Analysis of Dynamic Ankle- Foot Orthotic Performance Through Gait Analysis of Multiple Sclerosis Participant Walking Trials

Data Science Final Report

Julianne Jorgensen & Justin Poh

30 April 2015

SECTION 1: BACKGROUND & INTRODUCTION

Multiple Sclerosis (MS) is a progressive, disabling disease that causes damage to the nervous system and widespread dysfunction[1] [2]. The disease includes sensory, cognitive and motor impairments. The most common motor impairment is difficulty walking[1]. By using statistics to analyze the gait of participants with MS, the changes in gait and common walking patterns can be measured.

However, analyzing gait for participants with MS presents an unusual set of challenges. Often, the walking disability caused by MS is the result of muscle weakness, spasticity, loss of balance, sensory deficit and fatigue[3]. Consequently, participants with MS might demonstrate unusual movements that makes performing analyses of their gait data more difficult. Some examples of such unusual movements include the inability to fully lift one's foot completely off the ground (Toe Drag), the inability to fully support one's foot off the ground (Foot Drop), or the outward swinging of the leg in order to take another step while walking (Leg Circumduction). If a participant demonstrates leg circumduction while walking, they might also lift their hip (Hip Hike) or lean their trunk (Trunk Lean) in order to swing their leg around while walking[3].

Despite these challenges, by correcting and stabilizing deviations in joint alignment, proper bracing can help participants establish a more stable and energy efficient gait. This bracing can be provided by ankle foot orthotics (AFOs) which can improve gait by increasing joint stability and joint alignment[4]. In addition, dynamic AFOs are sometimes used to provide power at the ankle to help MS participants further achieve better gait while walking.

As part of the process of developing such an orthotic and understanding its effects on users, there are several key questions about the performance of the dynamic AFO compared to the normal AFO or the participant's own shoes:

---

[1] http://www.nature.com/nature/journal/v399/n6738supp/full/399a040.html

[2] http://www.hindawi.com/journals/msi/2013/645197/

[3] http://www.nationalmssociety.org/Symptoms-Diagnosis/MS-Symptoms/Walking-%28Gait%29,-Balance-Coordination

[4] http://www.oandp.com/articles/2007-05_05.asp

1. **Does the dynamic AFO enable better symmetry during gait?** An alternative phrasing of this question is: Does the impaired limb better match the healthier limb when the orthotic is dynamic?

2. **Is knee extension greater with the dynamic AFO compared to the shoes and normal AFO conditions?** There is a concern that the dynamic AFO is causing knee overextension compared to baseline standing condition during walking using the dynamic AFO.

3. **Is leg circumduction reduced with the dynamic AFO and normal AFO compared to the participant's own shoes?** Analyzing leg circumduction would compare sideways movement of the hips and knee angles. This requires 3D analysis of all joint angles.

Over the course of this report, we will present our analysis of the data in order to provide answers to these three key questions. We will start by discussing the characteristics of the data and then discuss the method we developed for extracting gait cycles from the data. Then, we will perform analyses on the data and, finally, present the answers we have found to the above questions along with the evidence to support those claims.

SECTION 2: DATA EXPLORATION

The gait data we worked with was originally captured using a motion capture suite called Qualysis[5]. Reflective markers are placed on the participant and the participant walks around a loop several times. As they walk through one particular section of the loop, the locations of the reflective markers placed on the participant are tracked by 8 cameras. Using the data from these cameras, the Qualisys software computes the X, Y, and Z spatial coordinates for each identifiable reflective marker for every frame of video from the cameras to create a 3D visualization of the participant walking. The efficacy of the 3D overlay can then be characterized by what is known as the residual, R. The X, Y, Z, and R coordinates are then compiled for each marker into a data file.

In this section, we discuss the structure of the data, our method for importing that data into Python and validating the data pipeline, as well as some basic characterizations that informed the analyses we performed to answer the three key questions.

SECTION 2.1: DATA STRUCTURE

The data for each participant varies in size from 35 megabytes to as large as 85 megabytes. The total size of all data for all participants is 942.1 megabytes.

The data we obtained is packaged in MATLAB structure files. There is one structure file for each trial for each footwear condition for each participant. The full data set contains data for 14 participants. Each participant has three different footwear conditions (i.e. normal AFO, dynamic AFO and Shoes), for each footwear condition, walks through the data capture space approximately 10 times, although this number can vary between participants because of walking speed. Therefore, each participant has approximately 30 MATLAB structure files associated with him or her.

[5] "Gait Analysis & Rehabilitation." Gait Analysis & Rehab – Qualisys Motion Capture Systems. <http://www.qualisys.com/applications/biomechanics/gait-analysis-and-rehabilitation/>.

Within each MATLAB structure file, the format is identical and is summarized in the figure below:

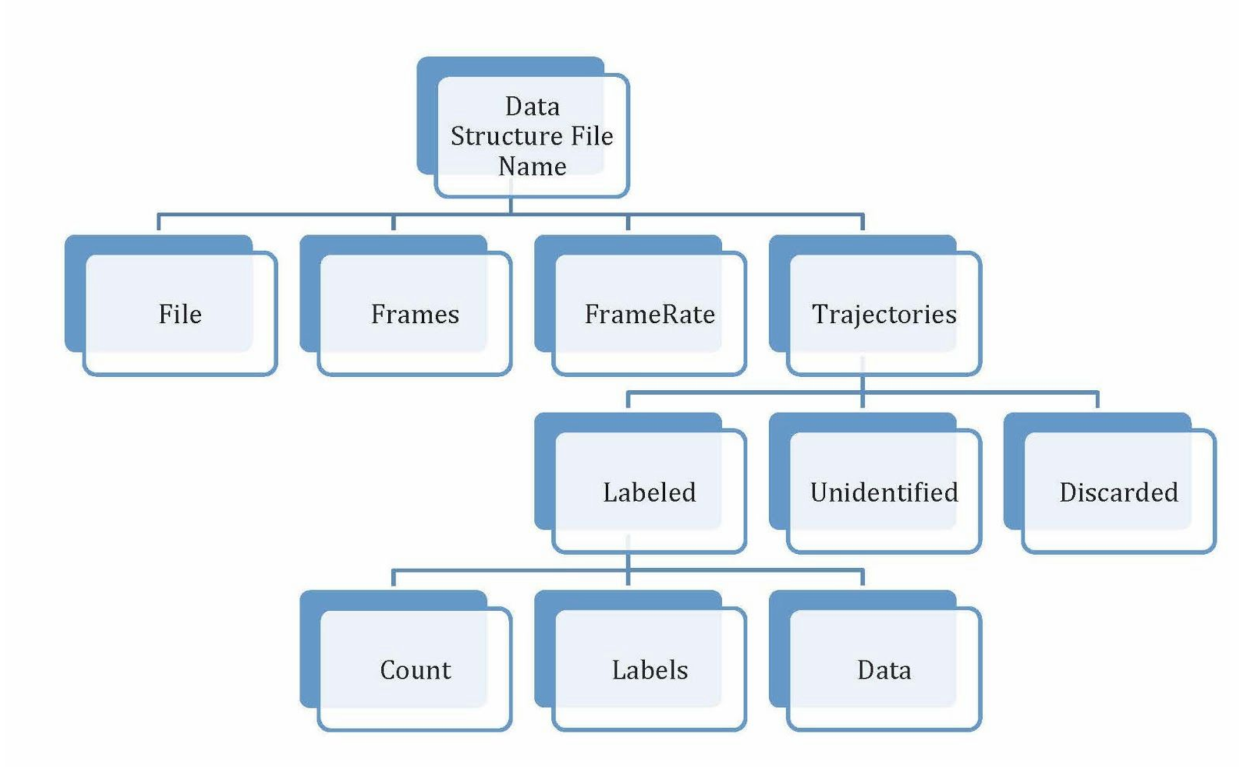

Figure 2.1: Flowchart showing the attributes of each MATLAB structure file for each trial for each participant starting at the top level of the structure file.

At the top level of each MATLAB structure file, there are four attributes: “File”, “Frames”, “FrameRate”, and “Trajectories”. “File” gives the file path of the structure file, which we do not use. “Frames” gives the number of frames of video captured to generate that MATLAB structure file. “FrameRate” contains the frame rate of the motion capture cameras and “Trajectories” is an attribute that contains another MATLAB structure file.

Three attributes within the “Trajectories” structure file are themselves MATLAB structure files: “Labeled”, “Unidentified” and “Discarded”. We only use the data contained in the attribute “Labeled” because “Unidentified” and “Discarded” contain data for markers that were either unidentifiable by name, or were otherwise discarded from the labeled data set because they were partially blocked from the view of the cameras.

Within the “Labeled” structure file, there are three more attributes: “Count”, “Labels” and “Data”. “Count” indicates how many of the motion capture markers were found. “Labels” contains the list of names of the markers that were identifiable by the motion capture system, and “Data” is the actual array of gait data.

In the data array itself, each row represents a motion capture marker, each column represents a spatial coordinate or residual (i.e. X, Y, Z, or R) and each slice in the third dimension represents one frame or time point. For example, one data array for participant 15 has the shape 21 X 4 X 3879. This shape indicates that for this trial for participant 15, there are 21 identified markers and 3879 video frames of data for the X, Y and Z coordinates and R. This arrangement of the data structure makes sense since the Qualisys software would take one slice of the data frame and determine the 3D location of each marker, then move on to the next slice, and index with time to get the new 3D location of each marker, recreating a 3D visualization of the walking motion. The number of rows is different for each trial and each participant because the markers identified by Qualisys differ between trials. Equally, the depth of the array also varies between trials and participants because the depth of the array depends on each participant's walking speed and therefore differs for each trial and participant.

SECTION 2.2: DATA PIPELINE, CLEANING AND VALIDATION

Since the data analysis tools we needed were readily accessible in Python, we began by developing an automated process to convert the MATLAB structure file into a set of Pandas dataframes to work with the data within the Python environment. We call this our pipeline.

The pipeline first extracts four 2D arrays, each representing either the X, Y, Z, or R coordinate, for all markers and all time points. Each 2D array is then saved to a separate Comma Separated Values (CSV) file. In addition, the pipeline saves the list of marker names as a separate CSV file. Once all the CSV files have been created by MATLAB, we can read those CSV files in Python and store them in custom objects. We then create a template Pandas dataframe where the rows represent time points and each column represents one of the 23 total possible markers. All elements in this template dataframe are initialized as NaN. We then copy this template dataframe four times in order to have a data frame each for X, Y, Z, and residuals. The imported gait data is then added one marker at a time to the the appropriate dataframe and the appropriate column based on the marker with which the data is associated. This process then results in four final Pandas dataframes that have the same number of columns and for which missing marker data shows up as NaNs instead of simply being missing from the dataframe. We then associate all four dataframes with a custom "Trial" object that stores the data, as well as characteristics such as participant number and trial number as attributes. With this "Trial" object, we can properly tag and label the data and attributes at all times. The process is repeated for every set of trial data for every participant. Based on the results from MATLAB, there are a total of 326 structure files, for which 18 of these or approximately 5.5%, do not contain any data.

In order to ensure that our pipeline was accurately transferring the data, we used two methods to validate the pipeline. First, we visualized 5 selected sets of data in both Python and MATLAB to ensure that they were visually similar. One example is shown below:

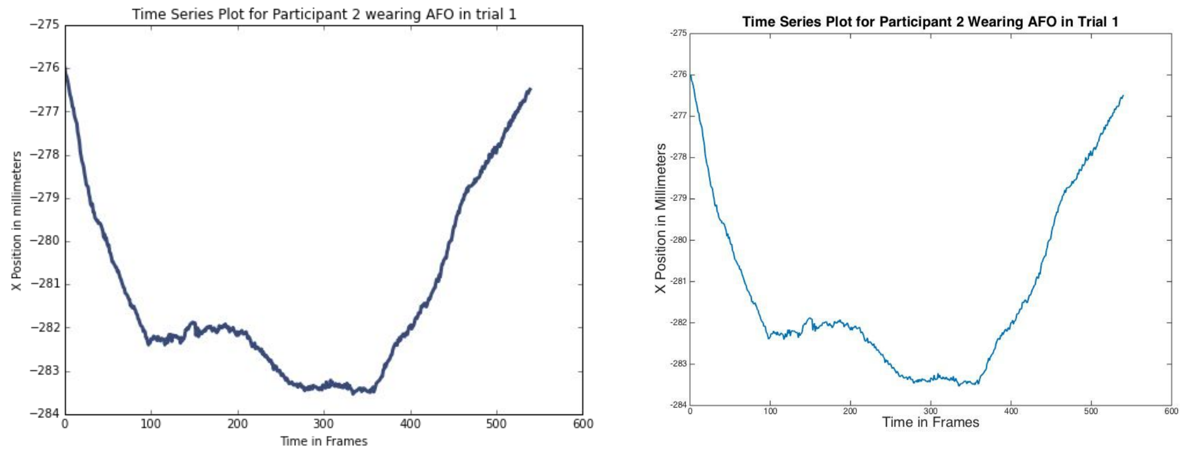

Figure 2.2: Time Series Plot for participant 2 wearing normal AFO in Trial 1 as plotted in Python (left) and in MATLAB (right); Since they are identical, the data for that trial is valid.

