

User Intent Identification in a Lower-Extremity Exoskeleton via the Mahalanobis Distance

Taylor M. Gampon, James P. Schmiedeler, and Patrick M. Wensing

Abstract—Existing strategies for controlling lower-limb robotic exoskeletons place different emphasis on the user’s intentions considered at various resolutions, from high-level goals (increase speed) to mid-level actions (increase stride length) to low-level joint behaviors (increase hip flexion). While sensors onboard the exoskeleton sense the human only indirectly, via the human-robot interface, they offer advantages over more direct methods in terms of the time required to don the device. In this study, exoskeleton users, both able-bodied and having spinal cord injury, were asked to perform changes in their intended gait speed. Onboard sensor measurements were used offline to test an intent identification algorithm based on the Mahalanobis distance. The algorithm’s goal is to identify an intent change and correctly classify its type, but not to realize that change via the exoskeleton. The algorithm correctly identified instances in which the user desired to walk faster or slower than the nominal speed in the device. For able-bodied subjects, the average delay between the known intent change and correct identification by the algorithm was 0.63 s. This delay for non-able-bodied subjects was 0.93 s on average. These proof-of-concept results show that intent identification based on the Mahalanobis distance is possible, while analysis of the approach suggests areas for further improvement.

I. INTRODUCTION

Control of robotic exoskeletons has been a recent topic of debate [1], [2], with applications spanning worker assistance in manufacturing to rehabilitation following neuro-muscular injury [3], [4]. All face similar constraints and challenges associated with physical human-robot interaction - the robot must ensure user safety in addition to mediating and enacting some agreed-upon goal. For lower extremity exoskeletons, this goal, in general, is to walk, but may include further details about the nature of the gait. Any user preference regarding the way in which the exoskeleton should move is referred to as “intent.” This work is specifically focused on determining, but not enacting the user’s intended gait speed.

Instead of explicitly identifying user intentions, many exoskeletons tightly couple human intent with the control law [5], [6]. For instance, in individuals with asymmetric gait, Complementary Limb Motion Estimation (CLME) [7] determines the more affected leg’s motion from that of the less affected leg assuming common intent on both sides updated every stride. Control is even further interwoven with the user’s intentions for direct volitional control strategies in which exoskeleton joint motion is caused by the excitation of electrical signals from the user’s muscles (electromyography

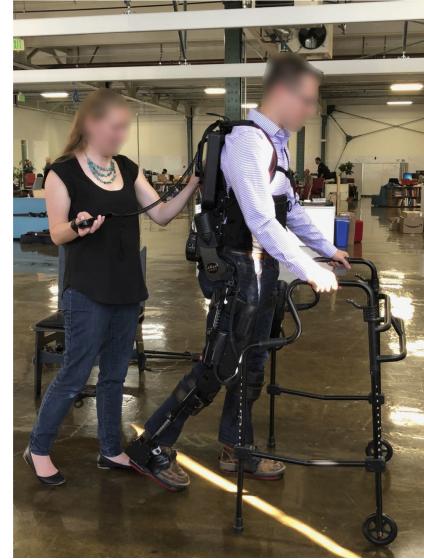


Fig. 1. Able-bodied exoskeleton user walking in EksoGT.

- EMG) [8], [9] or brain (electroencephalography) [10]–[12]. While more directly sensing user intentions, direct volitional control requires additional sensors to be placed on the human body, which can consume valuable therapy time in clinical settings. Furthermore, EMG-based methods can be challenging for individuals with chronic spinal cord injury (SCI) due to lack of ability to reliably control muscle contraction. Finally, EMG-based methods can be grouped with other methods like admittance/impedance shaping and gravity compensation that react to intended joint-level motions, but do not attempt to decipher the user’s higher-level motion goals (i.e., to change gait speed) [13]–[15].

Recent work has shown that changes in user-intended walking speed cause detectable differences in the output of sensors already embedded in the exoskeleton [16]. Specifically, joint motor positions and commanded motor currents vary in systematic ways when the user wishes to speed up or slow down. These results suggest that a model could be trained on the expected values of the onboard sensor readings during constant-intent walking. During testing, new data that are statistically different from the expected values would indicate that an intent change has likely occurred, though the accuracy of such an indicator may depend on the AAD used. One statistical tool for identifying outliers of a multivariate Gaussian random variable is the Mahalanobis distance [17]. Unlike hypothesis testing in which multiple samples of the dependent variable are required, the Mahalanobis distance

*This work was funded by the National Science Foundation

*All authors are with the Department of Aerospace and Mechanical Engineering, University of Notre Dame, 365 Fitzpatrick Hall of Engineering, Notre Dame, Indiana tgambon@nd.edu

allows a single data point to be compared to the distribution in question. Since data from the exoskeleton arrive one time-step at a time, this is the critical reason for using the Mahalanobis distance. This metric is the multivariate analog of expressing how many standard deviations a sample is from the mean of a univariate distribution. Where the standard deviation scales the univariate distance by the amount of variance in the distribution, the Mahalanobis distance takes into account both the variance and covariances of each dimension of the distribution. This is particularly advantageous for exoskeleton sensor data because the position and current draw at one joint are likely to be highly correlated with those at other joints.

