

OPTIMIZING SENSORY STIMULATION IN HUMANS AFTER SPINAL CORD INJURY

A Dissertation
Presented to
The Academic Faculty

by

Jason Martin White

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
Department of Biomedical Engineering

Georgia Institute of Technology
And
Emory University
December, 2016

Copyright© 2016 by Jason White

OPTIMIZING SENSORY STIMULATION IN HUMANS AFTER SPINAL CORD INJURY

Approved by:

Dr. Stephen P. DeWeerth, Co-Advisor
Department of Biomedical Engineering
Georgia Institute of Technology

Dr. Keith E. Tansey, Co-Advisor
Department of Neurology
Emory University

Dr. Lena Ting,
Department of Biomedical Engineering
Georgia Institute of Technology

Dr. T. Richard Nichols
Department of Physiology
Georgia Institute of Technology

Dr. Boris I. Prilutsky
Department of Physiology
Georgia Institute of Technology

Dr. Magnus Egerstedt
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Date Approved January 8th, 2016

The contractile energy of the entire muscle can surpass that of the stimulus. It is thus that a tiny spark ignites a great mass of gunpowder, the energy of which is prodigious.

-Felice Fontana (1760)

Mémoires sur les parties sensibles et irritables du corps animal

To Livia, whose love and support made all this possible.

ACKNOWLEDGEMENTS

There are many, many people without whom this research would have been impossible, and I cannot possibly mention everyone in a few short pages. Nevertheless, I will attempt to show my appreciation to some of the people who directly supported me on this research.

I would like to thank my advisors, Steve DeWeerth and Keith Tansey, for all the work they did on my behalf, including reviewing drafts of manuscripts, pushing my boundaries academically, and supporting me both intellectually and financially throughout this long road.

I would like to thank my fellow lab-mates, Gareth Guvanasen, Jeonghee Kim, Hangu Park, and Edgar Brown, who gave great feedback on my research over the years. Hangu also was willing to give me feedback on an early draft of this document.

I would like to thank my undergraduate mentee, Ravinderjit Singh, who ran many of the simulations required to validate my optimization algorithm on the GA Tech PACE cluster.

I would like to thank the current and former members of the Hulse Spinal Cord Injury Lab at the Shepherd Center. Notably, Lauren McCollough, Joy Bruce, Leslie VanHiel, and Adam were instrumental. They took primary responsibility for recruiting subjects, preparing subjects for the experiments, and providing feedback on my approach. They assisted with managing the IRB documents and process. I would like to thank Edelle Field-Fote, and Shepherd Center more generally, for allowing me to recruit from their patient population, use their facilities, and for all the support they provided me.

I would to thank my funding sources. Georgia Tech and the biomedical engineering department supported my research through bridge funding. Khalifa University supported

my research for >2 years. NIH NINDS supported my research, and the founding grant from the Hulse family for the Hulse Spinal Cord Injury helped to purchase much of the equipment I used for this research.

I would like to thank Vafa Bayat, for his detailed reading and feedback on an earlier version of this document.

I would like to thank my kids, Aiden and Marcus, for the joy they have provided during the hard times, and the opportunities to explain what those squiggly lines on the computer screen were.

I would like to thank my parents, Douglas and Carolynne White, for all their support over the years to get to this point. They created me, supported me emotionally and financially through many years of primary and secondary education, and gave greatly of their time and support to foster my interest in science and engineering. More specifically, they offered and completed a reading of every chapter in this document. There would certainly be many more formatting and grammatical errors if not for their assistance.

Lastly, none of this work, not even my status as a graduate student, would have been possible without the support of my wife, Livia Van. She has supported me throughout all the triumphs and frustrations inherent in research and life. She has consistently given me feedback and advice regarding both big and small issues in my research. She took care of me and the kids so that I could write. She supported our family financially during my graduate work. She read and gave feedback on my writing. I am lucky to have such a great partner in this life.

TABLE OF CONTENTS

Acknowledgements	v
List of Tables	xiii
List of Figures	xiv
List of Symbols and Abbreviations	xv
Summary	xvi
Chapter 1: Introduction	1
1.1. Motivations	1
1.2. Problem Statement.....	4
1.3. Contributions	4
1.4. Overview of Remaining Chapters	5
Chapter 2: Relevant Background	6
2.1. The Physiology of Walking	6
2.1.1. Supraspinal Role	7
2.1.2. Central Pattern Generators	8
2.1.3. Sensory Afferents.....	11
2.2. Spinal Cord Injury.....	17
2.2.1. Spinal Cord	17
2.2.2. Pathophysiology	19
2.2.3. Clinical Assessment of Spinal Cord Injury	20

2.2.4.	Epidemiology.....	20
2.2.5.	Measuring the Quality of Walking after SCI	21
2.2.6.	Existing Approaches for Recovery.....	22
2.2.7.	Sensory Input after Spinal Cord Injury	24
2.3.	Sensory Stimulation and Walking	28
2.3.1.	Potential for Fine-tuned Control and Assistance.....	30
2.3.2.	Sensory Stimulation with Functional Electrical Stimulation	30
2.3.3.	Sensory Stimulation with Physical Therapy	31
2.3.4.	Spinal Cord Stimulation.....	32
2.3.5.	Cutaneous Nerve Stimulation.....	36
2.3.6.	Time-Varying Effects in the Nervous System from Stimulation	54
2.4.	Optimization	55
2.4.1.	Derivative-Based Optimization	56
2.4.2.	Derivative-Free Optimization	56
2.4.3.	Global Optimization	57
2.4.4.	Optimization with Noise	58
2.4.5.	Response Surface Based Optimization	60
2.4.6.	Time-Varying Optimization	64
2.4.7.	Closed-Loop Experimentation in Neuroscience	66
Chapter 3:	A system to optimize sensory Stimulation in Humans after Spinal Cord Injury	
		70
3.1.	Introduction	70

3.2.	System Design	75
3.2.1.	Intervention: Sensory Electrical Stimulation.....	76
3.2.2.	Assessment: A Real-Time Walking Metric.....	78
3.2.3.	Optimization: Closing the Loop.....	79
3.3.	Experimental Results	86
3.3.1.	Experimental Protocol	86
3.3.2.	Randomized Sampling vs Optimized Sampling	89
3.3.3.	Confidence Analysis.....	90
3.3.4.	Sensitivity Analysis.....	92
3.4.	Discussion.....	93
3.4.1.	Optimization was Successful.....	94
3.4.2.	Limitations of Study	95
3.5.	Conclusions.....	96
Chapter 4:	Optimizing Sensory Stimulation was Successful	97
4.1.	Introduction	97
4.2.	Methods	100
4.2.1.	Subjects and Preparation	101
4.2.2.	Stimulation	102
4.2.3.	Data Collected.....	105
4.2.4.	Force Metric	106
4.2.5.	Optimization	107

4.2.6.	Statistics.....	108
4.3.	Results	109
4.3.1.	Comparison of Stimulation Protocols.....	109
4.3.2.	Optimal Stimulation Parameters	112
4.3.3.	Time-Variance.....	113
4.3.4.	Model Fit: Sensory Site and Injury Details	115
4.4.	Discussion.....	116
Chapter 5:	Insights from Sensory Stimulation	120
5.1.	Introduction	121
5.2.	Methods	124
5.2.1.	Subject Recruitment and Preparation	125
5.2.2.	Sensory Stimulation	127
5.2.3.	Data Collected and Post Processing	131
5.2.4.	Experimental Protocol	132
5.2.5.	Features.....	132
5.2.6.	Statistics.....	133
5.3.	Results	134
5.3.1.	Forces	135
5.3.2.	Short-Term Force Analysis.....	138
5.3.3.	Multiple Feature Analysis	140
5.3.4.	Muscle Activation (EMG)	141

5.3.5.	Assistive Force vs EMG	143
5.3.6.	History Dependence.....	145
5.3.7.	Individual Differences	147
5.3.8.	Onset Response	149
5.4.	Discussion.....	151
5.4.1.	Reductions in Assistive Force.....	152
5.4.2.	Muscle Activation and Force	153
5.4.3.	Individuality and History.....	155
5.4.4.	Onset-Response and Termination-Response.....	156
5.4.5.	Conclusions.....	158
Chapter 6:	Efficient Global Optimization of Time-Varying and Noisy Systems	159
6.1.	Introduction	160
6.2.	Overview of Approach	163
6.2.1.	Measuring Time-variance	163
6.2.2.	Measuring Noise	165
6.2.3.	Time-varying, Noisy, Efficient Global Optimization (TVN-EGO).....	165
6.2.1.	Illustrative Example	171
6.3.	Simulated Experiments.....	173
6.3.1.	Alternative Algorithms	173
6.3.2.	Test Functions.....	175
6.3.3.	Empirical Comparisons.....	180

6.4. Discussion.....	188
6.4.1. Dimensional Scalability	190
6.4.2. Computational Efficiency	191
6.4.3. Possible Extensions	192
6.5. Conclusions.....	193
Chapter 7: Conclusions.....	194
References	191

LIST OF TABLES

Table 4-1. Subject Details and Stimulation Sites tested on each subject. LEMS: Lower extremity motor score. AIS: ASIA Impairment Scale. ¹This subject was able to get voluntary movement in a non-key muscles greater than 4 segments below the injury, but had no other voluntary control. ²LEMS was unavailable for these subjects. TSCS: Transcutaneous Spinal Cord Stimulation. PTNS: Posterior Tibial Nerve Stimulation. CPNS: Common Peroneal Nerve Stimulation. 106

Table 5-1. Subject Details. LEMS: Lower extremity motor score. AIS: ASIA Impairment Scale. ¹This subject was able to get voluntary movement in a non-key muscles greater than 4 segments below the injury, but had no other voluntary control. ²LEMS was unavailable for these subjects. 126

Table 5-2. All experimental sessions completed, organized by subject and site stimulated. *This subject had a substantial response to tonic stimulation, despite the stimulus amplitude being set to only 22% of the CMAP. Params: Parameters. Stim: Stimulation. BWS: Body weight support. CMAP: Common motor action potential (the amplitude at which all muscles are recruited). 130

Table 6-1. Summary of variables used in TVN-EGO. Each variable is shown with a brief description. 167

LIST OF FIGURES

- Figure 3-1: Overview of the optimization approach. Each iteration consisted of the algorithm completing one loop. In the Assessment stage, force and angle data from each step in the Lokomat was converted into a measure of the quality of stepping (the force metric). In the Optimization stage, a maximum likelihood model was fit, and the optimum was found within that model. In the Intervention stage, the identified model-optimal stimulation parameters were converted into a pulse train which was then applied to the individual in the Lokomat (at the distal tibial nerve), and measurements were made again, starting a new iteration.76
- Figure 3-2: Pulse Train Parameterization. The pulse train was parameterized by three parameters: the frequency (stimuli per second), the duration (length of the pulse train, as a percentage of the gait cycle), and the gait phase (time in the gait cycle at the start of the pulse train). Steps were defined as beginning at the swing-to-stance transition on the right foot.77
- Figure 3-3: Experimental validation. The experiment consisted of a human SCI subject walking 496 steps in the Lokomat. The first half of all plots (first 261 steps, uncolored background) is the randomized-sampling period, which initialized the space for the optimization period (“optimizing” in the right half, consisting of the final 245 steps, gray background). The top panel (panel A) shows the frequency, phase, and duration of the pulse train vs the step number, the tested values are colored according to their metric value (i.e. how good the output was) in the bottom panel. The middle panel (panel B) shows the actual pulse trains (each hash is a pulse in the pulse train) from the stimulation parameters in the top panel. The pulse trains are shown with individual stimuli represented as hash marks. Every pulse train was repeated four times, but shown only once in the plots above (for ease of visualization). The bottom panel (panel

C) shows the force metric in response to the pulse trains vs the step number, with only the median for each iteration shown. In the force metric, a value of 1.0 corresponds to the average response for no stimulation. Most pulse trains reduced the force metric (improved stepping), but some combinations increased the force metric to above 1 (worse than no stimulation).....88

Figure 3-4. The optimal stimulation parameters, confidence regions, and input-to-output mapping. A model was created from the experimental data, including the expected value and standard deviation at every point in space, and the covariance between every combination of two points in the space. From this stochastic model, confidence regions (95% and 99%) of the location of the global minimum were calculated using a Monte Carlo simulation (Panel A on the left). To visualize the space near the optimal stimulation parameters, a 2D slice of the expected value space was made (Panel B on the right). A duration of 36% was chosen, as it is the optimal duration. The areas with better (i.e. lower) force metrics are concentrated at the highest frequency tested (top part of panel B has both the lightest colors), with starting times (phases) near the stance-to-swing transition. While most areas in the parameter-space reduced the force metric (improved stepping), some combinations increased the force metric to above 1 (worse than no stimulation).....91

Figure 3-5. Sensitivity of the optimum. The optimum is shown with the large dot. The solid line shows the estimated response, the shaded area above and below the dark line shows the standard deviation. The dotted red line shows the second derivative, calculated at the optimum, and the values shown above the optimum report the sensitivity in both normalized (by the range of allowed values), and by standard units (e.g. Hz for frequency).....93

Figure 4-1. Conceptual Overview. The subject was walking in the Lokomat, and I recorded the amount of force assistance applied to maintain the walking pattern. This

force assistance was summarized into the force metric. A model was created of all previous stimulation parameter tests and the next test point was selected via an optimization within the model. The next test point was translated into individual stimuli, depending on the parameterization being used (tonic or pulse-train stimulation). The stimuli were delivered to one stimulation site in each session (on both legs in the case of pulse-train stimulation), and more force metrics data points were recorded, beginning the loop again. The angles of the limbs were used to define the step cycle, and the steps were divided at the beginning of stance on the right leg. Pulse trains were parameterized using the gait phase of the beginning of the train, the frequency of stimuli, and the duration (as a percentage of one complete step cycle). Tonic stimulation was parameterized by the frequency of stimuli.101

Figure 4-2. Optimal Stimulation as compared to no, random, and standard stimulation.

Panel A shows an example step (Step #388), from subject #6, with a TSCS frequency of 100 Hz. Hip and knee forces are shown. Dark blue shows the current step, with portions of the previous and next step shown in light blue. This steps are combined (in order) with the other 488 steps from the same session in panel B. The bottom plot of panel B shows the TSCS frequencies which produced the force metric values in the main panel. The right plot of panel B shows the distribution of force metric values according to stimulation condition, where pink is optimal stimulation. In the top plot of panel C, the session from panel B is combined with 14 other TSCS sessions to produce the force metric distribution for the entire dataset. The second plot in panel C shows the distribution across all posterior tibial stimulated steps and sessions. The third plots in panel C shows the distribution across all common peroneal stimulated steps and sessions. In all three stimulation sites the optimal stimulation (colored shading) was statistically significantly lower than all three other conditions (no, standard, and random stimulation). The number of steps in each set was 6,740 (TSCS), 8,165 (posterior tibial stimulation), and 2,640

(common peroneal stimulation). I tested significance using a Wilcoxon Rank Sums test, and quantified effect size (“d” in the above) using Cliff’s Delta (range of -1 to 1, with 1 being the largest difference). Std. = Standard, Rnd. = Random.111

Figure 4-3. Distribution of Optimal Stimulation Parameters. The time-varying optima from each session were used to estimate the probability distribution of the optimal stimulation parameters in that session (the colored lines in the background of each plot). Then, the average of the distributions of all sessions was calculated to estimate the global distribution of optimal stimulation parameters (black line in the foreground in each panel). Transcutaneous spinal cord stimulation is on the top row. Posterior tibial nerve stimulation is on the second row. Common peroneal nerve stimulation is on the third row. The first column shows frequency of stimulation for all three stimulation sites. The second and third columns show phase of stimulation and duration of the pulse train respectively.113

Figure 4-4. Examples of Time-Variance. Stimulation Frequency is shown by the color of the circles, black is zero or low frequency stimulation, yellow is high frequency stimulation. These plots show different ways that subjects may vary their response in time. In all the plots pay special attention to the same color circles (same frequency of stimulation) at different times. In the top left, an example plot shows a non-time-varying dataset. The top right shows a habituating subject. The bottom left plot shows a subject requiring more and more force assistance. The bottom right shows a subject requiring less and less force assistance.115

Figure 4-5. Model performance compared to injury type. The performance of the optimization algorithm is limited by the ability to model the data. Percent variance explained is a standard measure of the quality of a model. The left scatter plot shows the results for all three stimulation sites organized by injury details (AIS classification, followed by LEMS). In some subjects, LEMS was not available, and they were excluded

from this figure. The right three plots show the distribution of explained variance based on the sensory stimulation site. TSCS: Transcutaneous spinal cord stimulation. PTNS: Posterior Tibial Nerve Stimulation. CPNS: Common Peroneal Nerve Stimulation. AIS: ASIA Impairment Scale. LEMS: Lower Extremity Motor Score.116

Figure 5-1. Overview of methods. In each experimental session, one stimulation site was studied. Tonic stimulation was applied to the transcutaneous spinal cord stimulation (TSCS) site. Pulse train stimulation was applied to the common peroneal and posterior tibial nerve stimulation sites. The subject walked in the Lokomat while each set of stimulation parameters was applied for 4–10 steps. The angles of the Lokomat joints and the amount of force the Lokomat applied were recorded, along with electromyography from four muscles on each leg. Each session included 130–200 random-stimulation-parameter steps. EMG: Electromyography. Quad: Quadriceps muscle. Ham: Hamstring muscle. Tib Ant: Tibialis Anterior. Tri Sur: Triceps Surae.125

Figure 5-2. Force Features vs Stimulation Parameters. Panel A shows how distributions were created. First, the feature was calculated (mean-squared) for each step for each sensory stimulation site. The best and worst 20% of the feature-values were collected separately. For each feature-value, I cross-referenced the stimulation-parameters, and created best-20% and a worst-20% probability distributions of stimulation parameters for each feature. The left column of the lower plots (panel B) shows the specific force features calculated for each row. The dark blue shows when the force is included in the feature. The light blue shows when it is excluded from the feature. The black line shows the joint angle. In panel C, each column shows one stimulation parameter. The best and worst 20% are plotted with a blue fill and a dotted red line, respectively. In all cases the distribution of the best and worst 20% were normalized to the overall dataset to eliminate sampling bias. To interpret these plots, flatter distributions mean less predictive power.

Larger peaks and lower valleys mean more predictive power. As in other plots, stance was defined to be the first 60% of gait, and swing was defined to be the last 40%.137

Figure 5-3. Short-term Force Features. I created short-term force-features to analyze the effect of TSCS frequency and cutaneous nerve stimulation phase on specific times/phases of gait. The force from the hip and knee were combined, and a distribution of the stimulation parameters for the best-20% was created (panel A) for both TSCS-frequency and PTNS-phase. The profile of that distribution was extracted for every feature and plotted in panel B. In the left column of Panel B, the black line shows the timing of each of the 20 overlapping, short-term force-features (each delayed 5% in the gait cycle from the preceding one). The second column of panel B shows TSCS-frequency distributions for each feature. The third and fourth columns show the PTNS-phase distributions for each feature. The two columns show the same data, with the contralateral phase shown delayed by half of the gait cycle. On the bottom of panel B, there are 4 icons of people, to clarify the stimulated leg as compared to the measured leg. A red, dotted line was added to assist the reader visualize the peaks. St: Stance. 139

Figure 5-4. TSCS Frequency vs both hip-swing and hip-stance forces. The left column shows the two force features compared (hip-swing and hip-stance force features). The middle column on the top and bottom rows show the distribution of those two features individually. They are combined in the middle column center plot to show how the hip-stance and hip-swing forces relate to one another. In the cases where both features have good values (bottom 33%), the stimulation frequency required to produce those points are plotted on the right column.140

Figure 5-5: Selected EMG features. The left column (panel A) shows the feature being calculated. In each case a time-window is applied to one EMG channel. Each row shows one feature. On the feature plot, dark blue shows the portion of the EMG channel that is included in the calculation, and light blue shows the excluded EMG data. The columns to

the right (panel B) show the stimulation parameter distribution of the best and worst 20% of the data with respect to each feature (row), calculated in the same manner as the force distributions in the previous figures, with one difference – EMG features are “better” when they are larger (assistive force features are “better” when they are smaller). In the columns showing the stimulation parameters, some plots were excluded due to too much stimulation artifact on those channels while recording EMG (marked with an ‘X’). These excluded channels were identified empirically. On the distribution plots, dark-blue fill shows the best 20%, and red, dotted line shows the worst 20%. As in other plots, stance was defined to be the first 60% of gait, and swing was defined to be the last 40%. EMG: Electromyography. Quad: Quadriceps muscle. Ham: Hamstring muscle. TBA: Tibialis Anterior. TrSu: Triceps Surae. 143

Figure 5-6. Force-to-EMG Comparison. In Panel A, the tibialis-anterior EMG-feature and combined-joint force feature are shown. The EMG feature summarizes the early swing activation of the tibialis anterior. The force was calculated from the same time-period in the gait cycle to compare whether lower force and higher EMG tended to be associated. In panel B, the results for TSCS are shown. In the first plot of panel B, the scatter plot comparing the EMG metric (TBA early-swing EMG) to the force metric (both joints, early-swing force) is shown. The 20% best (lowest) points in the force metric are extracted (blue) from the rest of the dataset (grey), and the resulting distribution of those 20%-best points is shown in the 2nd plot in terms of the EMG metric. Panel C shows the results from posterior tibial nerve stimulation. TBA: Tibialis Anterior. 145

Figure 5-7. Comparison of Stimulation parameters to History as a parameter. Panel A, at the top, shows the feature I compared (hip-swing force). In Panel B (the bottom two rows), TSCS (row 1) and posterior tibial nerve stimulation (row 2) are shown, along with the distribution of the best feature values with respect to experimental history. History, in

the far right column of Panel B, was measured by the number of steps the subject had already performed. 147

Figure 5-8. Individual Variation. Panel A (left column) shows the features. The top 4 features are force features (hip-stance, hip-swing, knee-stance, and knee-swing), and the bottom 2 features are EMG features (tibialis-anterior-early-swing and triceps-surae-stance). Panel B (2nd column) shows the distribution of TSCS frequency from the subject population as whole. Panel C (3rd–5th columns) show three example of individual distribution of TSCS frequency that differ from the population distribution. 149

Figure 5-9. Onset and termination response of tibialis anterior posterior EMG evoked from PTNS (pulse trains). In panel A, it shows the onset response. The red, striped bar shows the timing of the beginning of the pulse train. The median and 75th percentile are shown. Panel B shows the termination response. The red, striped bar in panel B shows the end of the pulse train. Panel C shows the onset response compared to the frequency. 10 Hz bands of frequency were collected, forming 10 bands between 0 and 100 Hz (e.g. a band including 20-30 Hz stimulation). Panel D shows the termination response versus stimulation frequency. All waveforms were normalized prior to combining into percentiles. 151

Figure 6-1. Illustrative example of TVN-EGO across 3 iterations during optimization of a 1D function. The first column show iteration 1, the second column shows iteration 2, and the third column shows iteration 3. The plots in the top row show the current function (dark blue line), the noise (filled light blue area), and the previous iteration's function (red line). The second row shows the samples thus far, along with the actual function. The third row shows the model's estimate of the function (dark red line) at the current time and the standard deviation in the estimate (filled pink area). The actual function is shown in the background (dark blue line). The fourth row shows the estimate of expected improvement for all points, with the selected maximum highlighted with a circle. That fifth

row shows new sample of the function, in terms of both the estimate of what the value would be, and the actual amount measured. 172

Figure 6-2. Performance of optimization algorithms in the context of increasing time-variance. The rows show different test functions, with the number of dimension labeled. The columns show the three optimization algorithms. In each plot, individual experiments (simulations) are shown as black circles, and the trend line, calculated from logistic function, is fit. The trend line is shown in blue. In all cases, the time-variance is normalized with respect to the spatial standard deviation (as discussed in the “Approach” section). The performance is the percentage of time the estimated best is within 10% of the actual best across all input dimensions between the 80th and 120th iterations. 183

Figure 6-3. Performance of optimization algorithms in the context of increasing noise. The rows show different test functions, with the number of dimension labeled. The columns show the three optimization algorithms. In each plot, individual experiments (simulations) are shown as black circles, and the trend line (blue), calculated from logistic function, is fit. The noise is normalized (as discussed in the “Approach” section). The performance is the percentage of time the estimated best is within 10% of the actual best across all input dimensions between the 80th and 120th iterations. 186

Figure 6-4. Performance in the context of noise and time-variance across multiple test functions and dimensions. Each plot shows a heat map of a specific test function. One-dimensional test functions are on the top row, two-dimensional test functions are on the second row, three- and four-dimensional test functions are on the bottom row. In all cases success was defined based on the performance metric of the number of iterations in which the estimated optimum was within 10% of the actual optimum. 188

LIST OF SYMBOLS AND ABBREVIATIONS

ASIA: American Spinal Injury Association

AIS: ASIA Impairment Scale

CPNS: Common Peroneal Nerve Stimulation

EGO: Efficient Global Optimization

PTNS: Posterior Tibial Nerve Stimulation

SCI: Spinal Cord Injury

SKO: Sequential Kriging Optimization

SPSA: Simultaneous Perturbation Stochastic Analysis

TSCS: Transcutaneous Spinal Cord Stimulation

SUMMARY

Sensory stimulation has shown promise in improving human walking after spinal cord injury (SCI). Previous studies have demonstrated some improvement with open-loop, non-individualized sensory stimulation, but after SCI, there are many unique, individual changes in sensorimotor processing. These changes make a priori identification of the best sensory stimulation pattern difficult for any given individual. Real-time optimization provides a solution to this individuality problem, through optimizing sensory stimulation parameters for a given subject in on-line (in real-time). In this research, I developed an approach to optimize sensory stimulation to maximally assist human walking after incomplete SCI. To do so, I had to develop and validate a novel optimization algorithm for globally-optimizing noisy, time-variant, black-box systems, while maximizing the information gained from each test (experiment). I optimized sensory stimulation across a range of SCI subjects, across multiple sensory stimulation sites, and with different stimulation parameterizations. In all subjects and stimulation sites, the optimal stimulation protocol produced better walking (i.e. less external force assistance was required) than three alternative stimulation protocols: an industry-standard stimulation protocol, a no-stimulation protocol, and a random-stimulation protocol. The optimization approach minimized the total force required from an assistive orthosis, and post-hoc analysis of the optimization sessions produced a better understanding of how stimulation parameters affected specific gait features (e.g. hip forces during swing).

Transcutaneous spinal cord stimulation (TSCS) frequency had divergent effects on the stance and swing phases – high frequencies tended to assist with swing, but low frequencies tended to assist with stance. For the two peripheral nerve stimulation sites (posterior tibial and common peroneal nerves), the optimal gait-phase for stimulation was generally after mid-stance and before early swing. There was some variability within

this time-range depending on the specific feature under study. Experimental history (i.e. time spent walking/time spent being stimulated) proved to be as important a predictor as any of the stimulation parameters.

CHAPTER 1: INTRODUCTION

This chapter briefly discusses the motivations, problem statement, and specific contributions from my research on optimizing sensory stimulation to assist humans to walk after spinal cord injury walk.

1.1. Motivations

The research presented in this dissertation was motivated by the following: (1) Sensory stimulation can assist people with spinal cord injury (SCI) to walk. (2) Reflex responses, however, have significant variability in the general population, and that variability is increased after SCI. (3) Real-time optimization is a new approach in medicine, that can be applied to identify the best stimulation pattern on an individual basis.

Spinal cord injury has a worldwide prevalence of 250-750/1,000,000 people (Wyndaele and Wyndaele 2006, Chiu et al. 2010), corresponding to an estimated 1.8-5.4 million people in the world, including approximately 270,000 spinal cord injured individuals in the United States (Center 2012). Of those cases, approximately 22% are complete paraplegics and 21% are incomplete paraplegics. Even with current clinical interventions, many people with SCI do not achieve the level of ambulation they would like (Anderson 2004). There is little we can offer to those with complete or the most severe incomplete injuries to recover the ability to walk.

In uninjured individuals, spinal neurons integrate sensory feedback from the legs with descending commands from the brain to produce walking (Kiehn and Butt 2003, Rossignol, Dubuc, and Gossard 2006). After SCI, spinal neurons are (partially or completely) disconnected from the brain, leading to the loss of voluntary locomotion. However, the injury generally leaves the sensory connections to the spinal neurons

intact, and the effects of sensory input on those neurons actually become stronger (Lhermitte 1919, Little and Halar 1985, Zehr, Komiyama, and Stein 1997, Rossignol and Frigon 2011). This connection remains sufficient for sensory input to recruit the central pattern generators (CPGs), which generate the walking pattern (Kiehn et al. 1998, Edgerton et al. 2001).

Sensory inputs from the legs have numerous and substantial effects on walking (Rossignol, Dubuc, and Gossard 2006). Many studies have shown that Sensory stimulation can both recruit and modulate the CPG (Marchetti, Beato, and Nistri 2001, Hultborn 2001, Schomburg et al. 1998, McCrea 1998, 2001, Roberts et al. 1998, Sillar and Roberts 1992). Studies have explored the use of sensory stimulation to assist people with SCI to walk, and have found it effective across multiple contexts and performance metrics (Ladouceur and Barbeau 2000b, Ladouceur and Barbeau 2000a, Granat et al. 1993, Ragnarsson 2008).

Despite the possibilities of sensory stimulation, sensory feedback to assist, produce, or control a walking pattern has not been rigorously optimized previously. If the correct sensory stimulus pattern can be identified, the walking pattern will not just be improved, but can be maximally improved.

However, identifying the optimal stimulation parameters is complicated by the variability in the SCI population. Even neurologically normal individuals have significant variability in their responses to sensory input. Identified sources of reflex-response variability include age (Evans, Harrison, and Stephens 1990, Gibbs et al. 1999), level of activity (Loeb 1993), and athletic history (Gruber et al. 2007, Wolpaw and Tennissen 2001, Maffiuletti et al. 2001). After SCI, that variability is compounded by additional sources, including SCI level (Poirrier et al. 2004, Dietz et al. 1999), severity (Rossignol et al. 1996), and time since injury (Hiersemenzel, Curt, and Dietz 2000).

Therefore, I pursued an approach based on individual, real-time optimization of sensory stimulation. Real-time optimization of clinical interventions is relatively unexplored, and the best practices remain unidentified. Real-time optimization, in this context, means to measure the performance of an intervention and modify the parameters of that intervention based on performance metrics. This process must be done quickly (e.g. in seconds), such that tens or hundreds of rounds of improvements can be made to the intervention within a relatively short time-period (e.g. a visit to the doctor's office). These approaches are rare in both neuroscience and clinical research, and researchers have only begun to explore the possibilities in the last two decades. A few previous studies have utilized real-time optimization in the fields of functional electrical stimulation (Kostov et al. 1995, Lynch and Popovic 2008), neurophysiological responses (Földiák 2001, Edin et al. 2004, O'Connor, Petkov, and Sutter 2005, Benda et al. 2007, Lewi, Butera, and Paninski 2008), brain-machine interfaces (Gürel and Mehring 2012, Fruitet et al. 2012, Fruitet et al. 2013), and epidural spinal cord stimulation (Desautels 2014, Desautels et al. 2015).

Beyond just identifying the optimal sensory stimulation patterns, my objective was to further characterize human sensorimotor responses after SCI. Human sensorimotor responses are under-studied as compared to animal sensorimotor responses. The optimization process in my research required testing many different stimulation patterns across multiple subjects and multiple stimulation sites. Ultimately, this dataset included >20,000 stimulated-steps, providing the opportunity to perform in-depth analysis to discover how specific stimulation patterns produced responses and affected specific features of gait.

1.2. Problem Statement

The goal of this research was to develop an approach to optimize sensory stimulation to maximally-assist spinal-cord-injured individuals to walk, identifying the best stimulation patterns for both individuals and the entire subject population. A secondary goal was to map the relationship of sensory-stimulation to specific features of gait.

1.3. Contributions

There were four main contributions of my research:

- I created an approach to individually optimize sensory stimulation in real-time to assist individuals to walk after SCI (Chapter 3). This approach was demonstrated in a system that could stimulate with multiple types of stimulus patterns, analyze walking in real-time, and maintain an up-to-date model of the sensory-input-to-motor-output mapping.
- I identified the optimal sensory stimulation patterns for individuals and for the entire subject population (Chapter 4). The optima produced are more precise than previous reports in the literature. Multiple stimulation sites were compared for their efficacy in assisting their walking pattern.
- I calculated the ideal sensory stimulation patterns for a range of gait-features that might be of interest to a clinical researcher through post-hoc analysis of force and EMG data (Chapter 5).
- I developed and validated a new algorithm for time-varying, noisy global optimization, exploring the limitations of this approach with a suite of test functions across multiple levels of noise and time-variance (Chapter 6). I empirically identified the limits of the approach, and compared the algorithm with two other potential algorithm choices.

In addition, there were several minor contributions. My study represents one of the largest transcutaneous spinal cord stimulation studies to date. I developed a new walking metric. I analyzed an onset stimulus response to cutaneous nerve stimulation that has never been thoroughly studied. My comparisons between optimization algorithms will better quantify the differences between the various classes of optimization approaches. Lastly, I demonstrated a novel model space for optimization for rotationally symmetric problems.

1.4. Overview of Remaining Chapters

The remainder of the thesis is organized into a background chapter, four experimental/analysis chapters, and a short conclusion. The background chapter (chapter 2) introduces the reader to the relevant topics in order to understand the context of the work. The first experimental chapter (chapter 3) introduces the system that was used to optimize, how it works, and a validation involving one subject. Chapter 4 discusses the results across all subjects and all sensory stimulation sites, quantifying the success of individual optimization. Chapter 5 presents an in-depth analysis of how sensory stimulation affects different aspects of gait, exploring stimulation parameters at multiple sensory stimulation sites with respect to how the parameters affect different gait features. Chapter 6 validates the optimization algorithm that was developed for this system, comparing the algorithm against alternatives, and quantifying its limitations.

CHAPTER 2: RELEVANT BACKGROUND

This chapter describes the relevant literature for the dissertation. It is split into four sub-chapters: the physiology of walking, spinal cord injury, sensory stimulation and walking, and optimization. In each section, more literature exists than can reasonably be covered in this chapter, but every attempt was made to include the findings necessary to understand the context of this work.

2.1. The Physiology of Walking

In uninjured individuals, the movements and motor commands for walking are generated by three parts of the nervous system working in concert: the brain, central pattern generators in the spinal cord, and sensory feedback. The brain supplies the drive, such as turning on the system to start walking forwards. Central pattern generators (CPGs) in the spinal cord produce oscillating patterns for walking. Sensory feedback modulates those patterns by reporting the biomechanical interactions between the body and the environment to the central nervous system (Rossignol, Dubuc, and Gossard 2006). This section will review the physiology of walking in the context of this thesis, but for more depth, there are many excellent reviews on walking in recent years, for example: Zehr and Duysens (2004), Kiehn and Butt (2003), and Rossignol, Dubuc, and Gossard (2006).

In the most basic sense, walking can be divided into the swing and stance phases. The stance phase is defined as the period of time when the foot is touching the ground, and the swing phase is defined as the period of time when the foot is “swinging” through the air. In humans, the stance phase begins when the heel touches the ground and ends when the toe lifts off. The precise timing of these phases vary, depending on the speed

of walking and the individual, but the stance phase is approximately 60% of the gait cycle, and the swing phase is approximately 40% of the gait cycle.

Much of the progress in understanding how the nervous system implements the walking task began with Sir Charles Sherrington (Sherrington 1906, 1910) and Graham Brown (Brown 1911a, b, 1914) in the early 1900s . Animal models were used to understand the necessary and sufficient parts of the nervous system required to walk. The use of animal models has continued throughout the 1900s, and as a result of this effort, substantial progress has been made in understanding how we walk.

Since the studies on walking have mostly been based on animal models, scientists currently understand quadrupedal walking (e.g. cat walking) better than human walking. There are differences between bipedal walking and quadrupedal walking (Rossignol and Frigon 2011), but, fortunately, many observations have been found to apply to both (Dietz 2002, Duysens and Van de Crommert 1998, Van de Crommert, Mulder, and Duysens 1998).

2.1.1. Supraspinal Role

The brain's role in walking is complex. In the 1960s, a portion of the brain called the mesencephalic locomotor region (MLR) was identified (Shik, Severin, and GN 1966). If the MLR is stimulated with a constant frequency, an animal will begin walking (Jordan, Pratt, and Menzies 1979). However, the MLR does not directly connect to the spinal cord, rather, it communicates with the spinal cord through the reticulospinal tract (Steeves and Jordan 1980, Garcia-Rill and Skinner 1987). The importance of the reticulospinal neurons in activating locomotion appears to be evolutionarily conserved, with examples from lampreys (Di Prisco et al. 2000), mammals (Garcia-Rill and Skinner 1987), and some limited results supporting its role in humans (Hathout and Bhidayasiri

2005, Thornton et al. 2002). Recent work has divided cells in the MLR into areas associated with exploratory, appetitive, and defensive locomotion, based on the behavioral context in which each subset of cells is active (Jordan 1998).

Other regions of the brain are involved in “skilled” walking (e.g. ladder-walking). For example, the pyramidal tracts (from the cortex) are much more active during skilled walking tasks (Marple - Horvat and Armstrong 1999).

2.1.2. Central Pattern Generators

Graham Brown, in an illustrative series of experiments, showed that the mammalian spinal cord could produce locomotor-like patterns (called fictive locomotion) entirely without supraspinal or sensory input, demonstrating unequivocally that walking is not just a series of reflexes and that the spinal cord contained the neural circuitry necessary for walking (Brown 1911a). He proposed an organization of neurons based on mutual inhibition called half-center oscillators (Brown 1914), which still inspires many theoretical models (McCrea and Rybak 2007, 2008). Since the early 1900s, substantial work has been done to better understand the central pattern generators (CPGs), reviewed in Grillner (1981), Duysens and Van de Crommert (1998), and Kiehn and Butt (2003). In general, the CPG is modeled as flexor and extensor oscillators on each side of the spinal cord, with mutual inhibition between the flexor and extensor on the same side, and inhibition between the flexor-flexor and extensor-extensor parts between opposite sides.

CPGs drive many of the rhythmic movements we perform. CPGs produce rhythmic movement for both the arms and legs during walking, and there may even be a common core of CPGs connected to both sets of limbs in both humans and animals (Dietz 2002, Zehr et al. 2007). CPGs not only support normal walking, but they also produce the

patterns for other types of walking, such as backwards walking (Lamb and Yang 2000). There is an extensive literature on CPGs involved in breathing, including the demonstration of an incredible robustness of the pattern generation even if a substantial number of neurons are lost (Von Euler 1983), and a great flexibility in the patterns a small group of neurons can produce (Lieske et al. 2000).

2.1.2.1. Central Pattern Generators in Humans

Much of the early work studying the CPG was completed in animals, but it has been more difficult to demonstrate their existence directly in humans. Therefore, researchers have sought out examples of the spinal cord producing rhythmic motor output without any supraspinal input or phasic sensory input.

Scientists have observed many specific subjects who had a complete loss of supraspinal input, but they still produced rhythmic movement (involuntarily). As with many other areas of medicine, large-scale wars have progressed our understanding though the regrettably large number of injured soldiers. After WWI, a soldier with a complete spinal cord injury was shown to produce rhythmic output (Lhermitte 1919); and after WWII, an example was documented of stepping movement in a soldier after complete spinal cord injury (Kuhn 1950). In earlier observations, rhythmic movements were seen following human decapitation (Luys 1893), and later, following brain death (Hanna and Frank 1995). Although spontaneous rhythmic motion after complete human SCI has remained rare, other researchers have noted specific examples (Bussel et al. 1988, Calancie et al. 1994, Bussel et al. 1996).

In a different type of reduced supraspinal input, some researchers have studied the immature nervous systems of infants and even fetuses still in the womb. The premise behind these studies is that infants have immature brains, which will not send coherent

commands to the spinal cord yet, so if the spinal cord produces rhythmic movement, it can be taken as evidence of a human central pattern generator. Fetuses have been observed to produce rhythmic leg movement while still in the womb (Rayburn 1995). Infants cannot walk independently, but they can walk with weight support, although their walking pattern looks different than the walking pattern of an adult (Forssberg 1985). Human infant walking is more similar to a cat's walking and sensory input plays a larger role (Pang and Yang 2000). There are many other studies supporting the existence of active CPGs in human infants (Yang, Stephens, and Vishram 1998, Pang and Yang 2002), including walking in both forward and backward directions (Lamb and Yang 2000)

Researchers have also attempted to directly activate the human CPG through either tonic spinal cord stimulation or tonic sensory input. Tonic stimulation is stimulation that is always on, unchanging on the timescale of the walking pattern. Tonic inputs show that the nervous system can produce a rhythmic output when no rhythm is present on the input, supporting the existence of CPGs. In a series of experiments using tonic leg vibration in healthy subjects, researchers were able to generate involuntary stepping-like movements (Selionov et al. 1997, Kazennikov et al. 1997). In a number of studies on humans with complete spinal cord injury, tonic spinal cord stimulation can produce stepping-like movement with reciprocal muscle activity (Rosenfeld et al. 1995, Gerasimenko and Makarovskiy 1996, Gerasimenko et al. 1996, Shapkova, Shapkov, and Mushkin 1997, Dimitrijevic, Pinter, and Sherwood 1997, Dimitrijevic, Gerasimenko, and Pollo 1997, Harkema et al. 2011).

Lastly, some characteristics of reflexes have been interpreted to give evidence for a human CPG. Contralateral and ipsilateral reflexes show an inhibitory relationship (Bussel et al. 1989), similar to Brown's half-center oscillator theory of how CPGs work. In humans with and without spinal cord injury, many reflexes are modulated during the

walking cycle (Yang and Stein 1990), even to the point that they can be completely reversed (Duysens et al. 1992), giving evidence that the spinal neural circuitry is capable of complex patterned modulation, and that sensory inputs are strongly coupled to the CPGs.

2.1.3. Sensory Afferents

Most sensory modalities have been shown to affect walking. From the head, visual (Sherk and Fowler 2001, Kennedy et al. 2003) and vestibular (Kennedy et al. 2003) inputs can affect gait trajectories. From the legs, cutaneous (Zehr, Komiyama, and Stein 1997, Zehr and Stein 1999, Rossignol, Dubuc, and Gossard 2006) and proprioceptive (Pearson 1995, Rossignol, Dubuc, and Gossard 2006) afferents have numerous and substantial effects on walking.

In the early 1900s, Sir Charles Sherrington began the early work of understanding how reflexes affect walking (Sherrington 1906, 1910). Due to the extensive role sensory feedback has in walking, some early researchers believed that walking could be explained as a series of reflexes (Sherrington 1906, 1910) before the discovery of the CPG (Brown 1911a).

Reflex stepping is not the primary means of generating a walking pattern, and animals can learn to walk in an altered fashion entirely without sensory input (Goldberger 1977, 1988). However, sensory feedback can still substantially affect the walking pattern (Rossignol, Dubuc, and Gossard 2006). Sensory inputs can change the phase of gait and aid phase transitions (Duysens and Pearson 1980, Conway, Hultborn, and Kiehn 1987). Sensory feedback can even change the phase in highly reduced preparations (paralyzed/fictive locomotion) where no other phasic sensory input enters the spinal cord (Andersson et al. 1978, Conway, Hultborn, and Kiehn 1987, LaBella, Niechaj, and

Rossignol 1992). In one particularly impressive study, the CPG was substantially impaired through hemisection of the spinal cord, and supraspinal input was also removed, but sensory feedback was still capable of driving locomotion (Kato 1989).

Sensory input is important to the normal activation of muscles in walking, beyond just correcting for environmental perturbations. (Yang, Stein, and James 1991). In complicated walking tasks (e.g. ladder-walking, uphill/downhill walking, etc), performance is impaired and sometimes eliminated without sensory feedback (Bouyer and Rossignol 1998, Rossignol et al. 2002, Bouyer and Rossignol 2003b, Panek et al. 2014).

The following sections discuss the role of specific sensory modalities in walking. I summarize the results regarding load receptors (e.g. Golgi tendon organs), positional afferents (e.g. muscle spindle fibers), and cutaneous afferents. These are followed by a discussion of cross-afferent synergies and contributions from other limbs.

2.1.3.1. Load Receptors in Walking

Animals, including humans, have load receptors throughout their body that signal how much weight or force is being carried by muscles and joints. Load receptors are very important in walking (Dietz and Duysens 2000).

In cats, ankle loading inhibits bursts in the flexor muscles (flexor bursts) and can extend the stance phase longer (Duysens and Pearson 1980), modulating the CPG. Other studies have found that ankle load receptors can reset the CPG (Conway, Hultborn, and Kiehn 1987). It has even been proposed that ankle load receptors are the most important sensory input for walking in decerebrate cats (Pearson 2008).

In humans with SCI, loading is important to generate normal gait. In one experiment, loading was required to generate correct EMG in one population (Dietz, Müller, and Colombo 2002). Another human SCI study found that if one leg is loaded and held stationary, rhythmic movements in the contralateral leg can induce rhythmic muscle activations in the stationary leg (Ferris et al. 2004).

Other studies have found that some loading (body weight support, BWS) improves many aspects of walking in humans who have suffered neural injury. In spastic paretic human subjects, BWS improved EMG, straightened the trunk, and produced better knee alignment (Visintin and Barbeau 1989). Another study in spastic paretic humans found that loading combined with training can help improve subjects even after their performance has plateaued with other methods (Barbeau, Danakas, and Arsenault 1993). In a study on human stroke subjects, load-bearing training produced improvements in functional balance, motor recovery, overground walking speed, and overground walking endurance (Visintin et al. 1998). The improvements in performance from load-bearing training (called body-weight-supported treadmill training, BWSTT) continue beyond the end of training in the SCI population (Wernig, Nanassy, and Müller 1998). This result led to body-weight supported walking becoming a major tool for physical therapy in the SCI population. Lastly, in human infants (with their immature supraspinal centers) load-bearing locomotion improves the walking pattern (Pang and Yang 2000).

2.1.3.2. Positional Afferents in Walking

Humans and other mammals, have several groups of proprioceptors that report the leg position and velocity to the nervous system. The most well-studied subclass of positional afferents are the large, myelinated afferents from the muscle, known as muscle spindle

fibers. There are also other types of positional afferents from the muscles, joints, and skin. Together these afferents signal the locations and movements of the body to the nervous system (Rossignol, Dubuc, and Gossard 2006).

Positional afferents, in particular hip afferents, are very important for walking. In spinalized cats, hip afferents were found to be the single most important afferent for walking (Pearson 2008). It is impossible to walk without moving one's hips. In human infants, with their immature supraspinal centers, hip movements are very important to help generate the walking pattern (Pang and Yang 2000). Human infant studies have found that hips triggered gait phase changes in both forward and backward walking (Pang and Yang 2002).

After hip afferents, ankle afferents are the most important. While only allowing ankle movement will not generate anything like normal EMG (Dietz, Müller, and Colombo 2002), ankle afferents have been found to affect gait transitions in the cat (Conway, Hultborn, and Kiehn 1987).

2.1.3.3. Cutaneous Foot in Walking

The cutaneous afferents of the foot have many roles in walking. Their roles were recently reviewed in Zehr and Duysens (2004), Rossignol, Dubuc, and Gossard (2006), and Panek et al. (2014). Many of the experiments performed to understand their role have been denervation studies (elimination of peripheral nerves) in animal models. Generally, researchers have found that animals, including humans, can walk without cutaneous afferents, but that their role becomes more important in complicated walking tasks, such as ladder walking, obstacle avoidance, or precise foot placement (Bouyer and Rossignol 1998, Panek et al. 2014). In one study, cats could not perform ladder walking for the first 1-3 days after denervation, but by weeks 3-7, ladder walking had

recovered, likely through substitution from supraspinal centers (Bouyer and Rossignol 2003b). However, even in the retraining studies, cats never walked in the same way as uninjured cats. For example, they no longer walked on the plantar surface of their foot, instead, adopting a claw-like foot position (Rossignol et al. 2002). In rats, acute hypothermic anesthesia resulted in altered walking kinematics (Varejao and Filipe 2007). In humans, hypothermic anesthesia also produces altered kinematics (Lin and Yang 2011).

In early studies, obstacle avoidance was identified as the primary role for cutaneous afferents during walking. In cats, obstacle avoidance is impaired if a local anesthetic is applied to eliminate cutaneous afferent signals (Prochazka, Sontag, and Wand 1978). Another study found kinematic adjustments made when cutaneous afferents were excited using a mechanical stimulus (Forssberg, Grillner, and Rossignol 1977).

Later, the role of cutaneous afferents has become more nuanced. For example, it was identified that the role may be gait phase dependent, such that cutaneous afferents assisted with obstacle avoidance during swing and stabilization during stance (Zehr, Komiyama, and Stein 1997). Cutaneous afferents have been found to contribute substantially to the reflex regulation of balance and movement in mammals (Rossignol, Dubuc, and Gossard 2006).

Although the role of cutaneous afferents may not be as strong as those of the hip afferents or load receptors in affecting gait changes, cutaneous afferents have been found to affect gait changes. Cutaneous foot afferents can affect the rhythm of walking in cats (Duysens 1977a) in different ways depending on whether the stimulus is delivered during stance or swing (Duysens and Pearson 1976). Lastly, cutaneous afferents have been shown to control the rhythmic movement in a human SCI subject (Bussel et al. 1988, Bussel et al. 1996)

Cutaneous afferents affects other parts of the nervous system involved in walking.

Cutaneous afferents can even modulate the reticulospinal system (Drew, Cabana, and Rossignol 1996), which contains the centers for the control of walking (Jordan, Pratt, and Menzies 1979, Steeves and Jordan 1980). Cutaneous afferents also appear to have direct access to the CPG (Lundberg 1979, Duysens and Van de Crommert 1998). The relationship of the CPG to cutaneous reflexes is discussed more thoroughly in section 2.3.5.10.

2.1.3.4. Cross-Afferent Synergies

Much of the research on sensory afferents during walking has focused on reductionist approaches, attempting to identify the specific role of a specific type of afferent during a well-defined, but limited task. However, sensory afferents are built to work together. Multiple sensory signals tend to reinforce one another. In studies on human infants, loading and hip afferents appear to work together. The more load there was, the more hip flexion extended stance and delayed the following swing (Pang and Yang 2000). Human infants are very similar to spinalized cats in this regard.

Cutaneous afferents have also been shown to scale corrective reflexes from other modalities (Bolton and Misiaszek 2009), and loading can change cutaneous reflexes (Bastiaanse, Duysens, and Dietz 2000). Cutaneous afferent signals from the bottom of the foot activate interneurons in concert with Group I afferents from the plantaris muscle (Hultborn 2001).

2.1.3.5. Contributions from Other Limbs

Sensory afferents affect both the ipsilateral limb and the contralateral limb, in some cases even affecting the other set of limbs (i.e. arms can affect legs and vice versa).

Stimulation of cutaneous afferents can produce an ipsilateral flexion reflex and a crossed (contralateral) extension reflex. (Sherrington 1910, Duysens, Loeb, and Weston 1980). There are many other forms of crossed coordination, including crossed flexor reflexes (Duysens and Loeb 1980).

In humans with SCI, loading of a stationary limb, combined with rhythmic motions on the contralateral limb, will produce rhythmic muscle activations in the stationary limb (Ferris et al. 2004). In human infants, changes in swing on one side affected changes in stance on the other (Pang and Yang 2001).

Coordination between arms and legs in quadrupedal animals shows task-dependence, and there is some evidence that humans do too, as precise hand movements can override the coupling (Dietz 2002).

2.2. Spinal Cord Injury

Spinal cord injury (SCI) is one of the most well-known examples of damage to the central nervous system. This section will briefly review the spinal cord, what happens when the spinal cord is injured, and how modern medicine approaches SCI. In all cases, there will be a particular focus on how walking is affected, and what happens with respect to sensory input.

2.2.1. Spinal Cord

To understand spinal cord injury, it is necessary to (briefly) review the spinal cord. From Principles of Neural Science (Kandel et al. 2014):

The spinal cord is the most caudal part of the central nervous system and in many respects the simplest part. It extends from the base of the skull to the first lumbar vertebra. The spinal cord receives sensory

information from the skin, joints, and muscles of the trunk and limbs, and contains the motor neurons responsible for both voluntary and reflex movements. Along its length the spinal cord varies in size and shape, depending on whether the emerging motor nerves innervate the limbs or trunk; it is thicker at levels that innervate the arms and legs.

The spinal cord is divided into a core of central gray matter and surrounding white matter. The gray matter, which contains nerve cell bodies, is typically divided into dorsal and ventral horns (so-called because the gray matter appears H-shaped in transverse sections). The dorsal horn contains an orderly arrangement of sensory relay neurons that receive input from the periphery, whereas the ventral horn contains groups of motor neurons and interneurons that regulate motor neuronal firing patterns. The axons of motor neurons innervate specific muscles. The white matter is made up in part of rostral-caudal (longitudinal) ascending and descending tracts of myelinated axons. The ascending pathways carry sensory information to the brain, while the descending pathways carry motor commands and modulatory signals from the brain to the muscles.

The spinal cord is where many of the reflexes related to walking are processed. It is where the locomotor CPG(s) resides. Most signals to/from all parts of the body below the neck travel via the spinal cord, and those signals are modulated, changed, combined, or filtered by the neurons there.

2.2.2. Pathophysiology

Injuries to the spinal cord come from many sources. They can be caused by an object cutting the spinal cord (e.g. bullet or knife), bruising (e.g. from a fall), ischemia (e.g. if a clot forms preventing blood to the spinal cord), or inflammation (e.g. from an adverse reaction to a medication).

When a spinal cord is injured, there is the so-called primary injury, which includes the damage directly caused by the trauma, for example: sharp bone fragments, contusion, compression and concussion (Tator 1990). Following the primary injury, there are signs of inflammation within 15 minutes (Tator 1990), necrotic zones within a few days (Tator 1990), and eventual degeneration of white matter tracts weeks after injury (Park, Velumian, and Fehlings 2004). Previously, this secondary injury was thought to be primarily caused by disruption of the vasculature (Tator 1990), but more recent research has focused on excitotoxicity (an excessive concentration of excitatory neurotransmitters) as the main cause (Park, Velumian, and Fehlings 2004).

Although the damage from secondary injury may spread much farther than the initial insult, ultimately, the lesion will stabilize. Most often, the injury is incomplete, and some pathways to/from the brain remain undamaged. Sometimes, the injury is complete, and there are no functional connection remaining to the brain (there may still be nonfunctional connections remaining).

For the purposes of this dissertation, the most important change after spinal cord injury is that the oscillatory circuitry associated with generating a walking pattern in the spinal cord is disconnected (partially or completely) from the brain, impairing or eliminating the ability to walk. Although clinical interventions have improved substantially in the last 20 years, many do not achieve the level of ambulation they would like (Anderson 2004),

and there is little we can offer to individuals with complete or the most severe incomplete injuries to help them recover the ability to walk.

