Improving Human-in-the-Loop Optimization Algorithms using Machine Learning Techniques

Thu Nguyen, *tmn33@stanford.edu* Katherine Poggensee, *ktpognc@stanford.edu*
and Guan Rong Tan, *guanrong@stanford.edu*

Abstract—Assistive device technology is rapidly improving with the advent of *human-in-the-loop optimization* algorithms. These algorithms find an optimal set of parameters to reduce the energy cost of movement. However, energy cost is a difficult quantity to measure and is thus the limiting factor in the optimization time. In an effort to both improve the time to optimize parameters and to take the device out of the lab, we were interested in predicting cost from other easier-to-measure features. We fit models to both update online in a real-world application (i.e., linear regression) and capture the complexity of the data (i.e., neural networks). Each model performed reasonably well, and the feature and example sets were further adjusted to test different aspects of the complexity in the human-machine system. To counteract the high bias, the feature space was reduced in the linear models. To understand the contribution of the intersubject variability, neural networks were trained both for the general population and individual subjects. These models show promise for use in these different applications, but more data needs to be collected before we can implement this on the real system.

1 Introduction

Roboticists have been trying to assist human movement for decades with very little success [1], [2]. Devices are created to assist patients with movement disorders regain normal function [3] and to improve therapy [4]. Robotic augmentation is also an area of great interest [5].

Designing effective assistive devices poses different challenges to designing effective control for a robot. Humans are variable both between subjects and even within a single subject [6], [7], [8]. Although models can capture some basic features of movement in spite of this variability [9], [10], the complexity of the human is too vast to accurately predict the response to a novel assistive device. Researchers have previously attempted to use these and other physics-based models to assist individuals [11], again with little success. The first assistive devices to decrease the energy cost of walking were published earlier this decade in rapid succession [5], [12], [13]. These algorithms were the result of intensive biomechanics research and meticulous hand-tuning. The time it takes to hand-tune parameters is unwieldy, especially when working with patient populations who cannot walk as long as healthy subjects.

Human-in-the-loop optimization strategies have had recent success in improving the energy cost of walking using assistive devices [14], [15]. Using stochastic, model-free optimization algorithms, the device parameters are optimized specifically to reduce the energy cost of walking. The parameters are adjusted online using a device emulator system [16]. The energy is measured as the human walks and is used in the cost function of the optimization algorithm, which guarantees the efficacy of the parameters and allows for a customized set of parameters for each subject. These processes are sample-efficient and have shown large success in healthy populations. The number of samples can be reduced through better optimization methods, but is constrained by the time it takes to quantify the energy cost of walking. Energy cost is estimated by fitting a first-order dynamical model [17] to two minutes of metabolic data. The time cost to getting one measurement then is 2 minutes, as a lower bound if metabolic energy is used in the cost function.

This project was proposed to find alternative strategies to estimating metabolic energy. Predicting metabolic energy through other means could reduce the amount of time it takes to measure the cost of each control parameter set and could also allow for the commercialization of these devices by instead measuring features that can be obtained through the device or other wearable sensors. We used biomechanics and device data from a similar human-in-the-loop optimization study to find models (i.e., linear models and neural networks) for the energy cost of each control parameter set.

2 Related Work

Metabolic cost is often held as the target metric for assistive devices [5], [12], [13], [14]. This likely comes from the idea that humans naturally tune their mechanics in energy-optimal ways [18]. Metabolic cost is difficult to measure: it is noisy and has slow dynamics [17]. Many devices, therefore, use a proxy to define control strategies [19]. Some initial research has been done to estimate metabolic cost using wearable devices [20]. While these models show promise in estimating metabolic cost using simple proxies, it is not guaranteed that they will be valid when a device is introduced. Biomechanics of movement change with the addition of any device [21], including simple devices like a treadmill [22].

There have been some successes in fitting models to other areas of movement, such as kinematic variables [8], [23], [24], [25], [26] and muscle activity [27], [28]. Predicting kinematic measures using simple models like linear regression [23] work well because subjects follow reliable trajectories that obey the laws of physics. We had previously attempted to use linear regression to predict muscle activity, similar to [23], with little success ($R^2 < 0.3$). Physiological

features, like muscle activity or metabolic cost, tend to be more difficult to predict with simple models because the complex dynamics are still not well-understood.