The preliminary study reported herein develops a statistical model for the onboard sensor measures of the EksoGT exoskeleton (Fig. 1), assuming they follow a multivariate Gaussian distribution at every time step in the gait cycle during constant-intent walking. After establishing the validity of this assumption, the model is built by estimating the appropriate Gaussian distributions at each time step of each gait phase based on samples from a training data set. Subsequently, data are compared to the model using the Mahalanobis distance to determine the likelihood that user intent has changed. Data for both model training and model evaluation are taken from experiments reported in [16]. The results are compared across exoskeleton users with and without a chronic SCI and across crutch use and walker use.

II. METHODS: DATA COLLECTION

Data were collected while three able-bodied (AB) and three non-able-bodied (NAB, chronic SCI) exoskeleton users made changes in their intended gait speeds. All subjects gave informed consent to participate in the study approved by the IRB of the University of Notre Dame (Protocol 18-04-4650). One NAB subject had an injury at T5, one from T8 to L2, and the third was not reported. No further demographic details were collected. All subjects had substantial experience using the EksoGT exoskeleton, by Ekso Bionics (Fig. 1). The EksoGT has four powered degrees of freedom (DOF): two revolute hip joints and two revolute knee joints, all in the sagittal plane. Each ankle joint contains a passive spring with an adjustable center point to allow for flexion/extension. The onboard computer records at 500Hz the joint (motor encoder) angle and commanded motor current for each powered DOF. Each foot contains two resistive force sensors - one at the toe and one at the heel.

For all trials, the EksoGT employed a finite-state control strategy. This strategy prevents collapse of the joints during stance and has predefined trajectories for swing, providing corrective input at the joints when the user deviates from them. In double support, the controller provides vertical support to both legs, and the user must shift weight to the leading leg to initiate swing of the trailing leg. Experiments were performed across a six-meter, flat, straight walkway. Subjects donned the exoskeleton and were allowed to perform several walkway warm-up laps to adjust to the exoskeleton before

beginning trials. Although no minimum or maximum warm-up time was enforced, most subjects either did not use the warm-up time or completed only a couple of lengths to warm up. Trials were completed with two AADs - crutches and a walker. When the AAD changed, subjects were again given the option to complete warm-up laps to adjust to the new equipment. Each official trial began with the subject walking naturally in the device across the walkway before receiving a single verbal command to speed up, slow down, or make no change in gait speed near the halfway point along the walkway. This protocol ensured that the subject had several steps to get up to steady-state speed prior to the intent command and several steps after the intent command before needing to slow down to stop.

One trial set consisted of three repetitions of each command with each AAD (trial order varied with subject) for a total of eighteen trials per subject. The command order was pseudo-randomized within each trial set so that subjects could not anticipate the upcoming command. The eight joint values (motor positions and commanded motor currents for both left and right knee and hip joints) were tracked by the intent recognition algorithm to determine an intent change. Trials with the “no change” (NC) command were used to train a constant intent model. Trials with the “speed-up” (SU) and “slow-down” (SD) commands were used to test the intent identification algorithm.

III. METHODS: MODEL BUILDING

A. Modeling Assumptions

Consider the vector of eight joint-measurement values as a random variable, $\mathbf{x} \in \mathbb{R}^8$. This work assumes that the variable follows a multivariate Gaussian distribution whose mean and covariance depend on the position in the gait cycle when the user is walking with constant intent according to

$$\mathbf{x}(p, t) \sim \mathcal{N}(\boldsymbol{\mu}(p, t), \boldsymbol{\Sigma}(p, t)), \quad (1)$$

where p represents the gait phase

$$p = \begin{cases} 1 & \text{if Right Double Support} \\ 2 & \text{if Right Single Support} \\ 3 & \text{if Left Double Support} \\ 4 & \text{if Left Single Support} \end{cases}$$

and $t \in \mathbb{N}$ represents the number of time steps since that phase began. Right/Left Double Support is defined as the double support phase in which the right/left leg is leading. As an example, the distribution of \mathbf{x} at the twentieth time step within the Left Double Support phase would be described by $\mathcal{N}(\boldsymbol{\mu}(3, 20), \boldsymbol{\Sigma}(3, 20))$. This model in Eq. (1) implies that $\mathbf{x}(p, t)$ comes from a cyclostationary random process [18], meaning that the distribution at each combination of p and t is stationary as long as the user is walking with constant gait intent. Furthermore, each mean and covariance can be estimated by finding the sample mean and covariance at each time step for a set of training observations. Update equations