2.2.3. Clinical Assessment of Spinal Cord Injury

Spinal cord injury is split into several categories. The most common method of division is the International Standards for Neurological Classification of Spinal Cord Injury ASIA Impairment Scale (ISNCSCI AIS), which divides SCI into AIS A, B, C, D, or E (Maynard et al. 1997):

- AIS A: No sensory or motor function is preserved
- AIS B: Some sensory, but no voluntary motor control
- AIS C: Some motor control is preserved
- AIS D: More motor control is preserved
- AIS E: Someone who is neurologically normal after an SCI.

In addition, the test also delineates where the injury occurred (i.e. which spinal segment). This dissertation focused on those with some preserved motor function (AIS C), and injuries above the lumbar enlargement (the part of the spinal cord where neural circuitry related to walking resides).

2.2.4. Epidemiology

There are an estimated 270,000 people with SCI in the United States, with 10,000-12,000 new cases of SCI each year (Center 2012). There is no reliable census on the size of the SCI population in the entire world, but reports have shown a worldwide prevalence (people alive with the condition) of ~250-750/1,000,000 people and an incidence (new cases/year) of ~20-60/1,000,000 (Wyndaele and Wyndaele 2006, Chiu et al. 2010). There is no reliable census on the size of the SCI population in the entire

world, but if the documented prevalence is extrapolated to the world population, there are ~1.8-5.4 million people with SCI worldwide.

In the US, SCI often affects people in the prime of their life, and it disproportionately affects males. Approximately 43% of the SCI population is paraplegic (functional loss does not include the arms), and among those 21.6% are complete paraplegics, while 21.4% are incomplete paraplegics (Center 2012).

2.2.5. Measuring the Quality of Walking after SCI

For both basic and clinical research, assessing human or animal gait quantitatively after an intervention can give a great deal of insight into whether the intervention helped or not. In humans, most clinical researchers rely on tests that measure how far a subject can walk in a specific amount of time (e.g. 6-minute walking test), or how fast they can walk a specific distance (e.g. 10-meter walking test) (van Hedel, Wirz, and Dietz 2005).

In animals, there are many approaches, primarily based on the walking pattern (i.e. kinematics). The Basso-Beattie-Bresnahan (BBB) scale for rats (Basso, Beattie, and Bresnahan 1995), and the Basso Mouse Scale (BMS) (Basso et al. 2006) are commonly used scales to categorically identify locomotion quality, involving the identification of specific kinematic features of rat and mouse gait. Other researchers have used even simpler approaches, such as the height of the foot in rats (Courtine et al. 2009) and cats (De Leon et al. 1999). The pattern of the hip and knee angles are also used by researchers (Wernig and Müller 1992, Field-Fote and Tepavac 2002, Courtine et al. 2009).

Electromyography (EMG) has been explored in many studies. In general, multiple traces are shown without quantification in a single metric (e.g. (Minassian et al. 2004, Harkema et al. 2011)), but EMG-envelope-based metrics have also been developed (Fung and

Barbeau 1989). One of the major difficulties with EMG-based metrics is that there is a large variability in EMG patterns, even among the uninjured population (Arsenault, Winter, and Marteniuk 1986, Winter and Yack 1987).

2.2.6. Existing Approaches for Recovery

The recovery from spinal cord injury includes a wide range of processes. The nervous system itself will mitigate some of the damage, and some recovery is possible without intervention. However, outcomes are better if physical therapy is incorporated. Other approaches have been pursued, including functional electrical stimulation and stem cells, with differing degrees of success. These processes and approaches are discussed in detail below.

Immediately after spinal cord injury, there is a period of spinal shock, with flaccid muscle tone, loss of reflexes, and a general muscle paralysis below the injury (Holaday and Faden 1982, Hiersemenzel, Curt, and Dietz 2000). Spinal shock generally progresses to spasticity over a period of months, where the muscles are tonically active (Hiersemenzel, Curt, and Dietz 2000). There is also some spontaneous recovery of function over the months or years after injury (Schwab and Bartholdi 1996). In humans, most of the recovery is in the first 3 months, but some spontaneous recovery will continue up to 1-1.5 years after the injury (Fawcett et al. 2007). Other mammals have a similar course, generally with animals demonstrating more functional improvement than humans post-injury (Rossignol et al. 1996, Belanger et al. 1996, Rossignol et al. 1999, Côté and Gossard 2004).

This functional improvement is the result of two processes: remyelination and sprouting. In SCI, much of the myelin (a form of insulation in the nervous system) is lost, even on connections that remain intact. As part of the recovery, the nervous system remyelinate the remaining connections (Gensert and Goldman 1997). Axonal sprouting also occurs

after SCI, and is a major contributor to the recovery of function in the remaining tracts (reviews: (Schwab and Bartholdi 1996, Raineteau and Schwab 2001)). Sprouting has been found to parallel functional recovery in multiple rat injury models (Weidner et al. 2001, Fouad et al. 2001), and the remaining corticospinal tracts will even sprout onto cross-lesional propriospinal tracts to route around the injury (Bareyre et al. 2004).

Training has been found to assist with the functional recovery. In cats, it was discovered that locomotor training can even help a cat perform load-bearing locomotion after a complete SCI (Barbeau and Rossignol 1987). A similar approach (body-weight-supported training) was tried shortly after in human SCI, with positive results (Wernig and Müller 1992). Studies found that body-weight-supported training would help once the spinal shock phase was complete (Dietz et al. 1998).

This training proved to be strenuous for physical therapists to perform for extended sessions, and several robotic solutions have appeared (Colombo, Wirz, and Dietz 2001, Veneman et al. 2007, Banala et al. 2009). The Hocoma Lokomat was an early commercial device to assist with training (Colombo et al. 2000, Colombo, Wirz, and Dietz 2001). It provides body-weight support, and robotically moves the legs to match the walking pattern of an uninjured human.

Another approach to address the reduced supraspinal control after an SCI is to provide that control via external devices. The approach to control muscles directly using external electrical stimulation is called functional electrical stimulation (FES). The study of FES began in the early 60s (Liberson et al. 1961), and has had some success with assisting SCI subjects to walk (Ragnarsson 2008), with additional health benefits beyond walking (Nash et al. 1997, Fornusek and Davis 2008), but issues with rapid fatigue are an ongoing problem (Thrasher, Graham, and Popovic 2005, Ragnarsson 2008).

There has been a great deal of research on pharmaceuticals, stem-cells, and biomaterials to explicitly cure spinal cord injury through regrowth of the lost connections. Substantial progress has been made in understanding the problem (Bradbury and Carter 2011, McCall, Weidner, and Blesch 2012, Gage and Temple 2013, Kabu et al. 2015), but no cure exists yet.

2.2.7. Sensory Input after Spinal Cord Injury

As discussed previously, walking relies on three systems: supraspinal connections, central pattern generator(s), and sensory feedback. After SCI, the role of sensory feedback in walking is even more important (Rossignol and Frigon 2011). Many reflexes change in amplitude or even reverse directions in some cases, and the injury can cause the sub-lesional spinal cord to alter the functional connections in highly individual ways.

2.2.7.1. Sensory Input is More Important after SCI

Some of the earliest studies on SCI came from soldiers suffering bullet wounds that hit their spinal cords. An early finding from this population was that cutaneous reflexes tend to be larger after SCI (Lhermitte 1919). Later, researchers attributed the amplitude increase in cutaneous reflexes to a loss of supraspinal suppression (Zehr and Duysens 2004). In the cat, these reflexes can also be normalized with training (Côté and Gossard 2004).

Cats can recover locomotion after SCI, even complete SCI, but sensory input is required for functional walking (Barbeau and Rossignol 1987). Several studies have attempted to answer what role sensory signals have in retraining the spinal cord after SCI. In general, these studies have found that loss of sensory input after SCI had a weaker effect than a sensory loss before the SCI, supporting the role of sensory feedback in recovery

(Carrier, Brustein, and Rossignol 1997, Bouyer and Rossignol 2003a). Partial denervation studies found that locomotion could be recovered after SCI without sensory feedback, but it was much slower than cats with normal sensory systems (Bouyer and Rossignol 1998). It also depended on the degree of partial denervation. If all cutaneous nerves were removed, that was sufficient to completely prevent recovery of locomotion (Bouyer and Rossignol 2003a, Bouyer and Rossignol 1998). Although it is not entirely clear why the recovery of walking is so closely tied to sensory input, studies have found that sensory afferents, after SCI, sprout new connections in the central nervous system (Helgren and Goldberger 1993).

In other situations with reduced supraspinal inputs, similar results have been found. In human infants, sensory afferents have a larger role (Pang and Yang 2000), and stimulating large proprioceptors after decerebration in cats produces larger responses than in non-decerebrated cats (Whelan and Pearson 1997). In subjects with SCI, augmenting sensory feedback (stimulating the peroneal nerve) during training improves multiple measures of walking in humans with SCI (Ladouceur and Barbeau 2000b, Ladouceur and Barbeau 2000a).

Sensory signals can even help recover locomotion if the CPG is impaired. In a powerful example of the sensory input, one study isolated a hemispinal cord (substantially impairing the CPG and eliminating supraspinal input), and they found that peripheral sensory signals were sufficient to recover locomotion (Kato 1989).

2.2.7.2. Changes in Reflexes after SCI

There are many changes in reflexes after SCI, including amplitude-changes, sign-reversals, and appearance of new phases in the reflexes.

In animal studies, researchers have documented many of these changes. In complete injuries (spinalization) of the cat, the overall qualitative change is a reduction in inhibitory responses and an increase in excitatory responses (Frigon and Rossignol 2008). In partial SCI (of the cat), the responses were more complicated, but tended toward more excitation (Frigon et al. 2009). The dependency of the reflex responses to phase of gait becomes stronger, such that fully reversed reflexes can be evoked (Nakajima, Kamibayashi, and Nakazawa 2012). Interestingly, these changes in the reflexes partially normalize after training (Côté and Gossard 2004).

After human SCI, similar changes have been found. In general, reflexes become larger in amplitude (cutaneous reflexes:(Lhermitte 1919), H-reflexes (Little and Halar 1985)). Long-latency reflex phases appear (Roby-Brami and Bussel 1987, Dietz et al. 2009). As in the animal models, the phase-dependency of reflexes increases (Yang and Stein 1990, Zehr, Komiyama, and Stein 1997). Lastly, there is an increase in habituation (a reduction in the response to an identical stimulus after repeated application) after SCI (Schindler-Ivens and Shields 2000).

Specific changes have been identified in the neural processing of the spinal cord post-injury. Interneuronal connections in the spinal cord change after SCI (Gazula et al. 2004, Bareyre et al. 2004, Courtine et al. 2008). In some cases, inhibitory pathways can simply become excitatory through a reversal of ion gradients in some neurons (Boulenguez et al. 2010).

2.2.7.3. Individualized Sensorimotor Processing after SCI

Much of the research presented here discusses the most common responses, but there is substantial variability in uninjured animal and human sensorimotor responses, and

that variability only increases after SCI. Even within a given individual, the response to stimuli can change over time.

Before injury, there is substantial variability in sensorimotor processing. During locomotion and reflex responses in the cat, there is significant inter-animal variability in terms of which muscles are active and how active they are (Loeb 1993). In human cutaneous nerve stimulation studies, the basic response varies substantially between studies (Brooke et al. 1997). Studies have shown that the physical activities (e.g. sports) subjects have participated in can affect reflex response (Nielsen, Crone, and Hultborn 1993, Maffiuletti et al. 2001, Wolpaw and Tennissen 2001, Gruber et al. 2007).

Sensorimotor processing is highly individualistic. Even in neurologically normal individuals, there is substantial variability in reflex responses based on age (Evans, Harrison, and Stephens 1990, Gibbs et al. 1999), level of activity (Loeb 1993), athletic history (Gruber et al. 2007, Wolpaw and Tennissen 2001, Maffiuletti et al. 2001), and other factors (Brooke et al. 1997). A number of studies have found further divergence after SCI in sensorimotor responses, including SCI level (Poirrier et al. 2004, Dietz et al. 1999), severity (Rossignol et al. 1996), and time since injury (Hiersemenzel, Curt, and Dietz 2000). In one study, all animals underwent identical preparations, but half had increases in extensor EMG, while the other half had reductions (Belanger et al. 1996). In a human SCI study, the relationship between the early part of a reflex and a late part of that reflex varied across the subject population (Roby-Brami and Bussel 1987). In studies of human rehabilitation after SCI, one study observed that no two individuals showed similar patterns of coordination after an intervention (Field-Fote and Tepavac 2002), and another study showed that individuals with SCI did not use consistent strategies to engage in an incline walking task (Leroux, Fung, and Barbeau 1999).

There have been observations of individual SCI subjects in whom flexor stretches led to

excitation of extensors instead of flexors, reversing the normal relationship (VanHiel 2012).

Inappropriate motor recruitment can appear after SCI, often classified clinically under the umbrella term “spasticity”, including hyperreflexia (Schurch et al. 2000), hypertonia (Woolacott and Burne 2006), dyssynergias (Dykstra et al. 1988), and spasms (Young 1994, Noreau et al. 2000).

This variability has lead many researchers to treat human SCI as individually as feasible, such as the common approach of showing tables of individuals’ clinical details in research studies (for example in: Behrman and Harkema (2000) and Knikou (2005)). An approach that individually optimized treatment could be very effective in the SCI population.

2.3. Sensory Stimulation and Walking

Sensory stimulation can initiate, control, or modulate locomotion in nearly all species examined (Kiehn et al. 1998). Many studies have found sensory stimulation to be sufficient by itself to excite the CPG. For example, in phylogenetically lower mammals, sensory stimulus has been shown to produce long-lasting bouts of locomotion without any additional pharmacological agent (Sillar and Roberts 1992, Roberts et al. 1998), and in the isolated neonatal rat spinal cord, a walking-like alternation of activity can also be produced with nothing more than sensory stimulation (Marchetti, Beato, and Nistri 2001). Sensory stimulation can modulate or reset ongoing locomotor activity, demonstrating that sensory afferents have direct access to the CPG (Schomburg et al. 1998, McCrea 1998, 2001, Hultborn 2001). Hip (Andersson and Grillner 1983, Kriellaars et al. 1994) and ankle afferents (Conway, Hultborn, and Kiehn 1987) can entrain locomotion frequency, affect gait state transitions (Grillner and Rossignol 1978b, Pang and Yang

2000), or change the phase of gait (Hultborn et al. 1998). Even reflexes that do not explicitly change the walking cycle show substantial differences depending on where in the cycle they are induced (Capaday and Stein 1986, Capaday and Stein 1987). Beyond these results, sensory stimulation can drive locomotion even with only an impaired CPG and without any supraspinal input (Kato 1989).

When combined with pharmacology, load-bearing stepping can be evoked in many animals. In the acute spinalized cat, a combination approach involving L-DOPA (a precursor of dopamine) and sensory stimulation produced stepping (Jankowska et al. 1967).

Sensory stimulation is commonly implemented by electrically stimulating nerves with a series of identical pulses (same amplitude, same pulse width) at a specific frequency, starting at a specific time, and lasting for a specific duration (Granat et al. 1992, Collins, Burke, and Gandevia 2001, Field-Fote 2001, Field-Fote and Tepavac 2002, Wu et al. 2011). This type of stimulation is often called “pulse trains.” Pulse trains are simple to understand conceptually, and they are easy for most labs to implement.

It has long been established that longer pulse widths are slightly more selective for sensory axons across animal and human models (Erlanger and Blair 1938, Paillard 1955, Veale, Mark, and Rees 1973, Panizza, Nilsson, and Hallett 1989, Panizza et al. 1992). Therefore, many reflexes have been studied using a pulse width of 1 ms to selectively stimulate sensory fibers, including H-reflexes (Panizza, Nilsson, and Hallett 1989, Lagerquist and Collins 2008), flexor reflexes (Toorring, Pedersen, and Klemar 1981), transcutaneous spinal cord stimulation studies (Kitano and Koceja 2009), and muscle afferent stimulation (Wu et al. 2011). Not only are longer pulse widths more selective for sensory fibers, but the reverse is also true, as shorter pulse widths (0.05 to 0.4 ms) preferentially activate motor axons (Grill and Mortimer 1996).

2.3.1. Potential for Fine-tuned Control and Assistance

Beyond global changes in locomotion, sensory stimulation has been shown to have a nuanced role in locomotion, with the potential for fine-tuned control. Sensory stimulation can reset or modulate ongoing locomotor activity, demonstrating that these afferents have direct access to the CPG (Schomburg et al. 1998, McCrea 1998, Hultborn 2001, McCrea 2001). Hip (Andersson and Grillner 1983, Kriellaars et al. 1994) and ankle afferents (Conway, Hultborn, and Kiehn 1987) can entrain locomotion frequency, affect gait state transitions (Grillner and Rossignol 1978b, Pang and Yang 2000), or change the phase of gait (Hultborn et al. 1998). The stimulation of the largest myelinated afferents (Group I) inhibits flexor burst and causes gait phase changes during locomotion-like activity without additional sensory input (fictive locomotion) (Conway, Hultborn, and Kiehn 1987).

Even reflexes that do not explicitly change the walking cycle show substantial differences depending on where in the cycle they are induced (Capaday and Stein 1986, Capaday and Stein 1987).

2.3.2. Sensory Stimulation with Functional Electrical Stimulation

Functional Electrical Stimulation (FES) is the use of electrical stimulation to directly modify walking patterns. It is most often associated with muscle stimulation through nerve stimulation that targets motor axons, not sensory axons (Peckham and Knutson 2005, Lynch and Popovic 2008). However, some muscles (e.g. iliopsoas) are inaccessible to surface stimulation, leading the FES community to make use of sensory stimulation. The earliest mention of FES incorporated peroneal nerve stimulation, which innervates part of the cutaneous surface of the foot, to induce hip flexion (Liberson et al. 1961), and cutaneous nerve stimulation has been used since (Granat et al. 1992). However, the focus

has not been on recruiting the CPG, but rather just utilizing the most convenient path to initiate hip flexion.

Despite this limited use of sensory input, FES has primarily focused on direct muscle activation. This focus is understandable, as controlling motor output through muscles is much more straightforward than controlling motor output through the nervous system. However, this directly activating muscles is limited in many ways. There is a reversed-order recruitment of the motor units (Peckham and Knutson 2005), where the larger, faster-fatiguing muscle fibers are recruited first before the smaller, more fatigue resistant fibers. This reverse-order recruitment leads to more fatigue as well as poorer grading of muscle force than a natural recruitment would entail. The incorrect muscle fiber recruitment order leads to more fatigue as well as poorer grading of muscle force. A more natural recruitment order would require the spinal interneurons to activate the motoneurons, but the only neurological signals that feed these interneuronal networks are supraspinal and sensory signals. In SCI, supraspinal signals are limited, leaving only sensory signals to activate these networks.

Ragnarsson, 2008, wrote that “a major shortcoming of multichannel FES systems” is that normal motor activity depends on a “highly sophisticated sensory feedback”(Ragnarsson 2008). FES systems are generally limited to the peroneal nerve for sensory stimulation (Gallien et al. 1995), and other forms of artificial sensory feedback are minimized through the use of very short pulse widths.

2.3.3. Sensory Stimulation with Physical Therapy

Physical therapists have utilized artificial sensory stimulation, including nonelectrical varieties, such as tapping the muscles during walking for many years (Perry 1967). Other researchers have explored vibration for affecting the CPG (Calancie et al. 1994), or

attempted to normalize patients towards better walking using electrical stimulation of reflexes (Bussel et al. 1988, Fung and Barbeau 1994). Lastly, there have been some attempts to use other sensory afferents in humans, including the sartorius muscle afferents (Wu et al. 2011).

Sensory stimulation has been explored in combination with training with generally positive results. Peroneal nerve stimulation with training in humans with SCI improves various measures of walking (Ladouceur and Barbeau 2000b, Ladouceur and Barbeau 2000a, Field-Fote 2001, Field-Fote and Tepavac 2002).

2.3.4. Spinal Cord Stimulation

Spinal cord stimulation (SCS) focuses on nonfocal stimulation of the spinal cord, and both epidural and transcutaneous varieties exist. The earliest work focused on a reduction in pain from epidural stimulation (Shealy, Mortimer, and Reswick 1967), followed by epidural stimulation to reduce pain from peripheral vascular disease in a subject with multiple sclerosis (Cook and Weinstein 1973).

Transcutaneous Spinal Cord Stimulation (TSCS) was developed as a noninvasive alternative to epidural SCS. Transcutaneous SCS activates posterior roots in the same way as epidural stimulation, according to models (Minassian, Persy, Rattay, Dimitrijevic, et al. 2007, Ladenbauer et al. 2010, Danner et al. 2011). However, due to the distance from the target, and due to the paraspinal muscles located between the skin and spine, TSCS will tend to be less focal and unintentionally stimulate more musculature.

In both epidural SCS (Coburn 1985, Holsheimer 1998, Rattay, Minassian, and Dimitrijevic 2000, Holsheimer 2002) and transcutaneous SCS (Minassian, Persy, Rattay, Dimitrijevic, et al. 2007, Danner et al. 2011), the lowest threshold target of SCS are the

large myelinated sensory afferents in the posterior roots. Therefore, these large myelinated dorsal roots are the primary activation target for SCS.

2.3.4.1. Spinal Cord Stimulation may substitute for supraspinal drive

Recently, there has been a lot of interest in SCS, as it may be effective as a replacement for the supraspinal drive in the CPG (Herman et al. 2002, Gerasimenko et al. 2003b, Ichiyama et al. 2005a, Harkema et al. 2011). In animal work, stimulation of the dorsal columns was very effective in eliciting locomotion in decerebrate cats (Beresovskii and Bayev 1988). Tonic stimulation of the dorsal surface of spinal cord (in cats) can produce rhythmic motion in both fore- and hind-limbs (Avelev et al. 1997, Gerasimenko et al. 2003a), and similar results were found in rats (Ichiyama et al. 2005a). In humans, spinal cord stimulation (tonic/continuous stimulation) can produce reciprocal EMG activity (a weak form of stepping) when applied to complete SCI subjects (Rosenfeld et al. 1995, Gerasimenko and Makarovsky 1996, Gerasimenko et al. 1996, Shapkova, Shapkov, and Mushkin 1997, Dimitrijevic, Pinter, and Sherwood 1997, Dimitrijevic, Gerasimenko, and Pollo 1997, Shapkova and Schomburg 2000). It could also induce bilaterally alternating modulation of reflexes – a weak form of phase dependence similar to when a subject is walking (Shapkova and Schomburg 2000). Later work has found that it could assist with weight-bearing locomotion in specific cases (Herman et al. 2002, Harkema et al. 2011).

2.3.4.2. Most Effective Stimulus Parameters

Previous studies have found the most effective SCS frequencies for walking to be 20-50 Hz. In cats, researchers have identified 20-35 Hz (Gerasimenko et al. 2003a). In rats, 40-50 Hz appears to work best (Ichiyama et al. 2005b). Humans studies have identified 25-60 Hz (Dimitrijevic, Gerasimenko, and Pinter 1998), 20-60 Hz (Herman et al. 2002), 25-50 Hz (Minassian et al. 2004), 25-50 Hz (Minassian, Persy, Rattay, Pinter, et al.

2007), 30-40 Hz (Harkema et al. 2011), 30 Hz (Hofstoetter et al. 2013), 50 Hz (Hofstoetter et al. 2014), and 25-50 Hz (Danner et al. 2015) as the ideal stimulation frequencies. Most of these studies had few participants (often just one), and none attempted to use any sort of real-time optimization to find the best stimulation frequency.

Some of the studies identified other effects of frequency. 5-15 Hz can assist with standing (Minassian et al. 2004, Minassian, Persy, Rattay, Pinter, et al. 2007). In a similar result, tonic activation switched to other patterns more conducive to walking at 22.5 Hz (Danner et al. 2015). Stimulation frequencies above 50 Hz reduced EMG activity (Gerasimenko et al. 2003a), and can correct some motor pathologies (Hofstoetter et al. 2014). Another study found slower EMG bursts with increasing frequency in the 30-70 Hz range (Dimitrijevic, Gerasimenko, and Pinter 1998).

The best spinal segment to stimulate at was generally identified as L2 (the top of the lumbar enlargement), although some studies found effective results as low as L4-L5. In cats, stimulating the L4-L5 spinal segments was effective (Gerasimenko et al. 2003a). In rats, stimulating the L2 spinal segment was effective (Ichiyama et al. 2005b).

Most human epidural stimulation studies stipulated they targeted the “upper lumbar enlargement,” which is generally located at the spinal segments between L2 and L4 (Herman et al. 2002, Minassian, Persy, Rattay, Pinter, et al. 2007). Other targets included: L2 (Dimitrijevic, Gerasimenko, and Pinter 1998), “vertebral levels ranging from T10 to L1” - roughly T12-L5 spinal segments (Minassian et al. 2004), L2-L4 spinal cord segments (Danner et al. 2015), and L1-S1 spinal cord segments (Harkema et al. 2011).

In transcutaneous spinal cord stimulation, the targeting of spinal segments is more complicated as the spinal vertebrae levels are not the same as the spinal cord segmental levels in human adults. This discrepancy is due to the fact the spinal cord

does not grow at the same rate as the spine during development. In the human adult population, there is some variability in the exact location of the segments with respect to the spinal cord, and the end of the human spinal cord (the conus medullaris) can range from T12-L3, with a mean of the lower third of L1 (Saifuddin, Burnett, and White 1998). Therefore, the upper part of the lumbar enlargement should be under the T11-T12 vertebrae in most subjects. Most studies tended to follow that approach, specifying they were targeting the “T11-T12 interspinous space” (Minassian, Persy, Rattay, Dimitrijevic, et al. 2007), or “the T11 and T12 spinous processes” (Hofstoetter et al. 2013, Hofstoetter et al. 2014).

Spinal cord stimulation works best in combination with other sensory inputs. Epidural stimulation in the rat has greatly reduced effectiveness, for example: if one leg is deafferented, it will have a poor walking gait with epidural SCS (Lavrov et al. 2008). Even when combined with a pharmacological agent (Quipazine: a serotonergic agent), epidural-SCS-induced stepping still required a moving treadmill (Gerasimenko et al. 2007). In another study on spinalized rats, body-weight support (i.e. sensory input) was required to produce stepping (Ichiyama et al. 2005b).

2.3.4.3. Spinal cord Stimulation as a Reflex Model

SCS has been explored in reflex models as well. In the normal, awake rat, epidural stimulation produces early, middle, and late responses. Reflexes were observed in the vastus lateralis (part of the quadriceps muscle), semitendinosus (part of the hamstring muscle), tibialis anterior, and triceps surae (Gerasimenko et al. 2006). In humans, similar results have been found with TSCS (Minassian, Persy, Rattay, Dimitrijevic, et al. 2007). The reflexes from spinal cord stimulation are generally called posterior root motor

responses, and they can be used to confirm the position of the stimulating leads over the lumbar enlargement (Hofstoetter et al. 2014).

2.3.5. Cutaneous Nerve Stimulation

Cutaneous nerve stimulation is the stimulation of a mostly cutaneous nerve. Most nerves carry a mix of proprioceptive afferents, motor neurons, cutaneous afferents, and sympathetic efferents. Some nerves only carry cutaneous afferents and sympathetic efferents. In this section (and the remainder of the thesis), I will use the term cutaneous nerve to refer to those nerves that are all or mostly cutaneously focused (i.e. cutaneous afferents and sympathetic efferents).

Stimulation of cutaneous nerves of the ankle/foot can generate a variety of responses, depending on the nerve stimulated and on the timing of the stimulation in the gait cycle, including flexion in the ipsilateral leg and/or extension in the contralateral leg. These responses are modulated differently from H-reflexes (Zehr, Hesketh, and Chua 2001) and differently from background EMG (Komiya, Zehr, and Stein 2000). Most areas of the skin on the legs can produce these responses, but the largest responses tend to be from the foot and ankle (Duysens and Loeb 1980).

There is a rich literature on cutaneous nerve stimulation, and I will not be able to review all of it in the context of this dissertation. I will focus on the major findings, beginning with a brief review of the afferents stimulated by cutaneous nerve stimulation.

2.3.5.1. *Cutaneous Sensory System*

To understand cutaneous nerve stimulation, it is necessary to briefly review cutaneous afferents. For a more thorough review, please consult (McGlone and Reilly 2010, Abraira and Ginty 2013). The differentiation of cutaneous afferents began with Müller's work in

the 1800's exploring whether different feelings were carried on the same nerve (Müller 1843). Cutaneous feelings are often split into major groups: tactile, thermal, pain, itch, and pleasant touch (McGlone and Reilly 2010). Within each of those major groups, there are multiple subgroups (e.g. pain can be split into myelinated and nonmyelinated afferents, and further split by the precise response characteristics of the afferents). The cutaneous afferents can also be split based on whether the afferents innervate "hairy" skin or glabrous skin. The vast majority of the body is classified as hairy skin, even if only a very small amount of hair is there, but several important areas for the purpose of this research are glabrous, including the palms of the hand and the bottom of the foot. Entirely different sets of afferents (of each of the above types) innervate hairy and glabrous skin.

Cutaneous afferents serve different roles depending on both the existence of myelination and the size of the axon. Myelinated afferents tend to serve a tactile discriminatory role (e.g. a quarter or a penny), while unmyelinated afferents tend to serve an affective-motivational role (e.g. hand is itchy). The most well-studied afferent are the large myelinated afferents in glabrous skin. They are split based on how quickly they adapt to new stimuli into either fast-adapting (FA) and slow-adapting (SA). There are four types: (1) Pacinian corpuscles (FA), (2) Meissner's corpuscles (FA), (3) Merkel's disks (SA), and (4) Ruffini endings (SA). They are low-threshold mechanoreceptors, transducing mechanical forces on the skin into action potentials. Their axons tend to be large ($>10 \mu\text{m}$), with fast conduction speeds ($>80 \text{ m/s}$). They are called $A\alpha$ or $A\beta$ afferents, depending on their size, and have the lowest threshold to electrical stimulation – i.e. they will always be the first group stimulated. Hairy myelinated afferents are less well-studied, but also include varieties of both slow-adapting and fast-adapting.

There is also a group of lightly myelinated, smaller afferents called A δ . These include mechanoreceptors and the largest pain afferents. They are not as well-understood as the A α or A β afferents, but the pain afferents in this group carry the signals for “first-pain”, which is the earliest feeling of pain from a noxious stimulus. It is differentiated from “second-pain”, which is signaled by unmyelinated afferents.

Unmyelinated afferents are more diverse, and less well-characterized than the larger myelinated afferents (A α and A β). They include pain afferents (multiple types), temperature afferents (both cold and hot), itch afferents, and even some low-threshold mechanoreceptors (Abraira and Ginty 2013), which may subserve “pleasant” touch (McGlone and Reilly 2010).

These afferents are bundled into nerves, which can be stimulated electrically, as discussed in the following sections.

2.3.5.2. *Cutaneous Reflexes*

Substantial research has been conducted on cutaneous reflexes, their responses, and how they are organized. Cutaneous reflexes include both non-noxious varieties, often called cutaneomuscular reflexes, and noxious varieties, generally associated with withdrawal reflexes (Zehr and Duysens 2004, Rossignol, Dubuc, and Gossard 2006, Panek et al. 2014).

In most studies, cutaneous reflexes are induced by electrical stimulation of either skin or nerves, but there are a few studies that use mechanical stimuli (Forssberg, Grillner, and Rossignol 1975, Forssberg, Grillner, and Rossignol 1977). The evoked response includes both ipsilateral and contralateral effects (Tax, Van Wezel, and Dietz 1995). These responses are modulated differently than H-reflexes (Zehr, Hesketh, and Chua 2001), and

do not follow general EMG activation (Forssberg, Grillner, and Rossignol 1977, Komiyama, Zehr, and Stein 2000)

Both noxious (Schouenborg, Weng, and Holmberg 1994, Schouenborg 2002) and non-noxious (Nakajima, Sakamoto, Tazoe, et al. 2006, Nakajima et al. 2009) stimuli produce cutaneous reflexes, converge on the same spinal motor pathways (Rossi, Zalaffi, and Decchi 1996), and are both modulated by gait (Schouenborg, Weng, and Holmberg 1994, Schouenborg 2002, Nakajima, Sakamoto, Tazoe, et al. 2006). However, noxious stimulation tends to produce more ipsilateral flexion (Forssberg 1979).

The spinal organization of cutaneous reflexes is not fully understood, but there is evidence that the earliest phase of the reflex arrives at the motoneurons via di- and/or tri-synaptic connections. The longer-latency effects likely come from interneurons involved in walking and the CPG (Quevedo, Stecina, and McCrea 2005). There is also evidence that the system is organized as a set of modules, where each module performs a function related to the withdrawal efficacy during the task or context at which the stimuli is applied (Schouenborg, Weng, and Holmberg 1994, Schouenborg 2002).

There has been some debate over the role of cutaneous reflexes. The classic view is that they are stumble-corrective reactions (Zehr, Komiyama, and Stein 1997). However, research has shown that the stumble-corrective hypothesis is oversimplified (Duysens et al. 1992), and that the reflexes play a more general role in shaping the gait pattern (Nakajima, Sakamoto, Tazoe, et al. 2006, Rossignol, Dubuc, and Gossard 2006).

Supporting the view that these reflexes have more general roles is the fact that these reflexes also exist in the hands (Jenner and Stephens 1982), legs (Duysens and Loeb 1980), and torso (Kugelberg and Hagbarth 1958). Among these locations, most studies have used hand-based cutaneous reflexes as their model. Hand reflexes have also been

found to be task-specific (Evans, Harrison, and Stephens 1989, Datta, Harrison, and Stephens 1989, Nakajima, Sakamoto, Endoh, et al. 2006).

2.3.5.3. Cutaneous Reflex Phases

The evoked cutaneous reflex response consists of several excitatory and inhibitory phases. These phases can be seen in both animal models and human, but they tend to be easier to evoke in models with less supraspinal influence. Multiple areas of the body have versions of these reflexes, including the hands, legs, torso, and feet.

In the human foot, there are several identified phases in the response. I will discuss the general latencies, but it is important to note that there are significant differences in latency depending on the exact combination of nerve-stimulated and muscle-response-studied – these are only rough guidelines. First, there is a short-latency inhibitory phase (<40 ms latency: (Zehr and Stein 1999)), followed by a longer latency excitatory phase, often called P1 (~50-60 ms: (Roby-Brami and Bussel 1987), ~70 ms latency: (Zehr and Stein 1999), <65 ms after the first stimulus: (Brooke et al. 1997), ~50-60 ms: (Shahani and Young 1971), ~50-60 ms: (Pedersen 1954)). This excitatory phase is generally inconsistent in human adults (Yang and Stein 1990, Van Wezel, Ottenhoff, and Duysens 1997, Baken, Dietz, and Duysens 2005). Whether this first phase of the reflex is an onset-response or a post-stimulus response has not been well-studied, but there is evidence that this part of the reflex is connected to the beginning of the pulse train, not the end of the pulse train. For example, in longer duration pulse trains, this response may actually occur during the pulse train (Roby-Brami and Bussel 1987, Pearson and Collins 1993). The main reason for this lack of certainty is most likely that researchers use very short pulse trains (<20 ms duration), for which differentiation of the onset and offset responses is impossible.

After the conclusion of the pulse train (or after the P1 phase for very short pulse trains), there is another excitatory phase, often called P2. It can be found at a latency of ~70-80 ms (75-100 ms: (Roby-Brami and Bussel 1987), 70-120 ms: (Brooke et al. 1997), ~80 ms (Tax, Van Wezel, and Dietz 1995), 70-80 ms (Duysens et al. 1996), and 80-85 ms (Van Wezel, Ottenhoff, and Duysens 1997)). The P2 phase is the most reproducible of the phases in human adults (Yang and Stein 1990, Baken, Dietz, and Duysens 2005).

There is also an inconsistent later phase (P3) after P2 in the cutaneous reflexes (Brooke et al. 1997), following an inhibitory period (Rossignol, Dubuc, and Gossard 2006). One study explored the late phases very thoroughly (Roby-Brami and Bussel 1987). They found that they were more likely to be produced with longer-duration and higher-stimulation-strength pulse trains. Sometimes the latency of onset was as late as 450 ms, beyond the window many studies even explore.

Although this discussion has been about the human foot reflexes, nearly identical reflexes exist in the hands. The responses there are generally termed E1 or E2 instead of P1 or P2. The first phase of the cutaneous hand reflex (E1) has been found to be spinally-mediated, while the second phase (E2) is supraspinally mediated (Evans, Harrison, and Stephens 1989), although an alternate hypothesis has been proposed that E2 is also spinal, but requires cortical inputs (Jenner and Stephens 1982).

The physiological significance of these phases is not entirely clear, but a few observations have been made. P1 and P2 are modulated differently during gait (Baken, Dietz, and Duysens 2005, Rossignol, Dubuc, and Gossard 2006), and the contralateral response appears to be connected to P2 (Brooke et al. 1997, Rossignol, Dubuc, and Gossard 2006). In one study the contralateral and ipsilateral P2 response were observed to only be separated by 3 ms (Van Wezel, Ottenhoff, and Duysens 1997). Loss of supraspinal input after stroke is correlated with small or absent P2 reflexes during

walking (Zehr, Fujita, and Stein 1998). After SCI, P2 is reduced (Jones and Yang 1994), and patients with hereditary disorders of their corticospinal tract also have reduced P2 reflexes (Duysens et al. 2004). The recovery of P2 is correlated with walking recovery in people with SCI (Dietz et al. 2009). Interestingly, the analogous reflex in the hand (E2) is also associated with the development of skilled, rapid finger movements during adolescence (Evans, Harrison, and Stephens 1990).

P1 appears to be purely spinal (Baken, Dietz, and Duysens 2005). In human adults, it is difficult to elicit (Yang and Stein 1990, Van Wezel, Ottenhoff, and Duysens 1997, Baken, Dietz, and Duysens 2005). In cats it is somewhat easier to elicit the P1 phase (Rossignol, Dubuc, and Gossard 2006).

2.3.5.4. Local Sign

Cutaneous reflexes demonstrate a local sign in the motor response – i.e. the response may change direction depending on where the limb was when the stimulus was applied. For example, the motor response from a stimulus while the foot is in one position might be the opposite of what it would have been if the limb were in another location (i.e. reflex reversal). The differing response will even affect how the contralateral side responds to a stimulus.

This phenomenon of limb position affecting a stimulus response has been known since at least 1877, when Gergens noted that in dogs with partial decerebration, the scratch reflex in the hind-limb could be changed by the position of the limb (Gergens 1877).

Later, Magnus published a series of studies (Magnus 1909a, Magnus 1909b) exploring the phenomenon more thoroughly. He found that passive flexion promotes extension in the response, while passive extension promotes flexion. In the second study (on cat

tails), he found the direction of tail movement was in the direction of the muscles being stretched.

Since those early studies, this phenomenon has been observed in many species and many types of cutaneous reflexes. “Local sign” has been observed in humans (Kugelberg and Hagbarth 1958, Nakajima, Sakamoto, Endoh, et al. 2006), dogs (Magnus 1909a, Magnus 1909b), guinea pigs (Brown 1911b), rats (Schouenborg 2002), and cats (Grillner and Rossignol 1978a). It has been observed in multiple cutaneous reflexes, including human hand reflexes (Nakajima, Sakamoto, Endoh, et al. 2006), human abdominal reflexes (Kugelberg and Hagbarth 1958), and human foot reflexes (Nakajima, Sakamoto, Tazoe, et al. 2006, Nakajima et al. 2009). The local sign phenomenon is generally easier to observe in subjects or animal models with less supraspinal influence (Gergens 1877, Brown 1911b, Kugelberg and Hagbarth 1958, Grillner and Rossignol 1978a, Rossignol and Gauthier 1980).

The local sign phenomenon is not limited to just the ipsilateral side either. Cutaneous reflexes produce a contralateral response, and how the contralateral side responds is dependent on the initial position of the contralateral limb (Grillner and Rossignol 1978a, Duysens, Loeb, and Weston 1980). Similar to the ipsilateral response, researchers have found that the stretch receptors are the most important afferent serving the contralateral local sign in reflex responses (Rossignol and Gauthier 1980).

2.3.5.5. Nerve Specificity

Cutaneous nerve stimulation in the foot produce both nerve-specific and nerve-aspecific features (Zehr and Duysens 2004). These similarities and differences are often task-specific. For example, the sural, peroneal, and tibial nerve produce different facilitatory responses in the tibialis anterior at the end of the stance phase, but all three produce the

same suppressive response in the tibialis anterior at the end of the swing phase (Zehr and Duysens 2004).

Stimulation of the cutaneous nerves of the foot and leg produce very similar responses in many cases. During human standing there is a general lack of nerve specificity (Komiya, Zehr, and Stein 2000). In experiments on decerebrate and spinal cats, stimulation of different nerves produces responses with different amplitudes in the muscles, but those responses are modulated similarly (LaBella, Niechaj, and Rossignol 1992). During stance (in humans), the peroneal and tibial nerve were found to evoke similar, small responses in the biceps femoris (Van Wezel, Ottenhoff, and Duysens 1997). In human walking, stimulation of sural and tibial nerves showed similar results at both low and high stimulus intensities during the swing phase (Duysens et al. 1990).

There are also many differences in the responses evoked when different nerves are stimulated. Peroneal nerve stimulation produces more suppression of tibialis anterior than tibial nerve stimulation does (Zehr, Komiya, and Stein 1997). Peroneal nerve stimulation excited the biceps femoris, but tibial nerve stimulation suppressed it (Van Wezel, Ottenhoff, and Duysens 1997). The differences between nerves are often stronger during walking than standing (Komiya, Zehr, and Stein 2000), and within gait the differences are stronger during swing than stance (Zehr, Komiya, and Stein 1997).

With the similarities and differences between the nerves, the posterior tibial nerve has several advantages as a stimulation site. It produces responses across multiple muscle groups (Yang and Stein 1990). In cat studies, a branch of the posterior (or distal) tibial nerve has shown strong effects on the CPG even during fictive locomotion (Guertin et al. 1995). In another study, researchers found that a branch of this nerve was better than any other cutaneous nerve in the leg to reset the rhythm of fictive locomotion (Frigon, Sirois, and Gossard 2010).

2.3.5.6. Afferent-Population/Stimulus-Intensity Dependence

As the stimulation strength increases, the largest myelinated afferents ($A\alpha/A\beta$) are recruited first. These large afferents carry the signals for vibration and non-painful skin stretch. The next set of thresholds recruit the smaller myelinated afferents ($A\delta$) that include the signals associated with hair movement on the skin and “first-pain”. If stimulation strength is increased further, the stimulation will also recruit the unmyelinated afferents (C) that transduce “second-pain”, temperature, itch, and additional mechanical information from the skin. With each additional population, the cutaneous reflexes change.

Stimulation of low-threshold (Group I or $A\alpha/A\beta$) tends toward increasing extension in either the current stance phase or following stance phase, while stimulation of high-threshold myelinated afferents (Group II/III or $A\delta$) tends towards more flexion. In an intact cat model, noxious stimulation (mechanical stimulation) consistently produced increased flexion, while non-noxious stimuli produced flexion during swing and extension during stance (Forssberg 1979). In a decerebrate cat model, both types of stimulation extended stance if applied during stance, but the responses diverged during swing, with low-intensity stimulation producing earlier extension and high-intensity stimulation producing increased flexion (Duysens 1977b). In a fictive locomotion cat model, low-intensity stimulation (Group I) tends to prolong the extension phase during the late extension phase, or initiate an extensor burst during flexion. Stronger stimulation (Group II or III) produced a mix of results – sometimes resetting to flexion, other times to extension. (Schomburg et al. 1998). Results in human subjects are similar, where a mix of responses were found from low intensity stimulation, while high intensity stimulation tended toward suppression of extension activity (Duysens et al. 1992) or increased

flexion (Duysens et al. 1990), with some differences depending on precisely when during gait stimulation was applied.

The differences in responses have also been explored in terms of the phases of the cutaneous reflex. As might be expected with the additional recruitment of slower afferents (i.e. it takes longer for the signal to reach the nervous system), the latter phases of the reflex are affected more than the earlier phases. In a study on anesthetized cats, researchers found that stimulation of the higher threshold afferents caused the later parts of the response to become stronger – i.e. they made the late inhibitory phases more inhibitory, and the late excitatory phases more excitatory (LaBella, Kehler, and McCrea 1989). Similar results were found in human SCI subjects (Roby-Brami and Bussel 1987), with the additional observation that increased intensity (increased recruitment of afferents) led to an increase in the latency of the late phase of the cutaneous reflex.

The afferents that can produce flexion when stimulated are sometimes referred to as flexor reflex afferents. However, flexion can be produced by a wide range of afferents under different circumstances (including both cutaneous and proprioceptive afferents), as discussed above. Many afferents that can produce flexion are also active in normal locomotion, leading some researchers to disagree with the use of the term “flexor reflex afferents” to describe these populations (Lundberg 1979, Hultborn 2001).

2.3.5.7. Task Dependence

Cutaneous reflexes depend on the task being performed. One of the earliest reports (Lisin, Frankstein, and Rechtmann 1973) studied stroke patients, uninjured humans, and decerebrate cats. The study found that the sural nerve, which normally evoked flexion during standing, could evoke extension during walking. This result has been followed by

a substantial body of work, reviewed in Zehr and Stein (1999) and Dietz (2002). I will briefly review the relevant findings.

There are differences in reflex responses between static positions and dynamic movement (Brooke et al. 1997). During standing (in cats), low-strength stimulation often results in extension (stabilizing), while the responses during walking can change based on when the stimuli is applied (Duysens and Pearson 1976), and similar results have been found with mechanical stimulation (Forssberg, Grillner, and Rossignol 1977). In another study, standing reflex responses showed a lack of nerve specificity (Komiya, Zehr, and Stein 2000). Among the cutaneous reflex phases, the most reproducible response in humans was the middle latency response (50-90 ms), which was not reproducible in standing (Yang and Stein 1990). Standing reflexes and walking reflexes have differences, but the greatest similarity is between reflexes during standing and reflexes during the stance phase (Yang and Stein 1990). Some of these task differences are likely due to the addition of the CPG during dynamic movements (Dietz 2002).

However, the task-based differences in cutaneous reflexes are not limited to whether or not the CPG is recruited. In human hand studies, the responses change depending on what task is being performed, based on whether the fingers were used together or the fingers were isolated (Evans, Harrison, and Stephens 1989, Datta, Harrison, and Stephens 1989, Nakajima, Sakamoto, Endoh, et al. 2006). There are differences in human cutaneous responses in the feet between running and walking (Duysens et al. 1993). In humans, even tasks like standing on one foot or a tilted platform can change the reflex responses (Burke, Dickson, and Skuse 1991). Another study exploring human foot reflexes found that at least some of the changes in foot reflexes are dependent on ankle muscle activation (Nakajima et al. 2009).

2.3.5.8. *Gait Phase Dependence*

There is an extensive literature on the dependence of cutaneous reflexes on the phase of gait in both animals and humans. The changes in cutaneous reflexes during gait go beyond the basic differences based on where the foot is (local sign) or task-dependency of responses (Zehr and Stein 1999). I will briefly review the major findings in this section.

As with many subtopics of research on walking, gait phase dependence is better-studied in cats than in humans. In the late 1970s and early 1980s, the first major studies on the subject documented that reflexes changing during gait in animals. They found that stimulation during stance prolongs extension, and stimulation during swing prolonged flexion – in both cases resulting in changes to the gait cycle (Duysens and Pearson 1976, Duysens 1977b). In a particularly dramatic example (using a spinalized cat model), they were able to show completely opposite recruitment (reflex reversal) in different parts of the gait cycle depending on when a mechanical stimulus (a small straw) was applied to the paw. Agonists were recruited in one phase of the gait cycle, and the antagonists in the other phase (Forssberg, Grillner, and Rossignol 1975). When higher intensity stimulation is used (recruiting smaller cutaneous afferents), the balance shifted towards more shortening of the extensor bursts, and more prolonging of the flexor bursts (Duysens 1977b).

Gait phase dependence is at least partially caused by the CPG. Gait phase dependence remains during fictive locomotion (i.e. no leg movement, therefore no sensory inflow) (Andersson et al. 1978, Rossignol, Julien, and Gauthier 1981, Schomburg et al. 1998), even when the fictive cat is also decerebrated (limited supraspinal combined with a lack of sensory inflow) (LaBella, Niechaj, and Rossignol 1992). Later research confirmed that these results were not just limited to electromyography (EMG) through the study of gait

phase dependence in the kinematic response to cutaneous stimulation (Drew and Rossignol 1987). In addition, although there are differences between the nerves in terms of the amplitudes of responses in different muscles, the modulation of those responses tend to be similar (LaBella, Niechaj, and Rossignol 1992).

Gait phase dependence does not just affect the ipsilateral response; it also affects the contralateral response. Classically, when the ipsilateral leg is stimulated to induce flexion, the contralateral leg will extend. However, if the ipsilateral leg is stimulated during swing and the contralateral leg is therefore in stance, the contralateral leg will yield (Duysens and Stein 1978). Other researchers have noticed similar contralateral reflex reversals, calling the reversed reflex the crossed flexion reflex (Duysens, Loeb, and Weston 1980, Rossignol, Julien, and Gauthier 1981).

There has been some argument about whether the reflexes are only modulated or truly reversed (Duysens and Loeb 1980). Although this debate is partially a semantics issue, the fact remains that neurologically impaired animals show much stronger phase-dependence of their reflexes than uninjured animals (Forssberg, Grillner, and Rossignol 1975, Duysens and Loeb 1980), such that true reflex reversals can be observed in some cases.

The study of gait phase dependence in human cutaneous reflexes quickly followed the animal work, beginning with an observation that the nociceptive version of the cutaneous reflex was phase-dependent (Crenna and Frigo 1984). In the 1990s, substantially more work was completed, including observations of reflex reversal (Yang and Stein 1990, Zehr, Komiyama, and Stein 1997), differences between high and low intensity phase dependence (Duysens et al. 1990), and comparisons of stimulating different cutaneous nerves (Duysens et al. 1990, Van Wezel, Ottenhoff, and Duysens 1997, Zehr, Komiyama, and Stein 1997, Zehr and Duysens 2004).

Although cutaneous reflexes are phase dependent, they are uncoupled from background EMG. This uncoupling is different from the H-reflex (Zehr, Hesketh, and Chua 2001), which is known to be coupled to motor neuron excitability (Schieppati 1987). Similar results have been found in neurologically normal humans (Komiya, Zehr, and Stein 2000, Zehr et al. 2007).

Some researchers have attempted to answer why there is gait phase dependency in the reflexes – i.e. what purpose does it serve. During swing, the primary goal of the reflexes is most likely obstacle avoidance (Eng, Winter, and Patla 1994, Zehr, Komiya, and Stein 1997, Schillings et al. 2000). During stance, the primary role is probably gait stabilization (Zehr, Komiya, and Stein 1997). Other researchers have put forward arguments that the obstacle avoidance hypothesis is oversimplified, and that the open and closing of reflex circuits is more about CPG involvement than stumble correction (Duysens et al. 1992). It has also been shown that gait phase dependence occurs among interneuronal networks, and motoneurons are either facilitated or depressed together (De Serres, Yang, and Patrick 1995).

Although most of these studies have been performed during walking tasks, gait phase dependency affects many other rhythmic tasks that humans and animals perform. Studies have found similar (but not identical) gait phase dependence of the cutaneous reflexes between human walking and running (Duysens et al. 1992, Tax, Van Wezel, and Dietz 1995) and between overground walking and walking in a robotic orthosis (Nakajima, Kamibayashi, and Nakazawa 2012). There is phase dependence in cutaneous reflexes with human backward walking (Duysens et al. 1996), cycling (Brown and Kukulka 1993, Zehr, Hesketh, and Chua 2001), and modified forms of walking (Zehr et al. 2007).

2.3.5.9. Other Dependencies

In addition to local sign, nerve-specificity, afferent-population-dependence, task-dependence, and gait-phase-dependence, there are other factors that have been observed to affect cutaneous reflex responses.

Cutaneous reflex responses can change with age. A study of hand reflexes in humans showed age-based differences in children both with and without cystic fibrosis (Gibbs et al. 1999). In a study comparing human babies, children, and adults, they found that longer-latency phases of cutaneous finger reflexes tend to become more dominant with age, and that these longer-latency phases are also related to the ability to perform rapid finger movements (Evans, Harrison, and Stephens 1990).

Cutaneous reflexes are affected by the length of the pulse train applied, but only a few studies have explored this dependency. In one study, longer pulse trains had a similar effect as stronger pulse trains – they both delayed the late (P3) motor response to the stimulation (Roby-Brami and Bussel 1987). In fictive locomotion preparations, there have been mixed results. In chemically-induced locomotion (Conway, Hultborn, and Kiehn 1987) and “spontaneous locomotion” (Frigon, Sirois, and Gossard 2010), prolonged stimulation maintained muscle activity, impairing locomotion. However, in MLR-stimulated locomotion, prolonged stimulation did not show the same result (Gossard et al. 1994).

Frequency dependence has been noted extensively in the pain literature, but it is poorly studied in the reflex and locomotion literature. For example, wind-up is a well documented phenomenon in the literature (Mendell and Wall 1965), defined by (Herrero, Laird, and Lopez-Garcia 2000) as “a progressive, frequency-dependent facilitation of the responses of a neuron observed on the application of repetitive (usually electrical)

stimuli of constant intensity.” Wind-up is often observed at frequencies above 5-10 Hz in the flexor reflex (Arendt-Nielsen et al. 1994). Aside from noting that frequency-dependence exists, there are no studies reporting on the relative efficacy of using different high frequencies in sensory stimulation.

2.3.5.10. *Cutaneous Reflexes and the CPG*

Cutaneous reflexes are thought to be closely coupled to the CPG. Cutaneous reflexes can affect the gait pattern, even initiating it in some cases. In addition, the CPG can affect cutaneous reflexes, modulating the reflexes through the gait cycle.

Cutaneous reflexes can modulate, and even initiate, the CPG. Cutaneous afferents, among others, can initiate/prolong extension or initiate/prolong flexion, depending on when and how they are applied, resetting the CPG rhythm (Schomburg et al. 1998, McCrea 2001). Cutaneous reflexes can initiate locomotion in some cases. In early experiments (Jankowska et al. 1967, Lundberg 1979), single pulse trains on flexor reflex afferents (including cutaneous afferents) were shown to evoke brief periods of alternating activity in a cat model with dopamine. It has been proposed that flexor reflexes have direct access to the flexor-half-center part of the CPG (Duysens and Van de Crommert 1998).

The CPG modulates cutaneous reflexes. During walking tasks, longer-latency reflex responses can appear, and are likely caused by the sensory inputs on the CPG (Bussel et al. 1996, Quevedo, Stecina, and McCrea 2005). Cutaneous reflex reversal, the complete reversal of the reflex between gait phases, demonstrates that the CPG can strongly modulate the reflex pathways (Nakajima, Kamibayashi, and Nakazawa 2012). Gait phase dependence remains during fictive locomotion (i.e. no leg movement therefore no sensory inflow) (Andersson et al. 1978, Rossignol, Julien, and Gauthier

1981, Schomburg et al. 1998), even when the fictive cat is also decerebrated (limited supraspinal combined with the lack of sensory inflow) (LaBella, Niechaj, and Rossignol 1992). Later research confirmed that these results were not just limited to EMG through the study of gait phase dependence in the kinematic response to cutaneous stimulation (Drew and Rossignol 1987). In addition, although there are differences between the nerves in terms of the amplitudes of responses in different muscles, the modulation of those responses tend to be similar (LaBella, Niechaj, and Rossignol 1992).