As can be seen in the above figure, the graphs of the data in MATLAB and Python look visually similar and this was the observation for all 5 sets of data we visualized. Therefore, we conclude that the pipeline was importing data into Python accurately.

The second method we used to validate the pipeline was to compare summary statistics computed in Python and MATLAB on the same set of data. In this case, we chose to compute the mean, median, maximum and minimum values for each trial and each marker in MATLAB, then in Python, and compared the two. Because MATLAB outputs the summary statistics to a CSV with a 2 decimal point level of precision, we set a threshold equal to this level of precision (i.e. threshold = 0.01mm) to ensure all values were equal to within 0.01mm, at which point we considered the pipeline to be validated.

SECTION 2.3: DATA VISUALIZATION

SECTION 2.3.1: TIME SERIES DATA VISUALIZATION

After importing the data into python, we began exploring the data using basic data visualizations. These visualizations, time series and spaghetti plots, were used to understand the data: the range of values, change in patterns over time, and variability or similarity in these trends amongst participants and trials.

Because the Sacral marker, located at the base of the spine, was the easiest marker to conceptualize, we started by visualizing that marker. The graphs below show an example of the time series for the Sacral marker in the X, Y and Z axes:

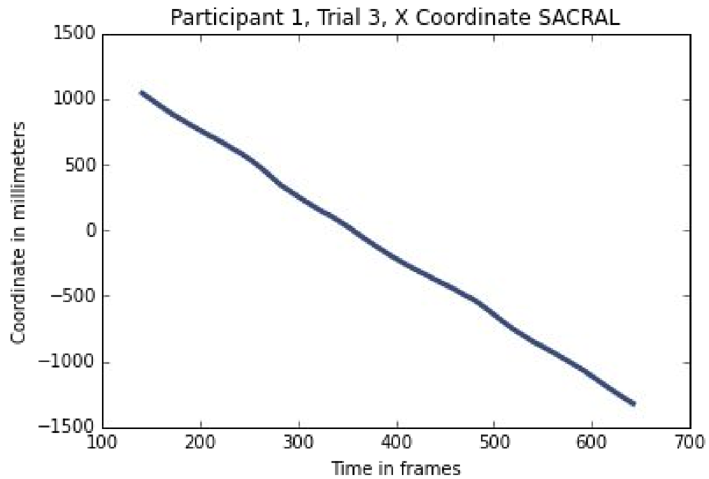

Figure 2.3: Time series plot of X coordinate for Sacral marker for participant 1 Trial 3; As expected, a line is obtained as the participant walks in the negative X direction

As one might expect, since the participant is walking in the -X direction when the gait data was collected, the graph for the X coordinate is a nearly straight line with a negative slope.

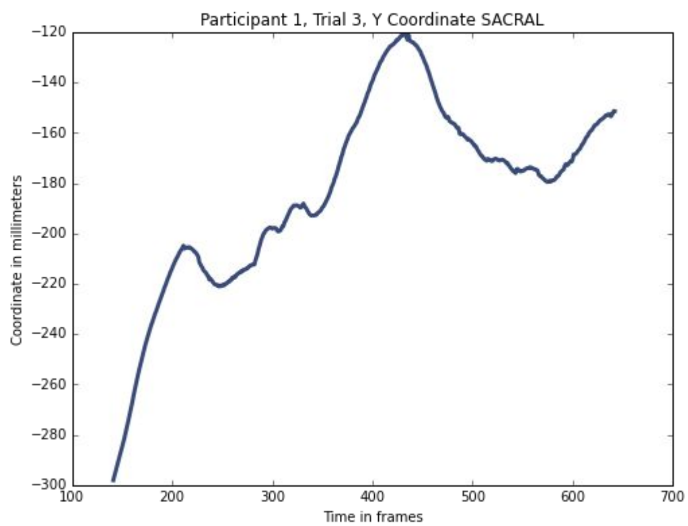

Figure 2.4: Time series plot of Y coordinate for Sacral marker for participant 1 Trial 3; As expected, there is some oscillatory behavior due to walking

As might be expected, in the Y direction (not the vertical axis and perpendicular to the direction of walking), there is some, but not much, oscillatory motion, since any motion in the Y axis would be the result of the person swaying from side to side. Possible reasons for a participant to sway from side to side include the participant using hip hike, or leg circumduction to swing their leg while walking.

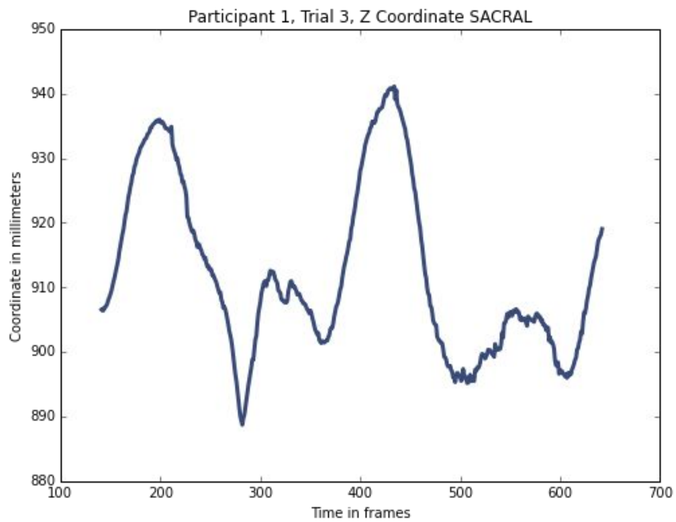

Figure 2.5: Time series plot of Z coordinate for Sacral marker for participant 1 Trial 3; As expected, there is some oscillatory behavior due to walking

In the Z coordinate (Figure 2.5), as one might expect, the graph shows signs of oscillatory motion, which makes sense given that walking is a cyclic motion. In addition, the numerical range of the values in the Z coordinate make sense since 880mm to 950mm above the ground plausibly corresponds to where the person's waist is located during walking. Additionally, since the two major peaks observed in Figure 2.5 correspond to the peaks observed in Figure 2.4, that suggests that the participant is moving their body side to side as well as up and down while walking, which are characteristic movements of hip hike and circumduction. This graph lends further credence to the theory that the participant is using hip hikes and circumduction to swing one leg through while walking.

SECTION 2.3.2: MASS DATA VISUALIZATION

In order to understand the variation between participants' data, we also compared all of the participants' Sacral data in the X, Y, and Z coordinates using spaghetti plots to allows us to visualize the range of the data before we began analysis on the walking data.

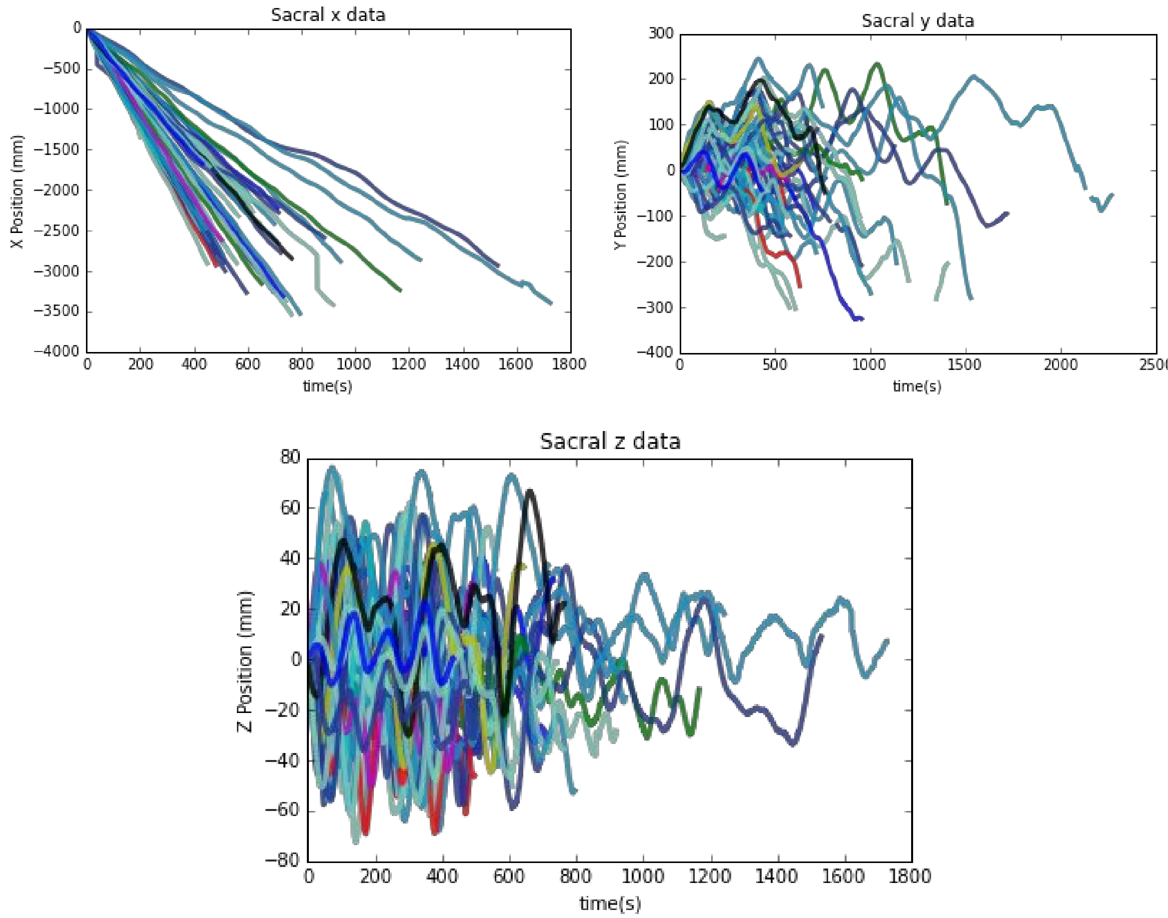

Figure 2.6: (Upper Left) *Time Series Plot for the Sacral marker in the X direction for all Participants wearing normal AFO devices in Trial 5.* (Upper Right) *Time Series Plot for the Sacral marker in the Y direction for all Participants wearing normal AFO devices in Trial 5.* (Lower Center) *Time Series Plot for the Sacral marker in the Z direction for all Participants wearing normal AFO devices in Trial 5.*

These graphs clearly show that there is a wide range of the start and end times of a participant's walk, as well as a range in the initial starting position, which suggests that any analysis we perform on this data will not be able to assume comparable durations of walk or similar starting points.

SECTION 2.4: ANALYSIS OF BASELINE DATA

In most gait analyses, it is standard procedure to normalize the walking trial gait data against a baseline data set, where the participant is standing still. Using this baseline, quantities such as joint angles can be calculated with respect to the baseline values. To do this, we first had to determine if the baseline data contained outliers or other seemingly anomalous data points that would affect the quality of the baseline data.

Out of the approximately ten trials each participant has for each footwear condition, the first three or four trials contain data of the participant standing still in the motion capture space. We use these trials as baseline data. Therefore, we began our analysis of the baseline data by just looking at the standard deviation (SD) of the data for all markers, space coordinates, and participants for the baseline trial (in this case we used Trial 1) to quantify the amount of variability present in the baseline data. The graph below shows the distribution of SDs for the data for all markers, space coordinates, and participants for trial 1:

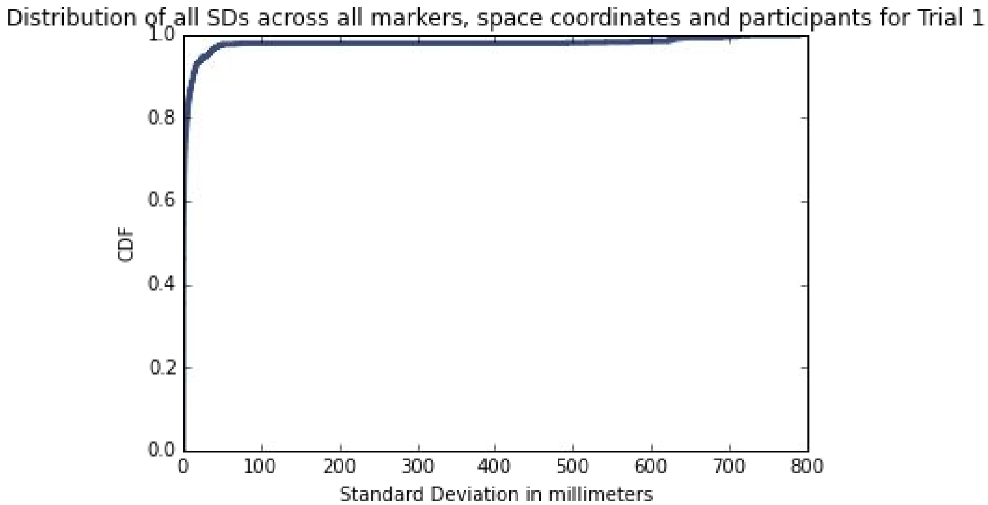

Figure 2.7: Cumulative Distribution Function (CDF) of Standard Deviations (SDs) across all markers, space coordinates, and participants for Trial 1. Since most of the data lies near SD=0mm, we can conclude most of the baseline data shows no significant standard deviation.