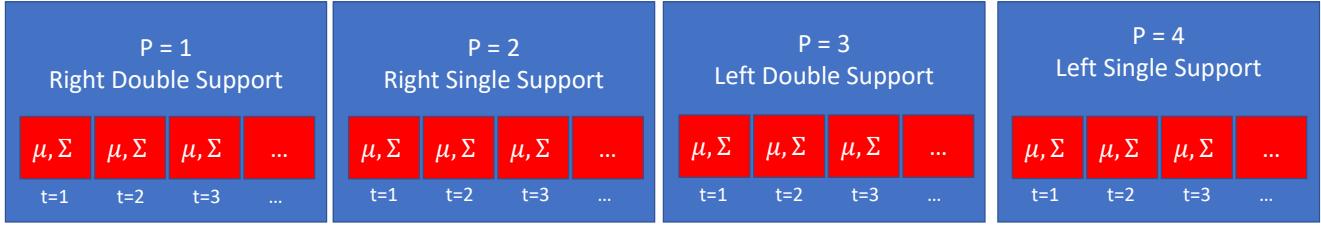


Fig. 2. Model structure. Each time step in each phase includes measurement mean $\mu \in \mathbb{R}^8$ and covariance $\Sigma \in \mathbb{R}^{8 \times 8}$.

estimate these means and covariances after the n^{th} training observation according to

$$\mu_n = \frac{(n-1)\mu_{n-1} + \mathbf{x}_n}{n}, \text{ and} \quad (2)$$

$$\Sigma_n = \frac{(n-1)\Sigma_{n-1} + (\mathbf{x}_n - \mu_n)(\mathbf{x}_n - \mu_n)^T}{n}. \quad (3)$$

B. Model Building

Data collected during NC trials with both AADs served as the training data to build two types of models - an individual model for each subject based on only that subject's trials and a single model trained on the trials of all subjects of the same type (AB/NAB). Training data were organized in the structure shown in Fig. 2 by determining to which of the four gait phases data from each time step belonged ($p = 1, 2, 3, 4$) based on an empirically derived set of rules. A data point was assigned to Right/Left Single Support ($p = 2/4$) if the right/left foot sensors indicated ground contact but the left/right foot sensors did not. If both left and right foot sensors indicated ground contact, the hip and knee angles were assessed to determine which leg was leading in double support ($p = 1$ or 3). A leg was determined to be leading if the hip angle was within 10 degrees of the knee angle because they are nearly the same only at the beginning of the leg's stance phase. For each model phase, t was tracked by counting the number of time steps since the phase's beginning, and it was reset to $t = 1$ when the incoming data switched to a new phase.

The number of observations n for each unique combination of p and t was stored to inform model updating per Eqs. (2) and (3). Once all training data were added, the model was reduced to include only time steps containing at least 10% of the maximum number of observations for each phase. This heuristic rule eliminated outlying observations in which the user was standing still or turning around, as opposed to walking steadily. Each full gait phase observation contained a variable number of time steps of data, even in the training data set. The standard deviation in the number of time steps in these observations was calculated to serve as the threshold of the timing offset that incoming data could have with respect to the model before being considered an indication of changed intent. For example, when the number of time steps in Right Double Support exceeded the mean plus one standard deviation of the model, the user was likely trying to slow down. If the number of time steps was fewer than the mean minus one standard deviation of the model, the

user was likely trying to speed up. Again, to reduce outliers in the calculation of the timing offset threshold, observations with fewer than 20 or greater than 500 time steps (fewer than 0.04 s or greater than 1 s) were not considered as full observations of the phase.

C. Assessing Modeling Assumptions

If the random variable \mathbf{x} follows a Gaussian distribution, the Mahalanobis distance,

$$M = (\mathbf{x} - \mu)^T \Sigma^{-1} (\mathbf{x} - \mu), \quad (4)$$

of samples from \mathbf{x} will follow a χ^2 distribution with the number of DOF equal to the dimension of \mathbf{x} . A quantile-quantile (Q-Q) plot is one way to test that a set of samples follows a given distribution. For a cumulative distribution function (CDF) $\Phi(x)$, quantiles are given by the inverse of Φ . That is, for any probability $\rho \in (0, 1)$, the quantile corresponding to ρ is given by $x = \Phi^{-1}(\rho)$ [17]. When the quantiles of two sets of samples are equal for all values of ρ in $(0, 1)$, the samples are from the same distribution, and the Q-Q plot will align with $y = x$. During model building, the Mahalanobis distances for all data at a single combination of p and t were computed. These distances served as a sample set for comparison to 1,000 random samples of the 8-DOF χ^2 distribution.