These methods illustrate the efficacy of machine learning algorithms to biomechanics data, even if they are not directly applicable. Supervised learning techniques have primarily resulted in classification techniques rather than continuous models [26], [27]. Reinforcement learning can generate control of assistive devices to modulate muscle activity [28]. While these results indicate that more sophisticated algorithms can be effective in describing human biomechanics, they fall outside of the scope of our project. We want to continue to use the existing human-in-the-loop optimization algorithm, so, with no actions to take, a reinforcement learning framework does not suit our problem. The optimization algorithm [29] uses a rank ordering update rather than a gradient descent approach on a continuous cost function. This is beneficial when working with noisy human data. Theoretically, there are ways that we could take advantage of this ordering and appropriately bin our targets to fit a classification model, but the noise in the target measurements is too large for this to be feasible.

3 Dataset and Features

The success of human-in-the-loop optimization strategies [14] prompted a follow-up study by the authors to determine the mechanisms behind the success. The idea behind the study was presented at the *Dynamic Walking* conference this past summer [30].

Subjects ($N = 10$) completed at least 36 trials as the exoskeleton control parameters were optimized. These control parameters, described in greater depth in [14], defined the assistance trajectory from the device. The four parameters were

- Peak magnitude of assistance (α_1)
- Timing of the peak assistance (α_2)
- Rise time, from 0 Nm to peak (α_3), and
- Fall time, from peak to 0 Nm (α_4).

These parameters are illustrated in Figure 1. Metabolic energy consumption, measured through indirect respirometry, was measured and used to optimize the control parameters. Standard biomechanics measurements were also collected. The data were analyzed and processed during the course of this project.

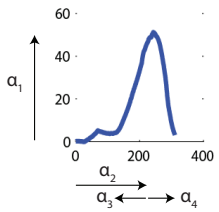

Fig. 1. Illustration of device control parameters. The magnitude (α_1) and timing (α_2) of the peak assistance define that node of the assistance trajectory. The time for the curve to rise from 0 Nm to peak assistance (α_3) and the time for the curve to fall from peak assistance to 0 Nm (α_4) define the shape of the curve.

There are some basic biomechanics terms needed to accurately describe our feature set. A single *stride* is typically defined from the point that the subject's foot strikes the ground until that same foot strikes the ground again. The period when the foot is on the ground is called *stance*, and the period when the foot is in the air is called *swing*. The motions of the ankle are primarily in one plane. These motions are called plantarflexion and dorsiflexion; *plantarflexion* is the motion where the toe points into the ground, and *dorsiflexion* is the opposite. The *vertical ground reaction forces* are the reaction forces of the ground on the subject; they are often characterized by two peaks during stance, with the valley at the middle of stance.

For this project, we used a subset of the measured data. The target labels were the energy cost associated with each set of parameters. The feature variables were the 4 control parameters described above and a selection of biomechanics measurements:

- Peak ankle plantarflexion angle
- Peak ankle dorsiflexion angle
- Ankle angle at the transition from stance to swing
- Vertical ground reaction force of the second peak at the end of stance
- Timing of that peak of ground reaction force
- Time spent in stance
- Time of a single stride, and
- Asymmetry between legs in stance and stride times [31].

Each one of these features was calculated at every stride on each leg (with the exception of asymmetry). The measurements at each stride were averaged for each set of control parameters. The timing of the peak ground reaction force and the time spent in stance were normalized to the stride time on that stride, so the values were nondimensional. For consistency with reporting standards in the literature [14], peak assistance magnitude (α_1) and the vertical ground reaction force magnitude were normalized by subject's body mass.

In total, we had $n = 20$ features. There were $m = 392$ total examples. For the linear regression analyses, $m = 294$ training examples and 98 test examples were used. The same distribution was used for the generalized neural network, and for the single subject neural network, $m = 38$ training examples and 9 test examples were used.

4 Methods and Experiments

4.1 Linear Regression

A recent study used linear regressions to predict the energy cost of walking from physiological and mechanical data collected from able-bodied individuals walking on a treadmill at different speeds on flat ground, inclined, and backwards in normal walking shoes [20]. Four linear models were trained on different feature subsets. The best set of features included heart rate, electrodermal activity, skin temperature, muscle activity, and accelerometer magnitudes which yielded an $R^2 = 0.93$.

