

**Human Prediction and Robotic Lower-Limb Prosthesis Planning for Safe Perturbation
Recovery During Motion**


by

Shannon M. Danforth

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Mechanical Engineering)
in the University of Michigan
2022

Doctoral Committee:

Associate Professor Ram Vasudevan, Chair
Assistant Professor Talia Y. Moore
Assistant Professor K. Alex Shorter
Associate Professor Leia Stirling



Shannon M. Danforth

sdanfort@umich.edu

ORCID iD: [0000-0002-5825-587X](https://orcid.org/0000-0002-5825-587X)

© Shannon M. Danforth 2022

DEDICATION

For my family: Mom, Dad, Katy, and Mark.

ACKNOWLEDGMENTS

First, I would like to thank my advisor, Ram. The subject matter of my Ph.D. was quite different than that of my B.S. in civil engineering, and I am grateful you recognized and supported my interest in studying human motion and wearable robots. Throughout my time as a Ph.D. student, you have always been available for discussion whenever I need advice or assistance and have an abundance of ideas for tackling the difficult problems we encounter. The collaborative environment you foster in ROAHM Lab has given me lifelong friends. I am thankful for the time you spend with each student in this lab to ensure we communicate our work in the clearest possible manner. Finally, I greatly appreciate your support for my other academic interests outside of my technical research projects. Thank you for encouraging and supporting me while I took extra classes for my policy certificate program and conducted independent policy research. I hope Ch. 7 of this document highlights how important this work was to me.

Thank you to the rest of my committee members for their time and valuable feedback. Talia, you have been a fantastic mentor throughout my time in graduate school, from the Stability Basin project to our more recent coral snake research. Your expertise in human and animal biomechanics has been incredibly valuable, and your passion for using robots to discover more about locomotion is inspiring. I have enjoyed attending your weekly lab meetings and learning about all kinds of fascinating animals. Collaborating with you on making figures and videos has been a highlight of my time at Michigan. Over the past several years, I've seen the members of your lab grow into strong researchers and excellent communicators through your mentorship. I look forward to hearing of your lab's discoveries about animal and robot locomotion!

Alex, I had the great fortune of presenting an early version of my Stability Basin research to you for my qualifying exam many years ago, and your insightful questions shaped my future research directions. I am happy that you can see the completed Stability Basin work as a member of my committee and continue to discuss my other research projects in the same thoughtful manner.

Leia, I have admired your research and teaching since taking your Wearable Sensors class a few years ago. I appreciated that you took the time to discuss ethical questions about wearable sensors and methods for contextual inquiry that we can employ as engineers. As a member of my committee, you have raised practical questions and offered useful suggestions for my recent work in human trip-recovery. I look forward to shaping the final version of this work together.

Thank you to Shobita Parthasarathy, professor in the Ford School of Public Policy and director of the Science, Technology, and Public Policy certificate program. Through two classes and one independent study, you've helped me develop a compelling research project about the policies and social implications of robotic prostheses. With the tools from science and technology studies, I feel equipped to generate technologies that make a positive societal impact, and I recommend the certificate program for any engineer.

To my labmates in ROAHM lab, thank you for making the lab an entertaining and supportive place to conduct research. I have made some of my best friends in this lab, and I can't wait to see what each of you accomplish. Patrick, you have been an incredible research partner and friend for the past six years. It is not easy developing computational tools for human motion—human datasets are never quite as consistent as we hope—but discovering the unique characteristics of individual movement is a challenge I was happy to undertake with you. Thank you for enthusiastically participating in three music videos for Dynamic Walking submissions; I think we perfected the art by the third one. I will miss sitting next to you in lab when we move onto our respective new jobs, but don't be surprised if you receive emails from me every so often asking for a sanity check on modeling a dynamic system. Xinyi, it has been such a pleasure mentoring you for the past year. You have met the challenge of learning about human biomechanics and wearable robotics with grace and skill. I am sad I will not have more time to work with you on our projects, but I can't wait to keep in touch as your future research projects take shape!

To Mom, Dad, Katy, and Mark: thank you for your constant love, support, and humor that has kept me motivated for the past six years. It's been a long time living across the country from you all, and I'm happy I'll be moving a bit closer soon. Your visits to Ann Arbor have been some of my best memories. Thank you to my extended family and Martin's family, who are always a joy to see during holidays and weddings.

To the friends I've made in Michigan during my time here, thank you for your laughter and company. I have enjoyed exploring this beautiful state with you and look forward to keeping in touch in the future. To Daphna and Hannah, my fellow engineering Ph.D. student friends, I am thankful to have experienced graduate school with you two by my side. I can't wait to hear about your fascinating research at your own dissertation defenses.

Thanks to Martin for your constant love and support through all six years of my time at the University of Michigan. In the past two years, I have also enjoyed having you as a co-worker in ROAHM Lab; thank you for all those early weekend mornings of helping with data collection. Looking forward to our next chapter!

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGMENTS	iii
LIST OF FIGURES	viii
LIST OF TABLES	xv
ABSTRACT	xvii
CHAPTER	
1 Introduction	1
1.1 Research Overview	3
1.2 Notation	4
1.3 Concepts	5
2 Human Subject Experiments	7
2.1 Cable Pulls and Foot Shifts during Sit-to-Stand	7
2.2 Trips during Level-Ground Treadmill Walking	10
3 Specific Aim 1: Quantifying Stability During Sit-To-Stand	16
3.1 Introduction	16
3.2 Methods	19
3.2.1 Classifying trials as successful/unsuccessful	19
3.2.2 Dynamic modeling	20
3.2.3 Controller Model	22
3.2.4 Computing Stability Basins	27
3.2.5 Evaluating Stability Basin accuracy	29
3.3 Results	33
3.4 Discussion	38
4 Specific Aim 2: Investigating Optimality in Perturbation Response During Sit-To-Stand	42
4.1 Introduction	42
4.2 Polynomial Inverse Optimal Control	43
4.3 Methods	46

4.3.1	Inverted Pendulum Model	46
4.3.2	Training the Lagrangians	47
4.3.3	Testing the Lagrangians	48
4.3.4	Comparison to LQR Controller	49
4.4	Results	50
4.5	Discussion	51
5	Specific Aim 3: Predicting Human Trip Response During Walking	56
5.1	Introduction	56
5.2	Trip-Recovery Experiment and Dataset	58
5.3	Gaussian Process Regression Model	60
5.3.1	Multi-Output GPR Distribution	60
5.3.2	Conditional Prediction	62
5.4	NARX Model	63
5.5	Pendulum Model with Moving Base	65
5.6	Training and Testing the Models	67
5.7	Results	68
5.8	Discussion and Conclusion	71
6	Specific Aim 4: Planning Trip-Recovery Trajectories for Robotic Lower-Limb Prostheses	73
6.1	Introduction	73
6.2	Trip Experiment and Recovery Strategies	76
6.3	Preliminaries	77
6.3.1	Planning Model	77
6.3.2	Phase Variable	80
6.3.3	Planning Strategy	81
6.4	Offline Computations	82
6.4.1	Parameterized Trip-Recovery Libraries	83
6.4.2	Training Predictive Models of Swing Hip Kinematics	86
6.4.3	Target Sets	87
6.5	Online Computations	88
6.5.1	Conditional Predictions and Reachable Sets of Human Hip States	90
6.5.2	Parameterized Reachable Sets of the Prosthesis Heel and Toe	92
6.5.3	Cost and Constraints for Optimization	95
6.5.4	Real-Time Optimization	99
6.6	Evaluation	100
6.6.1	Evaluating the Proposed Planning Framework	100
6.6.2	Comparison to Impedance-Based Control	101
6.7	Results	104
6.8	Discussion	104
7	Specific Aim 5: Towards a Technology Assessment for Robotic-Limb Prostheses	109
7.1	Introduction	109
7.2	Background	110

7.2.1	Some key terms and concepts	110
7.2.2	Three types of lower-limb prostheses	111
7.3	Why a technology assessment?	112
7.4	Current technology assessment methods	113
7.4.1	Delphi	114
7.4.2	Participatory TAs	115
7.4.3	Focus groups: “Publics-in-Particular”	116
7.4.4	Scenario planning	116
7.4.5	Analogical case studies	117
7.4.6	Discursive TA	118
7.4.7	Proposed assessment for robotic lower-limb prostheses	118
7.5	Preliminary Results	119
7.5.1	Existing concerns about robotic lower-limb prostheses	119
7.5.2	Analysis of discursive themes and standardized assessments	122
7.6	Discussion	125
8	Conclusion and Future Directions	127
8.1	Discussion of Contributions	127
8.2	Future Directions	129
8.2.1	Stability Basins	129
8.2.2	Inverse Optimal Control	130
8.2.3	Trip-Recovery Prediction	131
8.2.4	Trip-Recovery Planning in Robotic Lower-Limb Prostheses	132
8.2.5	Technology Assessment for Robotic Lower-Limb Prostheses	133
8.3	Concluding Remarks	134
APPENDIX		135
BIBLIOGRAPHY		143

LIST OF FIGURES

FIGURE

2.1	Example of each type of Sit-To-Stand trial.	9
2.2	Trip experiment setup, including tethers attached to each foot, brakes applied at short intervals, and load cells to measure the tether force.	11
2.3	This figure shows the (a) tripped heel, (b) contralateral swing heel, (c) tripped hip, (d) tripped knee, and (e) tripped ankle trajectories after a trip for one subject in our dataset. We observe changes in the kinematic trajectories immediately after trip onset, indicated by the vertical dashed line, and a return to baseline levels after 2-3 steps. . .	13
2.4	We analyzed the swing heel anterior-posterior position, swing heel height, and swing leg configuration throughout the tripped step to sort the trip recoveries by hand. The horizontal axis in all plots is a variable that represents progression through the swing phase, described in Ch. 6. We label (a) as an elevating strategy because the swing heel elevates (above average nominal behavior) after the trip occurs, with the heel striking in front of the body; (b) as a delayed lowering strategy because the swing heel initially elevates, then lowers behind the body with the toe striking the ground first; (c) as a lowering strategy because the swing heel immediately lowers behind the body after a trip with the toe striking the ground first; and (d) as an “incomplete arrest” because there is little difference from average nominal behavior.	14
2.5	The three trip-recovery strategies with corresponding swing heel trajectories from one subject (Subject ID 1) in our dataset. The horizontal axis is a variable that represents progression through the swing phase, described in Ch. 6. (a) A depiction of the swing leg during each strategy, adapted from Shirota et al. 2014. (b) A sample of anterior-posterior (AP) swing heel positions relative to the stance foot during each type of recovery compared to nominal (level-ground walking) data. Note that the swing heel reaches close to the same final position as nominal during the elevating strategy, whereas the swing heel travels backwards in the anterior-posterior position during the delayed lowering and lowering strategies. (c) A sample of swing heel heights during each type of recovery compared to nominal data. Note that the swing heel reaches ground level during the elevating trials, but does not reach the ground level during delayed lowering and lowering trials (aside from one instance). During the latter two recovery strategies, the swing toe typically strikes the ground first, so the heel stays above the ground.	15

3.1	An illustrative overview of Stability Basins for Sit-to-Stand. The Stability Basin represents the set of model states through time that will successfully arrive at a target set for a given individual and Sit-to-Stand strategy. Trajectories of the model are illustrated, where the times that perturbations are applied are denoted by the dashed lines. Trajectories that exit the Stability Basin are predicted to lead to stepping or sitting.	18
3.2	Subjects began from a seated position with their arms crossed against their chests. The subject's COM is illustrated by filled circles. Cable-pulls applied to the subject sometimes caused them to step or sit; otherwise, the trial was considered to be successful.	21
3.3	These figures show how we construct the controller models from data. The process is shown for u_x , and is identical for u_y . In Fig. 3.3(a), u_x evolves as a function of time t and state x , which has been depicted as a single dimension for this illustration. The grey lines represent the input trajectories of nominal and foot-shift trials, while the blue lines represent the input trajectories of perturbed successful trials. Note that we have not plotted the portions of the successful perturbed trajectories when the perturbations were active, because these segments are not used to train the controller models. An arbitrary time $t = 0.25$ is illustrated by the black outline in Fig. 3.3(a). The observed inputs from nominal, foot-shift, and perturbed successful trials at that time are plotted as points in Fig. 3.3(b - d). Figure 3.3(b) shows that the LQR controller utilizes linear feedback about a nominal trajectory. Because the LQR controller is generated by minimizing a cost function specified by the matrices Q and R , its feedback matrix does not necessarily coincide with a line of best fit. Figure 3.3(c) illustrates that the FF+FB controller is given by a line of best fit. Figure 3.3(d) shows how the upper and lower bounds $u_x^{lb}(t, x(t))$ and $u_x^{ub}(t, x(t))$ are parallel linear functions of state that encompass all of the successful inputs.	24
3.4	This figure illustrates how X_T (the target set) is formed from data for a given subject and strategy. The states (r_x and v_x) at 100% STS of each observed successful trial within a strategy are illustrated as blue dots. An example X_T represented by the shaded blue zonotope contains all of the states observed to lead to standing. The zonotope is parameterized by its center c and generator vectors $g^{(1)}$ and $g^{(2)}$. This example depicts a 2-dimensional X_T , but in reality they are 4-dimensional and encompass both horizontal and vertical components of the state space.	28
3.5	This figure illustrates how to check if a point is inside or outside of a zonotope. An example zonotope (shaded grey) is parameterized by its center c and generator vectors $g^{(1)}$ and $g^{(2)}$. A test point p is contained within the zonotope because the maximum absolute value of the coefficients on $g^{(1)}$ and $g^{(2)}$ is less than or equal to 1. Another test point q lies outside of the zonotope because the maximum absolute value of the coefficients on $g^{(1)}$ and $g^{(2)}$ is greater than 1.	30

3.6	An outline of the Stability Basin validation with examples provided by the computed Stability Basin for subject ID 1’s natural strategy under the Input Bounds controller. The horizontal projection of the Stability Basin is represented as the region encapsulated by the light grey borders. The projection of the target set X_T is shown as the dark grey region on the right side of each plot. The state trajectories of all of subject ID 1’s successful natural strategy trials <i>except one</i> are used to construct the Stability Basin on which the left-out trial is tested, as shown in Fig. 3.6(a). Then all of subject ID 1’s successful natural strategy trials are used to construct the Stability Basin on which the unsuccessful trials are tested, as shown in Fig. 3.6(b). State trajectories of a step and a sit exit the basin before the onset of failure. As detailed in Sec. 3.2.1, we define step initiation as the time when the toes of either foot move more than 0.0762 m (3 inches) in the anterior-posterior direction, and sit initiation as the time at which both v_x and v_y become negative.	32
3.7	This figure shows example nominal unnormalized horizontal acceleration trajectories for each Sit-to-Stand strategy for subject ID 8.	33
3.8	Maximum (Fig. 3.8(a)) and minimum (Fig. 3.8(b)) horizontal accelerations are significantly different between Sit-to-Stand strategies, with $p < 0.001$ when comparing the maximums or minimums of any two strategies under a two-sample t-test. Only successful trials are used in this analysis. For each strategy, each small circle represents data from a single Sit-to-Stand trial. The central mark within the box plot indicates the median horizontal acceleration. The bottom and top edges of the box indicate the the 25th and 75th percentiles, respectively. The whiskers extend to the most extreme horizontal accelerations not considered outliers, and the outliers are plotted individually as small circles with an ‘x’ through them.	34
3.9	The Stability Basins formed for Subject ID 1’s natural Sit-to-Stand strategy using the LQR controller model (Fig. 3.9(a)), FF+FB controller model (Fig. 3.9(b)), Input Bounds controller model (Fig. 3.9(c)), and the naive method (Fig. 3.9(d)) are shown. The horizontal and vertical projections of the Stability Basins are represented as the regions encapsulated by the light grey borders, where black outlines are used to emphasize the changes in shape of cross-sections over time. The projection of the target set X_T is shown as the dark grey region on the right side of each plot.	37
3.10	This figure compares the predictive accuracy of the SBs formed using the Input Bounds controller to three other methods. Results are presented for each Sit-to-Stand strategy, aggregated across subjects. The Input Bounds controller’s predictions for both successes and failures are near the ground truth, while the three other methods predict many more failures than were experimentally observed.	38
4.1	The inverted pendulum model used to model STS for inverse optimal control. The input u is a torque applied to the pendulum base.	47
4.2	Weights on all monomials of the best fitting (a) unperturbed degree 2, (b) unperturbed degree 4, (c) perturbed degree 2, and (d) perturbed degree 4 lagrangian for each subject’s data, normalized to have unit norm.	51
4.3	Weights on all monomials of the best fitting lagrangian for unperturbed (light green) and perturbed (light orange), averaged across all subjects. The coefficient weights have been normalized and are shown here for (a) degree 2 and (b) degree 4 lagrangians.	52

4.4	The unperturbed lagrangians, perturbed lagrangians, and LQR controller were tested on perturbed data. This figure shows an example simulation result for Subject ID 6 during a forwards cable-pull trial. The ground truth scaled angular positions (green), angular velocities (purple), and inputs (blue) are pictured in the solid line, with the bold segment indicating the time when the cable pull perturbation was active. (a) and (b) show the simulation results for lagrangians of degree 2, 4, and 6 for comparison. (c) shows the simulation result for the LQR controller.	53
5.1	In this chapter, we develop and compare predictions of sagittal-plane swing hip kinematics after a trip occurs. The three hip states of interest (height, anterior-posterior position, and flexion-extension angle) are shown in blue to the left. On the right, the stick figures show an example of one subject’s trip-recovery response from our experiment. Note that the knee location is determined by the hip flexion-extension angle and thus is plotted in addition to the hip positions. The black line shows the ground truth hip and knee locations, and the pink line shows the predicted response from our Gaussian process regression model.	57
5.2	Hip flexion-extension angle, height, and anterior-posterior (AP) position relative to the stance foot, plotted for the elevating, delayed lowering, and lowering strategies in orange, green, and purple, respectively. The thick black lines show the average nominal hip states. These data are from one subject in our dataset.	59
5.3	Diagram of the (a) open-loop configuration and (b) closed-loop configuration of the NARX model, where TDL stands for time delay layer.	64
5.4	Our model for predicting hip kinematics has a floating base located at the hip joint and a pendulum representing the swing leg. The model has horizontal and vertical forces and a torque applied at the hip.	66
5.5	Visualization of summary statistics for the (a) subject-specific and (b) generalized models. The central black line indicates the median NRMSE for each model; the bottom and top edges of each box are the 25th and 75th percentiles, respectively; the dashed whiskers extend to the limits of the data not considered outliers; and the outliers (determined using the interquartile range) are plotted individually as plus signs. The y-axis is a logarithmic scale due to the extreme magnitude of outliers.	68
5.6	An example of simulation results for one delayed lowering trial using the subject-specific (a) GPR, (b) NARX, and (c) pendulum models. The GPR model prediction in (a) includes a 95% confidence interval. The entire ground truth trial is shown in the black line, with the trip onset indicated by the arrow. The NRMSE for this trial, averaged across each response variable, is provided for each prediction method. The black dots represent the information provided to each model: 15 conditional points in (a); 15 feedback delay points in (b), and one initial condition for simulation in (c).	70

6.1	The framework developed in this chapter plans trajectories for a knee prosthesis during and after a stumble that ensure successful recoveries for a set of predicted human behavior. The kinematic model used in this work contains human and prosthesis subsystems, constructed from experimental data. This figure shows reachable sets for the model, which depend on both human and prosthesis behavior, after a toe catch perturbation. The gray shapes at the hip and knee represent the range of joint locations from the subject’s predicted swing hip response. The light orange and light purple shapes at the foot represent the set of prosthesis heel and toe locations for a range of parameterized prosthesis behavior. Our framework selects one heel and toe parameter in an optimization program, which produce the darker orange and purple subsets. The optimization program aims to (1) avoid toe scuffs and (2) ensure the heel strike location lands as close as possible to a <i>target set</i> representing desired foot contact location, shown by the green lines (in the callout), for each time step where ground contact is possible.	74
6.2	The planning model considered in this chapter consists of a transfemoral amputee (black links) and robotic prosthesis (blue links). Part (a) depicts the kinematic model used in our proposed framework. To compare to an impedance-based controller, we use a dynamic model, which includes the link masses and inertias and joint inputs shown in (b).	78
6.3	This figure shows example observed knee and ankle trajectories and parameterized knee and ankle trajectories for Subject ID 1. In (a) and (c), the observed ankle velocity trajectories are plotted for each trip recovery strategy. A k -value in $[-1, 1]$ is assigned to each trajectory based on its amplitude at 50% of swing phase. In (b) and (d), the parameterization of ankle velocity trajectories $q_3(t; k)$ is shown. The mean trajectories $q_{2,\text{mean}}^s(t)$ and $q_{3,\text{mean}}^s(t)$ are shown in black, and the gradient from dark to light is created by varying the value of k_2 or k_3 from -1 to 1 . We note that in each of these plots, the phase $t \in [0, 1]$ has been scaled the same for each strategy, highlighting differences in phase duration between each strategy.	85
6.4	This figure shows the target sets X_T computed from Subject ID 1’s data. The orange, green, and purple points represent the hip positions and velocities and foot strike locations at the end of each trial for elevating, delayed lowering, and lowering, respectively. The volumes encompassing each set of points are the zonotope target sets. Note that the target sets are defined over heel positions for the elevating strategy and over toe positions for the delayed lowering and lowering strategies. The target sets found for each strategy occupy distinct regions. As expected for the elevating strategy, X_T is characterized by more forward progression of the heel and hip positions, and a higher hip velocity. For the delayed lowering strategy, X_T occupies a region of low forward progression of the toe and hip positions, and negative hip velocities. The lowering strategy’s target set lies between these two.	89

6.5	This figure shows computed reachable sets for the human-prosthesis subsystem. The elevating trial is from Subject ID 8, the delayed lowering from Subject ID 1, and the lowering from Subject ID 10. The stick figures show the subjects' ground truth trip-recovery behavior. The reachable sets are plotted for every seven phase steps, which cause the gaps at the beginning when the leg swings forward quickly. Part (a) shows the reachable sets computed before the cost and constraints are formed for optimization. The human subsystem's reachable sets are plotted in gray. Note that because the human's hip flexion-extension angle influences the position of the knee, the knee position reachable sets are also plotted in gray. The heel and toe reachable sets $V_{\text{heel}}(t)$ and $V_{\text{toe}}(t)$, computed in Sec. 6.5.2, are shown in light orange and light purple, respectively. They describe all possible locations of the heel and toe throughout swing phase given the predicted sets of human motion and the range of parameterized prosthesis behavior. Part (b) shows the optimization results for each trial (Secs. 6.5.3 and 6.5.4). The dark orange and purple sets represent the possible heel and toe locations for a choice of $k \in K$ and the entire set of predicted human trip-recovery behavior. The optimization picks $k \in K$ to avoid undesired foot scuffs (toe scuffs for elevating trials, heel scuffs for delayed lowering and lowering) and to minimize the distance outside the target sets (shown as the green intervals) any time ground contact is possible (heel for elevating, toe for delayed lowering and lowering).	96
6.6	This figure shows example knee and ankle torques calculated from an average nominal trajectory T_{ref} via inverse dynamics (black solid curve), and the desired knee and ankle torques produced by the nominal impedance-based controller T_{des} (black dashed curve) for Subject ID 1. The number of modes is set to $n_m = 4$ in this case. Note that for the knee and ankle torque profiles, we use different mode divisions and parameters, as shown in (a) where the divisions are set at 9%, 25% and 64%, and in (b) where the divisions are set at 16%, 53% and 71%.	103
6.7	This figure shows results for three sample trials. Row (a) shows an elevating trial from Subject ID 8. The first column depicts the subject's observed trip-recovery behavior. The second column shows the optimization result from our proposed framework, evaluated using the subject's ground truth hip motion as described in Sec. 6.6. At the phase point where the heel contacts the ground, we determine the desired range of heel strike locations using the hip's AP position and velocity. The desired range is pictured as the green interval below the ground. Note that our proposed framework successfully places the heel in this range. In the third column, we see that the impedance-based controller also places the heel in the desired heel strike range, but that the toe scuffs the ground along the way. Rows (b) and (c) depict delayed lowering and lowering trials for Subject IDs 1 and 10, respectively. Note that for these two strategies, we aim to avoid heel scuffs and to place the toe in the desired range at the end of the step. In the middle column, our proposed framework shows the toe contacting the ground within the green range for both delayed lowering and lowering. In both (b) and (c), the impedance-based controller scuffs the heel on the ground, and the toe AP position lies outside the desired range.	106
7.1	The gait cycle follows one leg through the stance and swing phases. Image adapted from T. Stöckel et al. 2015.	110

7.2 An example of each prosthesis type. Image adapted from Ossur Total Knee, Ossur
Rheo Knee, and Ossur Power Knee images. 111

LIST OF TABLES

TABLE

3.1	Statistics collected from the perturbative sit-to-stand experiment are aggregated across subjects and summarized below. Note that the onset of failure (Step %STS and Sit %STS) often occurred after the defined trial end ($t = 1$), which is discussed in Sec. 3.2.5.1. Supplementary material that reports these statistics for individual subjects is available online.	34
3.2	The accuracy of the Stability Basins generated using the Input Bounds controller model are reported below. This table reports a tally of the correct predictions for each Sit-to-Stand strategy and trial type across subjects.	35
3.3	The accuracy of Stability Basins generated using the Input Bounds controller model are reported for cable pull trials at three different force levels of the applied perturbation. The results are aggregated across subjects and the three tested Sit-to-Stand control strategies.	35
3.4	The accuracy of the Stability Basins formed using the Input Bounds controller model were compared to Stability Basins formed using other methods. The results of the comparison are reported below, and are aggregated across subjects and sit-to-stand strategies. Supplementary material that reports these results for individual subjects and sit-to-stand strategies is available online.	36
3.5	The accuracy of the Stability Basins for each method are compared when dilating the LQR, FF+FB, and naive Stability Basins. The results are aggregated across subjects and STS strategies.	38
4.1	Maximum Euclidean distance between any two subjects' set of coefficients with unit norm, reported for unperturbed and perturbed lagrangians of degrees 2, 4, and 6. The maximum possible d between any two unit norm vectors is 2.	50
4.2	Normalized root mean square error E for the training and testing cases described in Sec. 4.3. The mean (standard deviation) is computed across all subjects and trials. For each case, the row corresponding to the degree with the minimum error is highlighted in gray.	52
5.1	Average NRMSE for all three prediction methods, with lowest subject-specific and generalized errors highlighted for each strategy.	69
5.2	Approximate f statistic between models, with * indicating statistical significance, and corresponding p -values.	69
5.3	Percent of ground truth trials that remain inside the 95% confidence bounds for the GPR model.	70

6.1	Results for our proposed framework compared to nominal impedance-based control. . .	105
A1	Operations with polynomial zonotopes may either be computed exactly or in an over-approximative fashion.	142

ABSTRACT

Falls are prevalent among older adults and people with lower-limb amputation of all ages. Because falls can lead to injury and loss of quality of life, there are several proposed ways to reduce fall risk. One recommendation is to identify those who are most likely to fall and provide targeted physical therapy, but existing methods for identifying fall risk have been unable to reliably predict who will fall. Another way to reduce falls is to assist the individual using a wearable robotic device, such as an exoskeleton or prosthesis, when stumbles occur. However, because there is uncertainty in predictions of how the human will respond, wearable robots are unable to *safely* assist with trip recovery.

To accurately identify who may become unstable during motion, this dissertation presents Stability Basins, which characterize individual dynamic stability during the Sit-to-Stand motion. Stability Basins, formed using a dynamic model with a model for the individual's control strategy, encompass all stable model states at each point during the motion. In this document, Stability Basins are validated using data from an 11-subject experiment where subjects were pulled by motor-driven cables as they stood up from a chair. The Stability Basins' accurate characterization of stability during motion shows promise for identifying fall risk.

Another way to predict human response to perturbation (e.g., a push or trip) is to find an individual's underlying objective during motion and use it to inform predictions. Given the assumption that humans are optimizing some cost such as metabolic energy during motion, *inverse optimal control* finds the underlying objective function corresponding to observed data. This dissertation presents results from applying an inverse optimal control formulation to the 11-subject Sit-to-Stand dataset. Results suggest that subjects place priority on the position and velocity of their center of mass rather than input torques during both perturbed and unperturbed Sit-to-Stand, and that the underlying cost function can be used to effectively simulate perturbation response.

Accurate and quick predictions of trip recovery during walking are necessary for planning safe trip-recovery motion in wearable robots. A 16-subject experiment was conducted where subjects were tripped via tethers attached to their feet while walking on a treadmill. In this document, three different models (Gaussian process regression, neural network, and dynamic model) were trained on the trip-recovery data and evaluated for prediction accuracy and computation time. The Gaussian process regression models outperformed the other two in both criteria, highlighting their

potential for use in wearable robot trip-recovery planning.

Despite the success of the Gaussian process regression models, predictions of human motion will never be perfectly accurate. To safely plan motion for trip recovery, wearable robots must account for prediction uncertainty. This dissertation presents a framework for planning trip-recovery motion in a robotic knee-ankle prosthesis after a trip occurs that accounts for a *set* of predicted human behavior. The framework solves an optimization problem to ensure that the foot is placed close to a desired location at the end of the step while avoiding unsafe foot scuffs before touch-down. The approach is demonstrated in simulation using data from the 16-subject trip experiment.

In the final chapter, this dissertation argues that robotics researchers should conduct *technology assessments* to address factors that may influence the real-world adoption of a technology, and presents preliminary results from a technology assessment of robotic lower-limb prostheses. Together, the frameworks, tools, and technology assessment included in this dissertation advance the practical and safe use of autonomy to reduce falls and fall-related injuries.

CHAPTER 1

Introduction

An estimated 1 in 3 adults over age 65 fall each year, with 20-30% of these falls resulting in moderate to severe injuries [1]. Approximately one third of people who experience falls develop a fear of falling and begin to avoid activities [2], which leads to functional decline, social isolation, and decreased quality of life [3]. Additionally, falls are costly for the healthcare system; in 2015, medical bills for falls exceeded \$50 billion in the United States, with Medicare and Medicaid covering 75% [1]. Common risk factors for older adults include lower body weakness; difficulty walking and balancing (sometimes due to medication); vision problems; and hazards such as clutter or uneven steps [4].

People with lower-limb loss¹ constitute another group with increased risk of falling. There are an estimated 830,000 to 1 million people with lower-limb loss in the United States [5]. The majority use passive lower-limb prostheses, and exhibit instances of falling comparable to older adults over 85 [6, 7]. Younger individuals with lower-limb amputation report falling at higher rates than other age groups, likely because they engage in riskier activities. These falls do not always lead to injury, but when older amputees' fall and injury data are included, the rates of fall-related injury for amputees is higher than that of older adults [8]. Across all ages, increased risk of falling can also lead to decreased confidence during everyday household tasks or community ambulation [9]. The causes of falls in community-dwelling amputees are understudied, but researchers postulate that the causes are similar to older adults [9].

In both older adult and amputee groups, fall prevention measures have not been overwhelmingly successful. One challenge is identifying who is most at risk of falling and in what situations they will become unstable. Clinicians use outcome measures such as Timed Up and Go or the Berg Balance Scale to assess fall risk in both groups [10, 11, 9]. These assessments struggle to identify individuals most at risk of falling [12, 13, 14]. The Timed Up and Go test, for example, reduces fall risk to one number. While the simplicity of this metric has advantages for clinicians who must

¹In this document, I use both person-first language (“person with limb loss/amputation”) and identity-first language (“amputee”) as suggested by the American Psychological Association; see <https://www.apa.org/pi/disability/resources/choosing-words>.

see multiple patients per day, it may also neglect valuable information about *how* the individual accomplishes everyday tasks (i.e., their motor control strategy).

After identifying at-risk individuals, fall prevention training might not be enough to significantly reduce the incidence of falls. Fall prevention measures for older adults span a broad range of interventions, including strength and balance training; home interventions such as anti-slip shoes; reviewing and adjusting medications; and surgery to improve vision [15]. Of these, exercise training and home interventions have the best success in reducing falls. Two studies have implemented fall-prevention exercise programs for lower-limb amputees with promising results [16, 17]. However, clinicians may not have the time to facilitate the consistent exercise training needed to effectively prevent falls. One study suggests that the maximum fall reduction rate to be achieved by fall prevention measures is 30-40% [18].

Clemens et al. urge readers to explore other efforts to mitigate fall-related injury besides the fall prevention measures above [9]. One such example is fall arrest training, which teaches individuals to land falls in safer ways. Another suggestion is using wearable robots (exoskeletons and prostheses) that can assist with stumble recovery [19]. Despite their potential to provide assistance via an external power source, wearable robots also add safety and comfort challenges due to the powerful torques they can produce. These devices must include features to avoid disrupting motion or leading to injuries if torques are applied at the wrong time [20]. Including such safety features is difficult because the robotic device must first predict how the human will respond during a stumble, then choose appropriate behavior in real-time. Because predictions will never be perfect, a planner that accounts for uncertainty would further increase safety.

This dissertation presents methods for characterizing stability during motion, predicting how humans respond to perturbations during motion, and planning wearable robotic behavior to assist with disturbances. Although external perturbations are not the sole cause of falls in older adults and amputees, I focus on unexpected disturbances such as pushes or trips during motion. My work analyzes data from perturbative Sit-to-Stand and level-ground walking experiments, as difficulty maintaining balance during these tasks has been correlated with falls in older adults, and both are crucial for maintaining independence [21].

The technological methods presented in this dissertation aim to prevent falls and assist with fall recovery. As with all new technologies, it is important to consider the social and policy implications of the frameworks proposed in this document. Although we often view standardized technologies as the objectively “best” design, they are in fact a product of social, political, and economic forces [22]. Scholars in Science and Technology Studies have introduced *technology assessments* for anticipating and addressing the societal impacts of a new technology. Conducting such assessments early in a technology’s development can prompt actors to think critically about design choices, pilot studies, regulations, and marketing. Despite technology assessments’ promise

to improve technology adoption rates and facilitate trust among lay people, they are typically conducted by an outside party and largely ignored by the developers themselves. In the last chapter of this document, I provide an overview of technology assessment methods in the context of robotic lower-limb prostheses and present preliminary findings from my own assessment.

1.1 Research Overview

This document proposes solutions to the challenges of preventing falls, demonstrated on the Sit-to-Stand and walking tasks, through five specific aims.

Specific Aim 1: Quantifying Stability During Sit-To-Stand. None of the established outcome measures for fall prediction have been shown to reliably predict falling, likely because a person's stability depends on their chosen control strategy and varies throughout a given dynamic task. In Ch. 3, I present a quantitative metric for time-varying stability during the Sit-to-Stand motion. This framework, called the Stability Basin, uses individualized dynamic models to determine the basin of attraction for successful Sit-To-Stand motion. Using a dataset where people were pulled forwards and backwards by cable pulls in the sagittal plane while standing up from a chair, I confirm that the Stability Basins can accurately predict when someone will become unstable. This work is published in [23].

Specific Aim 2: Investigating Optimality in Perturbation Response During Sit-To-Stand. Many researchers have hypothesized that humans and animals behave in a way that is optimal with respect to some criteria, such as energy cost or balance [24]. Knowing the cost function minimized by a human during motion can be useful for predicting responses to unexpected disturbances. Inverse optimal control uses data and a dynamic model to solve for an underlying cost function that is minimized. In Ch. 4, I implement an inverse optimal control formulation with Sit-to-Stand data to solve for individualized cost functions that subjects minimized during motion. I then use the cost functions to generate Sit-to-Stand trajectories and compare them to ground truth data. I found that all 11 subjects prioritized the same model states during perturbed and unperturbed Sit-to-Stand, although with variation in relative weighting between these states. The inverse optimal control formulation implemented here generated slightly more accurate predictions of perturbation response than an LQR controller that generated feedback about the nominal trajectory. This work was presented at *Dynamic Walking 2019*.

Specific Aim 3: Predicting Trip Response During Walking. Humans typically exhibit three types of recovery strategies when tripped during walking: Elevating, where the foot clears the obstacle and the heel strikes further ahead; delayed lowering, where the foot initially elevates, then lands backwards with the toe striking first; and lowering, where the foot immediately lowers behind the body with the toe striking first [25]. Within each of these three strategies, large variation

exists. Predicting the human response to these strategies can help assistive robotic devices choose recovery behavior that will work with the human’s chosen recovery strategy. In Ch. 5, I focus on predicting swing hip behavior: flexion-extension angle, anterior-posterior position, and height. I conducted a 16-subject experiment that introduced trips via tethers attached to subjects’ feet while they walked on a treadmill. I then fit three types of models to the perturbed swing hip data in each strategy and evaluate the models for prediction accuracy and computation time. This work is published in [26].

Specific Aim 4: Planning Trip-Recovery Trajectories for Robotic Lower-Limb Prostheses.

Because trips can occur at any point during gait, there is no predefined trip-recovery reference behavior for a robotic lower-limb prosthesis. Online optimization shows promise for planning prosthesis behavior after a trip in real-time, but existing methods are unable to account for the uncertainty in the human reaction [27]. In Ch. 6, I present a trip-recovery planning framework for wearable robotic lower-limb devices, specifically focusing on robotic knee-ankle prostheses. The framework conducts online optimization of knee and ankle trajectories during and after trips that account for predicted sets of human trip-recovery behavior, utilizing the model predictions from Specific Aim 3. As mentioned before, the prosthesis must work with the human’s chosen recovery strategy to avoid applying dangerous torques while the human is trying to recover from a trip. I first incorporate the prediction uncertainty to obtain sets of predicted swing hip behavior, then select knee and ankle reference trajectories for the robot that are safe for all swing hip behaviors in the predicted set.

Specific Aim 5: Towards a Technology Assessment for Robotic Lower-Limb Prostheses.

Robotic lower-limb prostheses are a relatively new technology, with most development restricted to the research setting. As these devices begin to enter commercial settings and prompt regulations and standards to be set for their performance, it is important for robotic prosthesis developers to consider what has informed their designs, how regulations and standards are set, and who is involved in the design and regulatory process. In Ch. 7, I provide a background about robotic lower-limb prostheses for a non-technical audience, then summarize several existing technology assessment methods and propose how they could be applied to robotic prosthesis technology. Finally, I present preliminary results for my chosen technology assessment and reflect on how my findings can influence future research in the field.

1.2 Notation

The following list describes notation that will be used throughout this document:

- Lowercase letters denote points, variables, and functions.

- Vectors are denoted as $[x_1, \dots, x_n]^\top$ or $(x_1, \dots, x_n)^\top$.
- Uppercase letters denote sets and matrices.
- \mathbb{R}^n denotes the n -dimensional set of real numbers.
- The unit circle is \mathbb{S}^1 .
- The set of 3×3 rotation matrices is $\text{SO}(3)$.
- Let $A \times B$ denote the Cartesian product of sets A and B.
- A state x has first and second time derivatives \dot{x} and \ddot{x} , respectively.

1.3 Concepts

This section details concepts that will appear throughout the document. Specific Aims 1 and 4 utilize *reachability analysis* for prediction and planning of human motion. Just as we can integrate a system's dynamics from an initial condition to simulate a trajectory, reachability analysis approximates the set of states that can be attained by the system from an initial set. The *forwards reachable set*, used in Specific Aim 4, is the set of all states that can be attained by the system starting from the initial set. The *backwards reachable set*, used in Specific Aim 1, is the set of all states that can arrive at some final set.

Because many set representations (e.g., level sets, polytopes, ellipsoids) are inefficient, reachability analysis is often time-consuming and restricted to lower-dimensional systems. However, *zonotopes* are a mathematical object that can be used to efficiently over-approximate reachable sets. First, let $U, V \subset \mathbb{R}^n$. The *Minkowski sum* is defined as:

$$U \oplus V = \{u + v \mid u \in U, v \in V\}. \quad (1.1)$$

A zonotope Z is a convex polytope that is closed under linear maps and Minkowski sums [28]. It describes the set of points that can be reached by a center, $c \in \mathbb{R}^n$, and linear combination of generators, $g^{(1)}, \dots, g^{(p)} \in \mathbb{R}^n$. The coefficient $\beta^{(i)}$ on each generator lies in $[-1, 1]$:

$$Z = \left\{ y \in \mathbb{R}^n \mid y = c + \sum_{i=1}^p \beta^{(i)} g^{(i)}, -1 \leq \beta^{(i)} \leq 1 \right\}. \quad (1.2)$$

In this document, we may also use the following shorthand notation:

$$Z = (c, \langle g_1, \dots, g_p \rangle). \quad (1.3)$$

More information about zonotopes is provided in Appx. A.

To compute reachable sets with zonotopes in Ch. 3, we use an open-source reachability toolbox called CORA [29, 30]. CORA represents reachable sets by using one zonotope at each finite time step (subintervals of a time interval $[0, T]$ of length Δt). Let t act as an index, so that $Z^{(t)}$ is the zonotope representing the reachable set over the time step containing t . Briefly, CORA works by linearizing the dynamics at each time step about the center of $Z^{(t)}$, and obtaining $Z^{(t+\Delta t)}$ by multiplying $Z^{(t)}$ by an overapproximation of the matrix exponential over that time step. It then expands $Z^{(t+\Delta t)}$ to account for the effects of inputs and linearization error.

CHAPTER 2

Human Subject Experiments

The research in this document utilizes data collected from two human subject perturbative experiments conducted at the University of Michigan. This chapter describes each experiment for reference throughout the document. Please note that both datasets are publicly available:

- Cable pulls and foot shifts during Sit-to-Stand: <https://doi.org/10.7302/mhjr-k798>
- Trips during level-ground treadmill walking: <https://doi.org/10.7302/pvhg-q324>

2.1 Cable Pulls and Foot Shifts during Sit-to-Stand

We conducted the Sit-to-Stand (STS) experiment primarily to investigate the effect of disturbances on individuals' stability. Therefore, the experiment introduced two different types of perturbations to nominal Sit-to-Stand motion. The first, shifting the feet forward or backwards from the subject's natural foot placement, acted as a perturbation to an individual's preferred initial condition. If the feet were shifted too far forward, the subject could be unable to generate enough momentum to successfully stand up with their natural strategy and need to sit back down. If the feet were shifted too close, the subject could need to step forward with one foot during the Sit-to-Stand motion to adjust their base of support. The second type of perturbation, sagittal-plane cable pulls applied to the subject's approximate midsection, elicited similar sitting/forward step responses. Both of these compensatory movements are considered failures of the subject's natural control strategy [31].

This experiment collected 948 trials from 11 participants: three female and eight male; ages 18-32; mean (standard deviation) height 1.70 (0.12) m; mean (standard deviation) body mass 65.4 (10.2) kg. The breakdown of trial type was 163 nominal, 194 foot-shifted, and 591 cable pull trials. Each subject gave their informed written consent, and had no physical or balance disorders which could affect their ability to perform Sit-to-Stand. The experimental protocol was approved by the University of Michigan Health Sciences and Behavioral Sciences Institutional Review Board, eResearch ID: HUM00020554.

Subjects began in a seated position on a stool with their arms crossed. The height of the stool was adjusted so that the subject's thighs were parallel to the ground. Subjects practiced standing up from the stool and were asked to find a comfortable foot position, which was then demarcated with a line of tape.

Subjects were instructed to perform three different Sit-to-Stand control strategies: their natural strategy, a Momentum-Transfer strategy, and a Quasi-Static strategy as done by Shia et al. [32]. Researchers have identified the Quasi-Static and Momentum-Transfer strategies as either end of a continuum of Sit-to-Stand strategies [33]. The Quasi-Static strategy involves shifting the feet closer to the chair and the hips forward, then standing up so that the center of mass remains inside the base of support. The Momentum-Transfer strategy involves shifting the feet further out from the chair, then using the forward momentum of the center of mass to propel the individual to standing. Hughes et al. found both of these strategies exhibited by older adult subjects [33], while younger adults tend to use the momentum transfer strategy [34]. Subjects watched a demonstration of each strategy, and then practiced each strategy a minimum of 10 times prior to data collection.

The following set of treatments were applied to each Sit-to-Stand strategy, illustrated in Fig. 2.1:

- Nominal trials: Subjects stood five times from a comfortable foot position using the specified control strategy.
- Foot-shift trials: Subjects varied their foot placement in 0.05-m (2-inch) increments in the anterior-posterior direction from their original position. The increments were demarcated by taped lines on the ground. Subjects stood up once at each increment, moving their feet backwards until their heels left the ground, and then forward until a strategy shift was observed. We noticed that subjects at extreme anterior foot positions attempted to stand up by lunging forward into a squatting position, and excluded these trials from our dataset because they represented a shift in strategy from the subject's nominal behavior.
- Cable pull perturbations: Subjects returned their feet to their original position. A cable system was attached to the subject's waist and connected to two high-torque motors. These motor driven cables applied impulses to the subject as they rose. The cable pulls were restricted to the anterior-posterior direction, and were applied either forward or backwards with variable timing and force. Specifically, three peak force levels – low, medium, and high – were calibrated to each subject. The low force level was designed to rarely induce stepping or sitting during Sit-to-Stand, while the high force level was designed to induce stepping or sitting approximately half of the time. Six trials were taken at each force level for each Sit-to-Stand strategy, with three pulling forwards and three pulling backwards, in random order.




Figure 2.1: Example of each type of Sit-To-Stand trial.

A 10-camera PhaseSpace motion capture system collected kinematic observations of 36 markers at 480 Hz. C-Motion's Visual3D biomechanics software was used to fit body segment models to each subject's data [35]. MATLAB was used for all subsequent analyses [36]. A 6th-order Butterworth filter with a cut-off frequency of 2 Hz was used to filter joint position trajectories.

For each natural and Momentum-Transfer STS trial, we determined the start and end of the trial through the following steps:

1. Manually segment each trial, starting several seconds before STS was initiated and ending several seconds after the subject reached standing.
2. Create an average nominal trial by aligning all of a subject's nominal trials by their peak horizontal center of mass (COM) velocity, then computing the mean.
3. Choose the start time of the average nominal trajectory based on when the horizontal acceleration exceeds 20% of its max. Choose the end time as when the vertical COM position first exceeds 99% of its max.
4. Nondimensionalize the segmented, average trial in time, described below.

The foot-shift and cable-pull trials were segmented to align with the average nominal trial’s peak COM accelerations. The Quasi-Static trials were segmented using a separate segmentation procedure because they lacked a distinct peak in their horizontal COM accelerations. More detail about the segmentation process can be found in Appendix B of [23].

Each STS trial was also nondimensionalized in time. Let $\tau \in [0, T] \in \mathbb{R}$ denote a time drawn from an interval during a segmented STS trial and $\tilde{x} : [0, T] \rightarrow X \subset \mathbb{R}^n$ and $\tilde{u} : [0, T] \times X \rightarrow U \subset \mathbb{R}^m$ be a state and input trajectory, respectively, of a chosen STS model with n states and m inputs. Then, the dynamics of this model are given by

$$\dot{\tilde{x}}(\tau) = f(\tilde{x}(\tau), \tilde{u}(\tau, \tilde{x}(\tau))). \quad (2.1)$$

To make accurate comparisons across trials, we normalize each state trajectory \tilde{x} by the trial’s length, so that each trajectory occurs over the interval $[0, 1]$, which we refer to as 0 to 100 %STS. We introduce a unitless time variable t , related to τ by $t = \frac{\tau}{T}$. This dimensionless variable allows us to introduce a normalized state trajectory $x : [0, 1] \rightarrow X$, where

$$x(t) = \tilde{x}(T \cdot t). \quad (2.2)$$

We scale accelerations and forces in (2.1) by a factor of T^2 following similar logic.

2.2 Trips during Level-Ground Treadmill Walking

This experiment aimed to study human responses to trips during level-ground treadmill walking. Instead of using physical obstacles to create trips, we used tethers attached to the subject’s feet and routed to the back of the treadmill through a custom braking device, shown in Fig. 2.2. The braking device was activated during the swing phase to interrupt the forward motion of the foot, thus inducing a trip-like perturbation. In each tether trip perturbation, the brake was applied for either 150, 250, or 350 ms, decided randomly during the experiment. These three trip durations were used in [25] to emulate different lengths of foot contact with an obstacle, either when during initial impact or when the foot gets caught and cannot immediately overcome an obstacle. Shirota et al. found that this technique resulted in a similar responses to trips as those from physical obstacles [25].

One advantage of the tethers is that the subjects are unable to anticipate the trip while maintaining their normal field of view. In experiments where obstacles are released onto the treadmill, subjects often wear special goggles to limit their lower field of view [37, 38]. Although this strategy discourages subjects from anticipating the obstacles, restricting the lower field of view may

introduce other balance and mobility challenges.




Figure 2.2: Trip experiment setup, including tethers attached to each foot, brakes applied at short intervals, and load cells to measure the tether force.

We collected data from 16 participants: Eight female and eight male; ages 18-29; mean (standard deviation) height 1.77 (0.10) m; mean (standard deviation) body mass 69.6 (12.6) kg. Each subject gave their informed written consent, and had no physical or balance disorders which could affect their ability to walk on a treadmill. The experimental protocol was approved by the University of Michigan Health Sciences and Behavioral Sciences Institutional Review Board, eResearch ID: HUM00195042.

Subjects were fit with a harness in case of falls, then the tethers were attached to each foot. The subjects selected their desired level-ground treadmill (Bertec) walking speed in a warm-up that started at 1 m/s. After one minute, subjects could request to increase or decrease the speed in 0.2 m/s increments. The mean (standard deviation) chosen walking speed was 1.14 (0.16) m/s. Then, the perturbed trials commenced. In total, the experiment consisted of 100 20-second walking trials, divided into 4 sessions of 25 trials each. A custom targeting algorithm measured force plate values from the instrumented treadmill in real-time, estimating whether each leg was in stance or swing. The trips were randomized by the following variables:

- leg side (right or left),
- percentage of swing phase (10-80%),
- percentage of time into each 20-second trial (5-15 seconds), and

- duration of tether brake (150, 250, or 350 ms) to simulate varying contact duration with an obstacle.

Note, due to latencies in the signals sent from the computer to the tripping device and our measures to correct for them, we observed trips outside of the 10-80% swing phase range. A minimum of 10 seconds between each trip allowed time for the subject to recover, and 15 random trials did not include trips. We observed that subjects' kinematics returned to baseline behavior (by a visual analysis of lower-limb kinematics; example provided in Fig. 2.3) after 2-3 steps, indicating that the 10-second minimum break between trips gave subjects adequate time for recovery. However, quantifying baseline levels in future work could help to further verify this experimental choice. Load cells measured the force in each tether throughout the experiment, and a 26-camera Vicon motion capture system collected kinematic observations of 16 markers at 100 Hz. We applied a Woltring filter and fit lower-body segment models to each subject's data using Nexus Plug-in Gait [39], and we used MATLAB for all subsequent analyses [36]. A 6th-order Butterworth filter with a cut-off frequency of 6 Hz was used to filter joint position trajectories.

We analyzed these experimental data during the swing phase only, i.e. between toe-off and heel-strike events. We determined each gait event using force plate readings with a threshold of 30 N for the stance phase. In Ch. 5, we analyze the trip-recovery trajectories using a normalized and shifted time variable that aligns peak heel height. In Ch. 6, we use a phase variable to describe the progression of gait, which is more amenable to real-time planning and control in robotic prostheses. We refer to either normalized time or phase as $t \in [0, 1] \subset \mathbb{R}$ in their respective chapters.

Previous research by Shirota et al. investigated the effect of trip onset, brake duration, leg side, and the interaction between brake duration and onset in a similar trip-recovery experiment [25]. The authors found that leg side did not significantly influence strategy selection in the experiment subjects. In this document, we therefore combine the swing phase trajectories for the left and right foot into one dataset. Future work should conduct a similar multinomial logistic regression to verify that leg side did not influence strategy selection in our own experiment. This analysis could also include time into the experiment as another independent variable, to investigate if strategy selection changed over time.

Humans typically pick one of three strategies to recover from trips, shown in Fig. 2.5:

1. Elevating: The subject's swing leg elevates over the obstacle, and the swing heel contacts the ground in front of the obstacle.
2. Delayed Lowering: The subject's swing leg initially elevates, then lowers behind the obstacle. The swing toe contacts the ground behind the obstacle.
3. Lowering: The subject's swing leg immediately lowers behind the obstacle, with the swing toe contacting the ground first.




Figure 2.3: This figure shows the (a) tripped heel, (b) contralateral swing heel, (c) tripped hip, (d) tripped knee, and (e) tripped ankle trajectories after a trip for one subject in our dataset. We observe changes in the kinematic trajectories immediately after trip onset, indicated by the vertical dashed line, and a return to baseline levels after 2-3 steps.