Cutaneous reflex responses demonstrate a coupling with the CPG. The timing of the reflex phases in fictive, pharmacological, and lesion preparations suggests involvement of the CPG (Burke 1999). During walking tasks, there is a long-latency contralateral facilitation (from cutaneous stimulation) during walking, similar to the half-center oscillator model of the CPG (Bussel et al. 1989).

2.3.5.11. *Clinical Use of Cutaneous Nerve Stimulation*

Clinicians and clinical researchers have made use of cutaneous nerve stimulation in studies. It is either used as part of a functional electrical stimulation approach (Granat et al. 1993, Ragnarsson 2008), or independently of other stimulation (Ladouceur and Barbeau 2000b, Ladouceur and Barbeau 2000a). Clinical approaches have often used very short pulse trains (≤ 6 pulses) at high frequency (≥ 200 Hz) (Tax, Van Wezel, and Dietz 1995, Van Wezel, Ottenhoff, and Duysens 1997, Zehr, Komiyama, and Stein 1997, Zehr, Hesketh, and Chua 2001).

The combination of functional electrical stimulation (FES) and cutaneous afferent stimulation has a long history, going back to the first paper on FES (Liberson et al. 1961). In the context of FES, there are no non-invasive ways to stimulate the hip flexors

directly, so cutaneous nerve stimulation is a convenient way to produce hip flexion (Granat et al. 1993, Ragnarsson 2008).

Some researchers have also explored the combination of classical physical-therapy-based training with cutaneous nerve stimulation. It was found that training with sensory stimulation outperformed training without stimulation across multiple performance metrics (Ladouceur and Barbeau 2000b, Ladouceur and Barbeau 2000a).

2.3.6. Time-Varying Effects in the Nervous System from Stimulation

Most studies on sensorimotor responses assume that the subject's responses will remain basically the same for an entire experimental session. In many cases though, humans and animals undergo changes that are relevant at the timescale of experimental studies. The best studied of these phenomena are sensitization and habituation.

Habituation is the reduction in a response to a repeated, identical stimulus (Harris 1943). Habituation affects a great range of reflexes, across a great range of species, with examples documented from amoebas to insects to mammals (Harris 1943). Habituation is dependent on how fast the stimulus is repeated (Lombard 1887, Dimitrijevic and Nathan 1969, Schindler-Ivens and Shields 2000), and habituation is greater after SCI (Prosser and Hunter 1936, Schindler-Ivens and Shields 2000, Dimitrijevic and Nathan 1969).

A specific type of habituation documented in humans after severe SCI is EMG exhaustion (Dietz and Müller 2004). It consists of the decline and eventual elimination of EMG during an assisted motor task (such as walking). It has been observed in individuals with severe SCI walking in a robotic orthosis for as little as ten minutes (Dietz and Müller 2004). EMG exhaustion in specific subjects is also associated with reflex differences, such as the loss of the early phase of the flexor reflex (Dietz et al. 2009).

Reflexes can also undergo sensitization (or potentiation), an increase in the amplitude of the response to an otherwise identical stimulus. A well-known example is post-tetanic potentiation, where a reflex pathway can be amplified for a few minutes after the cessation of a tetanic (constantly on) stimulus (Lloyd 1949). Several types of potentiation occur in pain-related pathways, including wind-up (Herrero, Laird, and Lopez-Garcia 2000), hyperalgesia, and allodynia (Woolf 2011).

There are also other more complicated short-term changes in neural responses. In a human infant study, infants continued to step high even after removal of an electrical stimuli that was previously applied phasically. It took several steps before they reverted to normal (Pang, Lam, and Yang 2003).

These short-term changes occur at the timescale of seconds to minutes, which will cause subjects' responses to vary in time within experimental sessions. The responses will vary in time even more after SCI.

2.4. Optimization

Optimization is the branch of applied mathematics that attempts to solve how to best achieve some goal. The goal must be defined mathematically, then a search algorithm will move through the input space (i.e., tweak “dials”) to get the best output.

Because of the relative brevity of this section, whole classes of optimization algorithms are necessarily ignored, including evolutionary methods (with the exception of genetic algorithms), temporal difference methods, stochastic, linear, and convex programming, and many more. Despite this, I attempted to cover the relevant approaches for this dissertation.

2.4.1. Derivative-Based Optimization

In the simplest formulation, the optimization algorithm has access to an objective function as well as the gradient of that objective function with each sample. There is only one minima (or maxima). There is no noise, and the system is time-invariant.

A popular and representative algorithm for this type of problem is the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm (Broyden 1970, Fletcher 1970, Goldfarb 1970, Shanno 1970). BFGS uses estimates of both the 1st derivative (the gradient) and the 2nd derivative (the convexity) to estimate where to check the algorithm next. BFGS (and its variants) can optimize in systems with >1,000 variables, but they only solve single-minima/no-noise/time-stationary/gradient-available systems.

2.4.2. Derivative-Free Optimization

Although derivative-based algorithms (discussed above) are fast and scalable, there are many cases where the derivatives are unavailable. In some of those cases, derivatives can be calculated from finite differences, but if the function evaluations are expensive, the cost may be prohibitive. If function evaluations are noisy, derivatives based on finite differences will tend to be useless (Conn, Scheinberg, and Vicente 2009).

Derivative-free optimization (DFO) has been developed to address these types of problems, but it is very limited compared to derivative-based optimization. Derivative-based algorithms can scale to problems with 1000's of variables, or even 100,000's of variables in specific problem types. With derivative-free optimization, the scale of the problems can only go up to a few hundred variables in the easiest cases (Conn, Scheinberg, and Vicente 2009).

In addition, the solutions from derivative-free approaches tend to be of lower quality than those produced by derivative-based optimization, as the knowledge of a derivative of zero (or very close to zero) will give strong evidence of being at or near a minimum. Without that check, it is more difficult to be certain of how close a near-optimal solution is to the actual optimum.

The most well-known historical algorithm in this class is downhill simplex (Nelder and Mead 1965). Downhill simplex is easy to understand and implement, but suffers from several flaws, such as lacking a guarantee of convergence. In fact it is known to not converge in specific classes of problems (Woods 1985, McKinnon 1998).

Modern DFO algorithms are often based on trust-regions, which form an estimate of local curvature in a region (Conn, Scheinberg, and Toint 1998, Powell 2006).

2.4.3. Global Optimization

The central goal of global optimization is to find the best solution in a system that may have multiple local optima. It is generally applied in constrained or bounded systems.

There are many approaches to global optimization. Some are limited to specific problem classes, such as branch-and-bound, which requires a convenient way to calculate the upper and lower bound of a region of the space. For the purposes of this dissertation, I will limit the discussion to those that can be used without prior knowledge of the function. Some examples of popular global optimization algorithms in this class include simulated annealing (Kirkpatrick, Gelatt, and Vecchi 1983, Černý 1985), differential evolution (Storn and Price 1997), and particle swarm optimization (Eberhart and Kennedy 1995, Shi and Eberhart 1998).

Global optimization algorithms tend to require a large number of function evaluations to achieve a given level of performance. Particle Swarm optimization, for example, generally requires 10-100 function evaluations per iteration (Kennedy 2010). Good results from differential evolution (Storn and Price 1997) and simulated annealing (Dekkers and Aarts 1991) generally 1,000's of function evaluations – even for relatively low dimensional systems.

In general, global optimization algorithms have few guarantees regarding how quickly or even if they will actually find the optimum. Simulated annealing is one of the few with any guarantees at all, with only a weak guarantee that it will converge in finite time for a finite system under a very slow cooling schedule (Granville, Křivánek, and Rasson 1994).

2.4.4. Optimization with Noise

Optimization in the presence of noise has a rich literature. Because of the relative brevity of this section, whole classes of algorithms must be left out, but I hope to give a rough overview of the highlights of the field, with a particular focus on the algorithms' applicability to on-line/real-time optimization. One of the main difficulties with optimization in the presence of noise is that the error in the estimate at a specific point in space goes down by $1/\sqrt{N}$ where N is the number of evaluations near that point (Conn, Scheinberg, and Vicente 2009).

The simplest class of algorithms that are robust against noise are coordinate searches (or pattern searches). These algorithms are often not designed with noise in mind, but their simplicity causes them to be reasonably robust against noise. The most well-known example of this class is downhill simplex (Nelder and Mead 1965).

Another branch of algorithms for optimizing noisy systems is called stochastic approximation. Stochastic approximation attempts to directly address the issue of

estimating the derivatives of noisy systems, and perform a gradient descent approach (Spall 2005). Early versions used a finite difference approach, requiring many function evaluations per iteration (Kiefer and Wolfowitz 1952), but later versions attempt to estimate the gradient with only two samples per iteration, regardless of the dimensionality (Spall 1992), and achieve this goal with very little loss of performance (Spall 2005). This approach makes them primarily a local optimizer, but it is possible to extend stochastic approximation to global optimization approaches (Maryak and Chin 2001).

Genetic algorithms are a type of stochastic optimization that attempts to mimic a weak form of genetics (Golberg 1989). At each iteration there exists a group of possible solutions (called “individuals” or “chromosomes”). Depending on each solution’s fitness (performance with respect to an objective function), it increases the likelihood of producing “offspring” in the next generation through either a cross with another “fit” solution, or just mutation on the original solution. The cross and mutation processes are implemented using stochastic variables to add randomness to the process. Eventually, the process will terminate due to either running out of function evaluations or because some performance metric was achieved (Mitchell 1998). Genetic algorithms can be robust against noise, but they assume that multiple identical scenarios are available, which may not be possible in a real-life, time-varying system.

Another approach that shares many similarities with genetic algorithms is particle swarm optimization (PSO). PSO starts with a randomly placed set of particles with random velocities. In each iteration, all points are tested, and they are “accelerated” towards both that particles best and the global best achieved by any particle with a random weight (Kennedy and Eberhart 1995, Poli, Kennedy, and Blackwell 2007). The “vanilla” version of PSO is reasonably robust against noise, despite a lack of being explicitly designed for

that purpose (Parsopoulos and Vrahatis 2001). A version of PSO has been developed to explicitly handle noise, with better performance as a result (Pugh, Zhang, and Martinoli 2005). However, PSO is not frugal with function evaluations, and will often test tens of points per iteration.

2.4.5. Response Surface Based Optimization

Response surfaces are a type of model to estimate how a complex system responds to a set of input parameters. When used in optimization, “the technique often requires the fewest function evaluations of all competing methods” (Jones, Schonlau, and Welch 1998).

The idea behind response surface-based methods is to fit a stochastic model to the data (globally), then optimize within that model to identify the next point to test. During the first step, the model-fitting step, the algorithm approximates the space. Then, in the second step, the algorithm uses that model to estimate where to sample next.

The use of stochastic models for optimization started in the 1960s in several different disciplines independently: mathematical geology, global optimization, and statistics. In mathematical geology, the approach (called “kriging” in that literature) began in 1963 (Matheron 1963, Cressie 1990). In mathematical geology, these approaches were designed to find a mineral in a 2- or 3-dimensional space, generally with the assumption of noise in the measurements. In the field of global optimization, the approach also began in the early 1960s with “Bayesian global optimization” (Kushner 1964). Kushner’s approach used a stochastic process based on Brownian motion to guide sampling of an unknown function. It even included an early concept of time-variance in the response function. Kushner’s approach has been expanded on by many authors in the years following (for example: (Mockus, Tiesis, and Zilinskas 1978, Žilinskas 1992, Betr  1992,

Jones, Schonlau, and Welch 1998). Lastly, in the early 1970s a similar approach was developed in the statistics literature to model complex functions (Sacks and Ylvisaker 1970, Sacks et al. 1989).

2.4.5.1. Gaussian Process Models

A Gaussian process model is a model where all points are modeled as Gaussian random variables with a structured covariance (discussed below). These models have the advantage of easily incorporating concepts of uncertainty, standard deviation, and noise. There are two analogies that may help understand them. First, Gaussian process models are designed to interpolate or extrapolate untested points from tested points. Therefore, they therefore are a type of regression model. In fact, there is a direct relationship between the expected value at each point and a spline representation (Jones, Schonlau, and Welch 1998). Second, Gaussian process models estimate the objective function from previously-tested, nearby points. Thus, Gaussian process models can be interpreted as an extension to the simple nearest neighbors algorithm. In the nearest neighbor algorithm, the median of a set of the nearest points (often 7 or 9 points) are used to estimate the value of a point. Gaussian process models also estimate the value based on an estimate of the correlation between the test point and the nearby points.

2.4.5.2. Efficient Global Optimization

A popular optimization algorithm within the class of response surface methods is Efficient Global Optimization (EGO). It was designed to optimize computationally expensive computer models using Gaussian processes (Jones, Schonlau, and Welch 1998). EGO uses a DACE (Design-and-Analysis-of-Computer-Experiments) model (Sacks et al. 1989) at its core, and optimizes the maximum likelihood of improvement

(with respect to the objective functions) to select the next test point. EGO recently was proved to converge (Vazquez and Bect 2010), although the proof did not guarantee a specific rate of convergence.

Modifications have been made to EGO allowing it be robust against noise (Huang et al. 2006, Forrester, Keane, and Bressloff 2006) through regularization of the covariance matrix (discussed below). More recently, modifications have also been proposed to make it robust against time-variance (Desautels 2014, Morales-Enciso and Branke 2015), through the addition of time as an uncontrolled input variable.

2.4.5.3. Covariance Functions

For every pair of random variables, you can define a covariance. If the random variables are completely unrelated to one another (in statistical terms: independent), this covariance will be set to zero. Positive values mean that the two random variables tend to move in the same direction, negative values mean that the two variables tend to move in opposite directions.

To estimate covariance from data, the classic approach is to test multiple samples from each random variable, and calculate the covariance directly. In Gaussian process models, every data point comes from different, but related Gaussian probability distributions. Therefore, the majority of data point pairs will have nonzero covariances, leading to a requirement to estimate $N^2/2$ covariances, where N is the number of data points. In the most general case, it is impossible to directly estimate more parameters than data points, so a smaller set of parameters must be used to construct a function that models all potential covariances in the input space.

A covariance function can be defined that models how any two random variables relate to one another in a space, and satisfies the requirement of having relatively few

parameters as compared to the number of data points. There are many choices of what covariance function to use in the model. In EGO, the model used is derived from kriging (Sacks et al. 1989), using an approach that has relatively few parameters to estimate, but can flexibly model many different covariance structures.

There has been substantial work on alternative covariance functions. This covariance model has been extended to allow for noisy responses by incorporating a regularization term into the covariance matrix (Huang et al. 2006, Forrester, Keane, and Bressloff 2006). Some of this work has focused on functions that are nonstationary in space (Xiong et al. 2007, Toal and Keane 2012) (e.g. a chirp function has high frequency content in one part of space, and low frequency content in another).

2.4.5.4. What Metric to Use in the Selection of the Next Point

In stochastic-model-based optimization, there is a question of what quantity to maximize (or minimize) when searching for the next point. There are many choices. One obvious answer is to minimize expected value, but this approach can lead the algorithm to get stuck in local minima (Jones, Schonlau, and Welch 1998). There are many alternatives, and each one may be more or less appropriate, depending on the context in which the algorithm needs to work. One popular choice is the maximum expected improvement as compared to the current best (Mockus, Tiesis, and Zilinskas 1978, Jones, Schonlau, and Welch 1998). Maximum variance and “minimizing surprises” (minimize the maximum probability that a true value deviates from the prediction) can help the algorithm search the most unknown areas (Sasena, Papalambros, and Goovaerts 2002). The knowledge gradient approach attempts to maximize the expected improvement in the estimate of the best point (Ryzhov, Powell, and Frazier 2012). Other approaches have attempted to

minimize “regret”, which is the cumulative loss from not sampling at the optimal point (Srinivas et al. 2009).

In a modified version of EGO (Huang et al. 2006), the maximum expected improvement was augmented to encourage “in-searching” or searching in areas near points previously sampled. It is important to test multiple times in these areas, as the first point may have been corrupted by noise. This augmented expected improvement (Huang et al. 2006) has excellent properties when optimizing noisy, multi-modal systems.

2.4.6. Time-Varying Optimization

Time-varying optimization is comparatively less well-studied than other types of optimization, and the research that has been done is split across several subfields, including stochastic optimization, optimal control, machine learning, and operational research.

Some of the earliest work on time-varying algorithms was pursued by Harold Kushner in the 1960s (Kushner 1962, 1964). His approach was to increase the estimate of the standard error of previously sampled points dependent on how long it had been since they were sampled. This algorithm was limited to one variable. He extended this work to a stochastic-process-based sampling approach for one-dimensional problems.

Dynamic optimization (or optimal control) is the optimization of a system in the context of a (known) dynamic system (Chiang 2000). In general, dynamic optimization assumes the agent optimizing the system knows quite a bit about how the system functions (and how it might change with time). In cases where the dynamics are unknown, the approaches cannot be easily applied (Conn, Scheinberg, and Vicente 2009). System identification could be used, but it may take a substantial portion of the experimental budget, leaving few experiments to optimize.

Within machine learning, there is a subfield called active sensing that has explored time-varying systems. Active sensing is the study of how an artificial agent might choose to pursue sensory tests (run experiments) from a set of options in such a way to most efficiently learn about a system or context. This question of how to most effectively decide on which sensory tests to pursue is not identical to optimizing an objective function, but it is closely related. One of the classic problems in this field is the multi-armed bandit problem. In this problem, an agent must choose from a finite set of choices which may benefit the agent to different degrees, but the agent does not know in advance which choice is best. Some researchers have explored how an algorithm might approach this problem if the measure changed with time (Hazan and Kale 2010). The same research group has extended this approach to portfolio optimization (e.g. when to buy/sell stocks) as well (Hazan and Kale 2009, 2015).

There are few algorithms designed to work explicitly in general time-varying systems where the dynamics may be completely unknown (i.e. black-box time-varying systems). Some local optimization functions work reasonably well, but they often have steps in their algorithms that reduce the search space to a smaller region when the algorithm believes it is close to the objective. This type of focus may backfire in a time-varying system. Towards that end, a version of the downhill simplex algorithm was created for time-varying system, created by removing the step that shrunk the simplex (Xiong and Jutan 2003).

Recently, there have been recommendations and attempts to extend response-surface-methodology approaches to time-varying systems. Generally, the best approach has been to add time as an uncontrolled variable (Desautels 2014, Morales-Enciso and Branke 2015). However, these approaches were not designed to be robust against noise in the system as well.

Another problem with time-varying objective functions is how to define classes of problems or success of the algorithm. Previous tests have used very specific functions, such as the moving peaks function, to explore how well time-varying optimization worked (Morales-Enciso and Branke 2015), but such an approach cannot be easily generalized to larger problem classes. Time-series analysis includes tests to decide if a time series is stationary or not (Dickey and Fuller 1979, Priestley 1981) based on a time-series of a specific value, but those tests do not quantify precisely how much time-variance there is in a system. In the machine learning literature, one approach has been to put upper limits on total variance of a set of values to delineate problem classes (Hazan and Kale 2009).

2.4.7. Closed-Loop Experimentation in Neuroscience

One of the earliest, and most famous, examples of closed-loop experimentation in neuroscience was Hodgkin and Huxley's efforts to learn about membrane conductances using a voltage clamp (Hodgkin and Huxley 1952), a circuit which adjusted the current to maintain a constant voltage. There has been a great body of work since then using voltage clamps, patch clamps (Zhao et al. 2008), and more recently, dynamic clamps (Prinz, Abbott, and Marder 2004). Control loops like those above have become more common, and approaches which attempt to not only maintain a value in real-time, but optimize the system in real-time with respect to some metric have begun to be explored. In order to retain brevity and relevance for this dissertation, I will limit my discussion to approaches which used pulse-like stimulation of neural structures.

2.4.7.1. *Closed-Loop Stimulation*

In the context of neural stimulation, experimental preparations (within the same study) are rarely truly identical, and to make matters worse, their response may drift over the

course of an experiment. In other cases, more accuracy may be desired than can be achieved with an open-loop approach. Therefore, closed-loop approaches to neural stimulation have been developed, often using the formalism of control theory. I will briefly review the literature, for a more thorough review, please consult Sun and Morrell (2014) or Lynch and Popovic (2008).

Within the neuroscience community, the FES literature has explored the use of closed-loop stimulation for the longest time. FES is a natural fit, as it is similar to the types of approaches used in robotics. The earliest reports on closed-loop FES focused on electrical stimulation of agonist/antagonist pairs to control the elbow joint (Vodovnik, Crochetiere, and Reswick 1967), followed by the ankle muscles (Stanic and Trnkoczy 1974, Crago, Mortimer, and Peckham 1980). Later approaches have expanded upon what information the controller had access to, including gait phase (Pappas et al. 2001, Mansfield and Lyons 2003) and EMG (Sinkjaer et al. 2003, Dutta, Kobetic, and Triolo 2008). The earliest reports relied on simplistic bang-bang or proportional control (Vodovnik, Crochetiere, and Reswick 1967, Stanic and Trnkoczy 1974, Crago, Mortimer, and Peckham 1980). More recent approaches have begun to use more robust techniques, including sliding mode control (Jezernik, Wassink, and Keller 2004, Nekoukar and Erfanian 2011) and adaptive control (Blaya and Herr 2004)

Closed-loop stimulation has begun to be explored in spinal cord stimulation. One pair of studies used a relatively simple approach of measuring body position in order to judge when to apply spinal cord stimulation to reduce chronic pain (Schade et al. 2011, Schultz et al. 2012). Another approach for reducing chronic pain with SCS relied on measuring compound action potentials to close the loop (Parker et al. 2012). Closed-loop epidural spinal cord stimulation to locomotion has also been explored using a simple EMG-based metric (Desautels 2014, Desautels et al. 2015).

There has been a lot of interest in using closed-loop feedback to improve the termination of epileptic seizures. The first approaches relied on large, cumbersome systems external to the body (Peters et al. 2001, Kossoff et al. 2004, Osorio et al. 2005), but more recent approaches have begun to explore the closed-loop systems on implanted devices (Morrell 2011, Heck et al. 2014).

In addition to the above, other subfields of neurology, neuroengineering, and neuroscience have begun to make use of closed-loop stimulation. In Parkinson's Disease, some early results for closed-loop stimulation have shown positive results (Rosin et al. 2011, Little et al. 2013). In addition, some work has pursued simulations in preparation for future closed-loop deep-brain stimulation (Santaniello et al. 2011). In the neurophysiological literature, closed-loop approaches have been explored in the vestibular system (Micera et al. 2010), the whisker barrel cortex (Venkatraman et al. 2009), and motor axon stimulation (Bostock, Cikurel, and Burke 1998).

2.4.7.2. On-line/Real-time Optimization in the Neuroscience

One important subcategory of closed-loop approaches is on-line optimization. On-line optimization attempts to go beyond just leveraging the state of the system to use that state to find the best possible input to get the best output – where “best” is defined by an objective function. These type of approaches are sometimes called continuous optimization, real-time optimization, or dynamic optimization, adaptive sampling, active sensing, or active learning, depending on what algorithm is used and how the optimization is setup. Within these approaches, there is often little separation between understanding how the system responds to inputs and finding the optimal response. Generally, both system identification and optimization must be done at the same time. It is also important to note that it is impossible to cleanly separate studies into closed-loop

control, on-line system identification, and on-line optimization. Many studies fall into more than one category. With that caveat in mind, I will briefly discuss the history of on-line optimization in neuroscience.

On-line optimization is still relatively rare in neuroscience, but it has begun to gain traction in the last two decades, with most of the work in the last 10 years. As with closed-loop studies, the FES literature has produced some of the earliest studies (Kostov et al. 1995, Lynch and Popovic 2008). In neurophysiology research, various approaches have been pursued to find the best stimulus to get a specific response (Földiák 2001, Edin et al. 2004, O'Connor, Petkov, and Sutter 2005, Benda et al. 2007, Lewi, Butera, and Paninski 2008). In the study of brain-computer-interfaces (BCI), also called brain-machine-interfaces (BMI), several studies have used on-line optimization to improve how the system interacted with the user (Gürel and Mehring 2012, Fruitet et al. 2012, Fruitet et al. 2013). Lastly, some recent studies have explored the use of on-line optimization in epidural spinal cord stimulation (Desautels 2014, Desautels et al. 2015).

Although on-line optimization is not strictly new in neuroscience, neuro-engineering, or neurology, it is still quite rare, and the best approaches are still being explored. Beyond this dissertation, I expect there to be many studies on this subject in the near future.

CHAPTER 3: A SYSTEM TO OPTIMIZE SENSORY STIMULATION IN HUMANS AFTER SPINAL CORD INJURY

This chapter presents an approach and system I developed to optimize sensory stimulation for maximally assisting human walking after spinal cord injury (SCI). The apparatus includes a robotic gait orthosis (the Hocoma Lokomat) to record force output during standardized stepping with variable assistance. In order to validate the system, detailed results from a representative motor-incomplete SCI subject is presented. The system optimized the distal tibial nerve stimulation, a predominantly cutaneous nerve. The objective function (i.e. the variable that the algorithm was optimizing) was minimizing the assistive force required from the Lokomat – maximizing the contribution from the subject.

The chapter is organized into an introduction, system design, results, and discussion. The introduction section presents an abbreviated summary of current research and motivation for this chapter. The system design discusses the components of the system, the derivation of the objective function, and the optimization algorithm. The results present the one-subject in-depth validation, and the discussion interprets and extrapolates the findings.

3.1. Introduction

Spinal cord injury (SCI) damages the pathways connecting spinal neural networks to the brain, impairing the nervous system's ability to produce the correct motor output for many motor tasks, including walking.

Sensory connections to the spinal cord often remain intact after injury. These sensory input are very important in walking, and become even more important after an SCI.

While an SCI damages the descending pathways to spinal neural networks, the sensory connections to those networks generally remain intact below the injury site. Sensory inputs from the legs have both numerous and substantial effects on walking (Rossignol, Dubuc, and Gossard 2006). Hip afferent activity can change the stance-to-swing transition in cats (Pearson 2008), human infants (Pang and Yang 2000), and humans with SCI (Dietz, Müller, and Colombo 2002). Extensor loading affects gait transitions (Rossignol, Dubuc, and Gossard 2006), can assist in generating a walking pattern even with only partial rhythmic sensory inputs (Ferris et al. 2004), and can increase appropriate EMG during walking (Dietz, Müller, and Colombo 2002). The large role of sensory feedback in walking is even more pronounced after SCI (Rossignol and Frigon 2011).

Cutaneous reflexes can modulate, and even initiate, the CPG. Cutaneous afferents, among others, can initiate/prolong extension or initiate/prolong flexion, depending on when and how they are applied, resetting the CPG rhythm (Schomburg et al. 1998, McCrea 2001). Cutaneous reflexes can initiate locomotion in some cases. In early experiments (Jankowska et al. 1967, Lundberg 1979), single pulse trains on flexor reflex afferents (including cutaneous afferents) were shown to evoke brief periods of alternating activity in a cat model with dopamine. It has been proposed that flexor reflexes have direct access to the flexor-half-center part of the CPG (Duysens and Van de Crommert 1998). Cutaneous reflex responses also demonstrate a coupling with the CPG. The timing of the reflex phases in fictive, pharmacological, and lesion preparations suggests involvement of the CPG (Burke 1999). During walking tasks, there is a long-latency contralateral facilitation (from cutaneous stimulation) during walking, similar to the half-center oscillator model of the CPG (Bussel et al. 1989).

Clinical researchers have made use of cutaneous nerve stimulation in various studies to assist with walking. It is either used as part of a functional electrical stimulation approach (Granat et al. 1993, Ragnarsson 2008), or independently of other stimulation (Ladouceur and Barbeau 2000b, Ladouceur and Barbeau 2000a). Clinical approaches have often used very short pulse trains (≤ 6 pulses) at high frequency (≥ 200 Hz) (Tax, Van Wezel, and Dietz 1995, Van Wezel, Ottenhoff, and Duysens 1997, Zehr, Komiyama, and Stein 1997, Zehr, Hesketh, and Chua 2001). The combination of functional electrical stimulation (FES) and cutaneous afferent stimulation has a long history, including the first paper on FES (Liberson et al. 1961). In the context of FES, there are no non-invasive ways to stimulate the hip flexors directly, so cutaneous nerve stimulation is a convenient way to produce hip flexion (Granat et al. 1993, Ragnarsson 2008). Some researchers have also explored the combination of classical physical-therapy-based training with cutaneous nerve stimulation. It was found that training with sensory stimulation outperformed training without stimulation across multiple performance criteria (Ladouceur and Barbeau 2000b, Ladouceur and Barbeau 2000a).

Although sensory stimulation has proven to be useful, there is a substantial variability in motor patterns after SCI. Spinal reflexes can change in amplitude (Little and Halar 1985) or latency (Roby-Brami and Bussel 1987). Inappropriate motor activation can appear, often classified clinically under the umbrella term “spasticity”, including hyperreflexia (Schurch et al. 2000), hypertonia (Woolacott and Burne 2006), dyssynergias (Dykstra et al. 1988), and spasms (Young 1994, Noreau et al. 2000). To further complicate issues, sensorimotor responses vary significantly from individual to individual based on athletic history (Gruber et al. 2007, Wolpaw and Tennissen 2001, Maffiuletti et al. 2001), SCI level (Poirrier et al. 2004, Dietz et al. 1999), SCI severity (Rossignol et al. 1996), and time since injury (Hiersemenzel, Curt, and Dietz 2000). When studying how SCI subjects

motor patterns are reorganized after an intervention, it has been observed that no two individuals show equivalent patterns of coordination (Field-Fote and Tepavac 2002). All of these sources of variability imply that the optimal stimulus pattern for functional benefit may be highly individualized, and may even change over time within a given individual.

One way to resolve highly individualized sensorimotor-processing is to use closed-loop approaches that measure and change the input real-time based on the output. The earliest reports on closed-loop stimulation focused on functional electrical stimulation (FES) of agonist/antagonist pairs to control the elbow joint (Vodovnik, Crochetiere, and Reswick 1967), followed by the ankle muscles (Stanic and Trnkoczy 1974, Crago, Mortimer, and Peckham 1980). Later approaches have expanded upon what information the controller had access to, including gait phase (Pappas et al. 2001, Mansfield and Lyons 2003) and EMG (Sinkjaer et al. 2003, Dutta, Kobetic, and Triolo 2008). The earliest reports relied on simplistic bang-bang or proportional control (Vodovnik, Crochetiere, and Reswick 1967, Stanic and Trnkoczy 1974, Crago, Mortimer, and Peckham 1980). More recent approaches in FES have begun to use more robust techniques, including sliding mode control (Jezernik, Wassink, and Keller 2004, Nekoukar and Erfanian 2011) and adaptive control (Blaya and Herr 2004). Closed-loop stimulation has begun to be explored in spinal cord stimulation. One pair of studies used a relatively simple approach of measuring body position in order to judge when to apply spinal cord stimulation to reduce chronic pain (Schade et al. 2011, Schultz et al. 2012). Another approach for reducing chronic pain with SCS relied on measuring compound action potentials to close the loop (Parker et al. 2012). Closed-loop epidural spinal cord stimulation to locomotion has also been explored using a simple EMG-based metric (Desautels 2014, Desautels et al. 2015). In Parkinson's Disease, some early results for closed-loop stimulation have shown positive results (Rosin et al. 2011, Little et al. 2013).

In addition, some work has pursued simulations in preparation for future closed-loop deep-brain stimulation (Santaniello et al. 2011). In the neurophysiological literature, closed-loop approaches have been explored in the vestibular system (Micera et al. 2010), the whisker barrel cortex (Venkatraman et al. 2009), and motor axon stimulation (Bostock, Cikurel, and Burke 1998). Although these closed-loop approaches leverage the actual state of the systems, they do not generally improve their performance as more is learned about the system or model the system in an individualized way.

On-line optimization is a closed-loop approach that continuously improves as the algorithm's model of the system improves. These approaches are still relatively rare in neuroscience, but they have begun to gain traction in the last two decades, with most of the work in the last 10 years. As with closed-loop studies, the FES literature has produced some of the earliest studies (Kostov et al. 1995, Lynch and Popovic 2008). In neurophysiology research, various approaches have been pursued to find the best stimulus to get a specific response (Földiák 2001, Edin et al. 2004, O'Connor, Petkov, and Sutter 2005, Benda et al. 2007, Lewi, Butera, and Paninski 2008). In the study of brain-computer-interfaces (BCI), also called brain-machine-interfaces (BMI), several studies have used on-line optimization to improve how the system interacted with the user (Gürel and Mehring 2012, Fruitet et al. 2012, Fruitet et al. 2013). Lastly, some recent studies have explored the use of on-line optimization in epidural spinal cord stimulation (Desautels 2014, Desautels et al. 2015). However, as far as these authors could find, no optimization work has been completed in peripheral sensory stimulation, nor sensory stimulation as applied to walking.

To address this gap, I developed a system that can apply sensory stimulation to someone with incomplete SCI and analyze walking performance in real-time. The system addresses the variability of neural circuit function by individually optimizing the

intervention. Individual optimization of sensory stimulation allows for targeted exploration of many sensory stimulation patterns in an automated experimental paradigm. I tested this system against one subject to show that it can work, and performed confidence and sensitivity analysis on this subject's optimal stimulation parameters.

3.2. System Design

Conceptually, this system (Figure 3-1) consists of three major components: (1) *intervention* (electrically stimulating the distal tibial nerve), (2) *assessment* (measuring the reduction in assistive force), and (3) *optimization* (selecting the intervention parameters with the efficient global optimization algorithm). The intervention and assessment components interact directly with the subject, who is walking in the Lokomat (a robotic gait orthosis made by Hocoma Inc.). The optimization component closes the loop, allowing for continuous improvement.

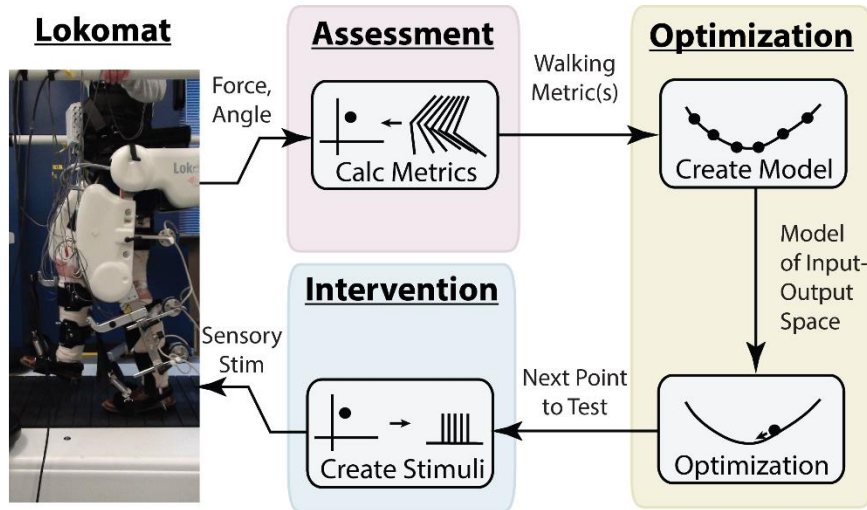


Figure 3-1: Overview of the optimization approach. Each iteration consisted of the algorithm completing one loop. In the Assessment stage, force and angle data from each step in the Lokomat was converted into a measure of the quality of stepping (the force metric). In the Optimization stage, a maximum likelihood model was fit, and the optimum was found within that model. In the Intervention stage, the identified model-optimal stimulation parameters were converted into a pulse train which was then applied to the individual in the Lokomat (at the distal tibial nerve), and measurements were made again, starting a new iteration.

3.2.1. Intervention: Sensory Electrical Stimulation

Many interventions can assist SCI subjects walk. Our approach required that the intervention be parameterized in a low dimensional space, that the parameters could be changed quickly (e.g. every four steps) during an experiment, and that it could be controlled by an automated system implementing the optimization algorithm. Sensory electrical stimulation meets all of those requirements, and it has shown promise in physical therapy and FES research. The distal tibial nerve, a mostly cutaneous nerve, was selected as a good candidate: It is easy to stimulate in the Lokomat, it can produce hip flexion, and it innervates the cutaneous surface of the bottom of the foot (an area important for walking).

The stimulation subsystem sends a pulse train (figure 3-2) to the distal tibial nerve once per complete step (steps are divided at swing-to-stance on the right foot). A real-time

embedded system (National Instruments CompactRIO) triggers a constant current stimulator (DS7A from Digitimer Ltd.) to produce the pulses. Pulse width is set to 1 ms and amplitude is set to the minimum that would produce a consistent flexion reflex. The pulse train is applied at the same time on each leg relative to the leg's phase of gait (e.g. stance or swing). For example, if the pulse train is applied on the left leg in stance, it will be applied on the right leg in stance as well, even though the right leg's stance is one second delayed from the other leg's stimulation.

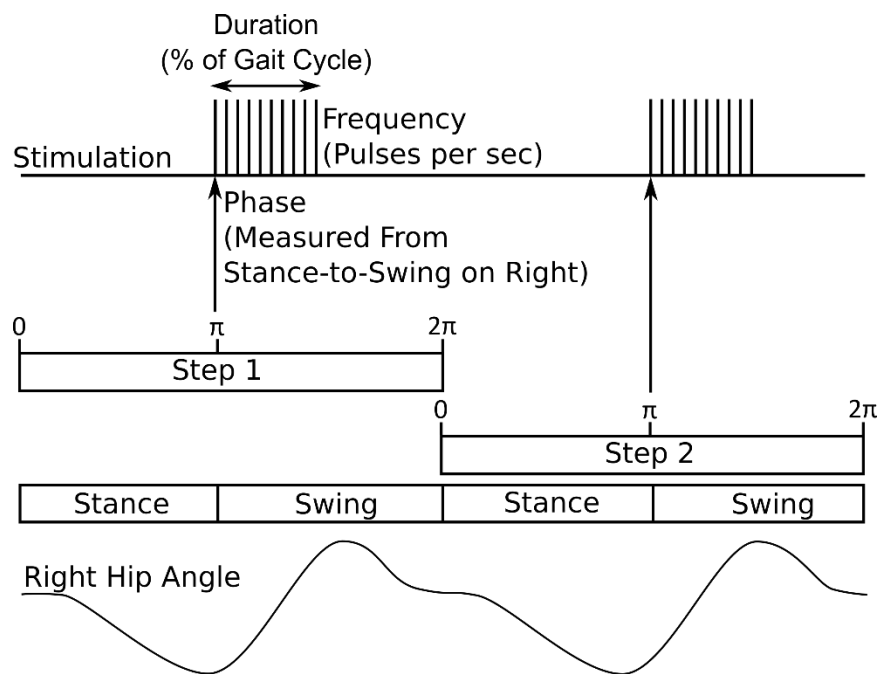


Figure 3-2: Pulse Train Parameterization. The pulse train was parameterized by three parameters: the frequency (stimuli per second), the duration (length of the pulse train, as a percentage of the gait cycle), and the gait phase (time in the gait cycle at the start of the pulse train). Steps were defined as beginning at the swing-to-stance transition on the right foot.

To parameterize pulse trains, I chose the following parameters (Fig. 2): gait phase (time within the gait cycle) at the start of the pulse train (e.g. mid-stance), stimulation frequency (e.g. 68 Hz), and pulse-train duration (e.g. 32% of the gait cycle). Using phase, frequency, and duration of the pulse train is common in many studies involving

sensory stimulation (e.g. (Perreault et al. 1995, Zehr, Komiyama, and Stein 1997, Field-Fote and Tepavac 2002)). For each of the parameters, I chose bounds based on the technical limitations of our system. Frequency is limited to 0-100 Hz. Pulse train duration is limited to 0-50% of the gait cycle. Phase forms a rotationally symmetric space, and we defined the space as one step-cycle mapping to $0-2\pi$. Together, gait phase, stimulation frequency, and pulse-train duration form a three-dimensional input space, from which a mapping to the output space is calculated.

3.2.2. Assessment: A Real-Time Walking Metric

In order to be part of an optimization loop, I needed to measure the quality of walking quantitatively (i.e. a walking metric). Our choice of walking metric was limited by two requirements: (1) the metric must be computed every step, and (2) the metric must produce viable results in people with even only minimal walking ability. Clinically, it is common to measure how far someone can walk within a time limit or how fast they can walk a specific distance, but such metrics do not meet the every-step requirement.

Initially, I explored EMG-based metrics, but four issues proved difficult to resolve. (1) Surface EMG is limited in what muscles are accessible. For example, the most important hip flexors, the iliopsoas, are inaccessible to surface EMG, particularly when the subject is wearing a body weight support harness over their pelvis. (2) EMG is sensitive to nearby electrical stimulation. EMG requires a large amount of amplification (500x), and the voltages of the stimuli are frequently $>1,000$ times the amplitude of the EMG signals, leading to amplifier saturation. This effect is distance-dependent, such that the closer the EMG electrode is to the stimulating electrode, the larger the effect. Therefore, even among the muscles that were recorded from, some had stimulation artifacts that were too large to be reliably removed by post-processing. (3) Surface EMG recordings are

poor predictors of muscle force or kinematic trajectories (De Luca 1997). (4) I found substantial inconsistencies in EMG patterns even in normal individuals in the Lokomat during our preliminary studies on this approach. With all of these downsides, I opted not to use EMG as our primary walking metric.

I chose, instead, to minimize the assistive force from the Lokomat as our primary outcome metric. To achieve a walking pattern, the Lokomat provides robotic assistance, applying varying amounts of force depending on what is required to get a limb segment to the correct place in space at the correct time. This assistive force is negatively related to the muscle force that the person is producing, thus it can be considered an error metric (i.e. the less force the person produces, the more force the Lokomat must produce). Incidentally, this force-error-analogy is still accurate even if the subject has spasticity, which the Lokomat would have to counter or resist in order to achieve the desired walking pattern. In addition, the use of the Lokomat's assistance allowed people with only minimal walking ability to participate. To create the walking metric, the assistive force is summarized with a mean-squared approach across all four actuated joints (hip and knee on left and right). I normalize this metric to have a value of 1 for the average response to the no stimulation condition, providing an intuitive visualization for which stimulation patterns are better or worse than no stimulation.

3.2.3. Optimization: Closing the Loop

To complete the loop (figure 3-1), an optimization algorithm was required that could explore the sensory-input-to-motor-output mapping as efficiently and thoroughly as possible. Most optimization algorithms are computationally efficient but not efficient in terms of "function calls" (i.e. individual tests of a set of parameters within a data collection session). For example, particle swarms, genetic algorithms, and simulated

annealing can require thousands of function calls. In our case, computational resources were not significantly limited. Each iteration lasts 4 steps, allowing the algorithm approximately 8 seconds to find the next stimulation pattern to test. However, the number of function calls is limited by both the number of steps a person can take in 15-20 minutes (450-600 steps), and patient safety issues (i.e. before they might develop pain, fatigue, or other clinical issues related to their SCI or being in the robotic orthosis).

There are good algorithms for local optimization in the presence of measurement noise (e.g. (Barton and Ivey Jr 1996)), but they require assumptions regarding the shape of the space (e.g. no multiple minima). A priori, there was no way guarantee only a single local minimum, so I pursued global optimization approaches. In addition, most optimization algorithms require gradients, which are not readily available in a black-box system like a human subject, and the system cannot spare the extra function calls required to compute the derivatives.

I chose an algorithm based on Efficient Global Optimization (EGO) (Jones, Schonlau, and Welch 1998), incorporating a modification to deal with measurement noise (Huang et al. 2006). I also made other modifications specific to our problem (rotationally symmetric space and time-variance, discussed below). Our optimization approach creates a meta-model (or response surface) of the space. Response surface approaches are effective when function calls (experiments) are limited in some manner (e.g. financially expensive, time consuming, etc.) and gradients are not available (Jones, Schonlau, and Welch 1998).

The optimization algorithm has two periods: an open-loop randomized sampling period (to initialize the parameter space), and a closed-loop optimized sampling period. At the beginning of the randomized sampling period, the algorithm generates 60 randomized stimulation parameters (inputs, or x-points) across the 3 input-dimensions (i.e. gait

phase of the start of the pulse train, stimulation frequency, and pulse-train duration), using a Latin hypercube (McKay, Beckman, and Conover 1979) to initially seed the 3D input space. The system tests these randomized x-points against the SCI subject in the Lokomat, producing force metric values (outputs, or y-points). Each x-point is repeated 4 times, but the first (the parameter-changing transition step) is removed from further analysis, yielding 3 tests per x-point.

After the randomized x-points are tested, the closed-loop optimization period begins. In each iteration, the algorithm constructs a Design-and-Analysis-of-Computer-Experiments (DACE) stochastic model (Sacks et al. 1989) of the space using the distance and correlation of the data gathered previous to that iteration. Our approach diverges from the classic DACE model by allowing the space to be rotationally symmetric in the gait-phase dimension (the 2nd dimension), because gait is periodic. To account for the rotationally symmetry, I calculate the minimum of the distance (d_k) in either rotational direction (the min-term for $k = 2$ in equation 1).

$$d_k(\mathbf{x}_i, \mathbf{x}_j) = \begin{cases} \min \begin{pmatrix} |x_i(k) - x_j(k) + 1| \\ |x_i(k) - x_j(k)| \\ |x_i(k) - x_j(k) - 1| \end{pmatrix}, & k = 2 \\ |x_i(k) - x_j(k)|, & k \neq 2 \end{cases} \quad (1)$$

The min statement is equivalent to finding the closest distance between one point and the other point in a rotationally symmetric space. This approach also diverges from the classic DACE model by allowing the system to be time-variant. I model time as an uncontrolled, but measureable input, thus treating it like another stimulus-parameter dimension. Therefore, the temporal distance (d_t , equation 2) is calculated as the amount of steps that have passed between the two x-points (i and j) being compared, normalized to a scale of $0 \rightarrow 1$ by dividing by the total steps (N_x).

$$d_t(i, j) = \frac{|i - j|}{N_x} \quad (2)$$

Aside from the two modifications discussed above (in equations 1 and 2), the remainder of the algorithm follows the standard approach (Huang et al. 2006, Jones, Schonlau, and Welch 1998). Briefly, the combined distance (d) is calculated between two x-points (x_i and x_j). The input-dimension-specific distances (d_k) and the time-dimension distance (d_t) are exponentiated by parameters p_k and p_δ respectively, then multiplied by θ_k and δ respectively. These parameters, as discussed later, are used to adjust the importance of the input- and time-dimensions to achieve the best model fit possible (discussed later).

$$d(x_i, x_j) = \sum_{k=1}^3 \theta_k (d_k(x_i, x_j))^{p_k} + \delta (d_t(i, j))^{p_\delta} \quad (3)$$

Once the combined distances (d) are calculated, I create a correlational matrix \mathbf{R} where the element at i, j is $r(x_i, x_j)$, as described in equation 4. The diagonal of the matrix ($r(x_i, x_j)$ for $i = j$) is set to 1. The off-diagonal terms utilize the exponential distance, which will range from $0 \rightarrow 1$. An additional regularization parameter (g) is added, allowing for measurement noise.

$$r(x_i, x_j) = \begin{cases} 1, & i = j \\ g e^{-d(x_i, x_j)}, & i \neq j \end{cases} \quad (4)$$

From these individual terms ($r(x_i, x_j)$) for every pairwise combination of x-points, the correlation matrix \mathbf{R} is created. This matrix (\mathbf{R}) is used to estimate the mean ($\hat{\mu}$) and variance ($\hat{\sigma}^2$) of the model using generalized least squares (equations 5 and 6). In those equations, “ \mathbf{y} ” is a vector containing the measurements, and “ $\mathbf{1}$ ” is a vector of ones.

$$\hat{\mu} = \frac{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{y}}{\mathbf{1}^T \mathbf{R}^{-1} \mathbf{1}} \quad (5)$$

$$\hat{\sigma}^2 = \frac{(\mathbf{y} - \mathbf{1}\hat{\mu})^T \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}\hat{\mu})}{N_x} \quad (6)$$

Once the mean ($\hat{\mu}$) and variance ($\hat{\sigma}^2$) are estimated, the algorithm computes the model likelihood, $p(\mathbf{y}|\mathbf{R})$. For convenience the algorithm maximizes a value proportional to the model likelihood (equation 7), instead of directly computing the model likelihood. It maximizes this measure by varying the model parameters (i.e. $g, \theta_k, p_k, \delta, p_\delta$) to select the best model.

$$p(\mathbf{y}|\mathbf{R}) \propto \frac{1}{\det(\mathbf{R})^{\frac{1}{N_x}} \hat{\sigma}^2} \quad (7)$$

After the algorithm identifies the model with the maximum likelihood given the observed data, as in (7), the algorithm must find the point that improves most on the current optimum, using the model (equations 8–12). However, this approach is complicated by the fact that same input may produce different outputs (i.e. measurement noise). Our algorithm, like the Sequential Kriging Optimization (SKO) (Huang et al. 2006), defines a utility function $u(\mathbf{x})$, in equation 11, to find the current optimum (\mathbf{x}^{**}).

In the first step (equation 8) of finding the next sample, estimates are made of the actual values ($\hat{Y}(\mathbf{x})$) for each x-point tested thus far, without measurement noise.

$$\hat{Y}(\mathbf{x}) = \hat{\mu} + g \mathbf{r} \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}\hat{\mu}) \quad (8)$$

In equation 9 and 10, the error in the estimate of $\hat{Y}(\mathbf{x})$ is calculated. First, the “controllable” variance ($\hat{\sigma}_Z^2$, i.e. the variance from changes in the input- or time- dimensions), as opposed to the total variance ($\hat{\sigma}^2$) or noise variance ($\hat{\sigma}_E^2$) is calculated from the parameter g . As discussed previously, g defines the relationship between the

total variance ($\hat{\sigma}^2$) and the “controllable” variance by controlling the regularization of the correlation matrix. Then the expected error in $\hat{Y}(\mathbf{x})$ is calculated, denoted $s^2(\mathbf{x})$.

$$\frac{\hat{\sigma}_Z^2}{\hat{\sigma}^2} = \frac{\hat{\sigma}_Z^2}{\hat{\sigma}_Z^2 + \hat{\sigma}_E^2} = g \quad (9)$$

$$s^2(\mathbf{x}) = \hat{\sigma}_Z^2 - \begin{bmatrix} 1 \\ \hat{\sigma}_Z^2 \mathbf{r}(\mathbf{x}) \end{bmatrix}^T \begin{bmatrix} 0 & \mathbf{1}^T \\ \mathbf{1} & \hat{\sigma}^2 \mathbf{R} \end{bmatrix} \begin{bmatrix} 1 \\ \hat{\sigma}_Z^2 \mathbf{r}(\mathbf{x}) \end{bmatrix} \quad (10)$$

Then, a utility function ($u(\mathbf{x})$) is calculated for every \mathbf{x} -point tested thus far in equation 11. This utility function combines both our estimates of which points produce better values ($\hat{Y}(\mathbf{x})$), and our uncertainty in those estimates ($s(\mathbf{x})$). The parameter c allows for a tradeoff between predicted objective values and predicted uncertainty. Following the SKO algorithm (Huang et al. 2006), our algorithm uses a value of 1 for the parameter c . The current optimum, \mathbf{x}^{**} , is the \mathbf{x} that maximizes that utility function is calculated in equation 12.

$$u(\mathbf{x}) = -\hat{Y}(\mathbf{x}) - cs(\mathbf{x}) \quad (11)$$

$$\mathbf{x}^{**} = \arg \max_{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N} (u(\mathbf{x})) \quad (12)$$

Once the current optimum is identified, the expected improvement ($EI(\mathbf{x})$) can be estimated for any point in the space, including previously untested points, as shown in equations 13 and 14. Equation 14 is only an expanded version of equation 13, showing how the probability density function (PDF) and the cumulative distribution function (CDF) are used to calculate the expected improvement ($EI(\mathbf{x})$). Following SKO (Huang et al. 2006), I also augmented the expected improvement to reduce the avoidance of selecting points near previous samples (equation 15).

$$EI(\mathbf{x}) = E[\max(\hat{Y}(\mathbf{x}^{**}) - \hat{Y}(\mathbf{x}), 0)] \quad (13)$$

$$EI(\mathbf{x}) = (\hat{Y}(\mathbf{x}^{**}) - \hat{Y}(\mathbf{x})) PDF\left(\frac{\hat{Y}(\mathbf{x}^{**}) - \hat{Y}(\mathbf{x})}{s(\mathbf{x})}\right) + s(\mathbf{x}) CDF\left(\frac{\hat{Y}(\mathbf{x}^{**}) - \hat{Y}(\mathbf{x})}{s(\mathbf{x})}\right) \quad (14)$$

$$AEI(\mathbf{x}) = EI(\mathbf{x}) * \left(1 - \sqrt{\frac{\hat{\sigma}_E^2}{s^2(\mathbf{x}) + \hat{\sigma}_E^2}}\right) \quad (15)$$

Now that the expected improvement of any potential point in the space can be calculated, the algorithm maximized that metric. I followed Huang et al.'s example of using a genetic algorithm to maximize the expected improvement (Huang et al. 2006). Although a branch-and-bound algorithm could also be use and is faster than a genetic algorithm (Jones, Schonlau, and Welch 1998), it required the parameter p_k to be constrained.

In some cases, the algorithm will select a point close to the current optimum; in other cases, it will select a point far from the current optimum. The goal of the algorithm that selects the next test point is to maximize the likelihood of improvement through measuring both uncertainty and expected values.

Once the next x-point is selected, it is applied to the system (three steps sequentially), and new outputs, or y-points, are recorded. Every fourth step, the stimulation parameters change, and the step on which the parameters change is removed from the data.

The algorithm here (with our modifications) has been validated with simulations in Chapter 6.

3.3. Experimental Results

In order to validate the system, I tested it with one spinal-cord-injured volunteer (32 years old, male, C6, motor incomplete, 5 years post-injury). The subject had a Lower Extremity Motor Score (LEMS) of 0: they could produce voluntary movement in a non-key muscle greater than 4 segments below the injury, but had no other voluntary control. The experiment was conducted at the Shepherd Center, an Atlanta-based rehabilitation hospital, with the approval of Shepherd Center's Institutional Review Board. Physical therapists confirmed the subject could participate safely in the study and the subject provided informed consent.

3.3.1. Experimental Protocol

A pair of bipolar stimulating electrodes were placed over the distal tibial nerve while the subject was supine on a mat. The stimulation current was set to the lowest level that could produce a consistent flexion reflex on each leg (20 mA on the right leg, 25 mA on the left leg). Pulse width was set to 1 ms.

The subject's legs were ace-wrapped to protect the skin from the Lokomat straps. The Lokomat body-weight-support harness was placed on the subject's torso and adjusted as appropriate. Then, the subject entered the Lokomat. The subject's pelvis was secured, the body-weight-support harness was attached to the body weight support system, the exoskeletal actuators were attached to the thigh and calf, and the front of the foot was lifted with a spring-loaded strap to avoid toe catch on the treadmill. Dynamic body weight support was set to 60% (the minimum level such that the subject could walk in the Lokomat). Walking speed was set to 2 kilometers per hour, which corresponded to approximately one complete step every two seconds. The experiment was begun once the subject was comfortable and ready.