We therefore chose to train multiple linear regression models using subsets of our features. We used features measured in the lab, rather than those from widely-marketable

sensors [20]. Because humans change their movement in response to device behavior [21], we were interested in first determining if subjects' energy cost correlated with these features, sensed with the highest precision, before fitting models with noisy data from mid-range sensors.

During our initial analysis, we used unweighted linear regression to try to fit a linear function relating the features to the target metabolic rate. By minimizing the least squares error $||\overrightarrow{y} - X\overrightarrow{\theta}||_2^2$, we have

$$X^T X \overrightarrow{\theta} = X^T \overrightarrow{y}. \tag{1}$$

Solving for θ, we obtained the closed form solution:

$$\overrightarrow{\theta} = (X^T X)^{-1} X^T \overrightarrow{y}. \tag{2}$$

We first trained a model using all of the features ($n = 20$). We then trained two additional models using the features selected through principal component analysis and neural networks.

4.2 Neural Networks

Unlike linear regression models, which can be trained online and for which there is a closed-form solution, we used a neural network to capture the complexity of the human-machine dynamics. Neural networks are non-linear with a large parameter space, allowing the possibility of a more complicated relationship between the features and the output.

In a neural network, features are separated as individual neurons in an input layer. These features are then propagated through subsequent hidden layer(s) and an output layer. During forward propagation, all outputs from the previous layer are individually passed to each of the neurons in the following layer, and manipulated by an activation function. The outputs from these activation functions are then passed on to the next layer and manipulated again by the new layer's activation functions. At the final layer (i.e., the output layer), the resulting values are the model prediction for the target.

A previous study had successfully used neural networks to identify human gait patterns using similar biomechanical features such as limb length and joint angle data [26]. They used a single hidden layer of 28 neurons with sigmoid activation functions and an output layer with linear nodes. They also had 10 gait features, and around 210 training examples. This was of a similar scale to our dataset containing 20 features and 294 training examples, which also required a continuous output (rather than classification with softmax). Hence, we also initialized our model to a single hidden layer with 28 neurons with sigmoid activation functions as follows:

$$\vec{a}^{[1]} = \frac{1}{1 + e^{-(W^{[1]}\vec{x} + B^{[1]})}} \tag{3}$$

and an output node with a linear regression activation function.

We eventually found optimal results when we tuned our network down to 15 features and 6 neurons. The process behind optimizing our model is described more rigorously in the following section.

4.3 Dimensionality Reduction

We reduced the dimensionality of the feature space using a modified principal component analysis and analysis of the neural network. We did not use any formal feature selection algorithms as they were outside of the scope of this class.

Because we are interested in reducing the number of sensors required to estimate metabolic cost, we ran principal component analysis (PCA) on our features to understand which contributed the most to the variance in the data. The human-in-the-loop optimization algorithm on this noisy data performs best when there is a certain degree of variance between the cost measures, so PCA seemed the logical choice. The PCA algorithm that we used first centered the data to a zero mean and then used singular value decomposition (SVD). We used SVD because our feature matrix was close to singular with each subjects' asymmetry very close to 0. In SVD, X is decomposed into $X = U\Sigma V^T$. In this analysis, V gives the principal directions and $U\Sigma$ gives the principal components. We used the MATLAB function `pca`.

We also analyzed the $W^{[1]}$ weights from our neural network to see which features were most highly weighted to reduce the number of features. After training the neural network, we summed the weights for each feature across all neurons and used those that had the greatest total weight.

5 Results and Discussion

TABLE 1
Results of the learned models. The linear regression models (LR) were computed on all subjects' data using a differing number of features; the neural networks (NN) were trained on all of the features for all and one subjects' data.

Model	Training			Test		
	m	MSE	R	m	MSE	R
LR	294	0.2196	0.81	98	0.4546	0.69
LR ($n = 5$)	294	0.4003	0.66	98	0.3787	0.63
NN (All subjects)	294	0.0667	0.95	98	0.3132	0.74
NN (1 subject)	38	0.0164	0.98	9	0.0749	0.95

5.1 Evaluation Metrics

We evaluated the performance of our models using mean squared error (MSE) and coefficient of correlation (R). We used MSE instead of root mean squared error because MSE is more analogous to variance than standard deviation. MSE was calculated as:

$$MSE = \frac{1}{m}\sum_{i=1}^{m}(Y_i - \hat{Y}_i)^2 \tag{4}$$

where m is the number of samples, Y_i is the measured metabolic cost, and $\hat{Y}_i$ is the predicted metabolic cost.