The trip-recovery data in this document was sorted into the three strategies by hand. To sort the trip-recovery strategies, we analyzed the anterior-posterior position and height of the swing heel compared to average nominal level-ground walking trajectories. We also looked at the swing leg configuration throughout the tripped swing phase. Sometimes, the trips did not induce a noticeable recovery, due to signal latency or the tether temporarily becoming displaced from between the brake pads. In this case, the recoveries were labeled as “incomplete arrest” as in [25] and were not analyzed. Examples of the sorting criteria are shown for elevating, delayed lowering, lowering,

and “incomplete arrest” strategies in Fig. 2.4

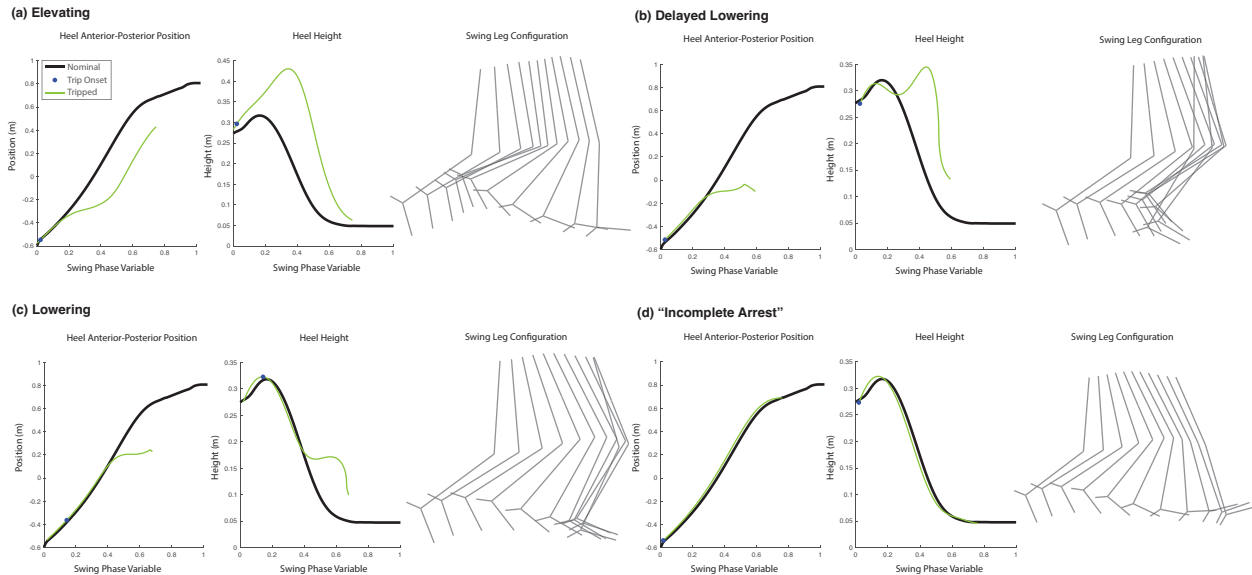


Figure 2.4: We analyzed the swing heel anterior-posterior position, swing heel height, and swing leg configuration throughout the tripped step to sort the trip recoveries by hand. The horizontal axis in all plots is a variable that represents progression through the swing phase, described in Ch. 6. We label (a) as an elevating strategy because the swing heel elevates (above average nominal behavior) after the trip occurs, with the heel striking in front of the body; (b) as a delayed lowering strategy because the swing heel initially elevates, then lowers behind the body with the toe striking the ground first; (c) as a lowering strategy because the swing heel immediately lowers behind the body after a trip with the toe striking the ground first; and (d) as an “incomplete arrest” because there is little difference from average nominal behavior.

In this document, we analyze over 900 trip trials across 16 subjects. Chs. 5 and 6 use slightly different total numbers because we re-sorted some trials for robotic prosthesis planning, and each chapter lists the number and breakdown across strategies.

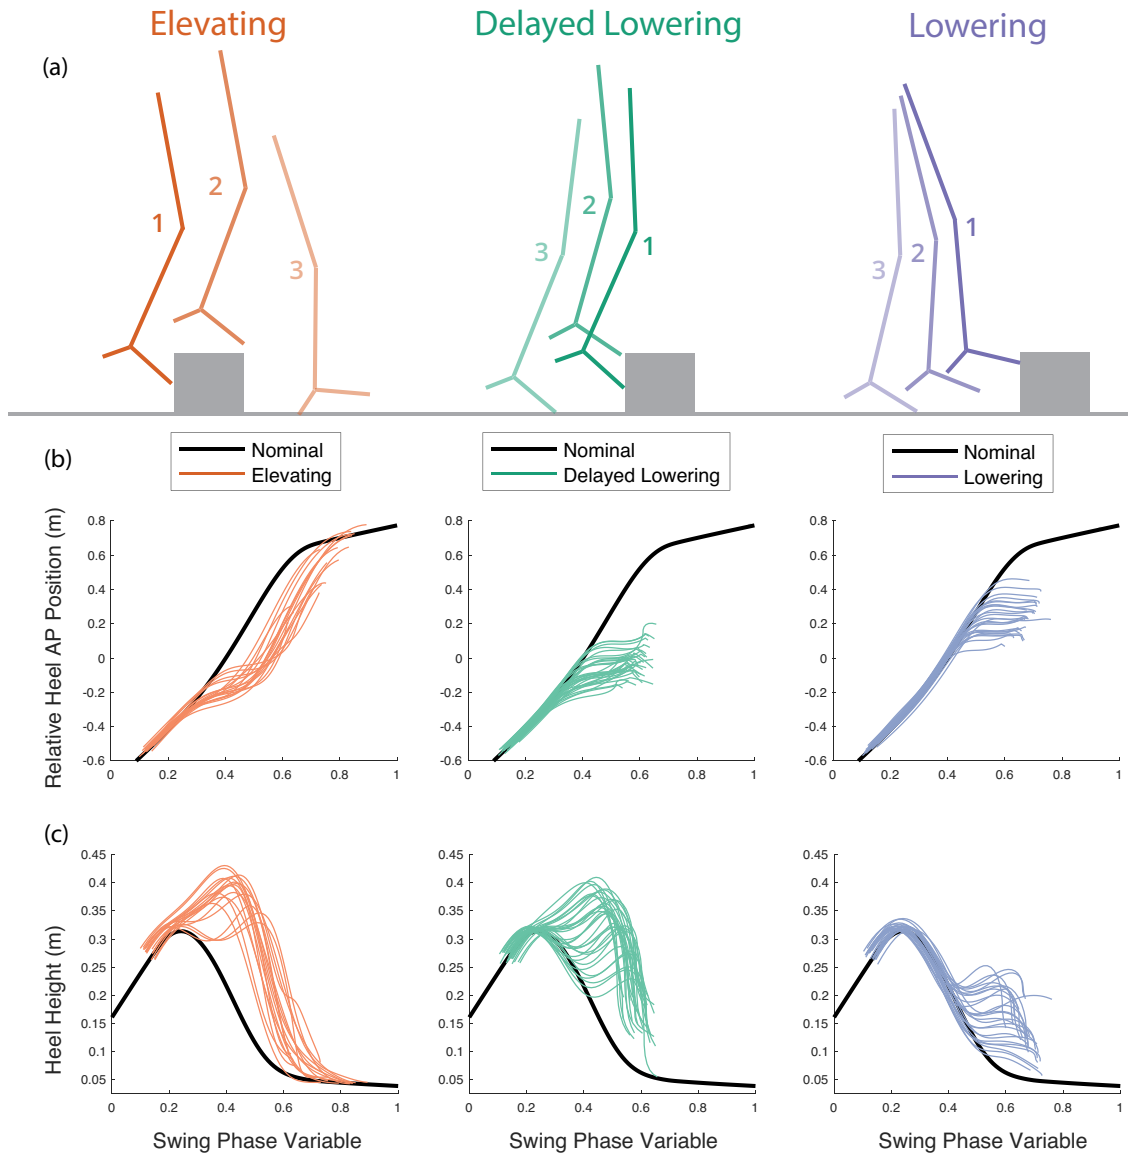


Figure 2.5: The three trip-recovery strategies with corresponding swing heel trajectories from one subject (Subject ID 1) in our dataset. The horizontal axis is a variable that represents progression through the swing phase, described in Ch. 6. (a) A depiction of the swing leg during each strategy, adapted from Shirota et al. 2014. (b) A sample of anterior-posterior (AP) swing heel positions relative to the stance foot during each type of recovery compared to nominal (level-ground walking) data. Note that the swing heel reaches close to the same final position as nominal during the elevating strategy, whereas the swing heel travels backwards in the anterior-posterior position during the delayed lowering and lowering strategies. (c) A sample of swing heel heights during each type of recovery compared to nominal data. Note that the swing heel reaches ground level during the elevating trials, but does not reach the ground level during delayed lowering and lowering trials (aside from one instance). During the latter two recovery strategies, the swing toe typically strikes the ground first, so the heel stays above the ground.

CHAPTER 3

Specific Aim 1: Quantifying Stability During Sit-To-Stand

This work is published [23]:

Holmes Patrick D., Danforth Shannon M., Fu Xiao-Yu, Moore Talia Y., and Vasudevan Ram 2020. Characterizing the limits of human stability during motion: perturbative experiment validates a model-based approach for the Sit-to-Stand task. *Royal Society Open Science*.

and has been modified for this document.

3.1 Introduction

Falls are the leading cause of injury in people over 75 years old, resulting in reduced quality of life, increased healthcare costs, and accident-related death [40, 41, 42]. If at-risk individuals can be identified prior to injury, the likelihood of falling can be significantly reduced through intervention such as physical therapy [43, 44]. Fall risk generally results from instability arising from neuromuscular deficiencies or external perturbations. Clinical assessments for identifying individuals who would benefit from preventative care are currently limited to questionnaires [45, 46] and non-perturbative motor assessments [10, 11, 47]. Self-reported information often has low reliability [48, 49], and current clinical motor performance tests have low fall-prediction rates, especially for active older adults [13, 50, 51, 14]. For widespread use, stability assessments must combine predictive power with minimal experimental and computation time.

Because older adults are more likely to fall while in motion [52], several studies suggest that quantifying dynamic stability may help identify biomechanical deficiencies associated with an increased risk of falling [53, 54, 55, 56]. Thus, a number of model-based methods have been developed to assess stability during walking [56]. Among these, variability measures [57] and

the maximum Lyapunov exponent [58] ranked highest overall in validity. However, these metrics only characterize a subject's ability to recover from small perturbations. Currently, it is unclear whether the most useful estimates of stability are provided by measuring an individual's ability to recover from small perturbations or by computing the largest possible perturbation that they can withstand without failure of a control strategy [56]. Unfortunately, these methods cannot yet be fully compared because a verified technique for computing the maximum perturbation from which a subject can recover does not exist. Additionally, the best-performing model-based methods are limited to periodic motion.

Stability during non-rhythmic motion is also of interest, especially because difficulty with aperiodic tasks such as Sit-to-Stand (STS) is strongly correlated with falls in older adults [21], and because these tasks are necessary for maintaining independence and quality of life. The motion of a person's center of mass (COM) during STS can be modelled as an inverted pendulum, which requires an appropriate amount of angular momentum at seatoff to successfully complete the task [59]. Drawing from this idea, metrics of stability for STS are generally based on an individual's initial COM velocity or acceleration [60, 61, 62]. However, a stability metric that considers data only at the onset of movement disregards valuable information about the control strategy used by the individual. In fact, human control strategies for STS lie on a spectrum ranging from Quasi-Static, in which little momentum is used and the body position is statically stable throughout the motion, to Momentum-Transfer, which is statically unstable and extensively uses momentum to achieve standing [33].

To account for the dynamic differences in STS motions under distinct control strategies, Shia et. al introduced the Stability Basin, computed using individualized pendulum models of STS with linear-quadratic regulator (LQR) controllers [32]. For this pendulum model, the Stability Basin at a given time is equal to the set of pendulum states that are able to successfully achieve standing under the specified controller. Though this Stability Basin under an LQR controller successfully distinguishes between less and more stable STS strategies, its ability to accurately identify the set of states that can arrive at standing without failure of a given control strategy is unverified.

To test whether the Stability Basin can accurately estimate an individual's stability for a particular task, perturbations must be introduced in a way that causes the states of the individualized dynamic model to exit the stable region. By stability, we mean the set of body configurations through time from which a subject can successfully stand up without switching from their chosen control strategy to stepping or sitting. An accurate prediction of instability corresponds to an experimentally observed failure of the individual's control strategy.

Here we use a perturbative STS experiment with 11 subjects, described in Ch. 2, Sec. 2.1, to validate stability predictions generated by the Stability Basin method. We use a dynamic model, presented in Sec. 3.2.2, and subject-specific controllers to form individualized Stability Basins,

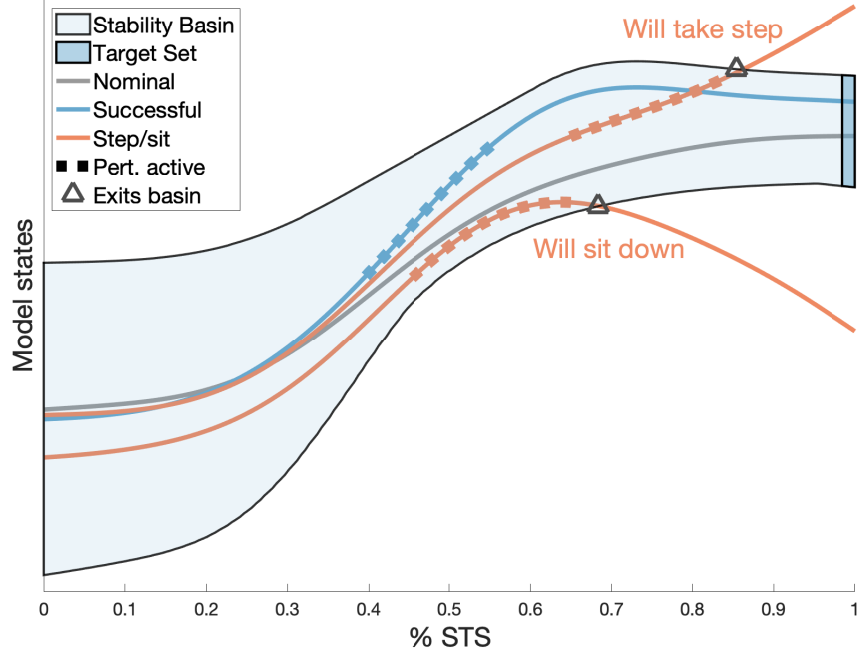


Figure 3.1: An illustrative overview of Stability Basins for Sit-to-Stand. The Stability Basin represents the set of model states through time that will successfully arrive at a target set for a given individual and Sit-to-Stand strategy. Trajectories of the model are illustrated, where the times that perturbations are applied are denoted by the dashed lines. Trajectories that exit the Stability Basin are predicted to lead to stepping or sitting.

detailed in Sec. 3.2.4. Then, we test whether each subject’s Stability Basin can accurately predict the failure of a control strategy to recover from the cable pull, described in Sec. 3.2.5. We consider common failure modes, such as a step or sit, to be observed failures of a STS control strategy [31]. In contrast to prior work that assumed that humans use LQR control during Sit-to-Stand while constructing the Stability Basin, we use a data-driven method to create bounds on the control input for a particular subject and control strategy, described in Sec. 3.2.3. Across all subjects and control strategies, the proposed method correctly predicts unsuccessful trials over 90% of the time, and correctly identifies successful trials over 95% of the time in a leave-one-out assessment. When compared to three other methods of assessing Sit-to-Stand stability, the proposed method outperformed each alternative. This perturbative validation experiment demonstrates that the Stability Basin method is accurate and reliable, showing promise for integration into current healthcare infrastructure.

3.2 Methods

The goal of this work is to characterize stability during Sit-to-Stand. We hypothesize that Stability Basins can accurately represent stability. The Stability Basin can be understood intuitively as follows. Given a dynamical model of Sit-to-Stand, kinematic observations of a subject performing Sit-to-Stand can be used to form trajectories of the model. The Stability Basin defines a boundary between “successful” and “unsuccessful” trajectories of the model, where “unsuccessful” implies the subject will have to take a step or sit back down. The Stability Basin is introduced pictorially in Fig. 3.1, with a formal definition and computational details given in Sec. 3.2.4. An accurate Stability Basin should predict when a subject’s Sit-to-Stand motion will be successful or unsuccessful based on model trajectories alone. To test the accuracy of Stability Basins, we performed a perturbative Sit-to-Stand experiment (Ch. 2, Sec. 2.1), occasionally inducing subjects to step or sit back down, and then compared Stability Basin predictions to experimentally observed outcomes.

This section describes our framework for computing and testing the accuracy of Stability Basins. We use kinematic observations of a subject performing a perturbative Sit-to-Stand experiment, described in Ch. 2, Sec. 2.1. Individualized biomechanical models of Sit-to-Stand are constructed for each subject (Sec. 3.2.2), and three controller models are detailed in Sec. 3.2.3. Our proposed controller, which uses strategy-specific input bounds, reflects the distinct range of inputs required for each control strategy [63]. After computing the Stability Basins (Sec. 3.2.4) for each controller model, we test whether the individual and strategy-specific Stability Basins correctly predict when a subject steps or sits down in response to perturbation, and when they are successful (Sec. 3.2.5). Finally, we perform a comparison to a naive method for estimating stability (Sec. 3.2.4.3).

3.2.1 Classifying trials as successful/unsuccessful

Two clear instances of a subject’s chosen control strategy (natural, Momentum-Transfer, or Quasi-Static) failing during Sit-to-Stand are taking a step and sitting back down. We refer to the trials in which these occurred as *steps* and *sits*, and collectively label them *unsuccessful*. At times, we will also refer to these trials as *failures*.

For step trials, the instance of failure is defined as the time when the subject moves the toes of either foot more than 0.0762 m (3 inches) in the anterior-posterior direction from its starting point, as measured by motion capture markers placed on the subject’s foot. Although previous studies measured ground reaction forces to identify the onset of steps [31], the distance threshold was used here to utilize motion capture data. For sit failures, the failure instance is defined as the time when both horizontal and vertical velocities of the COM become negative, i.e. the COM begins to move back towards the seat. If a subject both stepped and sat down, failure is defined as the earlier of the

two.

All foot-shift trials were considered successful. Subjects attempting the extreme anterior and posterior foot positions were either unable to initiate the Sit-to-Stand movement, or violated the parameters of acceptable trials stated in Ch. 2, Sec. 2.1. In either case, these trials were not included in the dataset.

If no step or sit is detected, the trial is considered *successful*. In this work, ‘successful trials’ refers to nominal trials, foot-shift trials, and successful perturbed trials. Note, this scheme classifies trials in which subjects rocked far onto their heels or toes to maintain balance as successful, so long as no step or sit occurred.

3.2.2 Dynamic modeling

Let $\tau \in [0, T] \subset \mathbb{R}$ denote a time drawn from a time interval. Without loss of generality, we assume each Sit-to-Stand trial starts at time $\tau = 0$. We use a Telescoping Inverted Pendulum Model (TIPM, Fig. 3.2) to model each subject’s COM motion in the sagittal plane during each Sit-to-Stand trial [64]. This model is capable of describing the sizable displacements of a subject’s COM in the horizontal and vertical directions that occur during Sit-to-Stand. The TIPM consists of a point mass of mass $m \in \mathbb{R}$ representing the subject’s COM. Individualized TIPMs are constructed for each subject by setting m as the subject’s mass.

We define the origin as the initial position of the ball of the subject’s foot. Specifically, the origin is defined as the mean initial position of motion capture markers attached to the metatarsophalangeal joints of the subject’s left and right feet, as estimated by Visual3D. We let the subscript $(\cdot)_x$ denote a quantity in the anterior-posterior direction, and the subscript $(\cdot)_y$ denote a quantity in the vertical direction. At each time $\tau \in [0, T]$, we denote the following quantities of the TIPM as:

- Positions: $\tilde{r}_x(\tau), \tilde{r}_y(\tau) \in \mathbb{R}$
- Velocities: $\tilde{v}_x(\tau), \tilde{v}_y(\tau) \in \mathbb{R}$
- Accelerations: $\tilde{a}_x(\tau), \tilde{a}_y(\tau) \in \mathbb{R}$
- Inputs: $\tilde{u}_x(\tau), \tilde{u}_y(\tau) \in \mathbb{R}$
- Cable-pull forces: $\tilde{d}_x(\tau), \tilde{d}_y(\tau) \in \mathbb{R}$

Let $\tilde{x} : [0, T] \rightarrow X \subset \mathbb{R}^4$ be a state trajectory, so that $\tilde{x}(\tau) \in X$ is the model’s state at time

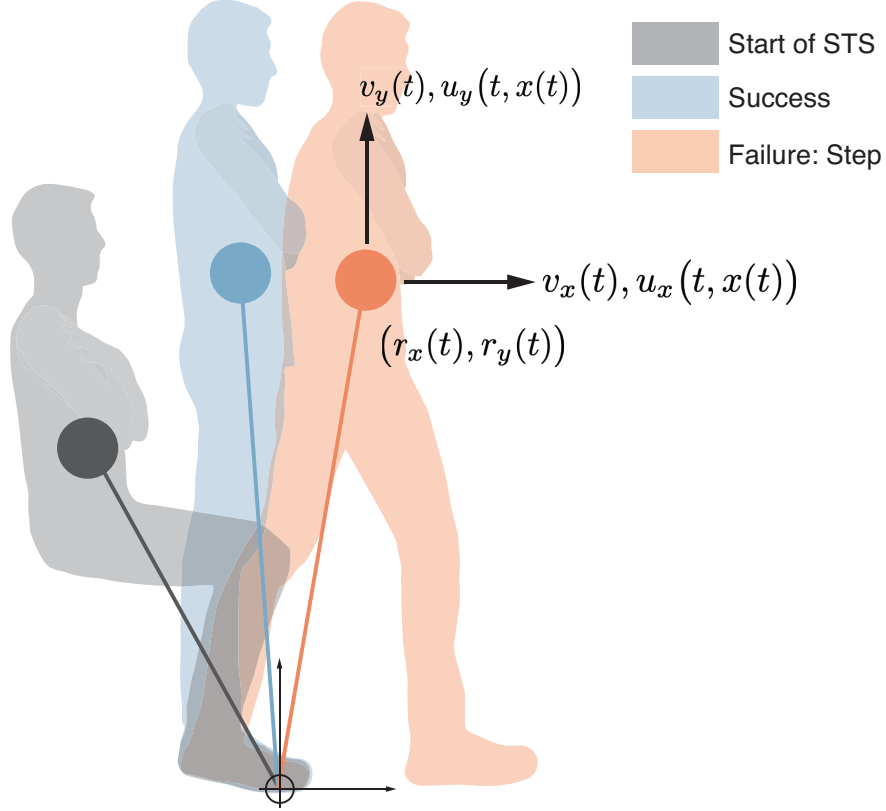


Figure 3.2: Subjects began from a seated position with their arms crossed against their chests. The subject's COM is illustrated by filled circles. Cable-pulls applied to the subject sometimes caused them to step or sit; otherwise, the trial was considered to be successful.

$\tau \in [0, T]$:

$$\tilde{x}(\tau) = \begin{bmatrix} \tilde{r}_x(\tau) \\ \tilde{v}_x(\tau) \\ \tilde{r}_y(\tau) \\ \tilde{v}_y(\tau) \end{bmatrix}. \quad (3.1)$$

Let $\tilde{u} : [0, T] \times X \rightarrow U \subset \mathbb{R}^2$ be an input trajectory, so that $\tilde{u}(\tau, \tilde{x}(\tau))$ is the input at that time and state:

$$\tilde{u}(\tau, \tilde{x}(\tau)) = \begin{bmatrix} \tilde{u}_x(\tau, \tilde{x}(\tau)) \\ \tilde{u}_y(\tau, \tilde{x}(\tau)) \end{bmatrix}. \quad (3.2)$$

The time derivative of $\tilde{x}(\tau)$ at time τ , denoted $\dot{\tilde{x}}(\tau)$, can be written as:

$$\dot{\tilde{x}}(\tau) = \begin{bmatrix} \tilde{v}_x(\tau) \\ \frac{1}{m} \left(\tilde{u}_x(\tau, \tilde{x}(\tau)) + \tilde{d}_x(\tau) \right) \\ \tilde{v}_y(\tau) \\ \frac{1}{m} \left(\tilde{u}_y(\tau, \tilde{x}(\tau)) + \tilde{d}_y(\tau) \right) - g \end{bmatrix}, \quad (3.3)$$

where g is the gravitational acceleration $9.81m/s^2$.

The time that it takes to complete Sit-to-Stand varies from trial to trial. To make accurate comparisons across trials, we normalize each TIPM state trajectory \tilde{x} by the trial’s length, so that each trajectory occurs over the interval $[0, 1]$, described in Ch. 2, Sec. 2.1. This normalization introduces a unitless time variable t and a normalized state trajectory $x : [0, 1] \rightarrow X$. For the rest of this chapter, we only consider normalized trajectories, unless explicitly stated. Let r_x, r_y, v_x and v_y refer to normalized position and velocity trajectories, and let a_x, a_y, u_x, u_y, d_x and d_y refer to normalized acceleration and force trajectories.

3.2.3 Controller Model

The Stability Basin computation, detailed in Sec. 3.2.4, requires a model for the TIPM controller. In this subsection, we describe and compare two existing controller models as well as our proposed controller model. All three controller models are used to compute Stability Basins that are compared during the validation phase described in Sec. 3.2.5.3. Before we describe these controller models in mathematical detail, we briefly describe their qualitative properties. The controllers are compared and contrasted in Fig. 3.3.

The controller models estimate the control input $u(t, x(t))$ at a given time t and state $x(t)$. In many previous studies, researchers have modelled human controllers as an open-loop “feed-forward” component plus a closed-loop linear “feedback” component [65, 66, 32], thus finding a single controller that best models the data. Following this paradigm, we develop an LQR controller model and a traditional feedforward plus feedback controller model (FF+FB). The main differences between these controllers are the specifics of their construction, detailed in Sec. 3.2.3.2 and Sec. 3.2.3.3. The LQR controller model is constructed by solving an optimal control problem to provide feedback about an average nominal trajectory, while the FF+FB controller model is constructed via linear regression. Both of these controller models return a single input at a given time and state.

Because the form of the controller employed by humans is unknown, it is difficult to assert with confidence that a single controller adequately captures human behavior. We propose that a more robust controller model may return a set of possible inputs rather than a single input at a given time

and a state. One way to achieve this is to bound the possible inputs at each time and state, and assume that the human could use any input within these bounds. We refer to this controller model as the *Input Bounds* controller.

The Input Bounds represent the range of time-varying inputs that are expected under a given Sit-to-Stand strategy. For example, consider a subject’s natural Sit-to-Stand strategy. At the beginning of the motion, we expect the subject to apply a large positive horizontal input to propel themselves forwards, while applying little positive vertical input. Halfway through the motion, the subject may apply a negative horizontal input to slow their forward velocity, as well as a large positive vertical input to rise towards standing. Therefore, the expected inputs applied by the subject are time-dependent. Moreover, because there is some variability in the way the subject performs Sit-to-Stand from trial to trial, we expect the inputs to fall within some range of values. The Input Bounds controller captures this variability by developing time-dependent upper and lower bounds on the allowable inputs.

Biomechanical constraints, such as ground reaction force limits or joint torque limits, constrain the maximum inputs that a subject can apply during Sit-to-Stand under any strategy. We want to be clear that the Input Bounds *do not* represent the range of all inputs that are biomechanically feasible, but rather the subset of inputs that represent the range expected under a given strategy. For example, consider a scenario in which a person begins standing up using their natural Sit-to-Stand strategy, but then stops halfway and remains in a crouched position. Certainly, the person can achieve this performance without violating biomechanical constraints. However, the person will have deviated from their natural Sit-to-Stand strategy, in which they normally stand all the way up. Accordingly, the input that they applied would have exited the Input Bounds we have developed for their natural strategy at the time that they decided to remain in a crouched position. For this reason, we develop the Input Bounds from observations of Sit-to-Stand rather than from biomechanical constraints.

3.2.3.1 Generating training data

Modelling human motion requires the LQR, FF+FB, and Input Bounds controller models to be estimated from observed data. In this work, we use kinematic observations of successful trials (i.e. nominal, foot-shift, and successful perturbed trials) to generate a dataset on which to train each control model. Note that the dataset on which each controller model is trained is identical. We fit the LQR, FF+FB, and Input Bounds controllers separately for each subject’s natural, Momentum-Transfer, and Quasi-Static control strategies.

Given a strategy, we first compute the inputs for each successful trial within that strategy via inverse dynamics using (3.3) [67, Chapter 5]. Let $x_i(t)$ and $u_i(t)$ denote the state and input at time t for the i^{th} successful trial computed via inverse dynamics, and let S denote the set of successful

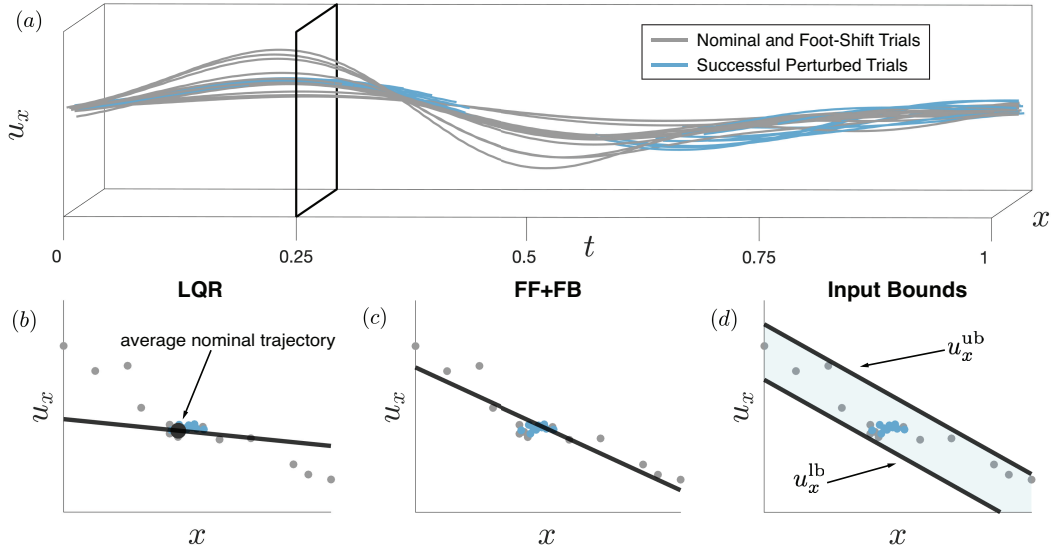


Figure 3.3: These figures show how we construct the controller models from data. The process is shown for u_x , and is identical for u_y . In Fig. 3.3(a), u_x evolves as a function of time t and state x , which has been depicted as a single dimension for this illustration. The grey lines represent the input trajectories of nominal and foot-shift trials, while the blue lines represent the input trajectories of perturbed successful trials. Note that we have not plotted the portions of the successful perturbed trajectories when the perturbations were active, because these segments are not used to train the controller models. An arbitrary time $t = 0.25$ is illustrated by the black outline in Fig. 3.3(a). The observed inputs from nominal, foot-shift, and perturbed successful trials at that time are plotted as points in Fig. 3.3(b - d). Figure 3.3(b) shows that the LQR controller utilizes linear feedback about a nominal trajectory. Because the LQR controller is generated by minimizing a cost function specified by the matrices Q and R , its feedback matrix does not necessarily coincide with a line of best fit. Figure 3.3(c) illustrates that the FF+FB controller is given by a line of best fit. Figure 3.3(d) shows how the upper and lower bounds $u_x^{lb}(t, x(t))$ and $u_x^{ub}(t, x(t))$ are parallel linear functions of state that encompass all of the successful inputs.

trials within a given strategy. Given a time t , we will frequently use the notation

$$\sum_{i \in S} (u(t, x_i(t)) - u_i(t))^2 \quad (3.4)$$

which means “the squared error of the estimated control input $u(t, x_i(t))$ from each observed control input $u_i(t)$, summed over all successful trials within a given strategy.”

Because many successful trials within a Sit-to-Stand strategy involved cable-pull perturbations, we make the following remark:

Remark 3.2.1. *When computing $u_i(t)$ via inverse dynamics for successful perturbed trials, we*

discard the portion of the trajectory during which the perturbation was active. This is possible because the perturbations were time-synchronized with the motion capture data.

3.2.3.2 LQR controller

Shia et al. proposed to model Sit-to-Stand using an LQR controller about a nominal Sit-to-Stand trajectory [32]. To build this controller, we first form a single average nominal trajectory \bar{x} for each of a subject's Sit-to-Stand control strategies by taking the mean of the five nominal Sit-to-Stand trajectories. We generate an open loop controller, $u_{ol} : [0, 1] \rightarrow U \subset \mathbb{R}^2$, for each average nominal trajectory via inverse dynamics. Then, we use LQR to design a linear feedback controller about the average nominal trajectory. We specify the quadratic cost function for LQR by a state weighting matrix $Q \in \mathbb{R}^{4 \times 4}$ and input weighting matrix $R \in \mathbb{R}^{2 \times 2}$, which are used to generate the feedback matrix $K(t) \in \mathbb{R}^{2 \times 4}$. [68, Chapter 16] At a given time and state, the input from the LQR controller can then be written as:

$$u(t, x(t)) = u_{ol}(t) - K(t)(x(t) - \bar{x}(t)) \quad (3.5)$$

We choose the Q and R that produces a controller that most closely fits the observed data. To do so, we wrap the choice of Q and R inside an optimization procedure that seeks to minimize the error between the LQR controller model and the observed control inputs over the time horizon $[0, 1]$:

$$\min_{Q,R} \sum_{i \in S} \int_0^1 (u(t, x_i(t)) - u_i(t))^2 dt \quad (3.6)$$

As in (3.4), S represents the set of observed successful trials within a Sit-to-Stand strategy, and $x_i(t)$ and $u_i(t)$ correspond to the observed states and input of the i^{th} trial at time t . The integral is approximated using 200 time steps. We specify the space over which Q is optimized:

$$Q \in \{M \in \mathbb{R}^{4 \times 4} | M_{i,j} = 0 \forall i \neq j, M_{1,1} = 1, M_{2,2} \in [0.1, 100], M_{3,3} \in [0.1, 100], M_{4,4} \in [0.1, 100]\} \quad (3.7)$$

and the space over which R is optimized:

$$R \in \{M \in \mathbb{R}^{2 \times 2} | M_{i,j} = 0 \forall i \neq j, M_{1,1} \in [1e-5, 1e-2], M_{2,2} \in [1e-6, 1e-1]\} \quad (3.8)$$

Q and R are constant diagonal matrices, where the spaces over each which is optimized were chosen to restrict the relative costs of states versus inputs, the relative costs between states, and the relative costs between inputs. In particular, no state can cost more than 100 times any other state, no input can cost more than 10 times another input, and the relative scale between states and inputs

was chosen based on the expected magnitudes of states and inputs.

3.2.3.3 FF + FB controller

We also test the efficacy of a traditional feedforward plus feedback controller. Let the input $u(t, x(t))$ at some time t and state $x(t)$ be described by the state feedback matrix $K(t) \in \mathbb{R}^{2 \times 4}$ and a feedforward component $\text{ff}(t)$:

$$u(t, x(t)) = \text{ff}(t) - K(t)x(t) \quad (3.9)$$

To generate $K(t)$ and $\text{ff}(t)$ from data, we use linear least squares to solve the following program:

$$\min_{\text{ff}(t), K(t)} \sum_{i \in S} (u(t, x_i(t)) - u_i(t))^2 \quad (3.10)$$

As in (3.4), S represents the set of observed successful trials within a Sit-to-Stand strategy, and $x_i(t)$ and $u_i(t)$ correspond to the observed states and input of the i^{th} trial at time t .

3.2.3.4 Input Bounds controller

The proposed Input Bounds controller differs from both the LQR and FF+FB controllers by returning a set of inputs at a given time and state, rather than a single point. To define the Input Bounds controller, we develop bounds on the TIPM control input. We denote a lower bound on the control input as u^{lb} , and an upper bound on the control input as u^{ub} . Note that

$$u^{\text{lb}} \leq u \leq u^{\text{ub}} \quad (3.11)$$

where the inequality is taken element-wise.

We present this approach more formally in the following definition:

Definition 3.2.1. *For the Input Bounds controller, we model the inputs u to the TIPM as a bounded component plus a linear feedback component, where the bounded component is drawn from a bounded set. Specifically, given time t and state $x(t)$, we define the bounds in (3.11) as*

$$u^{\text{lb}}(t, x(t)) = b^{\text{lb}}(t) - K(t)x(t) \quad (3.12)$$

$$u^{\text{ub}}(t, x(t)) = b^{\text{ub}}(t) - K(t)x(t) \quad (3.13)$$

where $b^{\text{lb}}(t)$ and $b^{\text{ub}}(t) \in \mathbb{R}^2$ are lower and upper bounds on the bounded component, and $K(t) \in \mathbb{R}^{2 \times 4}$ is a matrix of linear feedback gains.

We generate the parameters of the Input Bounds controller $b^{\text{lb}}, b^{\text{ub}}$, and K from data by solving a constrained linear least squares problem:

$$\begin{aligned} \min_{b^{\text{lb}}(t), b^{\text{ub}}(t), K(t)} \quad & \sum_{i \in S} \left(u^{\text{lb}}(t, x_i(t)) - u_i(t) \right)^2 + \left(u^{\text{ub}}(t, x_i(t)) - u_i(t) \right)^2 \\ \text{s.t.} \quad & u^{\text{lb}}(t, x_i(t)) \leq u_i(t) \quad \forall i \in S, \\ & u^{\text{ub}}(t, x_i(t)) \geq u_i(t) \quad \forall i \in S \end{aligned} \quad (3.14)$$

where the inequalities are understood element-wise. This is a quadratic program, which can be solved rapidly and efficiently.

In summary, at each instance in time, we compute the range of control inputs of a certain form that can be applied while explaining the data we observed, thereby modelling a range of inputs a human may use to recover from perturbation. Our method of Stability Basin computation, detailed in Sec. 3.2.4.2, is flexible to this type of controller specification, and in particular, we verify that this approach is not too conservative in Sec. 3.3. In effect, the Input Bounds controller model allows us to compute Stability Basins with a set of hypothetical controllers, instead of just one.

3.2.4 Computing Stability Basins

Three elements are required to form the Stability Basin for each of an individual's Sit-to-Stand control strategies:

1. The TIPM dynamics (3.3).
2. A controller model for the TIPM (e.g. (3.5), (3.9), (3.11)).
3. A *target set*, X_T , that encapsulates all model states observed to lead to successful standing.

Given these elements, we define the Stability Basin as follows:

Definition 3.2.2. *The Stability Basin $\subset [0, 1] \times X$ is a subset of times (0 - 100%STS) and model states from which a trajectory of the model obeying the dynamics (3.3) and controller model will arrive at the target set X_T .*

In the rest of this subsection, we explain how we compute the target set X_T from data. Then, we use this object in conjunction with a controller from Sec. 3.2.3 to generate the Stability Basins.

3.2.4.1 Generating the target set

The goal of the Sit-to-Stand motion is to arrive at a standing configuration. We define the target set X_T as a set that encompasses the states at time $t = 1$ observed to lead to successful standing. Note, the target set X_T does not represent quiet standing itself, which we consider a distinct task.




Figure 3.4: This figure illustrates how X_T (the target set) is formed from data for a given subject and strategy. The states (r_x and v_x) at 100% STS of each observed successful trial within a strategy are illustrated as blue dots. An example X_T represented by the shaded blue zonotope contains all of the states observed to lead to standing. The zonotope is parameterized by its center c and generator vectors $g^{(1)}$ and $g^{(2)}$. This example depicts a 2-dimensional X_T , but in reality they are 4-dimensional and encompass both horizontal and vertical components of the state space.

We define $X_T \subset X$ as a zonotope (Ch. 1, Sec. 1.3 and Appx. A) that encompasses the final states (i.e., the states at $t = 1$) of all successful trials observed for a given Sit-to-Stand strategy. Specifically, X_T is a zonotope with 4 generators computed from the final states as proposed by Stursberg [69, Section 3], where each generator is expanded by 5% to avoid observed states lying on the edge of the set. The target set generation process is depicted in Fig. 3.4.

3.2.4.2 Computing Stability Basins via reachability analysis

The reader may notice that the Stability Basin (Defn. 3.2.2) is defined as a set of states through time that satisfy some dynamics specified by an ordinary differential equation (ODE). Recalling the concepts introduced in Ch. 1, Sec, 1.3, the Stability Basin is a backwards reachable set: The set of states through time that arrive at the final target set X_T . To compute the Stability Basins, we use an open-source reachability toolbox called CORA [29, 30], introduced in Ch. 1, Sec, 1.3.

CORA represents the Stability Basin as a zonotope at each of a finite collection of *time steps*, which are subintervals of the interval $[0, 1]$ of length Δt . Each Stability Basin was formed using a time step Δt of 0.005 (i.e., 0.5% STS), so that 200 zonotopes represent the Stability Basin over the interval $[0, 1]$. Although CORA chooses the exact number of generators to use for each zonotope, we set the maximum number of generators as 800.

Generally, CORA expects the inputs to the system to be functions of time, allowing the feed-forward inputs and gain matrices for the LQR and FF+FB controllers to be specified at each time

step. However, CORA also allows the input to the system at each time step to be drawn from a set, allowing us to use the Input Bounds controller model and let the control input be defined as in (3.12) and (3.13). CORA in this case assumes the input at a given time and state can take any value within the Input Bounds to create the Stability Basin. Because the Stability Basin is the set of states through time that arrive at the final target set X_T , it can be computed as a backwards reachable set. In practice, this means we treat X_T as an initial set, and use CORA to flow the set *backwards* in time under the negative of the dynamics (3.3) and the controller model.

Stability Basins were generated for each controller described in Sec. 3.2.3 on a laptop computer with a 2.7 GHz Intel Core i7 processor. Stability Basins computed using the Input Bounds controller take 0.76 ± 0.014 seconds to compute.

3.2.4.3 Naive method for computing Stability Basins

Stability can also be estimated from observed perturbed trials by simply drawing a volume around the state trajectories of the observed successful trials. This method does not utilize reachability or a controller model, and relies on state trajectories alone. Here, we apply this method by generating a 4-dimensional zonotope at each time step that encloses all of the observed successful states. This is the same procedure that we employ to generate the target set X_T (described in Sec. 3.2.4.2 and illustrated in Fig. 3.4) applied at each instance in time over the course of the motion.

3.2.5 Evaluating Stability Basin accuracy

We hypothesize that the Stability Basins can predict when a subject must depart from their control strategy in response to perturbation to avoid falling. Previously in Sec. 3.2.1, we defined the failure of a control strategy as stepping or sitting back down and explained how we determine the onset of stepping or sitting. We now use the Stability Basins to predict whether or not a strategy failure will occur by checking if a Sit-to-Stand trajectory exits the Stability Basin at any point. Finally, we detail the evaluation procedure we used to test the accuracy of the Stability Basins' predictions.

3.2.5.1 Defining onset of failure

To test the predictive power of the Stability Basins, we first identify the onset of failure (as previously described in Sec. 3.2.1) during a Sit-to-Stand movement. We denote $t_f \in [0, 1]$ as the time of failure onset, and will test whether the Stability Basins predict failure prior to that time. Because we are using a procedure based on an average nominal trial for aligning and segmenting trials (detailed in Appendix B of [23]), it is possible for the onset of stepping or sitting to occur after the trial end time $t = 1$. The end times chosen by the segmentation procedure do not necessarily imply that the subject has reached standing, but are just a means to achieve a consistent segmentation of




Figure 3.5: This figure illustrates how to check if a point is inside or outside of a zonotope. An example zonotope (shaded grey) is parameterized by its center c and generator vectors $g^{(1)}$ and $g^{(2)}$. A test point p is contained within the zonotope because the maximum absolute value of the coefficients on $g^{(1)}$ and $g^{(2)}$ is less than or equal to 1. Another test point q lies outside of the zonotope because the maximum absolute value of the coefficients on $g^{(1)}$ and $g^{(2)}$ is greater than 1.

the data. We consider a trial with a step or sit occurring later than the trial's end to be unsuccessful, and set $t_f = 1$ in this case.

3.2.5.2 Checking trajectories

Given an observed trajectory x_i from a Sit-to-Stand trial, we use a linear program to determine whether that trajectory lies inside the Stability Basin at each time step. Recall that each zonotope $Z^{(t)}$ representing the Stability Basin is parameterized by its center $c^{(t)} \in \mathbb{R}^4$ and generator matrix $G^{(t)} \in \mathbb{R}^{4 \times p}$. At each time t , we represent the state $x_i(t)$ as a linear combination of the generator matrix's columns, and use a linear program to find the representation that has the smallest maximum coefficient $\beta_{\max}^{(t)} \in \mathbb{R}$:

$$\begin{aligned}
 & \min_{\beta_{\max}^{(t)} \in \mathbb{R}, \beta^{(t)} \in \mathbb{R}^p} \beta_{\max}^{(t)} & (3.15) \\
 & \text{s.t. } G^{(t)}\beta^{(t)} = x_i(t) - c^{(t)}, \\
 & |\beta^{(t)}| \leq \beta_{\max}^{(t)}
 \end{aligned}$$

where the absolute value and inequalities are applied element-wise. If $\beta_{\max}^{(t)}$ is less than or equal to 1, the point $x_i(t)$ is inside the Stability Basin at that time step, depicted in Fig. 3.5.

3.2.5.3 Evaluation procedure

The Stability Basins are used to predict the success or failure of every trial within each Sit-to-Stand strategy (natural, Momentum-Transfer, and Quasi-Static). The accuracy of each Stability Basin is then tested by comparing its predictions to the experimentally observed outcomes of each Sit-to-Stand trial.

Given the state trajectory x_i observed for a Sit-to-Stand trial, the linear programs described in (3.15) are used to predict success/failure. If a state trajectory x_i remains inside the Stability Basin at each time step, the trial is predicted to be successful. If x_i exits the Stability Basin at any time step, the trial is predicted to fail. To ensure that the Stability Basins' predictions are made fairly, we make the following restrictions:

1. For cable pull trials, the prediction uses only the portion of the Stability Basin after the onset of perturbation.
2. For trials in which failure was observed, the prediction uses only the portion of the Stability Basin before failure occurred (i.e., $t \leq t_f$ as in Sec. 3.2.5.1).

We evaluate the Stability Basins formed using the LQR, FF+FB, and Input Bounds controllers, as well as the naive method in Sec. 3.2.4.3. We perform separate types of evaluation for successful trials versus unsuccessful trials. For successful trials, we employ a *leave-one-out* procedure to avoid training and testing on the same data. After forming the target set, we leave one successful trial out of the controller (LQR, FF+FB, Input Bounds) training data in Sec. 3.2.3.1. For the naive method, one successful trial is left out of the zonotope generation at each time step. We then construct a Stability Basin using each controller/method to test the trial that was left out. We repeat this procedure until all successful trials have been tested.

For unsuccessful trials, we compute a single Stability Basin for each controller model/method, formed using all observed successful trials. We use this single Stability Basin to predict the outcome of each unsuccessful trial within a given strategy. An outline of the validation procedure is provided in Fig. 3.6.

To be clear, we do not follow a leave-one-out procedure when generating the target set X_T from data. As displayed in Fig. 3.4, the target set generation is similar to taking the convex hull of a set of 4-dimensional points. If we follow a leave-one-out procedure when generating the target set, then the points on the boundary of the original set will likely be on the outside of the new set when they are left out. This would mean that many successful trials whose final states lie on the edge of the original target set necessarily exit the Stability Basins, and therefore are automatically misclassified when following this leave-one-out procedure. For these reasons, this procedure would serve to validate the robustness of the target set generation, rather than the

For each **subject**, **strategy** (Natural, Momentum Transfer, Quasi-Static), and **controller/method** (Input Bounds, FF+FB, LQR, Naive):

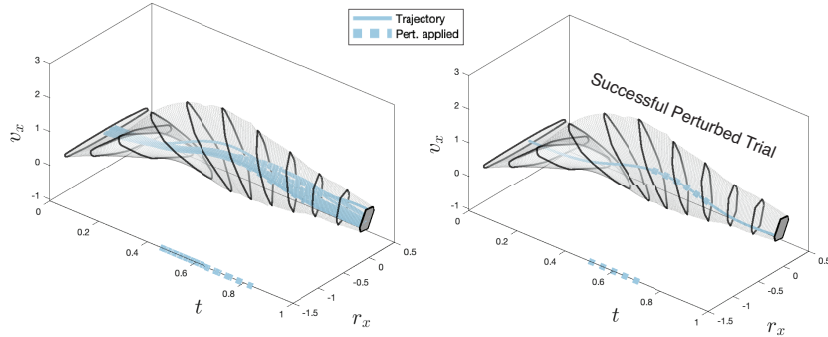
(a) Nominal, foot-shift, and successful perturbed trials

Step 1. Form standing set with all nominal, foot-shift, and successful perturbed trials (Fig. 4).

Step 2. Train controller (or naive method) with all nominal, foot-shift, and successful perturbed trials **except one** (Fig. 3).

Step 3. Form Stability Basin with standing set and controller, if applicable.

Step 4. Check that the left-out trial remains in the basin.



(b) Unsuccessful trials

Step 1. Form standing set with all nominal, foot-shift, and successful perturbed trials (Fig. 4).

Step 2. Train controller (or naive method) with **all** nominal, foot-shift, and successful perturbed trials (Fig. 3).

Step 3. Form Stability Basin with standing set and controller, if applicable.

Step 4. Check that the unsuccessful trial exits the basin before the step/sit occurs.

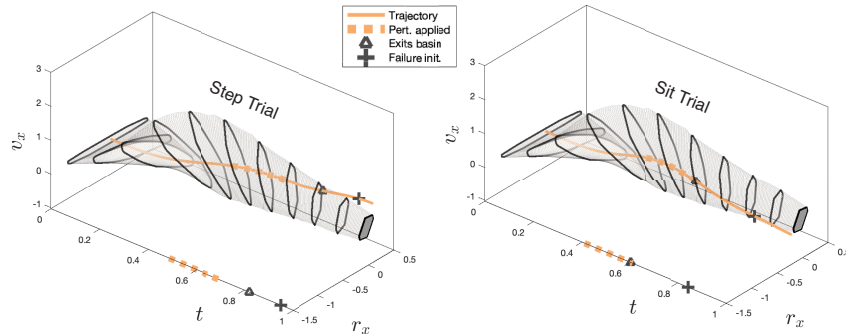


Figure 3.6: An outline of the Stability Basin validation with examples provided by the computed Stability Basin for subject ID 1’s natural strategy under the Input Bounds controller. The horizontal projection of the Stability Basin is represented as the region encapsulated by the light grey borders. The projection of the target set X_T is shown as the dark grey region on the right side of each plot. The state trajectories of all of subject ID 1’s successful natural strategy trials *except one* are used to construct the Stability Basin on which the left-out trial is tested, as shown in Fig. 3.6(a). Then all of subject ID 1’s successful natural strategy trials are used to construct the Stability Basin on which the unsuccessful trials are tested, as shown in Fig. 3.6(b). State trajectories of a step and a sit exit the basin before the onset of failure. As detailed in Sec. 3.2.1, we define step initiation as the time when the toes of either foot move more than 0.0762 m (3 inches) in the anterior-posterior direction, and sit initiation as the time at which both v_x and v_y become negative.




Figure 3.7: This figure shows example nominal unnormalized horizontal acceleration trajectories for each Sit-to-Stand strategy for subject ID 8.

Stability Basins themselves. However, collecting a larger number of samples should decrease the effect of missing data on the target set generation.

3.3 Results

The subjects performed Sit-to-Stand trajectories with similar maximum and minimum accelerations for each control strategy as the (different) subjects examined by Shia et al. [32]. Even though subjects were allowed to determine the speed with which each Sit-to-Stand was performed, each subject’s natural, Momentum-Transfer, and Quasi-Static Sit-to-Stand strategies showed distinct kinematic characteristics, depicted in Fig. 3.7. In particular, the maximum and minimum horizontal accelerations of the COM for the Momentum-Transfer $([-2.1, 2.2] \text{ m/s}^2)$ and Quasi-Static $([-0.8, 0.6] \text{ m/s}^2)$ strategies were significantly different from each other as well as from the natural strategy $([-1.3, 1.5] \text{ m/s}^2)$, depicted in Fig. 3.8, with p -values < 0.001 when comparing the successful trials within any two strategies. While the duration of natural and Momentum-Transfer trials were similar (average times of 1.25s and 1.13s), Quasi-Static trials lasted longer (average 2.23s), due in part to the low observed COM accelerations. The Stability Basin shape and volume are highly dependent on the kinematic and temporal characteristics corresponding to the strategy selected, as in previous work [32].