As can be seen in the above figure, the majority of participants' Trial 1 gait data in x, y and z show a standard deviation near zero, suggesting that the Trial 1 (baseline) data generally does not vary. This observation of the baseline data is desirable for a trial where the participant is standing still. Therefore, we can conclude that the data seems reliable enough to use to establish a participant's baseline since approximately 80% of the trials have a standard deviation equal to or near zero.

In addition, at the upper range of the CDF, there seems to be a small jump in the CDF at SD ≈ 650mm. This value is surprising because it is unlikely that a participant might sway by 650mm while trying to stand still. Therefore, this suggests that there might be some amount of measurement error due to equipment error affecting the data. To further investigate this, we can focus the graph on just the data sets with a standard deviation (SD) of more than 30mm:

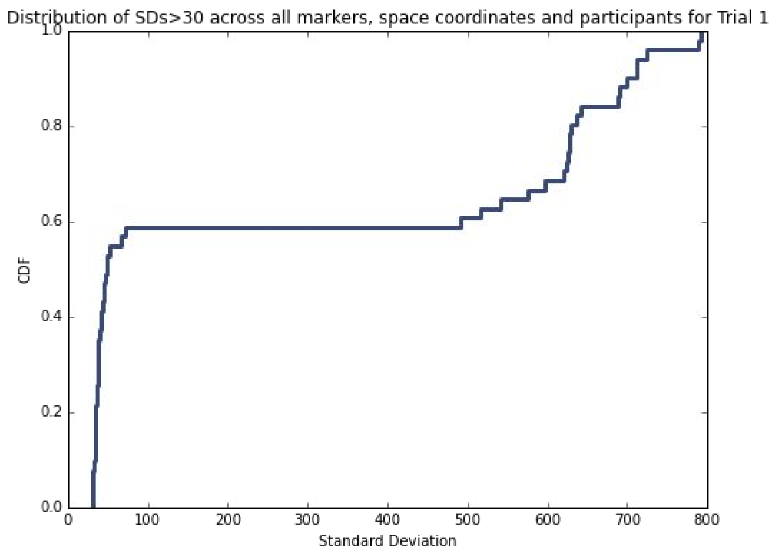

Figure 2.8: CDF of SDs greater than 30mm across all markers, space coordinates and participants for Trial 1. Since most of the data lies below SD=100mm, we can conclude that the baseline data is reasonably reliable.

As can be seen from the above graph, by zooming in to focus only on participants who have SDs larger than 30mm, we find that approximately 60% of this subset of the data have SDs of 100mm or less, which is likely to be acceptable, considering some participants may have a hard time standing still to within 30mm. These results thus demonstrate that the baseline data shows only small variability. However, it is concerning that the maximum calculated SD is 791.9mm or 800mm. This result seems to suggest there may be erroneous data points in the data set but they are not a frequent occurrence.

In order to be fully confident that the data was reliable, we further analyzed the data by computing the difference between median and mean to determine if there were any outliers present in the data. If the data did not contain outliers at either ends, then the median and mean values should be approximately the same. However, if there are significant outliers, one will be higher than the other. The graph below shows the CDF of all (Median-Mean) values for all markers, space coordinates, and participants for Trial 1:

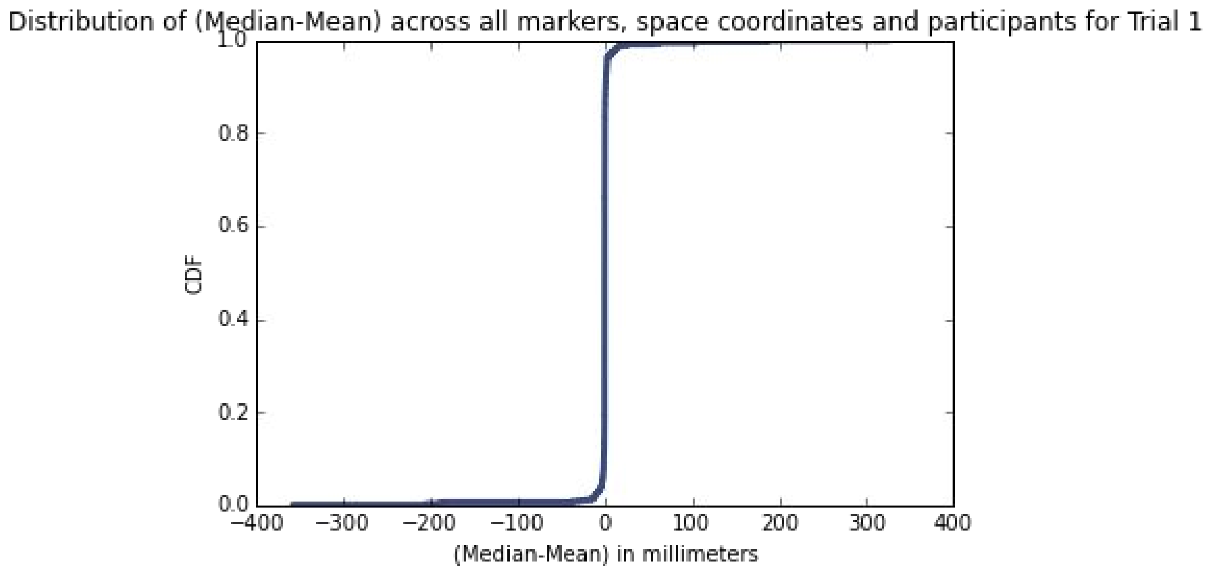

Figure 2.9: CDF of (median-mean) values across all markers, space coordinates, and participants for Trial 1. Since almost all values lie at the point where (median-mean)=0, we can conclude that there are generally no significant outliers in the data.

As can be seen in the above figure, the steepest slope occurs at about zero and that slope starts at approximately CDF=0 to about CDF=1. This result suggests that a large majority of participants have Median-Mean values that are essentially the same and their data does not have significant outliers.

Finally, we also recorded which participants' trials showed (median-mean) values larger than 10mm and discovered that those trials all belonged to participants 3, 10 and 15. Since all the trials with (median-mean) values larger than 10mm belong to only three participants, we can conclude that the majority of participants have means and medians that lie within 10mm of each other.

Overall, these results show that, for the vast majority of participants' trials, there are no significant outliers and the data shows small variability over time. Therefore, the first trial is suitable as a baseline measurement.

SECTION 3: GAIT CYCLE EXTRACTION

Another standard procedure when analyzing gait is to carry out the gait analysis over the course of a gait cycle, which is defined as the time between two heel strikes of the same foot. For this reason, extracting gait cycles from the data was an important step in order to perform meaningful analyses. For a person with a healthy gait, the heel strike can be found by searching for the point in the heel motion where the heel marker has the lowest z position, no velocity, and a maximum acceleration. Unfortunately, due to the unusual foot movements made by participants with MS, finding the heel strike using these factors alone is difficult.

Thus, for our gait data, we needed two pieces of information to extract a gait cycle: where the earliest heel strike occurs and how long the gait cycle is. Starting at the earliest detected heel strike, the gait cycle can be extracted from a walking trial if we know the period of the gait. By using the frequency of the gait to get the period of the gait, we can extract a gait cycle from our data.

To determine the period of a gait cycle, we initially performed autocorrelation on the gait data. However, autocorrelation requires more than three gait cycles to accurately identify the period and there are several participants whose gait data contains fewer than three gait cycles, or has missing data in a few of the cycles. Therefore, autocorrelation was not an ideal method for extracting a gait cycle alone. In unusual cases like these, we used the average peak frequency, as calculated by discrete fourier transform (DFT) analysis, to find the fundamental period of the gait cycle. Because this method is less precise, we only use it to find gait cycles when autocorrelation fails.

SECTION 3.1: FREQUENCY DETECTION

To most effectively find the period of a gait cycle, we analyzed the frequency content of the walking trials. This work was done to ensure that the fundamental frequency associated with the heel marker (the marker traditionally used for gait cycle analysis) was the same as the fundamental frequency for the other markers on the same leg during a given trial. Thus, we compared frequency and position plots for the Sacral marker (the highest position maker possible) and the heel (one of the lowest position markers possible).

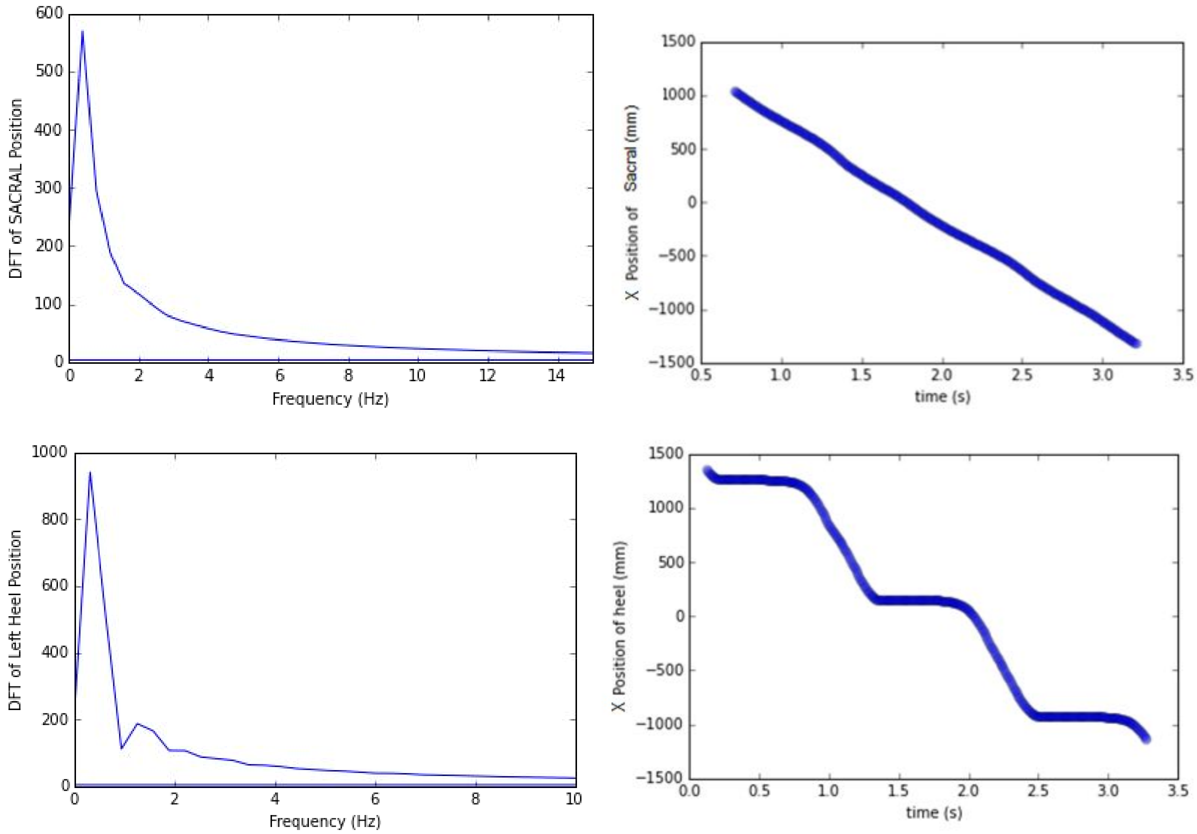

Figure 3.1: Frequency, as calculated with discrete fourier transform (DFT) analysis and position of left (unimpaired) Heel and Sacral markers

As can be seen in the figure above, the Left Heel and Sacral markers show the same fundamental frequency. However, secondary frequencies are detected for the left heel marker that are not in the Sacral marker. This finding would suggest that we should use the Sacral marker to find the fundamental walking frequency. Unfortunately, the Sacral marker data may be affected by unusual hip movements as described in section 1. Therefore, since the same fundamental frequency shows up in both markers, we can conclude that the heel marker will likely give us a more reliable measure of the walking frequency which can be used to determine the period of a gait cycle.