We used R instead of the coefficient of determination (R^2) because we wanted to know how well the target and predicted metabolic cost were correlated rather than the exact relationship between the two. R was calculated as:

$$R = \frac{\text{cov}(\hat{Y}, Y)}{\sigma_{\hat{Y}}\sigma_Y} \tag{5}$$

where cov is the covariance between the measured and predicted metabolic cost, $\sigma_{\hat{Y}}$ is the standard deviation of the

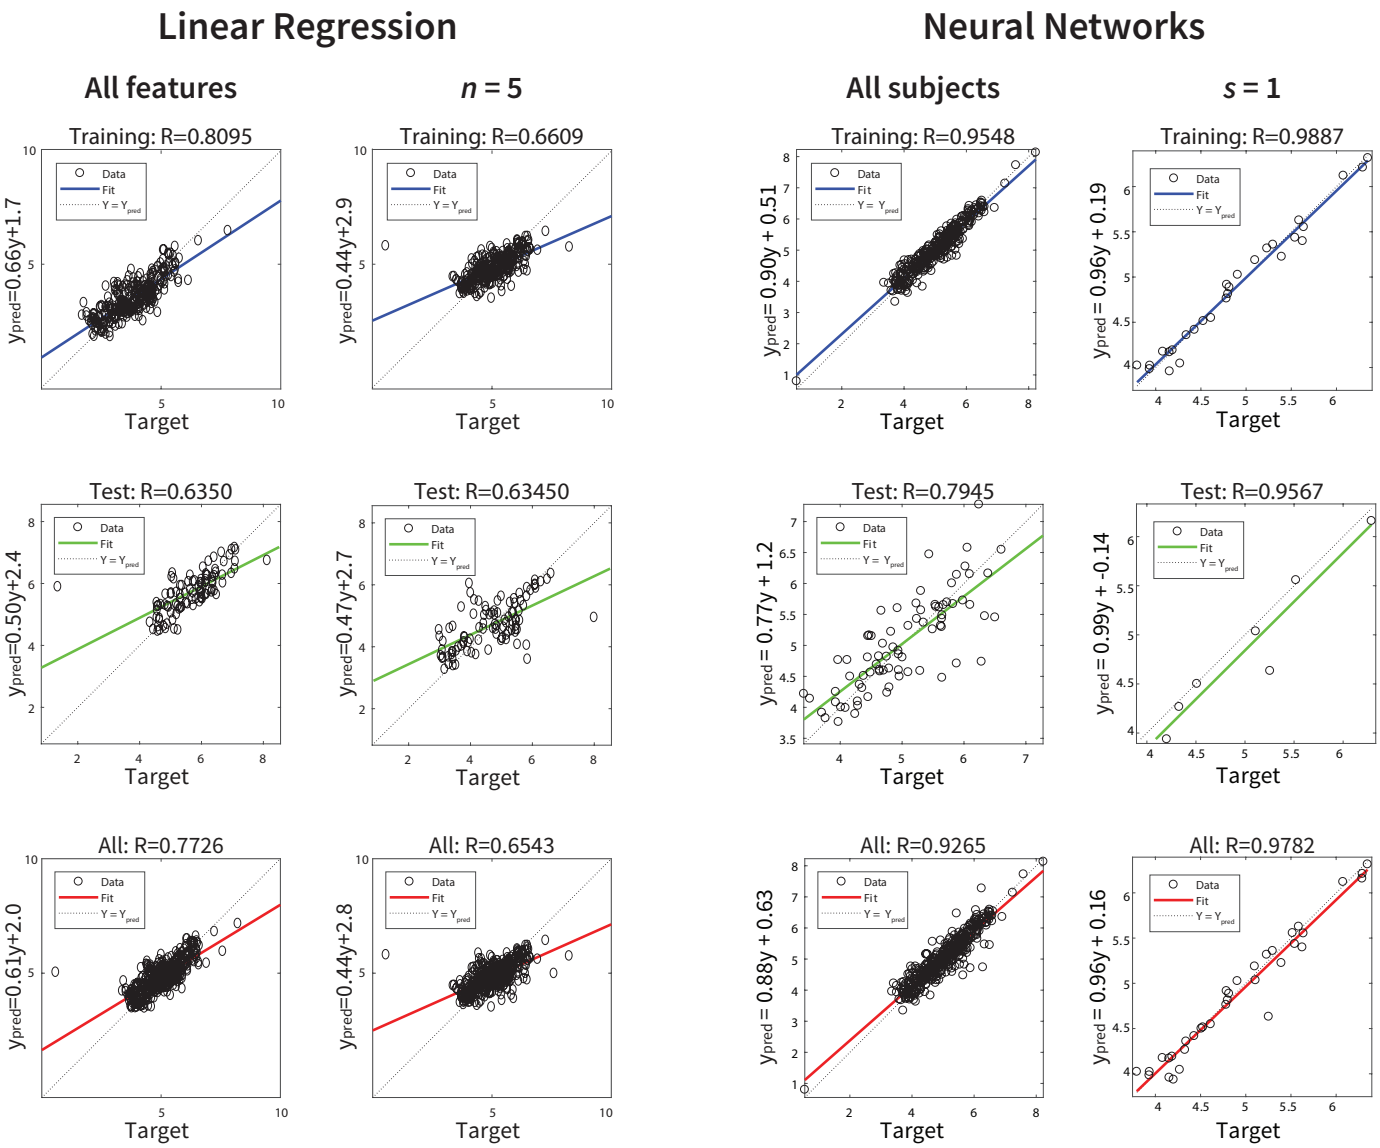

Fig. 2. Accuracy plots of the different models. For a given model, the predicted target values y_{pred} and the target values y were regressed. The resulting linear model is stated on the vertical axis of each plot. For a perfect fit, the model would be $y_{pred} = y$. The individual data points are also shown in the plot, as is the ideal line ($y = y_{pred}$). For each model, the training, test, and final accuracies are shown with corresponding metric. We used the correlation coefficient to determine correlation as a measurement of accuracy rather than a regressed fit. The two high-level models were linear regression and neural networks. Linear regressions were fit using all of the features and a selected subset; neural networks were trained on all subjects and an individual subject.

predicted metabolic cost, and σ_Y is the standard deviation of the measured metabolic cost.

We trained all of our models on a randomized 75% of the sample set. We then tested our models on the remaining 25% and ran a final test on the entire sample set. Historically, 70% of the sample set is used for training when there are less than a few thousand datasets. However, with our smaller dataset, we wanted to make sure we had enough training data to prevent overfitting, especially for the neural network, so we increased the percentage to 75%. We also tried other percentages and found that 75% yielded the best results.