The experiment was divided into four consecutive periods: 10 steps of no stimulation, 251 steps of randomized-stimulation parameters, 226 steps of the optimized-stimulation parameters, and 9 more steps with no stimulation (figure 3-3). In both the randomized-stimulation and optimized-stimulation periods, the stimulation parameters changed every 4 steps. In each group of 4 steps, the first step (the parameter-changing, transition step) was removed from further analysis to avoid artifacts. During the optimization period, the system automatically selected the stimulation parameters based on the closed-loop algorithm described previously (see optimization section above).




Figure 3-3: Experimental validation. The experiment consisted of a human SCI subject walking 496 steps in the Lokomat. The first half of all plots (first 261 steps, uncolored background) is the randomized-sampling period, which initialized the space for the optimization period (“optimizing” in the right half, consisting of the final 245 steps, gray background). The top panel (panel A) shows the frequency, phase, and du-ration of the pulse train vs the step number, the tested values are colored according to their metric value (i.e. how good the output was) in the bottom panel. The middle panel (panel B) shows the actual pulse trains (each hash is a pulse in the pulse train) from the stimulation parameters in the top panel. The pulse trains are shown with individual stimuli represented as hash marks. Every pulse train was repeated four times, but shown only once in the plots above (for ease of visualization). The bottom panel (panel C) shows the force metric in response to the pulse trains vs the step number, with only the median for each iteration shown. In the force metric, a value of 1.0 corresponds to the average response for no stimulation. Most pulse trains reduced the force metric (improved stepping), but some combinations increased the force metric to above 1 (worse than no stimulation).

Each step, the system applied pulse trains to each leg at the same gait phase for each leg. For example, if the right leg was stimulated at mid stance, the left leg was also stimulated at mid stance.

3.3.2. Randomized Sampling vs Optimized Sampling

There were qualitative differences between the randomized-sampling and optimized-sampling periods (figure 3-3). During randomized sampling, the sampling was spread (uniformly) across the entire parameter space without focus on any specific region of the parameter space. Most sets of stimulation parameters reduced the force metric (improved stepping), but some combinations increased the force metric to above 1 (worse than no stimulation). During optimized sampling, the system focused its sampling on stimulation parameters that produced better force metric values (i.e. lower force metric values), which tended to be stimulation patterns applied during the early swing phase. The algorithm did not always stay near the last iteration's sample. Instead, it always tried to maximize the likelihood of improvement over the entire space. As more samples were selected near the stance-to-swing transition, the chances that some hypothetical point will be better than all previous points in that area diminishes, and a point in a less explored region became more attractive to explore. In figure 3-3, as more data samples were gathered, the algorithm estimate of the optimum was near early swing in the phase space, and frequencies and durations in the top half of the frequency range.

The qualitative differences between the randomized-sampling and optimized-sampling periods were compared statistically (t-test) to determine if the optimization period sampled in a significantly better way than randomized sampling. The difference was statistically significant ($p < 0.001$), meaning that the optimization algorithm identified and sampled a better region (according to the force metric) of the stimulation-parameter space than a randomized sampling did.

The optimal pulse train begins just before the stance-to-swing transition (9% earlier) is high frequency (97 Hz), and lasts 36% of the gait cycle (~0.7 seconds). This optimum had the lowest (i.e. best) expected value for the force metric in the entire space.

3.3.3. Confidence Analysis

To quantify the accuracy of this optimum, I created confidence regions for both 95% and 99% confidences (Fig. 4a). To create the regions, the space was cut into 2.5% increments (40 divisions in each dimension), and used the model to estimate the response and standard deviation at every point. I also calculated the covariance between those points using the model. To increase the applicability and interpretability of this analysis, time-variance was ignored, so these values can be considered to represent the optimal stimulation parameters across the entire experiment. The estimates, standard deviations, and covariances at every increment created a large Gaussian process ($40 \times 40 \times 40 = 64,000$). Then, a Monte Carlo simulation was run, with 2,000,000 experiments. In each Monte Carlo experiment, I instantiated every point in the 3D space, and ordered the points by the force metric, constructing a probability distribution of the global minimum. This probability distribution was used to create confidence regions for both 95% and 99% confidences (Fig. 4a).

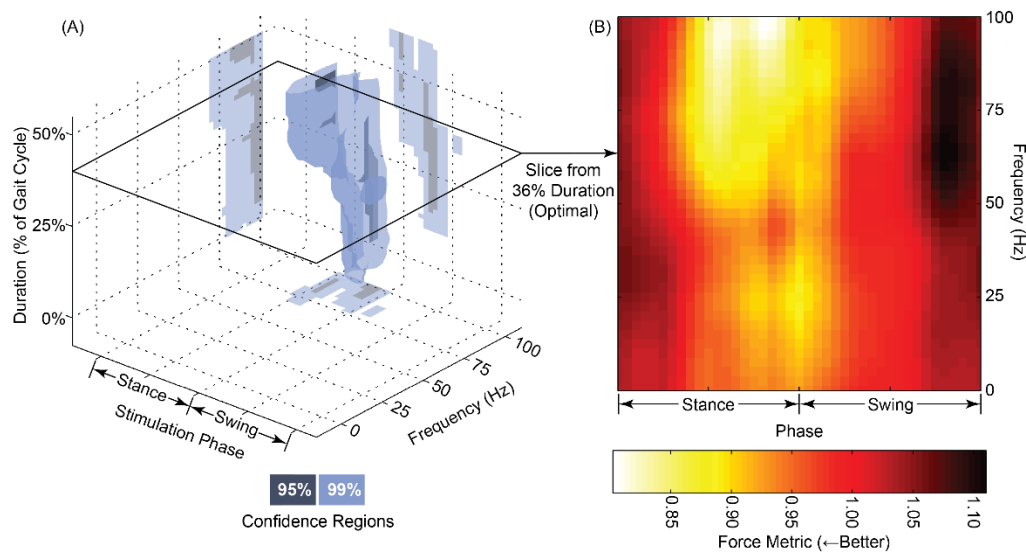


Figure 3-4. The optimal stimulation parameters, confidence regions, and input-to-output mapping. A model was created from the experimental data, including the expected value and standard deviation at every point in space, and the covariance between every combination of two points in the space. From this stochastic model, confidence regions (95% and 99%) of the location of the global minimum were calculated using a Monte Carlo simulation (Panel A on the left). To visualize the space near the optimal stimulation parameters, a 2D slice of the expected value space was made (Panel B on the right). A duration of 36% was chosen, as it is the optimal duration. The areas with better (i.e. lower) force metrics are concentrated at the highest frequency tested (top part of panel B has both the lightest colors), with starting times (phases) near the stance-to-swing transition. While most areas in the parameter-space reduced the force metric (improved stepping), some combinations increased the force metric to above 1 (worse than no stimulation).

The confidence showed a complicated distribution. There was one region focused on slightly shorter pulse-trains (e.g. 36% duration) starting later (e.g. 9% before the stance-to-swing transition), and another region with longer pulse-trains (e.g. 50% duration) starting earlier (~19% before the stance-to-swing transition). The maximum likelihood estimate of the global optimum is in the first region. Both regions are focused on the highest frequencies that the experiment tested.

In order to gain a better understanding of the input-to-output mapping, I calculated the expected value at every point. A slice was extracted from that 3D space at the optimal

duration (figure 3-4b) to focus on the frequency and phase relationship (those two parameters were identified the most sensitive in the next section). In addition to showing the good areas, this slice-visualization also shows the worst values. These worst values in the force metric were greater than 1, revealing that these stimulation parameters were worse than no stimulation at all. The worst pulse was a high frequency pulse train applied just before the swing-to-stance transition (~180 degrees away in the phase space).

3.3.4. Sensitivity Analysis

In order to test the quality of the hypothetical optimum I performed a series of tests, exploring both the optimum as compared to the rest of the space, and the local area near the optimum. In all cases, the estimates relied on the stochastic DACE model.

I analyzed the sensitivity of the optimum and model through perturbation analysis. The DACE model (the input-output model) was leveraged to perform a second-derivative test of the stimulation parameters at the optimum (figure 3-5). The second derivative evaluated at the optimum can be used to understand how tight the optimum is (i.e. how much better the optimum is than some other point near to the optimum) with respect to each stimulation parameter. Our analysis (figure 3-5) demonstrates that the sensitivity of the optimum (for at least this one test subject) using normalized units (normalized by the range of the stimulation parameter) and physical units (e.g. Hz for frequency). For our comparisons, normalized units were used, as they allow for easier comparison between different stimulation parameters.

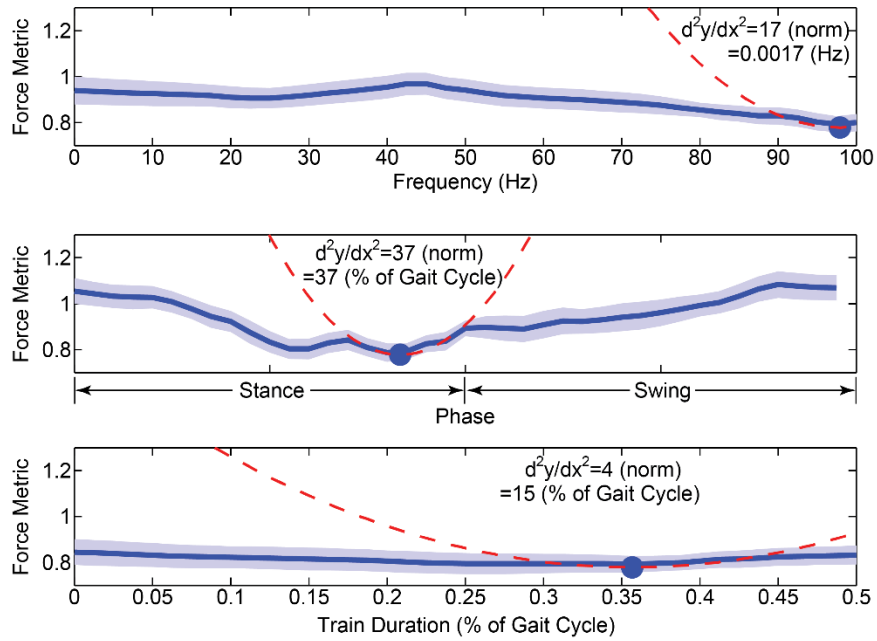


Figure 3-5. Sensitivity of the optimum. The optimum is shown with the large dot. The solid line shows the estimated response, the shaded area above and below the dark line shows the standard deviation. The dotted red line shows the second derivative, calculated at the optimum, and the values shown above the optimum report the sensitivity in both normalized (by the range of allowed values), and by standard units (e.g. Hz for frequency).

The analysis identified phase as the most sensitive parameter (37 in normalized units), followed by frequency (17 in normalized units), with duration of the pulse-train as the least sensitive parameter (4 in normalized units). These results (figure 3-5) mean that a small change in phase may make a big difference, but changes in duration do not affect the outcome variable (the force metric) very much.

3.4. Discussion

This chapter presented a novel approach: real-time modification of the walking pattern through optimizing sensory stimulation, demonstrating that individually optimized protocols for humans with SCI is feasible. Our approach dealt with the noisy, time-varying system embodied by a real human subject, and showed it can find a candidate optimum. Through the global optimization algorithm, I developed a stochastic nonlinear mapping of the space, allowing for confidence and sensitivity validations.

On-line optimization (or optimal experimentation) is a relatively new improvement on closed-loop approaches. It allows researchers to focus their experiments automatically on whatever they define as the most interesting subspace of their stimulation-parameter space. The approach detailed here worked in an actual system with a real human subject. The model allowed for time-variance, noise, and multiple local minima. More importantly, the approach identified a candidate optimum for this specific subject, which could be used in a rehabilitation setting on this specific subject.

3.4.1. Optimization was Successful

The optimal pulse train was a medium length pulse train (35% of the gait cycle) at a high frequency, starting just before the stance-to-swing transition. The majority of previous studies (e.g. Field-Fote and Tepavac (2002) and Granat et al. (1993)) utilize high-frequency, short-duration pulse trains starting at or slightly before the stance-to-swing transition. Our algorithm figured out a similar solution without previously being told what an effective pulse train would look like. The previous results have validated our approach (in one subject).

In addition to the pulse train discussed above, there was an additional local optimum of longer pulse trains, starting earlier (closer to mid-stance). These two local minima can be seen in the confidence and sensitivity plots. Although these two local minima represent pulse trains that start at different times, the timing of the middles of those pulse trains are similar, suggesting that the nervous system may “perceive” the middle of the pulse train as the independent parameter, rather than “perceiving” the starting of the pulse train as the independent parameter. However, this result is based on only one subject, and a detailed exploration of how the nervous system “perceives” pulse trains will be left to future studies.

I validated the protocol in two ways: confidence regions in the stochastic mapping and a sensitivity analysis exploring perturbations. The confidence regions and sensitivity are tight in the phase and frequency axes, but looser in the duration axis. In the frequency space, there was only a slight reduction in the performance from reducing the frequency, suggesting that the system is not particularly sensitive to the precise stimulation frequency, so long as it is high. Overall, these results suggest that when a pulse train is applied during the gait cycle is the most important, so long as the pulse train is high frequency. How long the pulse train lasts is a secondary consideration – with all associated caveats that, at this point, I can only say this is true for a single subject.

3.4.2. Limitations of Study

First, this study is meant to demonstrate a new system and show proof of principle – it was conducted in only one SCI subject. Second, the Lokomat affects the walking pattern through the constraints it puts on pelvic movements and ankle plantar flexion. Despite these constraints, I chose the Lokomat due to its advantages in a research setting. It allows people who cannot walk to achieve a stable walking pattern. The Lokomat also produces highly repeatable steps, which are very amenable to experimental optimization. Lastly, and most importantly for this work, the Lokomat outputs joint angles, joint forces, and body weight support in real time.

It is difficult to measure anything as complicated as walking on a 1-dimensional scale. I chose a force-based metric, because it was simple, easy to implement/understand, and it worked well for our purposes. In future work, I hope to explore other metrics, including those based on electromyography and/or kinematic descriptions of leg movements.

Although our choices of stimulation parameterization and walking metrics made the system approachable, they also affected what optimal stimulation patterns could be

found. However, the advantage of the methodology demonstrated in this paper is that it can be used with any low-dimensional parameterized intervention and any walking metric that can be defined on a step-by-step basis.

The approach detailed in this paper is ideal for a safe training environment where both the best and worst stimulation patterns can be explored safely. In the real world, an on-line optimization algorithm would need to be extrinsically safe – not allowing for occasional falls while walking across a street. Therefore, an additional term would need to be added to the optimization algorithm to encourage the next iteration to stay relatively local to its previous iteration, and to stay clear of known bad areas.

3.5. Conclusions

Our approach shows that real-time optimization of a human subject to assist their walking is feasible, even in the context of the noisy, nonlinear system that a human with SCI represents.

CHAPTER 4: OPTIMIZING SENSORY STIMULATION WAS SUCCESSFUL

This chapter presents the results of an optimized-stimulation study across a range of SCI subjects, across multiple sensory stimulation sites, and with different stimulation parameterizations. As in chapter 3, the algorithm optimized the stimulation parameters to minimize the force assistance required to maintain a walking gait. Across all subjects and stimulation sites, the optimal stimulation protocol produced better walking than three alternative stimulation protocols: an industry-standard stimulation condition, a no-stimulation condition, and a random-stimulation condition.

The chapter is organized into an introduction, methods, results, and discussion, following the traditional paper format. The introduction section presents an abbreviated summary of current research and motivation for this chapter. The methods section describes the subjects, subject preparation methods, and stimulation protocols. The results section presents the details of the findings across the subject population in the multiple stimulation sites, and the discussion section interprets the results.

4.1. Introduction

Typically, walking relies on three parts of the nervous system: central-pattern-generator (CPG) neural circuits in the spinal cord, supraspinal input, and sensory feedback. The CPG provides the basic motor pattern of walking (Duysens and Van de Crommert 1998). Communication with the brain initiates (Jordan 1998) and modulates locomotion (Marple - Horvat and Armstrong 1999). Sensory signals modulate the walking pattern by conveying information about biomechanical interactions between the body and the environment (Rossignol, Dubuc, and Gossard 2006).

Spinal cord injury (SCI) disrupts the pathways between the brain and the central pattern generator (CPG), impairing or eliminating the ability to walk. Despite substantial improvement from modern clinical interventions, most SCI subjects are not satisfied with the recovered level of ambulation they achieve post-injury.(Anderson 2004) For those with the most severe injuries, current treatments are generally not effective for the recovery of ambulation.

Sensory stimulation can initiate, control, and modulate locomotion in nearly all species examined (Kiehn et al. 1998). Two types of sensory stimulation have recently shown promise. First, cutaneous nerve stimulation has been used in conjunction with both functional electrical stimulation (Liberson et al. 1961, Granat et al. 1992) and with physical therapy to assist with training (Ladouceur and Barbeau 2000a, Ladouceur and Barbeau 2000b). Cutaneous reflexes are associated with the central pattern generator (Bussel et al. 1989, Burke 1999, Nakajima, Kamibayashi, and Nakazawa 2012), and studies have identified correlations between cutaneous reflexes and the recovery of locomotion after SCI (Dietz et al. 2009). Second, spinal-cord stimulation can assist the spinal cord to produce walking and walking-like behaviors (Gerasimenko et al. 1996, Herman et al. 2002, Harkema et al. 2011). Transcutaneous spinal cord stimulation (TSCS) has been developed as a non-invasive variety of spinal cord stimulation, and it has been shown to produce similar results as epidural stimulation (Ladenbauer et al. 2010, Danner et al. 2011, Hofstoetter et al. 2013). In both epidural (Rattay, Minassian, and Dimitrijevic 2000, Holsheimer 2002) and transcutaneous (Minassian, Persy, Rattay, Dimitrijevic, et al. 2007, Danner et al. 2011) spinal cord stimulation, the dorsal roots (sensory afferents) are activated at lower thresholds than any other spinal structures.

Sensorimotor processing is highly individualistic. Even in neurologically normal individuals, there is substantial variability in reflex responses based on age (Evans,

Harrison, and Stephens 1990, Gibbs et al. 1999), level of activity (Loeb 1993), athletic history (Wolpaw and Tennissen 2001, Gruber et al. 2007), and other variables (Brooke et al. 1997). A number of studies have found further divergence in sensorimotor responses after SCI, depending on SCI level (Dietz et al. 1999, Poirrier et al. 2004), SCI severity (Rossignol et al. 1996), and time since injury (Hiersemenzel, Curt, and Dietz 2000). In one study, all animals received identical injuries, but half had increases in extensor EMG, while the other half had reductions (Belanger et al. 1996), demonstrating individuality in even the most controlled conditions. In a human SCI study, the relationship between the early part of a reflex and a late part of that reflex varied significantly across the subject population (Roby-Brami and Bussel 1987). In studies of human rehabilitation after SCI, one study observed that no two individuals showed similar patterns of coordination after an intervention (Field-Fote and Tepavac 2002), and another study showed that individuals with SCI did not use consistent strategies to engage in an incline walking task (Leroux, Fung, and Barbeau 1999). There have been observations of individual SCI subjects in whom flexor stretches led to excitation of extensors instead of flexors, reversing the normal relationship (VanHiel 2012).

Real-time optimization can be applied to find the best sensory stimulation for an individual SCI subject by iteratively adjusting the sensory stimulation parameters to improve a walking metric. Real-time optimization has been applied to neurological applications previously. One of the early reports compared two approaches to algorithmically develop stimulation protocols for functional electrical stimulation (Kostov et al. 1995). Some neurophysiological studies have utilized real-time optimization to identify the best stimulus to get a specific response (Edin et al. 2004, O'Connor, Petkov, and Sutter 2005, Benda et al. 2007, Lewi, Butera, and Paninski 2008). Optimization and machine learning algorithms have been used to improve how brain-computer interfaces

interact with the user (Fruitet et al. 2012, Gürel and Mehring 2012). A recent study has used optimization to improve epidural spinal cord stimulation (Desautels 2014).

The goal of this work is to evaluate the performance of real-time optimization of sensory stimulation to assist walking in SCI subjects. I developed a system to optimize stimulation in real-time (Chapter 3), applied it to a group of SCI subjects, and tested multiple sensory stimulation sites. In each session, the system optimized the parameters to maximally improve walking in each subject individually. I hypothesized that this approach would produce optimal stimulation parameters that would outperform no-stimulation, random-stimulation, and clinical-standard stimulation protocols.

4.2. Methods

Sensory stimulation was optimized in real-time in a group of SCI subjects during a walking task in a Lokomat (Hocoma Inc.). Forces and joint angles were recorded (Figure 4-1). Lokomat-based walking is similar to normal walking, but not identical (Colombo, Wirz, and Dietz 2001). However, the Lokomat had three major advantages that allowed our study. (1) Assistive walking allowed us to recruit those subjects whose SCI was too severe to walk independently. (2) The Lokomat standardizes the step pattern, creating more favorable experimental conditions for testing different stimulation patterns. (3) Our Lokomat was equipped with an analog output box to access force traces in real-time, allowing analysis of how much assistance the SCI subjects required under different conditions.

The experiments were conducted at the Shepherd Center, an Atlanta-based rehabilitation hospital, with the approval of Shepherd Center's Institutional Review Board.

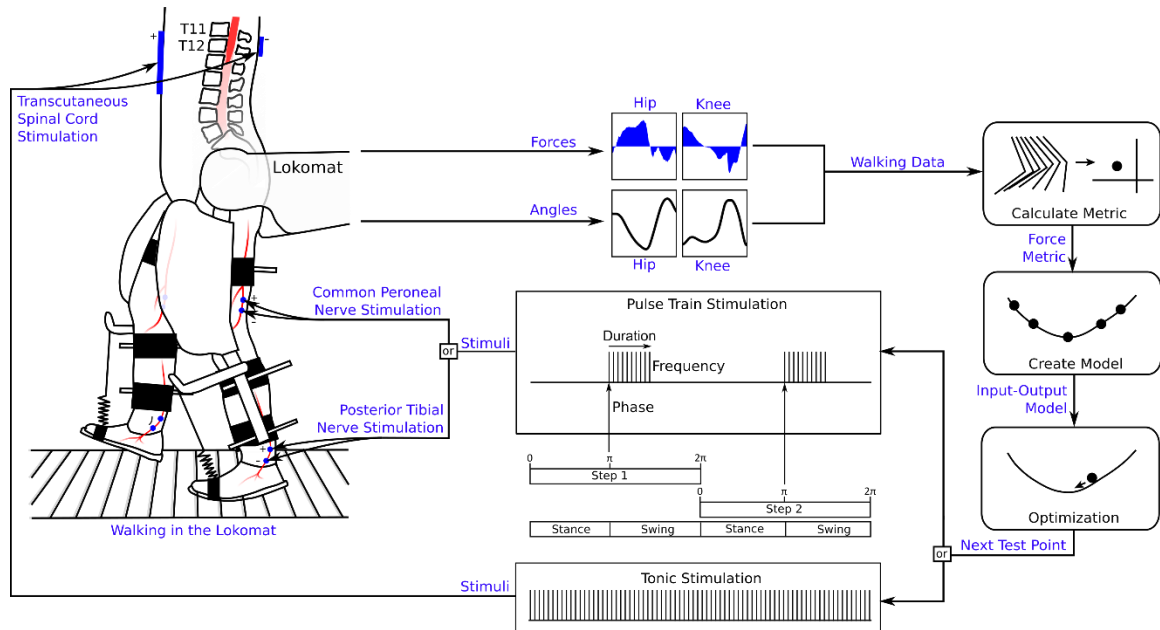


Figure 4-1. Conceptual Overview. The subject was walking in the Lokomat, and I recorded the amount of force assistance applied to maintain the walking pattern. This force assistance was summarized into the force metric. A model was created of all previous stimulation parameter tests and the next test point was selected via an optimization within the model. The next test point was translated into individual stimuli, depending on the parameterization being used (tonic or pulse-train stimulation). The stimuli were delivered to one stimulation site in each session (on both legs in the case of pulse-train stimulation), and more force metrics data points were recorded, beginning the loop again. The angles of the limbs were used to define the step cycle, and the steps were divided at the beginning of stance on the right leg. Pulse trains were parameterized using the gait phase of the beginning of the train, the frequency of stimuli, and the duration (as a percentage of one complete step cycle). Tonic stimulation was parameterized by the frequency of stimuli.

4.2.1. Subjects and Preparation

Subjects were recruited from the SCI population affiliated with Shepherd Hospital in Atlanta, consented on the study, and screened for safety in the Lokomat. Our subjects (Table 1) included both males and females, ranged from 18-57 years old, and were between 6 months and 10 years post-injury. Most of the subjects' injuries were AIS C's, but AIS A, B, and D were also included. The spinal level of the injuries ranged from C5 to T10.

After the subjects gave consent, they laid supine on a raised mat. EMG electrodes were applied (described in “Data Collected”). Stimulation electrodes were applied (described in “Stimulation”). Then, the subjects’ legs were wrapped to protect their skin from the Lokomat straps. The body weight support (BWS) harness was placed around their pelvis and torso and adjusted to fit. Then, they entered the Lokomat and their pelvis was secured. The BWS harness was attached to support system, and set to between 40–60%, aiming for the minimum BWS such that the subject could still walk in the Lokomat without exceeding the safety tolerances in the required force assistance. The Lokomat was augmented with a dynamic BWS system, which was used in all sessions to give smoother support. The actuators were attached to the thigh and calf. In order to prevent the subject from catching their foot on the treadmill, foot supports were attached, maintaining the angle of their ankle joint at slightly above 90 degrees with springs. Walking speed was set to ~2 kilometers per hour, which corresponded to approximately one complete step every two seconds. Guidance force was set to 100%.

4.2.2. Stimulation

Stimulation was applied to three different sensory stimulation sites: transcutaneous spinal cord stimulation (TSCS), posterior tibial nerve stimulation, and common peroneal nerve stimulation. For all stimulation sites, pulse width was set to 1 ms in order to be slightly more selective for sensory afferents (Veale, Mark, and Rees 1973). In all cases, the stimulation was applied with a DS7A (Digitimer Ltd.), triggered by a CompactRIO (National Instruments), running custom software. The number of experimental sessions was limited, so I prioritized posterior tibial nerve stimulation and transcutaneous spinal cord stimulation.

Transcutaneous spinal cord stimulation: TSCS was conducted per established protocol (Hofstoetter et al. 2013). Briefly, a pair of 2” bipolar stimulating electrodes was placed over the T11–T12 interspinous space, manually identified by palpation, with one electrode on either side of the spine, connected together to act as a single larger cathodal electrode. Before application of the electrodes, the area was rubbed with an abrasive gel to reduce the skin resistance. A ring of conductive gel was placed around the edge of the cathodal electrode to help eliminate skin problems from high-current stimulation. An additional larger pair of rectangular electrodes were placed over the lower anterior abdomen equidistant from the umbilicus. These were also connected together to act as a single larger anode. A constant-current stimulator (Digitimer DS7A) delivered monophasic rectangular pulses (1 ms pulse-width). Electrode placement was confirmed by eliciting posterior root-muscle reflexes in all EMG channels while the subject relaxed in a supine position. Amplitude was set to >60% of the current required to produce a reflex response in all EMG channels (the common mode action potential – CMAP). In our subject population, that amplitude range was from 22–110 mA, with most subjects at ~60 mA. The frequency of stimulation was optimized by the protocol discussed in the “Optimization” section.

Posterior Tibial Nerve: A pair of 1” stimulating electrodes was placed over the posterior tibial nerve while the subject was supine on a mat after the stimulation site was rubbed with an abrasive gel. The site was found by a trained physical therapist manually palpating the medial ankle, and confirmed with stimulation. The cathode (-) was always placed caudally, and the anode (+) was always placed rostrally. Setting stimulation current posed some difficulties. In many cutaneous reflex studies, stimulation current is often set to some multiple of the minimum perceptive threshold, but many of our subjects had impaired sensory perception, so an alternate approach was required. The

stimulation current (Table 1) was set to the lowest level that could produce a robust flexion reflex (while the subject was supine). This approach resulted in a range of 6–25 mA, with most subjects at ~12 mA.

Posterior tibial nerve and common peroneal nerve stimulation was applied using a pulse train, lasting less than 50% of the gait cycle. The pulse train was allowed to start at any arbitrary point in the gait cycle, last 10–50% of the gait cycle (duration), and stimulation between 1–100 Hz. Stimulation was applied to both legs at the same respective time in their gait cycles (out of phase in absolute time).

Common Peroneal Nerve: The procedure was very similar to that of the posterior tibial nerve, using the same electrodes, pre-rubbed with an abrasive gel, and applied while the subject was supine on a mat. The site was found by locating the lateral hamstring tendons on the lower thigh, and palpating slightly lateral to that tendon until the nerve was found. In some cases the electrodes had to be placed slightly lower in the popliteal fossa (crease in the back on the knee). In all cases the placement was confirmed with test stimuli. As with the posterior tibial nerve, I could not rely on perceptible threshold and instead found the lowest stimulation current (Table 1) that would produce a robust flexion reflex. The range of stimulus current used was 3–12 mA, with most subjects at ~8 mA.

PTNS and CPNS were applied using pulse trains (figure 4-1). The pulse train started at any arbitrary point in the gait cycle (phase), lasted 10–50% of the gait cycle (duration), with frequencies of 0–100 Hz. Both legs were stimulated at the same respective time in their gait cycles.

The *clinical-standard* stimulation protocols were defined based on previous reports in literature. For TSCS, effective stimulation frequencies are reported as 20-50 Hz

(Hofstoetter et al. 2013, Danner et al. 2015, Hofstoetter et al. 2014, Minassian et al. 2004). Therefore, I defined 35 Hz as clinical-standard. For PTNS and CPNS, studies generally utilize short, high-frequency stimulation, applied just before the stance-to-swing transition (Field-Fote and Tepavac 2002, Granat et al. 1993, Zehr, Komiyama, and Stein 1997). I defined “short” as 10% of the gait cycle, “high-frequency” as 85 Hz, and the clinical-standard gait phase as 50% of the gait cycle after heel strike (10% before swing).

4.2.3. Data Collected

Our system recorded the force that the Lokomat applied to the hip and knee joint on either side. The Lokomat was equipped with an analog output box, which enabled the recording of Lokomat-applied forces, body-weight-support (BWS), and joint angles. The CompactRIO was configured to receive these signals, along with the EMG channels. Our post-processing algorithms low-pass filtered the force, BWS, and angle signals at 20 Hz.

Table 4-1. Subject Details and Stimulation Sites tested on each subject. LEMS: Lower extremity motor score. AIS: ASIA Impairment Scale. ¹This subject was able to get voluntary movement in a non-key muscles greater than 4 segments below the injury, but had no other voluntary control. ²LEMS was unavailable for these subjects. TSCS: Transcutaneous Spinal Cord Stimulation. PTNS: Posterior Tibial Nerve Stimulation. CPNS: Common Peroneal Nerve Stimulation.

ID	Gender	Age (years)	Time since Injury (years)	AIS	Neurological Level	LEMS	TSCS Sessions	PTNS Sessions	CPNS Sessions
1	Female	30.2	1.5	A	T1	0	1		
2	Female	22.9	1.4	B	T4	0	1		
3	Male	32.1	5.0	C	C6	0 ¹	1	3	1
4	Male	57.1	3.5	C	T10	3	3	3	
5	Female	21.3	1.1	C	T3	6	1	1	1
6	Male	18.1	1.5	C	C6	7	1	4	
7	Male	21.1	0.6	C	T1	9	1	2	
8	Male	51.5	0.8	C	C7	15	2	2	
9	Male	41.5	4.1	C	C5	15	2	1	1
10	Male	42.4	9.5	C	T10	20		2	2
11	Male	40.6	1.6	C	T3	- ²	1	2	1
12	Male	49.5	1.9	D	C5	- ²	1	1	
	Ratio: 9M:3F	Mean: 35.7	Mean: 2.7	Range: A-D	Range: C5-T10	Range: 0-20	Total: 15	Total: 21	Total: 6

4.2.4. Force Metric

Our primary walking metric was to measure the amount of force the Lokomat required to maintain the subject in a walking gait. The choice of walking metrics was limited by two requirements: (1) the metric must be computed every step, and (2) the metric must produce viable results in people with even only minimal walking ability. To achieve a walking pattern, the Lokomat provides robotic assistance, applying varying amounts of force depending on what is required to get a limb segment to the correct place in space at the correct time. This assistive force is negatively related to the muscle force that the person is producing, thus it can be considered an error metric (i.e. the less force the person produces, the more force the Lokomat must produce). Incidentally, this force-

error-analogy is still accurate even if the subject has spasticity, which the Lokomat would have to counter or resist in order to achieve the desired walking pattern.

The use of the Lokomat's assistance allowed people with only minimal walking ability to participate in our study. The assistive force was chosen as our primary outcome metric, with less assistive force defined as better walking. The assistive force was summarized with a mean-squared operator across all four actuated joints (hip and knee on left and right) for every complete step (left-right). This metric was normalized to have a value of 1 for the average response to the no stimulation condition, providing an intuitive visualization for which stimulation patterns are better or worse than no stimulation. For statistical comparisons, I also corrected for time-variance through maintaining the average no stimulation near 1.0 through the experimental session using a Gaussian process model (the same model used for optimization).

4.2.5. Optimization

Our approach is discussed in detail in previously (Chapter 3). Briefly, our approach utilized a modified form of the sequential kriging optimization (SKO) algorithm (Huang et al. 2006). Our modification allows the use of time as an uncontrolled variable, thereby allowing time-variance in the optimization, in a similar manner as the N+1 approach described in (Morales-Enciso and Branke 2015).

Each experimental session began with open-loop sampling to seed the model space. Then, the optimization period began, and continued until the conclusion of the session. The same stimulation pattern was applied for 4–10 steps, corresponding to one iteration. Four steps were used for the three-parameter pulse-train stimulation patterns (PTNS and CPNS). Ten steps were used for the single-parameter tonic stimulation pattern (TSCS). Steps were divided at the onset of stance on the right leg (figure 4-1).

In post-processing, steps were divided into no-stimulation, random, clinical-standard, and optimal stimulation conditions. Each experimental protocol began with random stimulation patterns. Steps were identified as clinical-standard if they fell within a Cartesian distance of 0.33 in a normalized parameter space (all parameter ranges transformed to 0–1). For example, a 70 Hz, 20%-of-the-gait-cycle duration pulse train, with a gait phase of 40% (towards the end of stance) has a distance of 0.21 to the clinical-standard parameters (85 Hz, 10% duration, 50% gait phase), but a pulse train with a gait phase of 80% (during swing) would not be included. Optimal-stimulation steps were defined in the same way, but based on the distance to the optimum identified by the algorithm for that individual during that step. In some cases, individual steps could be considered in more than one protocol (e.g. if the algorithm identified 0 Hz as optimal for that step, the step would count as both a “no-stimulation” step and an “optimal” step).

4.2.6. Statistics

The data distributions often demonstrated skew, so Student T-Tests could not be used to test significance. Instead, the Wilcoxon Rank-Sums was used for all significance tests. Due to skew in the data distributions, Cliff’s Delta was chosen to measure effect size (Macbeth, Razumiejczyk, and Ledesma 2011). A Cliff’s Delta of 1 means every data point in the “better” group is better than every point in the other group. A Cliff’s Delta of 0 means that about the same number of points are better as are worse. A Cliff’s Delta of between 0 and 1 mean that most, but not all, points are better.

To estimate probability distributions the kernel density estimation function in Python’s scikit-learn library (v0.16.1) was used. For phase, the distribution was wrapped around the edge (i.e. the left-edge of the distribution was copied and placed beyond the bounds on the right). For other datasets, the data was reflected at the bounds to give more

accurate distribution-estimates for edge conditions (Cline and Hart 1991, Schuster 1985, Silverman 1986).

4.3. Results

Sensory stimulation patterns were tested at multiple sites in 12 subjects across 42 experimental sessions. In each session, our system recorded the force metric for every step and the stimulation parameters that produced each force metric value (figures 4-2A and 2B). The algorithm identified the optimal stimulation parameters that maximally reduced the force metric (described in methods). I quantified the probability distributions of the optimal stimulation patterns, in addition to no, random, and standard stimulation patterns. These protocols were then compared to identify the improvement due to individual optimization. To illustrate the optimal stimulation patterns, I reported the probability distributions of those parameters. Examples of time-variance in our responses are highlighted, and an analysis is shown that individual differences in the quality of optimization could not be easily derived from injury characteristics.

4.3.1. Comparison of Stimulation Protocols

In this analysis, our goal was to determine if on-line optimization produced better walking (as defined by a lower force metric) than no stimulation, random stimulation, or standard stimulation patterns. Our experiments included a random sweep of stimulation patterns for the first half of each session, allowing us to explore hypotheses on whether arbitrary stimulation patterns were equal or worse than the stimulation pattern identified by optimization (in terms of the force metric). The probability distribution of the force metric was analyzed (figure 4-2C) for each condition and statistically tested the differences.

The standard stimulation pattern was based on previous reports in the literature. For TSCS, the standard stimulation parameters was 35 Hz tonic stimulation. For common peroneal and posterior tibial stimulation, the standard stimulation was defined as short (<20% of the gait cycle), high-frequency (>70 Hz) pulse trains applied just before the stance-to-swing transition.

In all conditions (figure 4-2C), individually optimized stimulation produced statistically lower (better) force metrics than all other conditions, but the effect sizes varied based on which site was stimulated and what was being compared. In all three conditions, the median of random stimulation was nearly the same as that of no stimulation, but random stimulation tended to have a greater spread of values.

The improvement (effect size) from using optimal stimulation differed based on which stimulation site was being used (e.g. distal tibial vs TSCS). The improvements (from optimization) for common peroneal and posterior tibial nerves were much larger than those for TSCS (0.44–0.71 vs 0.12–0.22), as can be seen in figure 4-2C. Posterior tibial nerve stimulation showed slightly greater effect sizes than common peroneal nerve stimulation (0.47–0.71 vs 0.44–0.68).

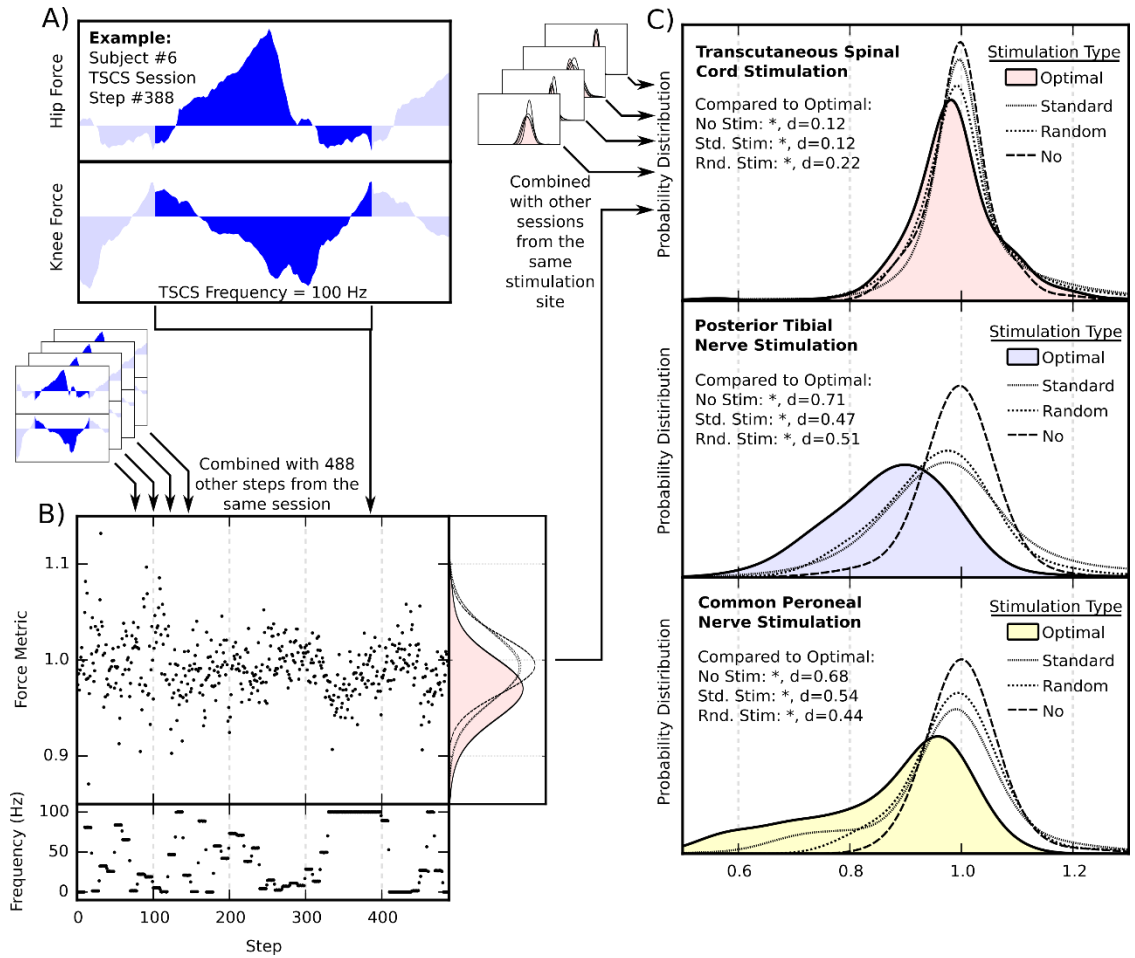


Figure 4-2. Optimal Stimulation as compared to no, random, and standard stimulation. Panel A shows an example step (Step #388), from subject #6, with a TSCS frequency of 100 Hz. Hip and knee forces are shown. Dark blue shows the current step, with portions of the previous and next step shown in light blue. This steps are combined (in order) with the other 488 steps from the same session in panel B. The bottom plot of panel B shows the TSCS frequencies which produced the force metric values in the main panel. The right plot of panel B shows the distribution of force metric values according to stimulation condition, where pink is optimal stimulation. In the top plot of panel C, the session from panel B is combined with 14 other TSCS sessions to produce the force metric distribution for the entire dataset. The second plot in panel C shows the distribution across all posterior tibial stimulated steps and sessions. The third plots in panel C shows the distribution across all common peroneal stimulated steps and sessions. In all three stimulation sites the optimal stimulation (colored shading) was statistically significantly lower than all three other conditions (no, standard, and random stimulation). The number of steps in each set was 6,740 (TSCS), 8,165 (posterior tibial stimulation), and 2,640 (common peroneal stimulation). I tested significance using a Wilcoxon Rank Sums test, and quantified effect size (“d” in the above) using Cliff’s Delta (range of -1 to 1, with 1 being the largest difference). Std. = Standard, Rnd. = Random.

4.3.2. Optimal Stimulation Parameters

Our goal in this analysis was to estimate the global probability distribution of optimal parameters for each sensory stimulation site. As part of the process of optimization, the algorithm estimates optimal stimulation parameters at every step. These parameters can be combined across the entire session to form a distribution of optimal parameters for that session. The distributions from each session can then be combined into global distributions of optimal stimulation parameters (figure 4-3). For the purposes of analyzing the distribution of these parameters, the complete model (future and past tests) was used to estimate the optimal stimulation parameters for every step in the experimental session.

The optimal parameters varied across subjects, and within the same subject it varied with time. However, when the distributions of all subjects' optimal stimulation parameters were combined, there was clustering in the data (peaks in figure 4-3).

Optimal stimulation parameters for TSCS were divergent (figure 4-3). A substantial minority (36%, subject IDs: 3, 4, 9, and 12) showed no stimulation or very low frequency stimulation as optimal, which was a permitted solution in the algorithm. Other subjects showed a spread across the non-zero frequencies.

The optima associated with posterior tibial stimulation also demonstrated clustering (figure 4-3). The optimal pulse trains tended to be high-frequency, with a peak distribution at 100 Hz. The optimal pulse trains tended to start in the late stance phase, with a peak at 44% of the gait cycle – approximately 2/3rds through the stance phase. Optimal pulse trains tended to be longer pulse trains, with the peak of the distribution at pulse trains that lasted 50% of the gait cycle (the highest value permitted by the

algorithm). Common peroneal nerve stimulation (figure 4-3) showed a similar distribution, with the peak of the phase of stimulation at 41% instead of 44%.

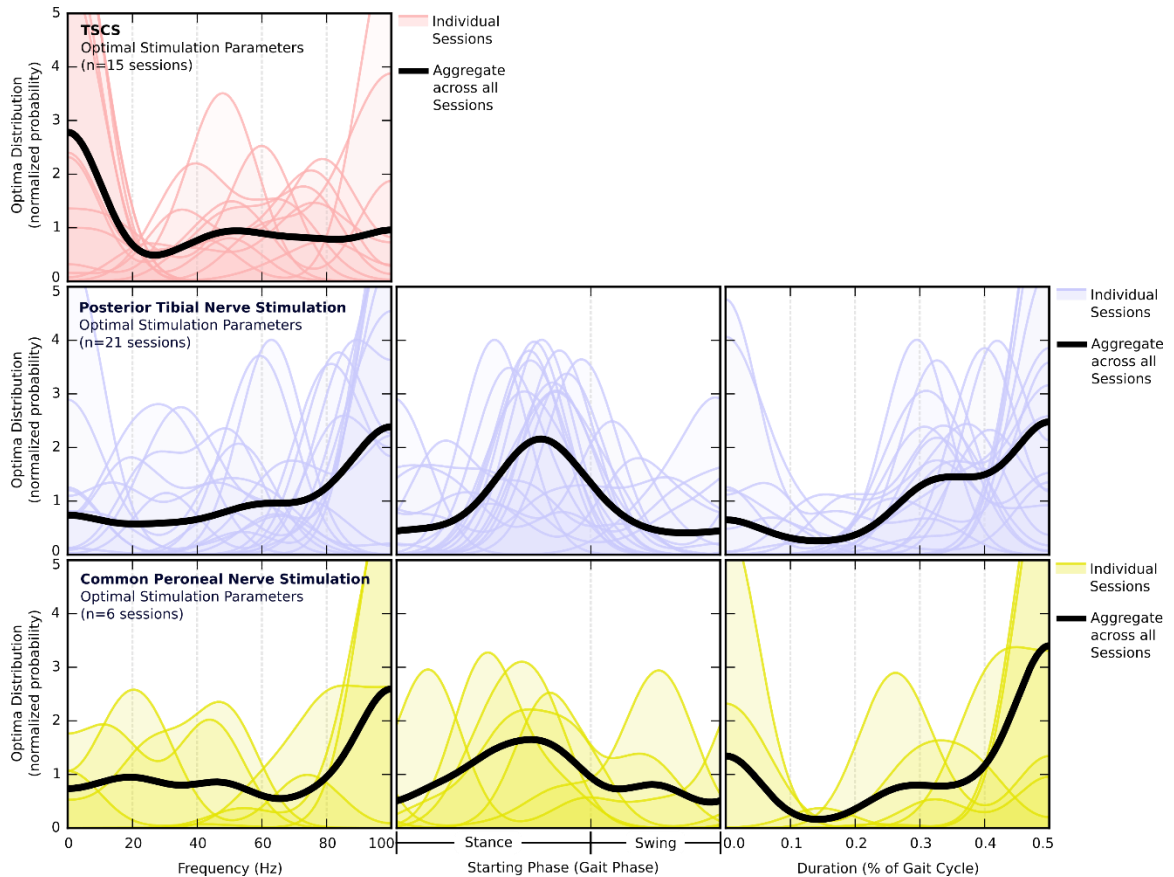


Figure 4-3. Distribution of Optimal Stimulation Parameters. The time-varying optima from each session were used to estimate the probability distribution of the optimal stimulation parameters in that session (the colored lines in the background of each plot). Then, the average of the distributions of all sessions was calculated to estimate the global distribution of optimal stimulation parameters (black line in the foreground in each panel). Transcutaneous spinal cord stimulation is on the top row. Posterior tibial nerve stimulation is on the second row. Common peroneal nerve stimulation is on the third row. The first column shows frequency of stimulation for all three stimulation sites. The second and third columns show phase of stimulation and duration of the pulse train respectively.

4.3.3. Time-Variance

In some of our subjects, the same input would produce different outputs at different times within a single experimental session (time-variance in the response). The qualitative types of time-variance differed between subjects. Examples of the types of

time-variance are shown in figure 4-4. In figure 4-4A, a non-time-varying example is shown – the same frequency of stimulation produces similar results at different times (with some noise in the process). Two of our subjects with the most severe injuries (AIS A and B) showed phenomena like that displayed in figure 4-4B. In figure 4-4B, there is a dynamic response to different stimulus frequencies in the first 50 steps, but after step 100, the responses slowly revert to the pre-experimental values, as the response is less and less affected by what frequency of stimuli is applied (i.e. habituation). In figure 4-4C, the subject requires more and more robotic force from the Lokomat to achieve a walking gait. In figure 4-4D, the subject requires less and less robotic force.

The response of most subjects changed only slightly with time. In those with time-variance, the only clear trend was that the most severely injured individuals (AIS A and B) tended towards habituation. For the less severely injured (AIS C and D) the trends were unclear. Regardless of the type of time-variance, our optimization algorithm accounted for these changes through the use of a non-stationary model.

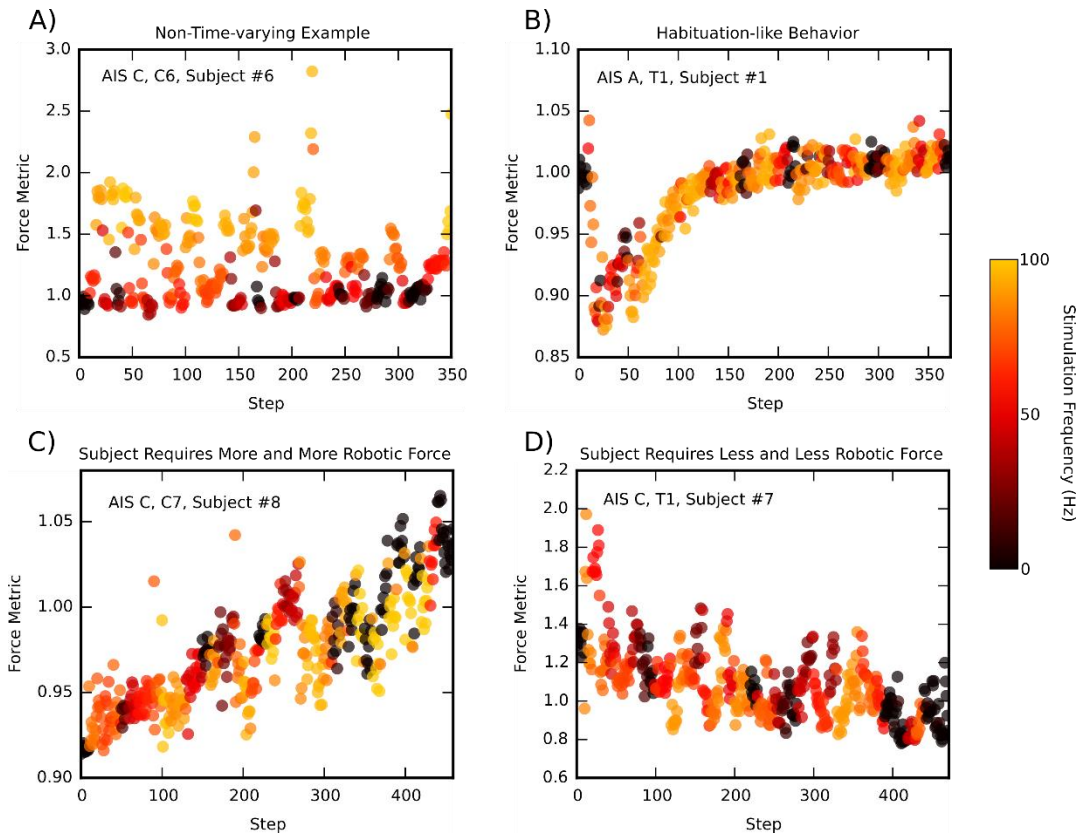


Figure 4-4. Examples of Time-Variance. Stimulation Frequency is shown by the color of the circles, black is zero or low frequency stimulation, yellow is high frequency stimulation. These plots show different ways that subjects may vary their response in time. In all the plots pay special attention to the same color circles (same frequency of stimulation) at different times. In the top left, an example plot shows a non-time-varying dataset. The top right shows a habituating subject. The bottom left plot shows a subject requiring more and more force assistance. The bottom right shows a subject requiring less and less force assistance.

4.3.4. Model Fit: Sensory Site and Injury Details

The quality of the model fit was analyzed to ascertain whether there were trends in optimizability with respect to injury types or sensory stimulation sites. The fit of the model provides a quantitative metric for this purpose, as the fit of the model determines how well the model can predict the output. Therefore, the fit of the model can tell how well the model can optimize and how confident we can be in any optimum produced.

There was substantial variability in our model's ability to predict the output (figure 4-5), with the percentage of variance explained ranging from 0.2–0.9 (20%–90%). The quality of the model fit was worse in the TSCS case (1D input) than in the posterior tibial or common peroneal nerve stimulation cases (3D input).

I found no statistically significant trends in our data as to the effects of injury type on the model fit (figure 4-5). In addition to plotting the data along AIS and LEMS, I also explored injury location (spinal segment) and time-since-injury, neither analysis demonstrated a statistically significant trend (data not shown).

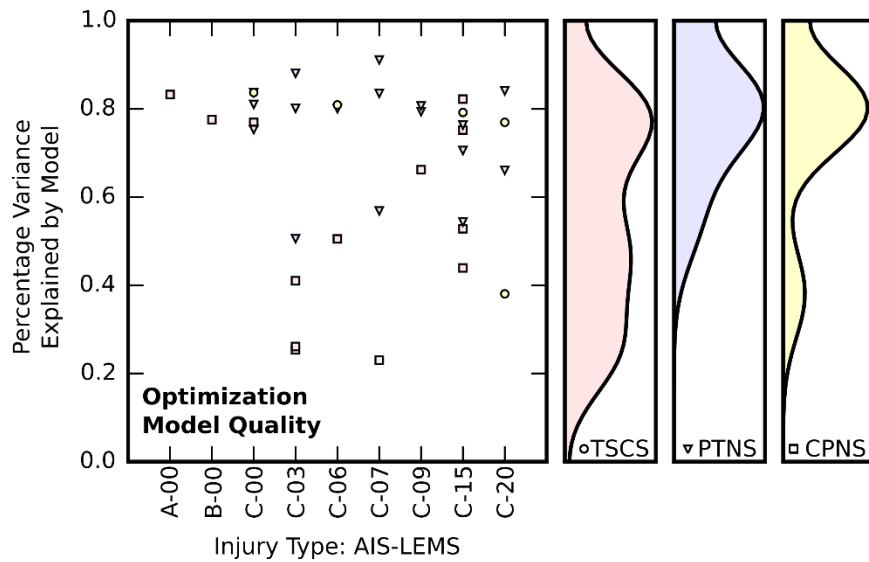


Figure 4-5. Model performance compared to injury type. The performance of the optimization algorithm is limited by the ability to model the data. Percent variance explained is a standard measure of the quality of a model. The left scatter plot shows the results for all three stimulation sites organized by injury details (AIS classification, followed by LEMS). In some subjects, LEMS was not available, and they were excluded from this figure. The right three plots show the distribution of explained variance based on the sensory stimulation site. TSCS: Transcutaneous spinal cord stimulation. PTNS: Posterior Tibial Nerve Stimulation. CPNS: Common Peroneal Nerve Stimulation. AIS: ASIA Impairment Scale. LEMS: Lower Extremity Motor Score.