While the fundamental frequency was comparable across markers on the unimpaired leg, we needed to ensure those frequencies would also be comparable across markers on the impaired leg. Therefore, we compared the differences in the frequency components between the unimpaired and impaired (or device and non-device) leg. In the figures below, these are represented by the left (unimpaired) and right (impaired) toe markers.

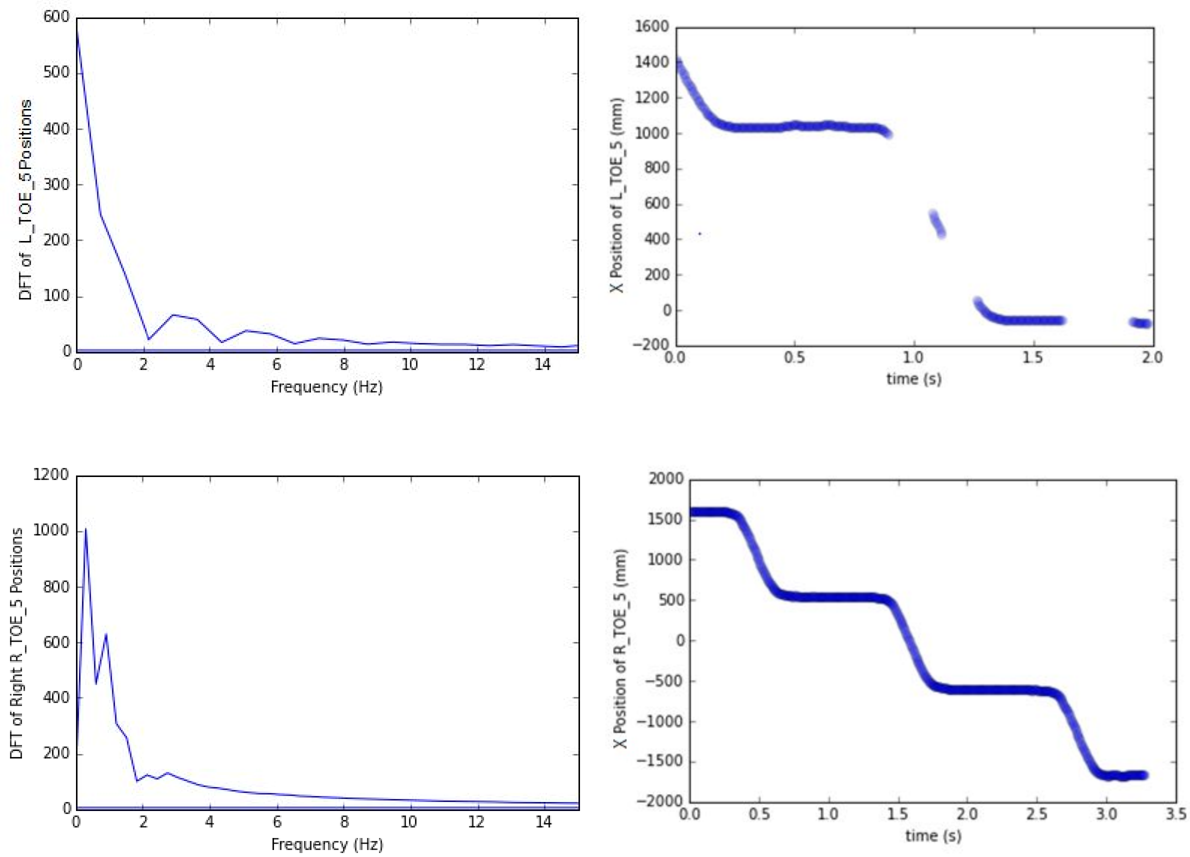

Figure 3.2: Frequency and position of the left and right Toe_5 marker. The lack of data in the left toe position causes the fundamental frequency to be zero. The right toe marker has a large secondary peak, but this peak is never taller than the initial fundamental walking frequency.

The right and left toe markers have significant differences in the peak frequencies that appear. We would expect the highest frequencies to occur in the toe, where there is the most degrees of freedom compared to the positions of the other markers. By comparing the fundamental frequencies from both of these graphs, we can ensure that secondary frequencies are never taller than the fundamental frequencies. Unfortunately, in cases like the L_Toe_5 marker (upper right graph), there is enough missing data that the fft module assumes that the fundamental frequency is zero. Thus, the fundamental frequency is chosen by selecting the frequency with the maximum DFT value. However, if that frequency is zero, that marker is ignored. If a large amount of analyses have a fundamental frequency of zero, we could un-bias the data first, which would remove zero as a fundamental frequency. However, because the majority of data does not have a zero fundamental frequency, these points were ignored. The non-zero frequencies for each marker with a maximum DFT value are then averaged between markers on the same leg to ensure that the frequency used to find the gait period is the fundamental frequency.

SECTION 3.2: HEEL STRIKE DETECTION

With the duration of the gait cycle quantified, the second piece of information needed to extract a gait cycle was to find the heel strike. We began by looking at the time series plots of the position, velocity and acceleration of the heel:

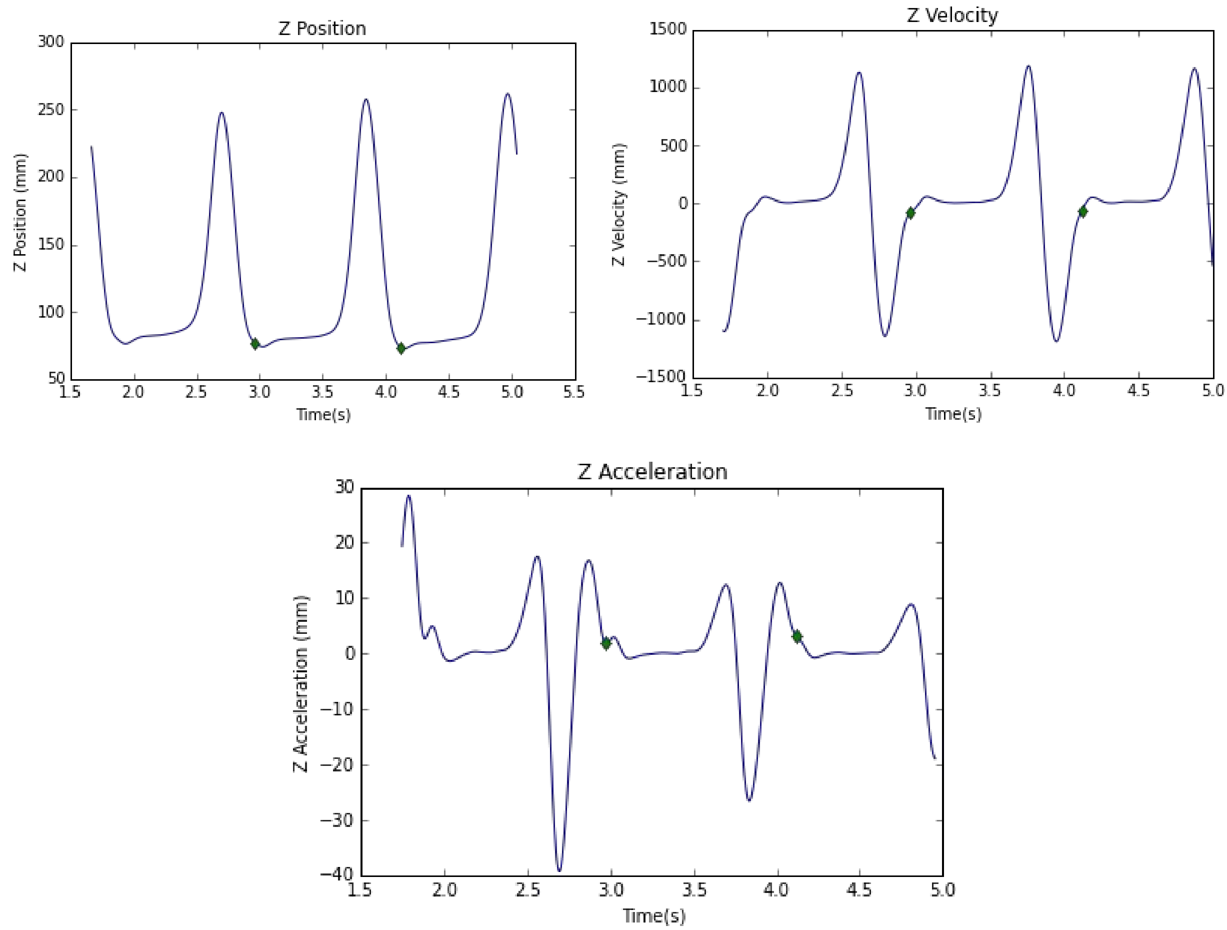

Figure 3.3: Position (top left), velocity (top right) and acceleration (bottom) of the unimpaired heel marker over time. The diamonds show where one would expect the heel strikes to be located.

Based on the above figures, at the points where we would expect the heel strike to occur, the position of the heel strike was not necessarily the lowest, the velocity at the heel strike was between -10 and 10, and the acceleration did change significantly, but was not the maximum. Therefore, the simplest metrics would not accurately allow us to detect the heel strike. To work around this, we gave each variable different weights that were individually tuned in order to take advantage of all of these factors and, using a convolution, maximized the value of the convolution at the point that a heel strike occurs, as seen in Figure 3.4 below:

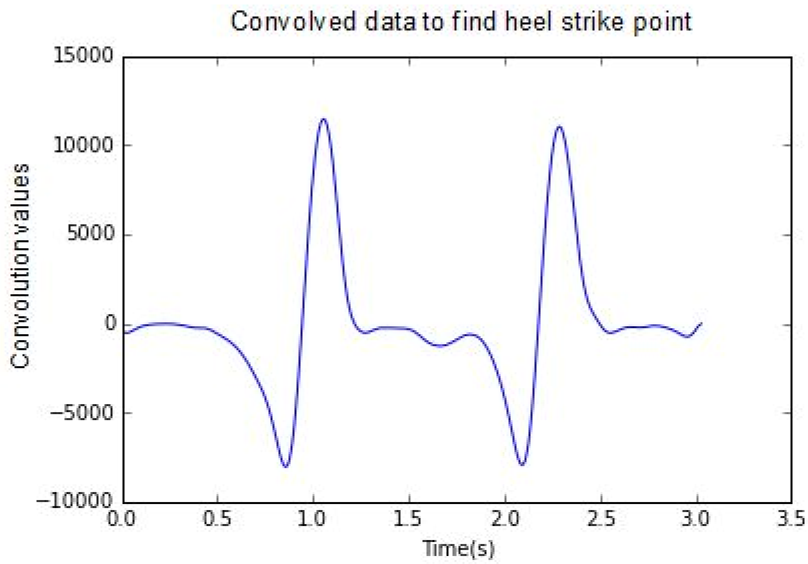

Figure 3.4: Data convolved to maximize heel strike in a walking trial over time.

With the convolution value at the heel strike point maximized, we can find at least a single heel strike as shown by the green dot in Figure 3.5 below:

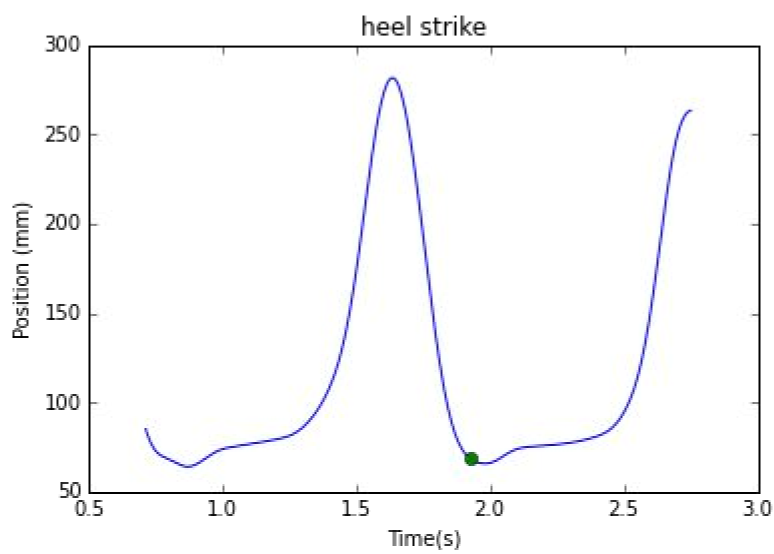

Figure 3.5: Initial heel strike identified (shown with the green dot). The heel strike occurs at the point where there is a maximum value from the convolved data and that point is overlaid over a plot of a walking trial.