IV. REAL-TIME INTENT IDENTIFICATION ALGORITHM

To simulate real-time intent identification, testing data were analyzed one data point at a time without reference to future time steps according to the block diagram in Fig. 3. For each data point, the current gait phase p and time step t were determined in the same manner as for model building. If this p and t combination existed in the model, the current data point was compared to that part of the model by assessing the Mahalanobis distance between it and the Gaussian distribution (Fig. 3). The Mahalanobis distances of samples from a normal distribution follow a χ^2 distribution. Therefore, evaluation of the χ^2 CDF for a Mahalanobis distance gives the probability that any other Mahalanobis distance for the given distribution will be smaller than the one being evaluated. Essentially, it provides the likelihood that another sample would be closer to the mean of that distribution than the point being evaluated. Further extrapolated, a value of one implies that the given data point is certainly not a sample from the model distribution and should be considered an indicator of changed intent. Likewise, a

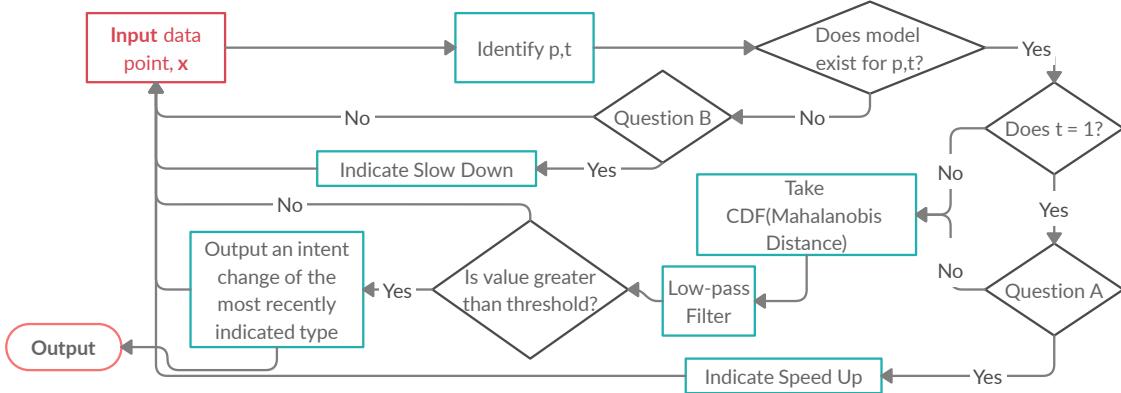


Fig. 3. Real-time intent identification flow chart. **Question A:** Is difference in number of time steps between previous phase and model for that phase greater than one standard deviation of number of time steps for this phase from training set? **Question B:** Is difference between current t & number of time steps in model for this phase greater than one standard deviation of number of time steps for this phase from training set?

value of zero implies that the data point is the exact mean value of the model distribution and should be considered an indicator that intent has not changed.

The Mahalanobis distance signal was filtered with a third-order low-pass Butterworth filter with a cut-off frequency of 2 Hz to avoid triggering intent changes more than twice per second. When a single time step of the filtered value exceeded a threshold, an intent change was signaled. The algorithm was tested with threshold values of 0.5 and 0.6, which correspond to Mahalanobis distances of roughly 7.5 and 8.5, respectively, in combination with the two variations of training data sets.

Recall, as a distance metric, the Mahalanobis distance is strictly positive, so it cannot differentiate between SU and SD cases. The *type* of change can, however, be determined by the timing offset between the most recent phase and the model. The type of intent change was determined based on timing changes at the end of each gait phase (“Question A” and “Question B” in Fig. 3). The result is that an intent change can be signaled at any time step, but the type of intent change (SU/SD) can only be modified once per gait phase (four times per stride). When a gait phase observation has the same number of time steps as the modeled training data, there is always a Gaussian distribution model available to take the Mahalanobis distance. In this case, if the filtered Mahalanobis distance rises above the threshold value at any point, an intent change is flagged, and the algorithm assigns the type of intent change as the most recently indicated type.

To assess the intent identification algorithm’s performance, the identification delay was calculated for each trial as the amount of time after the intent command was issued that the algorithm first identified an intent change of the correct type. The number of true positive, true negative, false positive, and false negative identification outcomes were tabulated to assess accuracy. To compare results with the current estimated walking speed, the step speed was calculated at the beginning of every double-support phase based on the

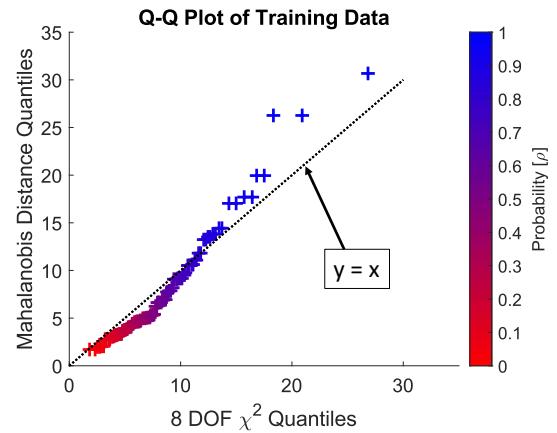


Fig. 4. Q-Q plot of sampled Mahalanobis distance & 8-DOF χ^2 distribution. Sample data marked by +.

geometric configuration of the exoskeleton and the time since the last heel-strike.