5.2 Linear Regression

The first linear model that we trained used all of the features. During training, the algorithm predicted metabolic cost with a MSE of 0.22 with a strong correlation between the predicted and measured metabolic cost ($R = 0.81$). When we tested this algorithm, the MSE increased to 0.4546 while the correlation coefficient decreased to 0.69 (Table 1). The large discrepancy between performance indicated that we overfit our model to the training set and therefore had high bias. To reconcile this, we chose to decrease the number of features used in the linear regression models.

Principal component analysis (PCA) was used to find which features contributed the most to the variance. The first principal component explained 99.98% of the variance in the data. In this principal component, all feature coefficients were close to 0, except for the timing of peak ground reaction force of both legs which had values of 0.7073 and 0.7069, respectively. We then ran a linear regression on these two main features but found that the predicted and measured metabolic costs were not well-correlated in the

training set ($R = 0.15$, MSE = 0.79) and test set ($R = 0.16$, MSE = 0.6311). While the timing of the peak ground reaction force was the most variable across control parameters and subjects, these features did not contribute to changes in metabolic cost. Because PCA uses the variance of the features and not their contribution to the variance in metabolic rate, using PCA to reduce the number of features was not an effective method. We therefore looked at neural networks to understand which features contributed to the variance of metabolic cost.