SECTION 3.3: DEFINING A GAIT CYCLE

After we have located the heel strike point and know the period of the gait cycle through detecting the frequency of the gait cycle, we can begin extracting the gait cycle from the heel strike point and calculate the length of the gait cycle using the detected fundamental frequency. This analysis can then allow us to find the approximate location of the second heel strike. We then take the 20 data points around that approximate location and search for a local maximum at that location and that maximum then corresponds to the next heel strike. The overall result is shown in the figure below:

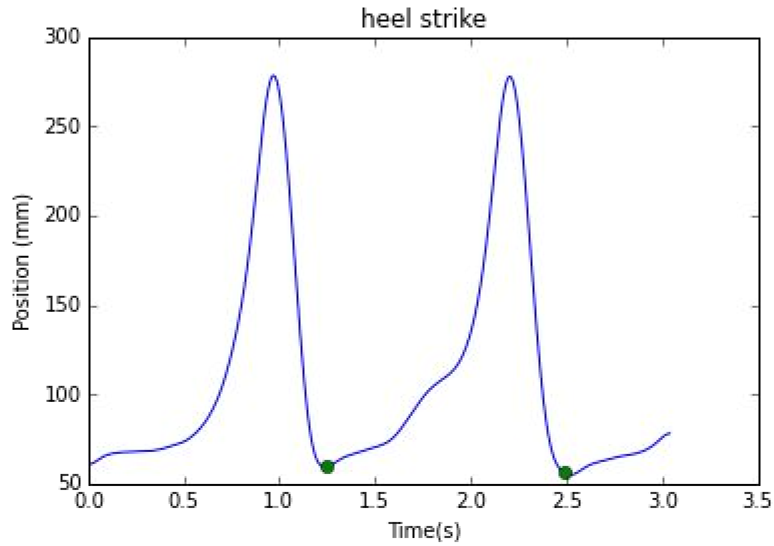

Figure 3.6: Graph of a walking trail showing the identified heel strike points. This graph is a visual confirmation that we can find successive heel strikes in a walking trial and can therefore extract a gait cycle.

Using the above method, all the analysis we run can be indexed by percentage of gait cycle to normalize the walking trials for all participants.

While the above methodology is functional for almost all cases, one failure mode of this method of detecting the gait cycle is the increased importance of the velocity and acceleration components. Because these two components are given increased importance, small fluctuations in the relative position that are not initially visually apparent may be given greater importance. Nonetheless, a full gait cycle is still captured, which leads us to believe that this method is fairly robust even with this failure mode.

SECTION 3.3: EXTRACTING GAIT CYCLE DATA

Once we had validated our method for extracting gait cycle data, we wrote a function that would take gait data as an input and return the heel strike locations in terms of the index location of the heel strike in the gait data. Using this information, we can then extract data frames that represent the data from all identified markers during that one gait cycle. As an example, a time series plot of the heel marker z axis data is shown in the figure below:

SECTION 4: ANALYZING SYMMETRY

The first key question is whether the dynamic AFO enables better symmetry during gait. To answer this question, we tried two methods of measuring symmetry: using cross-correlation as well as using frequency analysis.

SECTION 4.1: MEASURING SYMMETRY WITH CROSS-CORRELATION

One way we wanted to measure symmetry was using cross correlation of the left and right heel marker data over the course of a single gait cycle. Using this method, the better the left and right heel data matched each other, the larger the cross-correlation value and therefore the more symmetric the gait. However, simply performing a cross-correlation of the data from a single gait cycle could give us inaccurate results because the cross-correlation value increases not only when the shapes of the left and right heel data are similar but also when the curve for the right heel data has a higher peak than the curve for the left heel data or vice versa. This would be an undesirable result since we would want the cross correlation value to decrease to indicate lack of symmetry if the curves for the left and right heel data have peaks of different heights.

As such, to combat this problem, we further processed the left and right heel data for a single gait cycle by subtracting the mean and dividing by the standard deviation in order to normalize the gait data and cause the peaks of the two curves become similar. With this processing done, if the person’s gait was symmetric, the left and right heel data would have comparable periodicity. As a result, we would expect the left and right heel data to match perfectly and therefore the cross-correlation value should be large. To initially visualize this, the gait cycles for the left and right heel markers were overlaid in order to observe how much overlap there would be. An example is shown in the graph below:

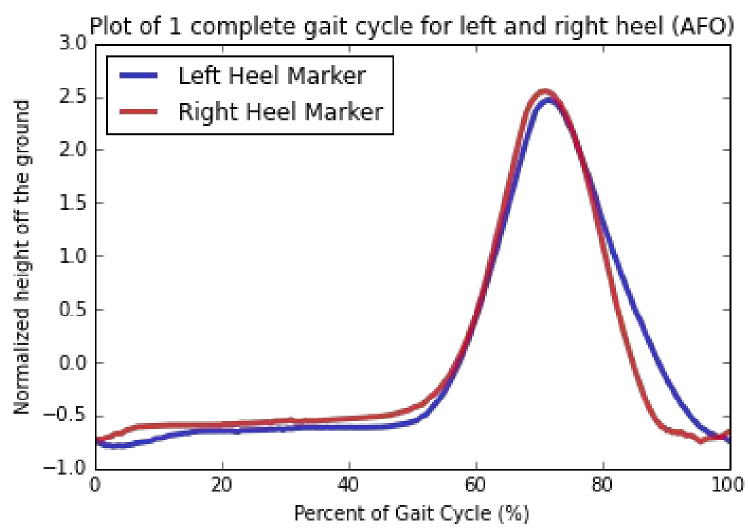

Figure 4.1: Overlaid right and left heel markers. The height data has been normalized by subtracting the mean and dividing by the standard deviation in order to ensure the peaks for the left and right heel are the same.

We then proceeded to carry out the cross correlation of the left and right heel marker data using the normalized NumPy correlate function. Cross-correlation is a 2D pattern identification that is commonly used to identify the degree to which multiple signals are correlated. This method can be employed to identify the offset of two different signals. If we treat the left and right leg as two signals, the cross-correlation of the two legs would demonstrate the walking symmetry by the vertical difference between the two signals from each foot.

After performing this cross-correlation for all participants and walking trials under the normal AFO and dynamic AFO conditions, we obtained two lists of cross-correlation values, one for the normal AFO condition and one for the dynamic AFO condition. We then performed a one-sided hypothesis test to determine if the mean cross-correlation value was indeed higher in the dynamic AFO condition compared to the normal AFO condition. If this was true, that would indicate that the left and right heel data was more symmetric in the dynamic AFO condition compared to the normal AFO condition.

The following summarizes the hypothesis test we performed:

Goal: To determine if the legs are more symmetric in normal AFO or dynamic AFO footwear conditions.

Test Statistic: Difference in mean cross-correlation value for the normal AFO and dynamic AFO conditions
Null Hypothesis: No difference in the mean cross-correlation value under either the normal AFO or dynamic AFO conditions

With this hypothesis test, we found that there was no statistically significant effect size between the normal AFO and dynamic AFO trials. This suggests that the dynamic AFO performs no better than the normal AFO in terms of changing the shape of the heel data. However, this test of symmetry is limited. If a person’s gait were not symmetric, there could be two key reasons why. Firstly, the person might not be moving their left and right leg at the same speed, which would change the shape of the left and right heel data. Secondly, the person might not be lifting their left and right heel up to the same height. As such, the hypothesis test above suggests that the dynamic AFO does not change the shape of the data compared to the normal AFO but it does not necessarily suggest that the dynamic AFO does not cause the data to be more symmetric. To conclude that, one would need to investigate how the dynamic AFO changes the height to which participants are lifting their left and right heel compared to the normal AFO. However, we did not have the opportunity to conduct that analysis and future work should further investigate

SECTION 4.2: UNDERSTANDING SYMMETRY THROUGH FREQUENCY ANALYSIS ACROSS FOOTWEAR CONDITIONS (AFO, dynamic AFO, SHOES)

Ideally, gait symmetry would be measured by comparing the gait cycles of the left and right leg for a participant to see how similar they are. However, because detecting gait cycles was such a major challenge during the early stages of the project, we attempted to work in parallel to find other ways to quantify symmetry without needing to rely on being able to detect a gait cycle. One hypothesis we had was that observed frequencies in the gait data would be an indicator of symmetry. This hypothesis is based in the theory that the frequencies that become apparent by conducting discrete fourier transform (DFT) analysis are inherently related to the walking speed and other movements in the walk. Therefore, participants whose gait showed fewer frequency peaks after DFT analysis likely have less added movement during their gait. Additionally, we posited that the closer the peak frequencies are between legs, the more symmetric the gait was.

Each participant has three different footwear conditions: normal AFO, dynamic AFO and Shoes. We used the NumPy fft module to perform DFT analysis on all of the left and right leg markers in order to understand the differences in observed frequencies between the right and left markers under each footwear condition. In addition, since participants are usually disabled in either the left or right leg but not both, our analysis also effectively allowed us to study differences between the healthy and disabled leg. We also performed this analysis on both the marker positions and velocities to see if different frequency patterns emerged if the analysis was run on the positions or velocities. The following sections will discuss our results for each of the three footwear conditions separately.

SECTION 4.2.1: SYMMETRY IN THE NORMAL AFO CONDITION

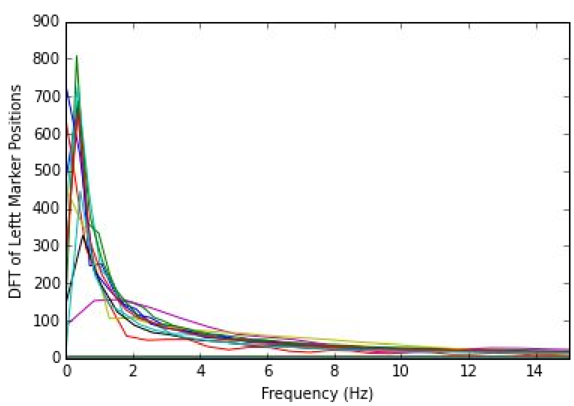

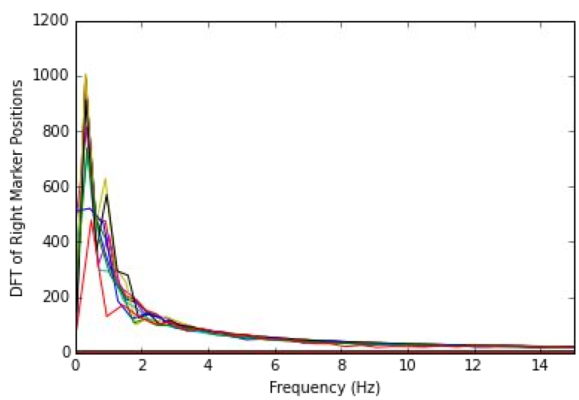

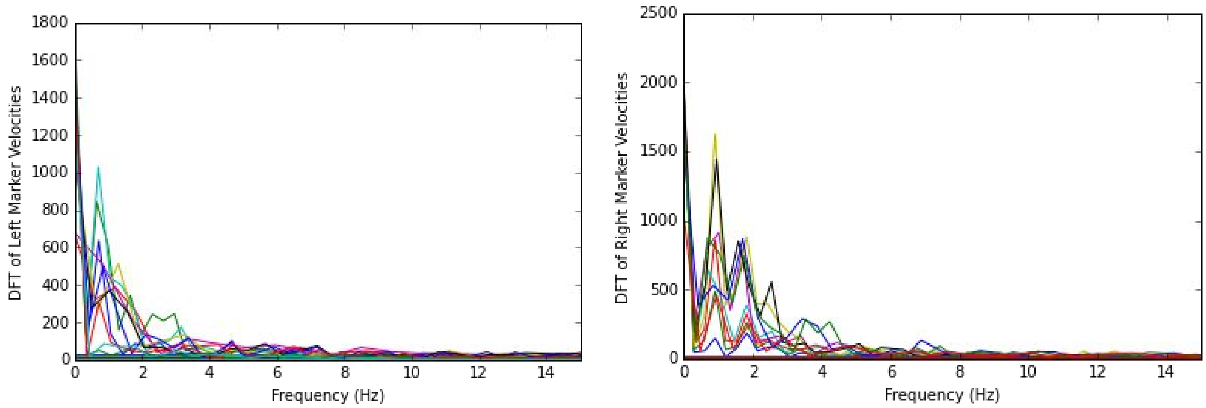

Figure 4.2: Differences between the frequencies in the left leg and in the right leg positions (upper graphs) and velocities (lower graphs). The right leg has a second major apparent frequency that is not seen in the left leg. There are also more higher frequencies in the right leg velocities compared to the left.