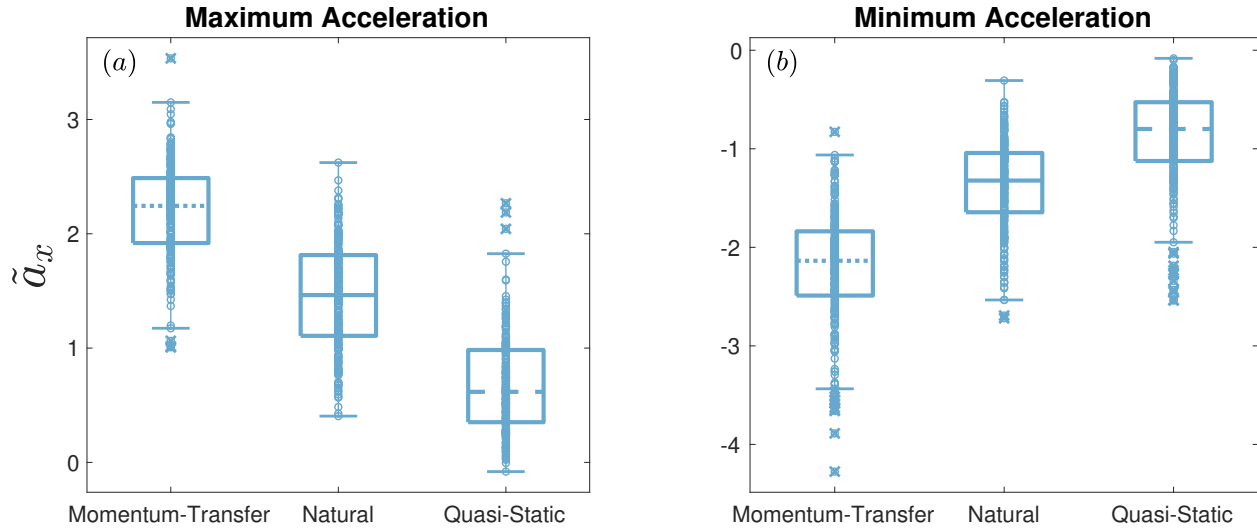


Figure 3.8: Maximum (Fig. 3.8(a)) and minimum (Fig. 3.8(b)) horizontal accelerations are significantly different between Sit-to-Stand strategies, with $p < 0.001$ when comparing the maximums or minimums of any two strategies under a two-sample t-test. Only successful trials are used in this analysis. For each strategy, each small circle represents data from a single Sit-to-Stand trial. The central mark within the box plot indicates the median horizontal acceleration. The bottom and top edges of the box indicate the the 25th and 75th percentiles, respectively. The whiskers extend to the most extreme horizontal accelerations not considered outliers, and the outliers are plotted individually as small circles with an ‘x’ through them.

When using the Momentum-Transfer strategy during the foot-shift trials, subjects were able to stand with their feet further forward (max. 0.199 ± 0.0699 m) than when using Quasi-Static (max. 0.0417 ± 0.0307 m).

Cable pull perturbations during Sit-to-Stand frequently induced a step or sit to avoid falling. Out of 591 cable pull trials, we observed failure in 198, or 33.5%. Specifically, we observed 42 steps and 29 sits for the natural strategy, 40 steps and 19 sits for Momentum-Transfer, and 39 steps and 29 sits for Quasi-Static. Fewer sits were observed for the Momentum-Transfer strategy than in both the natural and Quasi-Static strategies, likely due to the larger forward momentum at seatoff for the Momentum-Transfer strategy. The results are summarized in Tab. 3.1.

Table 3.1: Statistics collected from the perturbative sit-to-stand experiment are aggregated across subjects and summarized below. Note that the onset of failure (Step %STS and Sit %STS) often occurred after the defined trial end ($t = 1$), which is discussed in Sec. 3.2.5.1. Supplementary material that reports these statistics for individual subjects is available online.

Sit-to-Stand Strategy	Total	Cable Pull	Steps	Sits	Trial Time (s)	Pert. Onset (%STS)	Step (%STS)	Sit (%STS)
Natural	326	197	42	29	1.25 ± 0.23	0.50 ± 0.10	1.07 ± 0.21	0.96 ± 0.15
Momentum-Transfer	322	198	40	19	1.13 ± 0.21	0.47 ± 0.11	1.14 ± 0.33	0.99 ± 0.15
Quasi-Static	300	196	39	29	2.23 ± 0.89	0.51 ± 0.17	0.95 ± 0.21	0.85 ± 0.20

To assess the accuracy of each Stability Basin, we compared the predictions to the experimental observation for each trial according to the procedure described in Sec. 3.2.5.3. The aggregate prediction rates of the Stability Basins formed using the Input Bounds controller for each Sit-to-Stand strategy across subjects are given in Tab. 3.2. Across Sit-to-Stand strategies, nominal, foot-shift, and successful perturbed trials (Fig. 3.6(a)) remain inside the basin 95.73% of the time. Of the failure trials, 90.91% of steps are predicted prior to onset, while 92.21% of sits are predicted prior to onset (Fig. 3.6(b)). Using the visual trends in Tab. 3.3, we observed no consistent effect of perturbation force on the predictive power of the method.

Table 3.2: The accuracy of the Stability Basins generated using the Input Bounds controller model are reported below. This table reports a tally of the correct predictions for each Sit-to-Stand strategy and trial type across subjects.

Sit-to-Stand Strategy	Successful Trials	Step Trials	Sit Trials
Natural	248/255 – 97.25%	41/42 – 97.62%	27/29 – 93.10%
Momentum-Transfer	257/263 – 97.72%	33/40 – 82.50%	17/19 – 84.75%
Quasi-Static	213/232 – 91.81%	36/39 – 92.31%	27/29 – 93.10%
Combined	718/750 – 95.73%	110/121 – 90.91%	71/77 – 92.21%

Table 3.3: The accuracy of Stability Basins generated using the Input Bounds controller model are reported for cable pull trials at three different force levels of the applied perturbation. The results are aggregated across subjects and the three tested Sit-to-Stand control strategies.

Cable Pull Force Level	Successful Trials	Step Trials	Sit Trials
Low	170/176 – 96.59%	15/15 – 100%	4/5 – 80.00%
Medium	125/132 – 94.70%	34/38 – 89.47%	24/27 – 88.89%
High	80/85 – 94.12%	61/68 – 89.71%	43/45 – 95.56%

Overall, the Input Bounds Stability Basins’ failure predictions are over 90% accurate; only 17 out of 198 unsuccessful trials are incorrectly predicted to be successful. To understand these false successes, we measured the maximum Euclidean distance of each unsuccessful trajectory to its nearest successful trajectory. The false successes are 34.24% closer to their nearest successful trajectories than are the true failure trajectories. Additionally, the onset of failure with respect to the onset of perturbation for false successes averaged 9.47% earlier than in the true failures. Note that it is plausible that individuals switched control strategies for some subset of the trials (e.g., from Quasi-Static to natural), which could affect the accuracy of the strategy-specific Stability Basins. Because we only characterize trials that are sits or steps as failures, it is possible that an individual switched strategies to successfully stand but would not have reached standing with their original strategy.

Table 3.4: The accuracy of the Stability Basins formed using the Input Bounds controller model were compared to Stability Basins formed using other methods. The results of the comparison are reported below, and are aggregated across subjects and sit-to-stand strategies. Supplementary material that reports these results for individual subjects and sit-to-stand strategies is available online.

Basin Type	Successful Trials	Step/Sit Trials	False Successful Predictions	False Step/Sit Predictions
LQR	92/750 – 12.27%	197/198 – 99.49%	1/93 – 1.08%	658/855 – 76.96%
FF+FB	351/750 – 46.80%	195/198 – 98.48%	3/354 – 0.85%	399/594 – 67.17%
Input Bounds	718/750 – 95.73%	181/198 – 91.41%	17/735 – 2.31%	32/213 – 15.02%
naive	125/750 – 16.67%	195/198 – 98.48%	3/128 – 2.34%	625/820 – 76.21%

We compared the accuracy of Stability Basins computed using the Input Bounds controller to Stability Basins formed using other methods (Tab. 3.4). The Stability Basins formed using all three controllers as well as the naive method for Subject ID 1’s natural Sit-to-Stand strategy are shown in Fig. 3.9.

Stability Basins formed using the LQR controller proposed in previous work [32] correctly predict 12.27% of successful trials. Stability Basins formed using a traditional FF+FB controller correctly predict the outcome of 46.80% of successful trials. Finally, a naive method for estimating stability yields an accuracy of 16.67% for successful trials. Though each of the LQR, FF+FB, and naive methods yield nearly 100% prediction accuracy for trials where failure occurred, the low successful trial prediction rates mean that few trials remain within each Stability Basin, implying a severe underapproximation of the true stable region. We also give the false successful prediction and false failure prediction rates of each method in Tab. 3.4, and display these rates for each strategy in Fig. 3.10. Each of the LQR, FF+FB, and naive approaches for computing Stability Basins has a high false failure prediction rate, further showing that each underestimates the size of the true stable region.

Because the LQR, FF+FB, and naive methods all appear to underestimate the true stable region, we tested the accuracy of the Stability Basins generated by these methods when they are dilated by 5%. This was accomplished by multiplying each of their Stability Basins’ zonotope generators by 1.05, as in the target set generation. The classification rates for these methods do improve, though not substantially. As seen in Table 3.5, the successful trial classification rate rises around 3% for the LQR controller, around 6% for the FF+FB controller, and around 5% for the naive method. To determine what would happen if we extended this analysis further, we also evaluated the Stability Basins when the basins are dilated by 25%. This further improves the successful trial classification rate of the LQR, FF+FB, and naive methods. However, at this dilation level, the FF+FB has a similar unsuccessful trial classification rate as the Input Bounds controller (89.41% compared to 91.41%), but a 41% higher false step/sit prediction rate (56.83% compared to 15.02%). To be

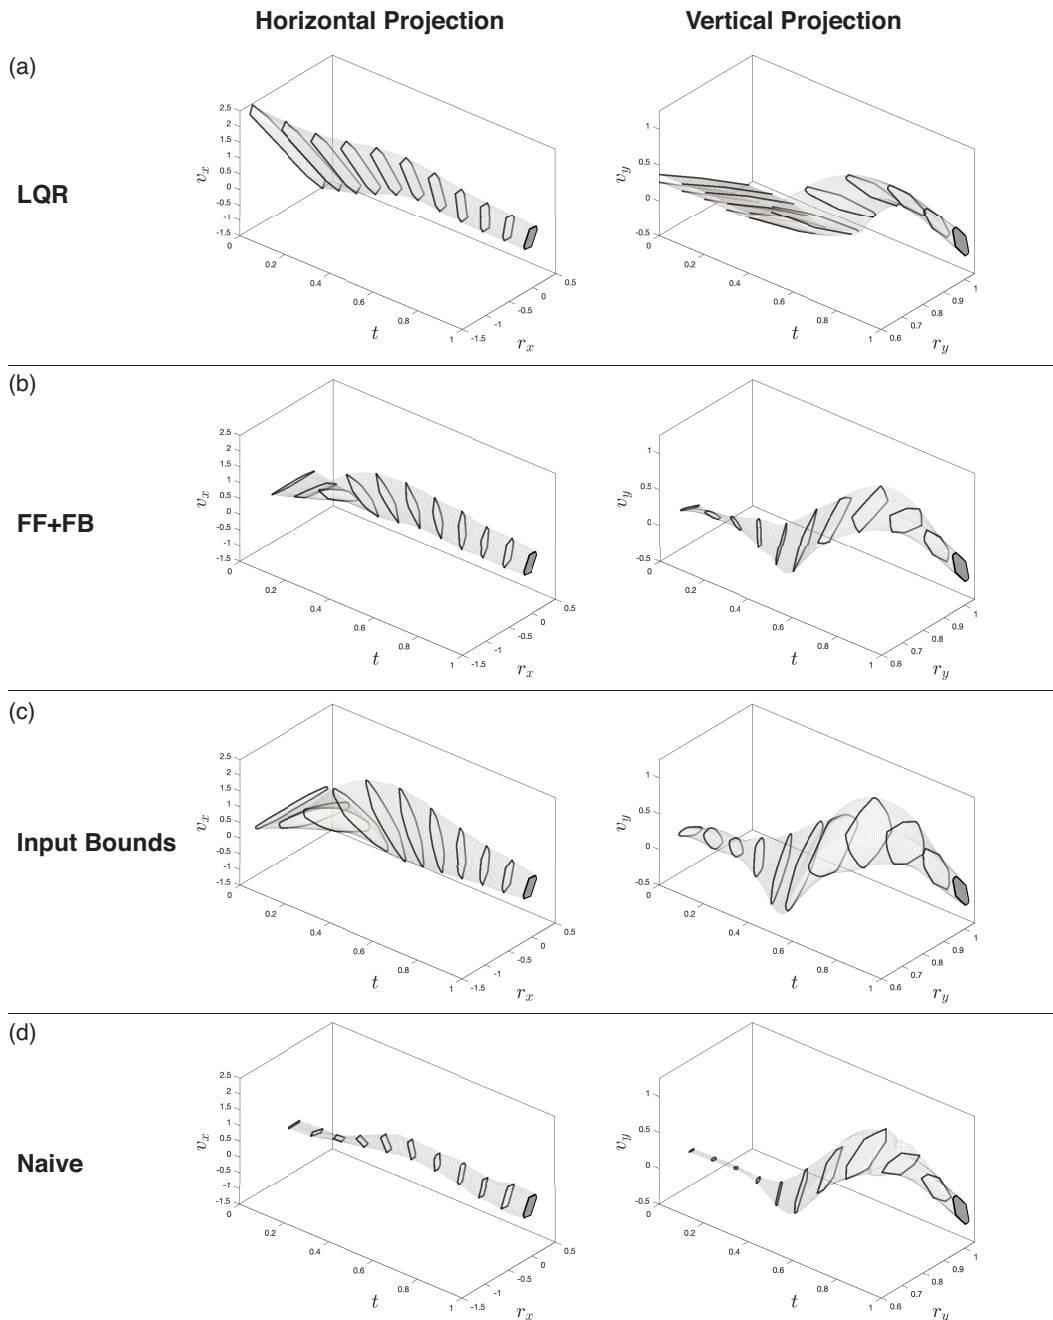


Figure 3.9: The Stability Basins formed for Subject ID 1's natural Sit-to-Stand strategy using the LQR controller model (Fig. 3.9(a)), FF+FB controller model (Fig. 3.9(b)), Input Bounds controller model (Fig. 3.9(c)), and the naive method (Fig. 3.9(d)) are shown. The horizontal and vertical projections of the Stability Basins are represented as the regions encapsulated by the light grey borders, where black outlines are used to emphasize the changes in shape of cross-sections over time. The projection of the target set X_T is shown as the dark grey region on the right side of each plot.




Figure 3.10: This figure compares the predictive accuracy of the SBs formed using the Input Bounds controller to three other methods. Results are presented for each Sit-to-Stand strategy, aggregated across subjects. The Input Bounds controller’s predictions for both successes and failures are near the ground truth, while the three other methods predict many more failures than were experimentally observed.

clear, we do not dilate the Stability Basins for our proposed controller model, the Input Bounds controller.

Table 3.5: The accuracy of the Stability Basins for each method are compared when dilating the LQR, FF+FB, and naive Stability Basins. The results are aggregated across subjects and STS strategies.

Basin Type	Successful Trials	Step/Sit Trials	False Successful Pred.	False Step/Sit Pred.
Input Bounds	718/750 – 95.73%	181/198 – 91.41%	17/735 – 2.31%	32/213 – 15.02%
LQR (no dilation)	92/750 – 12.27%	197/198 – 99.49%	1/93 – 1.08%	658/855 – 76.96%
FF+FB (no dilation)	351/750 – 46.80%	195/198 – 98.48%	3/354 – 0.85%	399/594 – 67.17%
naive (no dilation)	125/750 – 16.67%	195/198 – 98.48%	3/128 – 2.34%	625/820 – 76.21%
LQR (5% dilation)	113/750 – 15.07%	194/198 – 97.98%	4/117 – 3.42%	637/831 – 76.65%
FF+FB (5% dilation)	390/750 – 52.00%	192/198 – 96.97%	6/396 – 1.52%	360/552 – 65.22%
naive (5% dilation)	159/750 – 21.20%	195/198 – 98.48%	3/162 – 1.85%	591/786 – 75.19%
LQR (25% dilation)	160/750 – 21.33%	188/198 – 94.95%	10/170 – 5.88%	590/778 – 75.84%
FF+FB (25% dilation)	517/750 – 68.93%	177/198 – 89.39%	21/538 – 3.90%	233/410 – 56.83%
naive (25% dilation)	292/750 – 38.93%	189/198 – 95.45%	9/301 – 2.99%	458/647 – 70.79%

This implies that the Input Bounds controller still yields the most accurate estimate of stability. The Input Bounds controller maintains a high accuracy for successful trials, while minimizing incorrect step and sit predictions. It requires no dilation to achieve this performance. The higher false step/sit prediction rates for the LQR and FF+FB controllers despite dilation indicates that the shape of the Stability Basin constructed using the Input Bounds controller is a better approximation of the true stable region.

3.4 Discussion

In theory, stability can be directly characterized by perturbing a subject in an infinite number of ways and drawing a volume around the observed successful trials. Because this approach is

practically infeasible, several model-based metrics have been developed for assessing dynamic stability. Many, such as Floquet Multipliers and Lyapunov Exponents, are limited to analyzing small perturbations during periodic motion. Only a few metrics assess the largest perturbations that can be withstood during motion [56], or the stability of aperiodic motion. In this area, Fujimoto and Chou [70] constructed estimates of the stable region for Sit-to-Stand at a single time based on the extrapolated center of mass concept. We build on these methods by leveraging novel data-driven controller models and reachability analysis to estimate stability over the course of the Sit-to-Stand motion. By performing a perturbative experiment, we demonstrate that our Stability Basin-based approach accurately predicts stability across a diverse set of body morphologies, perturbations, and Sit-to-Stand control strategies.

Our data-driven controller models consist of a feedforward (“predictive”) and a linear feedback (“reactive”) component. Although there is evidence that humans employ feedforward and feedback terms for control [65], the exact form of these controllers remains unclear. Linear feedback applied about the COM has been sufficient to explain observed walking [71] and standing [72] behavior. However, we found the LQR controller model proposed by Shia et al. [32] as well as the classical FF+FB controller model had poor prediction rates for successful trials, implying that they both underestimate the true stable region of movement. In contrast, our Input Bounds controller model draws from a set of bounded inputs, plus a linear feedback term, that encapsulates inputs from observed successful trials, resulting in improved performance over models that return a single input.

The poor prediction rates of the LQR and FF+FB controller models for successful trials may be explained by considering the inputs that cause trajectories to exit the Stability Basin. A trajectory’s input not following a given controller model is not a sufficient condition for the trajectory to exit the Stability Basin generated for that controller model. However, it is a necessary condition: if a trajectory exits the Stability Basin, then the trajectory’s input must not have exactly followed the controller model at some point. Therefore, the accuracy of the Stability Basins strongly depends on how well each controller model fits the observed data. As can be seen in Fig. 3.3, the Input Bounds controller accounts for a wider range of inputs than the LQR and FF+FB controllers, and is able to encapsulate the variability present in the observed data. This in turn creates a Stability Basin which better matches the size and shape of the true stable region. This improvement in the accuracy of stability estimation is an important step towards clinical deployment. When applying the Stability Basin method to other Sit-to-Stand datasets in future work, one important question is how many successful trials are needed to appropriately estimate an individual’s Stability Basin. To investigate this question, we could compare Stability Basin accuracy between several cases where different samples of trials have been removed from the dataset described in this paper.

In addition to increased accuracy, the computationally-tractable, individualized, and time-

varying aspects of the Stability Basin approach make it well suited for widespread deployment and integration with existing clinical methods. The rapid and automated computation of Stability Basins from kinematics alone minimizes the temporal and financial investment in performing this analysis, especially if coupled with computer vision techniques for pose estimation from digital cameras [73]. This can enable longitudinal studies examining how stability changes over time, and whether a smaller Stability Basin is correlated with fall risk, to be performed both accurately and cheaply. Although we make use of perturbative data for computing Stability Basins here, one advantage of our proposed method is that it only relies on data from successful trials, and does not require a subject to be perturbed to failure. Thus, the Stability Basin method shows promise for characterizing the stability of subjects already at a high risk of falling or injury, such as elderly or rehabilitating subjects.

The TIPM and Input Bounds models are computed for each individual, enabling examination of a wide range of body morphologies and Sit-to-Stand control strategies. Stability throughout the duration of a task can be studied by measuring the size and location of cross-sections of the Stability Basin at particular times. For example, cross-sections of the Stability Basin that examine COM position at seat-off reveal the viable starting postures for each control strategy [32, 70]. In particular, Quasi-Static Stability Basins were previously shown to include COM positions near one's feet at seat-off, whereas Momentum-Transfer Stability Basins can extend much further. We propose that in longitudinal studies with periodic data collection, sudden changes in Stability Basin volume overall or in a portion of the movement can indicate how injury or age contribute to instability. By combining the Stability Basin approach with other clinical techniques, such as electromyography, clinicians can identify muscles associated with unstable portions of motion to target preventative and rehabilitative care. Such detailed analyses can provide deeper insight into the mechanisms underlying motor control. Lastly, although Stability Basins have been shown to better distinguish between more and less stable Sit-to-Stand control strategies than the Region of Stability, future work could compare the Stability Basins to another time-varying stability metric such as the extrapolated center of mass [59].

Our approach can also be used in real time to inform controller design for wearable robotic devices because it provides time-varying boundaries on the region inside which an individual is stable. Determining the state at which a robotic system is unable to recover from a perturbation is key to designing dynamically stable controllers [74, 75, 76]. If a prosthetic device or exoskeleton detects that its user has exited the Stability Basin and become unstable, it can adapt its control scheme to aid in recovery. One interesting area for future work is to characterize the timing of Stability Basin predictions so the wearable robot could have knowledge of *when* the step or sit would be initiated. Furthermore, just as current tuning of controllers for prosthetic devices and exoskeletons focuses on minimizing metabolic cost, [77, 78], wearable robotic controllers can be

optimized to increase stability by maximizing the volume of the Stability Basins for the combined robot-human system. Thus, the high time-varying accuracy and low investment in computation demonstrate that the Stability Basin approach has the potential to increase access of personalized preventative and rehabilitative mobility care and inform the design of robotic systems for stable movement.

CHAPTER 4

Specific Aim 2: Investigating Optimality in Perturbation Response During Sit-To-Stand

This work has been submitted as an extended abstract and presentation:

Danforth Shannon M., Holmes Patrick D., and Vasudevan Ram 2019. Inverse optimal control with Sit-To-Stand data. *Dynamic Walking*, Canmore, Canada.

and has been modified for this document.

4.1 Introduction

A common assumption is that humans and animals behave optimally during motion by minimizing some cost, such as metabolic energy [24]. Many claims about underlying optimality criteria involve (1) selecting a dynamic model for the task, (2) hand-picking a cost function to produce input trajectories for a model in an optimal control problem, and (3) simulating model states and comparing the results to human data [79, 80]. However, there is no way to confirm this hand-picked cost is the objective function minimized by the subject unless an exhaustive (and impossible) search of every possible cost function is investigated. Inverse optimal control (IOC) aims to extract a cost function from a set of observed state and input trajectories. IOC has been implemented in various forms for a range of human motions, including manipulation tasks [81], seated posture [82], and walking [83]. In this work, we apply IOC to the Sit-to-Stand (STS) motion, previously unstudied in IOC.

Our use of IOC with perturbed Sit-to-Stand data has two motivations. First, the structure of the underlying objective may be useful in studying human motor control. For example, the cost functions could be related to physical quantities or exhibit patterns across subjects and conditions (perturbed versus unperturbed). Second, the cost functions from IOC can help predict human response to perturbation. The recovered cost functions can be used in an optimal control problem

with a dynamic model to generate a control input trajectory for simulation. We can then compare the simulated model states to test data.

In this chapter, we implement an IOC formulation developed by [84] using data from the 11-subject STS experiment described in Ch. 2. We model the STS motion as an inverted pendulum model with a torque applied at the base, and approximate the cost functions using polynomials (degree 2, 4, and 6) of the model states and input. We compare the resulting objective functions between subjects and perturbed conditions, which suggest that optimality criteria is similar in both cases. Next, we use the objective functions to generate control inputs for the pendulum model and simulate left-out STS trials to investigate this method’s ability to produce accurate predictions of perturbation response during STS. We compare the simulation results to using an LQR controller model and find that the IOC formulation produces more accurate results.

4.2 Polynomial Inverse Optimal Control

This section describes the IOC formulation and implementation we used on the STS dataset, developed by [84]. We first provide a general overview of the IOC problem. Let our system dynamics be given by

$$\dot{x}(t) = f(x(t), u(t)) \tag{4.1}$$

with state and control constraints $x(t) \in X$, $u(t) \in U$, and $t \in [0, t]$. We also have a set of observed trajectories

$$(x(t; x_0), u(t; x_0)) \tag{4.2}$$

parameterized by $t \in [0, t]$ and initial condition $x_0 \in X$. The goal of IOC is to find a function (called a *lagrangian* in [84])

$$l : X \times U \rightarrow \mathbb{R} \tag{4.3}$$

such that all trajectories in $(x(t; x_0), u(t; x_0))$ are optimal trajectories for the direct optimal control problem:

$$\begin{aligned}
v(t, z) = \min_u \quad & \int_t^T l(x(s), u(s)) ds \\
\text{s.t.} \quad & \dot{x}(s) = f(x(s), u(s)), \\
& x(s) \in X, u(s) \in U, s \in [t, T], \\
& x(t) = z, \\
& x(T) \in X_T.
\end{aligned} \tag{4.4}$$

The function $v(t, z)$ is called a *value function*, which provides the value of the objective function at a solution.

The main challenge of IOC lies in uncovering the lagrangian function l . Many methods solve for the lagrangian by iteratively picking a candidate function and solving (4.4) until a local minimum is found [83, 85]. To exploit the problem structure, Pauwels et al. instead use explicit analytical optimality conditions for the direct optimal control problem, drawing from Hamilton-Jacobi-Bellman (HJB) theory [84]. The authors picked HJB theory for three main reasons: the HJB optimality equation provides a criterion to express optimality for the database trajectories in (4.2); the HJB condition can guide the choice of search space for candidate lagrangians; and relaxing the HJB optimality equation makes the problem amenable to considering polynomial lagrangians and value functions. To further understand these advantages, we first define a linear operator that acts on lagrangians and value functions:

$$\mathcal{L} : (l, v) \rightarrow l + \frac{\delta v}{\delta t} + \frac{\delta v^T}{\delta x} f. \tag{4.5}$$

We can use the \mathcal{L} operator to define a sufficient condition for optimality in the direct optimal control problem:

Proposition 4.2.1 (Pauwels et al.). *Suppose we have observed state and control trajectories $(x_0(s), u_0(s))$, with $s \in [t, T]$, such that*

$$\begin{aligned}
& \dot{x}_0(s) = f(x_0(s), u_0(s)), \\
& x_0(s) \in X, u_0(s) \in U, s \in [t, T], \\
& x_0(t) = z, \\
& x_0(T) \in X_T.
\end{aligned} \tag{4.6}$$

Suppose there exists a continuous lagrangian l_0 and continuously differentiable value function v_0

such that

$$\begin{aligned}\mathcal{L}(l_0, v_0)(s, x, u) &\geq 0, \forall (s, x, u) \in [t, T] \times X \times U, \\ v_0(T, x) &= 0, \forall x \in X_T, \\ \mathcal{L}(l_0, v_0)(s, x_0(s), u_0(s)) &= 0, \forall s \in [t, T].\end{aligned}\tag{4.7}$$

Then the state and control trajectories $(x_0(s), u_0(s))$, for $s \in [t, T]$, are optimal solutions of the direct optimal control problem.

Thus, the HJB optimality equation can incorporate the observed trajectories into an analytical criterion. The proof is provided in [84].

The IOC problem is based on a weakening of Proposition 4.2.1, therefore allowing for approximate solutions.

Proposition 4.2.2 (Pauwels et al.). *Let observed state and control trajectories $(x_0(s), u_0(s))$, with $s \in [t, T]$, satisfy (4.6). Suppose that there exist a real ϵ , continuous lagrangian l_0 , and continuously differentiable value function v_0 such that*

$$\begin{aligned}\mathcal{L}(l, v)(s, x, u) &\geq 0, \forall (s, x, u) \in [t, T] \times X \times U, \\ v_0(T, x) &= 0, \forall x \in X_T, \\ \int_t^T \mathcal{L}(l, v)(s, x_0(s), u_0(s)) ds &\leq \epsilon.\end{aligned}\tag{4.8}$$

Then, the input trajectory $(x_0(s), u_0(s))$, for $s \in [t, T]$, is ϵ -optimal for the direct optimal control problem.

With the certificate of ϵ -optimality, this proposition can be extended to multiple trajectories. The proof is provided in [84].

Lastly, the authors present a practical implementation by considering polynomial l and v candidates, polynomial dynamics, and assuming the observed state and input trajectories are indexed by a finite set, i.e. $\{(t_i, x_i, u_i)\}_{i=1, \dots, n}$:

$$\begin{aligned}\min_{l, v, \epsilon} \quad &\epsilon + \lambda \|l\|_1 \\ \text{s.t.} \quad &\mathcal{L}(l, v)(t, x, u) \geq 0, \forall (t, x, u) \in [0, T] \times X \times U, \\ &v(T, x) = 0, \forall x \in X_T, \\ &\frac{1}{n} \sum_{i=1}^n \mathcal{L}(l, v)(t_i, x_i, u_i) \leq \epsilon, \\ &\mathcal{A}(\mathcal{L}(l, v)) = 1.\end{aligned}\tag{4.9}$$

In this implementation, λ is a nonnegative regularization parameter and $\|\cdot\|_1$ is the ℓ_1 norm of a polynomial. The last constraint aims to mitigate trivial solutions (i.e., l and v with coefficients of zero) by applying any linear functional \mathcal{A} to $\mathcal{L}(l, v)$. This implementation is amenable to semidefinite programming. Polynomials are used because they can effectively represent general functions like l and v , and there is open-source software to solve semidefinite programs by approximating all functions as polynomials [86]. Because the polynomial lagrangian is a function of the model states, we can look at relative weights on each monomial to gain insights about which combinations of model states and inputs are prioritized during Sit-to-Stand.

4.3 Methods

4.3.1 Inverted Pendulum Model

In this chapter, we use data from a perturbative Sit-To-Stand experiment described in Ch. 2, Sec. 2.1. We fit individualized inverted pendulum models (IPMs) to each subject's data. This model consists of a point mass m located at a fixed distance L from the origin. We defined the origin as the initial position of the ball of the subject's foot. Individualized IPMs were constructed for each subject by setting m as the subject's mass and L as the average distance from the origin to the subject's center of mass (COM) across all trials. Fig. 4.1 depicts the pendulum model states and input.

The trials were aligned and nondimensionalized in time as described in Ch. 2, Sec. 2.1, with unitless time variable t . Therefore, each Sit-to-Stand trial starts at $t = 0$ and lies within $[0, 1]$. At each point t in a trial, the model quantities are:

- Pendulum angle from vertical, $\theta(t) \in \mathbb{R}$
- Pendulum angular velocity from vertical, $\dot{\theta}(t) \in \mathbb{R}$
- Torque input applied to pendulum base, $\tilde{u}(t) \in \mathbb{R}$.

Let $x : [0, 1] \rightarrow X \subset \mathbb{R}^2$ be a state trajectory, so that $x(t) \in X$ is the model's state at time $t \in [0, 1]$:

$$x(t) = \begin{bmatrix} \theta(t) \\ \dot{\theta}(t) \end{bmatrix} = \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix}. \quad (4.10)$$

The unitless time derivative of $x(t)$ at time t , denoted $\dot{x}(t)$, is written as:

$$\dot{x}(t) = \begin{bmatrix} x_2(t) \\ \frac{g}{L} \sin(x_1(t)) + \frac{1}{mL^2} \tilde{u}(t) \end{bmatrix}, \quad (4.11)$$

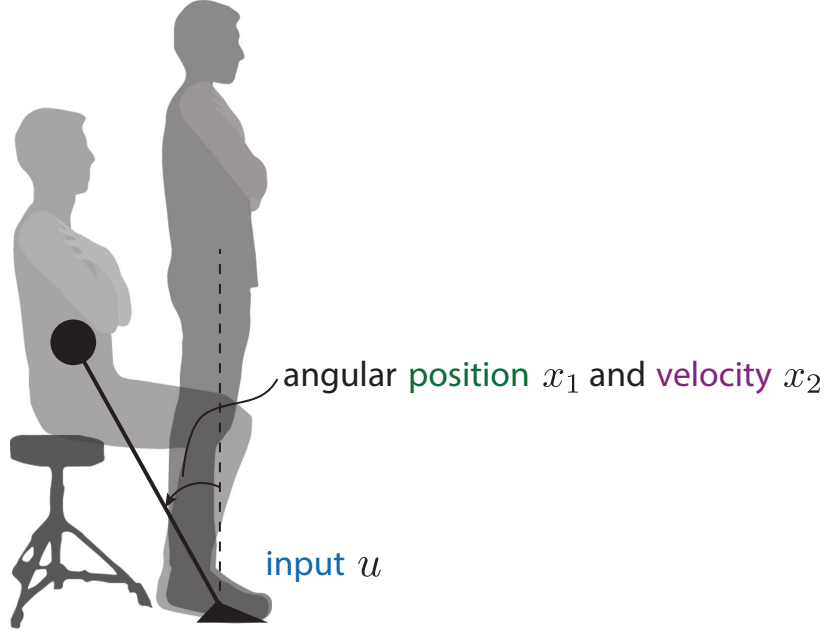


Figure 4.1: The inverted pendulum model used to model STS for inverse optimal control. The input u is a torque applied to the pendulum base.

where g is the gravitational acceleration 9.81 m/s^2 . Because the IOC formulation in (4.9) requires polynomial dynamics, we use a Taylor series expansion of $\sin(x_1(t))$ of degree 9, centered at zero. Our polynomial dynamics are given by:

$$\dot{x}(t) = \begin{bmatrix} x_2(t) \\ \frac{g}{L} \left(x_1(t) + \frac{1}{3!}x_1(t)^3 + \dots + \frac{1}{9!}x_1(t)^9 \right) + \frac{1}{mL^2}\tilde{u}(t) \end{bmatrix}. \quad (4.12)$$

For the remainder of this chapter, we define our control input as $u(t) = \frac{1}{mL^2}\tilde{u}(t)$.

4.3.2 Training the Lagrangians

We implement the IOC formulation in (4.9) using the polynomial system dynamics in (4.12), a polynomial value function v of degree 8, a polynomial lagrangian l of degree 2, 4, or 6, and a regularization term λ of 1×10^{-5} . The polynomial coefficients of l and v are constrained to $[-1, 1]$. We formulate the semidefinite program in the spotless toolbox [86] and solve it using MOSEK [87].

Recall that the STS experiment had three types of trials: nominal, forwards/backwards foot-shifted, and forwards/backwards cable pulls. The subjects also performed the experiment using three different strategies (natural, momentum transfer, and quasi-static); this work only uses data from the natural strategy. For this chapter, we refer to nominal and foot-shifted trials as

“unperturbed” and cable-pull trials as “perturbed.” We create perturbed and unperturbed subsets from the data to investigate the underlying cost functions. These training datasets form the $\{(t_i, x_i, u_i)\}_{i=1, \dots, n}$ in (4.9). When we train using perturbed data, we only include the portion of the trial before the perturbation started and after the perturbation ended. We exclude the portion during the perturbation because the external cable pulls affected the observed angles, angular velocities, and torques outside of the subject’s control, which would in turn affect the resulting lagrangian.

Because polynomials of higher degree can grow large if the states or inputs have magnitude greater than 1, we scale the angle, angular velocity, and control input to be within $[-1, 1]$. We first pick a scaling value for the angular position and velocity. Denote the scaled angular position and velocity as $x_{1,scaled}$ and $x_{2,scaled}$, respectively. We then find the scaled input associated with those trajectories, denoted as u_{scaled} . We iterate the process until $u_{scaled} \in [-1, 1]$, because the input typically has the largest magnitude. Because we deal exclusively with scaled values for the remainder of this chapter, we will refer to scaled states and inputs as $x(t)$ and $u(t)$, respectively.

As noted in Sec. 4.1, this study aimed to compare cost functions between subjects and conditions (perturbed and unperturbed). We are also interested in learning whether IOC can effectively simulate subjects’ response to perturbation during STS. If cost functions found from unperturbed data can effectively simulate perturbed responses, then the real-world applicability of this method could be increased, because perturbing subjects can be dangerous and time-consuming. We therefore use two separate training datasets, perturbed and unperturbed, for each subject. For each training subset and each subject, we compute one lagrangian formed from *all* data in the training subset (perturbed or unperturbed), which is used to analyze the characteristics of each subject’s objectives. We also use the unperturbed lagrangian to simulate each subject’s perturbed STS trials. For the perturbed lagrangian, to avoid training and testing on the same data, we also compute leave-one-out perturbed lagrangians formed from all data in the perturbed subset *except one*, which is left out for testing in simulation.

4.3.3 Testing the Lagrangians

Before detailing the tests conducted for each lagrangian found in Sec. 4.3.2, we first describe how to obtain simulation results given a lagrangian l found by IOC. Note that we have all of the pieces of the direct optimal control problem in (4.4) to solve for an optimal control input trajectory $u(t)$ for $t \in [0, 1]$. We define the target set X_T as

$$X_T = \{x(1) \mid x_1(1), x_2(1) \in [-1, 1]\}. \quad (4.13)$$

In other words, the target set X_T is equal to the scaled space X , and thus does not constrain solutions to end at a reasonable standing state. For simulating perturbed trials, we are interested in predicting the response to the cable pull. Let $t_{p,end} \in [0, 1]$ denote the time a perturbation ended for an example trial. Then, we define the initial condition for this trial simulation as $x(t_{p,end})$.

We solve (4.4) using a method for nonlinear optimal control synthesis that formulates a sequence of semidefinite programs with polynomial cost and dynamics [88]. Once the optimal control trajectory is recovered, we simulate the test trial from its initial condition using MATLAB's `ode45`. For any simulation evaluation, we calculate the root mean square error E between the simulated states and input ($\hat{x}(t)$ and $\hat{u}(t)$) and observed states and input ($x(t)$ and $u(t)$) for $t \in [t_{p,end}, 1]$. The error values are normalized by the total range of observed state or input values in that trial.

We first analyze the similarities in lagrangians between subjects and condition (perturbed and unperturbed). Treating the set of polynomial coefficients in the lagrangians as vectors, we normalize each to have a unit norm. We then use a Euclidean distance metric to measure the distance d between subjects' sets of coefficients. We compute the same distance metric between each subject's perturbed and unperturbed coefficients.

Next, to investigate whether unperturbed data can produce accurate predictions of perturbed STS trials, we use each subject's unperturbed lagrangians to simulate that subject's perturbed trials. We compute the normalized root mean square error E between simulated and ground truth data for each trial's states and input. Finally, for each subject, we use the leave-one-out perturbed lagrangians to simulate the left-out perturbed trial.

4.3.4 Comparison to LQR Controller

For comparison to the control inputs produced by IOC cost functions, we train LQR controllers using the perturbed data from each subject. LQR controllers have previously been used to model Sit-to-Stand motion [32, 23]. As in Ch. 3, we model input for this comparison as an LQR controller about a nominal Sit-to-Stand trajectory. We follow the same procedure for fitting the LQR controller as detailed in Ch. 3, Sec. 3.2.3.2, but for the simplified pendulum model in this chapter. We first form the average nominal angle and angular velocity trajectory, denoted as \bar{x} , for each of a subject's STS control strategies by taking the mean of the five nominal trials. We use this nominal trajectory to generate an open loop controller $u_{ol} : [0, 1] \rightarrow U \subset \mathbb{R}$ via inverse dynamics. In this chapter, our state weighting matrix is $Q \in \mathbb{R}^{2 \times 2}$ and input weighting matrix is $R \in \mathbb{R}^{1 \times 1}$. These two matrices are used to generate the feedback matrix $K(t) \in \mathbb{R}^{1 \times 2}$ [68, Chapter 16]. At a given time and state, the input from the LQR controller can then be written as:

$$u(t, x(t)) = u_{ol}(t) - K(t)(x(t) - \bar{x}(t)) \quad (4.14)$$

	Degree 2	Degree 4	Degree 6
Unperturbed	0.25	0.13	0.18
Perturbed	0.30	0.19	0.21

Table 4.1: Maximum Euclidean distance between any two subjects’ set of coefficients with unit norm, reported for unperturbed and perturbed lagrangians of degrees 2, 4, and 6. The maximum possible d between any two unit norm vectors is 2.

We choose the Q and R that produce a controller that most closely fits the observed perturbed data, solving the same optimization procedure described in Ch. 3. To avoid training and testing on the same data, we train and test the LQR controllers in a leave-one-out fashion.

4.4 Results

In this section, observed state and input trajectories are denoted as $x(t)$ and $u(t)$, respectively, and simulated state and input trajectories are denoted as $\hat{x}(t)$ and $\hat{u}(t)$, respectively, for $t \in [0, 1]$. The Euclidean distance d is computed between sets of coefficients treated as vectors. The normalized root mean square error, E , is computed between simulated and observed states and inputs.

Fig. 4.2 shows the coefficients of each subject’s degree-2 and -4 cost function plotted for each monomial of u , x_1 , and x_2 . The unperturbed lagrangian coefficients are plotted in Fig. 4.2(a)-(b), and the perturbed lagrangian coefficients are plotted in Fig. 4.2(c)-(d). Using visual trends in Fig. 4.2, we see that all 11 subjects have higher-magnitude weights on the same monomials; this trend holds for lagrangians of degree 6. For degree 2 lagrangians, the high weights were on x_1^2 (the squared angular position) and x_1x_2 (the multiplied angular position and velocity); all monomials that contained the input u had weights closer to zero. Recall that to compare the coefficients across subjects, we treated the set of polynomial coefficients in the lagrangians as vectors, then normalized each to have a unit norm. The maximum d between any two subjects’ sets of unit-normalized coefficients are reported for each data subset and lagrangian degree in Tab. 4.1. For comparison, the maximum possible d between any two unit norm vectors is 2.

Fig. 4.3 compares the mean coefficients for perturbed and unperturbed lagrangians (degree 2 and 4) across all trials. The maximum d between any one subject’s perturbed and unperturbed lagrangian coefficients was 0.29, 0.25, and 0.29 for degree 2, 4, and 6 lagrangians, respectively.

In the first case, we simulated responses to perturbed trials using each subject’s unperturbed lagrangian. The mean E -values between angular positions $\hat{x}_1(t)$ and $x_1(t)$, averaged across all subjects and trials, were 0.058, 0.12, and 0.16 for lagrangians of degree 2, 4, and 6, respectively. Mean E for angular velocity and input trajectories are summarized in Tab. 4.2 in the first row, and an example of simulation results is shown in Fig. 4.4(a).

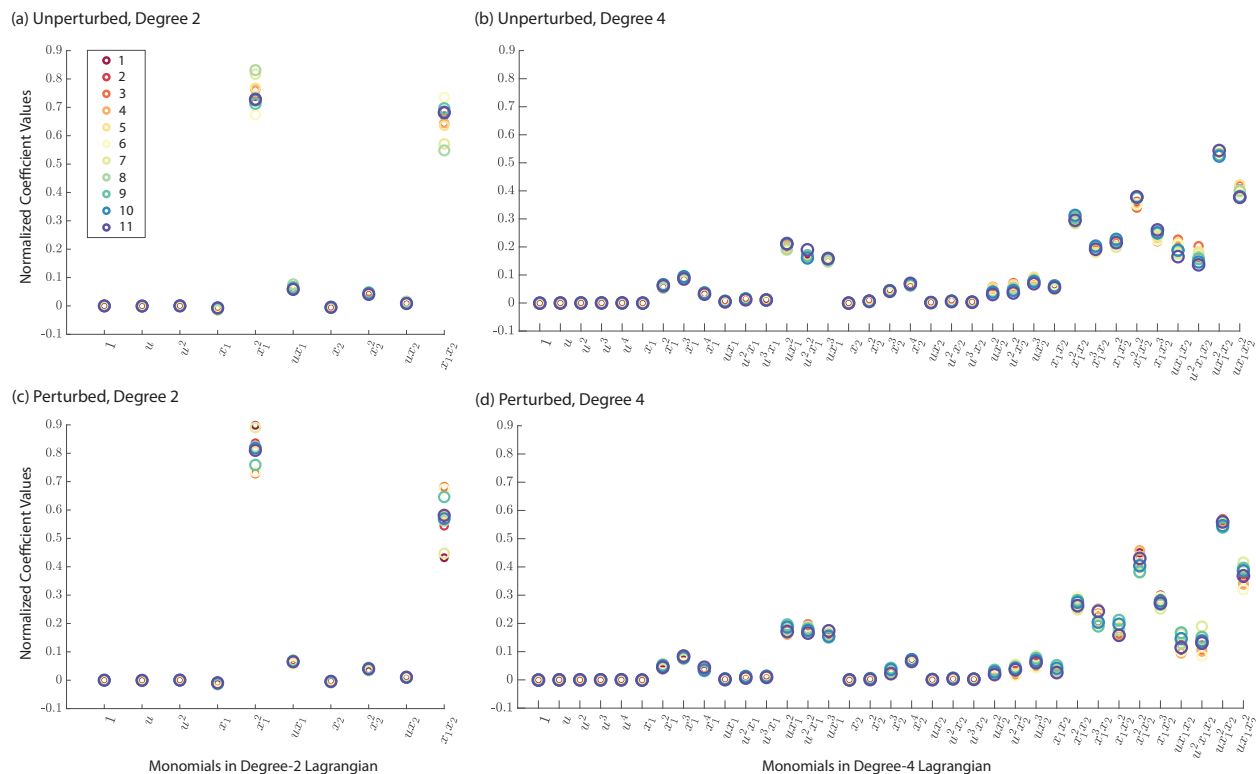


Figure 4.2: Weights on all monomials of the best fitting (a) unperturbed degree 2, (b) unperturbed degree 4, (c) perturbed degree 2, and (d) perturbed degree 4 lagrangian for each subject’s data, normalized to have unit norm.

In the second case, we simulated responses to perturbed trials using each subject’s leave-one-out lagrangians. The E -values between angular positions $\hat{x}_1(t)$ and $x_1(t)$, averaged across all subjects and trials, were 0.053, 0.16, and 0.15 for lagrangians of degree 2, 4, and 6, respectively. Normalized error for angular velocity and input is listed in Tab. 4.2 in the second row, and an example of simulation results is shown in Fig. 4.4(b).

We compared the IOC simulation results to an LQR controller that provided feedback about an average nominal trajectory. An example of simulation results using this controller is shown in Fig. 4.4(c). Results are summarized in Tab. 4.2 in the third row.

4.5 Discussion

The purpose of this study was to investigate optimality in STS using an inverse optimal control formulation. We divided the STS dataset into unperturbed and perturbed subsets, which produced lagrangians from unperturbed and perturbed data. We first analyzed similarities in these lagrangians between subjects and perturbed conditions. Next, we tested the lagrangians’ ability to simulate




Figure 4.3: Weights on all monomials of the best fitting lagrangian for unperturbed (light green) and perturbed (light orange), averaged across all subjects. The coefficient weights have been normalized and are shown here for (a) degree 2 and (b) degree 4 lagrangians.

Case	l degree	x_1 mean (stdev)	x_2 mean (stdev)	u mean (stdev)
Trained on unperturbed	2	0.058 (0.031)	0.15 (0.046)	0.13 (0.036)
	4	0.12 (0.062)	0.37 (0.16)	0.17 (0.051)
	6	0.16 (0.10)	0.47 (0.21)	0.21 (0.051)
Trained on perturbed (leave-one-out)	2	0.053 (0.029)	0.16 (0.045)	0.15 (0.035)
	4	0.14 (0.076)	0.39 (0.15)	0.19 (0.048)
	6	0.14 (0.070)	0.39 (0.13)	0.21 (0.039)
LQR controller	N/A	0.064 (0.054)	0.17 (0.091)	0.11 (0.033)

Table 4.2: Normalized root mean square error E for the training and testing cases described in Sec. 4.3. The mean (standard deviation) is computed across all subjects and trials. For each case, the row corresponding to the degree with the minimum error is highlighted in gray.

subjects’ responses to perturbations. We compared the simulation results to those produced by using an LQR controller about a nominal STS trajectory.

For all three degrees of lagrangians trained on unperturbed data, we see visual trends in the relative magnitude of coefficients on each monomial across all 11 subjects and between perturbed and unperturbed datasets. We used a Euclidean distance metric to quantify the differences in subjects’ lagrangian coefficients, but future work should use a statistical analysis to provide more insight into what differences between lagrangian coefficients is meaningful. The visual trends in coefficient values can still tell us about subject objectives during Sit-to-Stand, however. In degree 2 lagrangians, for example, subjects’ highest weights were on squared angular position and the multiplication of angular position and velocity, with near-zero weights on any monomial containing the input torque. This result suggests that subjects do not care about minimizing joint torque during Sit-to-Stand and instead may place priority on their COM position and velocity regardless of the required magnitude. Notably, the coefficients were close to zero for the monomial representing

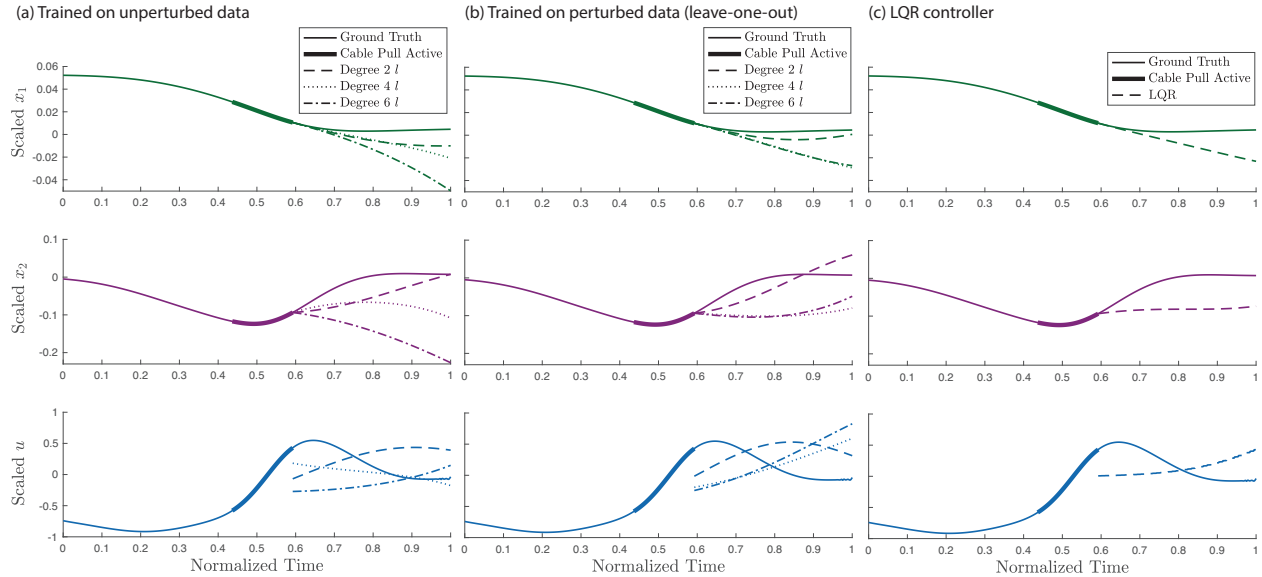


Figure 4.4: The unperturbed lagrangians, perturbed lagrangians, and LQR controller were tested on perturbed data. This figure shows an example simulation result for Subject ID 6 during a forwards cable-pull trial. The ground truth scaled angular positions (green), angular velocities (purple), and inputs (blue) are pictured in the solid line, with the bold segment indicating the time when the cable pull perturbation was active. (a) and (b) show the simulation results for lagrangians of degree 2, 4, and 6 for comparison. (c) shows the simulation result for the LQR controller.

mechanical work of the pendulum (ux_2), indicating that minimizing energy cost is not a priority during Sit-to-Stand.

Closely analyzing the variation of coefficients in each monomial between subjects and perturbed cases could also tell us more about small differences in behavior. For example, we see in Fig. 4.3(a) that subjects had, on average, a higher priority on the squared angular position and a lower priority on the combination of position and velocity for perturbed cases compared to unperturbed. This indicates that the objective to reach standing (i.e., minimizing the squared angle relative to vertical) is paramount after a perturbation, and velocity is not as much of a priority as it was during unperturbed motion. Future work should look at whether statistically significant differences in coefficient values produce differences in simulated behavior.

When lagrangians from unperturbed data were used to simulate perturbed response, the degree 2 lagrangian produced the lowest normalized root mean square error. Future work should perform a statistical analysis to investigate if the differences in prediction accuracy are statistically significant. Higher-degree lagrangians likely overfit to unperturbed data, thus performing poorly on the perturbed dataset. Visually (Fig. 4.4(a) provides an example), we observed that the simulated angular positions for degrees 4 and 6 often veered negative (i.e., the pendulum model leaned too far forward at the end of the task). Looking at the control inputs in Fig. 4.4(a), we see that the IOC

inputs are predominantly negative while the ground truth input is positive immediately after the perturbation. The subject applied a positive input to slow themselves down after a forwards cable pull, but the IOC formulation produced a negative or near-zero input at this time (depending on the lagrangian degree), causing the simulation to overshoot standing. The negative or near-zero input at this percentage of STS may be more typical during an unperturbed trial, where subjects do not have to apply as much compensatory torque to slow themselves down.