V. RESULTS

For trials with single user training data, AB subjects had an average of 47 non-outlying observations of each phase (max 52, min 39), and NAB subjects had an average of 74 (max 92, min 41). For multi-user training data, the cohort averages were 141 and 220 per phase for the AB and NAB cohorts, respectively (about 3 times each individual average).

The Q-Q plot for the sampled training data and the 8-DOF χ^2 distribution is shown in Fig. 4 for a single AB subject. When viewing all of the data, it is clear that the tail end does not align with the χ^2 distribution. For smaller values of the Mahalanobis distance (roughly less than 15, or the 0.9 quantile), there is much better alignment between the two distributions (Fig. 4). This level of agreement between the quantiles of the two sample sets indicates that the Gaussian assumption is likely appropriate for the majority of the data.

The Mahalanobis distance-based intent recognition output is summarized with the example in Fig. 5. The step speeds

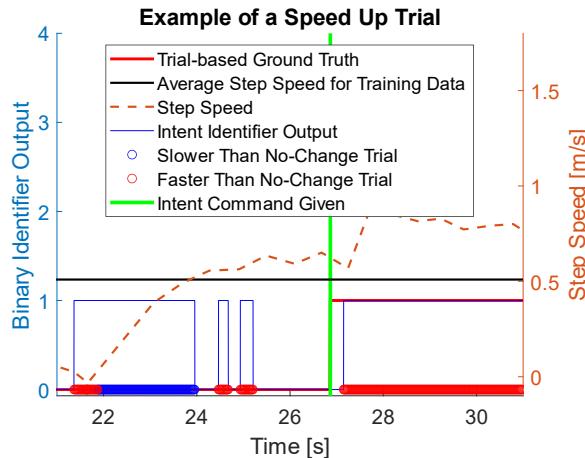


Fig. 5. Example intent identifier's binary output during AB SU trial consisting of 17 steps. Identifier trained on data only from this subject.

are plotted in orange against the y-axis on the right. The subject begins a trial from rest, so the first steps are quite slow. After getting up to speed, the subject reaches the average step speed from the training data, around 0.5 m/s. After the commanded intent change to speed up is given, as indicated by the “Trial-based Ground Truth” red line’s change from zero to one, the step speed increases slightly. The blue square wave showing the binary output of the intent identifier against the y-axis on the left indicates a difference between the new data and the model when the user first begins walking. The blue circles on the x-axis indicate that the steps are slower than the training data. In the middle of the trial, the identifier accurately determines that the user likely has the same intended gait speed as the training model by not indicating a difference. Several times before the commanded intent change, the identifier briefly identifies a speed-up intent change, and this is discussed further in the next section. Finally, after the intent command, the identifier recognizes a change and indicates that the user intends to walk faster than the model.

For each AB subject, the NC trials contributed on average 42 observations to each time step (max 53, min 9) to estimate each Gaussian model. For each NAB subject, the NC trials contributed on average 65 observations (max 105, min 10). NAB subjects had more observations because their smaller step lengths resulted in more steps per trial [16]. Table I shows the average time in seconds after the intent command was given that the intent identifier indicated an intent change of the correct type. These identification delays are also tabulated as a percentage of the average step duration for each subject group to indicate the likely number of steps from intent command to intent identification. Results are averaged by the types of user (AB/NAB), intent command (SU/SD), and AAD (crutches/walker). Training and testing were repeated for each of the four combinations of training set and indicator threshold.

The number of true/false positive/negative results are presented in Table II. Before an intent command is given, true negative and false positive are the only possible outcomes, so

these results are expressed as percentages of the total number of pre-command time steps. Similarly, the possible outcomes after an intent command is given are true positive and false negative, so these results were expressed as percentages of the total number of post-command time steps.

VI. DISCUSSION

A. Normality Assumptions

The Q-Q plot in Fig. 4 indicates that the training data are not entirely Gaussian distributed, but that the Gaussian assumption is valid even past the 0.75 quantile. In other words, at Mahalanobis distances lower than 15, the Gaussian assumption is likely appropriate, but not at larger values. This result does not invalidate the methodology because the intent change thresholds are set at the 0.5 quantile. Even in the region where the Gaussian assumption breaks down, the outliers of the Mahalanobis distance are above the $y = x$ line, which means that the identifier will consider high values to be at an even greater quantile than they truly are - magnifying the indication of an intent change.

B. Algorithm Structure

The intent identification algorithm successfully reported when users walked at speeds other than those in the training data and distinguished between speed-up and slow-down intentions. By indexing the model in time, the algorithm identified changes in both the shape of the data trajectories (via the Mahalanobis distance) and the timing (given by the speed-up/slow-down indicator).

