizons.




Figure A.4: Histogram showing the combined prediction error of the left and right foot positions 10 timesteps in advance.




Figure A.5: Histogram showing the combined prediction error of the left and right foot positions 25 timesteps in advance.




Figure A.6: Histogram showing the combined prediction error of the left and right foot positions 50 timesteps in advance.




Figure A.7: Histogram showing the combined prediction error of the left and right foot positions 100 timesteps in advance.




Figure A.8: Histogram showing the combined prediction error of the left and right foot positions 200 timesteps in advance.




Figure A.9: Right foot position predictions in the x-direction (side-to-side) predicted 0.1 seconds into the future.




Figure A.10: Right foot position predictions in the y-direction (forwards and backwards) predicted 0.1 seconds into the future.




Figure A.11: Right foot position predictions in the z-direction (up and down) predicted 0.1 seconds into the future.




Figure A.12: Right foot position predictions in the x-direction (side-to-side) predicted 0.5 seconds into the future.




Figure A.13: Right foot position predictions in the y-direction (forwards and backwards) predicted 0.5 seconds into the future.




Figure A.14: Right foot position predictions in the z-direction (up and down) predicted 0.5 seconds into the future.




Figure A.15: Right foot position predictions in the x-direction (side-to-side) predicted 1 second into the future.




Figure A.16: Right foot position predictions in the y-direction (forwards and backwards) predicted 1 second into the future.




Figure A.17: Right foot position predictions in the z-direction (up and down) predicted 1 second into the future.




Figure A.18: Right foot position predictions in the x-direction (side-to-side) predicted 2 second into the future.




Figure A.19: Right foot position predictions in the y-direction (forwards and backwards) predicted 2 second into the future.




Figure A.20: Right foot position predictions in the z-direction (up and down) predicted 2 second into the future.