When the IOC training and testing was conducted with perturbed trials, the simulation error was around the same as that produced by the unperturbed lagrangians. As with the unperturbed lagrangians, degree 2 performed best, with degrees 4 and 6 consistently overshooting (see Fig. 4.4(b) for an example). The mean angular velocity error slightly improved compared to unperturbed simulations in degree 2 and 6, but was larger for degree 4. It is surprising that lagrangians trained on perturbed data did not perform better than unperturbed lagrangians in simulation. However, as shown in Fig. 4.3, the unperturbed and perturbed lagrangians were quite similar.

The LQR controller trained on perturbed data demonstrated similar mean simulation accuracy to the IOC cases, though slightly lower. A closer look at Tab. 4.2 reveals relatively large standard deviations in angles and angular velocity errors compared to the best-performing lagrangians. Because the LQR controller used feedback about an average nominal trajectory, it likely performed well in cases where subjects' perturbation-recovery behavior resembled nominal behavior, and poorly when subjects deviated from nominal behavior during recovery.

The coefficient values may also be influenced by the formulation requirements. The coefficient weights in Fig. 4.2(b) and (d) are larger for higher-degree monomials, but this phenomenon may have resulted from scaling the states and inputs. High-degree monomials of state and input values that have been scaled to $[0, 1]$ would produce relatively small values in the objective function, so perhaps the IOC formulation exploited the scaling when determining these coefficient values. Future studies could investigate the relative weights in a different implementation that does not require scaling.

The inverted pendulum model used in this chapter only included angular position and velocity of the COM and a torque applied at the pendulum base, which is most similar to ankle kinematics and torque in the Sit-to-Stand motion. However, torques at the hip and knee are also important in Sit-to-Stand [89, 90]. Including these joints in our model would enable us to investigate the relative weights in the lagrangian between hip, knee, and ankle angles and torques. It would also provide another opportunity to validate the model by comparing the results to previous biomechanical analyses. We restricted this formulation to the simple pendulum model because semidefinite optimization programs become computationally difficult or impossible to solve as the dimension increases. Advances in computational resources and optimization software may enable us to implement this IOC formulation on a three-link pendulum model in future work.

In this chapter, we applied inverse optimal control to STS data. The lagrangians of degree 2, trained on either perturbed or unperturbed data, were able to simulate responses to perturbations with lower mean error than an LQR controller model. The results of this study suggest that the subjects in our study prioritize COM position and velocity rather than torques during Sit-To-Stand, and this objective remains similar between unperturbed and perturbed standing. Future work should investigate the use of different IOC formulations and dynamic models to evaluate consistency across formulations; uncover physical meanings in cost functions; and further investigate optimality in responses to perturbation during STS.

CHAPTER 5

Specific Aim 3: Predicting Human Trip Response During Walking

This work is published [26]:

Danforth Shannon M., Liu Xinyi, Ward Martin J., Holmes Patrick D., and Vasudevan Ram 2022. Predicting Sagittal-Plane Swing Hip Kinematics in Response to Trips. *IEEE Robotics and Automation Letters*.

and has been modified for this document.

5.1 Introduction

Wearable lower-limb robotic devices, such as exoskeletons or prostheses, offer advantages for people with mobility limitations by supplying positive work to joints [91, 92]. Control options for these devices typically focus on common tasks (e.g., walking on level ground), because baseline kinematics for these tasks are well-known [93]. If an unexpected event such as a trip occurs during walking, controllers struggle to effectively assist with recovery because it is unclear how the human will respond. We therefore need accurate predictions of human kinematics during these unexpected events. As illustrated in Fig. 5.1, this chapter focuses on predicting sagittal-plane swing hip response during and after trips.

When tripped, humans typically exhibit one of three recovery strategies: Elevating, where the swing foot clears the obstacle and the heel strikes in front of the body; delayed lowering, where the swing foot initially elevates, then lowers behind the body with the toe contacting the ground first; and lowering, where the swing foot immediately lowers behind the body with the toe contacting the ground first [25]. These strategies have been studied most often in able-bodied subjects, but have also been observed in studies with lower-limb amputees [19]. Each of the trip-recovery strategies is accompanied by distinct swing hip kinematics, which are of interest if the controller




Figure 5.1: In this chapter, we develop and compare predictions of sagittal-plane swing hip kinematics after a trip occurs. The three hip states of interest (height, anterior-posterior position, and flexion-extension angle) are shown in blue to the left. On the right, the stick figures show an example of one subject’s trip-recovery response from our experiment. Note that the knee location is determined by the hip flexion-extension angle and thus is plotted in addition to the hip positions. The black line shows the ground truth hip and knee locations, and the pink line shows the predicted response from our Gaussian process regression model.

aims to avoid toe scuffs or place the swing foot in a location that ensures stability at the end of the step. Furthermore, if the controller is planning trip-recovery behavior in real-time as in [27], the predictions must be generated before the human responds. Human perturbation reaction time is influenced by many factors, including psychological and neurological limitations [94], but studies have shown stretch reflex responses around 40 ms and voluntary responses around 160 ms [95].

Human trip reactions are difficult to model due to the trip’s unpredictability and the variability in the human reaction. Trips can occur throughout the swing phase, and human recoveries can differ based on the perturbation onset and duration [25]. Eveld et al. developed an algorithm for predicting which of the three strategies a human will use after tripping [96], but even within a given recovery strategy, large variation in human kinematics exists. No models have been developed to predict trip-recovery kinematics after a specific strategy is identified. Thatte et al. used a Gaussian process regression (GPR) model conditioned on previously measured points to predict swing hip

height and angle in response to “hip drop” perturbations [27]. Because the predictions were not the main focus of their paper, the authors did not provide information on training the model or report the model’s prediction accuracy. The authors trained separate GPR models for each hip state, but training one multi-output model could increase accuracy by accounting for correlation between states.

Two other classes of models have been used for predictions of human kinematics. Neural networks are often used to estimate current kinematic variable values given measurements of different variables [97], but a recent study used neural networks to predict future lower-limb kinematics during different walking speeds [98]. The second common models are dynamic models that generate kinematic perturbation-recovery behavior. For example, a pendulum model with linear feedback has been shown to accurately estimate foot placement after a lateral perturbation [99]. Another study induced simulated trips in a sagittal-plane dynamic model and simulated its response using limit cycle optimization methods [100]. These studies, while highlighting the potential of various models for kinematic predictions, do not compare prediction accuracy between model classes or generate the strategy-specific trip-recovery predictions of interest.

In this chapter, we evaluate hip response predictions using data from an able-bodied trip experiment with 16 subjects, described in Sec. 5.2. Given EVELD et al.’s previous work in predicting strategy selection [96], this work assumes knowledge of the subject’s trip-recovery strategy and focuses on predicting strategy-specific kinematics. First, we implement multi-output GPR models with conditional predictions of trip-recovery swing hip response variables (Sec. 5.3). We compare the GPR models to nonlinear autoregressive exogenous (NARX) neural networks and pendulum models with a moving base (Sec. 5.4 and 5.5, respectively). Recognizing the need for both safety and low computation time, we evaluate all prediction methods for both simulation accuracy and online prediction time. We also evaluate each model’s ability to generalize predictions across subjects, which would increase the prediction method’s real-world applicability. Out of all simulation methods and training cases, the subject-specific GPR models exhibited the highest simulation accuracy and the lowest computation time. MATLAB code with examples for training each model is available at https://github.com/roahmlab/swing_hip_trip_prediction.

5.2 Trip-Recovery Experiment and Dataset

This chapter uses the 16-subject trip-recovery dataset described in Ch. 2, Sec. 2.2. Recall from Ch. 2, Sec. 2.2 that we analyzed these experimental data during the swing phase only, i.e. between toe-off and heel-strike events. In this chapter, we form a shifted normalized time by first aligning all of a subject’s swing phase trajectories by peak heel height. After shifting the trials so the peaks align, we find the minimum and maximum shifted time across all trials for one strategy, then use

the values to normalize time. We define the normalized swing time separately for each recovery strategy. In this chapter, we analyze 963 total trips across the 16 subjects: 433 elevating, 310 delayed lowering, and 220 lowering. The higher number of elevating trials is likely due to the targeting algorithm, which elicited more trips earlier in the swing phase.

Each recovery strategy was associated with distinct hip motion, shown for one subject in Fig. 5.2. The trip-recovery trials' difference from nominal level-ground walking varies by response variable and recovery strategy. For example, lowering trials most closely resemble nominal motion, while both elevating and delayed lowering trials exhibit more divergence. Within each strategy and response variable, the variation between trip-recovery trajectories also differs. The elevating and delayed lowering hip angles, for example, show a relatively large amount of variation between trials, while the lowering hip angles appear more uniform.

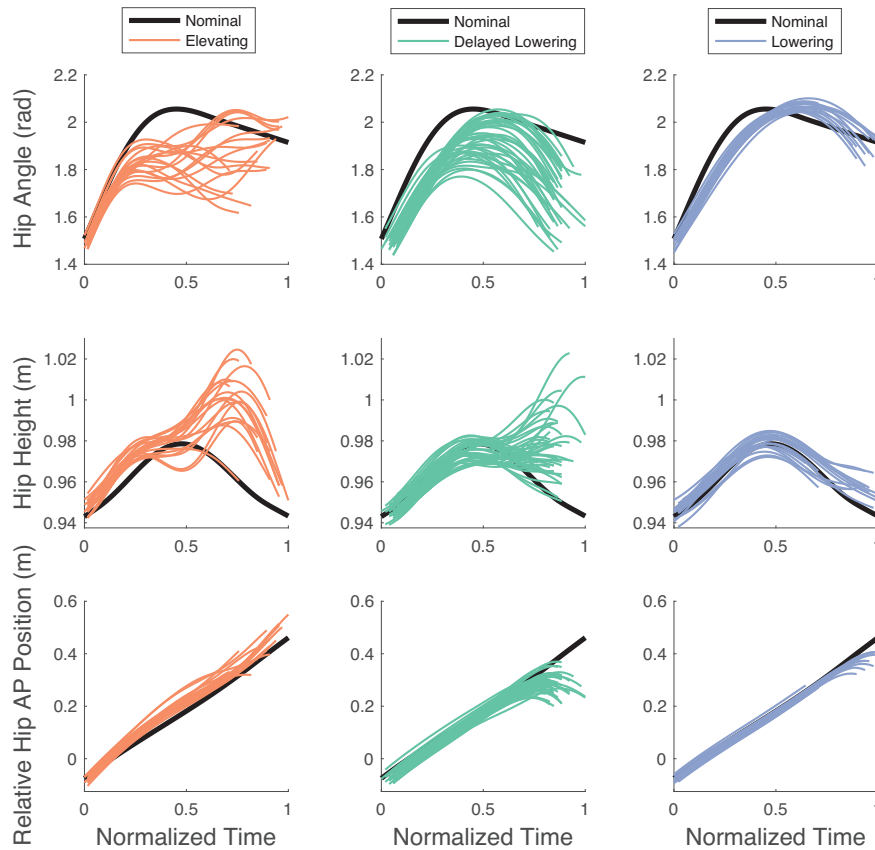


Figure 5.2: Hip flexion-extension angle, height, and anterior-posterior (AP) position relative to the stance foot, plotted for the elevating, delayed lowering, and lowering strategies in orange, green, and purple, respectively. The thick black lines show the average nominal hip states. These data are from one subject in our dataset.

5.3 Gaussian Process Regression Model

Given a set of noisy observations of a dependent variable, evaluated at certain values of an independent variable x , our prediction problem aims to find the best estimate of the dependent variable at a new value, x_* . Rather than assuming the underlying function is linear or another specific model, Gaussian process regression (GPR) assumes the observations can be represented as a sample from a multivariate Gaussian distribution. Note, we can always imagine a dataset with m observations to be a single point in \mathbb{R}^m drawn from an m -variate Gaussian distribution [101]. GPR models also provide uncertainty bounds on predictions.

For our predictions, we compute a separate multi-output GPR model for each recovery strategy, with outputs of hip flexion-extension angle, height, and AP position, described in Sec. 5.3.1. The GPR models are formed using a training dataset containing multiple trials for each recovery strategy. We then test trials that are excluded from the training dataset. For each test trial, we condition the trip-recovery prediction on its pre-trip values, described in Sec. 5.3.2. The training and testing procedure is described in Sec. 5.6.

5.3.1 Multi-Output GPR Distribution

Each observation in our training dataset (note, the dataset contains multiple trip trials) can be thought of as related to an underlying function \mathbf{f} through a Gaussian noise model:

$$\mathbf{y}_i = \mathbf{f}(x_i) + \boldsymbol{\epsilon}_i, \quad i = 1, \dots, n, \quad (5.1)$$

where $x_i \in [0, 1] \subset \mathbb{R}$ is the argument to \mathbf{f} and can be thought of as normalized time in our formulation, $\mathbf{y}_i \in \mathbb{R}^3$ is the output variable given by the hip flexion-extension angle, height, and AP position, and n is the number of points in the training dataset. The noise term $\boldsymbol{\epsilon}_i \in \mathbb{R}^3$ is a zero mean Gaussian distribution. We denote the covariance matrix of $\boldsymbol{\epsilon}_i$ as \mathbf{S} .

Let $k(x_i, x_j)$ denote the *covariance function*, which relates one input value x_i to another x_j . We use the rational quadratic kernel, a common covariance function for GPR models [102, Chapter 4], given by

$$k(x_i, x_j) = \left(1 + \frac{\|x_i - x_j\|_2}{2\alpha l^2} \right)^{-\alpha}, \quad (5.2)$$

where $\|\cdot\|_2$ denotes the ℓ^2 -norm, α is the scale mixture parameter, and l is the length scale of the kernel. Parameters α and l are estimated separately for each recovery strategy. Sec. 5.3.2 describes the procedure for estimating α and l .

We calculate the covariance matrix \mathbf{K} by evaluating (5.2) at all possible combinations of x_i for

$i \in \{1, \dots, n\}$:

$$\mathbf{K} = \begin{bmatrix} k(x_1, x_1) & k(x_1, x_2) & \cdots & k(x_1, x_n) \\ k(x_2, x_1) & k(x_2, x_2) & \cdots & k(x_2, x_n) \\ \vdots & \vdots & \ddots & \vdots \\ k(x_n, x_1) & k(x_n, x_2) & \cdots & k(x_n, x_n) \end{bmatrix}. \quad (5.3)$$

Recall that we aim to predict hip response values at a test input x_* . We therefore compute (5.2) between x_* and each observation x_i for $i = 1, \dots, n$:

$$\mathbf{K}_* = \left[k(x_*, x_1) \quad k(x_*, x_2) \quad \cdots \quad k(x_*, x_n) \right]^T, \quad (5.4)$$

and the covariance between x_* and x_* :

$$\mathbf{K}_{**} = k(x_*, x_*). \quad (5.5)$$

In a single-output GPR model, we would now have our data represented as Gaussian distribution with covariance matrix \mathbf{K} . The mean and variance of this distribution evaluated at test input x_* would form the best estimate of the output. A multi-output GPR model adds an additional challenge, because our response vector $\mathbf{y} = [\mathbf{y}_1^T, \dots, \mathbf{y}_n^T]^T$ is itself a Gaussian with zero mean and covariance matrix $\mathbf{C} \in \mathbb{R}^{3n \times 3n}$. Let the entry C_{ij}^{pq} be the covariance between the p^{th} and q^{th} response variable of vectors \mathbf{y}_i and \mathbf{y}_j , for each $p, q \in \{1, 2, 3\}$ and for each $i, j \in \{1, 2, \dots, n\}$. Recalling the three-dimensional noise term ϵ_i in (5.1), each entry of \mathbf{C} has two components: one corresponding to the cross-response covariance $\mathbf{C}_{\text{cross}}$, and one corresponding to the covariance of the noise term, $\mathbf{C}_{\text{noise}}$.

Calculating \mathbf{C} requires the estimation of significantly more hyper-parameters. Two assumptions are made to simplify the multi-output formulation [103]:

1. Within a cross-response pair indexed by (p, q) , the covariance function is formulated as a scalar-valued covariance matrix $k(x_i, x_j)$ scaled by a cross-response weight B_{pq} . The cross-response covariance is therefore be written as $[C_{\text{cross}}]_{ij}^{pq} = B_{pq}k(x_i, x_j)$. The B_{pq} entries construct a symmetric matrix \mathbf{B} , representing the pairwise covariance between different response variables.
2. The noise covariance matrix \mathbf{S} is assumed to be diagonal, because the covariances between responses have already been captured in the C_{cross} component. We therefore define $[C_{\text{noise}}]_{ij}^{pq} = \delta_{pq}\delta_{ij}S_{ij}$, where δ indicates a Kronecker delta.

With these assumptions, the covariance matrix \mathbf{C} is:

$$\mathbf{C} = \begin{bmatrix} B_{11}\mathbf{K} + S_{11}\mathbf{I}_n & B_{12}\mathbf{K} & B_{13}\mathbf{K} \\ B_{21}\mathbf{K} & B_{22}\mathbf{K} + S_{22}\mathbf{I}_n & B_{23}\mathbf{K} \\ B_{31}\mathbf{K} & B_{32}\mathbf{K} & B_{33}\mathbf{K} + S_{33}\mathbf{I}_n \end{bmatrix}, \quad (5.6)$$

where \mathbf{I}_n is the $n \times n$ identity matrix. We can also write \mathbf{C} as

$$\mathbf{C} = \mathbf{B} \otimes \mathbf{K} + \mathbf{S} \otimes \mathbf{I}_n, \quad (5.7)$$

where \otimes denotes the Kronecker product. Note that if \mathbf{S} , \mathbf{B} and \mathbf{K} are all positive definite, then so is \mathbf{C} , because both the Kronecker product and the summation on two positive definite matrices maintain positive definiteness.

Because \mathbf{B} is symmetric, constructing \mathbf{B} only requires the six elements from the upper triangle. Adding the three diagonal elements of \mathbf{S} , the multi-output formulation adds nine estimation parameters to the training process.

Our training dataset takes the form $\mathbf{y} = [\mathbf{y}_1^T, \mathbf{y}_2^T, \dots, \mathbf{y}_n^T]^T$ for n data points, and we aim to predict the response \mathbf{y}_* at a test input x_* . The probability of a certain prediction for \mathbf{y}_* follows a Gaussian distribution. Evaluating (5.4) and (5.5), we find the distribution's mean and covariance at x_* :

$$\begin{aligned} \boldsymbol{\mu}_* &= \mathbf{C}_*^T \mathbf{C}^{-1} \mathbf{y} \\ \boldsymbol{\Sigma}_* &= \mathbf{B} \otimes \mathbf{K}_{**} + \mathbf{S} - \mathbf{C}_*^T \mathbf{C}^{-1} \mathbf{C}_*, \end{aligned} \quad (5.8)$$

where $\mathbf{C}_* = \mathbf{B} \otimes \mathbf{K}_*$. The best estimate for \mathbf{y}_* and its associated uncertainty are given by $\boldsymbol{\mu}_*$ and $\boldsymbol{\Sigma}_*$. In our implementation, we evaluate the mean predictions and their covariance at m evenly-spaced values of x_* from 0 to 1, generating a $3m \times 1$ mean vector:

$$\boldsymbol{\mu} = [\boldsymbol{\mu}_1^T, \boldsymbol{\mu}_2^T, \dots, \boldsymbol{\mu}_m^T]^T. \quad (5.9)$$

The associated covariance matrix $\boldsymbol{\Sigma}$ has size $3m \times 3m$.

5.3.2 Conditional Prediction

The mean values in (5.9), computed from a training dataset, provide a prediction of hip responses at discretized points in normalized time. Given a test trip trial, we can use $\boldsymbol{\mu}$ to predict the hip response values after the trip onset. However, $\boldsymbol{\mu}$ does not account for the pre-trip points of a test trial, and therefore does not generate a continuous prediction from pre-trip points. To improve prediction accuracy, we condition the GPR prediction on pre-trip points in the test trial.

Our test trial begins and ends at normalized time values x_1 and x_2 , not necessarily spanning the entire normalized swing time range. To ensure the test trial is the same size as $\boldsymbol{\mu}$, we re-sample the test trial, and pad with NaN values if x_1 does not equal 0 and/or x_2 does not equal 1, to obtain a $3m \times 1$ test trajectory $\mathbf{y}_t = [\mathbf{y}_{t,1}^T, \mathbf{y}_{t,2}^T, \dots, \mathbf{y}_{t,m}^T]^T$.

For a given test trial, we denote the index of trip occurrence as $i_{trip} \in \{1, \dots, m\}$. We then divide \mathbf{y}_t into *completed* (pre-trip) and *future* (trip-recovery) trajectories:

$$\begin{aligned} \mathbf{y}_{t,c} &= [\mathbf{y}_{t,1}^T, \mathbf{y}_{t,2}^T, \dots, \mathbf{y}_{t,i_{trip}}^T]^T \\ \mathbf{y}_{t,f} &= [\mathbf{y}_{t,i_{trip}+1}^T, \mathbf{y}_{t,i_{trip}+2}^T, \dots, \mathbf{y}_{t,m}^T]^T, \end{aligned} \quad (5.10)$$

where subscript c indicates completed and f indicates future. We likewise partition the GPR mean and covariance into completed and future parts:

$$\boldsymbol{\mu} = \begin{bmatrix} \boldsymbol{\mu}_c \\ \boldsymbol{\mu}_f \end{bmatrix}, \quad \boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma}_{cc} & \boldsymbol{\Sigma}_{cf} \\ \boldsymbol{\Sigma}_{fc} & \boldsymbol{\Sigma}_{ff} \end{bmatrix}. \quad (5.11)$$

Then, our predicted future (trip-recovery) hip states and the prediction's uncertainty are given by the conditional mean and covariance:

$$\begin{aligned} \hat{\boldsymbol{\mu}} &= \boldsymbol{\mu}_f + \boldsymbol{\Sigma}_{fc} \boldsymbol{\Sigma}_{cc}^{-1} (\mathbf{y}_c - \boldsymbol{\mu}_c) \\ \hat{\boldsymbol{\Sigma}} &= \boldsymbol{\Sigma}_{ff} - \boldsymbol{\Sigma}_{fc} \boldsymbol{\Sigma}_{cc}^{-1} \boldsymbol{\Sigma}_{cf}. \end{aligned} \quad (5.12)$$

The prediction covariance can be used to form a confidence interval. As the number of conditional points grows, the inverse term in (5.12) takes longer to compute. In practice, to enable real-time predictions, we use at most 15 conditional normalized time points in our predictions.

For each recovery strategy, our multi-output GPR has 11 parameters to train: Parameters α and l from (5.2) and the nine parameters from the multi-output formulation. During training, we select parameters that minimize the root mean square error (RMSE) between the ground truth trip response and the conditioned GPR prediction across all response variables and training trials. One can solve this optimization problem using a nonlinear programming solver like MATLAB's `fmincon`. Although the training process may be time-consuming, the resulting GPR model can be saved and loaded for real-time conditional prediction.

5.4 NARX Model

For comparison to the GPR model in Sec. 5.3, we develop nonlinear autoregressive exogenous (NARX) models to predict hip response kinematics. Neural network methods, such as recurrent

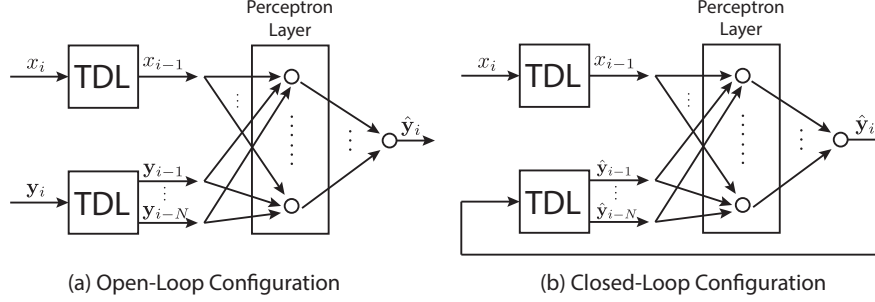


Figure 5.3: Diagram of the (a) open-loop configuration and (b) closed-loop configuration of the NARX model, where TDL stands for time delay layer.

neural networks (RNNs) and autoregressive networks (ARs), have been popular options for forecasting time series with unknown underlying functions [98]. It has been shown that any RNN network can be equivalent to a NARX neural network with a similar neuron transfer function [104]. Because NARX models exhibit computational advantages in training compared to RNNs, we use NARX models in this work.

We fit separate multi-output NARX models for each recovery strategy. We form the models using a training dataset and then test the models on trials that are excluded from training. Because the NARX model requires sequential time-series data for training, we concatenate the training trials into one time series. When combining the normalized time values from each trial into one vector, we add an offset to each subsequent trial, i.e., the first training trial’s normalized time points lie in $[0, 1]$, the second trial’s normalized time points lie in $[1, 2]$, and so on. For the hip response variables, we likewise concatenate the trajectories from each training trial into one time series (with no offset). Our concatenated dataset with n points has normalized time values $x_i \in \mathbb{R}, i = 1, \dots, n$. The $3 \times n$ response variable training data takes the form $\mathbf{y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n]$, with each $\mathbf{y}_i \in \mathbb{R}^3$.

The NARX model fits a function \mathbf{f} that relates each entry of the time series, \mathbf{y}_i , with previous values, $\mathbf{y}_{i,\text{prev}}$, and exogenous input, $x_{i,\text{ex}}$:

$$\mathbf{y}_i = \mathbf{f}(\mathbf{y}_{i,\text{prev}}, x_{i,\text{ex}}) + \boldsymbol{\varepsilon}_i, \quad (5.13)$$

where $\boldsymbol{\varepsilon}_i$ is an error term. At the i^{th} step, we specify the exogenous input to be the previous normalized time value, i.e. $x_{i,\text{ex}} = x_{i-1}$, and use N previous response values, i.e. $\mathbf{y}_{i,\text{prev}} = [\mathbf{y}_{i-N}, \mathbf{y}_{i-N+1}, \dots, \mathbf{y}_{i-1}]$.

We train the NARX networks using MATLAB’s `narxnet`. As shown in Fig. 5.3, there are two configurations, open-loop and closed-loop. When training each network, we use the open-loop configuration. At each index i of our training dataset, we input the delayed ground truth response $\mathbf{y}_{i,\text{prev}}$ and the exogenous input $x_{i,\text{ex}}$. The inputs first pass through a fully-connected layer

of perceptrons, also depicted in Fig. 5.3, which apply the sigmoid operation to yield intermediate outputs. For the j^{th} perceptron, the intermediate output \mathbf{o}^j at step i is given by

$$\mathbf{o}^j = \sigma\left((\mathbf{w}_x^j)^T \mathbf{x}_{i,\text{ex}} + (\mathbf{w}_y^j)^T \mathbf{y}_{i,\text{prev}} + \mathbf{b}^j\right), \quad (5.14)$$

where σ is the sigmoid function, $\mathbf{x}_{i,\text{ex}} = [x_{i,\text{ex}}, x_{i,\text{ex}}, x_{i,\text{ex}}]^T$, $\mathbf{b}^j \in \mathbb{R}^3$ is the bias for the j^{th} perceptron, and $\mathbf{w}_x^j \in \mathbb{R}^3$ and $\mathbf{w}_y^j \in \mathbb{R}^{3 \times N}$ are the weights for the j^{th} perceptron [105]. The intermediate outputs are then passed through the output linear perceptron, yielding a prediction $\hat{\mathbf{y}}_i$. The NARX network selects the bias and weights for each perceptron that minimize the RMSE between the true output \mathbf{y}_i and predicted output $\hat{\mathbf{y}}_i$, for $i \in \{1, \dots, n\}$.

For a test trial with m data points, we denote the normalized time and response as $x_t = [x_{t,1}, x_{t,2}, \dots, x_{t,m}]$ and $\mathbf{y}_t = [\mathbf{y}_{t,1}, \mathbf{y}_{t,2}, \dots, \mathbf{y}_{t,m}]$, respectively. Because we predict future values in the time series with the NARX model, the x_t vector is offset so that it occurs immediately after the concatenated training data. As in Sec. 5.3.2, the test trial has trip onset index i_{trip} and can be divided into completed and future partitions. We use the completed portion of the trial, $\mathbf{y}_{t,c} = [\mathbf{y}_{t,1}, \mathbf{y}_{t,2}, \dots, \mathbf{y}_{t,i_{\text{trip}}}]$, to inform post-trip predictions.

We use the closed-loop configuration, shown in Fig. 5.3(b), to generate predictions of the test trial for multiple future steps. The optimized bias and weight parameters from the training process are used in the closed-loop model. To predict a future post-trip value $\mathbf{y}_{t,i}$ where $i > i_{\text{trip}}$, the closed loop configuration uses previous predicted values, $\hat{\mathbf{y}}_{t,i,\text{prev}} = [\hat{\mathbf{y}}_{t,i-N}, \hat{\mathbf{y}}_{t,i-N+1}, \dots, \hat{\mathbf{y}}_{t,i-1}]$, as input. However, for the first N points after the trip occurs, there are not yet enough predicted values to form N previous inputs. Therefore, to predict the first post-trip point $\mathbf{y}_{t,i_{\text{trip}}+1}$, we use N previous points of $\mathbf{y}_{t,c}$; to predict $\mathbf{y}_{t,i_{\text{trip}}+2}$, we use $\hat{\mathbf{y}}_{t,i_{\text{trip}}+1}$ and $N - 1$ previous points of $\mathbf{y}_{t,c}$; and so on. We set the feedback delay N to be 15 normalized time steps, the same as the number of conditional points used for the GPR model.

5.5 Pendulum Model with Moving Base

To describe the dynamics of the swing hip variables, we use a pendulum model with a moving base, shown in Fig. 5.4. Although we focus on kinematics in this work, dynamic models have been used to effectively predict kinematic variables in previous research [99, 100]. A dynamic model could also prove useful if predictions of hip forces/torques are needed in a wearable robot's control scheme.

Our pendulum model consists of a mass m_1 at the hip joint location, which can translate in the horizontal and vertical direction, connected to a pendulum of length l and mass m_2 . The mass m_1 consists of the estimated torso, arms, and head, and the mass m_2 consists of the estimated swing

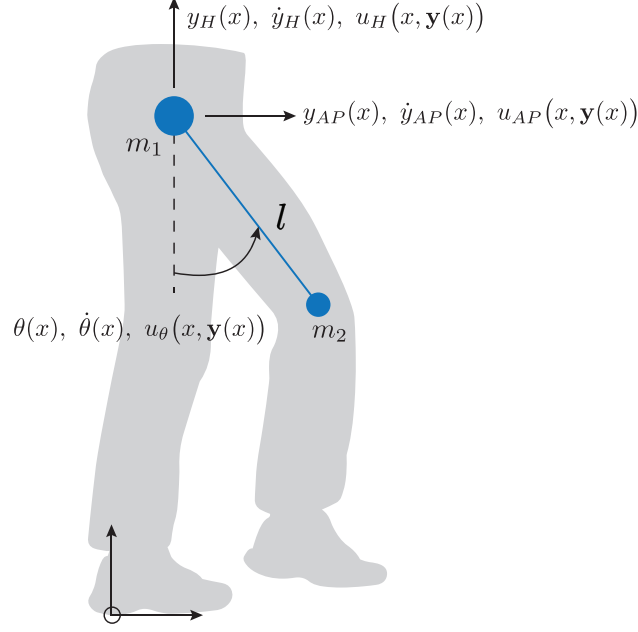


Figure 5.4: Our model for predicting hip kinematics has a floating base located at the hip joint and a pendulum representing the swing leg. The model has horizontal and vertical forces and a torque applied at the hip.

leg mass. The pendulum length l is the estimated swing thigh length. Masses and lengths were estimated using average distributions from anthropometric data [106, Chapter 4]. Because m_1 is considerably larger than m_2 , our model neglects the effect of the pendulum motion on the moving base. The model's origin is the location of the stance heel.

For this section, we define the normalized time as $x \in [0, 1] \subset \mathbb{R}$. At each point x , our model has anterior-posterior hip position $y_{AP}(x)$, hip height $y_H(x)$, and hip flexion-extension angle $\theta(x)$. The linear and angular velocities (with respect to normalized time) are denoted as $\dot{y}_{AP}(x)$, $\dot{y}_H(x)$, and $\dot{\theta}(x)$. Let $\mathbf{y}(x) = [y_{AP}(x) \ \dot{y}_{AP}(x) \ y_H(x) \ \dot{y}_H(x) \ \theta(x) \ \dot{\theta}(x)]^T$ denote the state vector. The model has three inputs: a horizontal force $u_{AP}(x, \mathbf{y}(x))$ acting on m_1 , a vertical force $u_H(x, \mathbf{y}(x))$ acting on m_1 , and a torque $u_\theta(x, \mathbf{y}(x))$ applied to the pendulum. The acceleration of the horizontal hip position is given by $\ddot{y}_{AP}(x) = \frac{1}{m_1}u_{AP}(x, \mathbf{y}(x))$ and the vertical hip acceleration is given by $\ddot{y}_H(x) = \frac{1}{m_1}u_H(x, \mathbf{y}(x)) - g$, where g is the gravitational acceleration $9.81m/s^2$. The hip flexion-extension angular acceleration is given as

$$\begin{aligned} \ddot{\theta}(x) = & -\frac{1}{l} \cos(\theta(x))\ddot{y}_{AP}(x) - \frac{1}{l} \sin(\theta(x))\ddot{y}_H(x) + \dots \\ & -\frac{g}{l} \sin(\theta(x)) + \frac{1}{m_2 l^2} u_\theta(x, \mathbf{y}(x)). \end{aligned} \quad (5.15)$$

We produce six first-order differential equations for the system using these accelerations.

Drawing from work that used pendulum models with linear feedback to accurately predict perturbation-response kinematics [99], our model’s control law consists of linear feedback gains to position and velocity error from each recovery strategy’s mean kinematics. One can solve for these feedback gains by using a nonlinear optimization procedure. As in Sec. 5.3.2, we use MATLAB’s `fmincon` to solve for feedback gains on position and velocity that minimize the RMSE between simulated and ground truth training trials.

5.6 Training and Testing the Models

Because we aim to test the significance of mean prediction errors across models, we select our training and testing datasets in accordance with the Combined 5×2 cv F Test framework [107]. This test has been shown to have higher power and lower false positives than previous statistical tests for comparing model performance. The training and testing process for each model is as follows:

1. Split each dataset into 50% testing, 50% training.
2. Fit the model using the training dataset.
3. For each test trial, use the trained model to simulate the trip-recovery hip states.
4. Compute the RMSE between predicted and ground truth trip-recovery trajectories. Normalize the RMSE for each response variable by the total range of response variable values in the test trial. Average these normalized RMSE (NRMSE) values across all three response variables.
5. Swap training and testing datasets; repeat steps 2)-4).
6. Repeat steps 1)-5) for five iterations.

For any two models, we compute an approximate f statistic to investigate whether the difference in NRMSE is statistically significant [107]. For the Combined 5×2 cv F Test, we conclude that the difference in mean NRMSE values between two models is statistically significant with 95% confidence if $f > 4.74$.

Individualized, data-driven models can produce accurate predictions of perturbation response, e.g. [23], but they necessitate collecting experimental (and often perturbative) data. Such experiments are infeasible for widespread use. We therefore conduct the Combined 5×2 cv F Test for both subject-specific and generalized cases. For the subject-specific case, we did not train and test models for strategies where subjects exhibited fewer than five trials (Subject 005 and 015’s

lowering trials and Subject 002’s elevating trials). For the generalized case, we first normalize the hip heights and AP positions by subject leg length, then train and test the models using all subjects’ data combined.

5.7 Results




Figure 5.5: Visualization of summary statistics for the (a) subject-specific and (b) generalized models. The central black line indicates the median NRMSE for each model; the bottom and top edges of each box are the 25th and 75th percentiles, respectively; the dashed whiskers extend to the limits of the data not considered outliers; and the outliers (determined using the interquartile range) are plotted individually as plus signs. The y-axis is a logarithmic scale due to the extreme magnitude of outliers.

Fig. 5.5(a) and (b) show box plots of the subject-specific and generalized model results, respectively. Due to the extreme values from the NARX model which often generated unstable predictions, we removed the outliers from all models for all following results and analyses. Table 5.1 presents the mean and standard deviation of NRMSE values for each strategy, model, and case across all subjects. In each strategy’s column, the minimum NRMSE is highlighted for both subject-specific (darker shade) and generalized (lighter shade) models. Example prediction results for each method, plotted for one subject’s delayed lowering trial, are shown in Fig. 5.6. Tab. 5.2 shows f statistics computed between each model within the subject-specific and generalized cases and their corresponding p -values.

The mean computation time to generate conditional GPR predictions, NARX predictions, and dynamic model predictions was 2.67 ms, 8.81 ms, and 25.1 ms, respectively, on a computer with

Table 5.1: Average NRMSE for all three prediction methods, with lowest subject-specific and generalized errors highlighted for each strategy.

Prediction Method	Case	Elevating mean (stdev)	Delayed Lowering mean (stdev)	Lowering mean (stdev)
GPR Model	Subject-Specific	0.110 (0.0333)	0.119 (0.0522)	0.0567 (0.0266)
	Generalized	0.151 (0.0500)	0.160 (0.0630)	0.0879 (0.0362)
NARX Model	Subject-Specific	0.660 (0.482)	0.500 (0.360)	0.320 (0.333)
	Generalized	0.262 (0.119)	0.321 (0.171)	0.141 (0.124)
Pendulum Model	Subject-Specific	0.241 (0.0618)	0.285 (0.0792)	0.277 (0.0782)
	Generalized	0.254 (0.108)	0.262 (0.135)	0.264 (0.147)

Table 5.2: Approximate f statistic between models, with * indicating statistical significance, and corresponding p -values.

Models	Case	$f(p)$ Elev.	$f(p)$ Del. Low.	$f(p)$ Lower.
GPR and NARX	Subj.-Spec.	63.7* (<0.001)	82.6* (<0.001)	29.1* (<0.001)
	Gen.	10.4* (0.009)	7.98* (0.02)	4.08 (0.07)
GPR and Pendulum	Subj.-Spec.	275* (<0.001)	83.5* (<0.001)	531* (<0.001)
	Gen.	17.4* (0.003)	82.6* (<0.001)	167* (<0.001)
NARX and Pendulum	Subj.-Spec.	45.6* (<0.001)	18.4* (0.002)	1.43 (0.4)
	Gen.	2.70 (0.1)	2.04 (0.2)	15.0* (0.004)

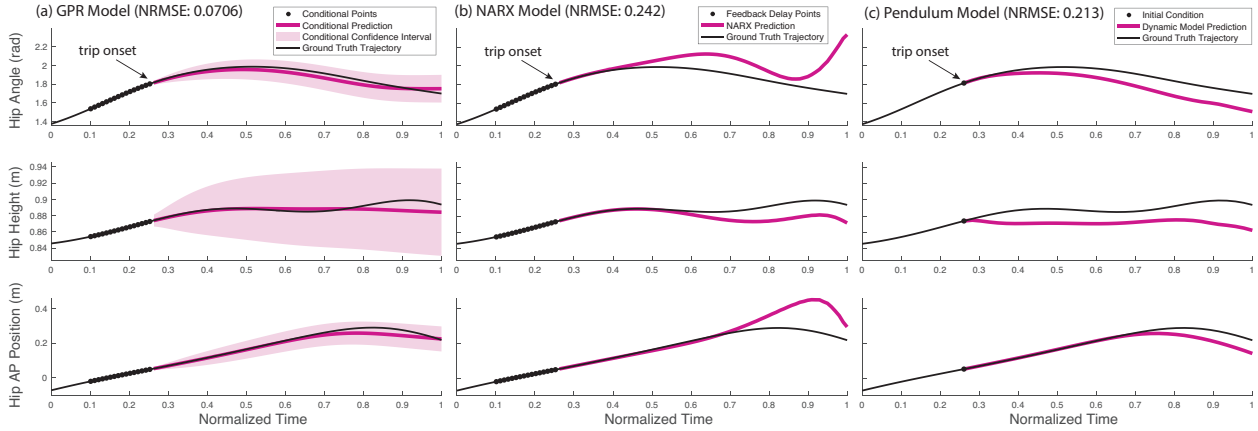


Figure 5.6: An example of simulation results for one delayed lowering trial using the subject-specific (a) GPR, (b) NARX, and (c) pendulum models. The GPR model prediction in (a) includes a 95% confidence interval. The entire ground truth trial is shown in the black line, with the trip onset indicated by the arrow. The NRMSE for this trial, averaged across each response variable, is provided for each prediction method. The black dots represent the information provided to each model: 15 conditional points in (a); 15 feedback delay points in (b), and one initial condition for simulation in (c).

a 3.5 GHz Intel Core i7 processor. For the GPR models, we checked whether the 95% confidence bounds, shown in Fig. 5.6(a), encompassed each ground truth trajectory. The percentages of trials that remained inside the confidence bounds, reported for each case (generalized and subject-specific), strategy, and response variable, are listed in Tab. 5.3.

Table 5.3: Percent of ground truth trials that remain inside the 95% confidence bounds for the GPR model.

Case	Strategy	Angle	Height	AP Position
Subject-Specific	Elevating	73.9	100.0	60.8
	Delayed Lowering	71.4	100.0	90.4
	Lowering	92.7	100.0	98.3
	Total	77.3	100.0	75.1
Generalized	Elevating	51.2	96.7	75.2
	Delayed Lowering	47.4	98.8	52.6
	Lowering	79.9	100.0	84.5
	Total	56.5	98.2	67.9

5.8 Discussion and Conclusion

Predicting human trip-recovery kinematics can help wearable lower-limb robots choose appropriate coordinating behavior. Previous studies that investigate predictive models for human kinematics report results for only one class of model [99]; use simulated trip data instead of human subject experimental data [100]; or use unperturbed data, where less variation is likely present [98]. This chapter investigates three modeling methods to predict swing hip kinematics in experimental trip data from 16 subjects. We restricted predictions to the sagittal plane, but we expect the models could be extended to include other hip motions (abduction-adduction, interior exterior rotation) in future work.

The subject-specific GPR produced the lowest mean error for each strategy out of all modeling methods. Across all strategies, we found that its NRMSE was statistically significant from that of the other two subject-specific models. The generalized GPR models produced the lowest mean NRMSE across generalized models, which was statistically significant from the NARX and pendulum models in all cases but one. Furthermore, the GPR model's lower computation time is useful for real-time trip-recovery control in wearable robots.

By analyzing visual trends in Tab. 5.1, we see that predictions had the highest accuracy in lowering strategies compared to elevating and delayed lowering. Looking at Fig. 5.2, we see that the lowering strategy trials for Subject ID 1 have less variation across all three response variables, which could explain the higher prediction accuracy. However, a statistical analysis should be conducted to confirm that the differences in error are statistically significant and to quantify the variation in hip kinematic data across all subjects and strategies.

The GPR model's confidence bounds offer another advantage over the other prediction methods. Most ground truth trajectories remained within the 95% confidence bounds, though more often in the subject-specific case. By visual trends in Tab. 5.3, we see that the percentages are lower for certain cases, strategies, and response variables (the lowest being the delayed lowering angle for generalized models). Future work will investigate the discrepancies between response variables and strategies.

The subject-specific NARX models frequently generated unstable predictions, especially for trips that occurred earlier in the swing phase. We speculate that the considerably large error in elevating and delayed lowering trials is a result of overfitting the network to relatively small datasets. Notably, the generalized NARX models' mean NRMSE decreased by approximately 50% from the subject-specific, likely due to larger training datasets.

The dynamic pendulum model produced higher mean NRMSE than the GPR model, with differences in NRMSE statistically significant from GPR in both subject-specific and generalized cases. We found that of the three response variables, the hip height contributed most to this high mean

NRMSE. Because the dynamic model applied feedback to mean kinematics, the high error could be explained by the relatively large variation in subjects' trip-recovery hip height trajectories. Another possible cause of higher prediction error is the dynamic model's use of average distributions for the mass and length of the torso and thigh. Because trip-recovery strategy may be influenced by anthropometry and our subjects' true mass and length distributions were likely not average, future work should examine the influence of these parameter values on prediction accuracy. One way to investigate this effect is to modify m_1 , m_2 , and l and see if prediction accuracy is statistically significant for different parameter values.

In this chapter, we compared relative prediction accuracy three types of predictive models. However, when using these predictions in wearable robot control, we are also interested in the required bounds on absolute prediction error values. In other words, what is the maximum allowable error in swing hip predictions that can still ensure successful trip recovery, user comfort, or safety? Peng et al. have recently explored human perception of exoskeleton actuation timing by quantifying the Just Noticeable Difference (JND): the actuation timing value that is actually perceivable by the subject [108]. A similar experiment could look at JND values for prostheses or exoskeletons during trip recovery that use a range of predicted hip states with varying accuracy.

Our results show the multi-output GPR model's promise for producing accurate predictions of kinematic responses to trips with low computation time and quantified prediction uncertainty. The generalized GPR models exhibited similar simulation accuracy to the subject-specific case, increasing the modeling method's real-time applicability. Although our experiment included only able-bodied subjects, individuals with lower-limb amputation have exhibited similar trip-recovery strategies in experiments [19]. In future work, we will investigate the similarities in swing hip kinematics during trip recoveries for those with and without lower-limb amputation. Furthermore, swing hip behavior would likely change based on how the prosthesis behaves during a trip. For example, hip kinematics may differ depending on whether a passive or robotic lower-limb prosthesis is in use. One way to address this challenge is to use wearable sensors, such as inertial measurement units, to collect swing hip kinematics during trips and continually update model predictions based on observed reactions.

CHAPTER 6

Specific Aim 4: Planning Trip-Recovery Trajectories for Robotic Lower-Limb Prostheses

A preliminary version of this work is available as a preprint [109]:

Danforth Shannon M., Holmes Patrick D., and Vasudevan Ram 2020. Trip Recovery in Lower-Limb Prostheses using Reachable Sets of Predicted Human Motion. *arXiv*.

and has been modified for this document.

6.1 Introduction

In 2005, an estimated 623,000 people were living with major lower-limb loss in the United States, with the number predicted to double by 2050 [5]. Most of these individuals use passive lower-limb prostheses, with instances of falling comparable to older adults over 85 [6, 7]. Powered lower-limb prostheses have the potential to lower metabolic cost and improve gait symmetry [110], and recent developments in their control allow them to assist users while walking up slopes, ascending/descending stairs, and walking backwards [111, 112, 113]. However, despite having the potential to assist users during stumbles, most powered prostheses do not include trip-recovery capabilities. Equipping powered prostheses with trip-recovery features could lead to fewer falls and increased balance confidence in lower-limb prosthesis users.

Before discussing why stumble-recovery control is difficult in prostheses, it is useful to understand their general control framework. During locomotion, robotic lower-limb prostheses use various sensor modalities, detailed by [114], in conjunction with a high-level controller to estimate the human's intent. Common high-level controllers include activity mode detection [93, 115], direct volitional control [116, 117, 118, 119], or a combination of both [113]. A mid-level controller then converts this estimated intent into a desired trajectory for a low-level controller to track [114].

Two challenges contribute to the difficulty of designing effective stumble-recovery controllers within this hierarchical control framework. The first is reliably detecting the stumble itself. In




Figure 6.1: The framework developed in this chapter plans trajectories for a knee prosthesis during and after a stumble that ensure successful recoveries for a set of predicted human behavior. The kinematic model used in this work contains human and prosthesis subsystems, constructed from experimental data. This figure shows reachable sets for the model, which depend on both human and prosthesis behavior, after a toe catch perturbation. The gray shapes at the hip and knee represent the range of joint locations from the subject’s predicted swing hip response. The light orange and light purple shapes at the foot represent the set of prosthesis heel and toe locations for a range of parameterized prosthesis behavior. Our framework selects one heel and toe parameter in an optimization program, which produce the darker orange and purple subsets. The optimization program aims to (1) avoid toe scuffs and (2) ensure the heel strike location lands as close as possible to a *target set* representing desired foot contact location, shown by the green lines (in the callout), for each time step where ground contact is possible.

stumble experiments, humans have exhibited responses with latencies under 100 ms [37, 120]. Researchers have shown that accelerometers mounted on the prosthesis can detect stumbles within this short time span, although false positives are common [121, 122]. The second challenge lies in choosing a recovery behavior for the prosthesis once the stumble has been detected. Difficulties in this area arise from the presence of uncertainty in estimates of human intent and the inability of existing mid-level controllers to generate trajectories for stumble recovery in real-time.

As a human trips and tries to recover, any uncertainty in the controller’s estimate of the human intent or current state could increase fall risk by producing unexpected torques [114]. This uncertainty is present in controllers that rely on training data and are at risk of misclassifying intent [123, 124, 125] as well as controllers that use an estimated human state to prescribe joint torques. Two widely-used prosthetic control methods that fall into the latter category are virtual constraint control [126, 119, 113] and finite-state control [111, 112, 124, 127, 128]. Both methods express prosthetic target trajectories and torques as functions of gait phase instead of time, allowing for generalization across tasks, speeds, and users. However, all of these control schemes prescribe prosthesis behavior based on a single (most likely) estimate of human behavior, which could prove disastrous when this estimate is wrong.

In addition to estimation uncertainty, the variability in the human’s trip response makes it difficult for mid-level controllers to generate trajectories that contribute to a successful stumble recovery. Humans typically respond to trips using three main strategies: elevating, delayed lowering, and lowering, described in Sec. 6.2. The choice of trip recovery strategy has been shown to vary based on the perturbation onset and duration [25], and further variation exists in kinematic trajectories within each strategy. Most often, controllers do not include a “trip” classification or strategy [114], and instead may enter a zero-torque mode to let the user recover on their own [129]. One promising strategy is to take the person’s current state as an initial condition and run an optimization problem, solving for a trajectory that gives them the best chance at recovering. Such a method would have an advantage to tracking a nominal trajectory because the result would be created online in response to a specific perturbation. Recent work has proposed a quadratic program for planning recoveries [27] to a “hip drop” perturbation in the swing phase, but the optimization program did not consider trips and generates knee and ankle trajectories for only the most likely human behavior.

A more robust approach would account for uncertainty by designing a controller that is suitable across a range of possible human behaviors. Such a controller could account for estimation uncertainty as well as the variation in human stumble recoveries. However, the high-dimensional, non-linear, hybrid nature of the human-prosthesis system makes set-based online optimization methods difficult. With these challenges in mind, we present a method for planning a prosthesis recovery strategy online that accounts for a range of human behavior. We draw this idea from Reachability-

based Trajectory Design (RTD), which uses reachable sets within a trajectory-planning framework to guarantee collision-free operation of autonomous robots [130]. Recently, RTD was applied to a 7-DOF robotic manipulator (ARMTD), composing reachable sets of each joint in a local coordinate frame into a reachable set of the robot in the global workspace [131].

In this work, we apply a similar reachability-based approach to a kinematic model of a human with a transfemoral, unilateral amputation and a robotic knee prosthesis. We use trip data from 16 able-bodied subjects (Sec. 6.2), adding a prosthetic knee and ankle to the swing leg in simulation. Offline, we fit predictive models of human swing hip behavior (Sec. 6.5.1). Online after a trip occurs, we condition the model on pre-trip trajectories to generate a prediction of the human trip response. We use the confidence bounds to form *reachable sets*: an estimate of all points the swing hip states can reach during trip recovery (Sec. 6.5.1). We compute similar reachable sets for the prosthesis knee and ankle for a set of *trajectory parameters*, which describe a wide range of parameterized prosthesis kinematic behavior. Because the prosthesis' position depends on the human's motion, the prosthesis reachable subsets are combined with the human reachable subsets online. We then solve a real-time optimization problem to select trajectory parameters for the prosthetic knee and ankle (corresponding to a knee and ankle recovery trajectory) that ensures successful trip recovery for the *set* of predicted human responses (Sec. 6.5.4). An overview of the framework can be found in Fig. 6.1. We compare this framework using real-world data in simulation to the prosthesis' nominal impedance-based controller (Sec. 6.6), and find our proposed framework greatly reduces incidences of foot scuffs and is able to place the foot closer to the desired location at the end of the step.

6.2 Trip Experiment and Recovery Strategies

We conducted a 16-subject trip experiment modeled after an experiment conducted by Shirota et al. [25], described in Ch. 2, Sec. 2.2. The experimental setup is shown in Fig. 2.2 in Ch. 2. Several studies have shown that when tripped, humans typically use one of three strategies to recover [132, 95, 25, 96]:

1. Elevating, where the swing foot clears the obstacle and the heel strikes in front of the body,
2. Delayed lowering, where the swing foot initially elevates, then lowers behind the body with the toe contacting the ground first, and
3. Lowering, where the swing foot immediately lowers behind the body with the toe contacting the ground first.

In this chapter, we analyze 430 elevating, 304 delayed lowering, and 223 lowering trips across the 16 subjects, for 957 trips total. We sorted the trip data into the three strategies by hand, and data from left and right swing legs were combined into one dataset. The heel height and anterior-posterior position for these strategies is plotted in Fig. 2.5 in Ch. 2. Note that each strategy diverges from the average nominal foot behavior, suggesting that a nominal lower-limb robotic prosthesis controller would be unable to assist with a successful trip recovery.