4.4. Discussion

This study validated a novel approach to optimizing sensory stimulation to assist walking after spinal cord injury across twelve subjects and three stimulation sites. Our analyses

compared four different stimulation protocols at each sensory stimulation site, and identified the optimal stimulation parameters across the population for each site and evaluated the time-variation of the subject's response to stimulation.

Optimized stimulation produced lower (better) force-metric values than alternate stimulation protocols. I compared optimized stimulation, no stimulation, random stimulation, and a standard stimulation protocol, and found that the optimized stimulation protocol was statistically better than all alternatives, but the degree of improvement depended on the stimulation site. The two peripheral nerve stimulation sites demonstrated larger improvement than the TSCS site (figure 4-2). This difference can be partially explained by the poorer quality of the model fit for TSCS (figure 4-5). Model fit constrains how well the algorithm can predict the value of untested stimulation parameters.

For both posterior tibial and common peroneal nerve stimulation, the optimal stimulation parameters were similar. The best time to stimulate (for both sites) was to start pulse trains after ~2/3rds of the stance phase had passed, using a high-frequency, long-duration pulse train. The vast majority of previous studies (e.g. Field-Fote and Tepavac (2002) and Granat et al. (1993)) utilize pulse trains starting at or slightly before the stance-to-swing transition (starting ~15–20% of the gait cycle later than our optimal trains), using high-frequency, short-duration pulse trains. Our results can be considered a refinement of previous results, as the ideal duration of pulse trains is rarely ever studied. If a longer pulse train starts earlier than a shorter pulse train, the middle of the two pulse trains will be at approximately the same time.

For TSCS, our results demonstrate that our subjects fall into two groups (figure 4-3). In one group (~36% of subjects), no TSCS was best – TSCS only made the walking worse. In the other group (~64% of subjects), more stimulation was better. There were no

statistically significant trends from injury details that would predict which group a subject would fall into. Some previous studies have found that TSCS frequencies in the 20-60 Hz range have demonstrated positive effects on walking (Dimitrijevic, Gerasimenko, and Pinter 1998, Herman et al. 2002, Minassian et al. 2004), which is similar to our second group. Our discovery of people who do not benefit from TSCS could be due to the larger subject population we recruited as compared to many previous studies.

The time-variance of the sensorimotor responses was qualitatively analyzed, and I found divergent patterns in our subject population. Some subjects required more and more force assistance, some subjects required less and less force assistance, and some subjects' responses remained stationary during the experimental session. In those with the most severe SCI (AIS A and B), there was a habituation-like effect where the dynamic range of the sensorimotor response followed an exponential-like decay towards having no response to any stimuli. Other researchers have noted similar behaviors after severe SCI, and referred to it as EMG exhaustion (Dietz and Müller 2004).

There are many ways that this approach could be improved with future work, including applying it to overground walking, utilizing multiple sensory sites simultaneously, and incorporating it into a wearable device. The Lokomat offers many advantages for our study, including standardized-stepping, assisting weak stepping, and outputting the required assistive force to calculate a force metric. However, it would be more powerful if this approach could be applied to overground walking. Overground walking would require the development of new walking metrics that measured the quality of walking in real-time (i.e. step-by-step walking metrics). Although our results suggest that it may be more useful to stimulate phasically (using pulse trains) than tonically (TSCS), the gains from using both stimulation types at the same time are likely better than either sensory stimulation site alone. Lastly, this approach could be incorporated into wearable devices.

It could optimize the subject's walking on an ongoing basis, continually learning about the subject and devising better stimulation patterns.

CHAPTER 5: INSIGHTS FROM SENSORY STIMULATION

In this chapter, I analyzed a set of gait features to determine a more nuanced view of how stimulation parameters affected gait. During the optimization process, a random sampling initialized the space, and that random sampling included approximately 10,000 stimulated-steps across all stimulation sites and the subject population, providing an opportunity for data mining from an unbiased dataset. As in Chapter 4, three sensory-stimulation sites were included: (1) transcutaneous spinal-cord stimulation (TSCS, targets the dorsal roots), (2) posterior tibial nerve stimulation (PTNS), and (3) common peroneal nerve stimulation (CPNS). Tonic stimulation (variable-frequency) was used for TSCS, and pulse-train stimulation (variable frequency, phase, and duration) was used for the two cutaneous nerves. Assistive force was recorded from the Lokomat, and a separate system recorded electromyography (EMG). TSCS frequency had divergent effects on the stance and swing phases, high frequencies tended to assist with swing, but low stimulation frequencies tended to assist with stance. The cutaneous nerve stimulation (pulse trains) demonstrated that higher frequency and longer durations were better for almost all gait features. The optimal phase for the pulse trains was after mid-stance and before early swing, with some variability depending on which feature is examined. History (i.e. time spent walking/time spent being stimulated) was as important a predictor as any stimulation parameters for every stimulation site. In addition, I found a previously uncharacterized onset response to cutaneous nerve stimulation.

The chapter is organized into an introduction, methods, results, and discussion, following the traditional paper format. The introduction section presents an abbreviated summary of current research and motivation for this chapter. The methods section describes the subjects, subject preparation methods, and stimulation protocols. The results section

presents the details of the findings with respect to different times in gait, forces, EMG, and the onset response. The discussion section interprets the results.

5.1. Introduction

Typically, walking relies on three parts of the nervous system: descending pathways from the brain, central-pattern-generators (CPGs) in the spinal cord, and sensory signals from the periphery. The descending connections from the brain initiate (Jordan 1998) and modulate locomotion (Marple - Horvat and Armstrong 1999). The CPGs provide the basic motor pattern of walking (Duysens and Van de Crommert 1998). The sensory signals convey information on the biomechanical interactions between the body and environment to the nervous system, allowing for closed-loop control of the walking pattern (Rossignol, Dubuc, and Gossard 2006).

Spinal cord injury (SCI) disrupts the descending pathways from the brain, impairing or eliminating the ability to walk. Although clinical interventions have improved substantially in the last twenty years, most SCI subjects are not satisfied with the recovered level of ambulation they achieve post-injury (Anderson 2004), and for those with the most severe injuries, current treatments are generally not effective for the recovery of ambulation.

Sensory stimulation can initiate, control, and/or modulate locomotion in nearly all species examined (Kiehn et al. 1998). Two types of sensory stimulation have recently shown particular promise. First, cutaneous nerve stimulation has been used in conjunction with both functional electrical stimulation (FES) to induce hip flexion (Liberson et al. 1961, Granat et al. 1992) and with physical therapy to assist with training (Ladouceur and Barbeau 2000b, Ladouceur and Barbeau 2000a). Cutaneous reflexes have long been known to be associated with the walking-CPGs (Lundberg 1979, Bussel et al. 1989, Burke 1999, Nakajima, Kamibayashi, and Nakazawa 2012), and recent studies have

found the normalization of the longer-latency phase of the reflex is associated with recovery of locomotion after SCI (Dietz et al. 2009). Second, spinal-cord stimulation to assist locomotion has been the subject of several recent studies showing that it can assist the spinal cord to produce walking and walking-like behaviors (Harkema et al. 2011, Ichiyama et al. 2005a, Herman et al. 2002, Gerasimenko et al. 1996).

Transcutaneous spinal cord stimulation (TSCS) has recently been developed as a non-invasive variety of spinal cord stimulation (SCS), and it has been shown to produce similar results as epidural SCS (Ladenbauer et al. 2010, Danner et al. 2011, Hofstoetter et al. 2013). In both epidural SCS (Coburn 1985, Holsheimer 1998, 2002, Rattay, Minassian, and Dimitrijevic 2000) and transcutaneous SCS (Minassian, Persy, Rattay, Dimitrijevic, et al. 2007, Danner et al. 2011), the dorsal roots (sensory afferents) are the primary neural structure activated, as the dorsal roots are activated at lower thresholds than any other part of the nervous system.

Sensory stimulation can affect features of gait in a stimulation-parameter-specific manner. The response to cutaneous nerve stimulation changes depending on when in the gait cycle they are applied (Forssberg 1979, Yang and Stein 1990, Zehr, Komiyama, and Stein 1997), what stimulus strength is used (Roby-Brami and Bussel 1987, LaBella, Kehler, and McCrea 1989, Duysens et al. 1990), how long the stimulation is applied (Roby-Brami and Bussel 1987), and which cutaneous nerve was stimulated (Duysens et al. 1990, Van Wezel, Ottenhoff, and Duysens 1997, Zehr, Komiyama, and Stein 1997, Zehr and Duysens 2004). Cutaneous nerve stimulation can induce both flexion and extension. Non-noxious stimulation tends to produce ipsilateral flexion during stance, and ipsilateral extension during swing (Yang and Stein 1990, Zehr, Komiyama, and Stein 1997). Noxious stimulation biases the ipsilateral response towards flexion (Forssberg 1979, Duysens et al. 1990) or suppression of extension (Duysens et al. 1992). Flexion

and extension can also be induced on the contralateral leg (contralateral to the cutaneous nerve stimulation) (Duysens, Loeb, and Weston 1980, Rossignol, Julien, and Gauthier 1981). Cutaneous nerve stimulation can modify the timing of gait, prolonging extension or beginning extension earlier (Schomburg et al. 1998). SCS is less well-studied than cutaneous nerve stimulation, but it has been shown to have a frequency-dependent effect on gait. Frequencies above 22.5 Hz were more conducive for walking (Danner et al. 2015). Stimulation frequencies above 50 Hz reduced EMG activity (Gerasimenko et al. 2003a), and can correct some motor pathologies (Hofstoetter et al. 2014). Another study found slower EMG bursts with increasing frequency in the 30-70 Hz range (Dimitrijevic, Gerasimenko, and Pinter 1998).

Despite the observed dependency of the responses on stimulation parameters, the best stimulation pattern to achieve specific effects is generally unknown. The optimal frequency of pulse trains is poorly studied, despite well-known frequency-dependent effects in the central nervous system, such as long-term potentiation (Randic, Jiang, and Cerne 1993) and wind-up (Herrero, Laird, and Lopez-Garcia 2000). The optimal duration of pulse trains is also poorly studied, despite observed differences in responses based on the duration (Roby-Brami and Bussel 1987, Conway, Hultborn, and Kiehn 1987, Gossard et al. 1994, Frigon, Sirois, and Gossard 2010). Even the phase of the stimulation in the gait cycle has not been optimized – forcing our clinical approaches to rely on approximate descriptions rather than quantitative comparisons of efficacy. SCS is much newer by comparison, but what frequency will optimally improve a specific feature of gait is unknown.

To answer how stimulation parameters affected different gait features, I collected data from twelve SCI subjects with a range of injuries (AIS A–D) using three stimulation sites. Electromyography and the variable assistive force were recorded from a robotic orthosis

during a walking task. I hypothesized that analyzing this data would allow us to learn what stimulation parameters were most effective for which gait features.

5.2. Methods

Subjects were stimulated using random stimulation-parameters during a walking task in a robotic orthosis (Lokomat, Hocoma Inc.). I recorded electromyography, force traces, and joint angles (figure 5-1). Robotic-orthosis-based walking is similar (but not identical) to normal walking in terms of the EMG pattern (Hidler and Wall 2005), the evoked responses (Nakajima, Kamibayashi, and Nakazawa 2012), and kinematic trajectories (Hidler, Wisman, and Neckel 2008). However, the Lokomat had two major advantages that allowed our study: (1) the use of assisted-walking allowed us to recruit those subjects whose SCI was too severe to walk independently, and (2) our Lokomat was modified to output force traces, allowing analysis of how much assistance the SCI subjects required to walk under different stimulation conditions.

The experiment was conducted at the Shepherd Center, an Atlanta-based rehabilitation hospital, with the approval of Shepherd Center's Institutional Review Board.




Figure 5-1. Overview of methods. In each experimental session, one stimulation site was studied. Tonic stimulation was applied to the transcutaneous spinal cord stimulation (TSCS) site. Pulse train stimulation was applied to the common peroneal and posterior tibial nerve stimulation sites. The subject walked in the Lokomat while each set of stimulation parameters was applied for 4–10 steps. The angles of the Lokomat joints and the amount of force the Lokomat applied were recorded, along with electromyography from four muscles on each leg. Each session included 130–200 random-stimulation-parameter steps. EMG: Electromyography. Quad: Quadriceps muscle. Ham: Hamstring muscle. Tib Ant: Tibialis Anterior. Tri Sur: Triceps Surae.

5.2.1. Subject Recruitment and Preparation

Subjects were recruited from the SCI population affiliated with Shepherd Hospital in Atlanta, consented on the study, and screened for safety in the Lokomat. Our subjects (Table 1) included both males and females, ranged from 18–57 years old, and were 0.5–10 years post-injury. Most of the subjects' injuries were AIS C's, but subjects were also included with injuries of AIS A, B, and D severities. The spinal level of the injuries was C5–T10.

Table 5-1. Subject Details. LEMS: Lower extremity motor score. AIS: ASIA Impairment Scale. ¹This subject was able to get voluntary movement in a non-key muscles greater than 4 segments below the injury, but had no other voluntary control. ²LEMS was unavailable for these subjects.

Subject ID	Gender	Age (years)	Time since Injury (years)	AIS	Neuro-logical Level	LEMS
1	Female	30.2	1.5	A	T1	0
2	Female	22.9	1.4	B	T4	0
3	Male	32.1	5.0	C	C6	0*
4	Male	57.1	3.5	C	T10	3
5	Female	21.3	1.1	C	T3	6
6	Male	18.1	1.5	C	C6	7
7	Male	21.1	0.6	C	T1	9
8	Male	51.5	0.8	C	C7	15
9	Male	41.5	4.1	C	C5	15
10	Male	42.4	9.5	C	T10	20
11	Male	40.6	1.6	C	T3	-
12	Male	49.5	1.9	D	C5	-
	Ratio: 9M:3F	Mean: 35.7	Mean: 2.7	Range: A–D	Range: C5–T10	Range: 0–20

After the subjects consented, the subject laid supine on a raised mat. EMG electrodes were applied (described in the “Data Collected and Post Processing” section).

Stimulation electrodes were applied (described in “Sensory Stimulation” section). Then, the subjects’ legs were wrapped to protect their skin from the Lokomat straps. The body weight harness was placed around their pelvis and torso and adjusted to fit. Then, they entered the Lokomat and their pelvis was secured. The body weight support harness was attached to support system. Body weight support (BWS) was set to 40–60%, aiming for the minimum BWS such that the subject could still walk in the Lokomat without exceeding the Lokomat’s maximum force tolerances for the required force assistance. Our Lokomat had a dynamic BWS addition, which was used in all sessions to give smoother support. The actuators were attached to the thigh and calf. In order to prevent the subject from catching their foot on the treadmill, I attached foot supports that held their ankle at slightly above 90 degrees with springs. Walking speed was set to 2

kilometers per hour, which corresponded to approximately one complete step every two seconds.

5.2.2. Sensory Stimulation

Stimulation was applied to three different sensory stimulation sites: transcutaneous spinal cord stimulation (TSCS), posterior tibial nerve stimulation (PTNS), and common peroneal nerve stimulation (CPNS). For all stimulation sites, pulse width was set to 1 ms. One-millisecond pulses are more selective for sensory afferents than shorter pulses (Panizza, Nilsson, and Hallett 1989, Lagerquist and Collins 2008, Kitano and Koceja 2009, Wu et al. 2011). In all cases, the stimulation was applied with a Digitimer DS7A, triggered by a CompactRIO (National Instruments), running custom LabVIEW software.

In general, I prioritized posterior tibial stimulation and TSCS, utilizing common peroneal nerve only when there was additional time. PTNS and CPNS are similar, and it was important to have as much data as possible in one of the two. The result that part of the posterior tibial nerve reflex response had been correlated with SCI recovery (Dietz et al. 2009), and that the posterior tibial nerve innervates the bottom of the foot (as opposed to the top of the foot) convinced us that it might be a better choice.

Transcutaneous spinal cord stimulation (TSCS): A pair of 2" bipolar stimulating electrodes was placed over the T11–T12 interspinous space, manually identified by palpation, with one electrode on either side of the spine, connected together to act as a single large cathodal electrode. Before application of the electrodes, the area was rubbed with an abrasive gel to reduce the skin resistance. A ring of conductive gel was placed around the edge of the cathodal electrode to dilute the precipitous drop off in conductivity at the edge. In previous studies in the lab, this approach helped to eliminate skin problems from electrical stimulation. An additional larger pair of rectangular

electrodes were placed over the lower anterior abdomen equidistant from the umbilicus (the “belly button”). These were also connected together to act as a single larger anode. A constant-current stimulator (Digitimer DS7A) delivered monophasic rectangular pulses of 1 ms pulse-width. Electrode placement was confirmed by eliciting posterior root-muscle reflexes in all EMG channels while the subject relaxed in a supine position. Amplitude was set to >60% of the current required to produce a reflex response in all EMG channels (the common mode action potential – CMAP). One subject showed a motor response to TSCS at a lower threshold (22% of CMAP), and I chose to include their data. The frequency of stimulation was optimized by the protocol discussed in the remainder of this section.

Posterior Tibial Nerve Stimulation (PTNS): A pair of 1” bipolar stimulating electrodes was placed over the posterior tibial nerve while the subject was supine on a mat after the stimulation site was rubbed with an abrasive gel. The site was found by a trained physical therapist manually palpating the medial ankle, and confirmed with stimulation. The cathode (-) was always placed caudally, and the anode (+) was always placed rostrally. Setting stimulation current posed some difficulties. In many studies, stimulation current is set to some multiple of the minimum perceptible threshold, but many of our subjects had impaired sensory perception, so an alternate approach was required. The stimulation current (Table 2) was set to the lowest level that could produce a robust flexion reflex while the subject was supine.

Common Peroneal Nerve Stimulation (CPNS): The procedure was very similar to that of the posterior tibial nerve. I used the same electrodes, pre-rubbed the skin with abrasive gel, and applied the electrodes while the subject was supine on a mat. The cathode (-) was always placed caudally, and the anode (+) was always placed rostrally. The site was found by locating the lateral hamstring tendon on lower thigh, and palpating

until the nerve was found. In some cases the electrodes had to be placed slightly lower in the popliteal fossa (crease in the back on the knee). In all cases the placement was confirmed with test stimuli. As with the PTNS, I could not rely on perceptive threshold and instead found the lowest stimulation current (Table 2) that would produce a robust flexion reflex.

Pulse trains were used in both PTNS and CPNS. The pulse train was allowed to start at any arbitrary phase (time) in the gait cycle, with a duration of 10–50% of the gait cycle and a frequency of 0–100 Hz. Stimulation was applied to both legs at the same respective time in their gait cycles (out of phase in absolute time).

Table 5-2. All experimental sessions completed, organized by subject and site stimulated. *This subject had a substantial response to tonic stimulation, despite the stimulus amplitude being set to only 22% of the CMAP. Params: Parameters. Stim: Stimulation. BWS: Body weight support. CMAP: Common motor action potential (the amplitude at which all muscles are recruited).

Subject Number	TSCS (n=15)			Posterior Tibial (n=21)			Common Peroneal (n=6)		
	Stim Amplitude (mA, Used / CMAP)	Lokomat Params (BWS, Speed (km/hr))	Steps (Random / Optimizing / Total)	Stim Amplitude (mA, Right/Left)	Lokomat Params (BWS, Speed (km/hr))	Steps (Random / Optimizing / Total)	Stim Amplitude (mA, Right/Left)	Lokomat Params (BWS, Speed (km/hr))	Steps (Random / Optimizing / Total)
1	48/60	60%, 2.0	59						
2	65/100	50%, 2.1	75						
3	90/120	60%, 2.0	268	25/20	60%, 2.0	261	15/15	60%, 2.0	249
				25/20	60%, 2.0	261			
				25/20	60%, 2.0	261			
4	90/150	40%, 2.0	310	12/8	50%, 2.0	261			
	110/160	50%, 2.0	310	12/15	40%, 2.0	261			
	110/160	50%, 2.0	233	12/8	50%, 2.0	214			
5	50/70	50%, 2.0	267	20/17	60%, 1.7	261	5/5	60%, 1.7	261
6	22*/100	60%, 2.0	310	20/20	40%, 2.0	261			
				20/20	40%, 2.0	261			
				20/20	40%, 2.0	170			
				9/18	60%, 2.0	261			
7	45/65	40%, 2.0	310	6/6	40%, 2.0	262			
				6/6	40%, 2.0	173			
8	60/100	40%, 2.0	135	10/10	40%, 2.0	262	15/15	40%, 2.0	262
	80/130	50%, 2.0	141						
9	75/110	40%, 2.0	310	10/8	40%, 2.0	261			
	60/100	40%, 2.0	309	13/7	40%, 2.0	261			
10				6.5/6.5	40%, 2.0	261	4/4	40%, 2.0	262
				7/9	40%, 2.0	262	5/3	40%, 2.0	262
11	40/75	60%, 2.0	309	16/16	60%, 2.0	261	12/12	60%, 2.0	261
				16/16	60%, 2.0	262			
12	86/140	40%, 2.1	135	20/20	40%, 2.0	262			
	Total Steps: 3604			Total Steps: 5260			Total Steps: 1557		

5.2.3. Data Collected and Post Processing

Electromyography (EMG), hip/knee force (including body-weight support), and hip/knee angles were recorded.

EMG electrodes were placed over the quadriceps muscles (above the rectus femoris), the hamstring muscles (equal coverage of biceps femoris and semitendinosus), the tibialis anterior, and the triceps surae (equal coverage of the lower medial gastrocnemius and the soleus). I also attempted to record from hip adductors (aiming at the gracilis) and the gluteus maximus, but our results were inconsistent in those two muscles, so I did not report EMG from those muscles. I selected these muscle groups in order to get a representative flexor and extensor from each accessible joint. The major hip flexors (e.g. the iliopsoas) are difficult to access via surface recordings – only the femoral triangle exposes any part of the iliopsoas to surface EMG – and even that small area is inaccessible while the subject is wearing a weight bearing harness.

The recordings were made using a Motion Lab Systems MA300 EMG recording system. Each EMG site was recorded using a 20x preamplified bipolar electrode (500x total gain with the main amplifier). The amplifier included a hardware low-pass filter at 1 KHz to prevent aliasing. The system recorded the EMG traces using a National Instruments CompactRIO at 2 KS/s, 14 bits/sample. In software the signals were high-pass filtered at 100 Hz, and stimulus artifacts were removed by subtracting an estimate of the stimulus artifact. This estimate was created from the average post-stimulus response 0–10 ms after each pulse. Then, to produce the EMG waveforms, the filtered signals were full-wave rectified and decimated to 100 samples/step.

Our system recorded the force that the Lokomat applied to the hip and knee joint on either side. Our Lokomat was equipped with an analog output box, which enabled the

recording of Lokomat-applied forces, body-weight-support (BWS), and joint angles. The CompactRIO was configured to receive these signals, along with the EMG channels. The post-processing algorithms low-pass filtered the force, BWS, and angle signals at 20 Hz using a Butterworth filter (applied both forward and backwards in time to eliminated phase artifacts).

5.2.4. Experimental Protocol

Once the subject was in the Lokomat and walking at the correct speed and body-weight support, I began the study. The stimulation patterns were chosen from a uniform, sparse sampling of the stimulation-parameter space (a Latin Hypercube). This sampling approach had the additional advantage of temporally unbiasing the results.

Each stimulation pattern was applied for 4–10 steps before changing. Four steps were used for 3-dimensional pulse-train stimulation patterns, applied at the posterior tibial nerve and common peroneal nerve. Ten steps were used for the simpler 1-dimensional tonic stimulation pattern, applied at the TSCS site. For the purposes of both changing the stimulation patterns every 4–10 steps, and for analysis, steps were cut at the onset of stance in the right leg.

5.2.5. Features

To quantify the improvement in different aspects of walking, features were identified and quantified. These features include both EMG and force features, with varying degrees of granularity. The features all consisted of combining one or more channels (e.g. hip and knee forces) over a specific time-range of the gait phase (e.g. 10–50% of the gait cycle) with a mean-squared operator.

The primary force-features were hip-stance, hip-swing, knee-stance, and knee-swing. In all cases, stance was defined as the first 60% of the gait cycle, and swing was defined as the last 40%. The gait cycle was defined by hip angle, approximating the time when the heel struck the treadmill. Other force features (e.g. short-term force-features) included both channels (hip and knee) at specific times in the gait cycle.

Defining the EMG features proved to be more difficult. Our goal was to measure EMG at times in the gait cycle when muscles should be active, but there is a large variability in what constitutes normal EMG (Arsenault, Winter, and Marteniuk 1986, Winter and Yack 1987). Instead, I identified periods when muscles were active across the whole population. These included quadriceps-early-stance, quadriceps-stance-to-swing transition, hamstring-swing-to-stance, tibialis-anterior-early-swing, and triceps-surae-stance features. Superficially, some of these features seem counterintuitive (e.g. quadriceps-stance-to-swing), but other researchers have observed similar EMG activity in the same muscles at similar gait-times during overground, treadmill, and Lokomat-based walking (Hidler and Wall 2005, Lee and Hidler 2008, Nymark et al. 2005).

The feature-values were collected for every session-leg combination (e.g. session 23, left leg). These sets of feature-values were z-score-normalized (mean was subtracted, then the value was divided by the standard deviation).

5.2.6. Statistics

Our input-to-output model was a 4D-to-1D mapping (3 stimulation patterns and history). It is difficult to model every combination of the 4 parameters to estimate the output (a forward mapping), so I modeled the input with respect to the output (a reverse mapping), in other words, our analysis was focused on what combinations of stimulation

parameters produced the best and worst feature-values (described in the “Features” section).

In order to identify the best stimulation parameters for each feature, the dataset was split into the 20% best for that feature, and the remainder (80%). In some figures, the 20% worst data-points (with respect to the feature) were also separated. From these subsets, I cross-referenced stimulation-parameters for the included steps, and created probability distributions from the set of stimulation-parameters for each session. Then, the resulting distributions were averaged to create the population-as-a-whole distribution. In all cases, I normalized the 20%-best distribution by the overall distribution to form a marginal probability density. The marginal probability density eliminates any sampling bias, representing the actual relationship between stimulation-parameters and features more accurately. All distributions were calculated using the kernel density estimation routine in Python’s scikit-learn library (v0.16.1). The data was reflected at the bounds to give more accurate distribution-estimates for edge conditions (Cline and Hart 1991, Schuster 1985, Silverman 1986).

If the reader is unfamiliar with these techniques, the plots can be read in the following manner: If the resulting distribution is flat (uniform), the parameter had no predictive value (i.e. it did not affect the feature). If the peaks and valleys were large, the parameter had a large predictive value (i.e. a strong effect on the feature).

5.3. Results

To quantify the improvement in specific aspects of walking, EMG and force features were identified and quantified. These features summarized the activity at specific joints (e.g. hip) or specific muscle groups (e.g. quadriceps) over a functionally defined period of the walking gait (e.g. phase). The optimal (with respect to each feature) stimulation-

parameters from each stimulation site were identified. I also explored the timing of force assistance in gait using short-term force features to differentiate specific periods of gait. EMG and force features were compared to better understand how muscle activation affected the amount of assistive force required. Then, major sources of variability both within subjects (history) and across subjects (individuality) were identified and characterized. Lastly, I identified a previously-uncharacterized onset-response to cutaneous nerve stimulation.

5.3.1. Forces

The first analysis was designed to better understand how stimulation parameters affected specific joints at specific times during the gait cycle. The force traces for hip and knee were divided into different temporal segments (e.g. hip-stance). These features allowed us to quantify changes in specific aspects of the required robotic assistance (force), and what stimulation parameters were most effective.

First, it is necessary to clarify that the force measured was robotic assistance. The best walking occurred when the subject required the minimal assistance. Therefore, lower values of force at either joint during any phase of gait were better than higher values.

For TSCS, the best values (least force assistance) during stance and swing showed divergent behaviors with respect to frequency (figure 5-2C). Both hip-stance and knee-stance forces showed the best values at lower frequencies, peaking at 0 Hz. Both hip-swing and knee-swing showed the best values at higher frequencies, peaking at 100 Hz, the highest frequency our experimental protocol included.

Across all force features, the best values (least force assistance) for both cutaneous nerve stimulation sites (PTNS and CPNS) were similar (figure 5-2C). The ideal phase of stimulation peaked in the region between latter half of the stance phase and the early

third of swing phase for all force features for both cutaneous nerve stimulation sites. The ideal stimulation phase (peaks in phase plots in figure 5-2C) was earliest in the gait cycle for knee-stance (approximately 50% through the stance phase), and latest for hip-stance (following the stance-to-swing transition). Both hip-swing and knee-swing had similar peaks. For posterior tibial nerve stimulation, hip-stance force-assistance was minimized by stimulation that begun during the stance-to-swing transition ipsilaterally (corresponding to early stance contralaterally). Both hip-swing and knee-swing forces were minimized by pulse trains beginning approximately at mid-stance. Higher-frequency and longer-duration pulse trains tended to produce better (lower) force-feature values.

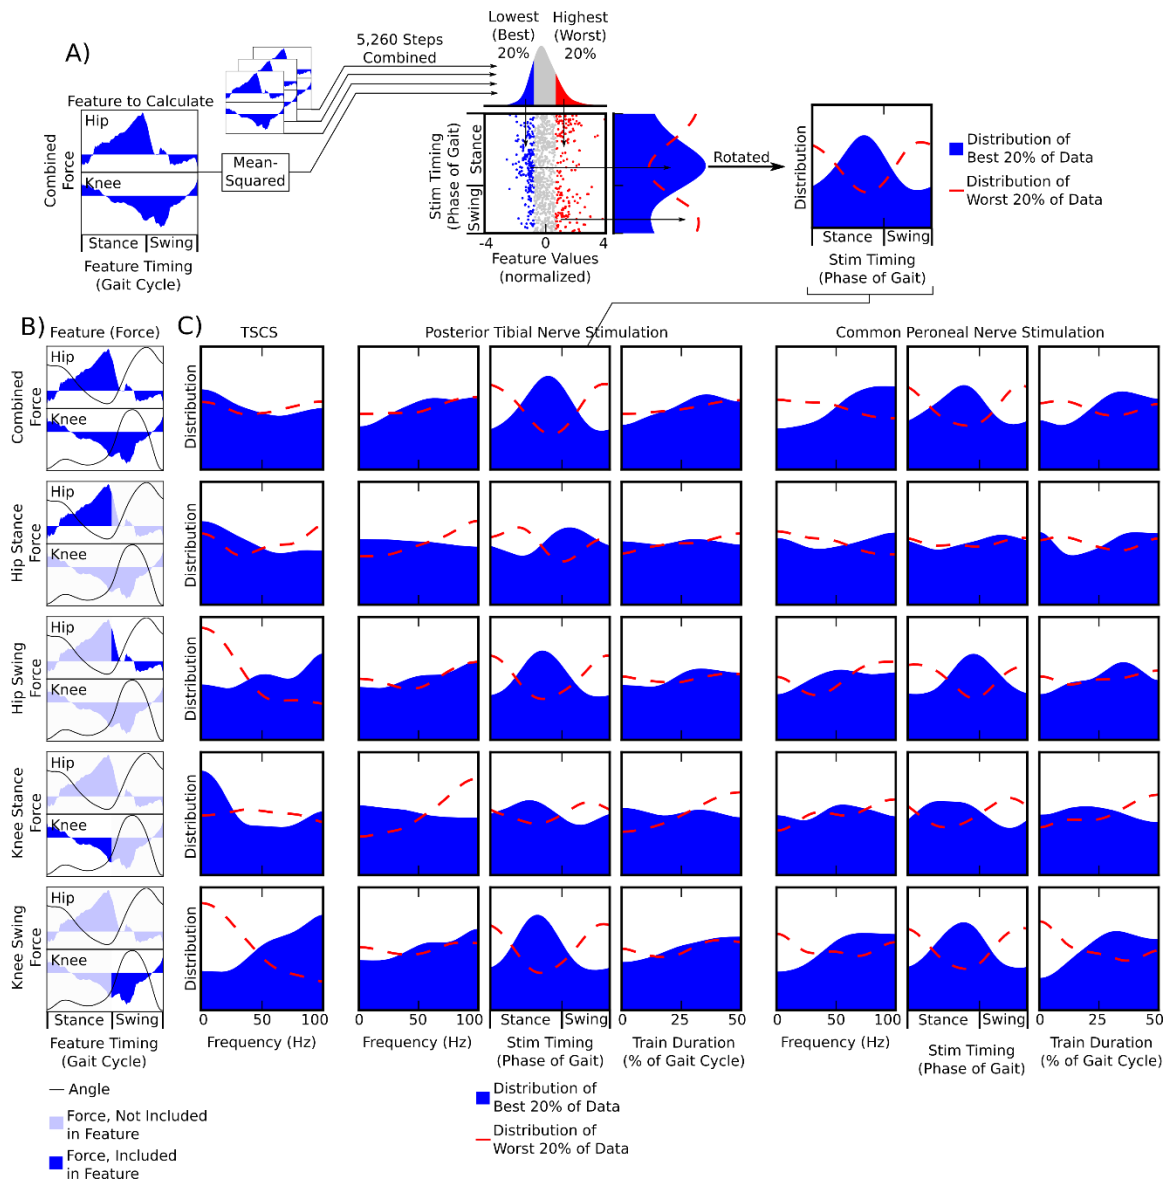


Figure 5-2. Force Features vs Stimulation Parameters. Panel A shows how distributions were created. First, the feature was calculated (mean-squared) for each step for each sensory stimulation site. The best and worst 20% of the feature-values were collected separately. For each feature-value, I cross-referenced the stimulation-parameters, and created best-20% and a worst-20% probability distributions of stimulation parameters for each feature. The left column of the lower plots (panel B) shows the specific force features calculated for each row. The dark blue shows when the force is included in the feature. The light blue shows when it is excluded from the feature. The black line shows the joint angle. In panel C, each column shows one stimulation parameter. The best and worst 20% are plotted with a blue fill and a dotted red line, respectively. In all cases the distribution of the best and worst 20% were normalized to the overall dataset to eliminate sampling bias. To interpret these plots, flatter distributions mean less predictive power. Larger peaks and lower valleys mean more predictive power. As in other plots, stance was defined to be the first 60% of gait, and swing was defined to be the last 40%.

5.3.2. Short-Term Force Analysis

In order to better understand the effects of stimulation on the timing of forces in gait, I split the gait cycle into 20 small, overlapping time-windows. This analysis can explore both the gait-phase dependence of the stimulation and possibly suggest the underlying mechanisms of the stimulation.

In figure 5-4, the TSCS results show that lower-frequency stimulation tended to assist more during stance, but higher frequency stimulation tended to assist near the stance-to-swing transition. The results in late swing were mixed. There are also qualitative differences between three periods: (1) stance phase, (2) stance-to-swing transition to the early swing phase, and (3) late swing phase.

The waveforms from PTNS demonstrated four distinct periods. The short-term force features from mid-stance to immediately before the stance-to-swing transition showed a peak in late stance that gradually delayed with the force-features. The peaks shifted quickly earlier for the short-term force features that included the stance-to-swing transition. Then, the peaks slowly delayed again during the swing phase, before becoming flat (i.e. ineffective) for short-term force features at the swing-stance transition. The slowly delaying peaks during the stance phase are actually after the force-feature they are measuring ipsilaterally, but just before the force-features contralaterally.

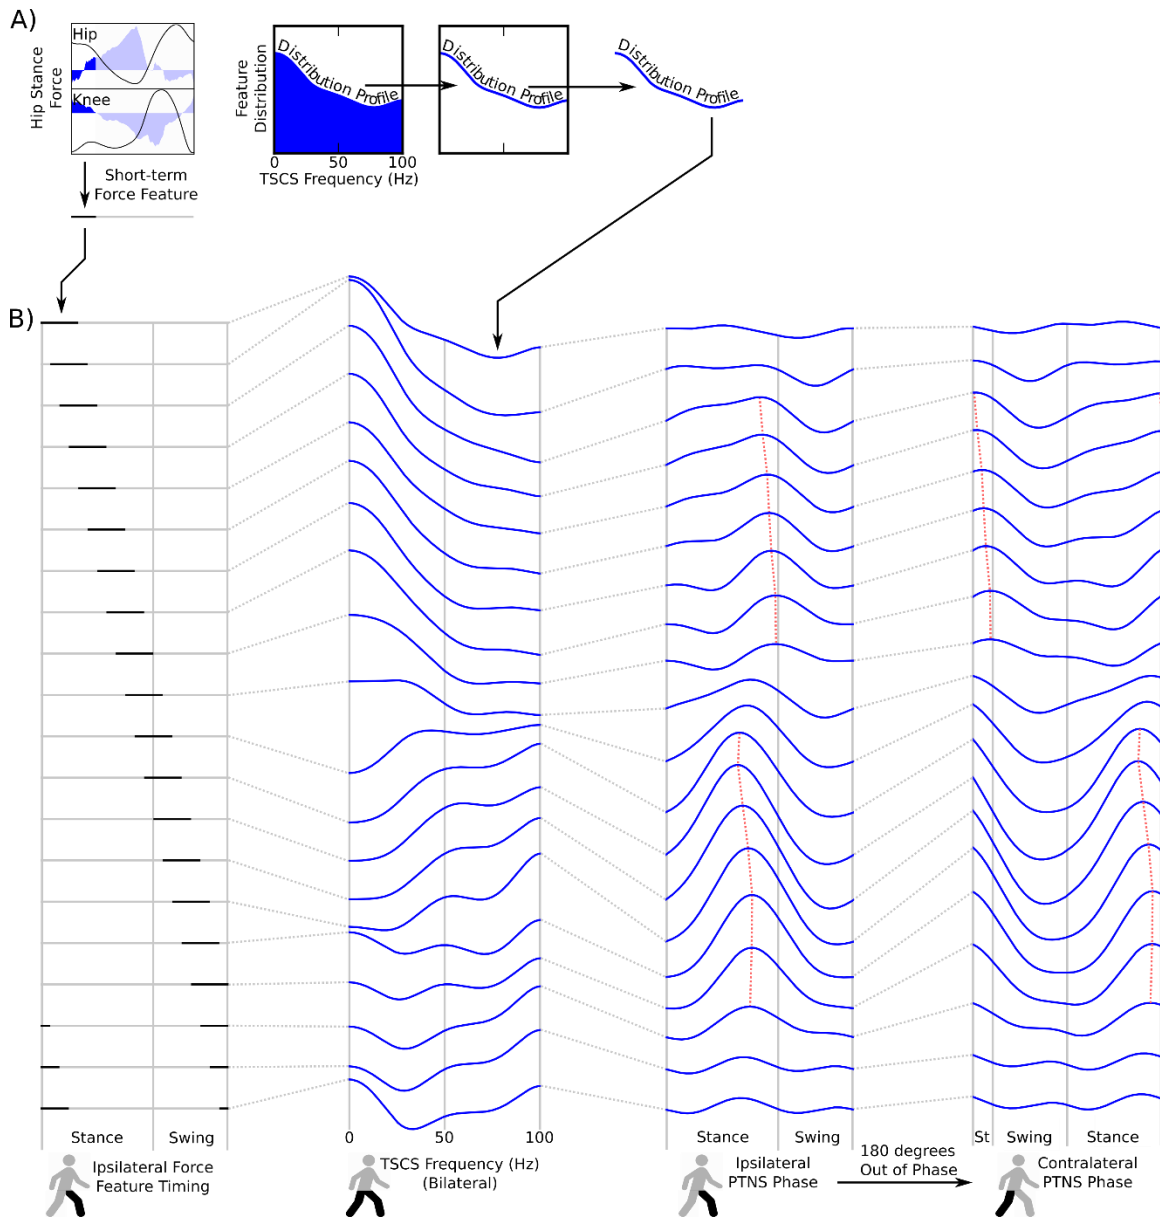


Figure 5-3. Short-term Force Features. I created short-term force-features to analyze the effect of TSCS frequency and cutaneous nerve stimulation phase on specific times/phases of gait. The force from the hip and knee were combined, and a distribution of the stimulation parameters for the best-20% was created (panel A) for both TSCS-frequency and PTNS-phase. The profile of that distribution was extracted for every feature and plotted in panel B. In the left column of Panel B, the black line shows the timing of each of the 20 overlapping, short-term force-features (each delayed 5% in the gait cycle from the preceding one). The second column of panel B shows TSCS-frequency distributions for each feature. The third and fourth columns show the PTNS-phase distributions for each feature. The two columns show the same data, with the contralateral phase shown delayed by half of the gait cycle. On the bottom of panel B, there are 4 icons of people, to clarify the stimulated leg as compared to the measured leg. A red, dotted line was added to assist the reader visualize the peaks. St: Stance.

5.3.3. Multiple Feature Analysis

In real scenarios, it is generally necessary to optimize against multiple parameters simultaneously. This section demonstrates an example of finding stimulation frequencies that produce good outcomes for both hip-stance and hip-swing forces.

Hip stance forces tend to be better at lower TSCS frequencies (figure 5-4). Hip swing forces tend to be better at higher frequencies. I tested whether there were TSCS frequencies that produced good (lowest 33%) values in both hip-stance and hip-swing forces. The distribution of those “good” frequencies peaked at approximately 52 Hz.

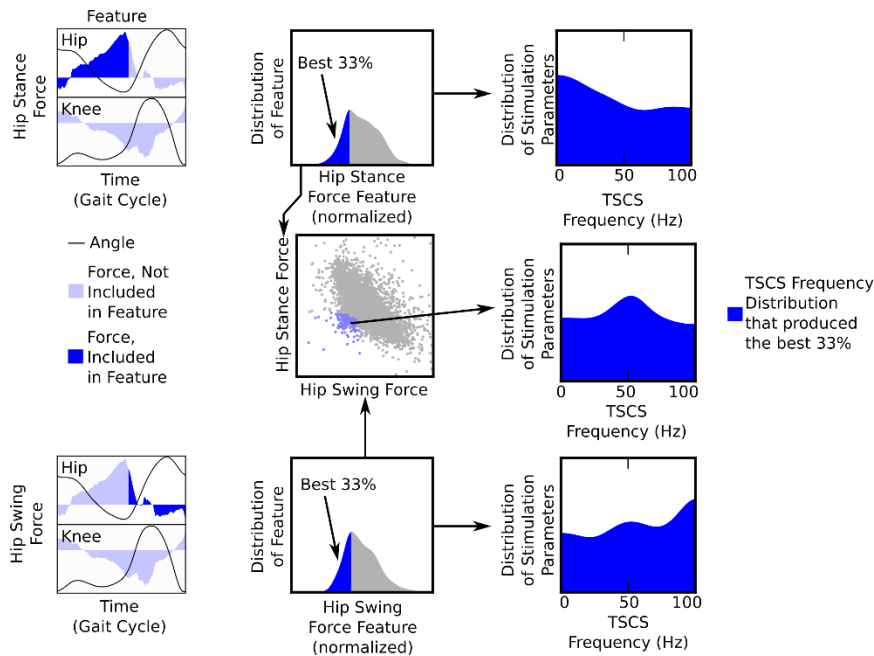


Figure 5-4. TSCS Frequency vs both hip-swing and hip-stance forces. The left column shows the two force features compared (hip-swing and hip-stance force features). The middle column on the top and bottom rows show the distribution of those two features individually. They are combined in the middle column center plot to show how the hip-stance and hip-swing forces relate to one another. In the cases where both features have good values (bottom 33%), the stimulation frequency required to produce those points are plotted on the right column.

5.3.4. Muscle Activation (EMG)

Muscle activation (via EMG) was analyzed in order to understand how muscle groups were affected by stimulation parameters. For this analysis, I identified specific times in gait when the muscle groups were active, and explored which stimulation parameters were most effective in increasing the activity during those times. This analysis was complicated by two factors: (1) surface EMG is limited in what muscles are accessible. For example, the most important hip flexors, the iliopsoas, are inaccessible to surface EMG, particularly when the subject is wearing a body weight support harness over their pelvis. (2) EMG is sensitive to nearby electrical stimulation. EMG requires a large amount of amplification (500x), and the voltages of the stimuli are frequently >1,000 times the amplitude of the EMG signals, leading to amplifier saturation. This effect is distance-dependent, such that the closer the EMG electrode is to the stimulating electrode, the larger the effect. Therefore, even among the muscles that were recorded from, some had stimulation artifacts that were too large to be reliably removed by post-processing.

For TSCS, stimulation affects flexors and extensors differently (figure 5-5). The triceps surae showed a negative relationship with frequency: higher TSCS frequencies tended to produce less triceps surae EMG. The results for tibialis anterior were mixed, and did not show a clear dependence on frequency.

For cutaneous nerve stimulation (figure 5-5), most channels demonstrated that there was more EMG when pulse trains were applied in the region between mid-stance and early swing, and higher-frequency, longer-duration pulse trains led to better muscle activation. These longer pulse-trains are much longer than pulse trains normally used

clinically. PTNS demonstrated stronger (larger peaks in the distribution) frequency-dependence and duration-dependence than CPNS did.

The cutaneous nerve stimulation sites showed that muscle groups were affected differently (figure 5-5). Hamstring muscles demonstrated the best response was from pulse trains applied in early swing. These pulse trains corresponded to stimulating during the period in which hamstring-feature was tested. Tibialis anterior and quadriceps muscles showed a similar effect from PTNS. Quadriceps showed divergent behavior between the two cutaneous nerve sites. CPNS was the most effective in generating quadriceps activation at stance-to-swing transition. PTNS and CPNS showed opposite phase distributions for quadriceps activation in early stance.

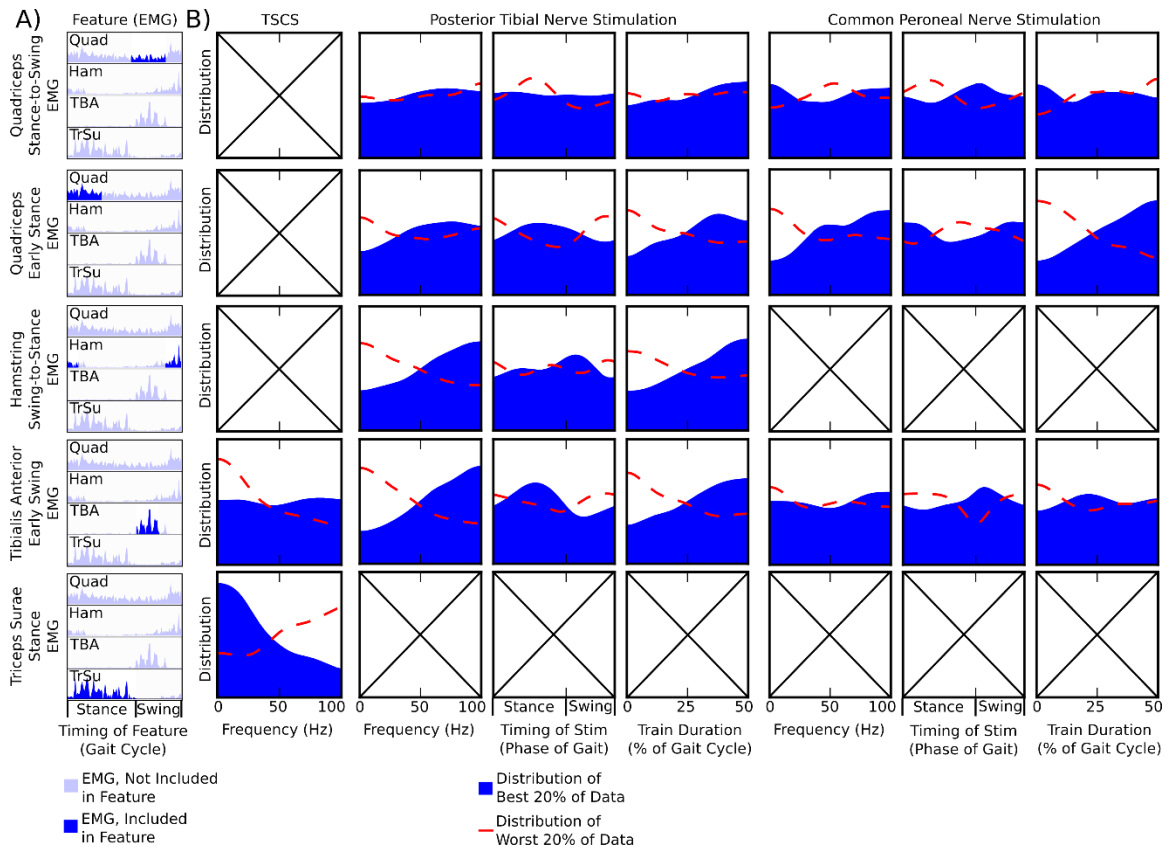


Figure 5-5: Selected EMG features. The left column (panel A) shows the feature being calculated. In each case a time-window is applied to one EMG channel. Each row shows one feature. On the feature plot, dark blue shows the portion of the EMG channel that is included in the calculation, and light blue shows the excluded EMG data. The columns to the right (panel B) show the stimulation parameter distribution of the best and worst 20% of the data with respect to each feature (row), calculated in the same manner as the force distributions in the previous figures, with one difference – EMG features are “better” when they are larger (assistive force features are “better” when they are smaller). In the columns showing the stimulation parameters, some plots were excluded due to too much stimulation artifact on those channels while recording EMG (marked with an ‘X’). These excluded channels were identified empirically. On the distribution plots, dark-blue fill shows the best 20%, and red, dotted line shows the worst 20%. As in other plots, stance was defined to be the first 60% of gait, and swing was defined to be the last 40%. EMG: Electromyography. Quad: Quadriceps muscle. Ham: Hamstring muscle. TBA: Tibialis Anterior. TrSu: Triceps Surae.

5.3.5. Assistive Force vs EMG

In this analysis, our goal was to understand whether less assistive force was associated with more or less muscle activation. This analysis is complicated by the limits of surface EMG and the difficulties of recording EMG while stimulating on a nearby location on the

body. Therefore, the analysis had to test EMG locations far from the joints at which the force was measured (e.g. hip force compared to triceps surae activation), or EMG on muscles which cross multiple joints (e.g. rectus femoris).

The results are mixed, even within each stimulation site. Lower TSCS frequencies are both associated with better (lower) forces during stance, and better (higher) triceps surae EMG during stance (data not shown). Tibialis anterior results are inconclusive. Better forces during the stance-to-swing transition were associated with lower EMG in the tibialis anterior (figure 5-6B).

For PTNS, better (more) EMG tended to increase in parallel with better (less) force across most features. In figure 5-6C, the best (lowest) 20% of force values at the stance-to-swing transition were associated with the highest tibialis anterior EMG results. The quadriceps-stance force (data not shown) followed a similar trend, i.e. the best 20% in the force feature tended to have high quadriceps EMG values during stance. The hamstring-swing-to-stance EMG-feature followed the reverse relationship.

In common peroneal nerve stimulation (data not shown), the tibialis anterior results were reversed as compared to the posterior tibial nerve stimulation data – the best 20% of force values tended to have lower EMG in the tibialis anterior. However, the results for the quadriceps-stance-EMG feature were similar to that of the posterior tibial nerve stimulation.

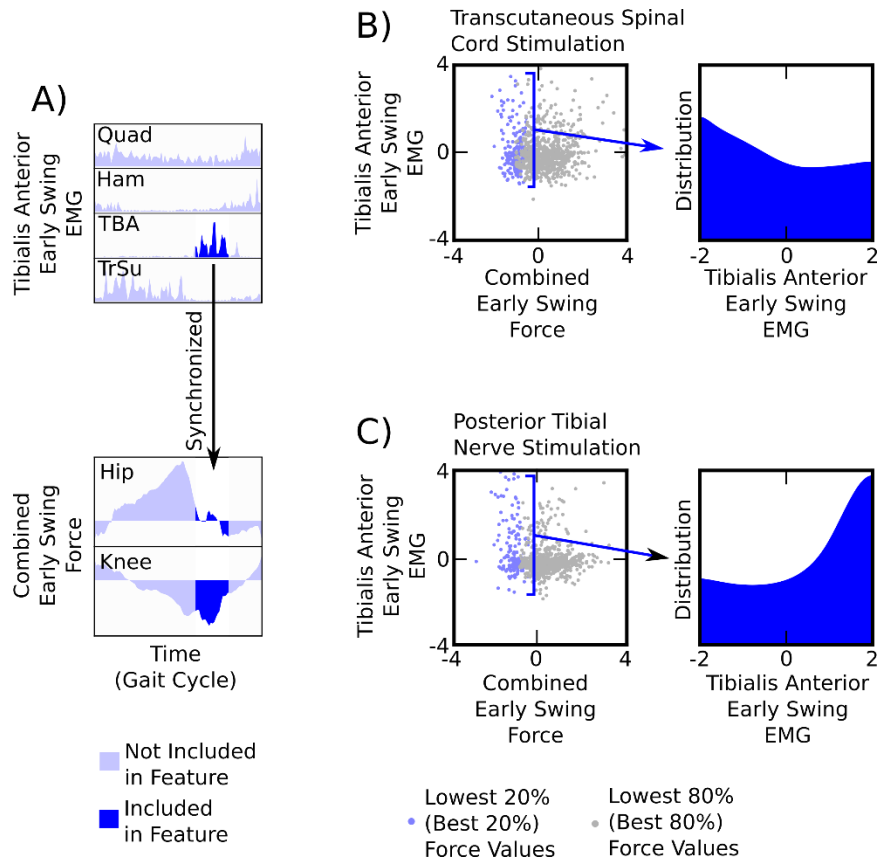


Figure 5-6. Force-to-EMG Comparison. In Panel A, the tibialis-anterior EMG-feature and combined-joint force feature are shown. The EMG feature summarizes the early swing activation of the tibialis anterior. The force was calculated from the same time-period in the gait cycle to compare whether lower force and higher EMG tended to be associated. In panel B, the results for TSCS are shown. In the first plot of panel B, the scatter plot comparing the EMG metric (TBA early-swing EMG) to the force metric (both joints, early-swing force) is shown. The 20% best (lowest) points in the force metric are extracted (blue) from the rest of the dataset (grey), and the resulting distribution of those 20%-best points is shown in the 2nd plot in terms of the EMG metric. Panel C shows the results from posterior tibial nerve stimulation. TBA: Tibialis Anterior.

5.3.6. History Dependence

In this study, history had as large an effect on the features as the current stimulation pattern. In this context, history refers to the number of steps the subject has been in our experimental setup, subject to sensory stimulation and walking in the Lokomat. In order to explore this phenomenon I analyzed history as though it were a stimulation pattern.