After training the neural network, we found that the five features with the largest weights were:

1) Peak magnitude of assistance,
2) Timing of the peak assistance,
3) Left ankle angle at the transition from stance to swing,
4) Right ankle angle at the transition from stance to swing, and
5) Stance time of the right leg.

Using these features, we ran another linear regression and found that the training and test sets performed similarly. The performance of the training set was $R = 0.66$ and MSE = 0.4004 while the testing set was $R = 0.63$ and MSE = 0.3787 (Table 1). Since the training and test sets yielded similar results, we were able to successfully reduce overfitting.

Previous studies indicate that assisting subjects during a period called *push-off* will yield the greatest reductions in energy cost [5], [12], [13], [14]. In fact, these results were the basis for the parameterization of the control [14]. This period is the time during stance that the foot exerts the greatest force to propel the foot forward [32]. Thus, it is unsurprising that the timing of the peak ground reaction force contributed the largest variance to the feature set. We observed during experiments that the optimization algorithm was determined by the peak magnitude of assistance and the timing of that assistance, so it was expected that these control parameters contributed the greatest total weight to the energy cost. We also noticed in this particular experiment the ankle angle at the transition from stance to swing had an effect on the performance of the hardware, so, while this may not have an effect in general, for this experiment and this hardware, we believe that subjects were unconsciously aware of this angle and modulated it accordingly. Most people are right-dominant and thus tend to favor this side in an uncomfortable environment [31]. Hence, changes in stance time on the dominant leg may be correlated with changes in energy cost.

The slope of the accuracy plots ranged from 0.44 to 0.66 where the target slope for a perfect prediction was 1 (Figure 2). However, when looking at the majority of the data, it appears as though most of the data lies around the $y = y_{pred}$ line with a few outliers strongly affecting the slope of the best-fit line. We know from previous literature that observed metabolic cost tends to lie in a small range of values [14], illustrated by the dense cluster of points. The fact that the linear regression is accurate within this dense cluster is more promising than the correlation coefficient or the slope of the best-fit line.

5.3 Neural Networks

The initial setup of our neural network with 1 hidden layer of 28 neurons and 20 features yielded results that suggested high bias in our model, as our test sets yielded starkly worse results than the training set. We pruned the number of features by a weighting metric, described in Section 4.3, until the performance decreased. We also reduced the number of neurons in the hidden layer to 6 neurons by the same method, to yield the best performance.

Based on literature [6] and qualitative observations during the experiment, we were also concerned about the variability between subjects. Following the reporting standards in the literature, some features were normalized by body mass (Section 3). However, variability is not easy to model [6], [7] and should not be treated as noise. To determine if the subject variability had a significant effect on the results, we trained separate neural networks for each individual subject and averaged the results over all the subjects. This resulted in a huge improvement in results, especially in the training accuracy. As seen in Table 1, our MSE errors and R values saw decent improvements dropped from 0.0667 to 0.95 and 0.95 to 0.97 in the training set. However, the most drastic improvements were in the test set where MSE and R values improved from 0.3132 to 0.0749 and 0.74 to 0.95, respectively.

This drastic improvement suggests that we cannot assume a generalized model for this population. The only identifying traits collected for subjects were mass and subject number. Mass is already used to normalize the data, with results discussed above. Subject number as a feature may improve the performance of the model, but, as it has no bearing on the individual, would make classifying new subjects impossible.

The experiment from which the data were taken was designed to understand the role of customization in the control algorithm [14], [30]. Therefore, the improvement in the neural network performance is an interesting result that confirms our hypothesis, and the issue of individualization will be considered outside of the scope of this project.

6 Conclusions

We used two different machine learning architectures to predict metabolic cost from standard biomechanics features. As in existing literature [20], we trained linear regression models to generate easily implementable models that could be used and trained online. These regressions performed well, considering the noise in the target measurements, and performance was not drastically affected when the feature size was reduced to one-fourth the size. Neural networks provided a more accurate model, but would be difficult to train online on a commercial device. These models could be used in a clinic, where having accurate subject-specific models are more important. Biomechanics experiments are time-consuming, so the data size is often small ($N = 10$). Before we implement these predictive models on commercial devices, we will collect more subjects' data to improve the training of the model.

CONTRIBUTIONS

All authors contributed equally to this project, poster, and report.