6.3 Preliminaries

This section introduces the preliminary components of our trip-recovery planning framework: the planning model, phasing variable, and planning strategy.

6.3.1 Planning Model

This chapter uses a 2D model in the sagittal plane to analyze the swing leg motion, although the framework can generalize to three dimensions. Because we focus on trip recovery during the same step as the trip occurrence and when the prosthesis is in swing, our model is in single support, with the intact limb in stance and the prosthesis in swing. The prosthetic leg is attached to the human thigh, and is composed of a powered knee and ankle joint.

6.3.1.1 Kinematic Model for Proposed Framework

We model the prosthetic leg as a kinematic chain with two revolute joints, which rotate about the same flexion-extension (mediolateral) axis. This kinematic model is illustrated in Fig. 6.2(a). We treat the kinematic tree as composed of a *human subsystem* (the intact hip and upper thigh) and *prosthesis subsystem* (the prosthetic upper thigh, knee, and lower leg). Note that because our trip experiment collected data from able-bodied subjects, we create these human and prosthesis subsystems in simulation.

We define an inertial reference frame with a coordinate system attached to the ground plane and located at the intact limb’s stance heel position, which we assume to be constant during the opposite limb’s swing phase. Let the coordinates $(y, z)^T \in \mathbb{R}^2$ denote the anterior-posterior and vertical components of positions in this coordinate frame, respectively. To be clear, the y -direction is forward and the z -direction is up (in a 3D frame, x would point in the mediolateral direction, specifically to the right so that the frame is right-handed, but our model is restricted to the sagittal plane). We assume that the ground plane is flat and $z = 0$ for all points on the ground.

The movement of the prosthesis in the inertial reference frame is largely determined by the movement of the thigh to which the prosthesis is attached. We treat the thigh as a rigid link




Figure 6.2: The planning model considered in this chapter consists of a transfemoral amputee (black links) and robotic prosthesis (blue links). Part (a) depicts the kinematic model used in our proposed framework. To compare to an impedance-based controller, we use a dynamic model, which includes the link masses and inertias and joint inputs shown in (b).

attached to the amputee's hip, which is described by three coordinates. We let $q_h \in \mathbb{R}^2 \times \mathbb{S}^1$ be the configuration of the thigh, with

$$q_h = \begin{bmatrix} q_y \\ q_z \\ q_1 \end{bmatrix} \quad (6.1)$$

where $(q_y, q_z)^\top \in \mathbb{R}^2$ describe the 2D position of the hip and $q_1 \in \mathbb{S}^1$ is an angle describing the flexion/extension of the hip. Next, let $q_p \in \mathbb{S}^2$ represent the two flexion-extension prosthetic joint angles:

$$q_p = \begin{bmatrix} q_2 \\ q_3 \end{bmatrix} \quad (6.2)$$

where q_2 is the knee angle and q_3 is the ankle angle. Finally, let $q = (q_h, q_p)^\top \in \mathbb{R}^2 \times \mathbb{S}^3 = Q$ be the full configuration of the planning model with *configuration space* Q . Let Q_h denote the configuration space of the hip configuration, and Q_p denote the configuration space of the prosthesis. Let W denote the *workspace*, i.e. the set of all points reachable by any point on the prosthetic leg in any configuration. We also define l_{th} as the length of the thigh segment, l_{sk} as the length of the prosthetic shank segment, l_{hl} as the length from the ankle to the heel, and l_{toe} as the length from the ankle to the toe. We assume a constant angular offset θ_{heel} between the two links

comprising the foot, as can be seen in Fig. 6.2(a). When we write link lengths in braces (i.e. $\{l_{\text{th}}\}$), we mean the point $\{(l_{\text{th}}, 0)^\top\}$ (i.e. the vector of length l_{th} pointing from the hip to knee in a local y - z coordinate frame).

In our trip-recovery planning framework, we are primarily concerned with position trajectories of the prosthetic heel and toe through the workspace. These positions depend on all the configuration variables defined above, as well as the lengths of the links of the kinematic tree. We therefore define the *forward occupancy* maps of the heel and toe, which define the volume in space occupied by the heel and toe for given values of q . We compute these maps by “stacking” the motion of the prosthetic leg on top of the human hip. The forward occupancy maps of the heel $\text{FO}_{\text{heel}} : Q \rightarrow \mathcal{P}(W)$ and $\text{FO}_{\text{toe}} : Q \rightarrow \mathcal{P}(W)$, where \mathcal{P} denotes the power set, are given by:

$$\text{FO}_{\text{heel}}(q) = \left\{ \begin{bmatrix} q_y \\ q_z \end{bmatrix} \right\} \oplus R_1(q_1)\{l_{\text{th}}\} \oplus \prod_{j=1}^2 R_i(q_i)\{l_{\text{sk}}\} \oplus \prod_{j=1}^3 R_i(q_i)R(\theta_{\text{heel}})\{l_{\text{hl}}\} \quad (6.3)$$

$$\text{FO}_{\text{toe}}(q) = \left\{ \begin{bmatrix} q_y \\ q_z \end{bmatrix} \right\} \oplus R_1(q_1)\{l_{\text{th}}\} \oplus \prod_{j=1}^2 R_i(q_i)\{l_{\text{sk}}\} \oplus \prod_{j=1}^3 R_i(q_i)\{l_{\text{toe}}\}. \quad (6.4)$$

The notation $\prod_{j=1}^2 R_i$ is understood as $R_1 R_2$, and $R_j(q_j) \in \text{SO}(3)$ is a rotation matrix corresponding to a rotation by q_j about the flexion-extension axis. The notation \oplus is a Minkowski sum, where $U \oplus V = \{u + v \mid u \in U, v \in V\}$ for $U, V \subset \mathbb{R}^n$. Similarly, $R(\theta_{\text{heel}})$ is a constant rotation matrix representing the offset between the ankle angle and the heel link. Although the maps FO_{heel} and FO_{toe} yield single points for a given configuration $q \in Q$, we have written them using set notation because later we overapproximate these maps using sets.

We may also want to project these forward occupancy maps onto the ground plane, because we specify constraints in terms of heel or toe strike locations on the ground. We therefore project these forward occupancy maps onto the ground plane and vertical axis. Although we overapproximate FO_{heel} and FO_{toe} using sets later in this chapter, we abuse notation and let $\pi_y : \mathbb{R}^2 \rightarrow \mathbb{R}$ project points onto the y -axis and $\pi_z : \mathbb{R}^2 \rightarrow \mathbb{R}$ project points onto the vertical axis. Then, the position of the swing heel in the anterior-posterior direction is $\pi_y(\text{FO}_{\text{heel}}(q))$, and the height of the swing heel is $\pi_z(\text{FO}_{\text{heel}}(q))$.

6.3.1.2 Dynamic Model for Impedance-Based Control Comparison

As mentioned in Sec. 6.1, we compare our proposed planning framework to a nominal impedance-based controller. Because impedance-based control generates desired torques for prosthetic joints, we use a dynamic swing leg model for this controller, shown in Fig. 6.2(b). For each subject in our dataset, we use subject-specific mass and inertia for the human subsystem (estimated using

average distributions from anthropometric data) [106, Ch. 4], then use the dimensions and mass of the University of Michigan Open-Source Leg for the prosthesis subsystem [133]. The thigh, shank, and foot links of the dynamic model have masses m_{th} , m_{sk} , and m_{ft} , respectively, and inertias I_{th} , I_{sk} , and I_{ft} . Using these quantities, the model’s dynamics at a given point in time are represented by the standard manipulator equations:

$$M(q)\ddot{q} + C(q, \dot{q})\dot{q} + G(q) = T, \quad (6.5)$$

where $M(q)$ is the inertia matrix, $C(q, \dot{q})\dot{q}$ are the Coriolis and centrifugal terms, $G(q)$ are forces due to gravity, and the model’s inputs are $T = (T_y, T_z, T_1, T_2, T_3)^\top$. In simulation, we use inverse dynamics to solve for the inputs to the human subsystem T_y, T_z and T_1 which generate the observed kinematics q_h and \dot{q}_h of the hip. The prosthetic leg generates the knee and ankle torques T_2 and T_3 . We compare two controllers for the prosthetic leg in subsequent sections. We note that there are no ground reaction forces present in the model because we only simulate the leg during the swing phase.

Simulating these dynamics produces trajectories of the knee and ankle angles q_2 and q_3 based on the torques that are applied to the model. We can use these prosthesis angles to calculate FO_{heel} and FO_{toe} as in (6.3) and (6.4).

6.3.2 Phase Variable

In our planning framework, we use a *phase variable* rather than time to describe the evolution of system trajectories. Phase variables allow joint angles, angular velocities, and torques to be expressed as a function of gait phase, which can better generalize across walking speeds and subjects compared to time. In this chapter, we use a phase variable to align groups of swing phase trajectories for creating trip-recovery libraries in Sec. 6.4.1; fitting predictive models in Sec. 6.4.2; computing reachable sets in Sec. 6.5.2; and developing the nominal controller for comparison in Sec. 6.6.

In robotic lower-limb prostheses, the phase variable can be measured via wearable sensors during gait and used to prescribe corresponding prosthesis behavior in real-time [113]. These phase variables must be monotonically increasing or decreasing through the segment of interest. Previous research has used the thigh angle as a phase variable for walking gait, but trip-recovery adds an additional challenge. In delayed lowering and lowering strategies, for example, subjects initially move the thigh forward during swing phase, but then reverse the swing leg’s direction and place the swing foot behind the body after the trip occurs.

We found that the anterior-posterior (AP) position of the pelvis relative to the AP position of the stance foot followed more consistent patterns across trip-recovery strategies than the swing thigh

angle. Note that the pelvis is not included in the swing leg model in Sec. 6.3.1. For this section, we denote the relative AP pelvis position as y_{pel} . We begin measuring y_{pel} values as soon as the swing phase is detected:

Assumption 6.3.1. *The prosthesis immediately detects the onset of swing phase, and the relative AP pelvis position y_{pel} is measured at n evenly spaced time points (determined by the sensor sampling rate) until touchdown.*

Despite improvements from the swing thigh angle, y_{pel} does not monotonically increase throughout the swing phase. During the swing phase, y_{pel} typically first increases monotonically, then decreases monotonically. We denote the phase variable as t , and define the first phase point t_1 as the observed y_{pel} at toe-off. For a given measurement index $i \in \{2, \dots, n\}$, we ensure t_i is not decreasing from the previous measured point by adding the absolute difference between the current AP pelvis position $y_{\text{pel},i}$ and the previous value, $y_{\text{pel},i-1}$. Then, we define the phase variable t_i as:

$$t_i = \begin{cases} y_{\text{pel},i}, & \text{if } i = 1 \\ y_{\text{pel},i-1} + |y_{\text{pel},i} - y_{\text{pel},i-1}|, & \text{otherwise} \end{cases} \quad (6.6)$$

We normalize each t_i to be in $[0, 1]$ using the minimum and maximum observed phase values across nominal and tripped trials. In this chapter, we use subject-specific bounds for normalization, but it is also possible to use generalized bounds formed from the entire trip-recovery dataset.

6.3.3 Planning Strategy

During unperturbed walking, there are several nominal controllers (e.g., volitional [113] or finite-state control [127]) that a powered prosthesis may use to prevent falls. However, when a person trips, a different prosthesis response may be necessary to keep the user from falling. The goal of our trip-recovery framework is to quickly design recovery trajectories in response to perturbation that lead to a successful recovery.

We define the phase at which swing foot contact occurs:

Definition 6.3.2. *Swing foot contact occurs at $t_f \in [0, 1]$. Specifically, t_f is the first $t \in [0, 1]$ such that $\pi_z(\text{FO}_{\text{heel}}(q(t))) \leq 0$ or $\pi_z(\text{FO}_{\text{toe}}(q(t))) \leq 0$.*

Recall that $\pi_z(\text{FO}_{\text{heel}}(q(t)))$ and $\pi_z(\text{FO}_{\text{toe}}(q(t)))$ give the vertical position of the swing heel and toe at configuration $q(t)$, respectively, so t_f is the first time that the swing heel or toe breaks the ground plane.

In unperturbed walking, heel strike signals a transition to the next double support phase. However, for the delayed lowering and lowering trip recovery strategies, it is often the swing toe that

first makes contact with the ground. Therefore we more generally say that a transition to double support occurs when either the heel or toe strikes the ground.

Our framework only considers swing phase trajectories from the instance a stumble is detected until t_f :

Assumption 6.3.3. *The trip begins at $t_{\text{trip}} \in [0, 1]$ and is immediately detected, and the initial state of the prosthesis is known, i.e. $q_p(t_{\text{trip}}) = q_{p0}$ and $\dot{q}_p(t_{\text{trip}}) = \dot{q}_{p0}$. After the phase $t_f \in [0, 1]$ when foot contact and a transition to double support occurs, the prosthesis’ nominal controller resumes control.*

We include immediate trip detection in our assumption given previous trip-detection work by other researchers [121, 122]. Furthermore, we plan strategy-specific trip-recovery trajectories:

Assumption 6.3.4. *At trip onset, we immediately know which strategy (elevating, delayed lowering, or lowering) that the individual will use to recover.*

We justify this assumption given Eveld et al.’s high accuracy in predicting trip-recovery strategy selection [96].

After a stumble occurs, we immediately apply our proposed framework to plan trip-recovery trajectories for the knee-ankle prosthesis. If the optimization returns a feasible solution, the prosthesis tracks this desired trajectory for $t \in [t_{\text{trip}}, t_f]$. Otherwise, if no safe trajectory can be found in the current planning iteration, the prosthesis continues to follow the nominal trajectory.

Because the exact instance of touchdown t_f is unknown due to uncertainty in the human trip response, our trip-recovery framework must generate prosthesis trajectories for a time horizon that includes all possible t_f values. Elevating trials can last until $t = 1$, but delayed lowering and lowering trials end much earlier in the normalized swing phase. Likewise, trials often start after $t = 0$. To avoid planning prosthesis trajectories for phase ranges where we have no observed data available, we set a strategy-specific initial and final phase bound for planning. Let $s \in \{1, 2, 3\}$ index the strategies (elevating, delayed lowering, and lowering, respectively). Then we define t_{lb}^s as the minimum observed phase value in strategy s and t_{ub}^s as the maximum observed phase value in strategy s . Note that t_{lb}^s was determined from pre-trip phase points, so $t_{\text{trip}} \geq t_{\text{lb}}^s$ for all trials. In our framework, once strategy s is detected, we plan a trajectory for $t \in [t_{\text{trip}}, t_{\text{ub}}^s]$, then execute the plan until foot contact occurs, i.e. when $t \in [t_{\text{trip}}, t_f]$.

6.4 Offline Computations

This section describes the offline computations in our framework. For each component, we first introduce the theory and then describe our specific implementation for this chapter.

6.4.1 Parameterized Trip-Recovery Libraries

Data from able-bodied subjects are often used to describe desired prosthesis motion. For example, Rezazadeh et al.'s level-ground walking controller uses average able-bodied knee and ankle angles from [106] as a reference trajectory for the prosthesis to track [113]. Likewise, we use data from our able-bodied trip experiment to create *trip-recovery libraries*: a collection of knee and ankle angle trajectories that the prosthesis can choose from online to ensure a successful trip recovery. We parameterize the trip-recovery library to make the planning framework amenable to real-time computation. Each trajectory in the library is represented by a *trajectory parameter* that describes the evolution of the prosthesis' generalized coordinates over time.

6.4.1.1 Theory

We define trajectory parameters as follows:

Definition 6.4.1. *Let $K \subset \mathbb{R}^{n_k}$ be a compact set of n_k trajectory parameters, where each $k \in K$ maps to a desired trajectory of the generalized coordinates $q_p : [0, 1] \rightarrow Q_p$. Let $q_p(t; k)$ be the prosthesis configuration at phase $t \in [0, 1]$ parameterized by $k \in K$. $\dot{q}_p(\cdot; k)$ is a continuous function with respect to phase.*

As an example parameterization, we may choose $k = (k_2, k_3) \in K \subset \mathbb{R}^2$ to parameterize a desired mapping from step height to trajectories of the knee and ankle joints. The goal of our framework is to select a trajectory parameter $k \in K$ that yields trajectories of the prosthetic joints which lead to successful trip recovery across a range of predictions of human motion.

These parameterized trajectories are kinematic, not dynamic:

Assumption 6.4.2. *The prosthesis is capable of perfectly tracking the trajectories specified by the trajectory parameterization.*

Because the prosthesis is in swing, its motors only have to move its own mass when tracking these trajectories. In Sec. 6.5.2, we ensure commanded trajectories start with the correct initial post-perturbation positions. This makes tracking the trajectories easier, as there is no initial error between the desired and actual positions of the robotic leg.

6.4.1.2 Implementation

Recall that we described three distinct trip recovery strategies in Sec. 6.2. Let $s \in \{1, 2, 3\}$ index the strategies, and let s^r refer to the r -th observed trial within that strategy. Let $q_2^{s^r}(t)$ and $q_3^{s^r}(t)$ be the observed knee and ankle joint angles at phase t observed for a specific trial. We fit a library

of parameterized trip recoveries to these observed trials. Specifically, fit polynomial functions of t and k to the joint angle trajectories observed in each recovery strategy.

First, we compute a *mean* angle trajectory for each of the observed strategies for the knee and ankle joints. We refer to these as $q_{2,\text{mean}}^s$ and $q_{3,\text{mean}}^s$. As a result, we define

$$q_2^s(t; k) = q_{2,\text{mean}}^s(t) + a_0 + a_1k_2 + a_2k_2^2 + a_3k_2^3 + a_4t + a_5tk_2 + a_6tk_2^2 + a_7t^2 + a_8t^2k_2 + a_9t^3 \quad (6.7)$$

$$q_3^s(t; k) = q_{3,\text{mean}}^s(t) + b_0 + b_1k_2 + b_2k_2^2 + b_3k_2^3 + b_4t + b_5tk_2 + b_6tk_2^2 + b_7t^2 + b_8t^2k_2 + b_9t^3 \quad (6.8)$$

where we use linear regression to solve for the coefficients $a_0, \dots, a_9, b_0, \dots, b_9$ that minimize the difference from a strategy's observed trials in the least-squares sense. In other words, we find the coefficients which minimize the squared difference between $q_2^s(t; k)$ and $q_3^s(t; k)$ and each observed $q_2^{s^r}(t)$ and $q_3^{s^r}(t)$ for $r = 1, \dots, n_s^r$. To employ linear regression to fit these coefficients, we assign the values of k_2 and k_3 based on the observed trajectories. For a given trial, we assign the k_2 and k_3 values based on the amplitude of the knee and ankle angles at 50% of the strategy's phase range, $0.5(t_{\text{ub}}^s + t_{\text{lb}}^s)$.

WLOG, we normalize k_2 and k_3 so that $k_2 \in [-1, 1]$, $k_3 \in [-1, 1]$. To be clear, $k = (k_2, k_3)$ and therefore $K = [-1, 1] \times [-1, 1]$. The angle q_1 corresponds to the individual's intact hip angle and is not parameterized by k .

These polynomials are fit separately to each s th observed trip recovery strategy. We fit the polynomials in $[t_{\text{lb}}^s, t_{\text{ub}}^s]$ for the s th trip recovery strategy. In this chapter, we fit subject-specific trip-recovery libraries, but generalized libraries are also possible due to similarities in knee and ankle trajectories across subjects.

Because not all trials last the same length, we observed unnatural jumps in the computed knee and ankle mean trajectories near t_{lb}^s and t_{ub}^s . Therefore, we extrapolate each observed trip trial to span $[t_{\text{lb}}^s, t_{\text{ub}}^s]$ with the constant initial or final observed value before taking the mean knee and ankle angles. When fitting the polynomials, we remove the extrapolated values from each trial and replace them with the computed mean at each phase point. In other words, we assume that trajectories that started later than t_{lb}^s or ended earlier than t_{ub}^s converge to the mean at t_{lb}^s and t_{ub}^s . Example knee and ankle trajectory parameterizations are displayed in Fig. 6.3.

Some trip trials in our dataset did not follow the ground contact patterns described by previous researchers. For example, we observed trip-recovery trajectories where subjects elevated their swing foot and moved their swing leg forward, but then contacted the ground with the ball or toe of their swing foot rather than the heel. In these trials, the knee and ankle angles follow different patterns than the typical trajectories observed during each strategy. Because we aim to have the

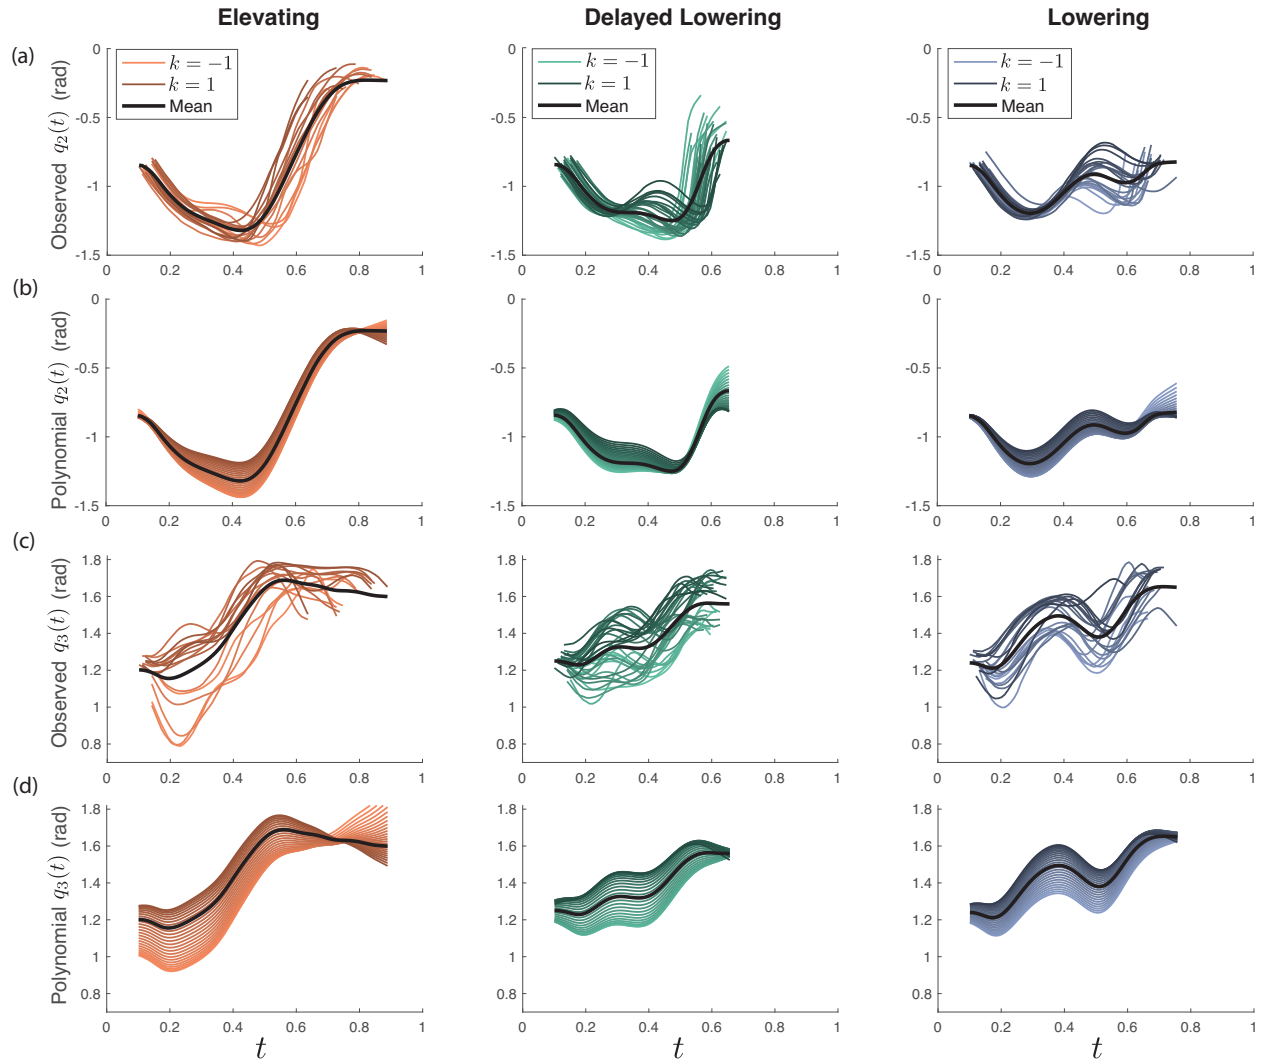


Figure 6.3: This figure shows example observed knee and ankle trajectories and parameterized knee and ankle trajectories for Subject ID 1. In (a) and (c), the observed ankle velocity trajectories are plotted for each trip recovery strategy. A k -value in $[-1, 1]$ is assigned to each trajectory based on its amplitude at 50% of swing phase. In (b) and (d), the parameterization of ankle velocity trajectories $q_3(t; k)$ is shown. The mean trajectories $q_{2,\text{mean}}^s(t)$ and $q_{3,\text{mean}}^s(t)$ are shown in black, and the gradient from dark to light is created by varying the value of k_2 or k_3 from -1 to 1 . We note that in each of these plots, the phase $t \in [0, 1]$ has been scaled the same for each strategy, highlighting differences in phase duration between each strategy.

prosthesis always contact the ground with the heel during elevating and the toe during delayed lowering and lowering, we remove the trials with the opposite foot contact when training the trip-recovery libraries.

6.4.2 Training Predictive Models of Swing Hip Kinematics

The human’s swing hip motion affects the location of the prosthetic heel and toe as described in (6.3) and (6.4). Because it is impossible to perfectly predict human behavior, we want to choose a desired trajectory of the prosthesis (represented by $k \in K$, as in (6.7) and (6.8)) that is robust to a range of plausible human trip-recovery responses. This section describes our approach for obtaining this predicted range of human behavior.

6.4.2.1 Theory

In our model of the swing leg, the human response is composed of the states in q_h : hip height, AP position, and flexion-extension angle. As in Assum. 6.3.3, we assume the state of the human subsystem at trip onset is known:

Assumption 6.4.3. *The initial state of the human is known, i.e. $q_h(t_{\text{trip}}) = q_{h0}$ and $\dot{q}_h(t_{\text{trip}}) = \dot{q}_{h0}$.*

We then generate a prediction of q_h states for $t \in [t_{\text{trip}}, t_{\text{ub}}^s]$, where t_{ub}^s is the upper phase bound for strategy $s \in \{1, 2, 3\}$. To obtain a *range* of predicted q_h states, our framework necessitates a model that generates a confidence interval or buffer around the most likely value of each state at each t . These models can be trained offline, then used online given a specific q_{h0} and t_{trip} (Sec. 6.5.1).

6.4.2.2 Implementation

We fit multi-output Gaussian process regression (GPR) models to the states in q_h . In previous work, we found that GPR models produce higher prediction accuracy and lower online computation time compared to a neural network and dynamic model (Ch. 5 and [26]). Because we know the trip-recovery strategy that the individual will select (Assum. 6.3.4), we fit separate GPR models to each strategy. We train subject-specific GPR models for this chapter. In Ch. 5, we found that generalized GPR models (fit to multiple subjects’ data) produced predictions with lower accuracy than subject-specific models. Improving the generalized model accuracy in future work can help increase the real-world applicability of our proposed framework.

The phase variable and system states in our trip-recovery dataset are measured at the Vicon motion capture system’s sampling rate (Sec. 6.2), producing a strategy-specific training dataset of n points for fitting each GPR model. For the i th measurement ($i = 1, \dots, n$), let t_i denote

the measured phase and $q_{h,i}$ denote the measured swing hip states. GPR models represent the relationship between each input t_i and output $q_{h,i}$ by a Gaussian noise model. For a particular strategy $s \in \{1, 2, 3\}$ and at a point of interest $t_* \in [t_{\text{lb}}^s, t_{\text{ub}}^s]$, the model produces a 3×1 mean μ_* , representing the model’s mean prediction of the states in q_h at $t = t_*$, and a 3×3 covariance matrix Σ_* . Evaluating the model’s mean and covariance at m evenly spaced values of t_* from t_{lb}^s to t_{ub}^s gives us a $3m \times 1$ mean prediction vector $\mu = [\mu_1^T, \mu_2^T, \dots, \mu_m^T]^T$ and its associated $3m \times 3m$ covariance Σ . The details for constructing and training the multi-output Gaussian noise model are described in Ch. 5.

6.4.3 Target Sets

The goal of our framework is to plan safe trip-recovery trajectories. To characterize “safe recoveries” from a trip, we create *target sets* which represent states at swing foot touchdown that are known to lead to successful trip recovery. To ensure that the prosthesis behavior aligns with the human’s intent, the prosthesis foot placement at the end of the step should also depend on some or all of the human subsystem states. The definition of a successful trip recovery could be formed in many ways, for example, using formal definitions of walking stability such as those described by [56]. In this paper, we examine combinations of human subsystem states and swing foot positions at touchdown.

6.4.3.1 Theory

Definition 6.4.4. A target set X_T is a set composed of a subset of human subsystem states q_{h,X_T} , their velocities \dot{q}_{h,X_T} , and swing heel/toe positions on the ground plane at the time of foot contact that lead to successful trip recovery. The target sets may differ for each trip recovery strategy. Its formulation in terms of heel or toe positions depends on which portion of the foot first makes contact with the ground under that strategy. Recall that in our formulation, the foot contact phase value is t_f and the y swing heel/toe position for a given configuration $q(t)$ is $\pi_y(\text{FO}_{\text{heel}}(q(t)))$ or $\pi_y(\text{FO}_{\text{toe}}(q(t)))$. If, at the time of foot contact t_f ,

$$\begin{bmatrix} q_{h,X_T}(t_f) \\ \dot{q}_{h,X_T}(t_f) \\ \pi_y(\text{FO}_{\text{heel}}(q(t_f))) \end{bmatrix} \in X_T \quad \text{or} \quad \begin{bmatrix} q_{h,X_T}(t_f) \\ \dot{q}_{h,X_T}(t_f) \\ \pi_y(\text{FO}_{\text{toe}}(q(t_f))) \end{bmatrix} \in X_T \quad (6.9)$$

the trip recovery is considered successful.

6.4.3.2 Implementation

Similar to previous work on the construction of target sets for the Sit-To-Stand motion [23], we build target sets from data observed to lead to successful walking following a trip (i.e., observed trips where subjects did not lose balance, fall, or undergo a second trip following the trip). Recent work has shown that the position of the swing foot at heel strike following perturbation is predicted well by a linear map of center of mass positions and velocities [99]. Because we collected only lower-body data in our trip experiment, we use the swing hip instead of the center of mass. We select the hip anterior-posterior position and velocity as our subsystem of human states to include in the target set, i.e. $q_{h,X_T} = q_y$ and $\dot{q}_{h,X_T} = \dot{q}_y$.

To form the target set over observed q_{h,X_T} , \dot{q}_{h,X_T} , and foot contact positions, we use *zonotopes*, geometric objects that can efficiently represent sets (introduced in Ch. 1, Sec. 1.3). A zonotope Z is a convex polytope that is closed under linear maps and Minkowski sums [28]. It describes the set of points that can be reached by a center, $c \in \mathbb{R}^n$, and linear combination of generators, $g_1, \dots, g_{n_g} \in \mathbb{R}^n$. The coefficient β_i on each generator lies in $[-1, 1]$:

$$Z = \left\{ y \in \mathbb{R}^n \mid y = c + \sum_{i=1}^{n_g} \beta_i g_i, -1 \leq \beta_i \leq 1 \right\}. \quad (6.10)$$

Each X_T is formed using zonotopes with 3 generators each, which encompass all anterior-posterior hip positions and velocities and foot contact locations observed to lead to successful walking. This mirrors the method found in [23] for generating zonotopes from sets of points. We fit separate target sets for each subject and strategy, and enlarge each one by 20% to add a buffer to states lying on the edge of the computed volume. Target sets computed for Subject ID 1 are displayed in Fig. 6.4.

6.5 Online Computations

When a trip occurs, we assume that the trip is immediately detected and the individual’s chosen recovery strategy is known, as stated in Assums. 6.3.3 and 6.3.4. We then execute a series of online computations. First, we generate *reachable sets* of the predicted human hip states: the set of all states that can be attained by the system starting from the initial set. We then compute parameterized reachable sets of prosthesis knee and ankle angles and combine the human and prosthesis sets to obtain reachable sets for the prosthesis heel and toe. We formulate cost and constraints with these heel/toe reachable sets, and solve an optimization problem to select $k \in K$ to ensure a successful trip recovery. As in Sec. 6.4, we first introduce the theory for each step, then describe our implementation.

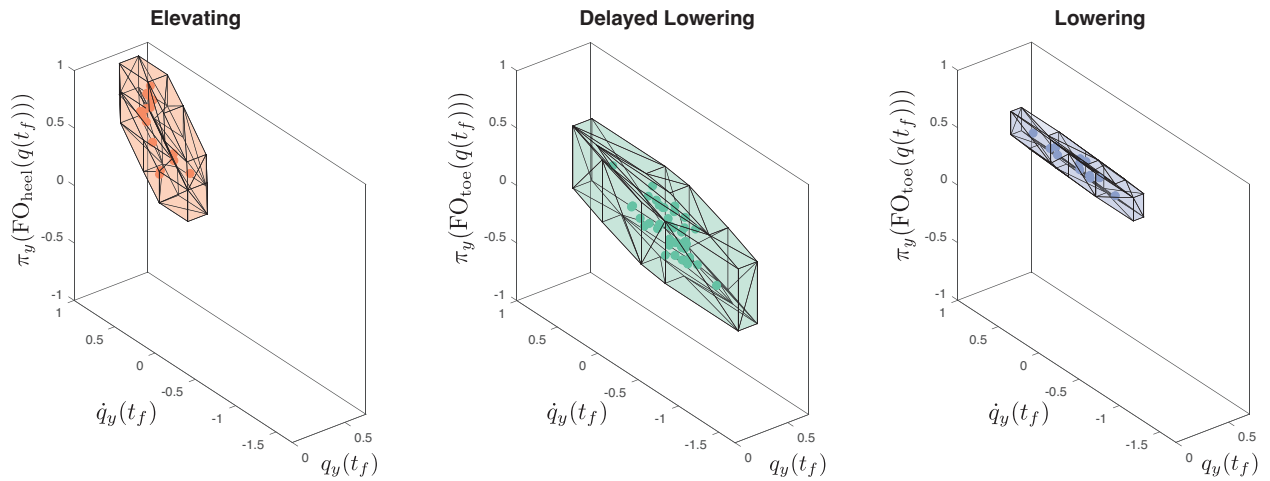


Figure 6.4: This figure shows the target sets X_T computed from Subject ID 1's data. The orange, green, and purple points represent the hip positions and velocities and foot strike locations at the end of each trial for elevating, delayed lowering, and lowering, respectively. The volumes encompassing each set of points are the zonotope target sets. Note that the target sets are defined over heel positions for the elevating strategy and over toe positions for the delayed lowering and lowering strategies. The target sets found for each strategy occupy distinct regions. As expected for the elevating strategy, X_T is characterized by more forward progression of the heel and hip positions, and a higher hip velocity. For the delayed lowering strategy, X_T occupies a region of low forward progression of the toe and hip positions, and negative hip velocities. The lowering strategy's target set lies between these two.

6.5.1 Conditional Predictions and Reachable Sets of Human Hip States

6.5.1.1 Theory

In Sec. 6.4.2, we fit predictive models of the states in q_h . Online, after a trip occurs, the model must generate a prediction informed by the human's pre-trip configuration. Such a prediction could be simulated from an initial condition in the case of a dynamic model; pre-trip feedback points for a neural network; or conditional points for a probabilistic method:

Assumption 6.5.1. *In addition to knowing the human's exact state at the time of the trip (Assum. 6.4.3), we have access to all previous observed q_h points in the current swing phase trajectory.*

We generate the prediction of q_h states from $[t_{\text{trip}}, t_{\text{ub}}^s]$, where $s \in \{1, 2, 3\}$ indicates a specific recovery strategy.

As mentioned in Sec. 6.5.1, we also need a measure of prediction uncertainty when generating the online predictions. Let $\hat{\mu}_y(t) - C_y(t)$ and $\hat{\mu}_y(t) + C_y(t)$ denote the lower and upper bounds of prediction uncertainty for the anterior-posterior hip position at t , where $C_y(t)$ represents half the width of the confidence interval. Likewise, let subscripts $C_z(t)$, $C_{\dot{y}}(t)$, and $C_1(t)$ indicate similar confidences for the hip height, anterior-posterior velocity, and angle, respectively. The online predictions of q_y and their associated prediction uncertainties form the reachable set of hip y translations for $t \in [t_{\text{trip}}, t_{\text{ub}}^s]$:

$$\mathcal{J}_y(t) = \left\{ y \in \mathbb{R} \mid y = q_y(t), y \in [\hat{\mu}_y(t) - C_y(t), \hat{\mu}_y(t) + C_y(t)] \right\}. \quad (6.11)$$

Likewise, we define the reachable set of hip z translations for $t \in [t_{\text{trip}}, t_{\text{ub}}^s]$ as

$$\mathcal{J}_z(t) = \left\{ z \in \mathbb{R} \mid z = q_z(t), z \in [\hat{\mu}_z(t) - C_z(t), \hat{\mu}_z(t) + C_z(t)] \right\}. \quad (6.12)$$

Because we also use predictions of \dot{q}_y to define a successful recovery (Sec. 6.4.3), we define reachable sets of hip anterior-posterior velocity values for $t \in [t_{\text{trip}}, t_{\text{ub}}^s]$ (Note, these can be formed by differentiating the predicted q_y trajectory):

$$\mathcal{J}_{\dot{y}}(t) = \left\{ \dot{y} \in \mathbb{R} \mid \dot{y} = \dot{q}_y(t), \dot{y} \in [\hat{\mu}_{\dot{y}}(t) - C_{\dot{y}}(t), \hat{\mu}_{\dot{y}}(t) + C_{\dot{y}}(t)] \right\}. \quad (6.13)$$

The prediction and uncertainty of q_1 produces a reachable set of hip flexion-extension angles for $t \in [t_{\text{trip}}, t_{\text{ub}}^s]$, which we define as follows.

$$\mathcal{J}_1(t) = \left\{ p \in \mathbb{R} \mid p = q_1(t), p \in [\hat{\mu}_1(t) - C_1(t), \hat{\mu}_1(t) + C_1(t)] \right\}. \quad (6.14)$$

6.5.1.2 Implementation

At trip onset, we generate conditional predictions using the GPR models in Sec. 6.4.2. Recall that we sampled the mean and covariance for strategy s 's GPR model at m evenly spaced phase points from $[t_{\text{lb}}^s, t_{\text{ub}}^s]$. The trip trial does not necessarily start at $t = t_{\text{lb}}^s$ or end at $t = t_{\text{ub}}^s$. When the trip occurs at t_{trip} , we find i_{trip} , the index of the m phase points that most closely corresponds to t_{trip} . We re-sample the trip trial's pre-trip points, and pad with NaNs at the beginning if needed, so that the trial has i_{trip} evenly spaced points that align with the GPR model's sampled mean μ and covariance Σ . We denote this re-sampled pre-trip trajectory as $q_{h,c}$, where the c stands for *completed*. Likewise, we define the GPR model's mean prediction for the completed points: $\mu_c = [\mu_1^T, \mu_2^T, \dots, \mu_{i_{\text{trip}}}^T]^T$. The n remaining points in the GPR model's mean prediction are labeled as *future* points: $\mu_f = [\mu_{i_{\text{trip}}+1}^T, \mu_{i_{\text{trip}}+2}^T, \dots, \mu_m^T]^T$. We also partition the covariance matrix Σ into completed and future parts:

$$\Sigma = \begin{bmatrix} \Sigma_{cc} & \Sigma_{cf} \\ \Sigma_{fc} & \Sigma_{ff} \end{bmatrix}. \quad (6.15)$$

We use the completed and future portions of μ and Σ to obtain a conditional prediction $\hat{\mu}$ and its uncertainty $\hat{\Sigma}$:

$$\begin{aligned} \hat{\mu} &= \mu_f + \Sigma_{fc} \Sigma_{cc}^{-1} (q_{h,c} - \mu_c) \\ \hat{\Sigma} &= \Sigma_{ff} - \Sigma_{fc} \Sigma_{cc}^{-1} \Sigma_{cf}. \end{aligned} \quad (6.16)$$

We separate the multi-output mean into n -dimensional vectors $\hat{\mu}_y$, $\hat{\mu}_z$, and $\hat{\mu}_1$ for the hip anterior-posterior position, height, and angle. To obtain the hip anterior-posterior velocity prediction, we numerically differentiate $\hat{\mu}_y$ to produce $\hat{\mu}_{\dot{y}}$. We form the confidence interval for each variable using the prediction covariance.

We use a 60% confidence interval to compute the q_h reachable sets in this chapter to avoid making the heel and toe reachable sets too large, but a higher confidence interval is possible if a more conservative planner is desired. We do not include confidence bounds on the velocity prediction (i.e., $\mathcal{J}_{\dot{y}}$ is a trajectory rather than a set), but our framework allows for uncertainty in this prediction as well.

We overapproximate the hip reachable sets using zonotopes, introduced in Sec. 6.4.3. However, we represent these zonotopes as polynomial zonotopes, which are a more complex and expressive set representation amenable to the computations we utilize in the next section. Polynomial zonotopes are similar to zonotopes except that each generator may be multiplied by a product of indeterminates. We provide a more detailed overview of polynomial zonotopes in Appx. A.4.

Briefly, a polynomial zonotope $\langle P \rangle$ is defined by generators $g_i \in \mathbb{R}^n$, an indeterminate vector

$x \in [-1, 1]^{n_x}$, and n_g exponents $\alpha_i \in \mathbb{N}^{n_x}$

$$\langle P \rangle = \{z \in \mathbb{R}^n \mid z = \sum_{i=0}^{n_g} g_i x^{\alpha_i}, x \in [-1, 1]^{n_x}\} \quad (6.17)$$

where the exponentiation is applied elementwise. Comparing (6.17) with the zonotope definition (6.10), we see that the set definitions are similar except that polynomial zonotope generators may be multiplied by exponentiated indeterminates.

Throughout this section, we will use the following notation for polynomial zonotopes. We let the bracket notation $\langle \cdot \rangle$ denote a polynomial zonotope. We write $\langle P \rangle = \mathcal{PZ}(g_i, \alpha_i, x)$ to say $\langle P \rangle$ is a polynomial zonotope created from generators g_i , exponents α_i , and indeterminate vector x . Note that by letting $\alpha_0 = 0$ (a vector of all zeros), the corresponding generator g_0 functions like a zonotope center.

We represent \mathcal{I}_y , \mathcal{I}_z , \mathcal{I}_{ij} , and \mathcal{I}_1 using one polynomial zonotope per time interval, denoted as $\langle q_y \rangle_i$, $\langle q_z \rangle_i$, $\langle q_{ij} \rangle_i$, and $\langle q_1 \rangle_i$ for a given phase index i . Each of these is comprised of a center and a single generator with a single indeterminate. At each discretized phase step $i \in \{1, \dots, n\}$, the polynomial zonotope center g_0 is given by the mean conditional prediction at that phase value, and the generator value is the conditional prediction's 60% confidence interval: We write

$$\langle q_y \rangle_i = \hat{\mu}_y + C_y x_y \quad (6.18)$$

$$\langle q_z \rangle_i = \hat{\mu}_z + C_z x_z \quad (6.19)$$

$$\langle q_1 \rangle_i = \hat{\mu}_1 + C_1 x_1 \quad (6.20)$$

$$\langle q_{ij} \rangle_i = \hat{\mu}_{ij} + C_{ij} x_{ij}, \quad (6.21)$$

where the indeterminates x_y , x_z , x_1 and x_{ij} are each contained in $[-1, 1]$. Lastly, we let $\langle q_{yz} \rangle_i = \langle q_x \rangle_i \times \langle q_y \rangle_i$.

6.5.2 Parameterized Reachable Sets of the Prosthesis Heel and Toe

Equipped with reachable sets of predicted hip motion, we now describe our method for using the trip-recovery libraries computed in Sec. 6.4.1 to compose reachable sets of the prosthetic knee and ankle joint motion that are parameterized by $k \in K$. We then combine the hip, knee, and ankle reachable sets to create reachable sets of the heel and toe through the workspace. To be clear, the heel and toe reachable sets are a function of both the hip predictions (and their uncertainty) and the trajectory parameters $k \in K$. The following subsections will rely on these reachable sets to enforce a ground clearance constraint and desired foot placement objective.

6.5.2.1 Theory

As in Sec. 6.5.1, we compute reachable sets of the prosthetic joint angles. Unlike the hip reachable sets, the prosthetic reachable sets are parameterized by the prosthetic joint angle trajectories described in Sec. 6.4.1. Recall that the prosthetic knee and ankle angles are described by q_2 and q_3 . The reachable set for trajectory parameter $k \in K$ and time $t \in [t_{\text{trip}}, t_{f,\text{max}}(k)]$ of q_2 and q_3 are defined as:

$$\mathcal{J}_j(t; k) = \left\{ p \in \mathbb{R}^2 \mid \text{s.t. } q_j \text{ as in Def. 6.4.1, } p = \tilde{q}_i + q_j(t; k) \right\} \quad (6.22)$$

where $j = 2, 3$ and $\tilde{q}_i = q_j(t_{\text{trip}}; k) - q_{j,0}$ is the offset of the desired trajectory from the actual initial configuration $q_{j,0}$ at the time of the trip.

With definitions of each joint's reachable sets, we can now define the heel reachable set in the workspace as:

$$\begin{aligned} \mathcal{L}_{\text{heel}}(t; k) = \left\{ X \subset \mathcal{P}(W) \mid X = \text{FO}_{\text{heel}}(q(t; k)), \right. \\ \left. \begin{aligned} q(t; k) &= (q_y(t), q_z(t), q_1(t), q_2(t; k), q_3(t; k))^{\top}, \\ q_y(t) &\in \mathcal{J}_y(t), \quad q_z(t) \in \mathcal{J}_z(t), \quad q_1(t) \in \mathcal{J}_1(t), \\ q_j(t; k) &\in \mathcal{J}_j(t; k) \quad \forall j \in \{2, 3\} \end{aligned} \right\}, \end{aligned} \quad (6.23)$$

where $\mathcal{L}_{\text{heel}} : [t_{\text{trip}}, t_{\text{ub}}^s] \times K \rightarrow \mathcal{P}(W)$ and t_{ub}^s is the upper phase bound for strategy $s \in \{1, 2, 3\}$ corresponding to elevating, delayed lowering, and lowering, respectively. The toe reachable set $\mathcal{L}_{\text{toe}}(t; k)$ is similarly defined using $\text{FO}_{\text{toe}}(q(t; k))$.

6.5.2.2 Implementation

As in Sec. 6.5.1, we represent \mathcal{J}_j for $j = 2, 3$ using one polynomial zonotope per time interval. We create these polynomial zonotopes using polynomial zonotope arithmetic, which is detailed in Appx. A.4. First, we create polynomial zonotopes representing each time interval. Let $\langle T \rangle_i$ correspond to the i -th time interval ($i \in \{1, \dots, n\}$) given by $[(i-1)\Delta t, i\Delta t]$, where $\langle T \rangle_i = \frac{(i+(i-1))}{2}\Delta t + \frac{\Delta t}{2}x_t$, and the indeterminate $x_t \in [-1, 1]$. Next, we define $\langle K_j \rangle_i$ which represents trajectory parameters $K_j = [0, 1]$ as $\langle K_j \rangle_i = 0.5 + 0.5x_{k_j}$, with indeterminate $x_{k_j} \in [-1, 1]$. With these definitions, we are able to plug in $\langle T \rangle_i$ and $\langle K_j \rangle_i$ in place of t and k to the desired trajectories given in (6.7) and (6.8). Then, we add the initial condition offset \tilde{q}_i to account for a possible mismatch between the initial prosthesis configuration and the parameterized trajectory. This yields a sequence of polynomial zonotope versions of the trajectories $\langle q_2 \rangle_i$ and $\langle q_3 \rangle_i$ which

are guaranteed to overapproximate (6.22).

These polynomial zonotopes depend on both time and the trajectory parameter $k \in K$. One essential property of polynomial zonotopes is that one can obtain a subset by plugging in particular values to a polynomial zonotope's indeterminates. We call this process "slicing". Let k_* denote a particular value of k in $[-1, 1]^2$. Note that $k_* \in [-1, 1]^2$ corresponds directly to values of $x_k \in [-1, 1]^2$. We say that

$$\text{slice}(\langle P \rangle, x_k, k_*) = \{z \in \langle P \rangle_i \mid z = \sum_{i=0}^{n_g} g_i x^{\alpha_i}, x_k = k_*\}. \quad (6.24)$$

For brevity, we frequently use the notation

$$\langle P; k \rangle = \text{slice}(\langle P \rangle, x_k, k_*). \quad (6.25)$$

This operation is useful because it allows us to compute subsets of reachable sets corresponding to particular choices of trajectory parameters. For example, the set $\langle q; k \rangle_i$ where $j \in \{2, 3\}$ overapproximates the subset of reachable joint angles for the the i -th phase interval. We construct reachable sets of the motion of the heel and toe from these joint reachable sets, which are also *sliceable* by the trajectory parameters.

It is important that we overapproximate $\mathcal{L}_{\text{heel}}(t; k)$ and $\mathcal{L}_{\text{toe}}(t; k)$ to guarantee the avoidance of foot scuffs for the entire set of predicted human motion. We have formed polynomial zonotopes which overapproximate $\mathcal{J}_y(t; k)$, $\mathcal{J}_z(t; k)$ and $\mathcal{J}_j(t; k)$ for $j = 1, 2, 3$, but now we use these sets to overapproximate the forward occupancy maps FO_{heel} and FO_{toe} . As seen in (6.3) and (6.4), these maps are formulated in terms of rotation matrices built from the cosine and sine of each rotational joint. Polynomial zonotopes offer a means to overapproximate these analytic functions using overapproximative Taylor expansions, detailed in Appx. A.4. We utilize this to generate polynomial zonotopes which overapproximate $\cos(\langle q_j \rangle_i)$ and $\sin(\langle q_j \rangle_i)$ for $j \in \{1, 2, 3\}$. Then, we combine these into polynomial zonotope versions of rotation matrices $\langle R_i \rangle_i$. We note that this operation preserves *sliceability*.

Once these rotation matrix polynomial zonotopes for all rotational joints have been formed, we compute their products to overapproximate the forward occupancy maps FO_{heel} and FO_{toe} at each time step. Here, we use the polynomial zonotopes $V_{\text{heel}}(t)$ and $V_{\text{toe}}(t)$ to overapproximate $\mathcal{L}_{\text{heel}}(t; K)$ and $\mathcal{L}_{\text{toe}}(t; K)$. Recalling that l_{th} , l_{sk} , l_{hl} and l_{toe} define the necessary lengths of our kinematic model, we may treat these as polynomial zonotopes with a center and no generators. Then, we compute $V_{\text{heel}}(t)$ and $V_{\text{toe}}(t)$ by multiplying the polynomial zonotopes $\langle R_j \rangle_i$ by these lengths in the order that they appear within the forward occupancy maps FO_{heel} and FO_{toe} . These polynomial zonotopes are then shifted by the set of possible hip positions given by $\langle q_{yz} \rangle_i$ by taking

a Minkowski sum:

$$V_{\text{heel}}(t) = \langle q_{yz} \rangle_i \oplus \langle R_1 \rangle_i l_{\text{th}} \oplus \prod_{j=1}^2 \langle R_j \rangle_i l_{\text{sk}} \oplus \prod_{j=1}^3 \langle R_j \rangle_i R(\theta_{\text{heel}}) l_{\text{hl}} \quad (6.26)$$

$$V_{\text{toe}}(t) = \langle q_{yz} \rangle_i \oplus \langle R_1 \rangle_i l_{\text{th}} \oplus \prod_{j=1}^2 \langle R_j \rangle_i l_{\text{sk}} \oplus \prod_{j=1}^3 \langle R_j \rangle_i l_{\text{toe}}. \quad (6.27)$$

Upon close inspection, these equations are identical to the forward occupancy equations found in (6.3) and (6.4), where we have replaced all rotation matrices with polynomial zonotopes and replaced the hip position with a set of possible positions. An example of the computed $V_{\text{heel}}(t)$ and $V_{\text{toe}}(t)$ are displayed in Fig. 6.5(a).

6.5.3 Cost and Constraints for Optimization

6.5.3.1 Theory

The goal of this work is to plan and execute prosthesis trajectories that robustly respond to trips. Importantly, we want to choose trajectory parameters k which lead to successful recovery across a range of plausible human trajectories. We formulate this as an optimization problem with a cost corresponding to foot placement and a constraint corresponding to undesired foot scuffs. The objective minimizes the distance between the reachable set of the swing heel/toe at touchdown corresponding to k and the target set X_T . The constraint prevents the prosthetic heel/toe from catching on the ground during swing. To be clear about the strategy-dependent role of the heel and toe in this placement cost and clearance constraint, we let $\mathcal{L}_{\text{placement}}(t; k) = \mathcal{L}_{\text{heel}}(t; k)$ and $\mathcal{L}_{\text{clearance}}(t; k) = \mathcal{L}_{\text{toe}}(t; k)$ for the *elevating* strategy, and let $\mathcal{L}_{\text{placement}}(t; k) = \mathcal{L}_{\text{toe}}(t; k)$ and $\mathcal{L}_{\text{clearance}}(t; k) = \mathcal{L}_{\text{heel}}(t; k)$ for the *delayed lowering* and *lowering* strategies.