In figure 5-7, the best feature output was analyzed to how many steps the subject had been walking in the experiment. I compared history to the stimulation parameters for both posterior tibial nerve stimulation and TSCS, using hip-swing force as an example feature. In TSCS, the “good” hip-swing force-values (or more accurately, the concentration of the lowest hip-swing force-values) showed a steep decline with respect to history. As more experimental time accrued, it was less and less likely that the hip-swing force would be “good”. In TSCS higher frequencies tended to produce better hip-swing forces. However, the differences between frequency and history show that (for at least this feature) history-as-a-stimulation-parameter was more salient than frequency.

The results from PTNS showed that history was also important. The distribution of “good” hip-swing forces demonstrated differences with respect to history, but phase was the most salient parameter (larger peaks and valleys). For the hip-swing force feature, history was approximately as predictive as frequency or duration.

Other features (not shown) demonstrated the importance of history. In some features, history was the most important parameter, and in others it was not. However, in all, history demonstrated that the response to stimulation was time-variant.

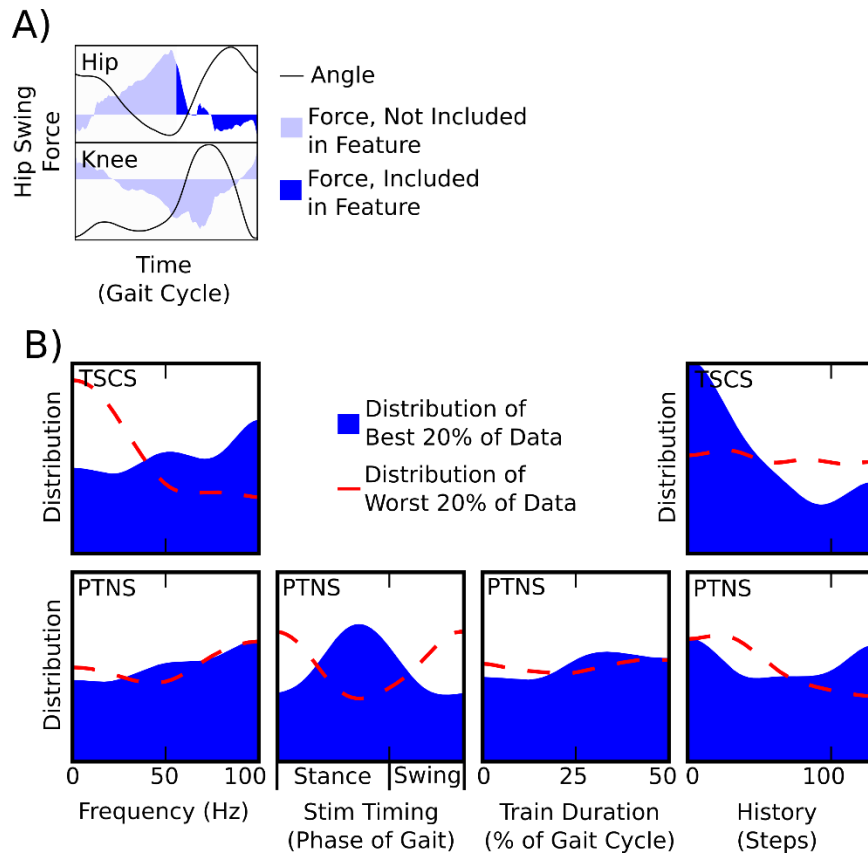


Figure 5-7. Comparison of Stimulation parameters to History as a parameter. Panel A, at the top, shows the feature I compared (hip-swing force). In Panel B (the bottom two rows), TSCS (row 1) and posterior tibial nerve stimulation (row 2) are shown, along with the distribution of the best feature values with respect to experimental history. History, in the far right column of Panel B, was measured by the number of steps the subject had already performed.

5.3.7. Individual Differences

The previous analyses focused on the entire population as a single group. I also observed individual differences, and this section demonstrates specific examples of how subjects diverge from the responses of the population as a whole. Our results show that there is a high degree of individuality in the data. I present three examples of subjects, demonstrating how they differ from one another, and how they differ from the population-as-a-whole.

In figure 5-8, the population distribution showed a negative relationship between TSCS frequency and hip-stance force (i.e. better hip-stance forces tended to be lower frequency). Subject 12 showed a divergent response, with “good” values at both extremes, while subjects 5 and 11 showed good responses from frequencies between 25-70 Hz, and subject 11 showed a particularly salient peak at approximately 40 Hz.

Hip-swing and knee-swing forces showed an opposite relationship as hip-stance in the population as a whole, and in the individuals (figure 5-8). The population distribution showed a positive relationship with higher frequencies. In the results from Subject 5 and 11, mid-frequencies produced bad feature-values, but in the results from subject 12, low-frequencies produced bad feature-values.

The EMG features for individuals also diverged from the population’s distribution (figure 5-8). Triceps Surae EMG was better at lower frequencies, while tibialis anterior showed mixed results in the population. However, the individual subjects showed more salient results. Subject 5 showed that both EMG features were at their best at TSCS frequencies of 35-50 Hz. Subjects 11 and 12, however, showed divergent patterns between better (higher) EMG output at the triceps surae and tibialis anterior.

Although this analysis was focused on TSCS, there are similar types of divergences in other stimulation sites (data not shown). PTNS demonstrates that stimulation phase is good for most features when applied in the region from mid-stance to the stance-to-swing transition. However, there were subjects whose “best” phase stimulate at was either earlier or later in this time period.

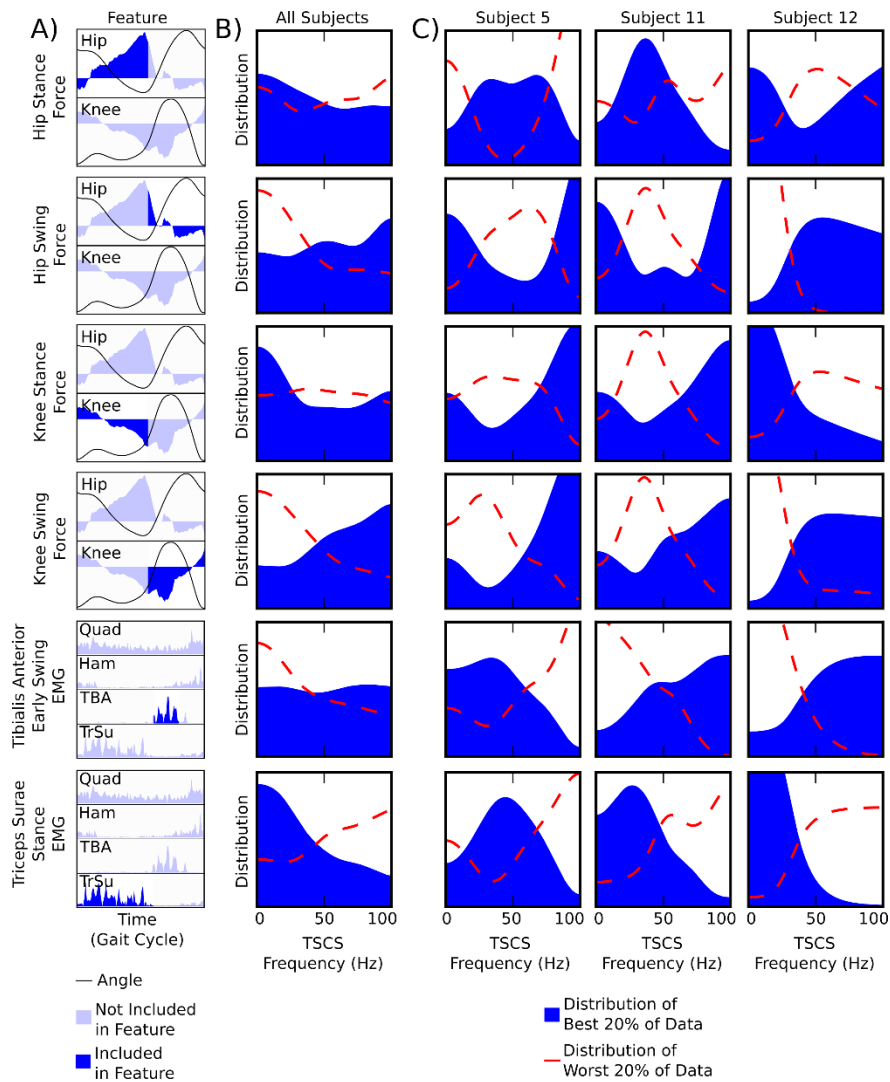


Figure 5-8. Individual Variation. Panel A (left column) shows the features. The top 4 features are force features (hip-stance, hip-swing, knee-stance, and knee-swing), and the bottom 2 features are EMG features (tibialis-anterior-early-swing and triceps-surae-stance). Panel B (2nd column) shows the distribution of TSCS frequency from the subject population as whole. Panel C (3rd–5th columns) show three example of individual distribution of TSCS frequency that differ from the population distribution.

5.3.8. Onset Response

In this analysis, our objective was to characterize the data from the perspective of the evoked responses (i.e. with respect to the timing of the pulse train), instead of from the perspective of steps (how the previous analyses were performed). This perspective is relevant, as different phases of the evoked response have been associated with the

central pattern generator, nervous system development, and recovery from spinal cord injury. Therefore, I identified and characterized the equivalent P1 and P2 phases of the evoked responses in the tibialis anterior.

The results (figure 5-9) demonstrate that there is an onset response and a termination response. The onset response peaks at approximately 57 ms after the pulse train begins. The termination response peaks at approximately 64 ms after the pulse train ends.

The onset response appears to be frequency-independent, but the termination response appears to be frequency-dependent. The onset response is visible at frequencies in the 0–10 Hz range, but the termination response is not visible at those frequency ranges.

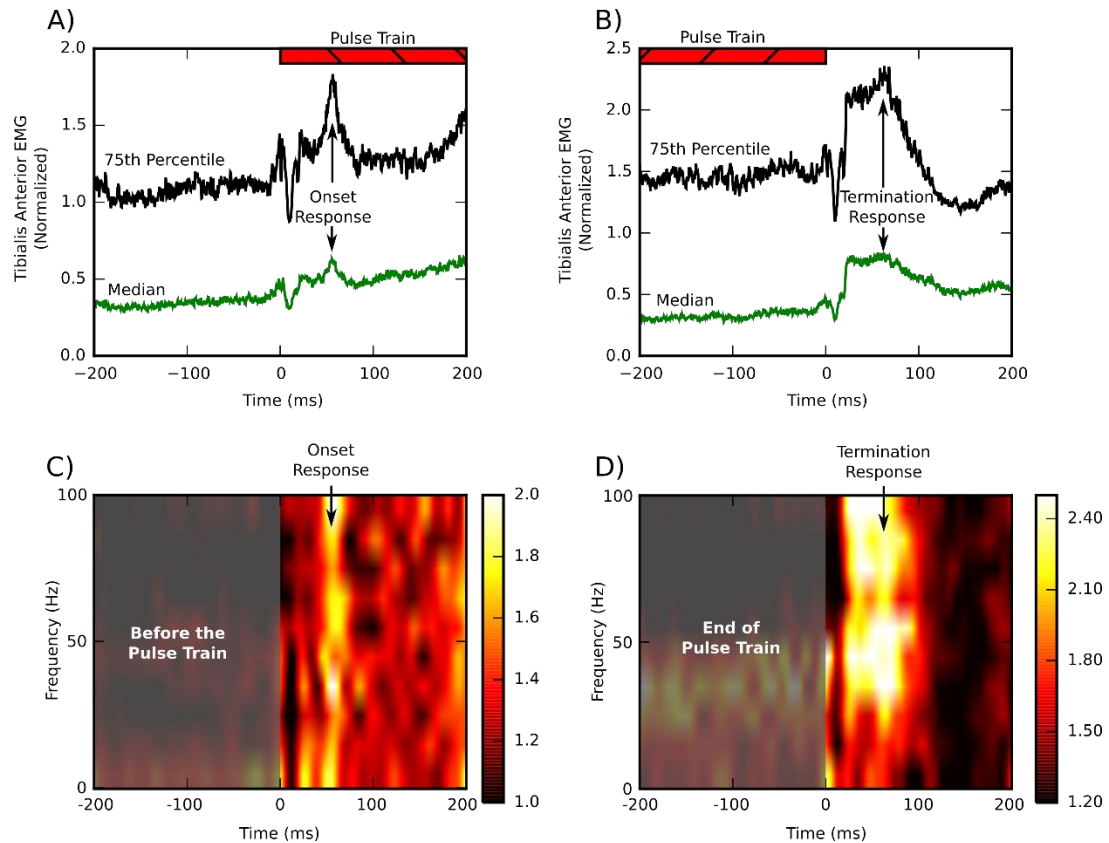


Figure 5-9. Onset and termination response of tibialis anterior posterior EMG evoked from PTNS (pulse trains). In panel A, it shows the onset response. The red, striped bar shows the timing of the beginning of the pulse train. The median and 75th percentile are shown. Panel B shows the termination response. The red, striped bar in panel B shows the end of the pulse train. Panel C shows the onset response compared to the frequency. 10 Hz bands of frequency were collected, forming 10 bands between 0 and 100 Hz (e.g. a band including 20-30 Hz stimulation). Panel D shows the termination response versus stimulation frequency. All waveforms were normalized prior to combining into percentiles.

5.4. Discussion

In this study, 10,000 random-parameter stimulated-steps were collected across three stimulation sites, two stimulation protocols, and twelve subjects. I analyzed how stimulation parameters affected specific features in gait, leveraging both EMG- and force-based analysis. The effect of history, and how different individuals were within the population proved to be significant factors in the results. Lastly, I characterized an onset-response to cutaneous nerve stimulation.

5.4.1. Reductions in Assistive Force

Our analysis included exploring how stimulation could reduce the required assistive force from a robotic orthosis (the Lokomat). This assistive force maintained the walking pattern, so less assistive force meant that the subject was either doing more correct muscle activation or less incorrect muscle activation.

The results show that TSCS frequency has a divergent effect on forces. Higher frequencies produce better swing phases, but lower frequencies produce better stance phases (figure 5-2). To get both good stance and swing forces, a frequency in the middle (approximately 50 Hz) produced the best response (figure 5-5) across both gait phases. When TSCS was compared to short-term force-features (figure 5-3), there were three distinct regions in gait. Early-stance to the stance-to-swing transition demonstrated the lower-frequency-is-better phenomenon. The early-swing to mid-swing demonstrated a higher-frequency-is-better phenomenon. Late-swing to swing-to-stance demonstrated mixed results, with weaker tendencies than the other time-periods. At first, these results may seem to suggest stimulating differently during swing and stance, but such an approach is complicated by the bilaterality of TSCS. If one leg is in the stance phase, the other is in the swing phase. Therefore, a modification to TSCS that allowed for unilateral stimulation may be most effective.

Cutaneous nerve stimulation showed a less divergent distribution. All force-features demonstrated peaks when the stimulation phase was between late stance and early swing. Of particular note was that stance forces could be improved by stimulation during the stance-to-swing transition. On the surface, this result would seem to violate causality (the stance-to-swing transition is immediately after the stance phase), but it is important to note that the contralateral leg was also being stimulated, and that stimulation was 180

degrees out of phase with the ipsilateral leg. The stimulation from contralateral leg would have reached the nervous system prior to the stance phase improvement. This phenomenon was likely due to the contralateral extension reflex. Swing forces could be assisted most by late stance pulse-trains.

The effects of stimulation frequency and pulse train duration on cutaneous nerve stimulation has not been thoroughly studied. Our results demonstrate that higher frequencies and longer pulse trains tended to produce better (lower) forces. Our experiments did not test pulse trains longer than 50% of the gait cycle, and there may be diminishing returns above 50% as the pulse-train-stimulation approaches tonic stimulation. Most previous studies are already using high-frequency stimulation, but long-duration stimulation is not widely implemented.

The short-term force features demonstrated that the best stimulation phase showed two periods of peaks delaying with their respective short-term force features (figure 5-3): during the stance phase and the early swing phase. The stance-phase delay phenomenon was likely due to the contralateral extension reflex, and the swing-phase delay was likely due to the ipsilateral flexion reflex.

5.4.2. Muscle Activation and Force

In order to better understand how stimulation affected muscle activation I tested EMG-features for all three stimulation sites. The EMG waveforms are much less consistent between people than the force waveforms were, and other researchers have found substantial variability in what constitutes “normal” EMG activation (Arsenault, Winter, and Marteniuk 1986, Winter and Yack 1987).

Muscle-activation features demonstrated the same types of divergences as the force features (figure 5-4). For TSCS, triceps surae (an extensor active during stance)

produced better EMG with low stimulation frequencies, similar to the stance force-features. The only flexor tested for TSCS (tibialis anterior) did not show clear results – TSCS may be ineffective for the tibialis anterior muscle.

Muscle activation from cutaneous nerve stimulation demonstrated similar results as force features (figure 5-4). The best phase to stimulate at was late-stance to early swing in almost all the features across both sites. The quadriceps-early-stance feature from CPNS was an exception to this rule, but this exception may be explained by two issues: (1) Quadriceps-muscle activation has been found previously to change during walking in the Lokomat – it is much more active at the stance-to-swing transition than in treadmill walking (Hidler and Wall 2005). The same stance-to-swing peak was evident in the quadriceps muscle activity in our data. (2) The CPNS proved to be more effective in generating a quadriceps-stance-to-swing feature than a quadriceps-early-stance feature. This difference could be due to the effects of gait phase dependence on the changes to walking induced by the Lokomat.

I also compared whether EMG activation was associated with force reduction, or was the relationship reversed. Our results were mixed. Some features showed that better (less) force led to better (more) EMG, and some showed the better (less) force led to worse (less) EMG. The triceps-surae-stance EMG evoked from TSCS is an example where better EMG and better force are seen together. Tibialis-anterior-stance-to-swing EMG evoked from TSCS is an example where better EMG and better force were divergent. However, the same feature and force combination (tibialis-anterior-stance-to-swing EMG and force at the same time) proved to be associated with PTNS.

Conclusions are difficult to draw from these mixed results, except that perhaps more EMG is not always better EMG.

5.4.3. Individuality and History

The response to stimulation varies both across individuals within a population, and within the same individual during the course of an experimental session.

Our results show that there is a large diversity within the population as to how stimulation affects specific features of gait. Some subjects had opposite responses to TSCS frequencies. Individuality in clinical populations is known, and there have been a number of sources of sensorimotor-response variability identified. In neurologically normal individuals, there is variability in reflex responses based on age (Evans, Harrison, and Stephens 1990, Gibbs et al. 1999), level of activity (Loeb 1993), and athletic history (Gruber et al. 2007, Wolpaw and Tennissen 2001, Maffioletti et al. 2001). A number of studies have found a further divergence after SCI in the sensorimotor responses, including SCI level (Poirrier et al. 2004, Dietz et al. 1999), severity (Rossignol et al. 1996), and time since injury (Hiersemenzel, Curt, and Dietz 2000). There have also been uncharacterized sources of variability. In one human SCI study, no two individuals showed similar patterns of coordination after an intervention (Field-Fote and Tepavac 2002), and in another human SCI study, individuals did not use consistent strategies to engage in an incline walking task (Leroux, Fung, and Barbeau 1999).

The implications of the variability in the SCI population complicates the study of sensorimotor responses. Fortunately, the population-as-a-whole still demonstrated trends in the data.

Within a given experimental session, the number of steps already completed affects the sensorimotor responses as much as the stimulation parameters do. This time-variant effect was observable in as few as 130 steps. Some specific types of time-variance have

been characterized previously, including habituation (Harris 1943) and EMG exhaustion (Dietz and Müller 2004). However, in some cases our results showed potentiation in the results, or biphasic behaviors which cannot be explained with habituation or habituation-like phenomena. This time-variance has many implications, potentially affecting how experiments should be designed for SCI subjects, and how analysis should be performed.

These sources of variability give strong incentive for the use of optimization approaches to individualize an intervention, and algorithms that are robust against time-variance.

5.4.4. Onset-Response and Termination-Response

In our results, pulse trains produce both an onset response and a termination response. Although long-duration pulse-trains have not been well-studied, two other papers have recorded onset-responses after the beginning of long pulse trains (Roby-Brami and Bussel 1987, Pearson and Collins 1993).

Classically, cutaneous nerve stimulation responses are divided into P1, P2, and (sometimes) P3 (Roby-Brami and Bussel 1987) phases, all occurring after the stimulation has concluded. Previous studies have found differences between the P1 and P2 phases. The P1 and P2 phases are modulated differently during gait (Baken, Dietz, and Duysens 2005, Rossignol, Dubuc, and Gossard 2006). The P1 phase is thought to be purely spinal (Baken, Dietz, and Duysens 2005), and it is difficult to elicit in human adults (Yang and Stein 1990, Van Wezel, Ottenhoff, and Duysens 1997, Baken, Dietz, and Duysens 2005). P2 is the most reproducible in human adults (Yang and Stein 1990, Baken, Dietz, and Duysens 2005), and equivalent of P2 in the hands (E2) has been linked to the development of skilled hand movements during adolescence (Evans,

Harrison, and Stephens 1990). P3 is more likely to be produced by higher-stimulation strength, longer-duration pulse trains (Roby-Brami and Bussel 1987).

Most previous studies have used very short pulse trains, such that it was impossible to differentiate the onset-response from the termination-response. I propose that the P1 response is actually an onset-of-stimulation response. The latency of our onset-response (~57 ms after the first stimulus) is consistent with the latency of the P1 response, as observed by other researchers (~50-60 ms: (Roby-Brami and Bussel 1987), ~70 ms latency: (Zehr and Stein 1999), <65 ms after the first stimulus: (Brooke et al. 1997), ~50-60 ms: (Shahani and Young 1971), ~50-60 ms: (Pedersen 1954)).

Researchers have observed that the P1 response was inconsistent in human adults (Yang and Stein 1990, Van Wezel, Ottenhoff, and Duysens 1997, Baken, Dietz, and Duysens 2005), and this inconsistency could be due to an inhibitory phase related to the termination-response overlapping with an excitatory phase related to the onset-response.

I also propose that the P2 response is a termination-of-stimulation response. The latencies of P2 in the literature are less consistent, but tend include 75-100 ms post-stimulus (75-100 ms: (Roby-Brami and Bussel 1987), 70-120 ms: (Brooke et al. 1997), ~80 ms (Tax, Van Wezel, and Dietz 1995), 70-80 ms (Duysens et al. 1996), and 80-85 ms (Van Wezel, Ottenhoff, and Duysens 1997)). However, those latency measurements are measured from the beginning of a pulse train that often lasts 10-30 ms. A time-period of 75-100 ms after the first pulse would be ~60-85 ms measured from the end of the pulse train, in line with our measured latency of ~64 ms after the termination. The termination response can also be seen in Roby-Brami and Bussel (1987), but that study did not explicitly compare their responses to the P1 or P2.

In our data, the termination-response was frequency-dependent, but the onset-response was not. As far as I found, no previous study has explored the multiphasic response to cutaneous nerve stimulation in terms of frequency-dependence. This difference will need to be further studied, but it suggests an interesting modulation of frequency-dependent processes in the nervous system.

5.4.5. Conclusions

I have demonstrated the optimal stimulation parameters for multiple features of gait, as evoked from multiple stimulation sites, across multiple subjects. I discovered that TSCS affects stance and swing differently, giving support for the development of bilateral TSCS. Through analysis of pulse trains, I found evidence that pulse trains during late-stance and early swing can assist both legs, through ipsilateral and contralateral reflex pathways. I identified an unappreciated source of intrasubject variability (history) and compared its predictive power to that of the actual stimulation parameters. Lastly, I proposed (and gave evidence for) and re-categorization of the classic model of evoked responses from cutaneous nerve stimulation may actually be driven by the onset and termination of the pulse train separately, and that those responses differ in their frequency-dependence.

CHAPTER 6: EFFICIENT GLOBAL OPTIMIZATION OF TIME-VARYING AND NOISY SYSTEMS

In this chapter, I present a novel approach for globally optimizing noisy, time-varying, black-box systems. This algorithm optimized the sensory stimulation to assist walking, as discussed in the previous chapters. It is based on modification of a previously existing algorithm, called efficient global optimization (EGO), which itself was derived from Kriging and Gaussian process models. I combined previously published modifications to EGO for noise and time-variance into a single algorithm, referred to as time-varying, noisy efficient global optimization (TVN-EGO). TVN-EGO was empirically validated with a set of test functions. The tests included nine test functions (representing problems with 1–4 dimensions) with multiple noise-levels and multiple time-variance-levels to identify the effective limits for time-varying, noisy optimization. As part of this validation, I proposed a new measurement of time-variance, based on the L2-norm of the derivative with respect to time. This time-variance metric was normalized by the spatial standard deviation (SSD). One-dimensional optimization proved to be successful when the time-variance was <5% of SSD, and when the noise standard deviation was <50% of the SSD. In higher dimensions, the performance was lower, but the algorithm was still successful in a 4 dimensional test function with noise of 5% of the SSD and time-variance of 0.5% of the SSD. In all cases, TVN-EGO outperformed the two local optimization algorithm by at least an order of magnitude (10x) in the systems with noise, time-variance, and both noise and time-variance.

The chapter is organized into four sections: the introduction, an overview of the approach, simulated experiments, and a discussion. The “overview of the approach” details how the algorithms were formulated, and the modifications that were implemented to previous versions of EGO. Then, the algorithm is compared to two other

standard optimization algorithms across a range of test functions to determine the limits of effectiveness.

6.1. Introduction

There are many costly-to-test systems with noise, time-variance, and unknown dynamics, and optimization in these systems could improve outcomes. Implanted electrodes for stimulating or recording the nervous system have to deal with the increasing presence of scar tissue and possibly changing neural connections nearby (Polikov, Tresco, and Reichert 2005). Agriculture must deal with difficult-to-predict climate changes that occur on a multi-year scale (Kurukulasuriya and Rosenthal 2013). In recent years, the macroeconomic environment for business regularly undergoes difficult-to-predict structural changes (D'Agostino, Gambetti, and Giannone 2013). These changes are occurring at timescales relevant to the speed at which new hypotheses can be tested against those systems.

There are many algorithms that exist for optimizing time-varying, noisy systems, but only when the dynamics are known. These approaches are generally classified as dynamic optimization or optimal control (Chiang 2000). The systems in these problems are sometimes colloquially called “white-box” or “gray-box” systems, to differentiate them from systems where the dynamics (and gradient) are unknown (“black-box” systems). In black-box optimization, algorithms based on dynamic optimization cannot be easily applied (Conn, Scheinberg, and Vicente 2009).

This author could find no other optimization algorithms explicitly designed for black-box, time-varying, noisy systems. There are relatively few optimization algorithms designed for just time-varying, black-box systems, even without noise, costly-to-test, or global-optimization requirements.

Particle swarm optimization (PSO) is a global optimization algorithm (Kennedy and Eberhart 1995, Poli, Kennedy, and Blackwell 2007), based on accelerating particles towards the measured optimum. The “vanilla” version of PSO is reasonably robust against noise, despite a lack of being explicitly designed for that purpose (Parsopoulos and Vrahatis 2001), and a version of PSO has been developed to explicitly handle noise, with better performance as a result (Pugh, Zhang, and Martinoli 2005). However, it appears that no variants of PSO have been designed for time-varying optimization. In addition, PSO is not frugal with function evaluations, and will often test 10-100 points per iteration (Kennedy 2010), which may be very expensive in some systems.

Genetic algorithms (Golberg 1989) can be robust against noise, but they assume that multiple identical scenarios are available (to test an entire generation at the same time), which may not be possible in a real-life, time-varying system. Good results from differential evolution (Storn and Price 1997) and simulated annealing (Dekkers and Aarts 1991) generally require thousands of function evaluations – even for relatively low dimensional systems. None of these approaches are easily used in time-varying problems.

Coordinate search (or pattern search) algorithms fulfill many of the requirements – they are designed for black-box systems, and algorithms exist for time-varying optimization. Although coordinate search algorithms were not generally designed for noisy measurements, their large step size makes them robust against small levels of noise. A version of the popular downhill-simplex algorithm (Nelder and Mead 1965) has been created to address time-varying systems (Xiong and Jutan 2003) through resampling and eliminating the shrinkage phase. Coordinate searches tend to be frugal with respect to function evaluations as well (Conn, Scheinberg, and Vicente 2009), but they are not global optimizers.

Stochastic approximation (SA) attempts to directly estimate the derivatives of noisy systems, and perform a gradient descent (Spall 2005). Early versions used a finite difference approach, requiring many function evaluations per iteration (Kiefer and Wolfowitz 1952). A later version, called simultaneous-perturbation stochastic approximation (SPSA), estimates the gradient with only two samples per iteration, regardless of the dimensionality (Spall 1992), and achieves this goal with very little loss of performance (Spall 2005). SPSA can handle noise, and is frugal with respect to sampling the space, but is not a global optimizer, nor can it handle time-varying systems. However, SPSA can be a global optimizer under certain conditions (Maryak and Chin 2001). For the purposes of the comparisons this chapter, I modified SPSA to improve its robustness to time-variance by eliminating the decay in the coefficients, similar to the modification made to downhill simplex to create dynamic simplex (Xiong and Jutan 2003).

Response-surface-based optimization algorithms are designed for both global optimization and frugality in the use of experiments. These approaches build a model of the response, and optimize within that model to guide sampling. They "often requires the fewest function evaluations of all competing methods." (Jones, Schonlau, and Welch 1998). Efficient global optimization (EGO) is a popular example of response-surface-based optimization algorithms. EGO relies on a stochastic model, specifically a Gaussian process model, and it has recently been proven to converge (Vazquez and Bect 2010), although the proof did not guarantee a specific rate of convergence. Modifications have been made to EGO allowing it be robust against noise (Huang et al. 2006, Forrester, Keane, and Bressloff 2006), and more recently, modifications have also been proposed to make it robust against time-variance (Desautels 2014, Morales-Enciso and Branke 2015), through the addition of time as an uncontrolled input variable.

I propose to combine these modifications to make a global optimization algorithm that is robust against both noise and time-variance: time-varying, noisy efficient global optimization (TVN-EGO).

Optimization algorithms can be robust against some quantity of noise and some quantity of time-variance, but the limits as to how much noise or time-variance they can handle are unknown. It is also unknown how different algorithms compare in their robustness to noise and time-variance. I set out to answer these questions with a series of simulations testing multiple levels of noise, multiple levels of time-variance, and different test functions. I compared TVN-EGO with Dynamic Simplex and SPSA to determine which algorithm was most robust to noise and time-variance.

6.2. Overview of Approach

This section presents the measurement of noise and time-variance, followed by a brief overview of time-varying efficient global optimization (TVN-EGO), and an illustrated example of a few iterations.

6.2.1. Measuring Time-variance

There is no standard way to measure how an entire input-output mapping function changes in time. For time-series models there are many approaches, such as continuity (i.e. how many orders of derivatives are well-defined, for example: C^1 , C^∞), or tests to determine if a time series is stationary or not (Dickey and Fuller 1979, Priestley 1981). In the machine learning literature, one approach has been to put upper limits on total variance of a set of values to delineate problem classes (Hazan and Kale 2009).

Norms of the derivative with respect to time appear to be the most intuitive metric.

Mathematical norms are functions that assign a distance to a multidimensional vector,

for example: Cartesian distance (as the bird flies) and “Manhattan” distance (driving between two points on a grid). A time-variance metric based on the L^∞ norm (i.e. $\max_i(x_i)$) would identify the largest single change between two iterations (similar to the approach in Hazan and Kale (2009)). A metric based on the $L1$ norm (i.e. $\sum |x_i|$, “Manhattan” distance) would give any deviation (even small ones) an equal weighting. I chose to use the $L2$ norm (i.e. $(\sum x_i^2)^{1/2}$, Cartesian distance), for many of the reasons that make it ideal for other applications. It is intuitive that larger “jumps” between iterations should be weighted more heavily, and very small changes should not matter as much.

An ideal property for the time-variance metric is scale-and-shift independence. The optimization algorithm modifies the data by shifting the mean or scaling the values as it is convenient, and the time-variance metric should remain constant through these shifts. Therefore, a normalization factor was needed based on the spatial variance of the function. There were two obvious choices for normalization of the time-variance metric: the range of the function, and the standard deviation of the function. As the standard deviation is similar to the $L2$ norm, and it includes more information of the input space than the range does.

$$TV_{unnorm} = \sqrt{\frac{1}{N_t - 1} \sum_{t_j} \frac{1}{N_x} \sum_{x_i} (x_i(t_j) - x_i(t_j - 1))^2} \quad (16)$$

$$\mu(t_j) = \frac{1}{N_x} \sum_{x_i} x_i(t_j) \quad (17)$$

$$k_{norm} = \sqrt{\frac{1}{N_t} \sum_{t_j} \frac{1}{N_x} \sum_{x_i} (x_i(t_j) - \mu(t_j))^2} \quad (18)$$

$$TV_{norm} = \frac{TV_{unnorm}}{k_{norm}} \quad (19)$$

In the above equations, TV_{unnorm} is the unnormalized time-variance metric, TV_{norm} is the normalized time-variance metric, k_{norm} is the constant of normalization (a standard-deviation-based metric of spatial variance), N_x is the number of x-points used in the summation, N_t is the number of iterations used in the summation, t_j refers to the j th iteration, x_i refers to the i th x-point, and is $\mu(t_j)$ the mean at time t_j . Due to the normalization by spatial variance, TV_{norm} is robust again scaling or shifting of the space.

6.2.2. Measuring Noise

In the simulated tests, Gaussian noise was added to each measurement. The variance of that noise is σ_{unnorm}^2 and the standard deviation is σ_{unnorm} . For the purposes of comparison, the noise measure was normalized with respect to the spatial standard deviation. This normalized noise measure is robust against scaling or shifting. Also, this measure is related to signal-to-noise ratio (SNR) – it the reciprocal of the square root of SNR.

$$\sigma_{norm} = \frac{\sigma_{unnorm}}{k_{norm}} \quad (20)$$

6.2.3. Time-varying, Noisy, Efficient Global Optimization (TVN-EGO)

Time-varying, noisy, efficient global optimization is a modification of a previously published algorithm, Efficient Global Optimization, or EGO (Jones, Schonlau, and Welch 1998). The modifications allow the algorithm to effectively optimize time-varying, noisy systems. The general approach is to form a probabilistic model of the input-output space using estimated correlations between data points. Then, using that model, the algorithm

will locate the point with the maximum likelihood of improving on the current best, among all potential points in the input space. TVN-EGO also includes a parameter to make it robust against measurement noise, utilizing the same matrix regularization approach as sequential kriging optimization (SKO) (Huang et al. 2006). In order to make the algorithm allow for time-variance, time was set as a uncontrolled input parameter (Morales-Enciso and Branke 2015).

In order to assist the reader understand this algorithm, it is presented here briefly in pseudocode:

```
data = Initialize_Space() // Sample a number of points
while (!finished) { // iterate
    model = Create_Model (data)
    next_point = Select_Next_Point (model)
    new_value = Test_against_system(next_point)
    data.append(next_point, new_value)
    finished = Check_if_finished()
}
```

6.2.3.1. Initialization

To start, $11 * x_{dim}$ open-loop x-points (inputs) are selected. For each 11 open-loop points, 9 are unique, and 2 are copies. This choice reflects the same approach used in SKO (Huang et al. 2006), but with two copies, instead of one. Two copies allows an initial estimate of both noise and time-variance in the open-loop period. The unique points are selected using a Latin hypercube to sparsely (and uniformly) fill the input space, with the duplicates chosen randomly. In practice, I created 200 potential Latin hypercubes, and chose the one with largest minimum-distance between points (i.e. “maximin”).

The open-loop test points are produced in a range of 0 to 1, and an affine mapping will scale and shift them to whatever range is required for the actual system. Then these

open-loop x-points were tested against the system one-by-one, producing y-points (outputs). Once the initial sparse sampling is complete, the algorithm begins iterations.

Table 6-1. Summary of variables used in TVN-EGO. Each variable is shown with a brief description.

Variables	Description
g	Ratio of accounted-for variance to all variance. $0.5 \leq g \leq 0.95$
θ_k	Correlation constant in the along the k^{th} input dimension. $0 \leq \theta_k \leq 10$
p_k	2 nd Correlation parameter along the k^{th} input dimension. $1 \leq p_k \leq 2$
N_x	Number of iterations
\mathbf{x}_i	Vector of input variables for the i^{th} iteration, both controlled and uncontrolled input variables
y_i	Output for the i^{th} iteration
\mathbf{y}	Vector of y_i 's
d	Distance in the input space
$r(\mathbf{x}_i, \mathbf{x}_j)$	Correlation between \mathbf{x}_i and \mathbf{x}_j
$\mathbf{r}(\mathbf{x})$	Vector of correlations between \mathbf{x} and all previous inputs
\mathbf{R}	Correlation Matrix
$\mathbf{1}$	Vector of ones
$\hat{\mu}$	Estimate of the mean
$\hat{\sigma}^2$	Estimate of the variance
$\hat{\sigma}_Z^2$	Estimate of accounted-for variance (from model)
$\hat{\sigma}_E^2$	Estimate of variance from noise
$p(\mathbf{y} \mathbf{R})$	Probability of the model given \mathbf{R}
$EI(\mathbf{x}_p)$	Expected Improvement for point \mathbf{x}_p
$\hat{Y}(\mathbf{x}_p)$	Estimate of system output from model at point \mathbf{x}_p
\mathbf{x}^{**}	Effective best previous solution (from historical data).
$s^2(\mathbf{x})$	Mean squared error of prediction at \mathbf{x}
$AEI(\mathbf{x}_p)$	Augmented Expected Improvement for point \mathbf{x}_p

6.2.3.2. Each Iteration: Model building

The first step in each iteration is to build a model of the input-to-output mapping from the data collected up to this point. As with EGO, TVN-EGO starts with a Design-and-Analysis-of-Computer-Experiments (DACE) stochastic model (Sacks et al. 1989) of the space, constructed using the distance (equation 21) and correlation between the data points. I augmented the classic model with time as an additional uncontrolled variable. It is scaled to a range of [0, 1] in the same manner as the other inputs.

$$d(x_i, x_j) = \sum_{k=1}^{x_{dim}} \theta_k |x_i(k) - x_j(k)|^{p_{\theta_k}} \quad (21)$$

$$x_{i,time} = \frac{time_i}{range(time)} \quad (22)$$

Theoretically, θ_k should be bounded only to $[0, \infty)$, but I found better performance by bounding θ_k to $[0.01, 10]$, as it reduced the chances of overfitting. In EGO (Jones, Schonlau, and Welch 1998) and SKO (Huang et al. 2006), p_{θ_k} was bounded either to $[0, 2]$ or just 2. The best results were found by bounding p_{θ_k} to $[1, 2]$.

Using the distance function defined above, a correlational matrix (\mathbf{R}) was created, defining the correlation between every data point tested thus far. Each element in the matrix is defined by the following equation (23):

$$r(x_i, x_j) = \begin{cases} \mathbf{1}, & i = j \\ g e^{-d(x_i, x_j)}, & i \neq j \end{cases} \quad (23)$$

This correlational matrix sets the diagonal to 1's, while the off diagonal elements are symmetrically set in the same manner as in SKO (Huang et al. 2006), with g acting to model the noise in the system and regularize the matrix.

In preparation to compute the likelihood of the model, the mean (equation 24) and variance (equation 25) were calculated using a generalized least square approach:

$$\hat{\mu} = \frac{1'R^{-1}y}{1'R^{-1}\mathbf{1}} \quad (24)$$

$$\hat{\sigma}^2 = \frac{(y - 1\hat{\mu})'R^{-1}(y - 1\hat{\mu})}{x_{len}} \quad (25)$$

These estimates are used to create a model likelihood metric L_M (equation 26), which is proportional to the actual model likelihood $p(y|\mathbf{R})$. The model parameters are adjusted in order to maximize L_M , similar to the approach was proposed in the original Design-and-Analysis-of-Computer-Experiments paper (Sacks et al. 1989).

$$\mathbf{p}(y|\mathbf{R}) \propto L_M = \frac{1}{\det(\mathbf{R})^{\frac{1}{N_x}} \hat{\sigma}^2} \quad (26)$$

A multi-start local optimization routine was used to maximize the model likelihood metric, but other approaches should work equally well. In the maximization of the model likelihood, some potential \mathbf{R} matrices will be poorly conditioned. These potential models were rejected by giving them a $(-\infty)$ for their model likelihood.

6.2.3.3. Each Iteration: Select Next Point

After the model was created, the algorithm selected the next point to test against the system. There are multiple ways of approaching this selection process, including selecting areas with maximal variance in the estimate (Sasena, Papalambros, and Goovaerts 2002), minimizing cumulative regret (Srinivas et al. 2009), following a “knowledge” gradient (Ryzhov, Powell, and Frazier 2012), or maximizing the expected improvement (Jones, Schonlau, and Welch 1998). I chose to follow SKO’s approach (Huang et al. 2006) of augmenting EGO’s expected improvement to allow for balancing areas of high variance and maximizing expected improvement.

Using the maximum likelihood model, the value of the output can be estimated at any point within the bounds using the following equation (27):

$$\hat{Y}(x) = \hat{\mu} + \mathbf{g}r\mathbf{R}^{-1}(y - 1\hat{\mu}) \quad (27)$$

The model also includes an estimate of the standard deviation at every potential point in the input space. In order to calculate it, the variance attributable to spatial transitions, $\hat{\sigma}_Z^2$ (as opposed to measurement error: $\hat{\sigma}_E^2$), must first be calculated (equation 28). Then the standard deviation of the estimate of the potential points can be calculated (equation 29).

$$\hat{\sigma}_Z^2 = g(\hat{\sigma}_Z^2 + \hat{\sigma}_E^2) = g\hat{\sigma}^2 \quad (28)$$

$$s^2(x) = \hat{\sigma}_Z^2 - \begin{bmatrix} \mathbf{1} \\ \hat{\sigma}_Z^2 r(x) \end{bmatrix}^T \begin{bmatrix} \mathbf{0} & \mathbf{1}^T \\ 1 & \hat{\sigma}_Z^2 R \end{bmatrix} \begin{bmatrix} \mathbf{1} \\ \hat{\sigma}_Z^2 r(x) \end{bmatrix} \quad (29)$$

In order to find the next point that will maximize improvement, the algorithm must first identify the current best, x^{**} . This best-point estimate is complicated by the measurement noise in the system. I followed the same approach as the SKO algorithm, which defined a utility function $u(x)$ (equation 30) to identify the current best, x^{**} .

$$u(x) = -\hat{Y}(x) - cs(x) \quad (30)$$

The utility function contains a parameter, c , allowing for a tradeoff between predicted objective values and prediction uncertainty. Following the SKO algorithm, I selected a value of 1 for c . Then, the utility function was maximized over the previously tested points (equation 31):

$$x^{**} = \mathbf{arg\,max}_{x_1, x_2, \dots, x_N}(u(x)) \quad (31)$$

Once the current best is defined, the value which maximally improves on that best can be calculated using the expected improvement (EI, equation 32). As in SKO (Huang et al. 2006), I opted for an augmented expected improvement (AEI), which will allow the algorithm to select replicates near the current best (equation 34).

$$EI(x) = E[\max(\hat{Y}(x^{**}) - \hat{Y}(x), 0)] \quad (32)$$

$$EI(x) = (\hat{Y}(x^{**}) - \hat{Y}(x)) \text{PDF}\left(\frac{\hat{Y}(x^{**}) - \hat{Y}(x)}{s(x)}\right) + s(x) \text{CDF}\left(\frac{\hat{Y}(x^{**}) - \hat{Y}(x)}{s(x)}\right) \quad (33)$$

$$AEI(x) = EI(x) * \left(1 - \sqrt{\frac{\hat{\sigma}_E^2}{s^2(x) + \hat{\sigma}_E^2}}\right) \quad (34)$$

I followed SKO's approach of using a genetic algorithm to maximize AEI (Huang et al. 2006). Branch-and-Bound may be more effective if the parameters p_{θ_k} and p_{δ} are constrained (Jones, Schonlau, and Welch 1998), and may be computationally faster than genetic algorithms for this step.

Once the next candidate point is identified, it is applied to the system. The response from the system is recorded, and a new iteration with one more data point to fit the model begins (section 6.2.3.2).

6.2.3.4. Stop Conditions

In EGO and SKO, the proposed stop conditions related to the estimated error remaining in the system. In a time-varying system that approach no longer makes sense, as the optimum can move, increasing the error at any point. In most cases, stop conditions will be driven by external factors, such as the end of a period of study.

6.2.1. Illustrative Example

Figure 6-1 demonstrates three iterations of the algorithm, showing the model fit, the expected improvement, the selection of the next point, and the new set of samples. The function used is a low-pass-filtered Gaussian process.

Figure 6-1 shows how a function changes with time in each iteration (after each sample), and demonstrates the process of fitting a model, and optimizing within that model to select the best point. In the specific example shown, iteration 14 has a much larger model uncertainty than the two previous iterations, due to the larger deviation between the two points near an input-value of zero. Previous deviations were smaller. Therefore, the model revised its estimate of both time-variance and noise.

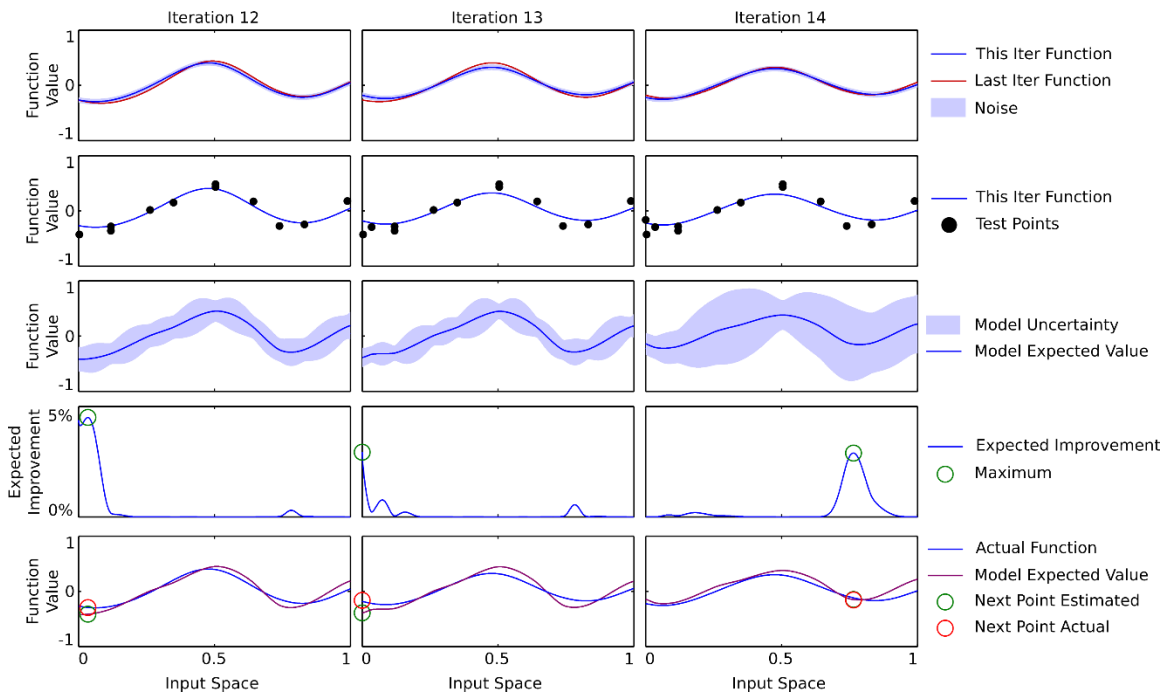


Figure 6-1. Illustrative example of TVN-EGO across 3 iterations during optimization of a 1D function. The first column show iteration 1, the second column shows iteration 2, and the third column shows iteration 3. The plots in the top row show the current function (dark blue line), the noise (filled light blue area), and the previous iteration's function (red line). The second row shows the samples thus far, along with the actual function. The third row shows the model's estimate of the function (dark red line) at the current time and the standard deviation in the estimate (filled pink area). The actual function is shown in the background (dark blue line). The fourth row shows the estimate of expected improvement for all points, with the selected maximum highlighted with a circle. That fifth row shows new sample of the function, in terms of both the estimate of what the value would be, and the actual amount measured.

6.3. Simulated Experiments

The performance of TVN-EGO and two alternative optimization algorithms (Dynamic Simplex and Simultaneous Perturbation Stochastic Approximation) were quantified in a series of experiments. Multiple levels of noise and time-variance were tested across nine test-functions, with coverage of problems with 1–4 input dimensions. This analysis demonstrated that TVN-EGO outperformed the alternatives, and the analysis identified the limits at which TVN-EGO can be effective in terms of noise and time-variance.

6.3.1. Alternative Algorithms

As discussed in the introduction, there are not any good alternatives for globally optimizing black-box, time-varying, noisy systems. Therefore, I compared TVN-EGO with two local optimization algorithms. Dynamic simplex is a time-varying version of the Nelder-mead downhill simplex algorithm. Simultaneous perturbation stochastic approximation (SPSA) attempts to follow the gradient, despite noise. I also proposed a minor modification to SPSA to allow it to handle time-varying systems, and test the resulting algorithm.

6.3.1.1. *Dynamic Simplex*

The dynamic simplex algorithm is a time-varying local optimization algorithm (Xiong and Jutan 2003), based on a modification of the popular downhill simplex algorithm (Nelder and Mead 1965). The major differences are that the simplex does not contract, and that there is no predefined stop condition.

A full summary of the algorithm can be found in the original paper (Xiong and Jutan 2003). Briefly, it is a non-derivative, local-optimization algorithm for time-varying systems. It is not explicitly robust against noise, but coordinate search algorithms tend to

have good noise robustness due to their design (Conn, Scheinberg, and Vicente 2009). Dynamic simplex begins by initializing the first simplex, with a defined step size (defined edge length for the simplex). In each iteration, multiple new simplexes are created by replacing the worst point in the previous simplex with a mirror-image across the other points in the simplex. All points in these multiple new simplexes are tested, and the best simplex is stored for the next iteration.

In order to test the dynamic simplex algorithm, a few possible parameter choices were explored, and the following parameters worked best. In each iteration, two additional simplexes were tested (as recommended in Xiong and Jutan (2003)). The space was scaled to bounds of 0 to 1. The initial simplex was set to a random equilateral simplex within the bounds of 0.1 to 0.9 in all dimensions, and the step size was set to 0.025. A small step size (0.025) will not converge as fast in the beginning, but it should increase the possibility of staying near the optimum as it moves.

6.3.1.2. Simultaneous Perturbation Stochastic Approximation

Stochastic approximation (SA) attempts to directly use gradient-descent in noisy systems through estimating the derivative (Spall 2005). There are two version of SA, one uses finite-differences in all dimensions in every iteration (Kiefer and Wolfowitz 1952), and the other uses one perturbation (2 points) in every iteration, regardless of the dimensionality (Spall 1992). This latter version, called simultaneous perturbation stochastic approximation (SPSA), is capable of following the gradient with very little performance loss per iteration as compared to the finite-difference version (Spall 2005). However, due to SPSA's 2 tests per iteration, it is much more frugal with experiments than finite-difference SA. As there are only two points tested to estimate the gradient, SPSA does not follow a strict gradient descent, instead it follows a more random course.

There are two coefficients in SPSA that normally decay (become progressively smaller) during the optimization session: (1) the distance of the jump to estimate the derivative (often called “ c_k ” in the SA literature), and (2) the velocity coefficient for gradient descent (often called “ a_k ” in the SA literature). In order to optimize time-variant systems, I eliminated the decay in those coefficients, leaving the values constant for the entire session. This approach resulted in a time-varying SPSA. The bounds for our test problem are always scaled to 0–1. The constant coefficient were 0.025 for the derivative-estimating jump (c_k) and 0.1 for the gradient velocity (a_k).

6.3.2. Test Functions

The three optimization algorithms were tested against a suite of test-functions at multiple dimensions, noise levels, and time-variance. There are a great number of test functions in the literature, but there is no standard set to test time-varying optimization function. Therefore, I selected standard test functions, and set one input-dimension to time in each case, leaving the remainder of the dimensions to the input space. In order to vary the levels of time-variance, I modified how fast the time-value changed in the time-dimension of the test function. The maximum time-variance tested for each function was achieved when the samples covered entire range of the time-dimension. In all cases the bounds for both input and time were set to [0, 1], and a scaling-correction (y_c), in some cases specific to the number of input dimension ($y_{c,d}$), is multiplied by the output to scale the spatial variance to 1.

$$0 \leq x_i \leq 1 \quad (35)$$

$$0 \leq t \leq 1 \quad (36)$$

Shifting Quadratic: The shifting quadratic is a simple test function, consisting of a convex quadratic with a minimum moving from 0.1 to 0.9 (in all input-dimensions) during the course of the optimization test. This test function is one of the simplest cases for systems that vary in time, so it was important to see how well TVN-EGO and the other algorithms would perform in this case.

$$dims_x \in [1, 2, 3, 4] \quad (37)$$

$$x_{center} = 0.8t + 0.1 \quad (38)$$

$$y = \sum (x_i - x_{center})^2 \quad (39)$$

$$y_c = [0.023487, 0.047254, 0.070498, 0.094379] \quad (40)$$

$$y_{c,d} = y_c(dims_x) \quad (41)$$

$$y = \frac{y}{\sqrt{y_{c,d}}} \quad (42)$$

Brannin: The Brannin function has two local minima for most of the time range. It is shaped roughly like a valley, and the global minimum moves from lower values to higher values as time progresses. The 2nd input-dimension was set to the time-dimension.

$$dims_x \in [1] \quad (43)$$

$$x_1 = 15x_1 - 5 \quad (44)$$

$$x_2 = 15t \quad (45)$$

$$a = 1, b = \frac{5.1}{4\pi^2}, c = \frac{5}{\pi}, r = 6, s = 10, u = \frac{1}{8\pi} \quad (46)$$

$$y_1 = a (x_2 - bx_1^2 + cx_1 - r)^2 \quad (47)$$

$$y_2 = s(1 - u)\cos(x_1) \quad (48)$$

$$y = y_1 + y_2 + s \quad (49)$$

$$y_c = 1988.3 \quad (50)$$

$$y = \frac{y}{\sqrt{y_c}} \quad (51)$$

Six-hump Camelback: In the time-varying case, the Six-hump Camelback function has a global minimum in the center, and local minima closer to the edges depending on the specific time-value under examination. The second input dimension was set to the time-dimension.

$$dims_x \in [1] \quad (52)$$

$$x_1 = 4x_1 - 1.6 \quad (53)$$

$$x_2 = 2t - 0.8 \quad (54)$$

$$y_1 = x_1^2(4 - 2.1x_1^2 + \frac{x_1^4}{3}) \quad (55)$$

$$y_2 = x_1x_2 \quad (56)$$

$$y_3 = x_2^2(-4 + 4x_2^2) \quad (57)$$

$$y = y_1 + y_2 + y_3 \quad (58)$$

$$y_c = 7.8585 \quad (59)$$

$$y = \frac{y}{\sqrt{y_c}} \quad (60)$$

Levy: The levy function has multiple local minima, very steep sides, and can be a variable number of dimensions. The final input-dimension was set to be time-dimension.

$$dims_x \in [1, 2, 3] \quad (61)$$

$$x_i = 10x_i - 5 \quad (62)$$

$$t = 10t - 5 \quad (63)$$

$$N = dims_x + 1 \quad (64)$$

$$w_i = 1 + \frac{x_i - 1}{4} \quad (65)$$

$$w_N = 1 + \frac{t - 1}{4} \quad (66)$$

$$y_1 = \sin(\pi w_1)^2 \quad (67)$$

$$y_2 = \sum_{i=1}^{N-1} (w_i - 1)^2 (1 + 10 \sin(\pi w_i + 1))^2 \quad (68)$$

$$y_3 = (w_N - 1)^2 (1 + \sin(2\pi w_N))^2 \quad (69)$$

$$y = y_1 + y_2 + y_3 \quad (70)$$

$$y_c = [6.9929, 14.575, 22.184] \quad (71)$$

$$y_{c,d} = y_c(dims_x) \quad (72)$$

$$y = \frac{y}{\sqrt{y_{c,d}}} \quad (73)$$

Hartman 4: The Hartman 4 function is a multimodal function with 4 input dimensions.