As was noted during the gait cycle discussion in section 3.1, while wearing an normal AFO, several frequencies higher than the fundamental frequency are observed for the markers on the right leg that are not observed in the left leg. Since this participant has a disabled left leg, the higher frequencies are likely again a bi-product of the additional movement needed to swing the leg through as well as an indication that the gait is not symmetric.

SECTION 4.2.2: SYMMETRY IN THE DYNAMIC AFO CONDITION

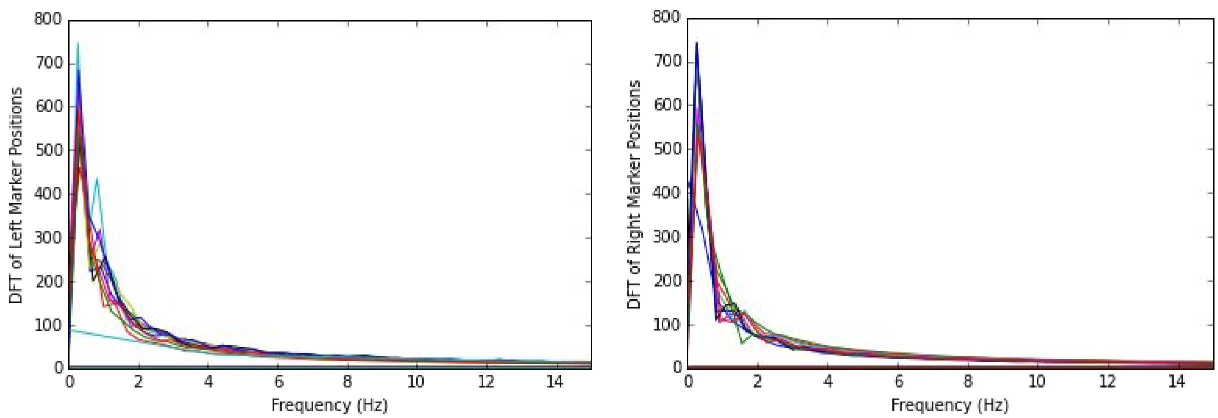

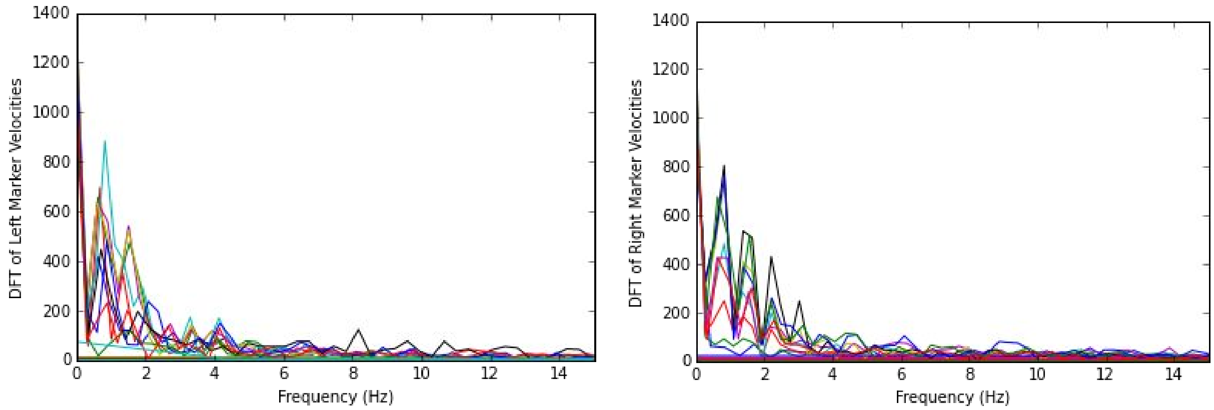

Figure 4.3: An overlay of the frequency distributions for all markers in the first dynamic AFO trial for the left and right leg positions (upper graphs) and velocities (lower graphs). As one of the legs are now powered by the dynamic AFO, all of the markers fall under similar frequency distributions.

By comparison, when the participant wears a dynamic AFO, only a single fundamental frequency is observed for all of the markers on both legs. Because of this similarity in the fundamental frequency for both position and velocity, it suggests that this gait is symmetric. This participant has an impaired left leg and the extra frequencies and differing distribution of fundamental frequencies indicate that the participant uses extra movement in walking. Because we have observed similarities in the frequency content of both legs in the dynamic AFO condition and did not observe such similarities in the normal AFO condition, we can conclude that the gait is likely more symmetric because both legs now exhibit the same peak frequency values.

SECTION 4.2.3: SYMMETRY IN THE SHOES CONDITION

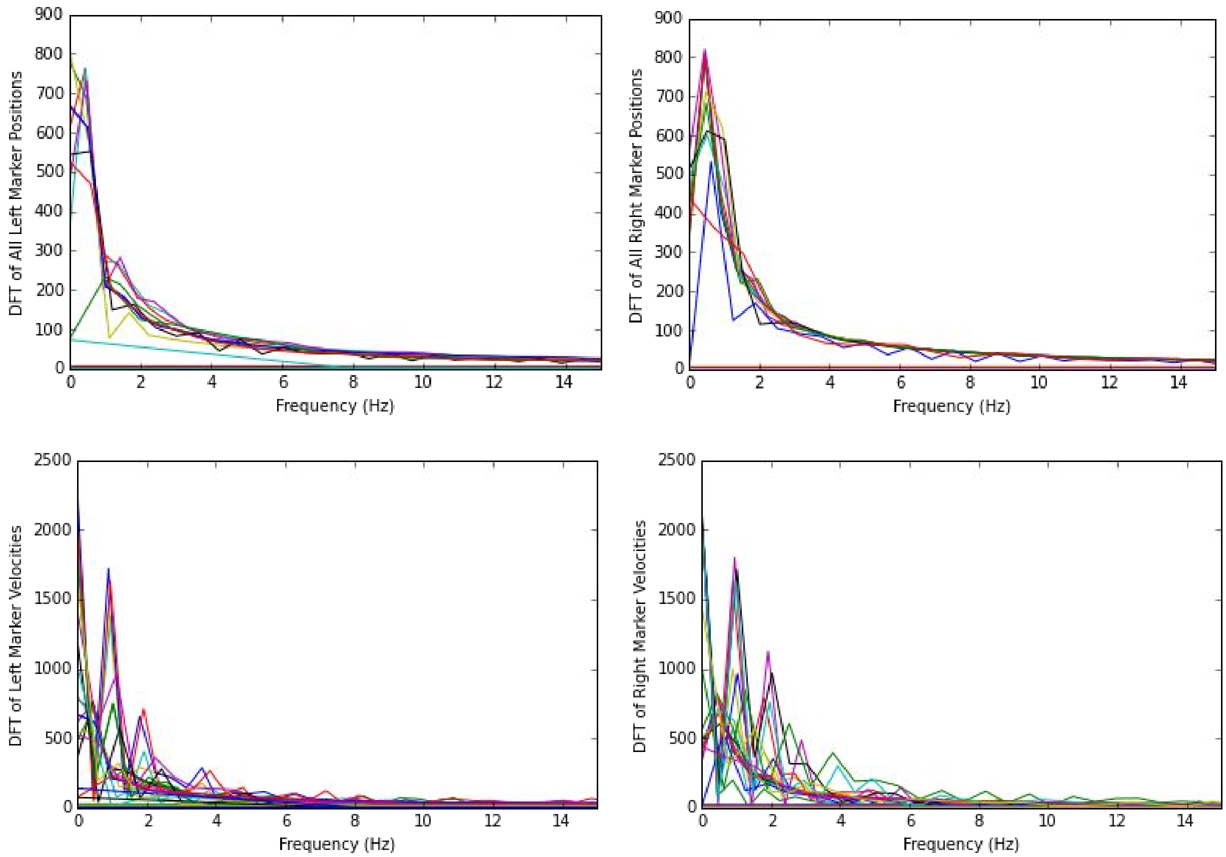

Figure 4.4: (Left) An overlay of the left marker position (upper graph) and velocity (lower graph) frequency distributions in the first shoes trial. (Right): An overlay of the right marker position (upper graph) and velocity (lower graph) frequency distributions in the first shoes trial.

The frequencies of the shoe positions and velocities are incoherent and varied between markers even on the same leg. The lack of a single fundamental frequency even between one leg, as well as the difference in the fundamental frequencies between legs, suggests that there is little symmetry between these legs. This variation is likely due to the amount of extra movement required to get the right leg to swing through. Additionally, the shoe condition has greater variation in the time spent on each leg as well as the length of the trial, which also contributes to the random frequencies in this graph.

SECTION 4.2.4: OVERALL SYMMETRY OBSERVATION

By comparing the left and right leg across footwear conditions, we can see an improvement in the alignment of peak fundamental frequencies between the impaired and unimpaired legs, from shoe through the normal AFO to the dynamic AFO shoe conditions. This result suggests that the greatest gait symmetry is achieved when a dynamic AFO is worn and the least gait symmetry occurs when participants wear their own shoes. Not only do all the markers exhibit the same characteristic frequencies, but also the number of higher frequencies are not apparent in the dynamic AFO condition. With fewer observed frequencies, we can plausibly conclude that the gait of both legs is now symmetrical and does not contain unusual leg movements that might contribute to the higher frequencies observed in the frequency curves for the other footwear conditions.

SECTION 5: ANALYZING OVEREXTENSION OF THE KNEES

Greater knee extension is sometimes referred to as knee hyperextension, which means that the portion of the leg below the knee is bent further forward than the baseline while the participant is standing still. This is unhealthy for a person because it is placing unnaturally high stresses on the knee.

To answer this question, we first developed a method for calculating joint angles before performing a hypothesis test to answer this question.

SECTION 5.1: CALCULATING JOINT ANGLES

In order to compute the joint angles in the simplest, most reliable way, we decided to use vectors and the cosine rule to calculate the knee angles given the spatial coordinates of the markers. To do this, we computed the direction vector from the Lat_Knee marker to the trochanteric (Troch) marker as well as the direction vector from the lateral knee lateral knee (Lat_Knee) marker to the lateral malleolus (Lat_Mal) marker on both legs. The diagram below shows where the markers are located:

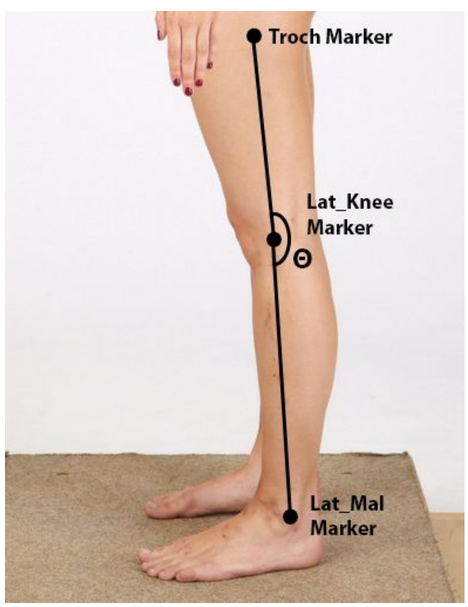

Figure 5.1: Diagram showing the positions of the Troch, Lat_Knee and Lat_Mal markers used to calculate the knee angle

With these two direction vectors defined, we can then compute the angle between them using the formula:

$$\mathbf{A} \cdot \mathbf{B} = \|\mathbf{A}\| \, \|\mathbf{B}\| \cos\theta,$$

Using this vector method, we can then calculate the knee angles observed for a participant. As an example, the time series of the joint angles calculated for the duration a walking trial is shown below:

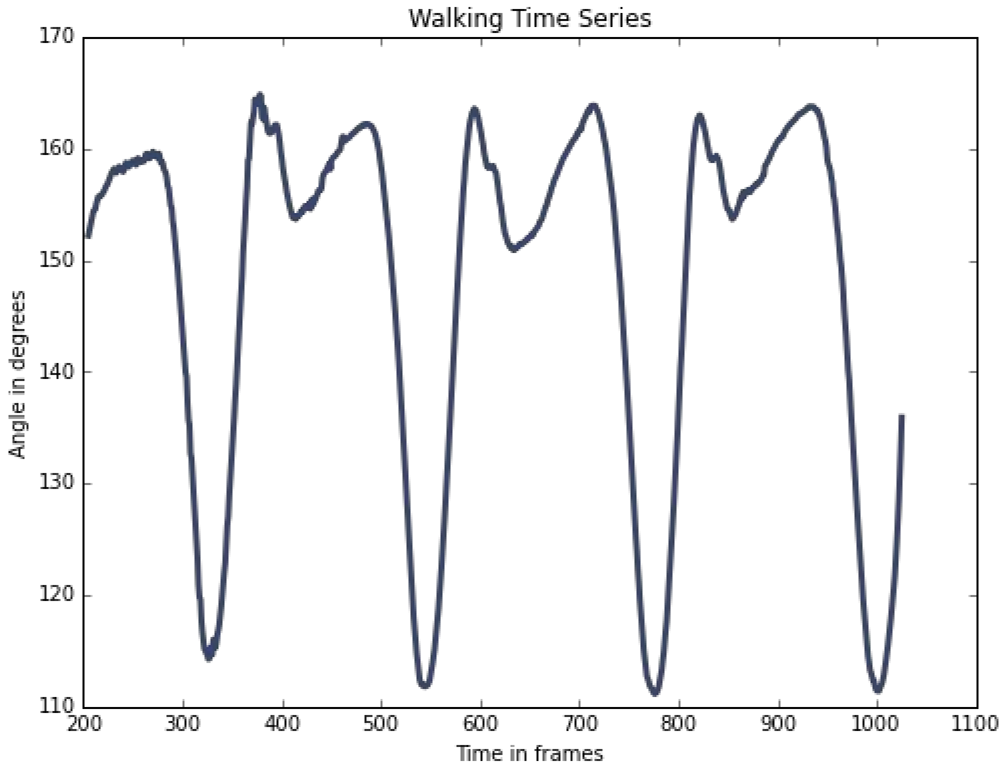

Figure 5.2: Visualization of knee joint angle while walking plotted as a time series. The plot shows a relatively clean data set with the expected periodic trend

As would be expected, the walking time series shows the expected periodicity and peaks but also some amount of missing data. However, the angles shown in the graph agree with expected range of possible joint angles during the course of walking. This indicated that the method we had developed for calculating joint angles was sufficient for our analyses.

SECTION 5.2: PROCEDURE FOR ANALYZING JOINT ANGLES

In order to perform analyses on joint angle data, it is standard practice to compare joint angles during walking to a participant's joint angles while standing, also known as their baseline joint angles. This is because every person's baseline joint angles will be slightly different and so normalizing against a person's baseline joint angles makes joint angles more comparable between people.

To compute this baseline, we use the vector method described above to calculate the joint angles observed over the duration of a trial where the participant is standing still (typically trial 1). From this time series of joint angles, we then select the median joint angle and treat that as the baseline angle of the knee for that participant. We can then take the joint angles while walking that was computed earlier and subtract the baseline joint angle. This thus gives us the change in joint angle from baseline while a participant is walking, which we can then perform analyses on. This procedure was used to calculate the joint angle changes for either the left or the right knee, whichever one the participant was wearing the orthotic on.

SECTION 5.3: MEASURING KNEE HYPEREXTENSION

Using the methods described above, we now had the change in knee angle from baseline while a participant is walking. We then proceeded to conduct a hypothesis test between the dynamic AFO and normal AFO conditions to determine if the dynamic AFO caused hyperextension of the knee compared to the normal AFO. The reason for only testing between dynamic AFO and normal AFO is that we assume that the regular orthotic (i.e. normal AFO) does not cause hyperextension and, therefore, we want to know if the dynamic AFO is causing any hyperextension that isn't already present by virtue of wearing an normal AFO. Because of the way the joint angles are calculated, the angle increases when the knee is bent forward and decreases as the knee is bent backward. So, more hyperextension occurs when the change in joint angle is greater.

A summary of this hypothesis test is described below:

Goal: To determine if the change in knee angle for someone wearing a dynamic AFO is indeed greater (i.e. hyperextended) than that for someone wearing an normal AFO

Test Statistic: Differences in mean joint angle
Null Hypothesis: There is no difference between the two groups

To perform this hypothesis test, we generate a model by stacking the data for the data sets (i.e. data set for the participant wearing a normal AFO and dynamic AFO) on top of each other and then shuffling them to simulate no difference between the two footwear conditions. The hypothesis test was conducted as a one-sided hypothesis test to determine when mean values are greater than or equal to the actual. We used mean to minimize the effect of outlying data points.The reason for using a one-sided test is that we don't care if the knee bends back more. We only care if the knee bends forward more. Thus, a one-sided test is appropriate for this application.

We performed this hypothesis test for every participant for every valid trial on either the left or right leg, whichever one the participant wears the orthotic on. We then plotted a histogram of the results:

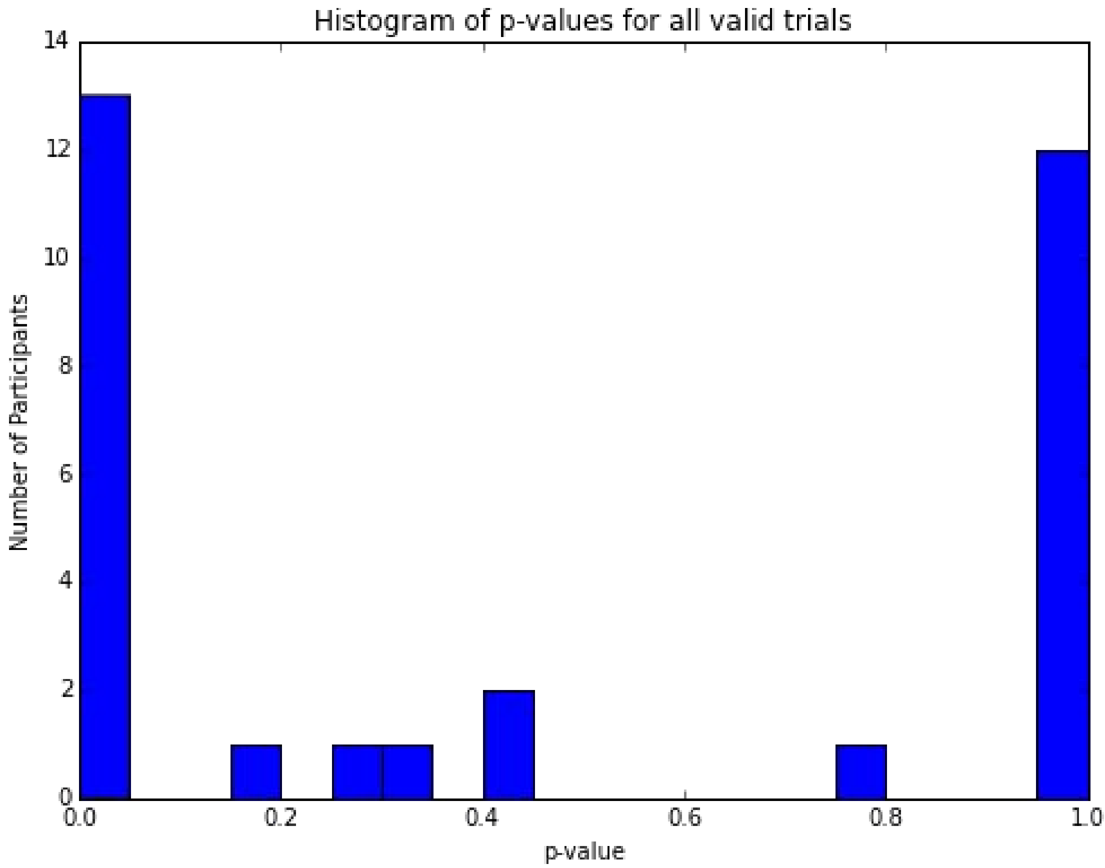

Figure 5.3: Histogram showing how many trials have each p-value. Based on the above graph, we can see that a 17 participants show p-values equal to or less than 0.05, indicating that the effect sizes observed are statistically significant.

As can be seen above, approximately 13 out of the 31 valid trials that were tested show p-values less than 0.05, indicating that, for those trials, the observed effect size is statistically significant. However, this also means that 18 of the 31 valid trials showed did not show any statistically significant difference in mean knee angle. Therefore, given this result, it is difficult to conclude if there is actually a generalizable difference in mean knee angle between the normal AFO and dynamic AFO conditions.

In terms of effect size, the actual observed effect size is approximately 3-6 degrees. This result could mean one of three things. Firstly, the effect size could be small and statistically insignificant. Secondly, the effect could be small and statistically significant. Finally, the effect size could be large and statistically significant. For trials that showed p-values less than 0.05, we know that the effect size is statistically significant. In addition, given the relatively small value of the effect size relative to the typical range of knee angle during walking, we do not suspect that the computation used to produce the data might be erroneous. That therefore leaves just one possible explanation of this outcome. That is, that the calculated joint angle is large enough to be notable but not large enough that we would expect the calculation to be wrong. Therefore, based on this hypothesis test, for the trials that showed p-values less than 0.05, we can conclude that there was a statistically significant change in the knee angle and therefore that the dynamic AFO is causing slight overextension of the knee compared to the normal AFO. However, because of the comparatively similar number of trials that showed statistically significant p-values compared to the number of trials that showed statistically insignificant p-values, we are unable to conclude that the dynamic AFO is causing overextension of the knee compared to the normal AFO in general.

SECTION 6: ANALYZING LEG CIRCUMDUCTION

The last key question is whether leg circumduction, which is typically observed in people with MS, is reduced when wearing the dynamic AFO or the normal AFO compared to when the participant is wearing their own shoes.

To investigate this, we employ a similar joint angle analysis as described in section 5 but restrict attention to just the plane perpendicular to the direction of walking (i.e. the Y-Z plane). We use the left and right Trochanteric (Troch) markers to create a horizontal reference direction vector and we use the Troch and Lateral Knee (Lat_Knee) markers on either the left or right leg to create a direction vector for that leg, depending on which leg the participant wears their orthotic on. The angle between those two direction vectors can then be used to measure the leg angle. The larger the leg angle, the further the leg is away from the centerline of the person and the smaller the angle, the closer the leg is to the centerline of the person.

To answer the key question, we again employ a one-sided hypothesis test. As with the hypothesis test conducted on the knee angles, the test is one-sided because we only care if the person is swinging their leg further out in one footwear condition compared to the other (i.e. if the leg angle is larger under one footwear condition compared to the other).

In order to be as thorough as possible, we performed two hypothesis tests in total, one to compare the shoes to the normal AFO condition and another to compare the shoes to the dynamic AFO condition. A summary of the goals for both hypothesis tests are shown below:

Goal for Test 1: To determine if the change in leg angle (i.e. leg circumduction) for someone wearing their an normal AFO is indeed less than that for someone wearing their own shoes.

Goal for Test 2: To determine if the change in leg angle (i.e. leg circumduction) for someone wearing their a dynamic AFO is indeed less than that for someone wearing their own shoes.

For all three hypothesis tests, the test statistic and null hypothesis was the same and can be stated as follows:

Test Statistic: Difference in mean leg angle
Null Hypothesis: There is no difference between the two groups

Each hypothesis test was performed for all participants and the p-values were plotted on histograms to show how many trials have each p-value. A discussion of the results is presented below.