C. Identification Delay

Trial numbers were insufficient to assess statistical significance of the values reported in Tables I & II (12 intent changes per subject). Trends in the results, however, suggest that the identification delay was always smallest for speed-up trials. This is surprising because previous analysis showed that subjects were able to reduce their speed more than they were able to increase their speed in the exoskeleton [16]. One possible explanation is that during speed-up trials, subjects significantly changed the trajectories of the onboard sensor measurements, but not in a way that produced a large increase in walking speed. This explanation highlights the promising possibility that this algorithm can identify an intent to change speed before the speed has actually changed. Another trend was that regardless of intent change type, identification delay was nearly always greater for walking with crutches than with a walker (except for AB subjects slowing down). Subjects indicated different AAD preferences during testing, but these preferences were not recorded. It is possible that walking with the non-preferred device increased variance and thus, negatively affected identification.

The identification delay was heavily dependent on the threshold for the change indication. For both individual and group models, the lower threshold of 0.5 resulted in reduced identification delay and correspondingly, reduced percentage of false negatives in Table II. It also, though, increased the number of false positives prior to the intent command being

TABLE I

AVERAGE TIME LAG FROM INTENT COMMAND, SPEED UP (SU) OR SLOW DOWN (SD), TO CORRECT INTENT CHANGE IDENTIFICATION FOR ABLE-BODIED (AB) & NON-ABLE-BODIED (NAB) SUBJECTS USING BOTH TYPES OF AMBULATORY ASSISTIVE DEVICE (AAD). TIMES ARE SHOWN FOR THE AVERAGE \pm ONE STANDARD DEVIATION AS WELL AS THE MEAN VALUE EXPRESSED AS A PERCENTAGE OF THE AVERAGE STEP DURATION.

Subject	AAD	Intent Change	Identification Delay [s]			
			Single User Training		Multi-user Training	
			Threshold at 0.6 [s \pm s, % avg. step dur.]	Threshold at 0.5 [s \pm s, % avg. step dur.]	Threshold at 0.6 [s \pm s, % avg. step dur.]	Threshold at 0.5 [s \pm s, % avg. step dur.]
AB	Walker	SU	0.15 \pm 0.12, 18%	0.10 \pm 0.11, 12%	0.35 \pm 0.26, 42%	0.27 \pm 0.21, 32%
		SD	0.68 \pm 0.24, 83%	0.68 \pm 0.24, 83%	0.82 \pm 0.42, 99%	0.81 \pm 0.43, 98%
	Crutches	SU	0.34 \pm 0.5, 41%	0.28 \pm 0.52, 34%	0.83 \pm 1.61, 101%	0.81 \pm 1.61, 98%
		SD	0.57 \pm 0.32, 69%	0.57 \pm 0.32, 69%	1.52 \pm 1.54, 184%	1.35 \pm 1.64, 164%
NAB	Walker	SU	0.22 \pm 0.32, 18%	0.19 \pm 0.30, 15%	0.73 \pm 0.58, 59%	0.22 \pm 0.21, 18%
		SD	1.15 \pm 0.87, 92%	1.04 \pm 0.87, 84%	1.42 \pm 0.96, 114%	1.34 \pm 0.98, 107%
	Crutches	SU	0.54 \pm 0.60, 43%	0.50 \pm 0.57, 40%	0.86 \pm 0.88, 69%	0.66 \pm 0.82, 53%
		SD	1.46 \pm 1.27, 117%	1.45 \pm 1.27, 116%	1.60 \pm 0.43, 128%	1.42 \pm 0.39, 114%

TABLE II

CONFUSION MATRIX VALUES (TRADITIONAL TRUE/FALSE (+)/(-) CLASSIFICATIONS DEFINED IN TEXT) FOR ABLE-BODIED (AB) & NON-ABLE-BODIED SUBJECTS (NAB) WITH INDIVIDUAL & GROUP TRAINING SETS USING TWO THRESHOLDS.

Subject	Training Data	Threshold	Speed Up				Slow Down			
			True (-)	False (+)	True (+)	False (-)	True (-)	False (+)	True (+)	False (-)
AB	Individual	0.6	44%	56%	88%	12%	47%	53%	94%	6%
		0.5	36%	64%	93%	7%	38%	62%	96%	4%
	Group	0.6	56%	44%	78%	22%	59%	41%	78%	22%
		0.5	45%	55%	85%	15%	49%	51%	84%	16%
NAB	Individual	0.6	55%	45%	64%	36%	54%	46%	65%	35%
		0.5	47%	53%	71%	29%	46%	54%	70%	30%
	Group	0.6	65%	35%	53%	47%	61%	39%	50%	50%
		0.5	56%	44%	60%	40%	50%	50%	59%	41%

given. The lower threshold reduces the deviation from the training set at which the identifier will allow the gait to be considered the same. Since it is possible that subjects walked faster than the training data even before the intent command to speed up was given, the identifier is possibly more accurate in identifying gait differences from the “trial-based ground truth”. After the intent command is given, however, there has definitely been an intent change. There should be greater confidence, then, given to results for true positives or false negatives than for false positives and true negatives.

The two sets of training data also had an effect on the identification delay. The model trained on all of the subjects generally had a larger identification delay and resulted in fewer positive results overall, both true and false. In fact, the multi-user training experiments each resulted in four trials in which the correct intent change type was never identified (two each for AB and NAB subjects). Both of these trends suggest that the increased variance in the training data set caused by the inclusion of data from multiple users caused the identifier to be less likely to indicate an intent change even when it would be accurate to do so.