Because the foot contact time t_f depends on both the reachable sets and choice of trajectory parameters, we let $t_{f, \text{max}}(k)$ be the last possible contact time for a given choice of trajectory parameters. As described in Sec. 6.4.3, we represent the target set using the swing hip AP position and velocity along with the foot AP position. We want to represent the set of “good” foot placement positions, which are valid for all reachable hip AP positions and velocities:

$$X_{\text{in}}(t) = \{x \in \mathbb{R} \mid (y(t), \dot{y}(t), x) \in X_T \forall y(t) \in \mathcal{I}_y(t), \dot{y}(t) \in \mathcal{I}_{\dot{y}}(t)\}. \quad (6.28)$$

We then define the subset of anterior-posterior positions of $\mathcal{L}_{\text{placement}}(t; k)$ that lie outside of the target set at any phase value where ground contact is possible and the associated phase values as

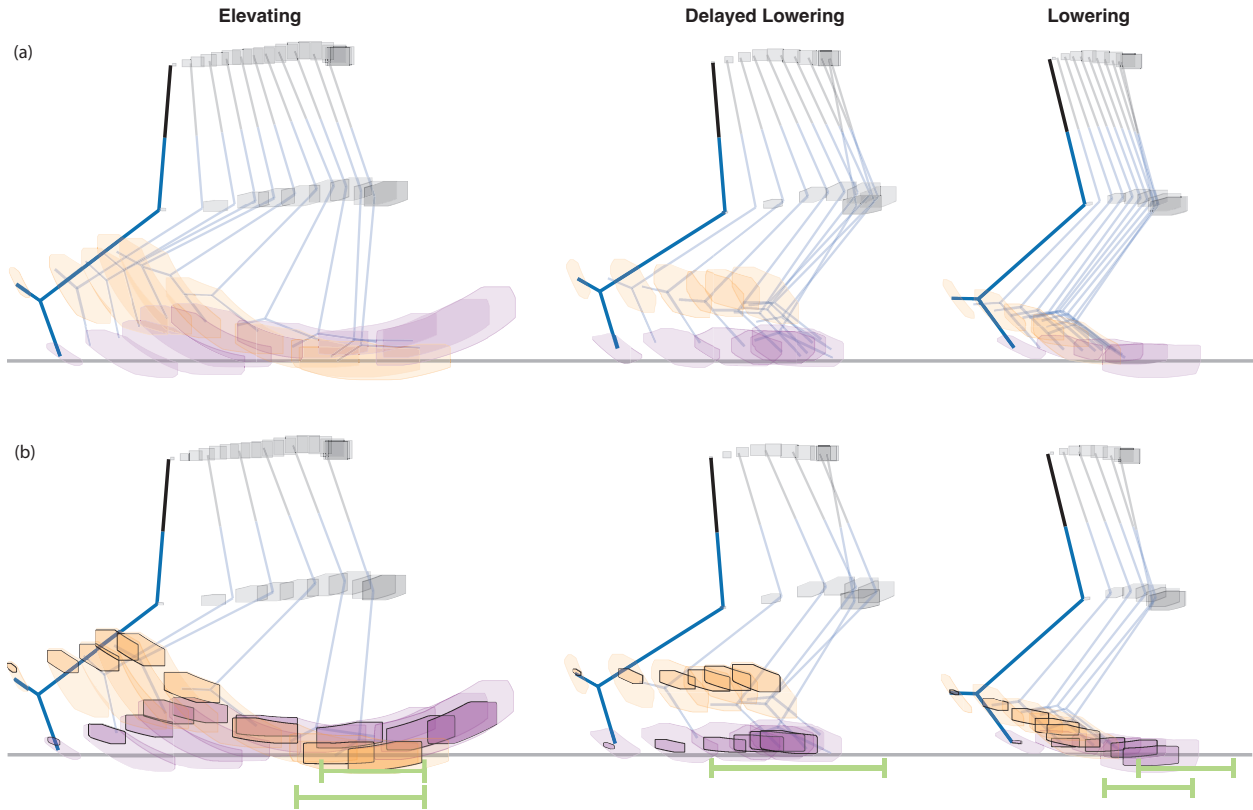


Figure 6.5: This figure shows computed reachable sets for the human-prosthesis subsystem. The elevating trial is from Subject ID 8, the delayed lowering from Subject ID 1, and the lowering from Subject ID 10. The stick figures show the subjects' ground truth trip-recovery behavior. The reachable sets are plotted for every seven phase steps, which cause the gaps at the beginning when the leg swings forward quickly. Part (a) shows the reachable sets computed before the cost and constraints are formed for optimization. The human subsystem's reachable sets are plotted in gray. Note that because the human's hip flexion-extension angle influences the position of the knee, the knee position reachable sets are also plotted in gray. The heel and toe reachable sets $V_{\text{heel}}(t)$ and $V_{\text{toe}}(t)$, computed in Sec. 6.5.2, are shown in light orange and light purple, respectively. They describe all possible locations of the heel and toe throughout swing phase given the predicted sets of human motion and the range of parameterized prosthesis behavior. Part (b) shows the optimization results for each trial (Secs. 6.5.3 and 6.5.4). The dark orange and purple sets represent the possible heel and toe locations for a choice of $k \in K$ and the entire set of predicted human trip-recovery behavior. The optimization picks $k \in K$ to avoid undesired foot scuffs (toe scuffs for elevating trials, heel scuffs for delayed lowering and lowering) and to minimize the distance outside the target sets (shown as the green intervals) any time ground contact is possible (heel for elevating, toe for delayed lowering and lowering).

$X_{\text{out}}(t; k)$:

$$X_{\text{out}}(t; k) = \left\{ x \in \pi_y(X) \mid X = \mathcal{L}_{\text{placement}}(t; k), \pi_z(X) \cap \{0\} \neq \emptyset, x \notin X_{\text{in}}(t) \right\}, \quad (6.29)$$

where the last condition implies a foot placement x is not valid for all reachable hip AP positions and velocities. Then, we define the objective function as the sum of X_{out} values:

$$f_{\text{placement}}(k) = \int_{t_{\text{trip}}}^{t_{f,\text{max}}(k)} \int_{X_{\text{out}}(t;k)} dx dt. \quad (6.30)$$

For the constraint, given $\mathcal{L}_{\text{clearance}}$, let K_u be the set of *unsafe* trajectory parameters for the prosthesis which is defined as follows.

$$K_u = \left\{ k \in K \mid t \in [t_{\text{trip}}, t_{f,\text{max}}(k)], X = \mathcal{L}_{\text{clearance}}(t; k), \exists z \in \pi_z(X) \text{ s.t. } z \leq 0 \right\}. \quad (6.31)$$

In short, K_u says that a trajectory parameter k is unsafe if the clearance reachable set can intersect the ground at any time *prior to contact* by the desired portion of the foot.

6.5.3.2 Implementation

We form the cost $f_{\text{placement}}$ and overapproximate K_u using a constraint function $h_{\text{clearance}}$. In the cost and constraint, we consider *slices* of $V_{\text{heel}}(t)$ and $V_{\text{toe}}(t)$ for $i \in \{1, \dots, n\}$ corresponding to particular choices of k (See Sec. 6.5.2 for the definition of slicing). As before, we let $V_{\text{placement}}(t) = V_{\text{heel}}(t)$ and $V_{\text{clearance}}(t) = V_{\text{toe}}(t)$ for the *elevating* strategy, and let $V_{\text{placement}}(t) = V_{\text{toe}}(t)$ and $V_{\text{clearance}}(t) = V_{\text{heel}}(t)$ for the *delayed lowering* and *lowering* strategies.

First, we discuss evaluating the objective function for desired foot placement. We want to ensure that whenever the foot can make contact before the ground (up until the maximum phase t_{ub}^s for strategy s), the distance between the foot placement location and the target set X_T (Def. 6.4.4) is minimized. This means that although we do not know exactly when the heel or toe will make contact with the ground, we can enforce its appropriate placement, as described by X_T .

For any $k \in K$, slicing $V_{\text{placement}}(t)$ by k returns the subset $V_{\text{placement, slc}}(t) \subset \mathbb{R}^3$. We can find the lowest point in this subset by projecting it onto the z -axis, then taking the infimum of the result. We define a function $h_{z,\text{min}} : \{1, \dots, n\} \times K \rightarrow \mathbb{R}$:

$$h_{z,\text{min}}(i, k) = \inf(\pi_z(V_{\text{placement, slc}}(t))) \quad (6.32)$$

where \inf is defined in the Appx. A.4 (A.26). Notice that $h_{z,\text{min}}(i, k) > 0$ for a given i and k implies that the sliced $V_{\text{placement, slc}}(t)$ is completely *above* ground. Therefore, contact with the

ground is possible for a given i and k whenever $h_{z,\min}(i, k) \leq 0$. At a given phase index i , the reachable hip AP positions and velocities determine the subset of X_T under consideration, which limits possible foot placement positions as represented by $X_{\text{in}}(t)$ in (6.28). We define the set $V_{\text{in}}(i)$:

$$V_{\text{in}}(i) = \left\{ x \in \mathbb{R} \mid (y, \dot{y}, x) \in X_T, \forall y \in \langle q_y \rangle_i, \dot{y} \in \langle q_{\dot{y}} \rangle_i \right\}, \quad (6.33)$$

which can be computed using set difference operations on polytopes. As in (6.29), we want to describe the set of undesirable reachable foot placement positions, which we then use to formulate our cost function. We do not know exactly where the foot will be placed, due to the overapproximative method for constructing reachable sets. To achieve a successful recovery, we want *all possible* foot placement positions to be close to the target set whenever ground contact is possible. To implement this, we consider the set of points that can contact the ground but lie outside $V_{\text{in}}(i)$:

$$V_{\text{out}}(i, k) = \left\{ x \in \mathbb{R} \mid x \in \pi_y(V_{\text{placement, slc}}(t)), \right. \quad (6.34)$$

$$h_{z,\min}(i, k) \leq 0,$$

$$\left. x \notin V_{\text{in}}(i) \right\}. \quad (6.35)$$

As in (6.30), we wish to approximate $\int_{V_{\text{out}}(i;k)}$, which measures the size of the set of unsafe foot positions. The set $\pi_y(V_{\text{placement, slc}}(t))$ can be overapproximated by a 1-D interval, and we represent $V_{\text{in}}(i)$ as a 1-D interval. When $h_{z,\min}(i, k) \leq 0$, we let $\int_{V_{\text{out}}(i;k)} dx$ be the size of $\pi_y(V_{\text{placement, slc}}(t))$ which does not overlap $V_{\text{in}}(i)$. For an interval $I \subset \mathbb{R}$, let

$$\sigma(I) = \sup(I) - \inf(I), \quad (6.36)$$

where $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ measures the range of an interval, and the \sup and \inf operations are introduced in (A.4) and (A.3) for intervals. Ideally, the possible foot placements $\pi_y(V_{\text{placement, slc}}(t))$ should be completely contained in $V_{\text{in}}(i)$, in which $V_{\text{out}}(i, k)$ is empty. Similarly, if $h_{z,\min}(i, k) > 0$, ground contact is impossible and $V_{\text{out}}(i, k)$ is empty. However, when $h_{z,\min}(i, k) \leq 0$, we measure the size of the set lying outside the target set as:

$$\int_{V_{\text{out}}(i;k)} dx = \sigma(\pi_y(V_{\text{placement, slc}}(t))) - \sigma(\pi_y(V_{\text{placement, slc}}(t)) \cap V_{\text{in}}(i)). \quad (6.37)$$

We then define $f_{\text{placement}}$ as the sum of maximum distances over all i , for $i \in \{l, \dots, m\} \subset$

$\{1, \dots, n\}$, where ground contact is possible:

$$f_{\text{placement}}(k) = \sum_{i=1}^m \left(\int_{V_{\text{out}}(i;k)} dx \right). \quad (6.38)$$

Recall that in delayed lowering and lowering trials, the toe contacts the ground at the end of the step. However, if the trip occurs early in the swing phase, the toe can be near the ground at the start of the trip as well, causing $h_{z,\min}(i, k) \leq 0$. To ensure that we are not trying to have the toe in X_T at the *beginning* of the swing phase, we do not enforce the cost for the first 30% of the planning phase, i.e. for any $i < 3n$. We chose 30% because we observed that subjects' computed reachable sets had cleared the ground by this percentage of swing phase.

For the heel/toe clearance constraint, we describe our representation of $K_{\text{clearance}}$, which identifies trajectory parameters which may cause unwanted collision between the foot and the ground. We define a similar function for the clearance height, $h_{z,\text{clearance},\min} : \{1, \dots, n\} \times K \rightarrow \mathbb{R}$

$$h_{z,\text{clearance},\min}(i, k) = \pi_z(\inf(V_{\text{clearance, slc}}(t))). \quad (6.39)$$

where $h_{z,\text{clearance},\min}(i, k) \leq 0$ for a given i and k implies that the sliced $V_{\text{clearance}}(t)$ *may* intersect with the ground. Therefore, letting

$$h_{\text{clearance}}(i, k) = -h_{z,\text{clearance},\min}(i, k) \quad (6.40)$$

and setting the condition $h_{\text{clearance}}(i, k) < 0$, we can ensure that the sliced $V_{\text{clearance}}(t)$ never intersects with the ground.

Lemma 6.5.2. *If $k \in K_{\text{clearance}}$, then there exists $t \leq t_{f,\max}(k)$ such that $h_{\text{clearance}}(i, k) \geq 0$.*

Recall that for elevating trials, we want to avoid toe scuffs for the entire step. If the trip occurs early enough in the swing phase where the toe is still on or near the ground at the beginning of the step, the constraint can return infeasible when in reality, toe contact is acceptable. We therefore do not enforce the foot scuff constraint for the first 30% of the planning phase, i.e. for any $i < 3n$.

6.5.4 Real-Time Optimization

6.5.4.1 Theory

We solve for the optimal trajectory parameter k as

$$k_{\text{opt}} = \operatorname{argmin}_{k \in K} \{ f_{\text{placement}}(k) \mid k \notin K_u \}. \quad (6.41)$$

Note that the cost function could be modified to include other user-specified objectives, such as minimizing the distance from a desired trajectory or the work done by the prosthesis. Unfortunately, the optimization (6.41) may not be feasible if every trajectory parameter is unsafe, i.e. $K = K_u$. In this case, the avoidance of foot scuffs can not be ensured, and our framework defaults to using the nominal trajectory in this case as described in Sec. 6.3.3.

Assuming the problem (6.41) was feasible and an optimal trajectory parameter has been identified, the control input for the prosthetic knee at time t is given by a tracking controller

$$u_p(t) = G_1(q_p(t) - q_{p,\text{des}}(t; k)) + G_2(\dot{q}_p(t) - \dot{q}_{p,\text{des}}(t; k)), \quad (6.42)$$

where G_1 and G_2 are user specified feedback gains, and we emphasize that the trajectory parameters produce desired joint trajectories to be tracked by the feedback controller.

6.5.4.2 Implementation

We implement (6.41) as a nonlinear program:

$$\begin{aligned} k_{\text{opt}} = \underset{k \in K}{\text{argmin}} \quad & f_{\text{placement}}(k) \\ \text{s.t.} \quad & h_{\text{clearance}}(i, k) < 0 \quad \forall i \in \{1, \dots, n\}. \end{aligned} \quad (6.43)$$

The plan lasts for the phase horizon $[t_{\text{trip}}, t_{\text{ub}}^s]$ for strategy $s \in \{1, 2, 3\}$ corresponding to elevating, delayed lowering, and lowering, respectively. As detailed in Sec. 6.3.3, the nominal prosthesis controller is assumed to resume control upon contact with the ground. Our proposed framework correctly accounts for the initial position of the prosthetic leg at the time of the trip, and utilizes reachable sets of both human and prosthetic subsystems to form the cost and constraint. If no safe trajectory can be found, our framework defaults to using the nominal controller. Example optimization results for each strategy are shown in Fig. 6.5(b).

6.6 Evaluation

6.6.1 Evaluating the Proposed Planning Framework

If the optimization problem in Sec. 6.5.4 is feasible, our framework produces k_2 and k_3 that correspond to knee and ankle angle trajectories, respectively, from $[t_{\text{trip}}, t_{\text{ub}}^s]$ for strategy s . These k -values avoid foot scuffs (toe scuffs for elevating, heel scuffs for delayed lowering and lowering) for the entire set of predicted hip motion, and aim to minimize the distance between the horizontal foot location (heel for elevating, toe for delayed lowering and lowering) and the target set at any

phase point where foot contact is possible. However, the reachable sets of q_h are predictions of human hip motion, so they may not actually encompass the human’s trip-recovery swing hip behavior. Recall that the human’s trip-recovery response lasts until $t_f < t_{ub}^s$. We evaluate the framework’s performance by computing FO_{heel} and FO_{toe} using the ground truth hip motion for q_h and the planning framework’s q_2 and q_3 from $[t_{trip}, t_f]$. To be clear, we evaluate trip-recovery planner using ground truth hip states from our able-bodied experiment described in Sec. 6.2 and simulated prosthesis knee and ankle states:

Assumption 6.6.1. *The prosthesis motion does not affect the human hip motion after a trip.*

This assumption is reasonable because in many trips, the recovery time is short enough that humans may not have time to react to prosthesis behavior.

In the simulated trip recovery, we first check if the foot (toe for elevating, heel for delayed lowering and lowering) scuffs the ground at any point from $[t_{trip}, t_f]$. Next, we look at the part of the foot where contact is desired at touchdown (heel strike for elevating, toe strike for delayed lowering and lowering). Because the prosthesis behavior is likely different from the human’s ground truth trip-recovery knee/ankle behavior, the phase value at touchdown is different from t_f . Let $t_{f,sim}$ denote the phase value at which the simulated heel/toe crosses the ground plane. We first find the ground truth hip states $q_h(t_{f,sim})$ and evaluate X_T at these states to produce a range of desired heel/toe strike locations. Next, we compare the anterior-posterior heel/toe strike location to the range of target set values. If the heel/toe strike lies inside the target set, we set the distance to the target set, d_{X_T} , to zero. Otherwise, we set d_{X_T} as the distance from the heel/toe strike point to the closest target set bound.

In our evaluation, we observed some trials where the foot (heel for elevating, toe for delayed lowering) came close to the ground at the end of the step, but did not contact the ground during $t \in [t_{trip}, t_{f,sim}]$. We would expect in this case that the human would continue moving the swing hip until touchdown occurs, but in our simulation, we only have human trip-recovery data from $[t_{trip}, t_f]$. We therefore compute the desired foot placement at t_f and evaluate d_{X_T} using the swing heel/toe’s horizontal position at $t = t_f$. In other trials, we were unable to evaluate a target set because the $q_h(t_{f,sim})$ states were outside of the bounds formed by the target set zonotope. This problem typically occurred when the foot contacted the ground much earlier than observed in any ground truth trials. For these cases, we found the first instance where the target set was defined for the trial, and evaluated the target set distance using this desired foot contact range.

6.6.2 Comparison to Impedance-Based Control

Without a trip-recovery controller, the prosthesis is likely to continue using its nominal level-ground walking controller after a trip occurs. We therefore compare our proposed trip-recovery

framework to a nominal impedance-based controller. In this control scheme, we divide the swing phase $t \in [0, 1]$ into $n_m = 4$ modes for each joint $j \in \{2, 3\}$. Each mode $i \in \{1, \dots, n_m\}$ has a phase interval $\tau_j^i = [t_{0,j}^i, t_{f,j}^i]$ of length $(t_{f,j}^i - t_{0,j}^i)$. Note that $t_0^1 = 0$, $t_f^{n_m} = 1$, and $t_0^i = t_f^{i-1}$ for $i = 2, \dots, n_m$. In each mode, we model the prosthetic joint j for $j \in \{2, 3\}$ as a spring-damper system with stiffness $\lambda_j^i \in \mathbb{R}^+$, damping coefficient $\alpha_j^i \in \mathbb{R}^+$, and equilibrium angle $\theta_{eq,j}^i \in \mathbb{R}$. Given prosthetic joint angular position and angular velocity $q_j, \dot{q}_j(t)$ at phase value $t \in \tau_j^i$, the desired joint torque for joint j in mode i is expressed as

$$T_{des,j}^i(t) = -\lambda_j^i(q_j(t) - \theta_{eq,j}^i) - \alpha_j^i \dot{q}_j(t). \quad (6.44)$$

To determine λ_j^i , α_j^i , and $\theta_{eq,j}^i$, we first compute a reference average nominal torque profile $T_{ref,j}^i : \tau_j^i \rightarrow \mathbb{R}$ via inverse dynamics. We define nominal swing phase steps as those that are three or more steps past the tripped step, and compute average nominal trajectories of the knee and ankle angles for $t \in [0, 1]$ for each subject. Next, we adopt the dynamic model described in Sec. 6.3.1 on this averaged nominal trajectory. Using the Recursive Newton-Euler algorithm, we solve the sequence $T_{ref,j}^i$ on each mode interval τ_j^i . We allow the mode intervals to be distinct for each joint j , but keep the same mode number n_m for both joints (i.e. $\tau_j^1, \dots, \tau_j^{n_m}$), with each $\tau_j^i = [t_{0,j}^i, t_{f,j}^i]$ for joint j and mode $i \in \{1, \dots, n_m\}$. Next, for each joint j , we solve for the intervals $\tau_j^1, \dots, \tau_j^{n_m}$ and the λ_j^i , α_j^i , and $\theta_{eq,j}^i$ of each mode that minimize the sum of the squared error between $T_{des,j}^i$ and $T_{ref,j}^i$. We formulate the problem as a bilevel optimization program:

$$\begin{aligned} & \min_{\tau_j^1, \dots, \tau_j^{n_m}} \sum_{i=1}^{n_m} \delta_j^i \\ \text{s.t. } & \delta_j^i = \min_{\lambda_j^i \in \mathbb{R}^+, \theta_{eq,j}^i \in \mathbb{R}^+, \alpha_j^i \in \mathbb{R}} \{SSE(T_{ref,j}^i, T_{des,j}^i)\} \\ & t_{0,j}^i, t_{f,j}^i \in [0, 1], \quad i = 1, \dots, n_m \\ & t_{0,j}^1 = 0, \\ & t_{f,j}^{n_m} = 1, \\ & t_{0,j}^i = t_{f,j}^{i-1}, \quad i = 2, \dots, n_m \\ & t_{0,j}^i < t_{f,j}^i, \quad i = 1, \dots, n_m \end{aligned} \quad (6.45)$$

where SSE calculates the integral of the squared difference of the two input functions on $[t_{0,j}^i, t_{f,j}^i]$. Notice that the above optimization program includes an outer optimization problem and a inner optimization problem. The outer optimization problem tries to find the best division of modes $[\tau_1^i, \tau_2^i, \dots, \tau_{n_m}^i]$ that minimizes the cost function, while the inner optimization tries to find the best system parameters for each mode. The inner optimization problem is a constrained linear

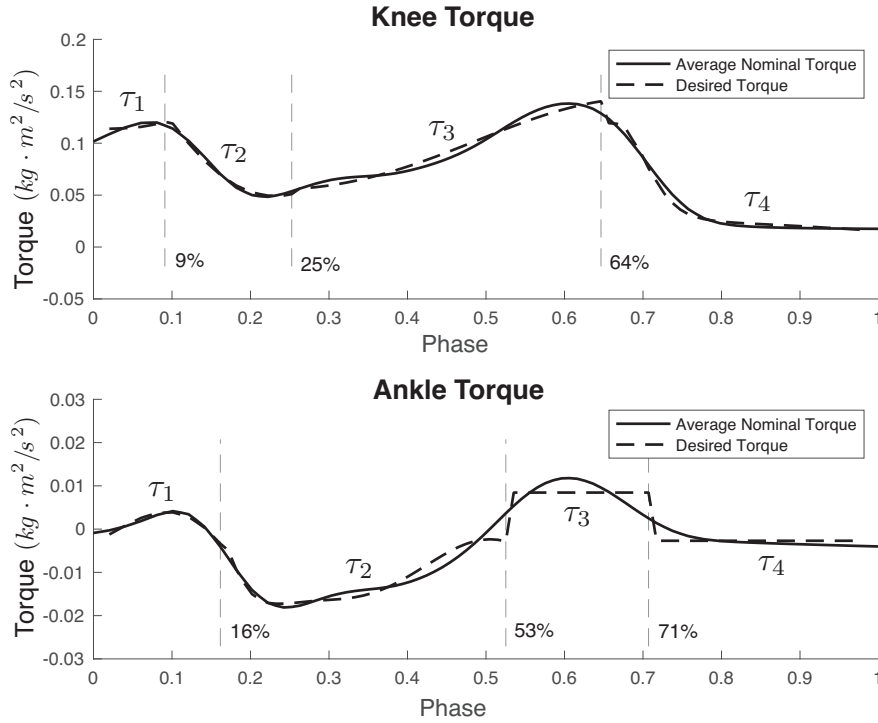


Figure 6.6: This figure shows example knee and ankle torques calculated from an average nominal trajectory T_{ref} via inverse dynamics (black solid curve), and the desired knee and ankle torques produced by the nominal impedance-based controller T_{des} (black dashed curve) for Subject ID 1. The number of modes is set to $n_m = 4$ in this case. Note that for the knee and ankle torque profiles, we use different mode divisions and parameters, as shown in (a) where the divisions are set at 9%, 25% and 64%, and in (b) where the divisions are set at 16%, 53% and 71%.

least-squares problem, so we use MATLAB's `lsqmin` function to solve the optimal parameters $\lambda_j^i, \theta_{eq,j}^i, \alpha_j^i$. Plugging in the inner loop optimal parameters, we solve the outer loop minimization problem with a genetic algorithm, using MATLAB's `ga` function. Fig. 6.6 shows average nominal torque T_{ref} and the desired torque T_{des} for Subject ID 1.

Given initial states $q_{0,j}$ and $\dot{q}_{0,j}$ at $t = t_{trip}$, we simulate prosthesis motion iteratively by first obtaining the desired torque from (6.44) and then applying it in the dynamic model in Sec. 6.3.1 to obtain the joint's angular position and velocity at the following phase value. Once the nominal prosthesis knee and ankle angles are found for $t \in [t_{trip}, t_f]$, we simulate the prosthetic joint motion with the human's ground truth hip trip-recovery behavior. We then conduct the same evaluation steps using the subject's ground truth hip trajectories and the nominal knee and ankle angle trajectories. In Sec. 6.3.3, we stated that if our framework fails to produce a feasible solution, the prosthesis defaults to its nominal controller. To avoid comparing the nominal controller to itself, we only run the evaluation for trials where our framework returns a feasible solution.

6.7 Results

The evaluation results are reported for each subject in Tab. 6.1. Fig. 6.7 shows comparisons of subjects' ground truth trip-recovery behavior, the proposed framework's result, and the impedance-based result for three examples.

Currently, our framework's planning time is on the order of seconds, which is much too long for planning a trip recovery. However, the long planning times result from a recent change in computing reachable sets. This new computation method helps reduce conservatism, ensuring the reachable sets of the heel and toe are not overapproximated by too large a buffer. If the reachable sets are too large, the method returns infeasible solutions more often. Although the computational changes improved feasibility, we are still in the process of reducing computation time. Past versions of our method implemented on a simulated trip dataset had a mean planning time of 94 ms when a feasible trajectory was found [109].

6.8 Discussion

Prosthetic trip-recovery controllers could help people with lower-limb amputations perform everyday tasks with more confidence. However, the uncertainty in estimating human behavior has made designing these controllers difficult. In this chapter, we present a method for planning trip-recovery trajectories online while explicitly accounting for a range of predicted human responses. Our framework constructs reachable sets of the prosthetic heel and toe that depend on both the human and prosthesis subsystems, then selects trajectory parameters for the prosthetic knee and ankle online to ensure that the entire system recovers successfully. To our knowledge, this work represents the first prosthetic controller framework that can plan trajectories online while accounting for a range of possible human behaviors.

Our proposed framework returned a feasible solution to the optimization program in Sec. 6.5.4 for 87.2% of trials across all 16 subjects. Looking at Tab. 6.1, we see that many subjects have feasibility rates above 95%, but the subject-wide average is brought down by some with lower rates, such as Subject ID 4. In investigating these low feasibility rates for some subjects, we observed that the low feasibility is often specific to one strategy (e.g., delayed lowering). In those cases, the subject's trip-recovery library for that specific strategy likely does not provide a large enough range of knee or ankle angles to lead to feasible solutions. We aim to generate feasible trip-recovery solutions as often as possible to avoid defaulting to the nominal impedance-based controller during a trip, so future work will investigate and address problems with the trip-recovery libraries for these subjects.

For the 834 trials where a feasible solution was returned, our proposed framework outperformed

Table 6.1: Results for our proposed framework compared to nominal impedance-based control.

Subject ID	Method	Feasible	Undesired Foot Scuffs	Mean d_{X_T} (cm)
0	Proposed	41/41 - 100.0%	0/41 - 0.0%	0.7
	Impedance-Based	N/A	25/41 - 61.0%	7.4
1	Proposed	75/77 - 97.4%	0/75 - 0.0%	2.6
	Impedance-Based	N/A	34/75 - 45.3	3.5
2	Proposed	29/39 - 74.4%	2/29 - 6.9%	1.4
	Impedance-Based	N/A	20/29 - 69.0%	11.9
3	Proposed	40/57 - 70.2%	1/40 - 2.5%	1.8
	Impedance-Based	N/A	24/40 - 60.0%	10.0
4	Proposed	26/42 - 62.0%	0/26 - 0.0%	1.2
	Impedance-Based	N/A	18/26 - 69.2%	11.5
5	Proposed	53/53 - 100.0%	0/53 - 0.0%	2.1
	Impedance-Based	N/A	7/53 - 13.2%	2.4
6	Proposed	65/76 - 85.5%	0/65 - 0.0%	9.1
	Impedance-Based	N/A	11/65 - 16.9%	4.9
7	Proposed	32/43 - 74.4%	1/32 - 3.1%	12.8
	Impedance-Based	N/A	17/32 - 53.1%	14.4
8	Proposed	70/73 - 95.9%	1/70 - 1.4%	1.4
	Impedance-Based	N/A	41/70 - 58.6%	8.8
9	Proposed	58/73 - 79.5%	0/58 - 0.0%	4.6
	Impedance-Based	N/A	8/58 - 13.8%	14.3
10	Proposed	64/72 - 88.9%	0/64 - 0.0%	0.6
	Impedance-Based	N/A	33/64 - 51.6%	8.9
11	Proposed	54/75 - 72.0%	0/54 - 0.0%	2.6
	Impedance-Based	N/A	33/54 - 61.1%	4.5
12	Proposed	68/71 - 95.8%	0/68 - 0.0%	0.8
	Impedance-Based	N/A	28/68 - 41.1%	6.7
13	Proposed	51/54 - 94.4%	0/51 - 0.0%	12.1
	Impedance-Based	N/A	28/51 - 54.9%	21.0
14	Proposed	61/63 - 96.8%	0/61 - 0.0%	1.6
	Impedance-Based	N/A	16/61 - 26.2%	7.6
15	Proposed	47/48 - 97.9%	1/47 - 2.1%	1.9
	Impedance-Based	N/A	24/47 - 51.1%	1.9
Total	Proposed	834/957 - 87.2%	6/834 - 0.7%	3.4
	Impedance-Based	N/A	367/834 - 44.0%	8.2

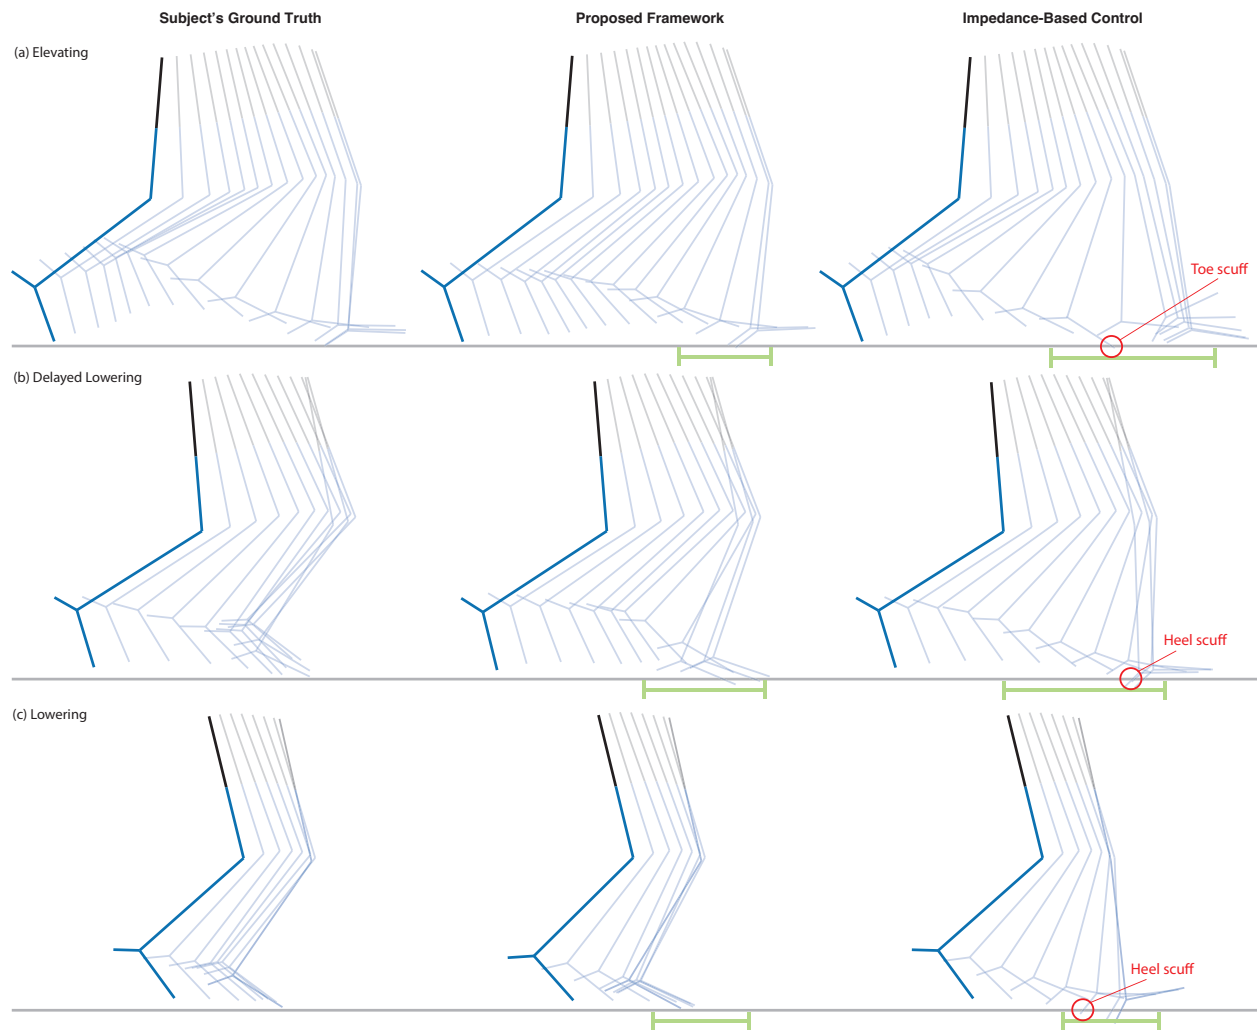


Figure 6.7: This figure shows results for three sample trials. Row (a) shows an elevating trial from Subject ID 8. The first column depicts the subject’s observed trip-recovery behavior. The second column shows the optimization result from our proposed framework, evaluated using the subject’s ground truth hip motion as described in Sec. 6.6. At the phase point where the heel contacts the ground, we determine the desired range of heel strike locations using the hip’s AP position and velocity. The desired range is pictured as the green interval below the ground. Note that our proposed framework successfully places the heel in this range. In the third column, we see that the impedance-based controller also places the heel in the desired heel strike range, but that the toe scuffs the ground along the way. Rows (b) and (c) depict delayed lowering and lowering trials for Subject IDs 1 and 10, respectively. Note that for these two strategies, we aim to avoid heel scuffs and to place the toe in the desired range at the end of the step. In the middle column, our proposed framework shows the toe contacting the ground within the green range for both delayed lowering and lowering. In both (b) and (c), the impedance-based controller scuffs the heel on the ground, and the toe AP position lies outside the desired range.

nominal impedance-based control in avoiding unsafe foot scuffs and placing the swing foot in a desired location at touchdown. The low percentage of foot scuffs (0.7%) that occurred in our planning framework highlights the effective performance of two aspects. First, the GPR models produced accurate predictions of human swing hip behavior. The six trials where foot scuffs did occur resulted from the swing hip motion exiting the 60% confidence bounds. Increasing the confidence interval could reduce the foot scuff percentage further; however, it would also cause larger heel and toe reachable sets, which may return more infeasible trials. Second, the parameterized trip-recovery libraries provided an effective range of knee and ankle trajectories that could avoid foot scuffs for the predicted set of hip behavior.

Of the feasible trials, our framework placed the foot closer on average to the target set compared to impedance-based control. However, for some subjects, especially Subject IDs 7 and 13, the mean distance from the target set is relatively large (over 12 cm). Because we are solving a nonlinear optimization problem, it is possible that the solver reached a local minimum that produced a high distance from the target set in these cases. Additionally, these subjects' trip-recovery libraries could have been unable to produce knee and ankle trajectories that could get closer to the target set, and tuning the libraries could address the issue.

In this work, we implemented the framework using swing hip data from able-bodied subjects attached to a robotic knee-ankle prosthesis in simulation. Despite the promise of this framework for planning safe trip recoveries online, several challenges must be addressed for the framework's successful implementation on prosthesis hardware. The assumptions made in this chapter are based on successfully implemented studies by other research groups (e.g., stumble detection work by [121, 122] and trip-recovery strategy selection work by [96]), but have not been tested with our trip dataset or implemented on our own hardware. Our predictions of swing hip responses are informed by data from able-bodied subjects, and we assume that the human swing hip motion remains unchanged throughout the swing phase regardless of prosthesis behavior. Although people with lower-limb amputation have displayed similar trip-recovery strategies to able-bodied subjects [19], the swing hip motion may be different. Furthermore, distinct prosthesis trip recovery behavior (e.g., the joint angles produced by nominal control compared to online optimization) could elicit different responses from individuals, especially for trips at the beginning of the swing phase when there is more time to react and plan. One way to address this challenge is to collect swing hip data while the trip-recovery framework is in use and continually re-train the GPR models with observed data. The current framework also assumes that our implemented target sets in Sec. 6.4.3 are sufficient for the system's successful trip recovery. However, the proposed framework can include other notions of target sets, informed by human stumble experiments or physics-based definitions of dynamic stability.

The theoretical trip-recovery planning framework presented here can generalize to walking

models with different dimensions, parameters and dynamics. The method could be applied to other types of trips or slips, regardless of whether the prosthesis is in stance or swing. Our optimization objective only considered the placement of the swing foot at the end of the tripped step, but the cost could also include goals such as minimizing the work done by the prosthesis. Our future work will generalize the framework in each of these areas with the goal of incorporating this framework into prosthetic hardware. The simulated prosthesis results with human trip-recovery data represent a promising step in developing prosthesis controllers that can increase confidence and safety by working alongside a human to recover from trips.

CHAPTER 7

Specific Aim 5: Towards a Technology Assessment for Robotic-Limb Prostheses

This work was completed as coursework and independent research in the University of Michigan Ford School of Public Policy's Science, Technology, and Public Policy certificate program, advised by Prof. Shobita Parthasarathy. It is written for a non-technical audience.

7.1 Introduction

An estimated one million people were living with lower-limb loss in the United States in 2005, and this number is expected to more than double by 2050 [5]. The most common cause of limb loss is vascular disease (e.g., diabetes or peripheral arterial disease), but other causes include congenital limb difference, traumatic injury, and cancer [134, 135]. The projected increase in limb loss reflects the growing prevalence of dysvascular disease in the United States [5, 136].

To maintain mobility, lower-limb amputees typically use a prosthesis and/or a wheelchair [137]. The decision of whether to use a prosthesis and the duration of daily prosthesis use depend on a number of factors such as age, gender, fitness level, amputation level, and satisfaction with prosthesis fit [138, 139]. Researchers have attempted to estimate percentages of prosthesis use among amputees, but obtaining sample populations that vary across factors is challenging. Generally, studies indicate a high percentage of prosthesis use, but with low overall satisfaction [140].

Many research groups across the academic, public, and private sectors focus on the design and control of *robotic* prostheses, which can more closely mimic the capabilities of a human joint than prostheses used by the participants in the surveys above. However, these devices are far from widespread adoption; most are not FDA approved [20]. As researchers continue to advance robotic prostheses in laboratories, they must think critically about the technology's impact on those with lower-limb loss to ensure its effective translation to the real world.

7.2 Background

7.2.1 Some key terms and concepts

A typical way to evaluate a lower-limb prosthesis is by studying its effect on an individual's walking gait. The *gait cycle*, shown in Fig. 7.1, describes one leg's transition from the *stance phase*, where weight is put on the leg, to the *swing phase*, where the leg lifts off the ground. To accomplish each stride in the gait cycle, a human's muscles produce *joint torques*, or rotational forces about the hip, knee, and ankle.




Figure 7.1: The gait cycle follows one leg through the stance and swing phases. Image adapted from T. Stöckel et al. 2015.

Lower-limb amputation's impact on mobility, cost, and health varies based on the characteristics of the limb amputation. A *transfemoral amputation* crosses the femur, above the knee, while a *transtibial amputation* crosses the tibia, below the knee [141]. *Hip disarticulation* refers to amputations where the femur is removed at the hip socket. A *unilateral amputation* refers to amputation on one leg, and *bilateral amputations* are on both legs. We refer to non-amputated joints and limbs as *intact*. The *residual limb* is the part of the leg remaining after limb loss, and the *socket* fits around the residual limb and attaches to the prosthesis. Typically, a *liner* holds the socket to the residual limb and provides extra comfort and protection.

In the United States, Medicare has established a system, adopted by many private insurers, called *K-levels* to determine eligibility for lower-limb prosthesis coverage [142]. The system ranges from K-0 to K-4, with K-0 indicating that the individual does not have the ability to ambulate safely with a prosthesis and K-4 indicating that the individual's ambulation skills exceed average. A prosthetist, physical therapist, or physician assigns a patient's K-level using one of several established outcome measures.

7.2.2 Three types of lower-limb prostheses

Many people with lower-limb loss use *passive* prostheses. A typical passive knee-ankle prosthesis includes a spring on the ankle joint and a damper (think of what slows down a swinging screen door) on the knee joint. Passive joints do not rotate on their own; to use this device, individuals must swing their residual limb forward to cause the lower leg to advance. The damper at the knee provides resistance as the lower leg swings forward, and the spring at the ankle provides push-off assistance as the leg leaves the ground [20]. Constant-friction knees, the simplest option, offer the same damping resistance throughout the gait cycle. More sophisticated models vary the resistance using mechanical components or compressed fluid to store and return energy [143].

Because passive devices cannot supply external power to the knee or ankle joint, individuals compensate by adjusting the behavior of their intact joints and limbs. For example, unilateral amputees often spend a longer time on their intact limb than their amputated limb in the stance phase [144]. This compensation strategy increases the forces exerted on the intact knee and hip, which leads to joint pain and osteoarthritis after years of repetitive use. Other common outcomes from compensation include back pain; osteoporosis in the residual limb due to decreased loading on the bones; and higher energy expenditure while walking.



Figure 7.2: An example of each prosthesis type. Image adapted from Ossur Total Knee, Ossur Rheo Knee, and Ossur Power Knee images.

Established in the late 1990s, *microprocessor-controlled prostheses* advance passive designs by actively modulating the amount of damping resistance provided by the knee. These devices have an

internal computer (microprocessor) that monitors the individual's gait cycle via force sensors and adjusts fluid resistance accordingly using hydraulics or pneumatics [145]. Many microprocessor-controlled prostheses are commercially available, including the Ottobock C-Leg and Ossur Rheo Knee [146, 147]. Studies generally indicate improved function and user satisfaction with these devices compared to passive prostheses [110, 148].

Robotic prostheses are designed to mimic the human neuromuscular system more closely than their predecessors. These devices have some sort of actuation (electric, pneumatic or hydraulic) at each joint and a wide variety of sensors to determine quantities such as joint angle, angular speed, and torque [20]. Just as with microprocessor-controlled prostheses, sensors allow the prosthesis to monitor the gait cycle and adjust accordingly, but robotic prostheses' motors enable a wider range of joint motion and do not require the user to supply the energy. Fig. 7.2 shows examples of commercially available passive, microprocessor-controlled, and robotic knees.

Only two robotic lower-limb prostheses are commercially available in the United States: the Ossur Power Knee and the Ottobock Empower ankle/foot [149, 150]. Because these devices span only one joint, they are categorized as Class I (low risk) medical devices in the United States Food and Drug Administration (FDA). Robotic knee-ankle prostheses, however, fall into the Class II (moderate to higher risk) category because they require coordination between two electronically connected joints [151]. To achieve Class II approval, medical devices must meet special requirements pertaining to labeling, guidance documents, and performance standards [152]. The FDA has not established such performance standards for lower-limb robotic prostheses, likely due to the absence of an analogous technology for comparison. Recent establishment of robotic lower-limb exoskeleton standards may speed up the process [153].

7.3 Why a technology assessment?

The researchers who develop robotic prostheses are eager to improve daily life for people with lower-limb loss. Their resulting prostheses, which utilize lightweight power sources, a wide variety of sensors, and artificial intelligence, have the potential to provide an unprecedented level of mobility assistance. Less often acknowledged, however, are these devices' unanticipated impacts. A recent podcast episode describes Britt Young, a woman with upper-limb loss, receiving a technologically advanced prosthetic hand: "At first, Britt was floored by her new bionic arm. . . [but] quickly realized this high-tech piece of machinery wasn't all that it was hyped up to be." The robotic hand was heavy, prone to overheating, and uncomfortable [154].

Because robotic lower-limb prostheses are largely in the development phase, now is a good time to carefully consider their potential impacts and barriers to widespread adoption. Conducting such a review is known as a *technology assessment* in Science and Technology Studies. Technology

assessments aim to “...provide an explicit mechanism for observing, critiquing, and influencing social values as they become embedded in innovations” [155]. Britt’s bionic arm, for example, reflected the engineers’ values of bionic technology and sophisticated robotic controllers, but failed to incorporate her values of comfort and functionality.

Researchers who work in laboratories that are disconnected from the amputee population they intend to assist can use this technology assessment to adjust the scope of their work or dedicate time and funding to mitigate a potential concern. A technology assessment of robotic prostheses can also benefit policymakers by, for example, recommending focus areas for FDA performance standards or identifying issues with current health insurance policies.

7.4 Current technology assessment methods

Technology assessment (TA) is a term first established by the United States Congress in the 1960s when assessing the impacts of technologies such as supersonic transport and genetic screening [156]. The U.S. Office of Technology Assessment (OTA) further standardized TAs during its existence through 1995. In these early TAs, scholars in Science and Technology Studies (STS) and experts in the technology being studied exclusively conducted the assessments. Their main goal was to provide advice to policy makers [157]. The OTA inspired many European governments to establish TA networks and institutions, such as the European Parliamentary Technology Assessment Network [156].

Since their establishment, TAs have evolved to include broader considerations and participants. Established TA categories include:

- Expert TA, a general term for TAs conducted by STS and technical experts,
- Participatory TA (pTA), a term for TAs that systematically include participants outside of the expert community, such as lay citizens, individual stakeholders, or representatives of civil society organizations [158],
- Constructive TA (CTA), which aims to influence the design stage of new technologies rather than regulatory practices [159],
- Discursive TA, which explores the language surrounding an innovation [160], and
- Health TA (HTA), a general term for TAs specifically pertaining to efficacy, safety, and cost of medical treatments, biomedical devices, and pharmaceuticals [161].

These categories may overlap for a given assessment. For example, NASA conducted a pTA to gauge lay citizen opinions about different ways to collect asteroid material [162]. This TA

was participatory due to its inclusion of lay people in the assessment process, but could also be considered constructive because it aimed to influence NASA mission design choices. Furthermore, many studies conduct assessments of new technologies without explicitly labeling them as TAs; such assessments may align with one or more of the above descriptions.

In this section, I introduce several methods for assessing technology, including some of the categories above. When evaluating the benefits and drawbacks of applying each assessment method to robotic lower-limb prosthesis technology, I ask the following questions:

- *What might I learn from implementing this method?* Each type of TA may reveal different societal outcomes based on its participants and content. A TA that centers on societal impacts of robotic lower-limb prostheses could focus on health and safety risks to amputees, barriers to widespread adoption, or ethical implications. Other TAs may consider alternatives to robotic lower-limb prostheses, such as increasing accessibility in the built environment.
- *Who would be represented in the assessment?* TA methods vary widely in the assessor's stakeholder involvement and level of expertise. A focus group for robotic lower-limb prostheses, for example, may specifically recruit lower-limb amputees and/or prosthetists to participate, but exclude roboticists who are aware of the most pressing technical barriers. On the other hand, both expert TA and participatory TA methods may not specifically recruit people with lower-limb loss, and risk conducting an assessment that ignores a major stakeholder.
- *What does this method enable us to do?* The information gained from a given TA method will influence in what space (academic, political) and at what point in the technology's development one can intervene to address negative consequences.
- *Does this method ignore concerns about access to robotic lower-limb prostheses?* There are several barriers to accessing technologically advanced lower-limb robotic prostheses, detailed in Sec. 7.5.1. Some assessment methods may begin with the assumption that everyone with a lower-limb amputation can access these devices, while others may include these barriers as a focus of the analysis.

7.4.1 Delphi

Delphi is a form of expert TA, originally developed in the 1950s by the RAND Corporation for the U.S. Air Force to predict the effects of future warfare [163]. Since then, public, private, and government sectors have used Delphi to gain information from a group of experts about any topic where there is imprecise or incomplete information, such as medical research or cybersecurity. A Delphi study is conducted by asking experts several rounds of questions. The participants submit

their answers anonymously and receive information and feedback from previous rounds to refine their answers. Possible aims of a Delphi study when applied to research or policy include arriving at a consensus; determining a range of possible responses or strategies; exposing underlying assumptions or values that lead to differing opinions; or educating the participants about related aspects outside of their field [164].

Applied to robotic lower-limb prostheses, a Delphi study could produce a consensus among anonymous experts about a chosen topic in the field. For example, experts could answer questions about the evaluation criteria for FDA approval of robotic knee-ankle prostheses; how to design devices for more effective clinical translation; or whether they should be focusing their research on lower-cost devices instead. In addition to technical experts (roboticists and biomechanists), the Delphi study should include clinical experts (physical therapists, occupational therapists, and prosthetists) to ensure that concerns about usability and activities of daily living are not ignored. The anonymity of the Delphi study could enable more candid responses than these experts would typically divulge in seminar talks or academic papers. Concerns about access could be central to the study or ignored, depending on the topic of discussion.

7.4.2 Participatory TAs

The pTA category broadly includes any TA with non-expert participants, but this section will focus on TAs conducted by lay people with no definitive position on the topic of discussion. Including publics¹ in a TA can foster greater trust of researchers and policy makers; clarify citizens' concerns and values with respect to a given technology; and gain insight by including non-traditional knowledge and skills in the conversation [165, 166].

Citizens' juries, also called citizens' panels, were concurrently established in the U.S. and Germany in the 1970s and have spread to many other countries, including Australia, India, the UK, and Brazil [167]. These juries typically consist of 12-20 lay citizens that meet over the course of several days to deliberate on a particular issue. The participants are given detailed background information and often meet with a panel of experts to ask questions. Then, after small- and full-group discussion and brainstorming, the panel releases a report that contains its recommendations [165]. Consensus conferences are similar to citizen panels but are typically smaller (10-15 participants) and place greater emphasis on the points of agreement between participants [168].

The Expert and Citizen Assessment of Science and Technology (ECAST) network, established in 2011, has developed a scalable and generalizable framework for conducting lay citizen TAs [169]. In this ECAST pTA method, 80-100 members of the public participate in a day-long workshop to deliberate a specific technology. The facilitators provide participants with background

¹The plural term "publics" highlights that citizens have diverse identities and perspectives, in contrast to "the general public."

material on the technology being discussed, then divide them into small groups of 6-8 people to complete a series of activities. Each group provides written rationale for their decision or opinion, and ECAST members analyze qualitative and quantitative data to determine citizens' values and viewpoints.

A pTA could cover a range of topics about robotic lower-limb prostheses, but some areas of interest might involve lay people's impressions of accessibility or prosthesis use cases. Would society welcome extensive infrastructure projects to make the built environment more accessible? Should we focus research funding on developing high-tech robotic prostheses, or improving low-cost passive devices? Just as with the Delphi method, concerns about access to robotic lower-limb prostheses may be ignored or discussed in-depth, depending on the topic choice.