- Thu Nguyen - Data extraction, linear regression, and PCA
- Katherine Poggensee - Data collection, processing, extraction, and overall guidance on the project and algorithms
- Guan Rong Tan - Data extraction and neural networks

REFERENCES

[1] R. Mosher, "Handyman to Hardiman," *SAE International*, p. 10, 1967. [Online]. Available: http://papers.sae.org/670088/

[2] H. Kazerooni, R. Steger, and L. Huang, "Hybrid Control of the Berkeley Lower Extremity Exoskeleton (BLEEX)," *The International Journal of Robotics Research*, vol. 25, no. 5-6, pp. 561–573, 2006. [Online]. Available: http://journals.sagepub.com/doi/10.1177/0278364906065505

[3] S. H. Collins and A. D. Kuo, "Recycling energy to restore impaired ankle function during human walking," *PLoS ONE*, vol. 5, no. 2, 2010.

[4] J. Hidler, D. Nichols, M. Pelliccio, K. Brady, D. D. Campbell, J. H. Kahn, and T. G. Hornby, "Multicenter randomized clinical trial evaluating the effectiveness of the Lokomat in subacute stroke." *Neurorehabilitation and neural repair*, vol. 23, no. 1, pp. 5–13, jan 2009.

[5] S. H. Collins, M. B. Wiggin, and G. S. Sawicki, "Reducing the energy cost of human walking using an unpowered exoskeleton," *Nature*, vol. 522, pp. 212–215, 2015. [Online]. Available: http://www.nature.com/doifinder/10.1038/nature14288

[6] D. A. Winter, "Kinematic and kinetic patterns in human gait: Variability and compensating effects," *Human Movement Science*, vol. 3, no. 1-2, pp. 51–76, 1984.

[7] T. M. Owings and M. D. Grabiner, "Measuring step kinematic variability on an instrumented treadmill: How many steps are enough?" *Journal of Biomechanics*, vol. 36, no. 8, pp. 1215–1218, 2003.

[8] S. Revzen and M. Kvalheim, "Data driven models of legged locomotion," in *SPIE Defense+ Security*, vol. 9467, 2015, p. 94671V.

[9] S. Kajita, F. Kanehiro, and K. Kaneko, "The 3D Linear Inverted Pendulum Mode: A simple modeling for a biped walking pattern generation," in *International Conference on Intelligent Robots and Systems*, no. 4, Maui, Hawaii, USA, 2001, pp. 239–246. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=973365

[10] S. W. Lipfert, M. Günther, D. Renjewski, S. Grimmer, and A. Seyfarth, "A model-experiment comparison of system dynamics for human walking and running," *Journal of Theoretical Biology*, vol. 292, pp. 11–17, 2012.

[11] H. van der Kooij, B. Koopman, and E. H. van Asseldonk, "Body weight support by virtual model control of an impedance controlled exoskeleton (LOPES) for gait training." *2008 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pp. 1969–1972, 2008.

[12] P. Malcolm, W. Derave, S. Galle, and D. De Clercq, "A Simple Exoskeleton That Assists Plantarflexion Can Reduce the Metabolic Cost of Human Walking," *PLoS ONE*, vol. 8, no. 2, pp. 1–7, 2013.

[13] L. M. Mooney, E. J. Rouse, and H. M. Herr, "Autonomous exoskeleton reduces metabolic cost of human walking during load carriage," *Journal of NeuroEngineering and Rehabilitation*, vol. 11, no. 1, p. 80, 2014.

[14] J. Zhang, P. Fiers, K. A. Witte, R. W. Jackson, K. L. Poggensee, C. G. Atkeson, and S. H. Collins, "Human-in-the-loop optimization of exoskeleton assistance during walking," *Science*, vol. 356, no. 6344, pp. 1280–1284, jun 2017. [Online]. Available: http://www.sciencemag.org/lookup/doi/10.1126/science.aal5054

[15] M. Kim, Y. Ding, P. Malcolm, J. Speeckaert, C. J. Siviy, C. J. Walsh, and S. Kuindersma, "Human-in-the-loop Bayesian optimization of wearable device parameters," *PLOS ONE*, vol. 12, no. 9, p. e0184054, sep 2017. [Online]. Available: http://dx.plos.org/10.1371/journal.pone.0184054

[16] J. Caputo and S. Collins, "An experimental robotic testbed for accelerated development of ankle prostheses," in *Robotics and Automation (ICRA), . . .* , Karlsruhe, Germany, 2013, pp. 2630–2635.