SECTION 6.1: COMPARING THE SHOES TO NORMAL AFO CONDITION

When comparing participants who were wearing shoes to participants wearing an normal AFO, we find that the majority of participants show effect sizes that are statistically significant. The histogram below shows how many trials have each p-value:

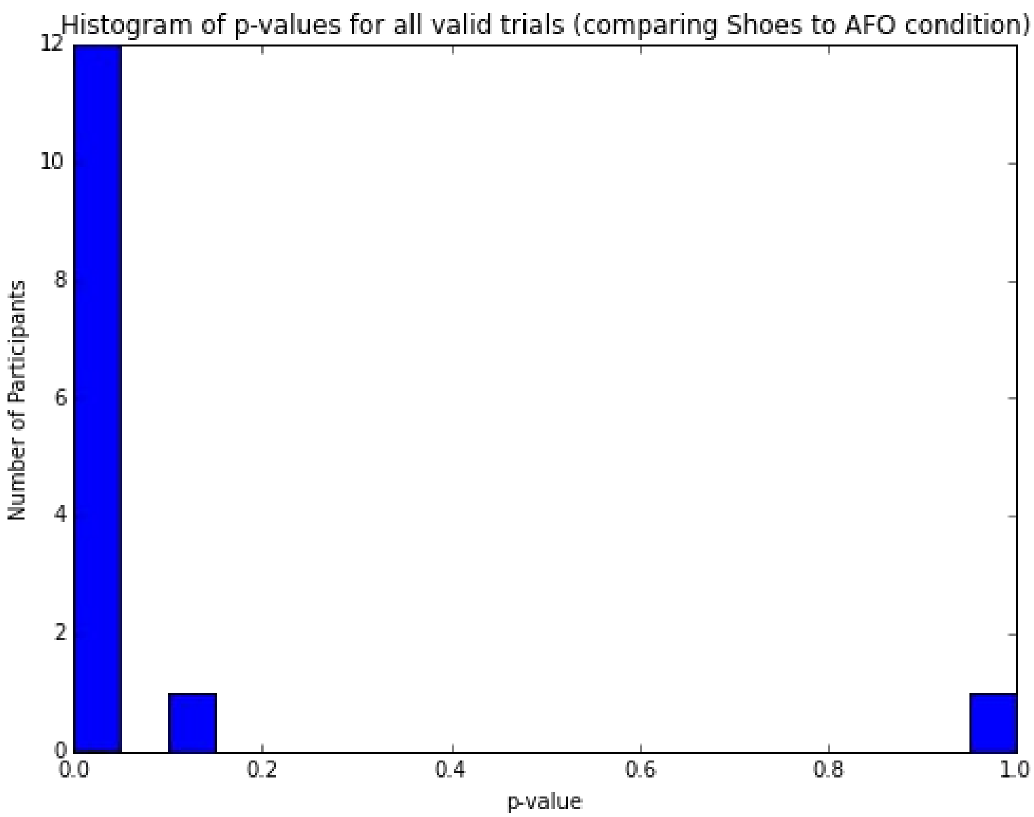

Figure 6.1: Histogram showing how many trials have each p-value (binned into bins 0.05 in size) when comparing the shoes to normal AFO footwear conditions. Based on the above graph, we can see that a 12 participants show p-values equal to or less than 0.05, indicating that the effect sizes observed are statistically significant.

As can be seen in Figure 6.1, 12 out of the 14 valid trials show a p-value of less than 0.05, which means the majority of the valid trials show effect sizes that are statistically significant.

In terms of effect sizes, we find that for trials that showed statistically significant p-values, participants wearing an normal AFO are swinging their leg out more by as little as 0.65 degrees and as much as 3.4 degrees. In this case, this positive effect size indicates that leg circumduction is increased (i.e. their leg is being swung out more) when wearing the normal AFO than when wearing their own shoes. The CDF below shows the distribution of effect sizes when comparing the shoes to the normal AFO condition:

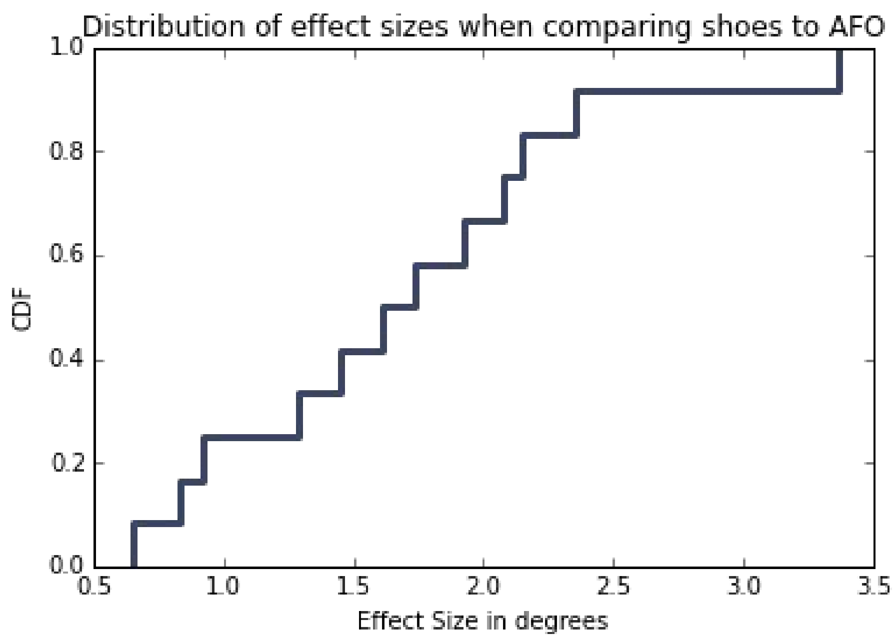

Figure 6.2: CDF showing the distribution of effect sizes for trials showing statistically significant p-values when comparing the shoes to normal AFO footwear conditions. Based on the CDF, we can see that the effect size is no larger than 3.4 degrees.

Based on the above figure, we can see that the maximum effect size is approximately 3.4 degrees. In addition, as can be seen by the steps in the graph shown in Figure 6.2, 11 out of the 12 trials with statistically significant p-values have an effect size of 2.5 degrees or less. When taking both this effect size distribution and the p-value distribution into account, we can conclude that leg circumduction is increased in general in participants who are wearing an normal AFO compared to participants wearing their own shoes.

SECTION 6.2: COMPARING THE SHOES TO DYNAMIC AFO CONDITION

When comparing participants who were wearing shoes to participants wearing a dynamic AFO, we also find that the majority of participants also show effect sizes that are statistically significant. The histogram below shows how many trials have each p-value:

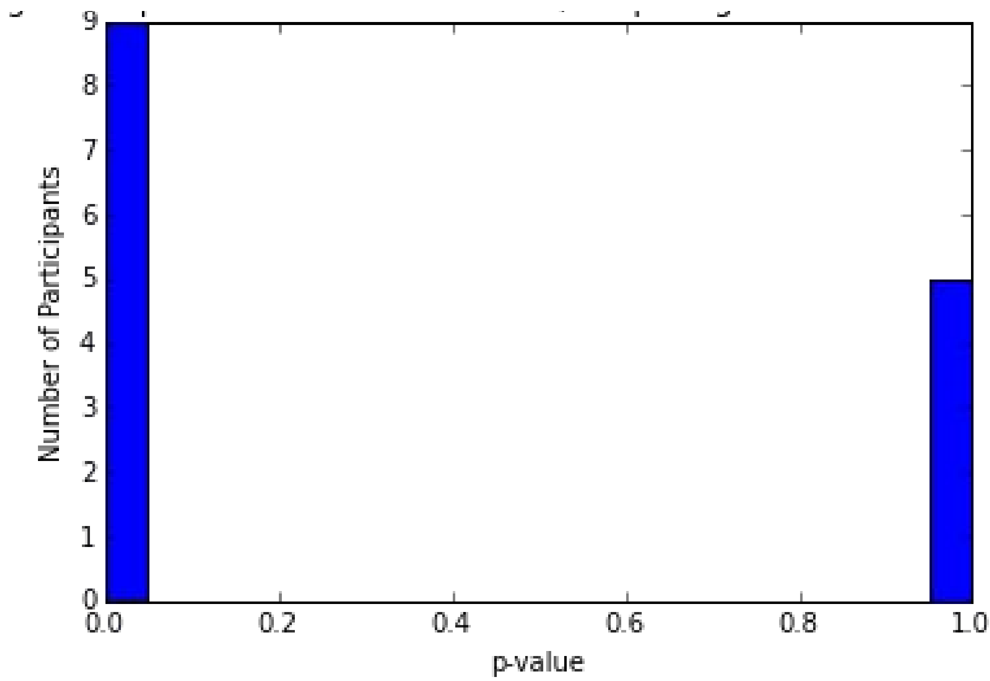

Figure 6.3: Histogram showing how many trials have each p-value (binned into bins 0.05 in size) when comparing the shoes to dynamic AFO footwear conditions. Based on the above graph, we can see that 9 participants show p-values equal to or less than 0.05, indicating that the effect sizes observed are statistically significant.

As can be seen in the figure above, 9 out of the 14 valid trials show a p-value of less than 0.05, which means the observed effect size is statistically significant. This shows that the majority of the valid trials show effect sizes that are statistically significant.

In terms of effect sizes, we find that for trials that showed statistically significant p-values, participants wearing a dynamic AFO are swinging their leg out more by as little as 0.5 degree and as much as 9.8 degrees. In this case, this positive effect size again indicates that leg circumduction is increased (i.e. their leg is being swung out more) when wearing a dynamic AFO than when participants wear their own shoes. The CDF below shows the distribution of effect sizes when comparing the shoes to the dynamic AFO condition:

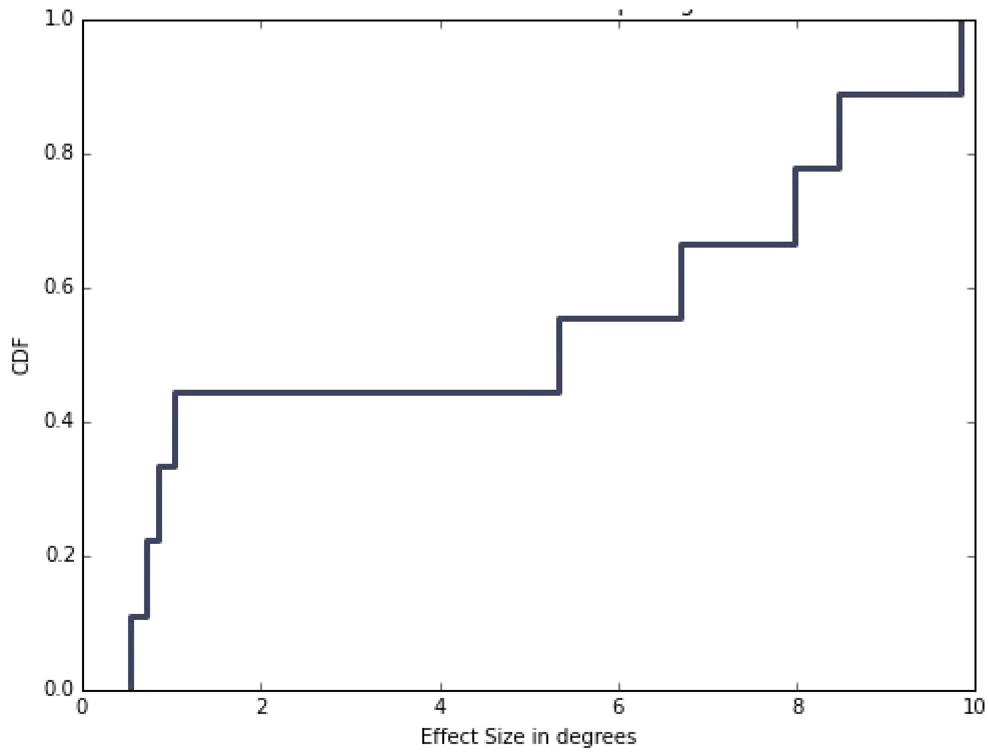

Figure 6.4: CDF showing the distribution of effect sizes for trials showing statistically significant p-values when comparing the shoes to dynamic AFO footwear conditions. Based on the CDF, we can see that the effect size is no larger than 9.8 degrees

Based on the above figure, we can see that the maximum effect size is 9.8 degrees. When taking both this effect size distribution and the p-value distribution into account, we can conclude that leg circumduction is increased in general in participants who were wearing a dynamic AFO compared to participants wearing their own shoes.

Overall, the above results therefore suggest that participants wearing an normal AFO or dynamic AFO are observed to have increased leg circumduction compared to when participants are wearing their own shoes.

SECTION 7: CONCLUSION AND FUTURE WORK

In conclusion, we began this report acknowledging that there were three key questions that are a key part of developing a dynamic AFO. The first key question was to determine if the dynamic AFO enables better symmetry during gait. The second key question was to determine if knee extension was greater with the dynamic AFO compared to the shoes and normal AFO conditions. Finally, the third key question was to determine if leg circumduction was reduced with the dynamic AFO and normal AFO compared to the participants own shoes. Based on our data, we were able to draw conclusions to answer all three questions. Firstly, based on results from cross-correlation and frequency analysis, we were able to conclude that the dynamic AFO does offer better symmetry during gait compared to the normal AFO. Secondly, based on our hypothesis test conducted on the calculated knee angles, we found that some trials did show statistically significant effect sizes of 3 - 6 degrees, suggesting that, for those trials, wearing the dynamic AFO was causing slight overextension of the knee compared to when the normal AFO was worn. However, because of the comparable number of trials that showed statistically insignificant effect sizes compared to those that showed statistically significant effect sizes, we were unable to conclude that the dynamic AFO was causing overextension of the knee compared to the normal AFO in general. Finally, based on our hypothesis test conducted on the calculated leg angles, we found that leg circumduction was increased when comparing participants who wore an normal AFO compared to when they wore their own shoes. Similarly, we also found that leg circumduction was increased when comparing participants who wore a dynamic AFO compared to when they wore their own shoes.

Future work on this project should include further analysis of knee extension and hip circumduction to produce more conclusive results than we have been able to produce at this point. In addition, further work could be done to determine the clinical significance of our findings for all three key questions.