7.4.3 Focus groups: “Publics-in-Particular”

Unlike the pTA methods above, which involve deliberation among the “public-in-general,” focus groups recruit “publics-in-particular,” or certain groups of publics most likely to use, prescribe, or assist with a new technology [170]. One example of TA focus groups, conducted to determine attitudes and beliefs regarding treatment of sickle cell disease via somatic genome editing, consisted of patients, parents, and physicians [171]. The groups' conversations revealed patients' motivators for and deterrents to participating in clinical trials; the information needed for patients and parents to make informed decisions about healthcare; and advice to the gene-editing research community on how to meaningfully engage with stakeholders. Other “publics-in-particular” assessments may not establish formal focus group sessions, but rely on other methods (e.g., fieldwork interviews spanning multiple years) for identifying and framing stakeholder values [172].

A focus group about lower-limb robotic prostheses should include people with lower-limb amputation, immediate family members and/or caretakers if applicable, and prosthetists, who prescribe and fit prostheses for patients. Ideally, the focus group could address concerns about access to robotic lower-limb devices by recruiting amputees with diverse financial status, K-level, and body type. A focus group may ask specific questions about robotic prosthesis adoption, including whether an amputee would be interested in using a robotic device at all; whether a prosthetist was willing to prescribe, tune, and troubleshoot the devices; and what information amputees, family members, and caretakers need to make informed decisions about prosthesis innovations.

7.4.4 Scenario planning

Scenario planning TAs acknowledge the uncertainty surrounding a new technology by envisioning multiple futures resulting from its interaction with society. Modern scenario planning exercises emerged in the U.S. and France following World War II, helping governments and mili-

taries develop policy and allocate research funding [173]. In the 1970s, companies such as the Royal Dutch/Shell group began to develop scenarios for strategic business planning [174]. STS researchers have since used scenario planning to assess potential outcomes of new technologies such as biomedical devices, nanotechnology, and ecological research [175, 176].

The most common type of scenario planning workshop uses a matrix approach, where participants identify two dimensions of uncertainty and then develop four mutually exclusive scenarios [177]. For example, a group at Arizona State University developed four scenarios about a household device for measuring biosignatures, using the axes of “responsibility for health” (e.g., control and interpretation of health data) and “value to society” [178]. Workshop participants have traditionally been technical experts and strategists, but recently STS scholars have pushed for “participatory scenario planning,” which includes non-expert stakeholders or publics in the scenario development [179].

A scenario planning exercise about robotic lower-limb prostheses could explore uncertainties with existence and embodiment: What does it mean to wear a “cyborg” limb, and how does prosthesis embodiment and user satisfaction change as the devices interface with humans in more advanced ways? In addition to technical experts, this scenario planning exercise should recruit people with lower-limb amputation. Issues of access to robotic prostheses may or may not be considered, depending on the dimensions identified at the workshop.

7.4.5 Analogical case studies

One can gain insights into the public reception of an innovation by considering previous technologies that are similar in ethical/social implications or intended function. When analyzing previous technologies, important considerations include who has responded to the innovation; whether the response was positive or negative; and through what avenues these responses were communicated [155]. These patterns of social responses can point out issues that technical researchers may not have considered and help generate policy recommendations, such as a moratorium on the technology or development of mandatory performance standards.

Recently, researchers at the University of Michigan studied the potential impacts of facial recognition technology in schools using an analogical case comparison approach [180]. By considering technologies similar in function (e.g., CCTV) and moral significance (e.g., biobanks), the analysis uncovered several potential negative impacts of facial recognition in schools, including the amplification of racial bias, normalization of surveillance, and stricter enforcement of social norms.

An analogical case study applied to robotic lower-limb prostheses may consider devices that have a similar aim of increasing mobility (wheelchairs, exoskeletons, orthoses); devices that have similar social implications (data monitoring in activity trackers); and other wearable medical de-

vices (insulin pumps or cochlear implants). The cases for consideration would be chosen by STS and technical experts, and the resulting analysis could identify policy directions in cyber security of wearable devices, insurance coverage, or criteria for FDA approval.

7.4.6 Discursive TA

Discourse analysis is the study of language and conversations (written, verbal, sign language) to reveal insights about the relationship between individuals and society [181]. Discursive TAs investigate the language surrounding new technologies, such as how a novel innovation is presented to stakeholders or interactions between experts and lay people, to gauge implications for the wider society. This approach is prominent in healthcare, where researchers have analyzed language in videotaped interviews and historical documents to study physician-patient interaction in primary care practices [182]. Discursive TAs have been conducted in the biomedical field as well, for example, to study patient reactions to new pharmaceutical developments [183].

Discursive TA researchers turn a critical lens towards the reported outcomes of participatory TAs, pointing out that conversation among lay people is likely facilitated by a technical or STS expert. The resulting lay citizen attitudes and beliefs are then widely viewed as true citizen values when they are instead a product of a social process. Additionally, pTA participants may be considering hypothetical situations, thus engaging with a new technology in unfamiliar terms that seem disconnected from their own daily life. Discursive TAs, therefore, aim to study participants' language and conversation "in its own right" by taking the interactional setting into account [183].

The language surrounding an innovation can "...shape the assumptions about what is real, what is possible, and what is desirable for a social group" [22]. Once a technology is standardized, any uncertainties in its design are generally forgotten. However, discursive strategies can expose "the loose ends and forgotten controversies which are hidden beneath an apparent consensus" [184]. Applied to robotic lower-limb prostheses, a discursive TA could ask what narratives are produced by robotic prosthesis developers, how they reflect the underlying values present in the design, and how developers' shared vision of robotic prostheses has become standardized over time.

7.4.7 Proposed assessment for robotic lower-limb prostheses

I propose conducting a discursive TA to study how robotic lower-limb prosthesis developers imagine the future of robotic prostheses; what narratives they produce or disregard to support their values; and how these narratives have been co-produced with standardized assessments for evaluating prosthesis performance. For this study, I plan to first analyze documents from academic studies, blog posts, and advocacy groups to gain a picture of the main concerns for robotic lower-limb prosthesis technology. Next, I will analyze documents and videos produced by (or with, in

the case of media interviews) lower-limb robotic prosthesis developers, sampling material from both the academic and commercial sectors. Selecting academic labs that have received a relatively high amount of citations and media attention and prosthesis companies that are top sellers of technologically-advanced prostheses will ensure I am analyzing the narratives most likely to shape society's shared visions.

In this study, I will look for these established discourses, identify new discourses, and identify disregarded or ignored discourses in the language surrounding robotic lower-limb prostheses. Emerson et al. present a coding strategy for analyzing themes in written material [185]: First, I will use an "open coding" process to identify "...as many ideas and themes as time allows" in the material. After selecting main themes, I will re-analyze each text in a "focused coding" stage, which involves "...further delineating subthemes and subtopics" and connecting data between research groups, intended audience, and sectors. Lastly, I will investigate how these discourses have been co-produced with standards and regulations. For example, are standardized assessments influencing what language is used to describe and market a new prosthesis design? Is the language further enforcing the criteria used for evaluating prosthesis performance? My final report will include both a summary of findings and recommendations for robotic lower-limb prosthesis researchers regarding their narratives and evaluations.

7.5 Preliminary Results

To focus the scope of my technology assessment, I first present my findings of evaluation criteria that have already been identified by academic studies, clinician experience, and the perspectives of people with lower-limb loss. These sources commonly highlight issues with safety, comfort, cost, accessibility, and clinical translation. Next, I present preliminary results from my discourse analysis of the narratives and standards surrounding robotic prostheses.

7.5.1 Existing concerns about robotic lower-limb prostheses

Safety: Safe operation is a major concern for robotic prostheses because the powerful motors can act independently of the human's motion. A robotic prosthesis relies on its sensors to detect the user's intent and produce the corresponding joint torques, but its sensing and prediction capabilities are far from infallible [20]. Applying the wrong torque to the knee or ankle joint could cause serious injury, especially during high-risk tasks such as stair descent or hiking.

Microprocessor-controlled knees have fewer safety concerns because they do not supply an external energy source. In fact, studies have shown that microprocessor-controlled knees reduce fall risk compared to passive prostheses [186]. However, microprocessor-controlled prosthesis com-

panies note that gait-disrupting malfunctions can still occur, such as loss of hydraulic/pneumatic resistance, locking up, or irregular side-to-side motion [187].

User comfort: Comfort is another important factor in the decision to adopt a new technology. Many K-2 (in laboratory settings; see accessibility concerns below) and K-3 amputees prefer microprocessor-controlled prostheses over passive devices [188, 189], but the responses to technologically advanced prostheses are not, of course, universal across lower-limb amputees. In one study where participants used a robotic ankle in daily life for two weeks, some reported that the device felt more natural and made walking easier while others noted it made walking harder, added to phantom limb pain, and was heavy, noisy, and “too controlling” [91].

Another comfort concern, socket pain, is present across passive, microprocessor-controlled, and robotic devices. One physical therapist writes: “Despite today’s high-tech equipment, we still have some really low-tech problems,” describing the heat rash, blisters, and painful ingrown hairs that result from liner and socket use on the residual limb [190]. An alternative to socket use is osseointegration, a surgical procedure to implant a metal anchor directly into the residual limb for prosthesis attachment [191]. Although this procedure addresses the painful socket effects, it also requires a surgery with long recovery time and increased infection risk.

Aside from physical comfort, the device’s appearance has implications for amputee comfort in social situations. As robotic prostheses add more sensors and torque-producing abilities, they tend to get larger and more “bionic,” resembling the robotic limbs in superhero movies. Britt Young notes: “It seems like...if you’re missing a limb, then you need to have a cyborg arm” [154], so many people with limb differences opt for the high-tech option to conform to peoples’ expectations. However, some amputees prefer simpler prostheses that offer cosmetic similarity to an intact limb rather than functional benefits.

The notion of prosthesis embodiment, or how people think of their prosthesis in relation to themselves (e.g., a part of their body versus a separate tool), is related to comfort and may affect prosthesis adoption rates. In a study about prosthesis embodiment in upper-limb devices, researchers found that embodiment may be a function of regular prosthesis wear rather than prosthesis design [192]. Research about lower-limb prosthesis embodiment found that embodiment can depend on many factors, such as amputation level, residual limb pain, visual appearance of the prosthesis, and mobility level [193]. Confirming these findings with further studies and incorporating these factors into the prosthesis design and fitting process could increase adoption rates.

Prosthesis cost and insurance coverage: Cost is a significant and complicated obstacle for individuals looking to upgrade their device. The out-of-pocket costs of a prosthesis vary along several axes: insurance providers’ contracts with prosthetists offices to determine device pricing; the premium and deductible associated with the patient’s insurance provider; and the amputee’s existing health status, smoking status, and K-level [194].

If an amputee is looking to buy a \$40,000 microprocessor-controlled knee with private insurance, a Premier (high-premium, low-deductible) Blue Cross Blue Shield plan would still require an out-of-pocket payment of \$3,500 [194]. Some more affordable plans, such as those provided by Aetna, may only supply \$2,000 per year for this type of medical equipment, leaving the patient with a \$38,000 bill. Medicare Plan B would cover 80% of the bill, which still leaves \$8,000. These coverage numbers² are for a non-smoking K-4 amputee, but the coverage may decrease or disappear completely as K-level decreases due to health issues or age. Prices for the two commercially available robotic lower-limb prostheses are not readily available online, but they are likely higher than microprocessor-controlled devices due to the addition of motors.

Despite what the high costs suggest, amputees do not purchase prostheses for life. Prosthesis sockets are typically replaced every 2-4 years [195], and microprocessor-controlled knees are most often replaced every 3-5 years [187]. The average lifespan of a robotic prosthesis has not yet been established due to their limited commercial use, but the sophisticated sensor array and motors will likely need frequent maintenance. Insurance providers are supposed to cover a replacement prosthesis if the device cannot perform its intended function, but the process is rarely straightforward. The Amputee Coalition warns its readers that they may have to spend days on the phone with insurance companies and encourages amputees to “be [their] own advocate” [137].

Accessibility: Advanced prostheses are inaccessible due to stringent body weight and fitness requirements. Microprocessor-controlled and robotic prostheses have weight limits that exclude people on either end of the operating range [196], and Medicare (and many private insurers) will not cover a microprocessor-controlled knee for patients under level K-3 [148]. These restrictions are concerning because clinicians utilize a wide variety of assessments to measure K-level, many of which are subjective. In a 2016 study, 67% of surveyed prosthetists did not believe that K-level could accurately capture a patient’s rehabilitation potential [197]. Furthermore, K-2 subjects have demonstrated preference for microprocessor-controlled knees in laboratory studies and clinical trials [188, 148].

Clinical translation: Because lower-limb prosthesis technologies are increasing in sophistication, difficulties arise when they move to a clinical setting. Robotic prostheses have a multitude of parameters (e.g., magnitude and timing of assistive joint torque) that must be adjusted based on user preference and observed gait benefits. This tuning process can take hours, which is unrealistic for busy clinicians [119]. In addition to time, the tuning process involves expertise; robotic prostheses are complex electrical and mechanical systems with which most clinicians have little experience [20].

These concerns highlight the need for a systematic evaluation of robotic lower-limb prostheses’ societal consequences. Such an assessment could help researchers mitigate the negative impacts of

²Note that these numbers are from a 2012 analysis of prosthesis costs in Connecticut.

robotic prostheses; provide regulation and legislation directions for policy makers; and facilitate awareness and deliberation among the lay public [198].

7.5.2 Analysis of discursive themes and standardized assessments

In disability studies, STS scholars have already identified several discourses surrounding assistive technologies. Ingunn Moser and John Law study how technologies articulate “personhood” in people with disabilities, viewing a person’s ability not as a “naturalized able body” but rather in the context of their relationships with other humans and objects [199]. They explore this concept in a case study of an assistive device called the Rolltalk, a wheelchair with features such as communication via speech synthesis, remote control of home appliances, and rearview cameras. The Rolltalk developers aim to “normalize” people with disabilities through use of the device, while the users themselves aim to be counted as “normal” by embracing socially accepted features of personhood. However, Rolltalk “...tends to constitute [forms of personhood] in specific ways, while allowing others to escape.” The specific tasks that Rolltalk enables reflect the designers’ values.

In another work, Moser explores four discourses that provide definitions of “human” and “normal” [200]. She uses descriptions from everyday life of a man named Olav to illustrate how each of four discourses constructs Olav, his abilities, and his disabilities. The *rehabilitation* discourse views differences from the “ideal human” as needing to be corrected using technology; the *prosthetics* discourse normalizes the prosthesis itself; the *actor-network* discourse defines ability and disability based on the individual’s entire human and technological network; and the *cyborg* discourse views the combination of human and technology as a single entity rather than a network. These discourses may overlap or contradict each other in the language surrounding a design.

Standard-setters in healthcare and biomedical fields produce benchmarks and regulations alongside the discourses surrounding technology. Despite the objective connotation that accompanies any health standard, Coreen McGuire notes that “...the thresholds of normalcy that we rely on for classification of health are more fluid than we might imagine” [201]. McGuire’s analysis of health measurement technology in Britain found that the people selected to define standards represent “easily recognizable groups” that reflect discrimination and bias. For example, designers of the 1920s artificial ear for testing telephone sound quality used data from ten “normal” males to tune the device; consequently, those with even a small amount of hearing loss had difficulty understanding telephone conversation. The shared vision of “normal” defined who could successfully use a telephone.

Rehabilitation discourse: My preliminary research analyzed documents from four academic research groups in the United States and two companies (Ossur and Ottobock) that sell robotic lower-limb prostheses in the United States. To narrow the scope of my preliminary analysis, I

selected academic professors that focus on prosthesis hardware design: Prof. Elliott Rouse at University of Michigan; Prof. Hugh Herr at Massachusetts Institute of Technology; Prof. Thomas Sugar at Arizona State University; and Prof. Michael Goldfarb at Vanderbilt University. Future work should add research groups that focus on prosthesis sensing and control to this analysis, as these factors also incorporate design narratives and influence user adoption. I analyzed two academic articles and two documents for a lay audience (e.g., university news article, TED Talk) for each academic research group. For the two companies, I analyzed the text on the sales page for the Ossur Power Knee and Ottobock Empower; the corresponding regulatory compliance documents if they were available (I could not find them on Ottobock’s website); the study summaries available for each device; and transcripts from one promotional video of each device.

In preliminary research, I have found examples of Moser’s rehabilitation discourse in all documents analyzed. The documents I analyzed include a notion of “normal” human movement, and in most cases, an aim to rehabilitate amputees back to this definition:

In this paper, we first review the biomechanics of a normal human ankle in walking, and then use this biomechanical study to motivate the prosthetic design [202].

This language conveys the goal of “rehabilitating” a lower-limb amputee’s ankle to mimic a “normal human ankle.” Likewise, Ossur markets their powered (robotic) knee using language such as:

Power is fundamental to the biomechanical systems that move us. If you’re living with limb loss or limb difference, that power is a missing ingredient. Restoring it can be transformational [149].

This marketing material presents normal biomechanical function as that which is created by human muscles, and reflects the vision of restoring that same power to prostheses.

Academic articles presenting designs of robotic lower-limb prostheses provide tables with design specifications determined from able-bodied characteristics, including maximum range of motion, degrees of freedom, peak torque, and peak velocity. One robotic ankle prosthesis mimics ankle torques from a 75-kg (165-lb) person with a walking speed of 1.25 m/s [202]. The paper from which the ankle torque data were drawn does not specify subject gender (10 able-bodied subjects, aged 23-29 years, mean 64 kg mass). The design specifications for one robotic knee prosthesis use a dataset with a near-even amount of male and female subjects and larger age range (20 able-bodied subjects, aged 22-72 years, mean 68.5 kg mass, 9 males and 11 females) [133]. Both articles describe problems with current amputee walking and propose their robotic prostheses as means to mitigate these issues.

The documents geared towards a lay audience (e.g., news articles) do not provide quantified specifications for the “normal” human; instead, they compare outcomes to pre-amputation or able-bodied functionality. A 2021 article about a robotic knee prosthesis, for example, notes its “...potential to provide the power and natural gait of a human leg” [203]. In a 2014 TED Talk, Hugh Herr describes his lab’s work to “study how humans with normal physiologies stand, walk and run,” which provides design motivation for bionic joints [204]. However, Herr’s TED talk is unique in the documents I analyzed because it also explores additional themes besides the rehabilitation narrative. For example, Herr proposes altering the built environment and technology:

A person can never be broken. Our built environment, our technologies, are broken and disabled. We the people need not accept our limitations, but can transcend disability through technological innovation [204].

We also see Moser’s cyborg theme in Herr’s talk, where the notion of bionic integration is presented positively:

We’re beginning the age in which machines attached to our bodies will make us stronger and faster and more efficient [204].

Standards for assessing prosthesis designs: The rehabilitation theme also appears in the types of assessments used to evaluate prosthesis designs. For example, researchers developed a controller for a robotic knee-ankle prosthesis to assist with stair ascent. To evaluate the controller, they recorded biomechanical data (knee and ankle angles, angular velocities, torques) of three amputees as they used the prosthesis to walk up stairs, and compared the results to average “healthy” (able-bodied) biomechanical data [205].

As noted in Sec. 7.2.2, the FDA has not yet established criteria for approving Class II robotic lower-limb prostheses. I have found that developers’ assessments of lower-limb prostheses reflect this lack of standardization. Research groups define different tasks (Sit-to-Stand, walking up an incline, stair ascent) for evaluation, and typically hand-pick human datasets to define “normal” human body characteristics for evaluation.

In contrast to lower-limb prostheses, assessments have been standardized across academic, commercial, and regulatory sectors for *upper-limb* robotic prostheses, which have a longer history of development and commercialization [206, 207, 208]. One such standardized assessment is the Box and Blocks Test, which was introduced by clinicians to evaluate human motor performance and was subsequently validated in the academic community [209, 210, 211]. The test involves moving small blocks between two compartments of a box using one hand, illustrating a collective vision of how a human or prosthetic hand should function. In 2014, the FDA defined this assessment as part of the criteria to evaluate robotic prostheses for commercial use [208]. De-

velopers wishing to commercialize their prosthesis therefore prioritize this assessment during the design stage.

The Box and Blocks Test assesses the performance of a one-handed task. However, for unilateral upper-limb amputees, one-handed tasks could be performed with the intact limb. Britt Young argues that people with unilateral upper-limb loss need upper-limb prostheses that can assist with *two-handed* tasks:

[Promotional] videos never show prosthesis users doing two-handed tasks. Instead, they do things most of them would only ever undertake with their other fleshy hand, like write or drink water...they forget to ask if these hands are actually useful [212].

Young's concern highlights the need for technology assessments in both the design and standardization phases to avoid this disconnect between assessment criteria and user needs.

The “rehabilitation” theme I identified in my preliminary research does not address the common concerns I described in Sec. 7.5.1. This mismatch between design criteria and stakeholder needs indicates the lack of human-centered design in robotics work. Human-centered design aims to improve human well-being and user satisfaction through applying concepts from human factors, ergonomics, and usability studies to the design of a technology [213]. Incorporating (or requiring) human-centered design techniques in lower-limb prosthesis development could help align stakeholder concerns and design objectives.

7.6 Discussion

Robotic lower-limb prostheses are still largely in the research and development phase. In this chapter, I introduce several technology assessment methods and argue that it is critical to analyze the discourse, standards, and regulations surrounding robotic lower-limb prostheses before they become widely commercialized. My preliminary results have identified the “rehabilitation” discourse to be common in text about robotic lower-limb prostheses, while the main concerns highlighted by stakeholders of safety, comfort, cost, accessibility, and clinical translation are typically not addressed. If the rejection rate of upper-limb robotic prostheses (which have widespread commercial availability) is to be avoided for lower-limb robotic prostheses, we must consider alternative discourses and standards in our research. My study of lower-limb prosthesis standards coincides with the FDA's recent standardized performance criteria for another type of lower-limb wearable robot: the exoskeleton [153]. This timing enables me to map how FDA regulations converge with narratives and assessments from the development sectors in my future work. As Moser notes, identifying these narratives and standards does not mean the technologies are fixed in place, but instead prompts “...possibilities for action, interruption, [and] interference” [200].

My own research in motion planning for robotic lower-limb prostheses reflects the rehabilitation narrative. In Ch. 6, I train trip-recovery libraries using data from an experiment with able-bodied subjects. The preliminary results of my technology assessment prompt me to acknowledge that able-bodied trip recovery behavior may not be effective and comfortable for lower-limb amputees. My future work with the trip-recovery framework will explore other methods for forming the libraries, ideally using a range of trip-recovery behavior desired by the prosthesis user rather than able-bodied subject data.

Despite the prevalence of the rehabilitation narrative in the documents that I have examined, there are some examples of existing research that specifically addresses the concerns in Sec. 7.5.1. For example, Amos Winter's group focuses on designing passive prosthetic knees that offer high performance and low cost [214]. To address problems with clinical translation and tuning time for robotic lower-limb prostheses, Robert Gregg's group has developed tools for generalizing human kinematics across a range of inclines and walking speeds, which can help minimize tuning time [215]. Regarding user comfort, Deanna Gates' group has surveyed individuals with upper-limb amputation regarding their attitudes towards novel prosthetic control techniques [216]. To ensure effective translation of robotic lower-limb prostheses, more research in these areas should be encouraged by funding sources.

Testimonials from people with upper-limb loss have shown us that the underlying values surrounding a design can influence user satisfaction and adoption rates [154]. Because narratives about collective technological futures "...play a critical role in framing what technology is made and why" [217], examining the discourse used by robotic prosthesis developers to create and enforce shared visions of the future can identify the underlying values embedded in robotic prosthesis design. My discursive TA will enable developers (including myself) to better communicate their innovations to prospective users; identify design characteristics and evaluation methods to enhance or change; and influence the regulations and standards used to approve a device for commercial use.

CHAPTER 8

Conclusion and Future Directions

Falls are common and dangerous for older adults and people with lower-limb amputation. Existing efforts to mitigate fall risk have been unsuccessful due to the difficulty of identifying those likely to fall and the challenges of safely planning the perturbation-recovery behavior of assistive robots. This dissertation made several contributions to human prediction and wearable robot motion planning for fall prevention. In this chapter, I provide a summary of the contributions described in this dissertation, as well as future research directions for each.

8.1 Discussion of Contributions

Chapter 2, Open-Source Datasets This dissertation includes two perturbative datasets that are publicly available. In the Sit-to-Stand experiment, subjects stood up from a chair while being pulled either forwards or backwards by motor-driven cables. This 11-subject experiment can be used by future researchers to investigate control strategies during center-of-mass perturbations. In the second experiment, subjects were tripped by tethers attached to their feet while walking on a treadmill. This is, to our knowledge, the largest dataset of tripped walking described in the literature and the only open-source trip dataset. Future researchers can use these data to learn more about able-bodied trip recoveries during walking. This dissertation analyzes data from the swing leg immediately after the trip until touchdown, but the gait patterns in subsequent steps following a trip are also of interest for human motor control and wearable lower-limb robot behavior.

Chapter 3, Characterizing Stability during Sit-to-Stand This dissertation describes a method for quantifying time-varying, individualized stability during the Sit-to-Stand motion, validated on the 11-subject Sit-to-Stand experiment. In Ch. 3, we use observed kinematic data to construct dynamic and controller models for each subject. We compute Stability Basins by approximating the set of model states that can successfully reach standing under the controller model. We show

that Stability Basins accurately predict (1) a subject’s instability (i.e., stepping or sitting) before the step or sit is initiated and (2) a subject’s successful completion of the Sit-to-Stand task.

Chapter 4, Inverse Optimal Control This chapter presents an implementation of an inverse optimal control formulation developed by [84] with perturbative Sit-to-Stand data. It is a common assumption that humans behave optimally during motion. Uncovering the objective during a specific task could provide insights for human motor control and generate accurate predictions of perturbation recovery. We modeled each subject’s Sit-to-Stand motion using an inverted pendulum with a torque applied to the base, and solved for a cost function of model states and control input that was minimized by the observed data. We found that across 11 subjects, inverse optimal control produced the highest weights on the center of mass angular position and velocity, with lower weights on the input torque. This visual trend was similar when trained on perturbed or unperturbed data, though relative weights between position and velocity changed. When we used these functions to generate pendulum control input trajectories using an optimal control formulation, the resulting simulated pendulum states performed with higher accuracy compared to an LQR controller model.

Chapter 5, Predicting Kinematic Responses to Trips In this chapter, we investigated three different models for predicting swing hip kinematics after a trip. We selected three response variables for prediction: Swing hip height, anterior-posterior position, and flexion-extension angle. Accurate predictions of these variables are crucial for determining correct knee and ankle trajectories (e.g., torques or angles) in a wearable lower-limb robot such that the system avoids unsafe foot scuffs and places the swing foot in a desirable location at touchdown. We found that of the Gaussian process regression, neural network, and dynamic models used in this chapter, the Gaussian process regression models produced predictions with the highest accuracy and lowest computation time. Furthermore, Gaussian process regression models trained on generalized data produced similar prediction accuracy to the individualized models, which is useful for real-world implementation.

Chapter 6, Fall Prevention via Prosthetic Trip Recovery This chapter introduces a trip-recovery planning framework for a powered prosthesis in the swing phase. After a trip is detected, the framework uses the swing hip predictions and uncertainty from Ch. 5 to generate a reachable set of predicted human behavior. Next, the framework chooses prosthesis knee and ankle trajectories that minimize a foot placement objective and satisfy a constraint to avoid foot scuffs for the entire set of predicted human motion. The framework is therefore able to incorporate set-based predictions of human motion and guarantees safety across sets of potential trajectories, instead of the single most likely one. The results, which displayed minimal foot scuffs and higher foot place-

ment accuracy compared to a nominal impedance-based controller, show promise for preventing falls with robotic lower-limb prostheses.

Chapter 7, Technology Assessment for Robotic Lower-Limb Prostheses Robotic lower-limb prostheses are a relatively new technology that have not yet reached widespread use. Therefore, identifying potential barriers to this technology’s real-world implementation will enable us as robotic prosthesis developers to change our research and policy objectives. This chapter presents an overview of technology assessments: systematic methods for evaluating a new technology’s political, social, and/or economic impact introduced by scholars in Science and Technology Studies. I present a research plan and preliminary results from a discourse analysis that focuses on the language and standards surrounding robotic lower-limb prostheses. The preliminary results indicate that the common narratives produced by prosthesis developers aim to “rehabilitate” amputees to a specific idea of “normal.” The rehabilitation narrative is not aligned with common stakeholder concerns of safety, comfort, cost, accessibility, and clinical translation. This discovery indicates several future directions that prosthesis researchers should explore.

8.2 Future Directions

The future directions for each chapter are motivated by the technology assessment in Ch. 7. They aim to address the challenges of real-world viability and user acceptance in each of the tools presented in this dissertation. Increasing real-world applicability in each of the chapters in this document involves conducting further validation studies; reducing assumptions; changing methods with stakeholder concerns in mind; and implementing robotic motion planning on hardware.

8.2.1 Stability Basins

The Stability Basin results in Ch. 3 are demonstrated for able-bodied young subjects. The validation in this dissertation confirmed that Stability Basins can accurately predict when a subject will become unstable (i.e., step or sit); however, they have not yet been shown to be a reliable predictor of fall risk.

To investigate Stability Basins’ effectiveness in identifying those at risk of falling, I propose a longitudinal study with older adults. The study’s hypothesis is that the volume of an individual’s Stability Basin is an indicator of fall risk. We will compute Stability Basins for 100 adults aged 85 or older and retain contact with the subjects over a five-year period. The subjects will be instructed to alert the study team of any falls and in what circumstances the falls occurred. After five years, we will compare the Stability Basins between fallers and non-fallers, analyzing Stability Basin volume

and shape. A significant difference in mean volume between the two groups would indicate the Stability Basins' promise for identifying fall risk. To mitigate falls, physical therapists could then provide targeted training to maximize patients' Stability Basin volume.

Several challenges must be addressed in the current Stability Basin framework before conducting a large-scale longitudinal study with older adults. The method for computing Stability Basins in Ch. 3 trains individual controller models using successful perturbed trials, i.e. trials where subjects were pulled either forwards or backwards by a motor-driven cable and reached standing without taking a step or sitting back down. However, applying such perturbations to adults aged 85 and older is a safety hazard, and perturbative experiments with a large subject population are time consuming.

Instead, we must develop a way to construct individual controller models using observed unperturbed data from subjects' daily lives. One option is to use cameras and a 3D pose estimation algorithm to record Sit-to-Stand joint motion in a home or clinical setting, and build individualized controller models using these data. The variation present in daily Sit-to-Stand (e.g., standing up from the couch or dinner table) could provide a similar range of input bounds as used in Ch. 3. Previous work has shown promise of using cameras and a 3D pose estimation algorithm to estimate loads during rehabilitation activities [218]. However, collecting such video data over an extended period of time poses security and privacy risks for subjects. If camera-based estimation is used to build individualized controllers, researchers must make every effort to ensure the data security and informed consent of participants.

The second challenge for the proposed longitudinal study involves clinical translation, because a successful study of fall risk in older adults needs an effective clinical partnership. We must develop the tools and training to ensure that busy clinicians can compute Stability Basins for study participants with minimal time and technological barriers. Addressing this challenge requires considering how to most easily incorporate Stability Basin research into a clinical schedule and designing open-source and intuitive software to run the computations.

8.2.2 Inverse Optimal Control

The primary aim of applying inverse optimal control to human motion datasets is to discover the underlying objectives during motion. In Ch. 4, we found polynomial objective functions of a pendulum model's states and input. Using the model's states and input, we connected the polynomial functions to physical meanings (a higher priority on center of mass kinematics than input torques). However, the scaling of states and inputs required by the polynomial formulation may have influenced some coefficient values. Overfitting may have also been a problem, especially for higher-degree lagrangians, because the Sit-to-Stand dataset described in Ch. 2 does not have

many observations per subject. Because Sit-to-Stand requires torques at the ankle, hip, and knee, future work should utilize computational resources to scale the inverted pendulum model to a three-link model, and compare the relative weights on each joint to existing biomechanical analyses of Sit-to-Stand.

Many other studies investigate humans' optimization of metabolic cost during walking, so applying the inverse optimal control formulation in [84] (or a different formulation) to a walking dataset may produce more interesting or intuitive results. Although we focused on studying tripped steps in this dissertation, the 16-subject walking experiment described in Ch. 2 includes nominal walking data for each subject. By nature of the task (it is less time consuming to record steps on a treadmill than the same number of Sit-to-Stand trials), we have more nominal and tripped step trials to train models with than for the Sit-to-Stand dataset. In future work, we will implement inverse optimal control with nominal walking data to investigate whether the results confirm findings from other studies about optimality in human walking. Then, we will compare the underlying nominal cost functions to those found with tripped steps to see if objective functions differ significantly during nominal walking and trip-recovery.

8.2.3 Trip-Recovery Prediction

This dissertation presented a method for predicting swing hip kinematics (height, anterior-posterior position, and flexion-extension angle) in the sagittal plane in the same step following a trip. We restricted our analysis to the sagittal plane and to the tripped swing phase because those hip response trajectories were needed for our trip-recovery planning framework in Ch. 6. However, we must address three challenges to effectively assist with perturbation recovery during walking in the real world. First, our work assumed that the human's chosen trip-recovery strategy (elevating, delayed lowering, or lowering) was known and trained strategy-specific models. We made this assumption given previous work that predicted strategy selection [96], but the strategy prediction algorithm was tested on a different trip dataset than our own. Our first task will be to verify that this algorithm generates accurate and quick predictions of strategy selection in our 16-subject trip dataset. If the strategy prediction algorithm proves inaccurate, we must either modify it to generalize to our dataset or generate predictive models of swing hip kinematics that are not strategy-specific.

Second, trips and trip recoveries may not be restricted to the sagittal plane. Humans may place their foot laterally or cross their swing foot over the midline when recovering from a trip, which would be associated with distinct lateral hip motion. Wearable robots should be aware of lateral hip motion predictions when planning joint behavior. Such predictions would be especially useful if the robotic ankle joint can provide actuation along the interior-exterior or abduction-adduction axes.

Third, this dissertation focuses on prediction during the same swing phase as the trip occurrence, but humans exhibit trip-recovery behavior different from nominal walking for several steps after a trip. A wearable robot should therefore predict human trip response for several steps after a trip. Future work will investigate how many steps subjects typically take when returning to nominal walking after a trip, and fit models to these multi-step data.

8.2.4 Trip-Recovery Planning in Robotic Lower-Limb Prostheses

Our proposed framework for planning trip recoveries in robotic knee-ankle prostheses is tested using able-bodied swing hip data from our 16-subject trip experiment and simulated prosthesis knee and ankle motion. The trip-recovery libraries and target sets are also trained using able-bodied data. Several challenges must be addressed to ensure the prosthesis trip-recovery framework is effective and comfortable for amputees.

First, swing hip trip-recovery behavior is likely different between able-bodied and amputee subjects. Shirota et al. tripped amputee subjects via retractable tethers and found that individuals exhibit similar trip-recovery strategies to those observed in able-bodied subjects [19], but the associated hip response variable trajectories for each strategy may be different. Furthermore, the subjects in Shirota's experiment were wearing passive prostheses, and swing hip behavior may change if a powered prosthesis is used. Even if recording trip-recovery data can produce accurate predictive models of human trip-recovery behavior, it would be impossible to conduct such an experiment with every prosthesis user. The challenge for future work is to generate accurate predictions of swing hip trip-recovery kinematics for people with lower-limb amputations. Because trip-recovery kinematics may evolve as an individual becomes accustomed to the trip-recovery controller, an online system that continually adds observed trip recoveries to the Gaussian process model training set may be a reasonable solution.

Second, the proposed framework selects prosthesis behavior from a parameterized library of knee and ankle trajectories that are formed using able-bodied data, and aims to place the foot close to a target set formed by able-bodied data. As noted in Ch. 7, although prosthesis developers often use able-bodied data as a standard for prosthesis performance, these trajectories may not be desirable or safe for trip recovery by an amputee. In future work, we will investigate ways to form the trip-recovery libraries and target sets using input from lower-limb amputees.

Third, the planning framework assumes that the human's swing hip motion remains unchanged throughout the tripped swing phase regardless of prosthesis knee and ankle behavior. This assumption may be reasonable for trips that occur at the end of the swing phase, where the human has little time to re-plan swing hip behavior. However, for trips that occur at the beginning of the swing phase, individuals typically have around 500 ms until touchdown, and prosthesis behavior

could affect the swing hip behavior. We propose adding re-planning to the trip-recovery framework. At trip onset, the prosthesis would generate and execute a plan that avoids foot scuffs and ensures correct ground placement; measure the human swing hip behavior while the plan is executed, update the Gaussian process regression model with new measured conditional points; and generate a new plan given the updated hip motion. This iterative planning would continue until touchdown occurs.

Fourth, as noted in Sec. 8.2.3 for the swing hip predictions, the trip-recovery planner must be extended to multiple steps after trip occurrence. The theoretical framework presented in Ch. 6 remains unchanged for planning other steps, but different data would be used for training the predictive models, trip-recovery libraries, and target sets.

Finally, a number of assumptions must be addressed for hardware implementation. We currently assume the trip is immediately detected; the current and previous states of the human are known at trip onset; and the human’s chosen strategy is known. Each of these assumptions must be implemented on prosthetic hardware for the framework to be effective. Previous work in each of these areas provides a guide for our own implementation [121, 122, 219, 96], but ensuring their reliable use on prosthesis hardware will be challenging.

While each of these challenges are being addressed, we will also work to generalize the framework to different tasks. For example, the trips studied in this dissertation were induced by tethers, but in reality, trips often occur from physical obstacles of varying heights and depths that the swing leg must avoid while recovering. If the prosthesis has access to vision information, we can add another constraint to avoid foot scuffs on the obstacle. Vision information could also be used to plan multi-step walking if an obstacle is detected ahead. In these cases, the framework could build human predictions, parameterized libraries, target sets, and costs/constraints using different sets of human data.

8.2.5 Technology Assessment for Robotic Lower-Limb Prostheses

The technology assessment presented in Ch. 7 is a starting point for re-shaping the narratives and standards surrounding lower-limb prostheses. However, to complete this qualitative work, I must analyze enough sources to reach *saturation*, where “...sufficient and redundant information for all aspects of the phenomenon under research has been gathered” [220]. One way to achieve saturation is through *triangulation*: using multiple methods and sources to confirm findings.

In my future work, I will gather hundreds of documents about robotic lower-limb prostheses from academia, industry, and regulatory and standard-setting organizations (e.g., FDA, ANSI). Document analysis will center the published narratives that shape our societal view of prostheses; explore what subject groups, measurements, tasks, and constraints are present in standards;

and investigate how actors justify the use of such standards for evaluation. I will supplement the document analysis with interviews with members of the academic, commercial, and regulatory settings. Interviews will help uncover the assumptions that underlie these narratives and provide an opportunity for triangulation.

For comparison, I will conduct a similar analysis for robotic upper-limb prostheses. Both upper- and lower-limb prosthesis technologies aim to replace a missing limb, but a closer look reveals differences in intended function, development timeline, adoption rates, and commercialization [221, 222]. Because upper-limb prostheses are commercially established and more data likely exists about their pricing, challenges, and regulations, comparing findings between the two technologies may help convince researchers and policymakers of the importance of addressing concerns.

8.3 Concluding Remarks

The research presented in this dissertation aims to prevent falls through predicting human response to perturbation and planning safe recovery motion in wearable lower-limb robots. Ch. 2 described two open-source human datasets that can be used to study perturbation response during Sit-to-Stand and walking. In Chs. 3 and 4, we characterize stability and investigate underlying optimality during center-of-mass perturbations during Sit-to-Stand. Chs. 5 and 6 present methods for predicting human responses to trips during walking and using the predictions to plan safe trip-recovery behavior in lower-limb prostheses. Lastly, Ch. 7 describes technology assessments that engineering researchers can apply to their own fields and presents preliminary results from my own assessment of robotic lower-limb prostheses.

Each framework and tool presented here shows promise for increasing safety in the real world, but is in an early stage of implementation. This chapter describes future to increase the real-world applicability of the content in each chapter. Importantly, the future work is informed by my analysis of the anticipated political, social, and economic impacts of assistive technology discussed in Ch. 7. By removing assumptions, developing methods to involve potential users in the design process, and thinking critically about clinical translation, the research presented in this dissertation can increase the safety, comfort, and effectiveness of technologies for reducing fall risk.

APPENDIX A

Intervals, Zonotopes, and Polynomial Zonotopes

A.1 Notation

The n -dimensional real numbers are \mathbb{R}^n , natural numbers \mathbb{N} , the unit circle is \mathbb{S}^1 , and the set of 3×3 orthonormal matrices is $\text{SO}(3)$. Upper case letters represent sets (or matrices), while lowercase letters represent vectors or scalars. Subscripts either index elements of a vector or elements of a set, depending on context. For a matrix A , we use A^\top to denote its transpose. Vectors are denoted as (x_1, \dots, x_n) unless the shape is critical to understanding, in which case we write $[x_1, \dots, x_n]^\top$. Let $U, V \subset \mathbb{R}^n$. For a point $u \in U$, $\{u\} \subset U$ is the set containing only u . The power set of U is $\mathcal{P}(U)$. The Minkowski sum is $U \oplus V = \{u + v \mid u \in U, v \in V\}$. The set-based cross product is defined as $U \otimes V = \{u \times v \mid u \in U, v \in V\}$. Let $W \subset \mathbb{R}^{n \times n}$ be a set of square matrices. Then, set-based multiplication is defined as $WV = \{Av, \mid A \in W, v \in V\}$. For vectors $a, b \in \mathbb{R}^3$, the cross product $a \times b$ can be expressed as $a^\times b$, where

$$a^\times = \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix}. \quad (\text{A.1})$$

Let $\mathbf{0}$ and $\mathbf{1}$ denote matrices of zeros and ones of appropriate size, and I_n be the $n \times n$ identity matrix.

A.2 Intervals

Various set representations are used throughout this dissertation. The most fundamental set representation we use is the interval.

An n-dimensional interval is a set of real numbers

$$[x] = \{y \in \mathbb{R}^n \mid \underline{x}_i \leq y_i \leq \bar{x}_i, \quad \forall i = 1, \dots, n\}. \quad (\text{A.2})$$

For succinctness, we sometimes denote an interval $[x]$ by $[\underline{x}, \bar{x}]$, where \underline{x} and \bar{x} are the infimum and supremum, respectively, of $[x]$. Let the operations inf and sup return the infimum and supremum of an interval:

$$\text{inf}([x]) = \underline{x} \quad (\text{A.3})$$

$$\text{sup}([x]) = \bar{x}. \quad (\text{A.4})$$

The Minkowski sum and difference of $[x]$ and $[y]$ are given by

$$[x] \oplus [y] = [\underline{x} + \underline{y}, \bar{x} + \bar{y}], \quad (\text{A.5})$$

$$[x] \ominus [y] = [\underline{x} - \bar{y}, \bar{x} - \underline{y}]. \quad (\text{A.6})$$

The product of $[x]$ and $[y]$ is

$$[x][y] = [\min(\underline{x}\underline{y}, \underline{x}\bar{y}, \bar{x}\underline{y}, \bar{x}\bar{y}), \max(\underline{x}\underline{y}, \underline{x}\bar{y}, \bar{x}\underline{y}, \bar{x}\bar{y})]. \quad (\text{A.7})$$

Given a scalar interval $[a]$ or interval matrix $[Y]$ multiplied by an interval matrix $[X]$, the element in the i^{th} row and j^{th} column of the product is

$$([a][X])_{ij} = [a][X]_{ij}, \quad (\text{A.8})$$

$$([X][Y])_{ij} = \bigoplus_{k=1}^n ([X]_{ik}[Y]_{kj}), \quad (\text{A.9})$$

where n is the number of columns of $[X]$ and number of rows of $[Y]$. Lastly, given two interval vectors $[x]$ and $[y] \in \mathbb{IR}^3$, their cross product is

$$[x] \otimes [y] = [x]^\times [y], \quad (\text{A.10})$$

where $[x]^\times$ is the skew-symmetric matrix representation of $[x]$ as in (A.1)

A.3 Zonotopes

Next, we introduce zonotopes, which are more expressive than intervals. A zonotope $Z \subset \mathbb{R}^n$ is a convex, centrally-symmetric polytope defined by a center $g_0 \in \mathbb{R}^n$ and generators $g_i \in \mathbb{R}^n$:

$$Z = \{z \in \mathbb{R}^n \mid z = g_0 + \sum_{i=1}^{n_g} g_i \beta_i, \quad \beta_i \in [-1, 1]\} \quad (\text{A.11})$$

where there are $n_g \in \mathbb{N}$ generators.

Lemma A.3.1. *A zonotope can be overapproximated by an interval.*

Proof. Consider $Z \subset \mathbb{R}^n$ with center c and n_g generators g_i . Then,

$$Z \subseteq [g_0 - \sum_{i=1}^{n_g} |g_i|, g_0 + \sum_{i=1}^{n_g} |g_i|] \quad (\text{A.12})$$

■

Zonotopes are closed under Minkowski sum. Given $Z_1 \subset \mathbb{R}^n$ with center g_0 and p generators g_i with indeterminates β_i , and $Z_2 \subset \mathbb{R}^n$ with center h_0 and q generators h_j with indeterminates λ_j :

$$Z_1 \oplus Z_2 = \{z \in \mathbb{R}^n \mid z = z_1 + z_2, z_1 \in Z_1, z_2 \in Z_2\} \quad (\text{A.13})$$

$$\begin{aligned} Z_1 \oplus Z_2 &= \{z \in \mathbb{R}^n \mid z = (g_0 + h_0) + \sum_{i=1}^{n_g} g_i \beta_i + \sum_{j=1}^q h_j \lambda_j, \\ &\quad \dots, \beta_i \in [-1, 1], \lambda_j \in [-1, 1]\}. \end{aligned} \quad (\text{A.14})$$

A matrix zonotope $M \subset \mathbb{R}^{n \times m}$ is a zonotope whose center $G_0 \in \mathbb{R}^{n \times m}$ and generators G_i are matrices:

$$M = \left\{ A \in \mathbb{R}^{n \times m} \mid A = G_0 + \sum_{i=1}^{n_g} G_i \beta_i \right\}. \quad (\text{A.15})$$

Note that a matrix zonotope M and zonotope Z of appropriate sizes can be multiplied together. Let M have p generators G_i with indeterminates β_i , and Z have q generators h_j with indeterminates

λ_j .. Then,

$$M \odot Z = \{z \in \mathbb{R}^n \mid z = Ay, A \in M, y \in Z\} \quad (\text{A.16})$$

$$M \odot Z = \left\{ z \in \mathbb{R}^n \mid z = G_0 g_0 + \sum_{j=1}^q G_0 g_j \lambda_j + \sum_{i=1}^{n_g} G_i g_0 \beta_i + \sum_{i=1}^{n_g} \sum_{j=1}^q G_i g_j \beta_i \lambda_j, \beta_i \in [-1, 1], \lambda_j \in [-1, 1] \right\}. \quad (\text{A.17})$$

However, the result is no longer a zonotope due to the products of indeterminates $\beta_i \lambda_j$. This motivates the need for a more complex set representation called polynomial zonotopes [223].

A.4 Polynomial Zonotopes

For a thorough introduction of polynomial zonotopes, we refer the interested reader to work by Kochdumper and Althoff [223]. We present an overview of the necessary definitions and operations here.

To ease the introduction of polynomial zonotopes, we first express zonotopes in a slightly different notation. Let the zonotope Z have center g_0 , n_g generators g_i , and indeterminates β_i . We introduce an indeterminate vector $x \in [-1, 1]^{n_g}$. The elements of x correspond to the indeterminates β_i . Next, we introduce the exponents $\alpha_i \in \mathbb{N}^{n_g}$. Letting $\alpha_1 = [1, 0, \dots, 0]$, $\alpha_2 = [0, 1, 0, \dots, 0]$, ..., $\alpha_{n_g} = [0, 0, \dots, 0, 1]$, we note that $x^{\alpha_1} = x_1, \dots, x^{\alpha_{n_g}} = x_{n_g}$ where the exponentiation is applied element-wise. Then, letting the center g_0 be associated with the exponent $\alpha_0 = [0, 0, \dots, 0]$ we may write the zonotope Z as

$$Z = \{z \in \mathbb{R}^n \mid z = \sum_{i=0}^{n_g} g_i x^{\alpha_i}, x \in [-1, 1]^{n_g}\}. \quad (\text{A.18})$$

Here, we have written Z as the set of points produced by the polynomial $p(x) = \sum_{i=0}^{n_g} g_i x^{\alpha_i}$ over the domain $x \in [-1, 1]^{n_g}$. We have our first *polynomial zonotope*, a set representation which includes zonotopes as a subset. Though in (A.18) we defined a zonotope where each α_i has only one nonzero element, more generally polynomial zonotopes may employ exponents $\alpha_i \in \mathbb{N}^{n_g}$ which are arbitrary. We refer to x^{α_i} as a *monomial*, and a *term* $g_i x^{\alpha_i}$ is produced by multiplying a monomial by the associated *generator* g_i . A polynomial zonotope $\langle P \rangle \subset \mathbb{R}^n$ is given by its coefficients, exponents, and indeterminates. We write $\langle P \rangle = \mathcal{PZ}(g_i, \alpha_i, x)$. Later in the text, we distinguish between “dependent” and “independent” terms:

Definition A.4.1. *Let α_i have only one nonzero element (the k -th element). If the k -th element is*

zero for all $\alpha_j, j \neq i$ then the term $g_i x^{\alpha_i}$ is independent. This implies that the term depends only on k -th indeterminate x_k , and that this indeterminate does not appear anywhere else. A term which does not meet these criteria is dependent.

The Minkowski Sum of two polynomial zonotopes $\langle P_1 \rangle \subset \mathbb{R}^n = \mathcal{PZ}(g_i, \alpha_i, x)$ and $\langle P_2 \rangle \subset \mathbb{R}^n = \mathcal{PZ}(h_j, \beta_j, y)$ follows polynomial addition:

$$\langle P_1 \rangle \oplus \langle P_2 \rangle = \{z \in \mathbb{R}^n \mid z = p_1 + p_2, p_1 \in \langle P_1 \rangle, p_2 \in \langle P_2 \rangle\} \quad (\text{A.19})$$

$$\langle P_1 \rangle \oplus \langle P_2 \rangle = \{z \in \mathbb{R}^n \mid z = \sum_{i=0}^{n_g} g_i x^{\alpha_i} + \sum_{j=0}^q h_j y^{\beta_j}\}. \quad (\text{A.20})$$

Similarly, we may write the matrix product of two polynomial zonotopes $\langle P_1 \rangle$ and $\langle P_2 \rangle$ when the sizes are compatible, i.e. elements in $\langle P_1 \rangle$ have the same number of columns as elements of $\langle P_2 \rangle$ have rows. Letting $\langle P_1 \rangle \subset \mathbb{R}^{n \times m}$ and $\langle P_2 \rangle \subset \mathbb{R}^{m \times k}$, we obtain $\langle P_1 \rangle \langle P_2 \rangle \subset \mathbb{R}^{n \times k}$:

$$\langle P_1 \rangle \langle P_2 \rangle = \{z \in \mathbb{R}^{n \times k} \mid z = p_1 p_2, p_1 \in \langle P_1 \rangle, p_2 \in \langle P_2 \rangle\} \quad (\text{A.21})$$

$$\langle P_1 \rangle \langle P_2 \rangle = \{z \in \mathbb{R}^{n \times k} \mid z = \sum_{i=0}^{n_g} g_i \left(\sum_{j=0}^q h_j y^{\beta_j} \right) x^{\alpha_i}\}. \quad (\text{A.22})$$

Furthermore, if $\langle P_1 \rangle \subset \mathbb{R}^3$ and $\langle P_2 \rangle \subset \mathbb{R}^3$, we implement a set-based cross product as matrix multiplication. We create $\langle P_1 \rangle^x \subset \mathbb{R}^{3 \times 3}$ as

$$\langle P_1 \rangle^x = \{A \in \mathbb{R}^{3 \times 3} \mid A = \sum_{i=0}^{n_g} \begin{bmatrix} 0 & -g_{i,3} & g_{i,2} \\ g_{i,3} & 0 & -g_{i,1} \\ -g_{i,2} & g_{i,1} & 0 \end{bmatrix} x^{\alpha_i}\} \quad (\text{A.23})$$

where $a_{i,j}$ refers to the j -th element of a_i . Then, the set-based cross product $\langle P_1 \rangle \otimes \langle P_2 \rangle = \langle P_1 \rangle^x \langle P_2 \rangle$ is well-defined.

We briefly note that the addition, multiplication and cross product of a polynomial zonotope with a constant vector or matrix is well-defined if the constant is appropriately sized. In this case, one constructs a polynomial zonotope with that constant vector or matrix as the center g_0 and no other generators, and applies the definitions above.