The identification delay was generally greater for NAB subjects. This trend is likely due to increased step-to-step variance for NAB subjects in both the training and testing data. For example, for a single combination of p and t , the trace of the covariance matrix of the model for an AB subject was 45, yet for an NAB subject was 72. Increased training data variance reduces the likelihood the identifier will signal an intent change, and increased testing data variance reduces

the accuracy of the flagged type of intent change. With a maximum delay of 1.6 seconds for multi-user training, the delay for NAB subjects was not excessive since 1.6 seconds represents 1.28 steps at the average step speed for these subjects. This result is particularly exciting, considering the challenges surrounding intent recognition for the SCI population. The only commercially available exoskeleton that considers user intent beyond shifting between finite state controllers is the HAL exoskeleton, which initiates gait cycles based on EMG [6]. Even so, it requires that users present motor function of the hip flexors and knee extensors, which is often not the case for individuals with complete SCI [19].

It is unclear in current literature what an appropriate or allowable delay in intent recognition would be when coupled with the intensity of the robot’s response and the distribution of false positives and false negatives. In terms of practical algorithm implementation, the threshold for intent change may need to be tuned to the user to achieve a comfortable level of device responsiveness. While the device may begin with a baseline set of training data for the given user type, it will need to prioritize user-specific data to achieve the desired level of identification delay and accuracy.

VII. ONGOING WORK

A key limitation of this algorithm is that the type of intent change can only be identified once per gait phase. An alternative to monitoring the timing offset of the most recently ended phase would be to build a library of the

expected sensor values for gaits at different speeds. When an intent change is indicated, the data could then be compared to the library to determine which new gait is most likely the intended gait. This strategy, however, would require gathering a large amount of training data to build the models of each gait speed, a task that may be impossible for an individual with SCI, but might be possible when pooling data across multiple individuals with user-specific tuning happening online. It is also not clear whether the sensor values for a user resisting the current gait speed would match those when the user and exoskeleton are walking in agreement at a different speed. An approach without these limitations would be to build only two secondary models, one each for the expected model-data error for a given type of intent change. Then, when an intent change occurs, the model-data difference could be compared to both the speed-up model and the slow-down model. The comparison with the smallest Mahalanobis distance would indicate which intent change type is most likely, and the timing offset at the end of the gait phase could confirm. This strategy would allow the type of intent change to be toggled at any time step, would require minimal additional training data, and would involve fusing two different indicators of the type of change.

Another limitation of this algorithm is that the model remains the same for all comparisons after training. Updating the model by adding in online data points (according to (2) and (3)) every time an intent change occurs would allow the algorithm to indicate an intent change, acclimate to the new walking speed, and then indicate a subsequent intent change in the same direction as the first. This ability to make subsequent speed changes of the same type would allow the user to achieve any walking speed as opposed to only being able to walk at three unique speeds (the same as, faster, or slower than the training set).

Finally, the algorithm, current or improved, requires validation in real-time experiments. Experiments conducted on a treadmill would facilitate recording of true gait speed and assessing the algorithm's sensitivity with respect to the size of the gait speed change. Once the gait intentions are identified, the next step is to figure out how the exoskeleton should respond. For instance, if the algorithm indicates that the user would like to speed up, what combination of increased step frequency and step length would be most comfortable for the user? How much should the device aim to speed up and over what time period should the speed increase be achieved? Finally, the gait trajectories that accomplish these actions will need to be calculated such that the speed transition is smooth and continuous, a task which was addressed by [20] using Fourier series and profile blending techniques. Future work will focus on tuning the exoskeleton response, the indicator threshold value, and model training scheme to maximize user comfort with respect to device responsiveness.

In its preliminary form, the intent identification algorithm based on the Mahalanobis distance has proven to be computationally simple, successful with minimal training data, and potentially viable for individuals with spinal cord injury. The methodologies described herein are platform independent,

though, so they provide an avenue for incorporating user intent information already collected by onboard sensors into intent-aware control schemes for any active assistance device, possibly even prostheses.