The final dimension was set to time, leaving 3 input dimensions.

$$x_4 = t \quad (74)$$

$$\alpha = [1.0, 1.2, 3.0, 3.2] \quad (75)$$

$$A = \begin{bmatrix} 10 & 3 & 17 & 3.5 & 1.7 & 8 \\ 0.05 & 10 & 17 & 0.1 & 8 & 14 \\ 3 & 3.5 & 1.7 & 10 & 17 & 8 \\ 17 & 8 & 0.05 & 10 & 0.1 & 14 \end{bmatrix} \quad (76)$$

$$p = 10^{-4} \begin{bmatrix} 1312 & 1696 & 5569 & 124 & 8283 & 5886 \\ 2329 & 4135 & 8307 & 3736 & 1004 & 9991 \\ 2348 & 1451 & 3522 & 2883 & 3047 & 6650 \\ 4047 & 8828 & 8732 & 5743 & 1091 & 381 \end{bmatrix} \quad (77)$$

$$y = - \sum_{i=1}^4 \alpha_i \exp \left(- \sum_{j=1}^4 A_{ij} (x_j - P_{ij})^2 \right) \quad (78)$$

$$y_c = 0.7271 \quad (79)$$

$$y = \frac{y}{\sqrt{y_c}} \quad (80)$$

Sine-Wave Quadratic Bowl: The sine wave parabola function is a stationary convex quadratic function with low-frequency waves moving along it. The global minima shifts from trough to trough.

$$y = \sum_{i=1}^{\dim_x} (x - 0.5)^2 + \sin(2\pi(0.5 * t + x_1)) \quad (81)$$

$$y_c = [0.5686, 0.5735] \quad (82)$$

$$y_{c,d} = y_c(\mathit{dims}_x) \quad (83)$$

$$y = \frac{y}{\sqrt{y_{c,d}}} \quad (84)$$

6.3.3. Empirical Comparisons

Any global optimization algorithm for noisy, time-varying systems must have limits on its potential effectiveness. I measured the limits empirically with a series of simulations, involving multiple test functions, different numbers of dimensions, multiple levels of noise, and multiple levels of time-variance. All tests in this section are taken from the 80th–120th iterations in simulated runs. This choice allowed the ongoing performance of the algorithm to be analyzed, as opposed to the transients that occur in the early portion of an optimization run. In addition, it allowed direct comparisons between 1D and 2-4D optimization simulations, as the higher dimensional simulations required a longer initialization period (discussed in the “Approach” section).

6.3.3.1. *Performance Metric*

Classically, the performance of optimization algorithms are judged with respect to how close their estimated optimum is to a known optimum after a specific number of iterations or function evaluations. This approach is more complicated in time-varying optimization because the optimum can move, resulting in sudden increases in the error. Therefore, I chose to measure the results over a period of time (the 80th – 120th function evaluations). This period allows the initial settling period for each algorithm to complete, so the performance metric primarily judges how well the algorithm can follow a moving optimum. Ten percent was specified as the maximum allowable error between the

estimated optimum and actual optimum as a “success” for a specific iteration. This 10% error was computed from Cartesian distance in the multidimensional systems. This 10% bound may appear loose by non-time-varying standards, but optimization in a system that is changing is much more challenging than optimizing in stationary systems.

An important observation about a distance-based performance metric is that it will become much harder in higher dimensional problems. In 2D systems, a 10% error is <1% of the area. In 3D system, it is <0.1% of the volume. Strictly speaking, it depends on whether the optimum is on an edge, in a corner, or in the middle of the space, but the geometric trend of increasing difficulty exists in all cases. Considering that the algorithm only has the opportunity to test 120 samples, finding a moving spot that is only 1/1000th of the volume (in 3D) is difficult.

6.3.3.2. Time-variance Limits

In order to determine the effective limits of time-varying optimization, a series of simulated experiments were conducted to determine at what time-varying-level optimization algorithms “break”. Success was defined based on the performance metric described previously.

In these experiments, the time-variance was varied by changing how fast the time-dimension changed in the test function (with zero noise). Practically, this mean the 120 samples covered a smaller time-range of the test function when time-variance was less. In all cases, the time-variance is reported in the normalized form (normalized by the spatial standard deviation), as described in the section on “Measuring Time-Variance”. Nine test functions were examined, across all three optimization algorithms (Figure 6-2). For each combination, 8 levels of time-variance were tested, and for each time-variance level, 10-13 experiments were performed at every level of time-variance tested. The data

was fit to the logistic function, $1/(1 + e^{-x})$, where x was mapped into the logarithmic space that the tests for time-variance were performed in.

The results (Figure 6-2) demonstrate that TVN-EGO outperformed the other two optimization algorithms by a large margin. For the 1D problems, Dynamic Simplex and SPSA performed well in the shifting quadratic function, but they had difficulty with the other two 1D functions due to the multiple local minima. TVN-EGO performed well on all 1D test functions, producing excellent results (approximately 100% success) even at the highest levels of time-variance tested. TVN-EGO even outperformed the other two algorithms on the shifting quadratic, which was the best case for those algorithms, considering the test function only has one minimum.

In the higher dimensional problems, the differences between the algorithms became more apparent (Figure 6-2). Both Dynamic simplex and SPSA demonstrated good results on the shifting quadratic in 2D, but both performed poorly on Levy-2D and the Sine-wave quadratic bowl. In both of these latter two test functions, multiple local minima played a role in the failures of Dynamic simplex and SPSA. In 2D, TVN-EGO outperformed the other two algorithms even for the single optimum case (shifting quadratic). In the 3D and 4D test functions, Dynamic Simplex produced very poor results (0% success in most cases), and SPSA only produced marginally better (<50% success in all cases). TVN-EGO continued to produce good results and time-variance levels of 1%–5%, depending on the test function. In the maximum time-variance case for the shifting quadratic, the optimum moves from one corner to the opposite corner in 120 samples, and TVN-EGO was still able to track the optimum with 80% success.




Figure 6-2. Performance of optimization algorithms in the context of increasing time-variance. The rows show different test functions, with the number of dimension labeled. The columns show the three optimization algorithms. In each plot, individual experiments (simulations) are shown as black circles, and the trend line, calculated from logistic function, is fit. The trend line is shown in blue. In all cases, the time-variance is normalized with respect to the spatial standard deviation (as discussed in the “Approach” section). The performance is the percentage of time the estimated best is within 10% of the actual best across all input dimensions between the 80th and 120th iterations.

6.3.3.3. Noise Limits

There has been a significant amount of research of optimization in noisy systems. Stochastic approximation (Spall 1992) and the modification to EGO (Huang et al. 2006) that I incorporated to handle noise are two examples. However, in these algorithms, it has been shown that they are robust against some noise, but the limits as to how much noise in what test function at what dimensionality have not been reported or compared. Therefore, a series of simulated experiments were conducted to determine the limits of effectiveness with respect to noise.

The noise was varied by adding Gaussian noise to each of the test functions (with zero time-variance). In all cases, the noise is reported in the normalized form (normalized by the spatial standard deviation), as described in the section on “measuring noise”. As with the experiments on time-variance, nine test functions were examined, across all three optimization algorithms (Figure 6-3). For each test-function/algorithm combination, 7 levels of noise were tested, and for each noise-level, 10-13 experiments were performed. The levels of noise tested are high, and in the most difficult tests, the standard deviation of the noise is approximately equal (10^0) to the standard deviation of the function itself. As in the time-variance figure, the data was fit to the logistic function, $1/(1 + e^{-x})$, where x was mapped into the logarithmic space that the levels of noise were represented in.

In the 1D test functions, TVN-EGO demonstrated good performance at noise-levels $>10x$ higher than the max noise-levels at which the other two algorithms showed good performance (Figure 6-3). This result is true for multiple definitions of “good performance”, including the noise-level at which success drops below 50%, 90%, or 100%. In the shifting quadratic test function, Dynamic Simplex and SPSA were above a

90% success rate at approximately 2% noise. TVN-EGO were above 90% success at 50% noise in the 1D shifting quadratic function. Dynamic simplex and SPSA had poor results with the other two 1D test functions.

At higher dimensions, the divergence between TVN-EGO and the other two algorithms was larger. In the 2D test functions, Dynamic Simplex and SPSA demonstrated 90% success only at noises below 0.1%. In the same test function, TVN-EGO demonstrated good performance at 20% noise. With the addition of noise, Dynamic Simplex and SPSA generally had poor performance at higher dimensions with the limited budget of samples our simulations allowed. TVN-EGO produced good results (90% success) even in the 4D test function with a noise level of 5%.




Figure 6-3. Performance of optimization algorithms in the context of increasing noise. The rows show different test functions, with the number of dimension labeled. The columns show the three optimization algorithms. In each plot, individual experiments (simulations) are shown as black circles, and the trend line (blue), calculated from logistic function, is fit. The noise is normalized (as discussed in the “Approach” section). The performance is the percentage of time the estimated best is within 10% of the actual best across all input dimensions between the 80th and 120th iterations.

6.3.3.4. Noise and Time-Variance

In the previous two sections, time-variance-based and noise-based failures were analyzed independently. In this section, the performance of TVN-EGO is analyzed in conditions with both noise and time-variance, as would be encountered in a real system.

For each test function, 7 levels of noise and 8 levels of time-variance were tested. For each noise-level/time-variance-level combination, 10-13 experiments were performed. The levels of noise and time-variance used were the same as in the previous two sections on noise and time-variance respectively. The data was fit with a multiplicative logistic function, $1/(1 + e^{-x_1}) * 1/(1 + e^{-x_2})$, where the x's were mapped into logarithmic spaces.

Figure 6-4 demonstrates that TVN-EGO can be successful in the context of both noise and time-variance. In the 1D test functions, results from shifting quadratic bowl demonstrate that TVN-EGO can perform well (>90% success) at time-variance levels on the order of 1% with noise levels at 10%. In practice that means that TVN-EGO can be successful in simpler 1D spaces, even if every location in the input-to-output mapping is changing by 1% after each sample, and each measurement is perturbed by 10% noise. The results are slightly worse in the higher dimensions, and slightly worse in the test functions with multiple local minima or more complicated structure. In the 4D shifting quadratic, TVN-EGO performs well (>90% success) at 3% noise and 0.5% time-variance.

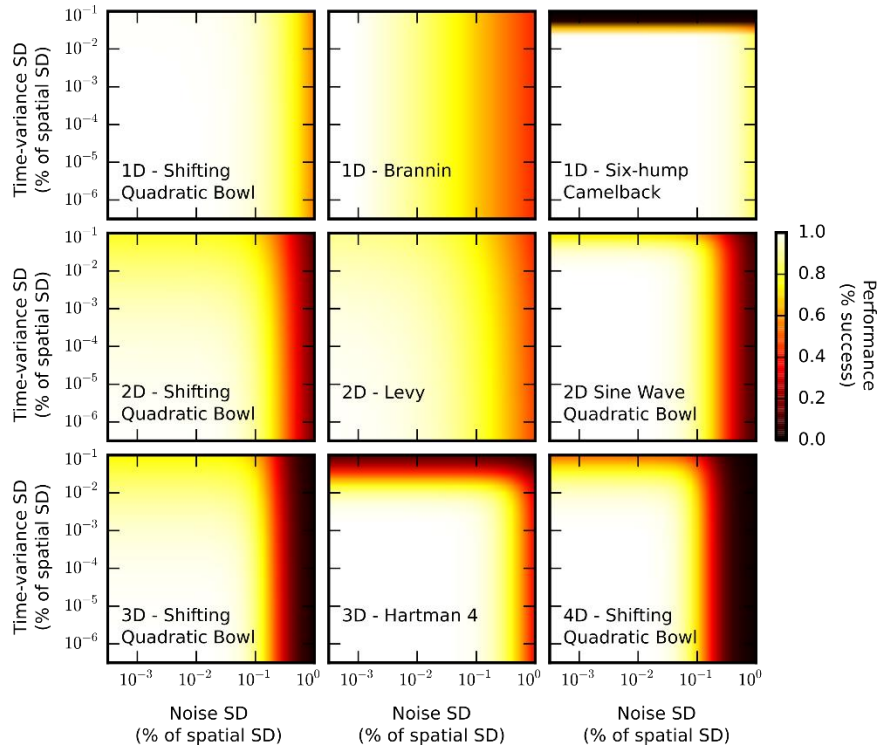


Figure 6-4. Performance in the context of noise and time-variance across multiple test functions and dimensions. Each plot shows a heat map of a specific test function. One-dimensional test functions are on the top row, two-dimensional test functions are on the second row, three- and four-dimensional test functions are on the bottom row. In all cases success was defined based on the performance metric of the number of iterations in which the estimated optimum was within 10% of the actual optimum.

6.4. Discussion

This chapter presents a validation of a novel optimization algorithm for noisy, time-varying systems, based on modifications to the EGO algorithm. These modifications allow for optimization in the presence of noise and time-variance. I proposed scale-and-shift-invariant measures of time-variance and noise, and compared the results of TVN-EGO to other algorithms at multiple levels of noise, time-variance, and dimensionality across several test functions. Through these simulations, I determined that TVN-EGO can produce good results at high levels of noise, time-variance, or both noise and time-variance.

In time-varying optimization, the amount of time-variance is very important. If the change in the function between iterations is much larger than the spatial variance of the function, it can be impossible to optimize. If there is no time-variance, traditional approaches, such as the initial EGO algorithm (Jones, Schonlau, and Welch 1998) or the noise-resistant version (Huang et al. 2006), can be used. In all test functions, the TVN-EGO algorithm produced good results at time-variance levels of 5% for 1D optimization, 1% at 2D, 0.5% at 3D/4D. In all cases, TVN-EGO outperformed the two local optimization algorithms. Dynamically simplex was specifically designed for time-varying systems, and SPSA was modified to handle time-varying systems. TVN-EGO even outperformed the local optimization algorithms in a simple shifting-quadratic test function that did not require global optimization.

In the optimization of noisy systems, TVN-EGO produced good results at very high levels of noise. In the 1D case, TVN-EGO was highly successful with noise at approximately 10% of the standard deviation of the function being optimized. TVN-EGO outperformed the two local optimization algorithms, often producing good results at 1-2 orders of magnitude more noise than either of the other two algorithms could. Dynamic simplex was not designed specifically for noise, but many coordinate search algorithms are robust against small amounts of noise. SPSA was designed explicitly for noise.

Through empirical testing the SPSA parameters were set to constant values that produced the best results, and the test functions were scaled in order to minimize any need to fine-tune those parameters. The differences in the performance of the local optimization algorithms and TVN-EGO are due to how the local optimization algorithms “forgets” previous points, but TVN-EGO continues to use the previous points to build a better model. There has been a significant amount of work on variants of EGO that are robust against noise (Huang et al. 2006, Forrester, Keane, and Bressloff 2006, Picheny,

Wagner, and Ginsbourger 2013), but most of these previous studies have only reported testing 1-2 noise levels. So far as this author is aware, this work represents the most thorough noise characterization of any algorithm based on EGO.

In systems with both noise and time-variance, TVN-EGO substantially outperformed the two local optimization functions. For 1D test functions, TVN-EGO could produce good results at noise levels of 50% of the spatial standard deviation, and time-variance levels of 3%. For higher dimensional test functions, TVN-EGO was still capable of producing good results with substantial noise (20%) and time-variance (1%) in systems with more than three input dimensions.

An important validation of any algorithm is to show how well it handles a real problem. The results of applying this algorithm to sensory stimulation are presented in chapters 3 and 4.

6.4.1. Dimensional Scalability

Dimensional scalability is a major issue in optimization, and time-varying optimization is no exception. In general, our results for similar test functions (e.g. shifting quadratic), showed that each dimension reduced performance in noise by a factor 0.5. The reduced performance from increased time-variance was less clear, but also evident. This reduction in performance is partially the result of the metric chosen, which represented a geometrically smaller volume with each additional dimension (e.g. $\leq 10\%$ in 1D case, $\leq 1\%$ in 2D, $\leq 0.1\%$ in 3D, etc.). To put it another way: there are simply more places to look in higher dimensional functions. In addition, the number of model parameters to estimate in TVN-EGO increases linearly ($2n$) with the number of dimensions, causing the model fit to be more difficult with each additional dimension.

Our results demonstrate that the best performance can be seen in 1D functions (success at 1% time-variance and 10% noise). If there are slightly lower levels of noise and time-variance (e.g. 0.5% time-variance and 3% noise), TVN-EGO can be successful at optimizing test functions up to 4 dimensions.

6.4.2. Computational Efficiency

Much of the focus of this paper has been on the high-degree of efficiency this algorithm has with regards to tests or experiments. It is also important to understand the computational requirements to run this algorithm, particularly if the reader is interested in real-time performance (as the application example required).

The major limiting factor for computational performance is fitting the model. It requires decompositions (generally Cholesky decompositions due to the symmetric, positive-definite matrix) to efficiently solve several equations of the form $x = A^{-1}b$. During a single iteration's worth of model-fit computation, hundreds or thousands of Cholesky decompositions need to be performed on NxN matrices, where N is the number of points tested thus far. Cholesky decompositions for dense matrices scale at $O(n^3)$ (Trefethen and Bau III 1997), which means the number of potential model fits that can be tested will become limited. On a recent (2013) Dell Precision workstation, an iteration involving 400 test points was completed in 4 seconds in MATLAB.

Both local optimization algorithms presented as alternatives in this paper are >1,000x more computationally efficient than TVN-EGO is, but their performance is 10-100x worse with respect to noise and time-variance. Specific use cases will need to determine their ideal tradeoff between these advantages.

6.4.3. Possible Extensions

There are extensions to TVN-EGO that could improve this algorithm, but will be left to future work. These extensions include (1) extending the approach to time-variance to multiple uncontrolled, but measurable inputs, (2) combining the time-variance approach with other modifications to EGO, (3) improving the computational performance, or (4) increasing the number of dimensions that can be handled by this approach.

Our approach to time-variance was to treat time as an uncontrolled, extrinsic input. The approach could be easily generalized to handle multiple uncontrolled extrinsic inputs, and our current equations are already written to handle such an approach.

A rich literature has developed regarding ways to improve EGO and closely related algorithms such as Gaussian process modeling and kriging. These modifications are generally designed to modify the sampling criteria (e.g. Sasena, Papalambros, and Goovaerts (2002)) or improve performance when the function has certain types of input-dimension non-stationarities (e.g. Xiong et al. (2007)). These modifications could be combined with our approach.

As mentioned previously, the major limiting factor in computational performance is the matrix decomposition (i.e. Cholesky decomposition) required to solve terms like $A^{-1}b$ in the maximum likelihood estimate of the model. In order to handle larger quantities of test points, changes to the covariance matrix to take advantage of sparsity in the matrix could greatly reduce the computational burden. The number of operations in sparse Cholesky decomposition increases much slower than $O(n^3)$, depending on precisely how sparse the matrix is (Davis and Hager 1999). Alternative approaches could also include mapping a large number of test points to a smaller kernel, and using that kernel in the model.

Lastly, this chapter only presented results up to 4 input-dimensions (and 1 time-dimension). Models could be constructed from problems with tens or hundreds of input-dimensions if the model-parameters were shared across input variables, similar to the approach used in convolving neural networks (LeCun and Bengio 1995, Simard, Steinkraus, and Platt 2003). Sharing could be applicable if a large subset of those dimensions were all of the same type or same types. For example, if they were values of a waveform after an event occurred (i.e. voltage measures in time), they could be considered to be the same measurement type, despite the waveform consisting of 10's or 100's of dimensions.

6.5. Conclusions

This chapter has introduced TVN-EGO, an optimization algorithm designed for noisy, time-varying, black-box systems. TVN-EGO proved to have good performance across test functions at high levels of both noise and time-variance in multi-dimensional systems with multiple local minima. The analysis has identified the limits of TVN-EGO, providing guidance on when it should be applied.

CHAPTER 7: CONCLUSIONS

In summary, I developed a novel approach to optimize a medical intervention in real-time. I showed that this approach outperformed open-loop approaches across multiple subjects, multiple stimulation sites, and multiple stimulation paradigms.

Algorithmic optimization of interventions represents a major shift in how medicine is practiced. This research is among the first steps in this direction. The approaches detailed in this dissertation can be adapted to most types of interventions, so long as the outcome can be measured quickly, and the intervention can be manipulated in closed-loop.

In the process of developing this system, I discovered that the human system is black-box, noisy, and time-varying. There were no pre-existing algorithms specifically designed for globally-optimizing noisy, time-varying, black-box systems. Therefore, I created the first such algorithm I am aware of, through the modification of an efficiency-focused optimization algorithm ("Efficient Global Optimization", designed for systems with expensive experiments). A focus on efficiency allowed my algorithm to maximize the gain per experiment, and handle time-variance at a higher level than other more targeted algorithms (e.g. local optimizers). In order to quantitatively test this algorithm, I created a novel metric for the time-variance in entire functions, building on work from time-series analysis. Optimization problems are common across disciplines. Therefore, this algorithm has the potential to impact many fields. Similar systems (i.e. expensive, noisy, time-varying, and black-box) exist across many fields, including medicine, economics, agriculture, and engineering.

The system that I developed was novel in several ways. The system was capable of measuring inputs related to gait, and acting on those inputs algorithmically in real-time.

The Lokomat is generally considered to be an intervention, but in this research, it was demonstrated that it can also be a quality-of-walking-measurement device.

The results from my studies have demonstrated several novel and significant findings with respect to the current clinical practice of sensory stimulation. For example, across nearly all gait features, longer-duration pulse trains produced better results, in contrast to the current approach of using short-duration pulse trains. History (i.e. how long the subject had participated in the experiment) proved to be as important a predictor as the stimulation parameters that were being used. The data generated from the stimulation parameter sweeps allowed me to produce a guide for how to maximally improve specific features of gait (muscle activation or force). This guide should assist clinicians with better ways to address specific shortcomings. Through the analysis of these features, I identified a novel gait-phase dependence from transcutaneous spinal cord stimulation, and better characterized how cutaneous nerve stimulation can assist gait. These discoveries have direct application to guide current clinical treatment and future research in SCI.

The specific approach of optimizing sensory stimulation to assist people with walking after SCI could be extended to wearable devices and overground locomotion. Sensory stimulation has already been combined with physical therapy and functional electrical stimulation with positive results, and optimizing that sensory stimulation can only improve those results. These explorations will continue to further the state of SCI medicine.

There are other neurological injuries that should be just as amenable to the optimized sensory stimulation. For example, stroke, traumatic brain injury, and multiple sclerosis share many similarities with SCI in how they affect motor control. The optimization of sensory stimulation could assist people with those conditions as well.

There are many minor ways in which my research may have impact. This study represents the largest clinical study of transcutaneous spinal cord stimulation. I developed a new metric for measuring the quality of walking after SCI, and this walking metric could be applied to other walking pathologies (e.g. Stroke, Parkinson's, etc.). I characterized the onset response to cutaneous stimulation, which had only been observed in two other papers before, without any analysis. My approach to modeling rotationally symmetric spaces is a novel way to represent this type of input-to-output mapping, and it could have application to many problems both inside and outside biology.

REFERENCES

- Abraira, Victoria E, and David D Ginty. 2013. "The sensory neurons of touch." *Neuron* no. 79 (4):618-639.
- Anderson, K. D. 2004. "Targeting recovery: priorities of the spinal cord-injured population." *J Neurotrauma* no. 21 (10):1371-83.
- Andersson, O., and S. Grillner. 1983. "Peripheral control of the cat's step cycle. II. Entrainment of the central pattern generators for locomotion by sinusoidal hip movements during "fictive locomotion."." *Acta Physiol Scand* no. 118 (3):229-39.
- Andersson, Olof, H. Forssberg, Sten Grillner, and M. Lindquist. 1978. "Phasic gain control of the transmission in cutaneous reflex pathways to motoneurons during 'fictive' locomotion." *Brain research* no. 149 (2):503-507.
- Arendt-Nielsen, Lars, Jannick Brennum, Soren Sindrup, and Peter Bak. 1994. "Electrophysiological and psychophysical quantification of temporal summation in the human nociceptive system." *European Journal of Applied Physiology and Occupational Physiology* no. 68 (3):266-273.
- Arsenault, A., D. Winter, and R. Marteniuk. 1986. "Is there a 'normal' profile of EMG activity in gait?" *Medical and Biological Engineering and Computing* no. 24 (4):337-343.
- Avelev, V, N Anissimova, E Khoroshikh, and Y Gerasimenko. 1997. Activation of central pattern generator in cat by epidural spinal cord stimulation. Paper read at Proceedings of the International Symposium on Brain and Movement St. Petersburg, Moscow.
- Baken, Bernke Christianne Maria, Volker Dietz, and Jacques Duysens. 2005. "Phase-dependent modulation of short latency cutaneous reflexes during walking in man." *Brain research* no. 1031 (2):268-275.
- Banala, Sai K, Seok Hun Kim, Sunil K Agrawal, and John P Scholz. 2009. "Robot assisted gait training with active leg exoskeleton (ALEX)." *Neural Systems and Rehabilitation Engineering, IEEE Transactions on* no. 17 (1):2-8.
- Barbeau, Hughes, and Serge Rossignol. 1987. "Recovery of locomotion after chronic spinalization in the adult cat." *Brain research* no. 412 (1):84-95.
- Barbeau, Hugues, Michel Danakas, and Bertrand Arsenault. 1993. "The effects of locomotor training in spinal cord injured subjects: a preliminary study." *Restorative neurology and neuroscience* no. 5 (1):81-84.
- Bareyre, F. M., M. Kerschensteiner, O. Raineteau, T. C. Mettenleiter, O. Weinmann, and M. E. Schwab. 2004. "The injured spinal cord spontaneously forms a new intraspinal circuit in adult rats." *Nat Neurosci* no. 7 (3):269-77.
- Barton, Russell R, and John S Ivey Jr. 1996. "Nelder-Mead simplex modifications for simulation optimization." *Management Science* no. 42 (7):954-973.
- Basso, D Michele, Michael S Beattie, and Jacqueline C Bresnahan. 1995. "A sensitive and reliable locomotor rating scale for open field testing in rats." *Journal of neurotrauma* no. 12 (1):1-21.
- Basso, D Michele, Lesley C Fisher, Aileen J Anderson, Lyn B Jakeman, Dana M Mctigue, and Phillip G Popovich. 2006. "Basso Mouse Scale for locomotion detects differences in recovery after spinal cord injury in five common mouse strains." *Journal of neurotrauma* no. 23 (5):635-659.

- Bastiaanse, CM, Jaak Duysens, and V Dietz. 2000. "Modulation of cutaneous reflexes by load receptor input during human walking." *Experimental brain research* no. 135 (2):189-198.
- Behrman, Andrea L, and Susan J Harkema. 2000. "Locomotor Training After Human Spinal Cord Injury: A Series of Case Studies." *Phys Ther* no. 80 (7):688-700.
- Belanger, Marc, Trevor Drew, Janyne Provencher, and Serge Rossignol. 1996. "A comparison of treadmill locomotion in adult cats before and after spinal transection." *Journal of Neurophysiology* no. 76 (1):471-491.
- Benda, J., T. Gollisch, C. K. Machens, and A. V. Herz. 2007. "From response to stimulus: adaptive sampling in sensory physiology." *Current opinion in neurobiology* no. 17 (4):430-6.
- Beresovskii, VK, and KV Bayev. 1988. "New locomotor regions of the brainstem revealed by means of electrical stimulation." *Neuroscience* no. 26 (3):863-869.
- Betró, Bruno. 1992. "Bayesian methods in global optimization." In *Operations Research'91*, 16-18. Springer.
- Blaya, Joaquin, and Hugh Herr. 2004. "Adaptive control of a variable-impedance ankle-foot orthosis to assist drop-foot gait." *Neural Systems and Rehabilitation Engineering, IEEE Transactions on* no. 12 (1):24-31.
- Bolton, David AE, and John E Misiaszek. 2009. "Contribution of hindpaw cutaneous inputs to the control of lateral stability during walking in the cat." *Journal of neurophysiology* no. 102 (3):1711-1724.
- Bostock, H., K. Cikurel, and D. Burke. 1998. "Threshold tracking techniques in the study of human peripheral nerve." *Muscle Nerve* no. 21 (2):137-58.
- Boulenguez, P., S. Liabeuf, R. Bos, H. Bras, C. Jean-Xavier, C. Brocard, A. Stil, P. Darbon, D. Cattaert, E. Delpire, M. Marsala, and L. Vinay. 2010. "Down-regulation of the potassium-chloride cotransporter KCC2 contributes to spasticity after spinal cord injury." *Nat Med* no. 16 (3):302-7.
- Bouyer, Laurent JG, and S Rossignol. 2003a. *Contribution of Cutaneous Inputs From the Hindpaw to the Control of Locomotion. II. Spinal Cats*. Vol. 90.
- Bouyer, Laurent JG, and Serge Rossignol. 1998. "The Contribution of Cutaneous Inputs to Locomotion in the Intact and the Spinal Cata." *Annals of the New York Academy of Sciences* no. 860 (1):508-512.
- Bouyer, Laurent JG, and Serge Rossignol. 2003b. "Contribution of cutaneous inputs from the hindpaw to the control of locomotion. I. Intact cats." *Journal of neurophysiology* no. 90 (6):3625-3639.
- Bradbury, Elizabeth J., and Lucy M. Carter. 2011. "Manipulating the glial scar: Chondroitinase ABC as a therapy for spinal cord injury." *Brain Research Bulletin* no. 84 (4-5):306-316.
- Brooke, JD, J Cheng, DF Collins, WE McIlroy, JE Misiaszek, and WR Staines. 1997. "Sensory afferent conditioning with leg movement: gain control in spinal reflex and ascending paths." *Progress in neurobiology* no. 51 (4):393-421.
- Brown, DA, and CG Kukulka. 1993. "Human flexor reflex modulation during cycling." *Journal of neurophysiology* no. 69 (4):1212-1224.
- Brown, T Graham. 1911a. "The intrinsic factors in the act of progression in the mammal." *Proceedings of the Royal Society of London. Series B, containing papers of a biological character*:308-319.
- Brown, T Graham. 1911b. "Studies in the Physiology of the Nervous System. VIII. Neural Balance and Reflex Reversal, with a Note on Progression in the Decerebrate Guinea-Pig." *Quarterly Journal of Experimental Physiology* no. 4 (3):273-288.
- Brown, T Graham. 1914. "On the nature of the fundamental activity of the nervous centres; together with an analysis of the conditioning of rhythmic activity in progression, and a

- theory of the evolution of function in the nervous system." *The Journal of Physiology* no. 48 (1):18-46.
- Broyden, Charles G. 1970. "The convergence of a class of double-rank minimization algorithms 2. The new algorithm." *IMA Journal of Applied Mathematics* no. 6 (3):222-231.
- Burke, David, Hugh G Dickson, and Nevell F Skuse. 1991. "Task-dependent changes in the responses to low-threshold cutaneous afferent volleys in the human lower limb." *The Journal of Physiology* no. 432 (1):445-458.
- Burke, RE. 1999. "The use of state-dependent modulation of spinal reflexes as a tool to investigate the organization of spinal interneurons." *Experimental brain research* no. 128 (3):263-277.
- Bussel, B, A Roby-Brami, O Rémy Nériss, and A Yakovleff. 1996. "Evidence for a spinal stepping generator in man." *Spinal Cord* no. 34 (2):91-92.
- Bussel, B, A Roby-Brami, A Yakovleff, and N Bennis. 1989. "Late flexion reflex in paraplegic patients. Evidence for a spinal stepping generator." *Brain research bulletin* no. 22 (1):53-56.
- Bussel, B., A. Roby-Brami, P. Azouvi, A. Biraben, A. Yakovleff, and J. P. Held. 1988. "Myoclonus in a patient with spinal cord transection. Possible involvement of the spinal stepping generator." *Brain* no. 111 (Pt 5):1235-45.
- Calancie, B., B. Needham-Shropshire, P. Jacobs, K. Willer, G. Zych, and B. A. Green. 1994. "Involuntary stepping after chronic spinal cord injury. Evidence for a central rhythm generator for locomotion in man." *Brain* no. 117 (Pt 5):1143-59.
- Capaday, C, and R B Stein. 1987. "Difference in the amplitude of the human soleus H reflex during walking and running." *J Physiol* no. 392 (1):513-522.
- Capaday, C, and RB Stein. 1986. "Amplitude modulation of the soleus H-reflex in the human during walking and standing." *The Journal of Neuroscience* no. 6 (5):1308-1313.
- Carrier, Lynda, Edna Brustein, and Serge Rossignol. 1997. "Locomotion of the hindlimbs after neurectomy of ankle flexors in intact and spinal cats: model for the study of locomotor plasticity." *Journal of neurophysiology* no. 77 (4):1979-1993.
- Center, National Spinal Cord Injury Statistical. 2012. "Spinal Cord Injury Facts and Figures at a Glance. Information Sheet." no. February 2012.
- Černý, Vladimír. 1985. "Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm." *Journal of optimization theory and applications* no. 45 (1):41-51.
- Chiang, AC. 2000. "Elements of dynamic optimization." *Illinois: Waveland Press Inc.*
- Chiu, Wen-Ta, Hsiao-Chiao Lin, Carlos Lam, Shu-Fen Chu, Yung-Hsiao Chiang, and Shin-Han Tsai. 2010. "Review paper: epidemiology of traumatic spinal cord injury: comparisons between developed and developing countries." *Asia-Pacific journal of public health* no. 22 (1):9-18.
- Cline, Daren B H, and Jeffrey D Hart. 1991. "Kernel estimation of densities with discontinuities or discontinuous derivatives." *Statistics: A Journal of Theoretical and Applied Statistics* no. 22 (1):69-84.
- Coburn, Barry. 1985. "A theoretical study of epidural electrical stimulation of the spinal cord- Part II: Effects on long myelinated fibers." *Biomedical Engineering, IEEE Transactions on* (11):978-986.
- Collins, D. F., D. Burke, and S. C. Gandevia. 2001. "Large involuntary forces consistent with plateau-like behavior of human motoneurons." *J Neurosci* no. 21 (11):4059-65.
- Colombo, G, M Wirz, and V Dietz. 2001. "Driven gait orthosis for improvement of locomotor training in paraplegic patients." *Spinal Cord* no. 39 (5):252-255.

- Colombo, Gery, Matthias Joerg, Reinhard Schreier, and Volker Dietz. 2000. "Treadmill training of paraplegic patients using a robotic orthosis." *Journal of rehabilitation research and development* no. 37 (6):693-700.
- Conn, Andrew R, Katya Scheinberg, and Luis N Vicente. 2009. *Introduction to derivative-free optimization*. Vol. 8: Siam.
- Conn, AR, K Scheinberg, and Ph L Toint. 1998. A derivative free optimization algorithm in practice. Paper read at Proceedings of 7th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, St. Louis, MO.
- Conway, BA, H Hultborn, and O Kiehn. 1987. "Proprioceptive input resets central locomotor rhythm in the spinal cat." *Experimental Brain Research* no. 68 (3):643-656.
- Cook, AW, and SP Weinstein. 1973. "Chronic dorsal column stimulation in multiple sclerosis. Preliminary report." *New York state journal of medicine* no. 73 (24):2868-2872.
- Côté, Marie-Pascale, and Jean-Pierre Gossard. 2004. "Step training-dependent plasticity in spinal cutaneous pathways." *The Journal of neuroscience* no. 24 (50):11317-11327.
- Courtine, G., B. Song, R. R. Roy, H. Zhong, J. E. Herrmann, Y. Ao, J. Qi, V. R. Edgerton, and M. V. Sofroniew. 2008. "Recovery of supraspinal control of stepping via indirect propriospinal relay connections after spinal cord injury." *Nat Med* no. 14 (1):69-74. doi: 10.1038/nm1682.
- Courtine, Grégoire, Yury Gerasimenko, Rubia van den Brand, Aileen Yew, Pavel Musienko, Hui Zhong, Bingbing Song, Yan Ao, Ronaldo M Ichiyama, and Igor Lavrov. 2009. "Transformation of nonfunctional spinal circuits into functional states after the loss of brain input." *Nature neuroscience* no. 12 (10):1333-1342.
- Crago, Patrick E, J Thomas Mortimer, and P Hunter Peckham. 1980. "Closed-loop control of force during electrical stimulation of muscle." *Biomedical Engineering, IEEE Transactions on* (6):306-312.
- Crenna, Paolo, and Carlo Frigo. 1984. "Evidence of phase-dependent nociceptive reflexes during locomotion in man." *Experimental neurology* no. 85 (2):336-345.
- Cressie, Noel. 1990. "The origins of kriging." *Mathematical geology* no. 22 (3):239-252.
- D'Agostino, Antonello, Luca Gambetti, and Domenico Giannone. 2013. "Macroeconomic forecasting and structural change." *Journal of Applied Econometrics* no. 28 (1):82-101.
- Danner, Simon M, Ursula S Hofstoetter, Brigitta Freundl, Heinrich Binder, Winfried Mayr, Frank Rattay, and Karen Minassian. 2015. "Human spinal locomotor control is based on flexibly organized burst generators." *Brain*:awu372.
- Danner, Simon M, Ursula S Hofstoetter, Josef Ladenbauer, Frank Rattay, and Karen Minassian. 2011. "Can the human lumbar posterior columns be stimulated by transcutaneous spinal cord stimulation? A modeling study." *Artificial organs* no. 35 (3):257-262.
- Datta, AK, LM Harrison, and JA Stephens. 1989. "Task-dependent changes in the size of response to magnetic brain stimulation in human first dorsal interosseous muscle." *The Journal of physiology* no. 418 (1):13-23.
- Davis, Timothy A, and William W Hager. 1999. "Modifying a sparse Cholesky factorization." *SIAM Journal on Matrix Analysis and Applications* no. 20 (3):606-627.
- De Leon, RD, JA Hodgson, RR Roy, and VR Edgerton. 1999. "Retention of hindlimb stepping ability in adult spinal cats after the cessation of step training." *Journal of Neurophysiology* no. 81 (1):85-94.
- De Luca, Carlo J. 1997. "The use of surface electromyography in biomechanics." *Journal of applied biomechanics* no. 13:135-163.

- De Serres, SJ, JF Yang, and SK Patrick. 1995. "Mechanism for reflex reversal during walking in human tibialis anterior muscle revealed by single motor unit recording." *The Journal of physiology* no. 488 (1):249-258.
- Dekkers, Anton, and Emile Aarts. 1991. "Global optimization and simulated annealing." *Mathematical programming* no. 50 (1-3):367-393.
- Desautels, Thomas A. 2014. *Spinal cord injury therapy through active learning*, California Institute of Technology.
- Desautels, Thomas, Jaehoon Choe, Parag Gad, Mandheerej Nandra, Roland Roy, Hui Zhong, Yu-Chong Tai, V Edgerton, and Joel Burdick. 2015. "An Active Learning Algorithm for Control of Epidural Electrostimulation."
- Di Prisco, Gonzalo Viana, Edouard Pearlstein, Didier Le Ray, Richard Robitaille, and Réjean Dubuc. 2000. "A cellular mechanism for the transformation of a sensory input into a motor command." *The Journal of Neuroscience* no. 20 (21):8169-8176.
- Dickey, David A, and Wayne A Fuller. 1979. "Distribution of the estimators for autoregressive time series with a unit root." *Journal of the American statistical association* no. 74 (366a):427-431.
- Dietz, V, S Grillner, A Trepp, M Hubli, and M Bolliger. 2009. "Changes in spinal reflex and locomotor activity after a complete spinal cord injury: a common mechanism?" *Brain* no. 132 (8):2196-2205.
- Dietz, V, M Wirz, A Curt, and G Colombo. 1998. "Locomotor pattern in paraplegic patients: training effects and recovery of spinal cord function." *Spinal cord* no. 36 (6):380-390.
- Dietz, V., K. Nakazawa, M. Wirz, and T. Erni. 1999. "Level of spinal cord lesion determines locomotor activity in spinal man." *Exp Brain Res* no. 128 (3):405-9.
- Dietz, Volker. 2002. "Do human bipeds use quadrupedal coordination?" *Trends in neurosciences* no. 25 (9):462-467.
- Dietz, Volker, and Jacques Duysens. 2000. "Significance of load receptor input during locomotion: a review." *Gait & posture* no. 11 (2):102-110.
- Dietz, Volker, and Roland Müller. 2004. "Degradation of neuronal function following a spinal cord injury: mechanisms and countermeasures." *Brain* no. 127 (10):2221-2231.
- Dietz, Volker, Roland Müller, and Gery Colombo. 2002. "Locomotor activity in spinal man: significance of afferent input from joint and load receptors." *Brain* no. 125 (12):2626-2634.
- Dimitrijevic, Milan R, Yuri Gerasimenko, and Michaela M Pinter. 1998. "Evidence for a spinal central pattern generator in humans." *Annals of the New York Academy of Sciences* no. 860 (1):360-376.
- Dimitrijevic, MR, Y Gerasimenko, and FE Pollo. 1997. Sustained epidural electrical stimulation of the spinalized lumbar cord in humans can facilitate or suppress spinal reflex and central pattern generator activities. Paper read at Proceedings of the International Symposium on Brain and Movement, St. Petersburg, Moscow.
- Dimitrijevic, MR, and PW Nathan. 1969. "Changes in the flexion reflex with repetitive cutaneous stimulation in spinal man." *Electroencephalography and clinical neurophysiology* no. 27 (7):721-722.
- Dimitrijevic, MR, M Pinter, and AM Sherwood. 1997. The effect of reduced suprasegmental input as a result of spinal cord injury on stepping movement induced by epidural stimulation of the lumbar enlargement in humans. Paper read at Proceedings of the International Symposium on Brain and Movement, St. Petersburg, Moscow.

- Drew, T, and S Rossignol. 1987. "A kinematic and electromyographic study of cutaneous reflexes evoked from the forelimb of unrestrained walking cats." *Journal of neurophysiology* no. 57 (4):1160-1184.
- Drew, Trevor, Thérèse Cabana, and Serge Rossignol. 1996. "Responses of medullary reticulospinal neurones to stimulation of cutaneous limb nerves during locomotion in intact cats." *Experimental brain research* no. 111 (2):153-168.
- Dutta, A., R. Kobetic, and R. J. Triolo. 2008. "Ambulation after incomplete spinal cord injury with EMG-triggered functional electrical stimulation." *IEEE Trans Biomed Eng* no. 55 (2 Pt 1):791-4. doi: 10.1109/TBME.2007.902225.
- Duysens, J. 1977a. "Fluctuations in sensitivity to rhythm resetting effects during the cat's step cycle." *Brain research* no. 133 (1):190-195.
- Duysens, J. 1977b. "Reflex control of locomotion as revealed by stimulation of cutaneous afferents in spontaneously walking preamammillary cats." *Journal of Neurophysiology* no. 40 (4):737-751.
- Duysens, J, and GE Loeb. 1980. "Modulation of ipsi-and contralateral reflex responses in unrestrained walking cats." *Journal of Neurophysiology* no. 44 (5):1024-1037.
- Duysens, J, GE Loeb, and BJ Weston. 1980. "Crossed flexor reflex responses and their reversal in freely walking cats." *Brain research* no. 197 (2):538-542.
- Duysens, J, and KG Pearson. 1976. "The role of cutaneous afferents from the distal hindlimb in the regulation of the step cycle of thalamic cats." *Experimental Brain Research* no. 24 (3):245-255.
- Duysens, J, and KJG Pearson. 1980. "Inhibition of flexor burst generation by loading ankle extensor muscles in walking cats." *Brain research* no. 187 (2):321-332.
- Duysens, J, and RB Stein. 1978. "Reflexes induced by nerve stimulation in walking cats with implanted cuff electrodes." *Experimental brain research* no. 32 (2):213-224.
- Duysens, J, AA Tax, L Murrer, and V Dietz. 1996. "Backward and forward walking use different patterns of phase-dependent modulation of cutaneous reflexes in humans." *Journal of Neurophysiology* no. 76 (1):301-310.
- Duysens, J, AAM Tax, M Trippel, and V Dietz. 1992. "Phase-dependent reversal of reflexly induced movements during human gait." *Experimental brain research* no. 90 (2):404-414.
- Duysens, J, AAM Tax, M Trippel, and V Dietz. 1993. "Increased amplitude of cutaneous reflexes during human running as compared to standing." *Brain research* no. 613 (2):230-238.
- Duysens, J, M Trippel, GA Horstmann, and V Dietz. 1990. "Gating and reversal of reflexes in ankle muscles during human walking." *Experimental brain research* no. 82 (2):351-358.
- Duysens, Jaak, BCM Baken, L Burgers, FM Plat, AR Den Otter, and HPH Kremer. 2004. "Cutaneous reflexes from the foot during gait in hereditary spastic paraparesis." *Clinical neurophysiology* no. 115 (5):1057-1062.
- Duysens, Jacques, and Henry WAA Van de Crommert. 1998. "Neural control of locomotion; Part 1: The central pattern generator from cats to humans." *Gait & posture* no. 7 (2):131-141.
- Dykstra, D. D., A. A. Sidi, A. B. Scott, J. M. Pagel, and G. D. Goldish. 1988. "Effects of botulinum A toxin on detrusor-sphincter dyssynergia in spinal cord injury patients." *J Urol* no. 139 (5):919-922.
- Eberhart, Russ C, and James Kennedy. 1995. A new optimizer using particle swarm theory. Paper read at Proceedings of the sixth international symposium on micro machine and human science.

- Edgerton, V. R., R. D. Leon, S. J. Harkema, J. A. Hodgson, N. London, D. J. Reinkensmeyer, R. R. Roy, R. J. Talmadge, N. J. Tillakaratne, W. Timoszyk, and A. Tobin. 2001. "Retraining the injured spinal cord." *J Physiol* no. 533 (Pt 1):15-22.
- Edin, F., C. K. Machens, H. Schutze, and A. V. Herz. 2004. "Searching for optimal sensory signals: iterative stimulus reconstruction in closed-loop experiments." *Journal of computational neuroscience* no. 17 (1):47-56.
- Eng, Janice J, David A Winter, and Aftab E Patla. 1994. "Strategies for recovery from a trip in early and late swing during human walking." *Experimental Brain Research* no. 102 (2):339-349.
- Erlanger, Joseph, and Edgar A Blair. 1938. "Comparative observations on motor and sensory fibers with special reference to repetitiousness." *American Journal of Physiology--Legacy Content* no. 121 (2):431-453.
- Evans, AL, Linda M Harrison, and JA Stephens. 1989. "Task-dependent changes in cutaneous reflexes recorded from various muscles controlling finger movement in man." *The Journal of physiology* no. 418 (1):1-12.
- Evans, AL, Linda M Harrison, and JA Stephens. 1990. "Maturation of the cutaneomuscular reflex recorded from the first dorsal interosseous muscle in man." *The Journal of physiology* no. 428 (1):425-440.
- Fawcett, JW, A Curt, JD Steeves, WP Coleman, MH Tuszynski, D Lammertse, PF Bartlett, AR Blight, V Dietz, and J Ditunno. 2007. "Guidelines for the conduct of clinical trials for spinal cord injury as developed by the ICCP panel: spontaneous recovery after spinal cord injury and statistical power needed for therapeutic clinical trials." *Spinal cord* no. 45 (3):190-205.
- Ferris, DP, KE Gordon, JA Beres-Jones, and SJ Harkema. 2004. "Muscle activation during unilateral stepping occurs in the nonstepping limb of humans with clinically complete spinal cord injury." *Spinal cord* no. 42 (1):14-23.
- Field-Fote, Edelle C. 2001. "Combined use of body weight support, functional electric stimulation, and treadmill training to improve walking ability in individuals with chronic incomplete spinal cord injury." *Arch Phys Med Rehabil* no. 82 (6):818-824.
- Field-Fote, Edelle Carmen, and Dejan Tepavac. 2002. "Improved Intralimb Coordination in People With Incomplete Spinal Cord Injury Following Training With Body Weight Support and Electrical Stimulation." *Physical Therapy* no. 82 (7):707-715.
- Fletcher, Roger. 1970. "A new approach to variable metric algorithms." *The computer journal* no. 13 (3):317-322.
- Földiák, Peter. 2001. "Stimulus optimisation in primary visual cortex." *Neurocomputing* no. 38-40 (0):1217-1222. doi: 10.1016/s0925-2312(01)00570-7.
- Fontana, F. 1760. "Mémoires sur les parties sensibles et irritables du corps animal."
- Fornusek, Ché, and Glen M Davis. 2008. "Cardiovascular and metabolic responses during functional electric stimulation cycling at different cadences." *Archives of physical medicine and rehabilitation* no. 89 (4):719-725.
- Forrester, Alexander I J, Andy J Keane, and Neil W Bressloff. 2006. "Design and analysis of "Noisy" computer experiments." *AIAA journal* no. 44 (10):2331-2339.
- Forsberg, H. 1979. "Stumbling corrective reaction: a phase-dependent compensatory reaction during locomotion." *Journal of Neurophysiology* no. 42 (4):936-953.
- Forsberg, H. 1985. "Ontogeny of human locomotor control I. Infant stepping, supported locomotion and transition to independent locomotion." *Experimental Brain Research* no. 57 (3):480-493.

- Forssberg, H, S Grillner, and S Rossignol. 1977. "Phasic gain control of reflexes from the dorsum of the paw during spinal locomotion." *Brain research* no. 132 (1):121-139.
- Forssberg, H, So Grillner, and So Rossignol. 1975. "Phase dependent reflex reversal during walking in chronic spinal cats." *Brain research* no. 85 (1):103-107.
- Fouad, Karim, Vera Pedersen, Martin E Schwab, and Christian Brösamle. 2001. "Cervical sprouting of corticospinal fibers after thoracic spinal cord injury accompanies shifts in evoked motor responses." *Current Biology* no. 11 (22):1766-1770.
- Frigon, Alain, Grégory Barrière, Hugues Leblond, and Serge Rossignol. 2009. "Asymmetric changes in cutaneous reflexes after a partial spinal lesion and retention following spinalization during locomotion in the cat." *Journal of neurophysiology* no. 102 (5):2667-2680.
- Frigon, Alain, and Serge Rossignol. 2008. "Adaptive changes of the locomotor pattern and cutaneous reflexes during locomotion studied in the same cats before and after spinalization." *The Journal of physiology* no. 586 (12):2927-2945.
- Frigon, Alain, Jennifer Sirois, and Jean-Pierre Gossard. 2010. "Effects of ankle and hip muscle afferent inputs on rhythm generation during fictive locomotion." *Journal of neurophysiology* no. 103 (3):1591-1605.
- Fruitet, Joan, Alexandra Carpentier, Maureen Clerc, and Rémi Munos. 2012. Bandit Algorithms boost Brain Computer Interfaces for motor-task selection of a brain-controlled button. Paper read at Advances in Neural Information Processing Systems.
- Fruitet, Joan, Alexandra Carpentier, Rémi Munos, and Maureen Clerc. 2013. "Automatic motor task selection via a bandit algorithm for a brain-controlled button." *Journal of neural engineering* no. 10 (1):016012.
- Fung, J, and H Barbeau. 1989. "A dynamic EMG profile index to quantify muscular activation disorder in spastic paretic gait." *Electroencephalography and clinical neurophysiology* no. 73 (3):233-244.
- Fung, J., and H. Barbeau. 1994. "Effects of conditioning cutaneomuscular stimulation on the soleus H-reflex in normal and spastic paretic subjects during walking and standing." *J Neurophysiol* no. 72 (5):2090-104.
- Gage, Fred H, and Sally Temple. 2013. "Neural Stem Cells: Generating and Regenerating the Brain." *Neuron* no. 80 (3):588-601.
- Gallien, P., R. Brissot, M. Eyssette, L. Tell, M. Barat, L. Wiart, and H. Petit. 1995. "Restoration of gait by functional electrical stimulation for spinal cord injured patients." *Paraplegia* no. 33 (11):660-4. doi: 10.1038/sc.1995.138.
- Garcia-Rill, E, and RD Skinner. 1987. "The mesencephalic locomotor region. II. Projections to reticulospinal neurons." *Brain research* no. 411 (1):13-20.
- Gazula, Valeswara-Rao, Melinda Roberts, Christopher Luzzio, Abbas F. Jawad, and Robert Gordon Kalb. 2004. "Effects of limb exercise after spinal cord injury on motor neuron dendrite structure." *J Comp Neurol* no. 476 (2):130-145. doi: 10.1002/cne.20204.
- Gensert, JoAnn M, and James E Goldman. 1997. "Endogenous progenitors remyelinate demyelinated axons in the adult CNS." *Neuron* no. 19 (1):197-203.
- Gerasimenko, Y, WB McKay, FE Pollo, and MR Dimitrijevic. 1996. Stepping movements in paraplegic patients induced by epidural spinal cord stimulation. Paper read at Soc Neurosci Abstr.
- Gerasimenko, YP, and AN Makarovskiy. 1996. "Neurophysiological evaluation of the effects of spinal cord stimulation in spinal patients." *Motor Control VII*:153-157.