Both set addition and multiplication can be complicated by the fact that $\langle P_1 \rangle$ and $\langle P_2 \rangle$ may share indeterminates. That is, the k -th element of x and the l -th element of y represent the same indeterminate. Therefore, it can be beneficial to bring both polynomial zonotopes to a *common representation* before performing these operations. Letting $\tilde{x} = x \cup y$, we may redefine the polynomial zonotopes $\langle P_1 \rangle = \langle g_i, \tilde{\alpha}_i, \tilde{x} \rangle$ and $\langle P_2 \rangle = \langle h_i, \tilde{\beta}_i, \tilde{x} \rangle$ where $\tilde{\alpha}_i$ and $\tilde{\beta}_j$ exponentiate the

appropriate elements of \tilde{x} .

One desirable property of polynomial zonotopes is the ability to obtain subsets by plugging in values of known indeterminates. For example, let's say a polynomial zonotope $\langle P \rangle$ represented a set of possible positions of a robot arm operating near an obstacle. It would be beneficial to know whether a particular choice of $\langle P \rangle$'s indeterminates yields a subset of positions that could collide with the obstacle. To this end, we introduce the operation of "slicing" a polynomial zonotope $\langle P \rangle = \mathcal{PZ}(g_i, \alpha_i, x)$ by evaluating an element of the indeterminate x . Given an indeterminate x_s and a value $\sigma \in [-1, 1]$, slicing yields a subset of $\langle P \rangle$ by plugging σ into the specified element x_s :

$$\text{slice}(\langle P \rangle, x_s, \sigma) \subset \langle P \rangle = \{z \in \langle P \rangle \mid z = \sum_{i=0}^{n_g} g_i x^{\alpha_i}, x_s = \sigma\}. \quad (\text{A.24})$$

In this document, we require a means to bound the elements of a polynomial zonotope. It is possible to efficiently generate these upper and lower bounds on the values of a polynomial zonotope through overapproximation. We define the `sup` and `inf` operations which return these upper and lower bounds, respectively, by taking the absolute values of generators. For $\langle P \rangle \subseteq \mathbb{R}^n$, these return

$$\text{sup}(\langle P \rangle) = g_0 + \sum_{i=1}^{n_g} |g_i|, \quad (\text{A.25})$$

$$\text{inf}(\langle P \rangle) = g_0 - \sum_{i=1}^{n_g} |g_i|. \quad (\text{A.26})$$

Note that for any $z \in \langle P \rangle$, $\text{sup}(\langle P \rangle) \geq z$ and $\text{inf}(\langle P \rangle) \leq z$, where the inequalities are taken element-wise. These bounds may not be tight because possible dependencies between indeterminates are not accounted for, but they are quick to compute.

The operations defined above (multiplication in particular) increases the number of generators required to represent a polynomial zonotope. In practice, successively applying these operations can become computationally intractable as the memory required to represent a polynomial zonotope increases. To combat this computational burden, it is possible to generate overapproximations of polynomial zonotopes using fewer generators by *reducing* a polynomial zonotope. If a polynomial zonotope $\langle P \rangle \subset \mathbb{R}^n$ has n_g terms, but a maximum of q terms are desired, excess terms can be overapproximated by an interval:

$$\text{reduce}(\langle P \rangle, q) = \left\{ z \in \mathbb{R}^n \mid z = \sum_{i=0}^{q-n} g_i x^{\alpha_i} + \sum_{j=1}^n [-b_j, b_j] \right\} \quad (\text{A.27})$$

where b_j is a vector of zeros except for the j^{th} element which is equal to $\sum_{i=q-n+1}^{n_g} |(g_i)_j|$. Notice

that $\text{reduce}(\langle P \rangle, q)$ always overapproximates $\langle P \rangle$, i.e. $\langle P \rangle \subseteq \text{reduce}(\langle P \rangle, q)$ [223, Prop. 16]. One can reorder the terms of the polynomial zonotope such that only certain desirable terms are replaced by intervals, e.g. to produce a tighter overapproximation. The output of `reduce` may be represented as a polynomial zonotope.

Though we have defined several basic operations like addition and multiplication above, it may be desirable to use polynomial zonotopes as inputs to more complicated functions. In fact, one can overapproximate any analytic function evaluated over a polynomial zonotope using a Taylor expansion, which itself can be represented as a polynomial zonotope. Consider the analytic function $f : \mathbb{R} \rightarrow \mathbb{R}$ (e.g. $\cos(\cdot)$, $\sin(\cdot)$) and polynomial zonotope $\langle P_1 \rangle = \mathcal{PZ}(g_i, \alpha_i, x)$, with each $g_i \in \mathbb{R}$. Then,

$$f(\langle P_1 \rangle) = \{y \in \mathbb{R} \mid y = f(z), z \in \langle P_1 \rangle\}. \quad (\text{A.28})$$

We generate a polynomial zonotope $\langle P_2 \rangle$ such that $f(\langle P_1 \rangle) \subseteq \langle P_2 \rangle$ using a Taylor expansion of degree $d \in \mathbb{N}$, where the error incurred from the finite approximation is overapproximated using a Lagrange remainder. The method follows the Taylor expansion found in the reachability algorithm in [223], which builds on previous work on conservative polynomialization found in [224]. Recall that the Taylor expansion about a point $c \in \mathbb{R}$ is

$$f(z) = \sum_{n=0}^{\infty} \frac{f^{(n)}(c)}{n!} (z - c)^n. \quad (\text{A.29})$$

Note that the error incurred by a finite Taylor expansion can be bounded using the Lagrange remainder R :

$$\left| f(z) - \sum_{n=0}^d \frac{f^{(n)}(c)}{n!} (z - c)^n \right| \leq R, \quad (\text{A.30})$$

where R is given by

$$R = \frac{M |z - c|^{d+1}}{(d+1)!}, \quad M = \max_{\delta} (|f^{(d+1)}(\delta)|), \quad \delta \in [c, z]. \quad (\text{A.31})$$

For a polynomial zonotope, the infinite dimensional Taylor expansion is given by

$$f(\langle P_1 \rangle) = \sum_{n=0}^{\infty} \frac{f^{(n)}(c)}{n!} (\langle P_1 \rangle - c)^n \quad (\text{A.32})$$

In practice, only a finite Taylor expansion of degree $d \in \mathbb{N}$ can be computed. Letting $c = g_0$

Table A1: Operations with polynomial zonotopes may either be computed exactly or in an overapproximative fashion.

Operation	Exact (x) or Overapproximative (o)
\oplus (Minkowski Sum)	x
\odot (Set Multiplication)	x
\otimes (Set Cross Product)	x
slice	x
sup/inf	o
reduce	o
cos/sin	o

(the center of $\langle P_1 \rangle$), and noting that $(z - c) = (\sum_{i=1}^{n_g} g_i x^{\alpha_i})$ for $z \in \langle P_1 \rangle$, we write

$$\langle P_2 \rangle := \{z \in \mathbb{R} \mid z = \sum_{n=0}^d \left(\frac{f^{(n)}(g_0)}{n!} \left(\sum_{i=1}^{n_g} g_i x^{\alpha_i} \right)^n \right) + [R]\} \quad (\text{A.33})$$

and the Lagrange remainder $[R]$ can be computed using interval arithmetic as

$$[R] = \frac{[M][\langle P_1 \rangle]^d}{(d+1)!}, [M] = f^{(d+1)}([\langle P_1 \rangle]) \quad (\text{A.34})$$

where $[\langle P_1 \rangle]$ is an overapproximation of $\langle P_1 \rangle$ by an interval that can be computed using the `sup` (A.25) and `inf` (A.26) operations. Note that $\langle P_2 \rangle$ can be expressed as a polynomial zonotope because all terms in the summation are polynomials of x , and an interval can be included using an independent generator.

BIBLIOGRAPHY

- [1] Anne Felicia Ambrose, Geet Paul, and Jeffrey M. Hausdorff. Risk factors for falls among older adults: A review of the literature. *Maturitas*, 75(1):51–61, 2013.
- [2] Bruno J. Vellas, Sharon J. Wayne, Linda J. Romero, Richard N. Baumgartner, and Philip J. Garry. Fear of falling and restriction of mobility in elderly fallers. *Age and Ageing*, 26, 1997.
- [3] G. A. R. Zijlstra, J. C. M. van Haastregt, J. Th. M. van Eijk, E. van Rossum, P. A. Stalenhoef, and G. I. J. M. Kempen. Prevalence and correlates of fear of falling, and associated avoidance of activity in the general population of community-living older people. *Age and Ageing*, 36(3):304–309, 03 2007.
- [4] Facts about falls. Centers for Disease Control and Prevention, Aug 2021.
- [5] Kathryn Ziegler-Graham, Ellen J. MacKenzie, Patti L. Ephraim, Thomas G. Trivison, and Ron Brookmeyer. Estimating the prevalence of limb loss in the united states: 2005 to 2050. *Archives of Physical Medicine and Rehabilitation*, 89(3):422 – 429, 2008.
- [6] William C. Miller, Mark Speechley, and Barry Deathe. The prevalence and risk factors of falling and fear of falling among lower extremity amputees. *Archives of Physical Medicine and Rehabilitation*, 82(8):1031 – 1037, 2001.
- [7] A. J. Blake, K. Morgan, M. J. Bendall, H. Dalloso, S. B. J. Ebrahim, T. H. D. Arie, P. H. Fentem, and E. J. Basse. Falls by elderly people at home: prevalence and associated factors. *Age and Ageing*, 17(6):365–372, 11 1988.
- [8] Christopher Kevin Wong, Stanford T Chihuri, and Guohua Li. Risk of fall-related injury in people with lower limb amputations: A prospective cohort study. *Journal of Rehabilitation Medicine*, 48(1):80–85, 2016.
- [9] Sheila Clemens, Charissa Doerger, and Szu-Ping Lee. Current and emerging trends in the management of fall risk in people with lower limb amputation. *Current Geriatrics Reports*, 9(3):134–141, 2020.
- [10] Katherine Berg. Measuring balance in the elderly: Preliminary development of an instrument. *Physiotherapy Canada*, 41:304–311, 11 1989.

- [11] Diane Podsiadlo and Sandra Richardson. The timed up and go: A test of basic functional mobility for frail elderly persons. *Journal of the American Geriatrics Society*, 39:142–8, 1991.
- [12] Matthew J. Major, Stefania Fatone, and Elliot J. Roth. Validity and reliability of the berg balance scale for community-dwelling persons with lower-limb amputation. *Archives of Physical Medicine and Rehabilitation*, 2013.
- [13] Anne Shumway-Cook, Sandy Brauer, and Marjorie Woollacott. Predicting the probability for falls in community-dwelling older adults using the timed up & go test. *Physical Therapy*, 80(9):896–903, 2000.
- [14] Susan W Muir, Katherine Berg, Bert Chesworth, and Mark Speechley. Use of the berg balance scale for predicting multiple falls in community-dwelling elderly people: A prospective study. *Physical Therapy*, 88(4):449–459, 2008.
- [15] L. D. Gillespie, M. C. Robertson, W. J. Gillespie, C. Sherrington, S. Gates, L. M. Clemson, and S. E. Lamb. Interventions for preventing falls in older people living in the community. *Cochrane Database of Systematic Reviews*, 2012.
- [16] Carol A Miller, Jennifer E Williams, Katey L Durham, Selena C Hom, and Julie L Smith. The effect of a supervised community-based exercise program on balance, balance confidence, and gait in individuals with lower limb amputation. *Prosthetics and Orthotics International*, 41(5):446–454, 2017. PMID: 28067123.
- [17] Zoe A. Schafer, John L. Perry, and Natalie Vanicek. A personalised exercise programme for individuals with lower limb amputation reduces falls and improves gait biomechanics: A block randomised controlled trial. *Gait & Posture*, 63:282–289, 2018.
- [18] A. John Campbell and M. Clare Robertson. Rethinking individual and community fall prevention strategies: a meta-regression comparing single and multifactorial interventions. *Age and Ageing*, 36(6):656–662, 01 2007.
- [19] Camila Shirota, Ann M. Simon, and Todd A. Kuiken. Transfemoral amputee recovery strategies following trips to their sound and prosthesis sides throughout swing phase. *Journal of Neuroengineering and Rehabilitation*, 12, 2015. Copyright - Copyright BioMed Central 2015; Last updated - 2016-08-07.
- [20] Michael Goldfarb, Brian E. Lawson, and Amanda H. Shultz. Realizing the promise of robotic leg prostheses. *Science Translational Medicine*, 5(210):210ps15–210ps15, 2013.
- [21] A. J. Campbell, M. J. Borrie, and G. F. Spears. Risk factors for falls in a community-based prospective study of people 70 years and older. *Journal of Gerontology*, 44(4):112–117, 1989.
- [22] Wenda K. Bauchspies, Jennifer Croissant, and Sal Restivo. *Science, Technology, and Society: A Sociological Approach*. Blackwell Publishing Ltd., 2008.

- [23] Patrick D Holmes, Shannon M Danforth, Xiao-Yu Fu, Talia Y Moore, and Ram Vasudevan. Characterizing the limits of human stability during motion: perturbative experiment validates a model-based approach for the sit-to-stand task. *Royal Society Open Science*, 7(1):191410, 2020.
- [24] R.M. Alexander. *Optima for Animals*. Princeton University Press, 1996.
- [25] Camila Shirota, Ann M. Simon, and Todd A. Kuiken. Trip recovery strategies following perturbations of variable duration. *Journal of Biomechanics*, 47(11):2679 – 2684, 2014.
- [26] S. M. Danforth, X. Liu, M. J. Ward, P. D. Holmes, and R. Vasudevan. Predicting sagittal-plane swing hip kinematics in response to trips. *IEEE Robotics and Automation Letters*, July 2022.
- [27] Nitish Thatte, Nandagopal Srinivasan, and Hartmut Geyer. Real-time reactive trip avoidance for powered transfemoral prostheses. In *Proceedings of Robotics: Science and Systems*, FreiburgimBreisgau, Germany, 2019.
- [28] A. Girard. *Reachability of uncertain linear systems using zonotopes*, pages 291–305. Springer, 2005.
- [29] M. Althoff. An introduction to cora 2015. In *Proc. of the Workshop on Applied Verification for Continuous and Hybrid Systems*, 2015.
- [30] M. Althoff and D. Grebenyuk. Implementation of interval arithmetic in CORA 2016. In *Applied Verification for Continuous and Hybrid Systems*, pages 91–105, 2016.
- [31] Patrick Riley, David E. Krebs, and Rita A. Popat. Biomechanical analysis of failed sit-to-stand. In *IEEE Transactions on Rehabilitation Engineering*, volume 5, pages 353–9, 1998.
- [32] V. Shia, T.Y. Moore, P.D. Holmes, R. Bajcsy, and R. Vasudevan. Stability basin estimates fall risk from observed kinematics, demonstrated on sit-to-stand task. *Journal of Biomechanics*, 2018.
- [33] M.A. Hughes, D.K. Weiner, M.L. Schenkman, R.M. Long, and S.A. Studenski. Chair rise strategies in the elderly. *Clinical Biomechanics*, 9(3):187 – 192, 1994.
- [34] Patrick O. Riley, Margaret L. Schenkman, Robert W. Mann, and W.Andrew Hodge. Mechanics of a constrained chair-rise. *Journal of Biomechanics*, 24(1):77–85, 1991.
- [35] Visual3D. C-Motion Inc., Germantown, MD, USA.
- [36] MATLAB Statistics and Machine Learning Toolbox, 2017. The MathWorks, Natick, MA, USA.
- [37] A.M. Schillings, B.M.H. Van Wezel, and J. Duysens. Mechanically induced stumbling during human treadmill walking. *Journal of Neuroscience Methods*, 67(1):11 – 17, 1996.

- [38] S.T. King, M.E. EVELD, A. Martínez, K.E. Zelik, and M. Goldfarb. A novel system for introducing precisely-controlled, unanticipated gait perturbations for the study of stumble recovery. *J Neuroeng Rehabil.*, 2019.
- [39] Nexus. Vicon, Oxford, UK.
- [40] Laurence Rubenstein. Falls in older people: Epidemiology, risk factors and strategies for prevention. *Age and Ageing*, 35 Suppl 2:ii37–ii41, 10 2006.
- [41] J A Stevens, P S Corso, E A Finkelstein, and T R Miller. The costs of fatal and non-fatal falls among older adults. *Injury Prevention*, 12(5):290–295, 2006.
- [42] S. Heinrich, K. Rapp, U. Rissmann, C. Becker, and H.-H. König. Cost of falls in old age: a systematic review. *Osteoporosis International*, 21(6):891–902, Jun 2010.
- [43] M Clare Robertson, Nancy Devlin, Melinda M Gardner, and A John Campbell. Effectiveness and economic evaluation of a nurse delivered home exercise programme to prevent falls. 1: Randomised controlled trial. *BMJ*, 322(7288):697, 2001.
- [44] Bas Bloem, I Boers, M Cramer, R.G.J. Westendorp, and W. Gerschlager. Falls in the elderly. i. identification of risk factors. *Wiener klinische Wochenschrift*, 113:352–62, 2001.
- [45] Kim Delbaere, Klaus Hauer, and Stephen Lord. Evaluation of the incidental and planned exercise questionnaire (ipeq) for older people. *British Journal of Sports Medicine*, 44:1029–34, 06 2009.
- [46] Klaus Hauer, Lucy Yardley, Nina Beyer, Gertrudis Kempen, N Dias, M Campbell, Clemens Becker, and Chris Todd. Validation of the falls efficacy scale and falls efficacy scale international in geriatric patients with and without cognitive impairment: Results of self-report and interview-based questionnaires. *Journal of Gerontology*, 56:190–9, 10 2009.
- [47] Lillemor Lundin-Olsson, Lars Nyberg, and Yngve Gustafson. Stops walking when talking as a predictor of falls in the elderly. *Lancet*, 349:617, 04 1997.
- [48] Neil D. Weinstein. Unrealistic optimism about susceptibility to health problems: Conclusions from a community-wide sample. *Journal of Behavioral Medicine*, 10:481–500, 11 1987.
- [49] Adrian Furnham. Response bias, social desirability and dissimulation. *Personality and Individual Differences*, 7:385–400, 12 1986.
- [50] Teresa Steffen, Timothy Hacker, and Louise Mollinger. Age- and gender-related test performance in community-dwelling elderly people: Six-minute walk test, berg balance scale, timed up & go test, and gait speeds. *Physical Therapy*, 82:128–37, 02 2002.
- [51] Uffe Laessoe, Hans Christian Hoeck, Ole Simonsen, Thomas Sinkjaer, and Michael Voigt. Fall risk in an active elderly population - can it be assessed? *Journal of Negative Results in Biomedicine*, 6:2, 02 2007.

- [52] Jeffrey Hausdorff, Dean A. Rios, and Helen K. Edelberg. Gait variability and fall risk in community-living older adults: A 1-year prospective study. *Archives of Physical Medicine and Rehabilitation*, 82:1050–6, 08 2001.
- [53] Thurmon E. Lockhart and Jian Liu. Differentiating fall-prone and healthy adults using local dynamic stability. *Ergonomics*, 51(12):1860–1872, 2008. PMID: 19034782.
- [54] Elizabeth Hsiao-Wecksler. Biomechanical and age-related differences in balance recovery using the tether-release method. *Journal of Electromyography and Kinesiology*, 18:179–87, 05 2008.
- [55] P. Hur, B. A. Duiser, S. M. Salapaka, and E. T. Hsiao-Wecksler. Measuring robustness of the postural control system to a mild impulsive perturbation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 18(4):461–467, Aug 2010.
- [56] Sjoerd Bruijn, Onno Meijer, Peter Beek, and Jaap Van Dieen. Assessing the stability of human locomotion: a review of current measures. *Journal of the Royal Society Interface*, 11, 01 2014.
- [57] Philippe Terrier and Olivier Dériaz. Kinematic variability, fractal dynamics and local dynamic stability of treadmill walking. *Journal of NeuroEngineering and Rehabilitation*, 8(1):12, Feb 2011.
- [58] J.B Dingwell, J.P Cusumano, D Sternad, and P.R Cavanagh. Slower speeds in patients with diabetic neuropathy lead to improved local dynamic stability of continuous overground walking. *Journal of Biomechanics*, 33(10):1269 – 1277, 2000.
- [59] At Hof, M.G.J. Gazendam, and W.E. Sinke. The condition for dynamic stability. *Journal of Biomechanics*, 38:1–8, 02 2005.
- [60] Yi-Chung Pai and James Patton. Center of mass velocity-position predictions for balance control. *Journal of Biomechanics*, 30(4):347 – 354, 1997.
- [61] Martin Simoneau and Philippe Corbeil. The effect of time to peak ankle torque on balance stability boundary: experimental validation of a biomechanical model. *Experimental Brain Research*, 165(2):217–228, Aug 2005.
- [62] Masahiro Fujimoto and Li-Shan Chou. Dynamic balance control during sit-to-stand movement: An examination with the center of mass acceleration. *Journal of Biomechanics*, 45(3):543 – 548, 2012.
- [63] Rachid Aissaoui and J Dansereau. Biomechanical analysis and modelling of sit to stand task: a literature review. In *IEEE Conference on Systems, Man, and Cybernetics*, volume 1, pages 141 – 146 vol.1, 02 1999.
- [64] Elisabetta Papa and Aurelio Cappozzo. A telescopic inverted-pendulum model of the musculo-skeletal system and its use for the analysis of the sit-to-stand motor task. *Journal of Biomechanics*, 32(11):1205 – 1212, 1999.

- [65] AD Kuo. The relative roles of feedforward and feedback in the control of rhythmic movements. *Motor control*, 6:129–145, 04 2002.
- [66] Emanuel Todorov. Optimality principles in sensorimotor control. *Nature Neuroscience*, 7(9):907–915, 2004.
- [67] D. A. Winter. *The Biomechanics and Motor Control of Human Gait: Normal, Elderly and Pathological*. Waterloo Biomechanics, 01 1991.
- [68] Peter Lancaster and L Rodman. *The Algebraic Riccati Equation*, volume 14. 01 1995.
- [69] Olaf Stursberg and Bruce H. Krogh. Efficient representation and computation of reachable sets for hybrid systems. In *Hybrid Systems: Computation and Control*, pages 482–497, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.
- [70] Masahiro Fujimoto and Li-Shan Chou. Region of stability derived by center of mass acceleration better identifies individuals with difficulty in sit-to-stand movement. *Annals of Biomedical Engineering*, 42(4):733–741, Apr 2014.
- [71] Yang Wang and Manoj 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*, 10(9):20140405, 2014.
- [72] Sukyung Park, Fay B. Horak, and Arthur D. Kuo. Postural feedback responses scale with biomechanical constraints in human standing. *Experimental Brain Research*, 154(4):417–427, Feb 2004.
- [73] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2D pose estimation using part affinity fields. In *CVPR*, 2017.
- [74] Nils Smit-Anseeuw, C. David Remy, and Ram Vasudevan. Walking with confidence: Safety regulation for full order biped models. *IEEE Robotics and Automation Letters*, 4(4):4177–4184, 2019.
- [75] Jinsun Liu, Pengcheng Zhao, Zhenyu Gan, Matthew Johnson-Roberson, and Ram Vasudevan. Leveraging the template and anchor framework for safe, online robotic gait design. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10869–10875, 2020.
- [76] Thomas Gurriet, Sylvain Finet, Guilhem Boeris, Alexis Duburcq, Ayonga Hereid, Omar Harib, Matthieu Masselin, Jessy Grizzle, and Aaron Ames. Towards restoring locomotion for paraplegics: Realizing dynamically stable walking on exoskeletons. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [77] J M Brockway. Derivation of formulae used to calculate energy expenditure in man. *Clinical Nutrition*, 41:463–71, 12 1987.
- [78] Jeffrey R Koller, Deanna Gates, Daniel Ferris, and C David Remy. 'body-in-the-loop' optimization of assistive robotic devices: A validation study. In *Robotics: Science and Systems*, 06 2016.

- [79] Gustavo Arechavaleta, Jean-Paul Laumond, Halim Hicheur, and Alain Berthoz. An optimality principle governing human walking. *Robotics, IEEE Transactions on*, 24:5 – 14, 03 2008.
- [80] Manoj Srinivasan and Andy Ruina. Computer optimization of a minimal biped model discovers walking and running. *Nature*, 439(7072), 01 2005.
- [81] Marcel Menner, Peter Worsnop, and Melanie Zeilinger. Predictive modeling by infinite-horizon constrained inverse optimal control with application to a human manipulation task. *arXiv:1812.11600v1*, 2018.
- [82] M Cody Priess, Richard Conway, Jongeun Choi, John M. Popovich, and Clark Radcliffe. Solutions to the inverse lqr problem with application to biological systems analysis. *IEEE Transactions on Control Systems Technology*, 23:1–1, 08 2014.
- [83] Katja Mombaur, Anh Truong, and Jean-Paul Laumond. From human to humanoid locomotion—an inverse optimal control approach. *Autonomous Robots*, 28(3):369–383, Apr 2010.
- [84] Edouard Pauwels, Didier Henrion, and Jean-Bernard Lasserre. Inverse optimal control with polynomial optimization. *Proceedings of the IEEE Conference on Decision and Control*, 2015, 03 2014.
- [85] Pieter Abbeel and Andrew Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *In Proceedings of the Twenty-first International Conference on Machine Learning*. ACM Press, 2004.
- [86] Anirudha Majumdar et al. *spotless*, 2015.
- [87] Mosek. Mosek ApS, Copenhagen, Denmark.
- [88] Pengcheng Zhao, Shankar Mohan, and Ram Vasudevan. Control synthesis for nonlinear optimal control via convex relaxations. In *2017 American Control Conference (ACC)*, pages 2654–2661, 2017.
- [89] Margaret K.Y. Mak, Oron Levin, Joseph Mizrahi, and Christina W.Y. Hui-Chan. Joint torques during sit-to-stand in healthy subjects and people with parkinson’s disease. *Clinical Biomechanics*, 18(3):197–206, 2003.
- [90] F. Sibella, M. Galli, M. Romei, A. Montesano, and M. Crivellini. Biomechanical analysis of sit-to-stand movement in normal and obese subjects. *Clinical Biomechanics*, 18(8):745–750, 2003.
- [91] J. Kim, J. Wensman, N. Colabianchi, and D.H. Gates. The influence of powered prostheses on user perspectives, metabolics, and activity: a randomized crossover trial. *Journal of neuroengineering and rehabilitation*, 18, 2021.

- [92] Louis N. Awad, Jaehyun Bae, Kathleen O'Donnell, Stefano M. M. De Rossi, Kathryn Hendron, Lizeth H. Sloot, Pawel Kudzia, Stephen Allen, Kenneth G. Holt, Terry D. Ellis, and Conor J. Walsh. A soft robotic exosuit improves walking in patients after stroke. *Science Translational Medicine*, 9(400), 2017.
- [93] F. Sup, A. Bohara, and M. Goldfarb. Design and control of a powered knee and ankle prosthesis. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pages 4134–4139, 2007.
- [94] Brian E. Maki and William E. McIlroy. Control of rapid limb movements for balance recovery: age-related changes and implications for fall prevention. *Age and Ageing*, 35(2):ii12–ii18, 09 2006.
- [95] Audrey Schillings, B M van Wezel, Theo W Mulder, and Jaak Duysens. Muscular responses and movement strategies during stumbling over obstacles. *Journal of Neurophysiology*, 83 4:2093–102, 2000.
- [96] Maura E. Eveld, Shane T. King, Leo G. Vailati, Karl E. Zelik, and Michael Goldfarb. On the Basis for Stumble Recovery Strategy Selection in Healthy Adults. *Journal of Biomechanical Engineering*, 143(7), 04 2021. 071003.
- [97] Mohsen Gholami, Christopher Napier, and Carlo Menon. Estimating lower extremity running gait kinematics with a single accelerometer: A deep learning approach. *Sensors*, 20(10), 2020.
- [98] Abdelrahman Zaroug, Alessandro Garofolini, Daniel T. H. Lai, Kurt Mudie, and Rezaul Begg. Prediction of gait trajectories based on the long short term memory neural networks. *PloS One*, 16(8), 2021.
- [99] Varun Joshi and Manoj Srinivasan. A controller for walking derived from how humans recover from perturbations. *Journal of The Royal Society Interface*, 16(157), 2019.
- [100] Behnam Miripour Fard. A comparison between virtual constraint-based and model predictive-based limit cycle walking control in successful trip recovery. *AUT Journal of Mechanical Engineering*, 4(2):169–182, 2020.
- [101] Mark Ebdn. Gaussian processes: A quick introduction. arXiv, 2015.
- [102] Carl Edward Rasmussen and Williams Christopher K I. *Gaussian processes for machine learning*. MIT Press, 2008.
- [103] Bo Wang and Tao Chen. Gaussian process regression with multiple response variables. *Chemometrics and Intelligent Laboratory Systems*, 142:159–165, 2015.
- [104] John Sum, W. K. Kan, and Gilbert H. F. Young. A note on the equivalence of narx and rnn. *Neural Computing & Applications*, 8:33–39, 1999.
- [105] Narxnet: Nonlinear autoregressive neural network with external input. Mathworks. The MathWorks, Natick, MA, USA.

- [106] David A. Winter. *Biomechanics and Motor Control of Human Movement*. John Wiley & Sons, 2005.
- [107] E Alpaydin. Combined 5 x 2 cv f test for comparing supervised classification learning algorithms. *Neural computation*, 11(8):1885—1892, 1999.
- [108] Xiangyu Peng, Yadrianna Acosta-Sojo, Man I Wu, and Leia Stirling. Actuation timing perception of a powered ankle exoskeleton and its associated ankle angle changes during walking. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 30:869–877, 2022.
- [109] Shannon M. Danforth, Patrick D. Holmes, and Ram Vasudevan. Trip recovery in lower-limb prostheses using reachable sets of predicted human motion. *CoRR*, abs/2010.11228, 2020.
- [110] Christopher Kevin Wong, John Rheinstein, and Michelle Stern. Benefits for adults with transfemoral amputations and peripheral artery disease using microprocessor compared with nonmicroprocessor prosthetic knees. *American Journal of Physical Medicine & Rehabilitation*, pages 804–810, 2015.
- [111] S. K. Au, J. Weber, and H. Herr. Powered ankle–foot prosthesis improves walking metabolic economy. *IEEE Transactions on Robotics*, 25(1):51–66, 2009.
- [112] M.S. Kobetic, Rudi, M.S. To, Curtis S., M.S. Schnellenger, John R., PhD. Audu, Musa L., M.S. Bulea, Thomas C., C.O. Gaudio, Richard, M.D. Pinault, Gilles, PhD. Tashman, Scott, and PhD. Triolo, Ronald J. Development of hybrid orthosis for standing, walking, and stair climbing after spinal cord injury. *Journal of Rehabilitation Research and Development*, 46(3):447–62, 2009.
- [113] S. Rezazadeh, D. Quintero, N. Divekar, E. Reznick, L. Gray, and R. D. Gregg. A phase variable approach for improved rhythmic and non-rhythmic control of a powered knee-ankle prosthesis. *IEEE Access*, 7:109840–109855, 2019.
- [114] Michael Tucker, Jeremy Olivier, Anna Pagel, Hannes Bleuler, Mohamed Bourri, Olivier Lamercy, Jose del R. Millan, Robert Riener, Heike Vallery, and Roger Gassert. Control strategies for active lower extremity prosthetics and orthotics: A review. *Journal of NeuroEngineering and Rehabilitation*, 12:1, 01 2015.
- [115] H. Huang, T. A. Kuiken, and R. D. Lipschutz. A strategy for identifying locomotion modes using surface electromyography. *IEEE Transactions on Biomedical Engineering*, 56(1):65–73, 2009.
- [116] Daniel P. Ferris, Keith E. Gordon, Gregory S. Sawicki, and Ammanath Peethambaran. An improved powered ankle–foot orthosis using proportional myoelectric control. *Gait & Posture*, 23(4):425 – 428, 2006.
- [117] Heike Vallery, Rainer Burgkart, Cornelia Hartmann, Jürgen Mitternacht, Robert Riener, and Martin Buss. Complementary limb motion estimation for the control of active knee prostheses. *Biomedical Engineering / Biomedizinische Technik*, 56(1):45 – 51, 01 Feb. 2011.

- [118] Carl D. Hoover, George D. Fulk, and Kevin B. Fite. The Design and Initial Experimental Validation of an Active Myoelectric Transfemoral Prosthesis. *Journal of Medical Devices*, 6(1), 03 2012. 011005.
- [119] R. D. Gregg and J. W. Sensinger. Towards biomimetic virtual constraint control of a powered prosthetic leg. *IEEE Transactions on Control Systems Technology*, 22(1):246–254, 2014.
- [120] A.M. Schillings, B.M.H. Van Wezel, TH. Mulder, and J. Duysens. Widespread short-latency stretch reflexes and their modulation during stumbling over obstacles. *Brain Research*, 816(2):480 – 486, 1999.
- [121] B. E. Lawson, H. Atakan Varol, F. Sup, and M. Goldfarb. Stumble detection and classification for an intelligent transfemoral prosthesis. In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, pages 511–514, 2010.
- [122] F. Zhang, S. E. D’Andrea, M. J. Nunnery, S. M. Kay, and H. Huang. Towards design of a stumble detection system for artificial legs. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19(5):567–577, 2011.
- [123] A. J. Young, A. M. Simon, N. P. Fey, and L. J. Hargrove. Classifying the intent of novel users during human locomotion using powered lower limb prostheses. In *2013 6th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 311–314, 2013.
- [124] H. A. Varol, F. Sup, and M. Goldfarb. Multiclass real-time intent recognition of a powered lower limb prosthesis. *IEEE Transactions on Biomedical Engineering*, 57(3):542–551, 2010.
- [125] B. Zhong, R. Luiz da Silva, M. Li, H. Huang, and E. Lobaton. Environmental context prediction for lower limb prostheses with uncertainty quantification. *IEEE Transactions on Automation Science and Engineering*, pages 1–13, 2020.
- [126] R. D. Gregg, T. Lenzi, N. P. Fey, L. J. Hargrove, and J. W. Sensinger. Experimental effective shape control of a powered transfemoral prosthesis. In *2013 IEEE 13th International Conference on Rehabilitation Robotics (ICORR)*, pages 1–7, 2013.
- [127] F. Sup, H. A. Varol, and M. Goldfarb. Upslope walking with a powered knee and ankle prosthesis: Initial results with an amputee subject. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 19(1):71–78, 2011.
- [128] Y. David Li and E. T. Hsiao-Wecksler. Gait mode recognition and control for a portable-powered ankle-foot orthosis. In *2013 IEEE 13th International Conference on Rehabilitation Robotics (ICORR)*, pages 1–8, 2013.
- [129] Dario Martelli, Federica Vannetti, Mario Cortese, Peppino Tropea, Francesco Giovacchini, Silvestro Micera, Vito Monaco, and Nicola Vitiello. The effects on biomechanics of walking and balance recovery in a novel pelvis exoskeleton during zero-torque control. *Robotica*, 32(8):1317–1330, 2014.

- [130] S. Kousik, S. Vaskov, F. Bu, M. Johnson-Roberson, and R. Vasudevan. Bridging the gap between safety and real-time performance in receding-horizon trajectory design for mobile robots. *International Journal of Robotics Research*, 9 2018. [View online](#).
- [131] Patrick Holmes, Shreyas Kousik, Bohao Zhang, Daphna Raz, Corina Barbalata, Matthew Johnson-Roberson, and Ram Vasudevan. Reachable sets for safe, real-time manipulator trajectory design. In *Robotics: Science and Systems*, 2020.
- [132] JJ Eng, DA Winter, and AE Patla. Strategies for recovery from a trip in early and late swing during human walking. *Experimental brain research*, 102(2):339—349, 1994.
- [133] A. F. Azocar, L. M. Mooney, L. J. Hargrove, and E. J. Rouse. Design and characterization of an open-source robotic leg prosthesis. In *2018 7th IEEE International Conference on Biomedical Robotics and Biomechanics (Biorob)*, pages 111–118, 2018.
- [134] Limb loss statistics. Amputee Coalition. Accessed 10/13/2021.
- [135] Douglas G. Smith. Congenital limb deficiencies and acquired amputations in childhood. Amputee Coalition. Accessed 10/13/2021.
- [136] National diabetes statistics report. United States Centers for Disease Control and Prevention, 2020. Accessed 10/13/2021.
- [137] Prosthetic faqs for the new amputee. Amputee Coalition, 2021. Accessed 10/13/2021.
- [138] S. Bilodeau, R. Hébert, and J. Desrosiers. Lower limb prosthesis utilisation by elderly amputees. *Prosthetics and Orthotics International*, 24(2):126–132, 2000. PMID: 11061199.
- [139] Christiane Gauthier-Gagnon, Marie-Claude Grisé, and Diane Potvin. Enabling factors related to prosthetic use by people with transtibial and transfemoral amputation. *Archives of Physical Medicine and Rehabilitation*, 80(6):706 – 713, 1999.
- [140] Liliana E Pezzin, Timothy R Dillingham, Ellen J MacKenzie, Patti Ephraim, and Paddy Rossbach. Use and satisfaction with prosthetic limb devices and related services 11 no commercial party having a direct financial interest in the results of the research supporting this article has or will confer a benefit on the author(s) or on any organization with which the author(s) is/are associated. *Archives of Physical Medicine and Rehabilitation*, 85(5):723–729, 2004.
- [141] Limb loss definitions. Amputee Coalition, 2008. Accessed 10/13/2021.
- [142] B. Mishler. What are k levels? Ottobock. Accessed 10/13/2021.
- [143] B. Dupes. Prosthetic knee systems. Amputee Coalition. Accessed 10/13/2021.
- [144] Robert S Gailey, Kerry Allen, Julie Castles, Jennifer Kucharik, and Mariah Roeder. Review of secondary physical conditions associated with lower-limb amputation and long-term prosthesis use. *Journal of rehabilitation research and development*, 45 1:15–29, 2008.
- [145] Computer controlled knees. Ottobock. Accessed 10/13/2021.

- [146] C-leg above knee prosthetic leg. Ottobock. Accessed 10/13/2021.
- [147] Rheo knee. Ossur. Accessed 10/13/2021.
- [148] Andreas Kannenberg, Britta Zacharias, and Eva Pröbsting. Benefits of microprocessor-controlled prosthetic knees to limited community ambulators: systematic review. *Journal of rehabilitation research and development*, 51(10):1469—1496, 2014.
- [149] Power knee. Ossur. Accessed 10/13/2021.
- [150] Empower. Ottobock. Accessed 10/13/2021.
- [151] D. Salisbury. Robotic advances promise artificial legs that emulate healthy limbs. Vanderbilt University Medical Center Reporter, 2013. Accessed 10/13/2021.
- [152] B. Sutton. Overview of regulatory requirements: Medical devices - transcript. United States Food and Drug Administration, 2011. Accessed 10/13/2021.
- [153] Powered lower extremity exoskeleton. United States Food and Drug Administration, 2021. Accessed 10/13/2021.
- [154] V. Le. The lows of high tech. 99% Invisible, 2021. Accessed 10/15/2021.
- [155] David H. Guston and Daniel Sarewitz. Real-time technology assessment. *Technology in Society*, 24(1):93–109, 2002. American Perspectives on Science and Technology Policy.
- [156] David Banta. What is technology assessment? *International journal of technology assessment in health care*, 25 Suppl 1:7–9, 07 2009.
- [157] Harro van Lente, Tsjalling Swierstra, and Pierre-Benoit Joly. Responsible innovation as a critique of technology assessment. *Journal of Responsible Innovation*, 4:254–261, 05 2017.
- [158] L. Klüver. Europta: European participatory technology assessment. The Danish Board of Technology, 2000. Accessed 11/10/2021.
- [159] K. Konrad. Constructive technology assessment – sts for and with technology actors. European Association for the Study of Science and Technology Review, 2017. Accessed 11/10/2021.
- [160] Ulrike Bechtold, Daniela Fuchs, and Niklas Gudowsky. Imagining socio-technical futures – challenges and opportunities for technology assessment. *Journal of Responsible Innovation*, 4(2):85–99, 2017.
- [161] Brian O’Rourke, Wija Oortwijn, and Tara Schuller. The new definition of health technology assessment: A milestone in international collaboration. *International Journal of Technology Assessment in Health Care*, 36:1–4, 05 2020.

- [162] David Tomblin, Zachary Pirtle, Mahmud Farooque, David Sittenfeld, Erin Mahoney, Rick Worthington, Gretchen Gano, Michele Gates, Ira Bennett, Jason Kessler, Amy Kamin-ski, Jason Lloyd, and David Guston. Integrating public deliberation into engineering systems: Participatory technology assessment of nasa’s asteroid redirect mission. *Astropolitics*, 15(2):141–166, 2017.
- [163] P. Davidson. Delphi method. Accessed 11/10/2021.
- [164] Murray Turoff. The design of a policy delphi. *Technological Forecasting and Social Change*, 2(2):149–171, 1970.
- [165] L. Carson. Consult your community: A guide to using citizens’ juries. Planning NSW, 2003. Accessed 11/10/2021.
- [166] R. Sclove. Reinventing technology assessment: A 21st century model. Woodrow Wilson International Center for Scholars, 2010. Accessed 11/10/2021.
- [167] A. Pekkonen. Citizens juries. CIVICUS. Accessed 11/10/2021.
- [168] Daniel Lee Kleinman, Maria Powell, Joshua Grice, Judith Adrian, and Carol Lobes. A toolkit for democratizing science and technology policy: The practical mechanics of organizing a consensus conference. *Bulletin of Science, Technology & Society*, 27(2):154–169, 2007.
- [169] Participatory technology assessment. Consortium for Science, Policy, and Outcomes at Arizona State University. Accessed 11/12/2021.
- [170] Mike Michael. Publics performing publics: of pigs, pips and politics. *Public Understanding of Science*, 18(5):617–631, 2009.
- [171] Anitra Persaud, Stacy Desine, Katherine Blizinsky, and Vence L Bonham. A crispr focus on attitudes and beliefs toward somatic genome editing from stakeholders within the sickle cell disease community. *Genetics in medicine : official journal of the American College of Medical Genetics*, 21(8):1726—1734, August 2019.
- [172] Aviram Sharma. ‘we do not want fake energy’: The social shaping of a solar micro-grid in rural india. *Science, Technology and Society*, 25(2):308–324, 2020.
- [173] Ron Bradfield, George Wright, George Burt, George Cairns, and Kees Van Der Heijden. The origins and evolution of scenario techniques in long range business planning. *Futures*, 37(8):795–812, 2005.
- [174] Celeste Amorim Varum and Carla Melo. Directions in scenario planning literature – a review of the past decades. *Futures*, 42(4):355–369, 2010. Learning the Future Faster.
- [175] Emmanuel Smadja. Four scenarios towards more ethical futures: A case study in nanoscale science and technology. *foresight*, 8:37–47, 11 2006.

- [176] Elisa Oteros-Rozas, Berta Martín-López, Tim Daw, Erin Bohensky, James Butler, Rosemary Hill, Julia Martin-Ortega, Allyson Quinlan, Federica Ravera, Isabel Ruiz-Mallén, Matilda Thyresson, Jayalaxshmi Mistry, Ignacio Palomo, Garry Peterson, Tobias Plieninger, Kerry Waylen, Dylan Beach, Iris Bohnet, Maike Hamann, and Sandra Vilardy. Participatory scenario planning in place-based social-ecological research: Insights and experiences from 23 case studies. *ECOLOGY AND SOCIETY*, 20:32, 12 2015.
- [177] Katrin Bock and Stefan Trueck. Assessing uncertainty and risk in public sector investment projects. *Technology and Investment*, 2:105–123, 01 2011.
- [178] C. Selin. Scenario development workshop report: The future of medical diagnostics. The Center for Nanotechnology in Society, 2008. Accessed 11/12/2021.
- [179] Rebecca Freeth and Scott Drimie. Participatory scenario planning: From scenario ‘stakeholders’ to scenario ‘owners’. *Environment: Science and Policy for Sustainable Development*, 58:32–43, 07 2016.
- [180] C. Galligan, H. Rosenfeld, M. Kleinman, and S. Parthasarathy. Cameras in the classroom: Facial recognition technology in schools. University of Michigan Science, Technology, and Public Policy, 2020. Accessed 11/12/2021.
- [181] D. Tannen. Discourse analysis—what speakers do in conversation. Linguistic Society of America. Accessed 11/12/2021.
- [182] Sara Shaw and Julia Bailey. Discourse analysis: What is it and why is it relevant to family practice? *Family practice*, 26:413–9, 07 2009.
- [183] M. Veen, B. Gremmen, H. te Molder, and C. van Woerkum. Emergent technologies against the background of everyday life: Discursive psychology as a technology assessment tool. *Public Understanding of Science*, 20(6):810–825, 2011. PMID: 22397087.
- [184] Michael Lynch. The discursive production of uncertainty: The OJ Simpson ‘Dream Team’ and the sociology of knowledge machine. *Social Studies of Science*, 28(5-6):829–868, 1998.
- [185] Robert M. Emerson, Rachel I. Fretz, and Linda L. Shaw. *Writing ethnographic fieldnotes*. The University of Chicago Press, 2020.
- [186] M. Jason Highsmith, Jason T. Kahle, Dennis R. Bongiorni, Bryce S. Sutton, Shirley Groer, and Kenton R. Kaufman. Safety, energy efficiency, and cost efficacy of the c-leg for transfemoral amputees: A review of the literature. *Prosthetics and Orthotics International*, 34(4):362–377, 2010. PMID: 20969495.
- [187] What to consider when it’s time for a new prosthetic knee. Accessed 10/15/2021.
- [188] Kenton R. Kaufman, Kathie A. Bernhardt, and Kevin Symms. Functional assessment and satisfaction of transfemoral amputees with low mobility (fastk2): A clinical trial of microprocessor-controlled vs. non-microprocessor-controlled knees. *Clinical Biomechanics*, 58:116–122, 2018.

- [189] Dale Berry, Mark Olson, and Kinley Larntz. Perceived stability, function, and satisfaction among transfemoral amputees using microprocessor and nonmicroprocessor controlled prosthetic knees: A multicenter survey. *JPO: Journal of Prosthetics and Orthotics*, 21:32–42, 01 2009.
- [190] N. LaReia. What are some of the long-term physical effects of using or not using a prosthesis? Amputee Coalition, 2010. Accessed 10/15/2021.
- [191] Osseointegration: An overview. Amputee Coalition. Accessed 10/15/2021.
- [192] Susannah M Engdahl, Sean K Meehan, and Deanna H Gates. Differential experiences of embodiment between body-powered and myoelectric prosthesis users. *Scientific reports*, 10(1):15471, September 2020.
- [193] Robin Bekrater-Bodmann. Factors associated with prosthesis embodiment and its importance for prosthetic satisfaction in lower limb amputees. *Frontiers in Neurobotics*, 14, 2021.
- [194] B. Morrison and D. Topping. Robotic prosthetic availability analysis. Worcester Polytechnic Institute, 2012. Accessed 10/15/2021.
- [195] Frequently asked questions. American Prosthetics and Orthotics. Accessed 10/15/2021.
- [196] Prosthetic devices, specialized, microprocessor or myoelectric limbs. United Healthcare. Accessed 10/15/2021.
- [197] Dylan Borrenpohl, Brian Kaluf, and Matthew J. Major. Survey of u.s. practitioners on the validity of the medicare functional classification level system and utility of clinical outcome measures for aiding k-level assignment. *Archives of Physical Medicine and Rehabilitation*, 97(7):1053–1063, 2016.
- [198] Ulrike Bechtold, Daniela Fuchs, and Niklas Gudowsky. Imagining socio-technical futures – challenges and opportunities for technology assessment. *Journal of Responsible Innovation*, 4(2):85–99, 2017.
- [199] Ingunn Moser and John Law. *Making Voices’: New Media Technologies, Disabilities, and Articulation*, chapter 18, pages 491–520. MIT Press, 2003.
- [200] Ingunn Moser. Against normalisation: Subverting norms of ability and disability. *Science as Culture*, 9(2):201–240, 2000.
- [201] Coreen A. McGuire. *Measuring difference, numbering normal: Setting the standards for disability in the interwar period*. Manchester University Press, 2020.
- [202] Samuel K. Au, Jeff Weber, and Hugh Herr. Biomechanical design of a powered ankle-foot prosthesis. In *2007 IEEE 10th International Conference on Rehabilitation Robotics*, pages 298–303, 2007.
- [203] Dan Newman. U-m, humotech partner to bring open-source bionic leg to research labs. University of Michigan: Michigan News, 2021. Accessed 11/15/2021.

- [204] Hugh Herr. The new bionics that let us run, climb, and dance. TED Video, 2014. Accessed 11/15/2021.
- [205] Elissa D. Ledoux and Michael Goldfarb. Control and evaluation of a powered transfemoral prosthesis for stair ascent. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(7):917–924, 2017.
- [206] Michael Leddy and Aaron Dollar. Preliminary design and evaluation of a single-actuator anthropomorphic prosthetic hand with multiple distinct grasp types. In *2018 7th IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob)*, pages 1062–1069, 08 2018.
- [207] Study summaries - upper limb prosthetics. Ottobock. Accessed 11/15/2021.
- [208] Clinical outcome assessments in composite endpoint for upper extremity prosthetics. United States Food and Drug Administration. Accessed 11/15/2021.
- [209] Virgil Mathiowetz, Gloria Volland, Nancy Kashman, and Karen Weber. Adult norms for the box and block test of manual dexterity. *The American journal of occupational therapy : official publication of the American Occupational Therapy Association*, 39:386–91, 07 1985.
- [210] Johanne Desrosiers, Gina Bravo, Réjean Hébert, Élisabeth Dutil, and Louise Mercier. Validation of the box and block test as a measure of dexterity of elderly people: Reliability, validity, and norms studies. *Archives of Physical Medicine and Rehabilitation*, 75(7):751–755, 1994.
- [211] Thomas Platz, Cosima Pinkowski, Frederike van Wijck, In-Ha Kim, Paolo di Bella, and Garth Johnson. Reliability and validity of arm function assessment with standardized guidelines for the fugl-meyer test, action research arm test and box and block test: a multicentre study. *Clinical Rehabilitation*, 19(4):404–411, 2005. PMID: 15929509.
- [212] Britt H. Young. I have one of the most advanced prosthetic arms in the world — and i hate it. Input Magazine, 2021. Accessed 10/13/2021.
- [213] Human centered design (hcd). NIST. Accessed 08/08/2022.
- [214] V. N. M. Arelekatti and Amos G. Winter. Design and Preliminary Field Validation of a Fully Passive Prosthetic Knee Mechanism for Users With Transfemoral Amputation in India. *Journal of Mechanisms and Robotics*, 10(3), 04 2018. 031007.
- [215] Emma Reznick, Kyle Embry, and Robert D Gregg. Predicting individualized joint kinematics over a continuous range of slopes and speeds. *Proceedings of the ... IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics. IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics*, 2020:666—672, 2020.

- [216] Susannah Engdahl, Breanne Christie, Brian Kelly, Alicia Davis, Cynthia Chestek, and Deanna Gates. Surveying the interest of individuals with upper limb loss in novel prosthetic control techniques. *Journal of neuroengineering and rehabilitation*, 12:53, 06 2015.
- [217] Jathan Sadowski and Roy Bendor. Selling smartness: Corporate narratives and the smart city as a sociotechnical imaginary. *Science, Technology, & Human Values*, 44(3):540–563, 2019.
- [218] Choong Hee Kim, Shannon M. Danforth, Patrick D. Holmes, Daphna Raz, Darlene Yao, Asheesh Bedi, and Ram Vasudevan. Automated camera-based estimation of rehabilitation criteria following acl reconstruction. *arXiv:1810.11087*, 2018.
- [219] Marion Mundt, Arnd Koeppel, Sina David, Tom Witter, Franz Bamer, Wolfgang Potthast, and Bernd Markert. Estimation of gait mechanics based on simulated and measured imu data using an artificial neural network. *Frontiers in Bioengineering and Biotechnology*, 8, 2020.
- [220] Lara Carminati. Generalizability in qualitative research: A tale of two traditions. *Qualitative Health Research*, 28(13):2094–2101, 2018. PMID: 30043686.
- [221] Knut Lechler, Bertrand Frossard, Lynsay Whelan, David Langlois, Roy Müller, and Kristleifur Kristjansson. Motorized biomechatronic upper and lower limb prostheses—clinically relevant outcomes. *PM&R*, 10(9S2):S207–S219, 2018.
- [222] Katherine Raichle, Marisol Hanley, Ivan Molton, Nancy Kadel, Kellye Campbell, Emily Phelps, Dawn Ehde, and Douglas Smith. Prosthesis use in persons with lower- and upper-limb amputation. *Journal of rehabilitation research and development*, 45:961–72, 02 2008.
- [223] Niklas Kochdumper and Matthias Althoff. Sparse polynomial zonotopes: A novel set representation for reachability analysis. *IEEE Transactions on Automatic Control*, 66(9):4043–4058, 2020.
- [224] Matthias Althoff. Reachability analysis of nonlinear systems using conservative polynomialization and non-convex sets. In *Proceedings of the 16th international conference on Hybrid systems: computation and control*, pages 173–182, 2013.