[17] W. Felt, J. C. Selinger, J. M. Donelan, and C. D. Remy, ""Body-In-The-Loop": Optimizing Device Parameters Using Measures of Instantaneous Energetic Cost," *Plos One*, vol. 10, no. 8, p. e0135342, 2015. [Online]. Available: http://dx.plos.org/10.1371/journal.pone.0135342

[18] J. C. Selinger, S. M. O'Connor, J. D. Wong, and J. M. Donelan, "Humans Can Continuously Optimize Energetic Cost during Walking," *Current Biology*, vol. 25, no. 18, pp. 2452–2456, sep 2015.

[19] K. E. Gordon and D. P. Ferris, "Learning to walk with a robotic ankle exoskeleton," *Journal of Biomechanics*, vol. 40, no. 12, pp. 2636–2644, 2007.

[20] K. A. Ingraham, D. P. Ferris, and C. D. Remy, "Using Wearable Physiological Sensors to Predict Energy Expenditure," in *2017 15th IEEE/RAS-EMBS International Conference on Rehabilitation Robotics (ICORR)*, London, UK, 2017, p. Submitted.

[21] D. McRuer, "Human dynamics in man-machine systems," *Automatica*, vol. 16, no. 3, pp. 237–253, 1980.

[22] S. J. Lee and J. Hidler, "Biomechanics of overground vs. treadmill walking in healthy individuals." *Journal of applied physiology (Bethesda, Md. : 1985)*, vol. 104, no. 3, pp. 747–755, 2008.

[23] Y. Wang and M. Srinivasan, "Stepping in the direction of the fall: the next foot placement can be predicted from current upper body state in steady-state walking," *Biology Letters*, vol. 10, no. 9, pp. 20 140 405–20 140 405, sep 2014.

[24] D. Engelhart, T. A. Boonstra, R. G. K. M. Aarts, A. C. Schouten, and H. van der Kooij, "Comparison of closed-loop system identification techniques to quantify multi-joint human balance control," *Annual Reviews in Control*, vol. 41, pp. 58–70, 2015.

[25] J. H. Pasma, D. Engelhart, A. B. Maier, R. G. K. M. Aarts, J. M. A. van Gerven, J. H. Arendzen, A. C. Schouten, C. G. M. Meskers, and H. van der Kooij, "Reliability of System Identification Techniques to Assess Standing Balance in Healthy Elderly," *Plos One*, vol. 11, no. 3, p. e0151012, 2016.

[26] J. H. Yoo, K. Y. Moon, D. Hwang, and M. S. Nixon, "Automated human recognition by gait using neural network," *2008 1st International Workshops on Image Processing Theory, Tools and Applications, IPTA 2008*, 2008.

[27] R. Begg and J. Kamruzzaman, "A machine learning approach for automated recognition of movement patterns using basic, kinetic and kinematic gait data," *Journal of Biomechanics*, vol. 38, no. 3, pp. 401–408, 2005.

[28] M. Hamaya, T. Matsubara, T. Noda, T. Teramae, and J. Morimoto, "Learning assistive strategies from a few user-robot interactions: Model-based reinforcement learning approach," *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2016-June, pp. 3346–3351, 2016.

[29] N. Hansen and S. Kern, "Evaluating the CMA Evolution Strategy on Multimodal Test Functions," 2004, pp. 282–291.

[30] K. L. Poggensee, K. Sreenath, and S. H. Collins, "Understanding the mechanisms behind human-in-the-loop optimization strategies," in *Dynamic Walking*, 2017. [Online]. Available: http://bit.ly/2BtRMuR

[31] D. S. Reisman, H. McLean, J. Keller, K. A. Danks, and A. J. Bastian, "Repeated split-belt treadmill training improves poststroke step length asymmetry," *Neurorehabilitation and neural repair*, vol. 27, no. 5, pp. 460–468, 2013.

[32] C. H. Soo and J. M. Donelan, "Coordination of push-off and collision determine the mechanical work of step-to-step transitions when isolated from human walking," *Gait and Posture*, vol. 35, no. 2, pp. 292–297, 2012.