REFERENCES

- [1] D. Losey, C. McDonald, E. Battaglia, and M. O'Malley, "A Review of Intent Detection, Arbitration, and Communication Aspects of Shared Control for Physical Human – Robot Interaction," *Appl Mech Rev*, vol. 70, no. 1, pp. 1–19, 2019.
- [2] J. Schmiedeler and P. Wensing, "Discussion of "A Review of Intent Detection, Arbitration, and Communication Aspects of Shared Control for Physical Human – Robot Interaction","" *Appl Mech Rev*, vol. 70, no. 1, pp. 015 503–1–015 503–3, 2018.
- [3] S. Banala, S. Agrawal, S. H. Kim, and J. Scholz, "Novel gait adaptation and neuromotor training results using an active leg exoskeleton," *IEEE/ASME Trans Mechatron*, vol. 15, no. 2, pp. 216–225, 2010.
- [4] I. Khanna, A. Roy, M. Rodgers, H. Krebs, R. Macko, and L. Forrester, "Effects of unilateral robotic limb loading on gait characteristics in subjects with chronic stroke," *J Neuroeng Rehabil*, vol. 7, no. 23, 2010.
- [5] M. Tucker, J. Olivier, A. Pagel, H. Bleuler, M. Bouri, O. Lambery, J. Millán, R. Riener, H. Vallery, and R. Gassert, "Control strategies for active lower extremity prosthetics and orthotics: a review," *J Neuroeng Rehabil*, vol. 12, no. 1, p. 1, 2015.
- [6] J. Contreras-Vidal, N. Bhagat, J. Brantley, J. Cruz-Garza, Y. He, Q. Manley, S. Nakagome, K. Nathan, S. Tan, F. Zhu, and J. Pons, "Powered exoskeletons for bipedal locomotion after spinal cord injury," *J Neural Eng*, vol. 13, no. 3, pp. 031 001–0 310 017, 2016.
- [7] H. Vallery, E. Van Asseldonk, M. Buss, and H. Van Der Kooij, "Reference trajectory generation for rehabilitation robots: Complementary limb motion estimation," *IEEE T Neur Sys Reh*, vol. 17, no. 1, pp. 23–30, 2009.
- [8] M. Ison and P. Artermiadis, "The role of muscle synergies in myoelectric control: Trends and challenges for simultaneous multifunction control," *J Neural Eng*, vol. 11, no. 5, 2014.
- [9] Z. Lu, K.-y. Tong, H. Shin, S. Li, and P. Zhou, "Advanced myoelectric control for robotic hand-assisted training: Outcome from a stroke patient," *Front in Neurol*, vol. 8, no. 3, pp. 1–5, 2017.
- [10] Y. He, D. Eguren, J. Azorín, R. Grossman, T. Luu, and J. Contreras-Vidal, "Brain-machine interfaces for controlling lower-limb powered robotic systems," *J Neural Eng*, vol. 15, no. 2, 2018.
- [11] J. Kang, U. Park, V. Gonuguntla, K. Veluvolu, and M. Lee, "Human implicit intent recognition based on the phase synchrony of EEG signals," *Pattern Recogn Lett*, vol. 66, pp. 144–152, 2015.
- [12] N. Bhagat, A. Venkatakrishnan, B. Abibullaev, E. Artz, N. Yozbatiran, A. Blank, J. French, C. Karmonik, R. Grossman, M. O'Malley, G. Francisco, and J. Contreras-Vidal, "Design and optimization of an EEG-based brain machine interface (BMI) to an upper-limb exoskeleton for stroke survivors," *Front Neursci*, vol. 10, no. 3, 2016.
- [13] U. Nagarajan, G. Aguirre-Ollinger, and A. Goswami, "Integral admittance shaping: A unified framework for active exoskeleton control," *Robot Auton Syst*, vol. 75, pp. 310–324, 2016.
- [14] N. Karavas, A. Ajoudani, N. Tsagarakis, J. Saglia, A. Bicchi, and D. Caldwell, "Tele-impedance based assistive control for a compliant knee exoskeleton," *Robot Auton Syst*, vol. 73, pp. 78–90, 2015.
- [15] D. Ragonesi, S. K. Agrawal, W. Sample, and T. Rahman, "Quantifying anti-gravity torques for the design of a powered exoskeleton," *IEEE T Neur Sys Reh*, vol. 21, no. 2, pp. 283–288, 2013.
- [16] T. Gambo, J. Schmiedeler, and P. Wensing, "Characterizing intent changes in exoskeleton-assisted walking through onboard sensors," *Proc IEEE Int Conf Rehab Rob*, vol. 6, pp. 471–476, 2019.
- [17] T. Tango, *Applied Multivariate Statistics with R*, 2010.
- [18] F. A. Zakaria, M. El-Badaoui, S. Maiz, F. Guillet, M. Khalil, K. Khalil, and M. Halimi, "Walking analysis: Empirical relation between kurtosis and degree of cyclostationarity," *ICABME 2013*, pp. 93–96, 2013.
- [19] O. Jansen, D. Grasmuecke, R. Meindl, M. Tegenthoff, P. Schwenkreis, M. Sczesny-Kaiser, M. Wessling, T. Schildhauer, C. Fisahn, and M. Aach, "Hybrid Assistive Limb Exoskeleton HAL in the Rehabilitation of Chronic Spinal Cord Injury: Proof of Concept; the Results in 21 Patients," *World Neurosurg*, vol. 110, pp. e73–e78, 2018.
- [20] S. Y. Shin and J. Sulzer, "An Online Transition of Speed-dependent Reference Joint Trajectories for Robotic Gait Training," *Proc IEEE Int Conf Rehab Rob*, pp. 983–987, 2019.