- Gerasimenko, Yury P, VD Avelev, OA Nikitin, and IA Lavrov. 2003a. "Initiation of locomotor activity in spinal cats by epidural stimulation of the spinal cord." *Neuroscience and behavioral physiology* no. 33 (3):247-254.
- Gerasimenko, Yury P, Ronaldo M Ichiyama, Igor A Lavrov, Gregoire Courtine, Lance Cai, Hui Zhong, Roland R Roy, and V Reggie Edgerton. 2007. "Epidural spinal cord stimulation plus quipazine administration enable stepping in complete spinal adult rats." *Journal of neurophysiology* no. 98 (5):2525-2536.
- Gerasimenko, Yury P, Igor A Lavrov, Gregoire Courtine, Ronaldo M Ichiyama, Christine J Dy, Hui Zhong, Roland R Roy, and V Reggie Edgerton. 2006. "Spinal cord reflexes induced by epidural spinal cord stimulation in normal awake rats." *Journal of neuroscience methods* no. 157 (2):253-263.
- Gerasimenko, Yury P., V. D. Avelev, O. A. Nikitin, and I. A. Lavrov. 2003b. "Initiation of locomotor activity in spinal cats by epidural stimulation of the spinal cord." *Neurosci Behav Physiol* no. 33 (3):247-54.
- Gergens, E. 1877. "Ueber gekreuzte Reflexe." *Pflügers Archiv European Journal of Physiology* no. 14 (1):340-344.
- Gibbs, John, Linda M Harrison, John A Stephens, and Andrew L Evans. 1999. "Cutaneomuscular reflex responses recorded from the lower limb in children and adolescents with cerebral palsy." *Developmental Medicine & Child Neurology* no. 41 (07):456-464.
- Golberg, David E. 1989. "Genetic algorithms in search, optimization, and machine learning." *Addion wesley* no. 1989.
- Goldberger, Michael E. 1977. "Locomotor recovery after unilateral hindlimb deafferentation in cats." *Brain research* no. 123 (1):59-74.
- Goldberger, Michael E. 1988. "Partial and complete deafferentation of cat hindlimb: the contribution of behavioral substitution to recovery of motor function." *Experimental brain research* no. 73 (2):343-353.
- Goldfarb, Donald. 1970. "A family of variable-metric methods derived by variational means." *Mathematics of computation* no. 24 (109):23-26.
- Gossard, J-P, RM Brownstone, I Barajon, and H Hultborn. 1994. "Transmission in a locomotor-related group Ib pathway from hindlimb extensor muscles in the cat." *Experimental brain research* no. 98 (2):213-228.
- Granat, M., J. F. Keating, A. C. B. Smith, M. Delargy, and B. J. Andrews. 1992. "The use of functional electrical stimulation to assist gait in patients with incomplete spinal cord injury." *Disability and Rehabilitation* no. 14 (2):93-97.
- Granat, MH, BW Heller, DJ Nicol, RH Baxendale, and BJ Andrews. 1993. "Improving limb flexion in FES gait using the flexion withdrawal response for the spinal cord injured person." *Journal of biomedical engineering* no. 15 (1):51-56.
- Granville, Vincent, Mirko Křivánek, and Jean-Paul Rasson. 1994. "Simulated annealing: A proof of convergence." *Pattern Analysis and Machine Intelligence, IEEE Transactions on* no. 16 (6):652-656.
- Grill, W. M., Jr., and J. T. Mortimer. 1996. "The effect of stimulus pulse duration on selectivity of neural stimulation." *IEEE Trans Biomed Eng* no. 43 (2):161-6. doi: 10.1109/10.481985.
- Grillner, S, and S Rossignol. 1978a. "Contralateral reflex reversal controlled by limb position in the acute spinal cat injected with clonidine iv." *Brain research* no. 144 (2):411-414.
- Grillner, S. 1981. "Control of locomotion in bipeds, tetrapods and fish." In *Handbook of Physiology. The Nervous System, Motor Control. Vol. 2.*, edited by V. B. Brooks, 1179-1236. Bethesda: American Physiological Society.

- Grillner, S., and S. Rossignol. 1978b. "On the initiation of the swing phase of locomotion in chronic spinal cats." *Brain Res* no. 146 (2):269-77.
- Gruber, Markus, Wolfgang Taube, Albert Gollhofer, Sandra Beck, Florian Amtage, and Martin Schubert. 2007. "Training-Specific Adaptations of H- and Stretch Reflexes in Human Soleus Muscle." *Journal of Motor Behavior* no. 39 (1):68-78.
- Guertin, P, MJ Angel, MC Perreault, and DA McCrea. 1995. "Ankle extensor group I afferents excite extensors throughout the hindlimb during fictive locomotion in the cat." *The Journal of Physiology* no. 487 (1):197-209.
- Gürel, Tayfun, and Carsten Mehring. 2012. "Unsupervised adaptation of brain-machine interface decoders." *Frontiers in neuroscience* no. 6.
- Hanna, Joseph P, and Jeffrey I Frank. 1995. "Automatic stepping in the pontomedullary stage of central herniation." *Neurology* no. 45 (5):985-986.
- Harkema, Susan, Yury Gerasimenko, Jonathan Hodes, Joel Burdick, Claudia Angeli, Yangsheng Chen, Christie Ferreira, Andrea Willhite, Enrico Rejc, and Robert G Grossman. 2011. "Effect of epidural stimulation of the lumbosacral spinal cord on voluntary movement, standing, and assisted stepping after motor complete paraplegia: a case study." *The Lancet* no. 377 (9781):1938-1947.
- Harris, J Donald. 1943. "Habitatory response decrement in the intact organism." *Psychological Bulletin* no. 40 (6):385.
- Hathout, Gasser M, and Roongroj Bhidayasiri. 2005. "Midbrain ataxia: an introduction to the mesencephalic locomotor region and the pedunculopontine nucleus." *American Journal of Roentgenology* no. 184 (3):953-956.
- Hazan, Elad, and Satyen Kale. 2009. On stochastic and worst-case models for investing. Paper read at Advances in Neural Information Processing Systems.
- Hazan, Elad, and Satyen Kale. 2010. "Extracting certainty from uncertainty: Regret bounded by variation in costs." *Machine learning* no. 80 (2-3):165-188.
- Hazan, Elad, and Satyen Kale. 2015. "An online portfolio selection algorithm with regret logarithmic in price variation." *Mathematical Finance* no. 25 (2):288-310.
- Heck, Christianne N, David King - Stephens, Andrew D Massey, Dileep R Nair, Barbara C Jobst, Gregory L Barkley, Vicenta Salanova, Andrew J Cole, Michael C Smith, and Ryder P Gwinn. 2014. "Two - year seizure reduction in adults with medically intractable partial onset epilepsy treated with responsive neurostimulation: Final results of the RNS System Pivotal trial." *Epilepsia* no. 55 (3):432-441.
- Helgren, Maureen E, and Michael E Goldberger. 1993. "The recovery of postural reflexes and locomotion following low thoracic hemisection in adult cats involves compensation by undamaged primary afferent pathways." *Experimental neurology* no. 123 (1):17-34.
- Herman, R, Jiping He, S D'luzansky, W Willis, and S Dilli. 2002. "Spinal cord stimulation facilitates functional walking in a chronic, incomplete spinal cord injured." *Spinal Cord* no. 40 (2):65-68.
- Herrero, Juan F, Jennifer MA Laird, and Jose A Lopez-Garcia. 2000. "Wind-up of spinal cord neurones and pain sensation: much ado about something?" *Progress in neurobiology* no. 61 (2):169-203.
- Hidler, Joseph M, and Anji E Wall. 2005. "Alterations in muscle activation patterns during robotic-assisted walking." *Clinical Biomechanics* no. 20 (2):184-193.
- Hidler, Joseph, Wessel Wisman, and Nathan Neckel. 2008. "Kinematic trajectories while walking within the Lokomat robotic gait-orthosis." *Clinical Biomechanics* no. 23 (10):1251-1259.
- Hiersemenzel, Lutz-Peter, Armin Curt, and Volker Dietz. 2000. "From spinal shock to spasticity." *Neurology* no. 54 (8):1574-1582.

- Hodgkin, Alan L, and Andrew F Huxley. 1952. "A quantitative description of membrane current and its application to conduction and excitation in nerve." *The Journal of physiology* no. 117 (4):500-544.
- Hofstoetter, Ursula S, William B McKay, Keith E Tansey, Winfried Mayr, Helmut Kern, and Karen Minassian. 2014. "Modification of spasticity by transcutaneous spinal cord stimulation in individuals with incomplete spinal cord injury." *The journal of spinal cord medicine* no. 37 (2):202-211.
- Hofstoetter, US, C Hofer, H Kern, SM Danner, W Mayr, MR Dimitrijevic, and K Minassian. 2013. "Effects of transcutaneous spinal cord stimulation on voluntary locomotor activity in an incomplete spinal cord injured individual." *Biomedical Engineering/Biomedizinische Technik*.
- Holaday, John W, and Alan I Faden. 1982. "Spinal shock and injury: experimental therapeutic approaches." *Advances in shock research* no. 10:95-98.
- Holsheimer, Jan. 1998. "Computer modelling of spinal cord stimulation and its contribution to therapeutic efficacy." *Spinal Cord* no. 36 (8):531-540.
- Holsheimer, Jan. 2002. "Which neuronal elements are activated directly by spinal cord stimulation." *Neuromodulation: Technology at the Neural Interface* no. 5 (1):25-31.
- Huang, D., T. T. Allen, W. I. Notz, and N. Zeng. 2006. "Global Optimization of Stochastic Black-Box Systems via Sequential Kriging Meta-Models." *Journal of Global Optimization* no. 34 (3):441-466.
- Hultborn, H. 2001. "State-dependent modulation of sensory feedback." *J Physiol* no. 533 (Pt 1):5-13.
- Hultborn, H., B. A. Conway, J. P. Gossard, R. Brownstone, B. Fedirchuk, E. D. Schomburg, M. Enriquez-Denton, and M. C. Perreault. 1998. "How do we approach the locomotor network in the mammalian spinal cord?" *Ann N Y Acad Sci* no. 860:70-82.
- Ichiyama, R. M., Y. P. Gerasimenko, H. Zhong, R. R. Roy, and V. R. Edgerton. 2005a. "Hindlimb stepping movements in complete spinal rats induced by epidural spinal cord stimulation." *Neuroscience letters* no. 383 (3):339-44.
- Ichiyama, RM, Yu P Gerasimenko, H Zhong, RR Roy, and VR Edgerton. 2005b. "Hindlimb stepping movements in complete spinal rats induced by epidural spinal cord stimulation." *Neuroscience letters* no. 383 (3):339-344.
- Jankowska, E., M. G. Jukes, S. Lund, and A. Lundberg. 1967. "The effect of DOPA on the spinal cord. 6. Half-centre organization of interneurons transmitting effects from the flexor reflex afferents." *Acta Physiol Scand* no. 70 (3):389-402.
- Jenner, JR, and JA Stephens. 1982. "Cutaneous reflex responses and their central nervous pathways studied in man." *The Journal of physiology* no. 333 (1):405-419.
- Jezernik, Sašo, Ruben GV Wassink, and Thierry Keller. 2004. "Sliding mode closed-loop control of FES controlling the shank movement." *Biomedical Engineering, IEEE Transactions on* no. 51 (2):263-272.
- Jones, C Allyson, and Jaynie F Yang. 1994. "Reflex behavior during walking in incomplete spinal-cord-injured subjects." *Experimental neurology* no. 128 (2):239-248.
- Jones, D. R., M. Schonlau, and W. J. Welch. 1998. "Efficient global optimization of expensive black-box functions." *Journal of Global Optimization* no. 13 (4):455-492.
- Jordan, Larry M. 1998. "Initiation of locomotion in mammals." *Annals of the New York Academy of Sciences* no. 860 (1):83-93.
- Jordan, Larry M, Carol A Pratt, and John E Menzies. 1979. "Locomotion evoked by brain stem stimulation: occurrence without phasic segmental afferent input." *Brain research* no. 177 (1):204-207.

- Kabu, Shushi, Yue Gao, Brian K Kwon, and Vinod Labhassetwar. 2015. "Drug delivery, cell-based therapies, and tissue engineering approaches for spinal cord injury." *Journal of Controlled Release*.
- Kandel, Eric R, James H Schwartz, Thomas M Jessell, Steven A Siegelbaum, and AJ Hudspeth. 2014. Principles of neural science.
- Kato, M. 1989. "Chronically isolated lumbar half spinal cord, produced by hemisection and longitudinal myelotomy, generates locomotor activities of the ipsilateral hindlimb of the cat." *Neuroscience letters* no. 98 (2):149-153.
- Kazennikov, O, V Selionov, Y Levik, and V Gurfinkel. 1997. Sensory facilitation of the vibratory evoked air-stepping in human. Paper read at Proceedings of the International Symposium on Brain and Movement. St. Petersburg, Moscow.
- Kennedy, James. 2010. "Particle swarm optimization." In *Encyclopedia of Machine Learning*, 760-766. Springer.
- Kennedy, James, and Russell Eberhart. 1995. Particle swarm optimization. Paper read at Neural Networks, 1995. Proceedings., IEEE International Conference on.
- Kennedy, Paul M, Anthony N Carlsen, J Timothy Inglis, Rudy Chow, Ian M Franks, and Romeo Chua. 2003. "Relative contributions of visual and vestibular information on the trajectory of human gait." *Experimental brain research* no. 153 (1):113-117.
- Kiefer, Jack, and Jacob Wolfowitz. 1952. "Stochastic estimation of the maximum of a regression function." *The Annals of Mathematical Statistics* no. 23 (3):462-466.
- Kiehn, O., and S. J. Butt. 2003. "Physiological, anatomical and genetic identification of CPG neurons in the developing mammalian spinal cord." *Prog Neurobiol* no. 70 (4):347-61.
- Kiehn, Ole, Ronald M. Harris-Warrick, Larry M. Jordan, Hans Hultborn, and Norio Kudo. 1998. "Preface, Neuronal Mechanisms for Generating Locomotor Activity." *Ann N Y Acad Sci* no. 860 (1):xiii-xiii.
- Kirkpatrick, Scott, C Daniel Gelatt, and Mario P Vecchi. 1983. "Optimization by simulated annealing." *science* no. 220 (4598):671-680.
- Kitano, Koichi, and David M. Kocejka. 2009. "Spinal reflex in human lower leg muscles evoked by transcutaneous spinal cord stimulation." *J Neurosci Methods* no. 180 (1):111-115. doi: 10.1016/j.jneumeth.2009.03.006.
- Knikou, Maria. 2005. "Effects of hip joint angle changes on intersegmental spinal coupling in human spinal cord injury." *Experimental Brain Research* no. 167 (3):381-393. doi: 10.1007/s00221-005-0046-6.
- Komiyama, Tomoyoshi, E Paul Zehr, and Richard B Stein. 2000. "Absence of nerve specificity in human cutaneous reflexes during standing." *Experimental brain research* no. 133 (2):267-272.
- Kossoff, Eric H, Eva K Ritzl, Jeffrey M Politsky, Anthony M Murro, Joseph R Smith, Robert B Duckrow, Dennis D Spencer, and Gregory K Bergey. 2004. "Effect of an external responsive neurostimulator on seizures and electrographic discharges during subdural electrode monitoring." *Epilepsia* no. 45 (12):1560-1567.
- Kostov, Aleksandar, Brian J Andrews, Dejan B Popovic, Richard B Stein, and William W Armstrong. 1995. "Machine learning in control of functional electrical stimulation systems for locomotion." *Biomedical Engineering, IEEE Transactions on* no. 42 (6):541-551.
- Kriellaars, D. J., R. M. Brownstone, B. R. Noga, and L. M. Jordan. 1994. "Mechanical entrainment of fictive locomotion in the decerebrate cat." *J Neurophysiol* no. 71 (6):2074-86.
- Kugelberg, E, and KE Hagbarth. 1958. "Spinal mechanism of the abdominal and erector spinae skin reflexes." *Brain: a journal of neurology* no. 81 (3):290-304.

- Kuhn, R. A. 1950. "Functional capacity of the isolated human spinal cord." *Brain* no. 73 (1):1-51.
- Kurukulasuriya, Pradeep, and Shane Rosenthal. 2013. "Climate change and agriculture: A review of impacts and adaptations."
- Kushner, Harold J. 1962. "A versatile stochastic model of a function of unknown and time varying form." *Journal of Mathematical Analysis and Applications* no. 5 (1):150-167.
- Kushner, Harold J. 1964. "A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise." *Journal of Fluids Engineering* no. 86 (1):97-106.
- LaBella, LA, JP Kehler, and DA McCrea. 1989. "A differential synaptic input to the motor nuclei of triceps surae from the caudal and lateral cutaneous sural nerves." *Journal of neurophysiology* no. 61 (2):291-301.
- LaBella, Lisa A, Andrzej Niechaj, and Serge Rossignol. 1992. "Low-threshold, short-latency cutaneous reflexes during fictive locomotion in the "semi-chronic" spinal cat." *Experimental brain research* no. 91 (2):236-248.
- Ladenbauer, Josef, Karen Minassian, Ursula S Hofstoetter, Milan R Dimitrijevic, and Frank Rattay. 2010. "Stimulation of the human lumbar spinal cord with implanted and surface electrodes: a computer simulation study." *Neural Systems and Rehabilitation Engineering, IEEE Transactions on* no. 18 (6):637-645.
- Ladouceur, M, and H Barbeau. 2000a. "Functional electrical stimulation-assisted walking for persons with incomplete spinal injuries: longitudinal changes in maximal overground walking speed." *Scandinavian journal of rehabilitation medicine* no. 32 (1):28-36.
- Ladouceur, Michel, and H Barbeau. 2000b. "Functional electrical stimulation-assisted walking for persons with incomplete spinal injuries: changes in the kinematics and physiological cost of overground walking." *Scandinavian journal of rehabilitation medicine* no. 32 (2):72-79.
- Lagerquist, Olle, and David F. Collins. 2008. "Stimulus pulse-width influences H-reflex recruitment but not Hmax/Mmax ratio." *Muscle & Nerve* no. 37 (4):483-489. doi: 10.1002/mus.20957.
- Lamb, Tania, and Jaynie F Yang. 2000. "Could different directions of infant stepping be controlled by the same locomotor central pattern generator?" *Journal of Neurophysiology* no. 83 (5):2814-2824.
- Lavrov, Igor, Grégoire Courtine, Christine J Dy, Rubia van den Brand, Andy J Fong, Yuri Gerasimenko, Hui Zhong, Roland R Roy, and V Reggie Edgerton. 2008. "Facilitation of stepping with epidural stimulation in spinal rats: role of sensory input." *The Journal of Neuroscience* no. 28 (31):7774-7780.
- LeCun, Yann, and Yoshua Bengio. 1995. "Convolutional networks for images, speech, and time series." *The handbook of brain theory and neural networks* no. 3361 (10).
- Lee, Song Joo, and Joseph Hidler. 2008. "Biomechanics of overground vs. treadmill walking in healthy individuals." *Journal of applied physiology* no. 104 (3):747-755.
- Leroux, Alain, Joyce Fung, and Hugues Barbeau. 1999. "Adaptation of the walking pattern to uphill walking in normal and spinal-cord injured subjects." *Experimental brain research* no. 126 (3):359-368.
- Lewi, Jeremy, Robert Butera, and Liam Paninski. 2008. "Sequential Optimal Design of Neurophysiology Experiments." *Neural Comput* no. 21 (3):619-687.
- Lhermitte, Jean. 1919. *La section totale de la moelle dorsale*. Bourges: Imprimerie Vve Tardy-Pigelet et fils.

- Liberson, W. T., H. J. Holmquest, D. Scot, and M. Dow. 1961. "Functional electrotherapy: stimulation of the peroneal nerve synchronized with the swing phase of the gait of hemiplegic patients." *Arch Phys Med Rehabil* no. 42:101-5.
- Lieske, SP, M Thoby-Brisson, P Telgkamp, and JM Ramirez. 2000. "Reconfiguration of the neural network controlling multiple breathing patterns: eupnea, sighs and gasps." *Nature neuroscience* no. 3 (6):600-607.
- Lin, SI, and WC Yang. 2011. "Effect of plantar desensitization on postural adjustments prior to step initiation." *Gait & posture* no. 34 (4):451-456.
- Lisin, VV, SI Frankstein, and MB Rechtman. 1973. "The influence of locomotion on flexor reflex of the hind limb in cat and man." *Experimental neurology* no. 38 (1):180-183.
- Little, J. W., and E. M. Halar. 1985. "H-reflex changes following spinal cord injury." *Arch Phys Med Rehabil* no. 66 (1):19-22.
- Little, Simon, Alex Pogosyan, Spencer Neal, Baltazar Zavala, Ludvic Zrinzo, Marwan Hariz, Thomas Foltynie, Patricia Limousin, Keyoumars Ashkan, and James FitzGerald. 2013. "Adaptive deep brain stimulation in advanced Parkinson disease." *Ann Neurol* no. 74 (3):449-457.
- Lloyd, David PC. 1949. "Post-tetanic potentiation of response in monosynaptic reflex pathways of the spinal cord." *The Journal of general physiology* no. 33 (2):147-170.
- Loeb, GE. 1993. "The distal hindlimb musculature of the cat: interanimal variability of locomotor activity and cutaneous reflexes." *Experimental brain research* no. 96 (1):125-140.
- Lombard, Warren Plympton. 1887. "The variations of the normal knee-jerk, and their relation to the activity of the central nervous system." *The American Journal of Psychology* no. 1 (1):5-71.
- Lundberg, A. 1979. Multisensory control of spinal reflex pathways. Paper read at Reflex Control of Posture and Movement: Proceedings of IBRO Symposium Held in Pisa, Italy, on September 11-14, 1978.
- Luys, J.B. 1893. *Le cerveau et ses fonctions*: F. Alcan.
- Lynch, Cheryl L, and Milos R Popovic. 2008. "Functional electrical stimulation." *Control Systems, IEEE* no. 28 (2):40-50.
- Macbeth, Guillermo, Eugenia Razumiejczyk, and Rubén Daniel Ledesma. 2011. "Cliff's Delta Calculator: A non-parametric effect size program for two groups of observations." *Universitas Psychologica* no. 10 (2):545-555.
- Maffiuletti, N. A., A. Martin, N. Babault, M. Pensini, B. Lucas, and M. Schieppati. 2001. "Electrical and mechanical H(max)-to-M(max) ratio in power- and endurance-trained athletes." *J Appl Physiol* no. 90 (1):3-9.
- Magnus, R. 1909a. "Zur regelung der bewegungen durch das zentralnervensystem." *Pflügers Archiv European Journal of Physiology* no. 130 (5):219-252.
- Magnus, R. 1909b. "Zur Regelung der Bewegungen durch das Zentralnervensystem." *Archiv für die gesamte Physiologie des Menschen und der Tiere* no. 130 (5-6):253-269. doi: 10.1007/BF01677966.
- Mansfield, Avril, and Gerard M. Lyons. 2003. "The use of accelerometry to detect heel contact events for use as a sensor in FES assisted walking." *Medical Engineering & Physics* no. 25 (10):879-885. doi: 10.1016/s1350-4533(03)00116-4.
- Marchetti, C., M. Beato, and A. Nistri. 2001. "Alternating rhythmic activity induced by dorsal root stimulation in the neonatal rat spinal cord in vitro." *J Physiol* no. 530 (Pt 1):105-12.
- Marple - Horvat, DE, and DM Armstrong. 1999. "Central regulation of motor cortex neuronal responses to forelimb nerve inputs during precision walking in the cat." *The Journal of physiology* no. 519 (1):279-299.

- Maryak, John L, and Daniel C Chin. 2001. Global random optimization by simultaneous perturbation stochastic approximation. Paper read at American Control Conference, 2001. Proceedings of the 2001.
- Matheron, Georges. 1963. "Principles of geostatistics." *Economic geology* no. 58 (8):1246-1266.
- Maynard, Frederick M, Michael B Bracken, GJFD Creasey, JF Ditunno, William H Donovan, Thomas B Ducker, Susan L Garber, Ralph J Marino, Samuel L Stover, and Charles H Tator. 1997. "International standards for neurological and functional classification of spinal cord injury." *Spinal cord* no. 35 (5):266-274.
- McCall, Julianne, Norbert Weidner, and Armin Blesch. 2012. "Neurotrophic factors in combinatorial approaches for spinal cord regeneration." *Cell and tissue research* no. 349 (1):27-37.
- McCrea, D. A. 1998. "Neuronal basis of afferent-evoked enhancement of locomotor activity." *Ann N Y Acad Sci* no. 860:216-25.
- McCrea, D. A. 2001. "Spinal circuitry of sensorimotor control of locomotion." *J Physiol* no. 533 (Pt 1):41-50.
- McCrea, David A, and Ilya A Rybak. 2007. "Modeling the mammalian locomotor CPG: insights from mistakes and perturbations." *Progress in brain research* no. 165:235-253.
- McCrea, David A, and Ilya A Rybak. 2008. "Organization of mammalian locomotor rhythm and pattern generation." *Brain research reviews* no. 57 (1):134-146.
- McGlone, Francis, and David Reilly. 2010. "The cutaneous sensory system." *Neuroscience & Biobehavioral Reviews* no. 34 (2):148-159.
- McKay, Michael D, Richard J Beckman, and William J Conover. 1979. "Comparison of three methods for selecting values of input variables in the analysis of output from a computer code." *Technometrics* no. 21 (2):239-245.
- McKinnon, Ken IM. 1998. "Convergence of the Nelder--Mead Simplex Method to a Nonstationary Point." *SIAM Journal on Optimization* no. 9 (1):148-158.
- Mendell, L. M., and P. D. Wall. 1965. "Responses of Single Dorsal Cord Cells to Peripheral Cutaneous Unmyelinated Fibres." *Nature* no. 206:97-9.
- Micera, S., J. DiGiovanna, A. Berthoz, A. Demosthenous, J. P. Guyot, K. P. Hoffmann, D. Merfeld, and M. Morari. 2010. A closed-loop neural prosthesis for vestibular disorders. Paper read at Neural Network Applications in Electrical Engineering (NEUREL), 2010 10th Symposium on, 23-25 Sept. 2010.
- Minassian, Karen, Bernhard Jilge, Frank Rattay, MM Pinter, H Binder, Franz Gerstenbrand, and Milan R Dimitrijevic. 2004. "Stepping-like movements in humans with complete spinal cord injury induced by epidural stimulation of the lumbar cord: electromyographic study of compound muscle action potentials." *Spinal cord* no. 42 (7):401-416.
- Minassian, Karen, Ilse Persy, Frank Rattay, Milan R Dimitrijevic, Christian Hofer, and Helmut Kern. 2007. "Posterior root-muscle reflexes elicited by transcutaneous stimulation of the human lumbosacral cord." *Muscle & nerve* no. 35 (3):327-336.
- Minassian, Karen, Ilse Persy, Frank Rattay, Michaela M Pinter, Helmut Kern, and Milan R Dimitrijevic. 2007. "Human lumbar cord circuitries can be activated by extrinsic tonic input to generate locomotor-like activity." *Human movement science* no. 26 (2):275-295.
- Mitchell, Melanie. 1998. *An introduction to genetic algorithms*: MIT press.
- Mockus, Jonas, Vytautas Tiesis, and Antanas Zilinskas. 1978. "The application of Bayesian methods for seeking the extremum." *Towards Global Optimization* no. 2 (117-129):2.
- Morales-Enciso, Sergio, and Juergen Branke. 2015. "Tracking global optima in dynamic environments with efficient global optimization." *European Journal of Operational Research* no. 242 (3):744-755.

- Morrell, Martha J. 2011. "Responsive cortical stimulation for the treatment of medically intractable partial epilepsy." *Neurology* no. 77 (13):1295-1304.
- Müller, Joh. 1843. *Elements of physiology*: Lea and Blanchard.
- Nakajima, Tsuyoshi, Kiyotaka Kamibayashi, and Kimitaka Nakazawa. 2012. "Somatosensory control of spinal reflex circuitry during robotic-assisted stepping." *The Journal of Physical Fitness and Sports Medicine* no. 1 (4):665-670.
- Nakajima, Tsuyoshi, Masanori Sakamoto, Takashi Endoh, and Tomoyoshi Komiyama. 2006. "Location-specific and task-dependent modulation of cutaneous reflexes in intrinsic human hand muscles." *Clinical neurophysiology* no. 117 (2):420-429.
- Nakajima, Tsuyoshi, Masanori Sakamoto, Toshiki Tazoe, Takashi Endoh, and Tomoyoshi Komiyama. 2006. "Location specificity of plantar cutaneous reflexes involving lower limb muscles in humans." *Experimental brain research* no. 175 (3):514-525.
- Nakajima, Tsuyoshi, Masanori Sakamoto, Toshiki Tazoe, Takashi Endoh, and Tomoyoshi Komiyama. 2009. "Location-specific modulations of plantar cutaneous reflexes in human (peroneus longus muscle) are dependent on co-activation of ankle muscles." *Experimental brain research* no. 195 (3):403-412.
- Nash, Mark S, Patrick L Jacobs, Berta M Montalvo, K John Klose, Rosalind S Guest, and Belinda M Needham-Shropshire. 1997. "Evaluation of a training program for persons with SCI paraplegia using the Parastep® 1 ambulation system: Part 5. Lower extremity blood flow and hyperemic responses to occlusion are augmented by ambulation training." *Archives of physical medicine and rehabilitation* no. 78 (8):808-814.
- Nekoukar, Vahab, and Abbas Erfanian. 2011. "Adaptive fuzzy terminal sliding mode control for a class of MIMO uncertain nonlinear systems." *Fuzzy Sets and Systems* no. 179 (1):34-49.
- Nelder, John A, and Roger Mead. 1965. "A simplex method for function minimization " *Computer journal* no. 7 (4):6.
- Nielsen, J, Clarissa Crone, and Hans Hultborn. 1993. "H-reflexes are smaller in dancers from The Royal Danish Ballet than in well-trained athletes." *European Journal of Applied Physiology and Occupational Physiology* no. 66 (2):116-121.
- Noreau, Luc, Pierre Proulx, Lisette Gagnon, Mélanie Drolet, and Marie-Thérèse Laramée. 2000. "Secondary Impairments After Spinal Cord Injury: A Population-Based Study." *American Journal of Physical Medicine & Rehabilitation* no. 79 (6):526-535.
- Nymark, Jennifer R, Suzanne J Balmer, Ellen H Melis, Edward D Lemaire, and Shawn Millar. 2005. "Electromyographic and kinematic nondisabled gait differences at extremely slow overground and treadmill walking speeds." *Journal of rehabilitation research and development* no. 42 (4):523.
- O'Connor, K. N., C. I. Petkov, and M. L. Sutter. 2005. "Adaptive stimulus optimization for auditory cortical neurons." *J Neurophysiol* no. 94 (6):4051-67.
- Osorio, Ivan, Mark G Frei, Sridhar Sunderam, Jonathon Giftakis, Naresh C Bhavaraju, Scott F Schaffner, and Steven B Wilkinson. 2005. "Automated seizure abatement in humans using electrical stimulation." *Annals of neurology* no. 57 (2):258-268.
- Paillard, Jacques. 1955. *Réflexes et régulations d'origine proprioceptive chez l'homme: étude neurophysiologique et psychophysique*, Arnette Paris.
- Panek, Izabela, Tuan Bui, Asher TB Wright, and Robert M Brownstone. 2014. "Cutaneous afferent regulation of motor function." *Acta Neurobiol Exp* no. 74:158-171.
- Pang, Marco YC, Tania Lam, and Jaynie F Yang. 2003. "Infants adapt their stepping to repeated trip-inducing stimuli." *Journal of neurophysiology* no. 90 (4):2731-2740.

- Pang, Marco YC, and Jaynie F Yang. 2000. "The initiation of the swing phase in human infant stepping: importance of hip position and leg loading." *The Journal of physiology* no. 528 (2):389-404.
- Pang, Marco YC, and Jaynie F Yang. 2001. "Interlimb co-ordination in human infant stepping." *The Journal of Physiology* no. 533 (2):617-625.
- Pang, Marco YC, and Jaynie F Yang. 2002. "Sensory gating for the initiation of the swing phase in different directions of human infant stepping." *The Journal of neuroscience* no. 22 (13):5734-5740.
- Panizza, M., J. Nilsson, and M. Hallett. 1989. "Optimal stimulus duration for the H reflex." *Muscle Nerve* no. 12 (7):576-9. doi: 10.1002/mus.880120708.
- Panizza, Marcela, Jan Nilsson, Bradley J Roth, Peter J Basser, and Mark Hallett. 1992. "Relevance of stimulus duration for activation of motor and sensory fibers: implications for the study of H-reflexes and magnetic stimulation." *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section* no. 85 (1):22-29.
- Pappas, I. P., M. R. Popovic, T. Keller, V. Dietz, and M. Morari. 2001. "A reliable gait phase detection system." *IEEE Trans Neural Syst Rehabil Eng* no. 9 (2):113-25. doi: 10.1109/7333.928571.
- Park, Eugene, Alexander A Velumian, and Michael G Fehlings. 2004. "The role of excitotoxicity in secondary mechanisms of spinal cord injury: a review with an emphasis on the implications for white matter degeneration." *Journal of neurotrauma* no. 21 (6):754-774.
- Parker, John L, Dean M Karantonis, Peter S Single, Milan Obradovic, and Michael J Cousins. 2012. "Compound action potentials recorded in the human spinal cord during neurostimulation for pain relief." *Pain* no. 153 (3):593-601.
- Parsopoulos, KE, and MN Vrahatis. 2001. "Particle swarm optimizer in noisy and continuously changing environments." *methods* no. 5 (6):23.
- Pearson, Keir G. 1995. "Proprioceptive regulation of locomotion." *Current opinion in neurobiology* no. 5 (6):786-791.
- Pearson, KG. 2008. "Role of sensory feedback in the control of stance duration in walking cats." *Brain research reviews* no. 57 (1):222-227.
- Pearson, KG, and DF Collins. 1993. "Reversal of the influence of group Ib afferents from plantaris on activity in medial gastrocnemius muscle during locomotor activity." *Journal of Neurophysiology* no. 70 (3):1009-1017.
- Peckham, P. H., and J. S. Knutson. 2005. "Functional electrical stimulation for neuromuscular applications." *Annu Rev Biomed Eng* no. 7:327-60.
- Pedersen, Ejner. 1954. "Studies on the central pathway of the flexion reflex in man and animal." *Acta psychiatrica et neurologica Scandinavica. Supplementum* no. 88:1.
- Perreault, MC, MJ Angel, P Guertin, and DA McCrea. 1995. "Effects of stimulation of hindlimb flexor group II afferents during fictive locomotion in the cat." *The Journal of physiology* no. 487 (Pt 1):211-220.
- Perry, Catherine E. 1967. "Principles and Techniques of the Brunnstrom Approach to the Treatment of Hemiplegia." *American Journal of Physical Medicine & Rehabilitation* no. 46 (1):789-812.
- Peters, Thomas E, Naresh C Bhavaraju, Mark G Frei, and Ivan Osorio. 2001. "Network system for automated seizure detection and contingent delivery of therapy." *Journal of clinical neurophysiology* no. 18 (6):545-549.

- Picheny, Victor, Tobias Wagner, and David Ginsbourger. 2013. "A benchmark of kriging-based infill criteria for noisy optimization." *Structural and Multidisciplinary Optimization* no. 48 (3):607-626.
- Poirrier, Anne-Lise, Yves Nyssen, Felix Scholtes, Sylvie Multon, Charline Rinkin, Géraldine Weber, Delphine Bouhy, Gary Brook, Rachelle Franzen, and Jean Schoenen. 2004. "Repetitive transcranial magnetic stimulation improves open field locomotor recovery after low but not high thoracic spinal cord compression-injury in adult rats." *Journal of Neuroscience Research* no. 75 (2):253-261.
- Poli, Riccardo, James Kennedy, and Tim Blackwell. 2007. "Particle swarm optimization." *Swarm intelligence* no. 1 (1):33-57.
- Polikov, Vadim S, Patrick A Tresco, and William M Reichert. 2005. "Response of brain tissue to chronically implanted neural electrodes." *Journal of neuroscience methods* no. 148 (1):1-18.
- Powell, Michael JD. 2006. "The NEWUOA software for unconstrained optimization without derivatives." In *Large-scale nonlinear optimization*, 255-297. Springer.
- Priestley, Maurice Bertram. 1981. "Spectral analysis and time series."
- Prinz, Astrid A, LF Abbott, and Eve Marder. 2004. "The dynamic clamp comes of age." *Trends in neurosciences* no. 27 (4):218-224.
- Prochazka, A, K-H Sontag, and P Wand. 1978. "Motor reactions to perturbations of gait: proprioceptive and somesthetic involvement." *Neuroscience letters* no. 7 (1):35-39.
- Prosser, C Ladd, and Walter S Hunter. 1936. "The extinction of startle responses and spinal reflexes in the white rat." *American Journal of Physiology*.
- Pugh, Jim, Yizhen Zhang, and Alcherio Martinoli. 2005. Particle swarm optimization for unsupervised robotic learning. Paper read at Swarm Intelligence Symposium.
- Quevedo, Jorge, Katinka Stecina, and David A McCrea. 2005. "Intracellular analysis of reflex pathways underlying the stumbling corrective reaction during fictive locomotion in the cat." *Journal of neurophysiology* no. 94 (3):2053-2062.
- Ragnarsson, K. T. 2008. "Functional electrical stimulation after spinal cord injury: current use, therapeutic effects and future directions." *Spinal Cord* no. 46 (4):255-74. doi: 10.1038/sj.sc.3102091.
- Raineteau, Olivier, and Martin E Schwab. 2001. "Plasticity of motor systems after incomplete spinal cord injury." *Nature Reviews Neuroscience* no. 2 (4):263-273.
- Randic, M, MC Jiang, and R Cerne. 1993. "Long-term potentiation and long-term depression of primary afferent neurotransmission in the rat spinal cord." *The Journal of neuroscience* no. 13 (12):5228-5241.
- Rattay, F., K. Minassian, and M. R. Dimitrijevic. 2000. "Epidural electrical stimulation of posterior structures of the human lumbosacral cord: 2. quantitative analysis by computer modeling." *Spinal Cord* no. 38 (8):473-489.
- Rayburn, William F. 1995. "Fetal movement monitoring." *Clinical obstetrics and gynecology* no. 38 (1):59-67.
- Roberts, A., S. R. Soffe, E. S. Wolf, M. Yoshida, and F. Y. Zhao. 1998. "Central circuits controlling locomotion in young frog tadpoles." *Ann N Y Acad Sci* no. 860:19-34.
- Roby-Brami, A., and B. Bussel. 1987. "Long-latency spinal reflex in man after flexor reflex afferent stimulation." *Brain* no. 110 (Pt 3):707-25.
- Rosenfeld, JE, AM Sherwood, JA Halter, and MR Dimitrijevic. 1995. Evidence of a pattern generator in paralyzed subject with spinal cord stimulation. Paper read at Soc Neurosci Abstr.

- Rosin, Boris, Maya Slovik, Rea Mitelman, Michal Rivlin-Etzion, Suzanne N Haber, Zvi Israel, Eilon Vaadia, and Hagai Bergman. 2011. "Closed-loop deep brain stimulation is superior in ameliorating parkinsonism." *Neuron* no. 72 (2):370-384.
- Rossi, A, A Zalaffi, and B Decchi. 1996. "Interaction of nociceptive and non-nociceptive cutaneous afferents from foot sole in common reflex pathways to tibialis anterior motoneurons in humans." *Brain research* no. 714 (1):76-86.
- Rossignol, S, C Chau, N Giroux, E Brustein, L Bouyer, J Marcoux, C Langlet, D Barthélemy, J Provencher, and H Leblond. 2002. "The cat model of spinal injury." *Progress in brain research* no. 137:151.
- Rossignol, S, T Drew, E Brustein, and W Jiang. 1999. "Locomotor performance and adaptation after partial or complete spinal cord lesions in the cat." *Progress in brain research* no. 123:349-368.
- Rossignol, S, and L Gauthier. 1980. "An analysis of mechanisms controlling the reversal of crossed spinal reflexes." *Brain research* no. 182 (1):31-45.
- Rossignol, S, C Julien, and L Gauthier. 1981. "Stimulus-response relationships during locomotion." *Canadian journal of physiology and pharmacology* no. 59 (7):667-674.
- Rossignol, S., C. Chau, E. Brustein, M. Belanger, H. Barbeau, and T. Drew. 1996. "Locomotor capacities after complete and partial lesions of the spinal cord." *Acta Neurobiol Exp (Wars)* no. 56 (1):449-63.
- Rossignol, Serge, Réjean Dubuc, and Jean-Pierre Gossard. 2006. "Dynamic sensorimotor interactions in locomotion." *Physiological reviews* no. 86 (1):89-154.
- Rossignol, Serge, and Alain Frigon. 2011. "Recovery of locomotion after spinal cord injury: some facts and mechanisms." *Annual review of neuroscience* no. 34:413-440.
- Ryzhov, Ilya O, Warren B Powell, and Peter I Frazier. 2012. "The knowledge gradient algorithm for a general class of online learning problems." *Operations Research* no. 60 (1):180-195.
- Sacks, Jerome, William Welch, Toby Mitchell, and Henry Wynn. 1989. "Design and Analysis of Computer Experiments." *Statistical Science* no. 4 (4):409-423. doi: 10.2307/2245858.
- Sacks, Jerome, and Donald Ylvisaker. 1970. "Statistical designs and integral approximation." *Proc. 12th Bienn. Semin. Can. Math. Congr.(R. Pyke, ed.):*115-136.
- Saifuddin, Asif, Sarah JD Burnett, and John White. 1998. "The variation of position of the conus medullaris in an adult population: a magnetic resonance imaging study." *Spine* no. 23 (13):1452-1456.
- Santaniello, Sabato, Giovanni Fiengo, Luigi Glielmo, and Warren M Grill. 2011. "Closed-loop control of deep brain stimulation: a simulation study." *Neural Systems and Rehabilitation Engineering, IEEE Transactions on* no. 19 (1):15-24.
- Sasena, Michael J, Panos Papalambros, and Pierre Goovaerts. 2002. "Exploration of metamodeling sampling criteria for constrained global optimization." *Engineering optimization* no. 34 (3):263-278.
- Schade, CM, DM Schultz, N Tamayo, S Iyer, and E Panken. 2011. "Automatic adaptation of neurostimulation therapy in response to changes in patient position: results of the Posture Responsive Spinal Cord Stimulation (PRS) Research Study." *Pain physician* no. 14 (5):407.
- Schieppati, Marco. 1987. "The Hoffmann reflex: a means of assessing spinal reflex excitability and its descending control in man." *Progress in neurobiology* no. 28 (4):345-376.
- Schillings, AM, BMH Van Wezel, TH Mulder, and Jaak Duysens. 2000. "Muscular responses and movement strategies during stumbling over obstacles." *Journal of Neurophysiology* no. 83 (4):2093-2102.

- Schindler-Ivens, Sheila, and Richard K Shields. 2000. "Low frequency depression of H-reflexes in humans with acute and chronic spinal-cord injury." *Experimental brain research* no. 133 (2):233-241.
- Schomburg, ED, N Petersen, I Barajon, and Hans Hultborn. 1998. "Flexor reflex afferents reset the step cycle during fictive locomotion in the cat." *Experimental brain research* no. 122 (3):339-350.
- Schouenborg, Jens. 2002. "Modular organisation and spinal somatosensory imprinting." *Brain research reviews* no. 40 (1):80-91.
- Schouenborg, Jens, Han-Rong Weng, and Hans Holmberg. 1994. "Modular organization of spinal nociceptive reflexes: a new hypothesis." *News in Physiological Sciences* no. 9:261-264.
- Schultz, David M, Lynn Webster, Peter Kosek, Urfan Dar, Ye Tan, and Mark Sun. 2012. "Sensor-driven position-adaptive spinal cord stimulation for chronic pain." *Pain Physician* no. 15 (1):1-12.
- Schurch, B., M. Stöhrer, G. Kramer, D. M. Schmid, G. Gaul, and D. Hauri. 2000. "Bolinum-A Toxin for Treating Detrusor Hyperreflexia in Spinal Cord Injured Patients: A New Alternative to Anticholinergic Drugs? Preliminary Results." *J Urol* no. 164 (3):692-697. doi: 10.1016/s0022-5347(05)67283-7.
- Schuster, Eugene F. 1985. "Incorporating support constraints into nonparametric estimators of densities." *Communications in Statistics-Theory and methods* no. 14 (5):1123-1136.
- Schwab, Martin E, and Deborah Bartholdi. 1996. "Degeneration and regeneration of axons in the lesioned spinal cord." *Physiological reviews* no. 76 (2):319-370.
- Selionov, VA, OV Kazennikov, YS Levik, and VS Gurfinkel. 1997. Experimental investigation of locomotor-like movements elicited by vibration in human. Paper read at Proceedings of the International Symposium on Brain and Movement, St. Petersburg, Moscow.
- Shahani, Bhagwan T, and Robert R Young. 1971. "Human flexor reflexes." *Journal of Neurology, Neurosurgery & Psychiatry* no. 34 (5):616-627.
- Shanno, David F. 1970. "Conditioning of quasi-Newton methods for function minimization." *Mathematics of computation* no. 24 (111):647-656.
- Shapkova, EY, and ED Schomburg. 2000. "Two types of motor modulation underlying human stepping evoked by spinal cord electrical stimulation (SCES)." *Acta physiologica et pharmacologica Bulgarica* no. 26 (3):155-157.
- Shapkova, HY, Yu T Shapkov, and A Yu Mushkin. 1997. Locomotor activity in paraplegic children induced by spinal cord stimulation. Paper read at Proceedings of the International Symposium on Brain and Movement. St. Petersburg, Moscow.
- Shealy, C Norman, J Thomas Mortimer, and James B Reswick. 1967. "Electrical inhibition of pain by stimulation of the dorsal columns: preliminary clinical report." *Anesthesia & Analgesia* no. 46 (4):489-491.
- Sherk, Helen, and Garth A Fowler. 2001. "Neural analysis of visual information during locomotion." *Progress in brain research* no. 134:247-264.
- Sherrington, Charles Scott. 1906. *The integrative action of the nervous system*. New York,: C. Scribner's sons.
- Sherrington, Charles Scott. 1910. "Flexion - reflex of the limb, crossed extension - reflex, and reflex stepping and standing." *The Journal of physiology* no. 40 (1-2):28-121.
- Shi, Yuhui, and Russell Eberhart. 1998. A modified particle swarm optimizer. Paper read at Evolutionary Computation Proceedings, 1998. IEEE World Congress on Computational Intelligence., The 1998 IEEE International Conference on.
- Shik, ML, FV Severin, and ORLOVSKI. GN. 1966. "Control of walking and running by means of electrical stimulation of mid-brain." *BIOPHYSICS-USSR* no. 11 (4):756-&.

- Sillar, K. T., and A. Roberts. 1992. "Phase-dependent Modulation of a Cutaneous Sensory Pathway by Glycinergic Inhibition from the Locomotor Rhythm Generator in *Xenopus* Embryos." *Eur J Neurosci* no. 4 (11):1022-1034.
- Silverman, Bernard W. 1986. *Density estimation for statistics and data analysis*. Vol. 26: CRC press.
- Simard, Patrice Y, Dave Steinkraus, and John C Platt. 2003. Best practices for convolutional neural networks applied to visual document analysis. Paper read at null.
- Sinkjaer, Thomas, Morten Haugland, Andreas Inmann, Morten Hansen, and Kim D. Nielsen. 2003. "Biopotentials as command and feedback signals in functional electrical stimulation systems." *Medical Engineering & Physics* no. 25 (1):29-40. doi: 10.1016/s1350-4533(02)00178-9.
- Spall, James C. 1992. "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation." *Automatic Control, IEEE Transactions on* no. 37 (3):332-341.
- Spall, James C. 2005. *Introduction to stochastic search and optimization: estimation, simulation, and control*. Vol. 65: John Wiley & Sons.
- Srinivas, Niranjan, Andreas Krause, Sham M Kakade, and Matthias Seeger. 2009. "Gaussian process optimization in the bandit setting: No regret and experimental design." *arXiv preprint arXiv:0912.3995*.
- Stanic, Uros, and Amadej Trnkoczy. 1974. "Closed-loop positioning of hemiplegic patient's joint by means of functional electrical stimulation." *Biomedical Engineering, IEEE Transactions on* (5):365-370.
- Steeves, John D, and Larry M Jordan. 1980. "Localization of a descending pathway in the spinal cord which is necessary for controlled treadmill locomotion." *Neuroscience letters* no. 20 (3):283-288.
- Storn, Rainer, and Kenneth Price. 1997. "Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces." *Journal of global optimization* no. 11 (4):341-359.
- Sun, Felice T, and Martha J Morrell. 2014. "Closed-loop neurostimulation: the clinical experience." *Neurotherapeutics* no. 11 (3):553-563.
- Tator, CH. 1990. "Review of experimental spinal cord injury with emphasis on the local and systemic circulatory effects." *Neuro-Chirurgie* no. 37 (5):291-302.
- Tax, AA, BMH Van Wezel, and V Dietz. 1995. "Bipedal reflex coordination to tactile stimulation of the sural nerve during human running." *Journal of neurophysiology* no. 73 (5):1947-1964.
- Thornton, Judith M, Tipu Aziz, David Schlugman, and David J Paterson. 2002. "Electrical stimulation of the midbrain increases heart rate and arterial blood pressure in awake humans." *The Journal of physiology* no. 539 (2):615-621.
- Thrasher, Adam, Geoffrey M Graham, and Milos R Popovic. 2005. "Reducing muscle fatigue due to functional electrical stimulation using random modulation of stimulation parameters." *Artificial organs* no. 29 (6):453-458.
- Toal, David John James, and Andy J Keane. 2012. "Non-stationary kriging for design optimization." *Engineering Optimization* no. 44 (6):741-765.
- Toorring, J., E. Pedersen, and B. Klemar. 1981. "Standardisation of the electrical elicitation of the human flexor reflex." *J Neurol Neurosurg Psychiatry* no. 44 (2):129-32.
- Trefethen, Lloyd N, and David Bau III. 1997. *Numerical linear algebra*. Vol. 50: Siam.
- Van de Crommert, Henry WAA, Theo Mulder, and Jacques Duysens. 1998. "Neural control of locomotion: sensory control of the central pattern generator and its relation to treadmill training." *Gait & posture* no. 7 (3):251-263.

- van Hedel, Hubertus J, Markus Wirz, and Volker Dietz. 2005. "Assessing walking ability in subjects with spinal cord injury: validity and reliability of 3 walking tests." *Archives of physical medicine and rehabilitation* no. 86 (2):190-196.
- Van Wezel, Bart MH, Frans AM Ottenhoff, and Jacques Duysens. 1997. "Dynamic control of location-specific information in tactile cutaneous reflexes from the foot during human walking." *The Journal of neuroscience* no. 17 (10):3804-3814.
- VanHiel, L., White, J., Braswell-Christy, J., Ford, M., Winchester, P., Bickel, C.S., Tansey, K., . 2012. Feasibility of Using a Robotic Assessment of Spasticity after Spinal Cord Injury. Paper read at American Society for Neurorehabilitation.
- Varejao, Artur SP, and Vitor M Filipe. 2007. "Contribution of cutaneous inputs from the hindpaw to the control of locomotion in rats." *Behavioural brain research* no. 176 (2):193-201.
- Vazquez, Emmanuel, and Julien Bect. 2010. "Convergence properties of the expected improvement algorithm with fixed mean and covariance functions." *Journal of Statistical Planning and inference* no. 140 (11):3088-3095.
- Veale, JL, RF Mark, and Sandra Rees. 1973. "Differential sensitivity of motor and sensory fibres in human ulnar nerve." *Journal of Neurology, Neurosurgery & Psychiatry* no. 36 (1):75-86.
- Veneman, Jan F, Rik Kruidhof, Edsko EG Hekman, Ralf Ekkelenkamp, Edwin HF Van Asseldonk, and Herman Van Der Kooij. 2007. "Design and evaluation of the LOPES exoskeleton robot for interactive gait rehabilitation." *Neural Systems and Rehabilitation Engineering, IEEE Transactions on* no. 15 (3):379-386.
- Venkatraman, S., K. Elkabany, J. D. Long, Yao Yimin, and J. M. Carmena. 2009. "A System for Neural Recording and Closed-Loop Intracortical Microstimulation in Awake Rodents." *Biomedical Engineering, IEEE Transactions on* no. 56 (1):15-22. doi: 10.1109/tbme.2008.2005944.
- Visintin, M, and H Barbeau. 1989. "The effects of body weight support on the locomotor pattern of spastic paretic patients." *The Canadian journal of neurological sciences. Le journal canadien des sciences neurologiques* no. 16 (3):315-325.
- Visintin, Martha, Hugues Barbeau, Nicol Korner-Bitensky, and Nancy E Mayo. 1998. "A new approach to retrain gait in stroke patients through body weight support and treadmill stimulation." *Stroke* no. 29 (6):1122-1128.
- Vodovnik, L, WJ Crochetiere, and JB Reswick. 1967. "Control of a skeletal joint by electrical stimulation of antagonists." *Medical and biological engineering* no. 5 (2):97-109.
- Von Euler, C. 1983. "On the central pattern generator for the basic breathing rhythmicity." *Journal of Applied Physiology* no. 55 (6):1647-1659.
- Weidner, Norbert, Arvin Ner, Nima Salimi, and Mark H Tuszynski. 2001. "Spontaneous corticospinal axonal plasticity and functional recovery after adult central nervous system injury." *Proceedings of the National Academy of Sciences* no. 98 (6):3513-3518.
- Wernig, A, A Nanassy, and S Müller. 1998. "Maintenance of locomotor abilities following Laufband(treadmill) therapy in para- and tetraplegic persons: follow-up studies." *Spinal Cord* no. 36 (11):744-749.
- Wernig, Anton, and S Müller. 1992. "Laufband locomotion with body weight support improved walking in persons with severe spinal cord injuries." *Spinal Cord* no. 30 (4):229-238.
- Whelan, PJ, and KG Pearson. 1997. "Comparison of the effects of stimulating extensor group I afferents on cycle period during walking in conscious and decerebrate cats." *Experimental brain research* no. 117 (3):444-452.
- Winter, D. A., and H. J. Yack. 1987. "EMG profiles during normal human walking: stride-to-stride and inter-subject variability." *Electroencephalogr Clin Neurophysiol* no. 67 (5):402-411. doi: 10.1016/0013-4694(87)90003-4.

- Wolpaw, J. R., and A. M. Tennissen. 2001. "Activity-dependent spinal cord plasticity in health and disease." *Annual review of neuroscience* no. 24:807-43.
- Woods, Daniel John. 1985. "An interactive approach for solving multi-objective optimization problems (interactive computer, nelder-mead simplex algorithm, graphics)."
- Woolacott, Adam, and John Burne. 2006. "The tonic stretch reflex and spastic hypertonia after spinal cord injury." *Experimental Brain Research* no. 174 (2):386-396. doi: 10.1007/s00221-006-0478-7.
- Woolf, Clifford J. 2011. "Central sensitization: implications for the diagnosis and treatment of pain." *Pain* no. 152 (3):S2-S15.
- Wu, Ming, Keith Gordon, Jennifer H Kahn, and Brian D Schmit. 2011. "Prolonged electrical stimulation over hip flexors increases locomotor output in human SCI." *Clinical Neurophysiology* no. 122 (7):1421-1428.
- Wyndaele, Michel, and Jean-Jacques Wyndaele. 2006. "Incidence, prevalence and epidemiology of spinal cord injury: what learns a worldwide literature survey?" *Spinal cord* no. 44 (9):523-529.
- Xiong, Qiang, and Arthur Jutan. 2003. "Continuous optimization using a dynamic simplex method." *Chemical Engineering Science* no. 58 (16):3817-3828.
- Xiong, Ying, Wei Chen, Daniel Apley, and Xuru Ding. 2007. "A non - stationary covariance - based Kriging method for metamodelling in engineering design." *International Journal for Numerical Methods in Engineering* no. 71 (6):733-756.
- Yang, Jaynie F, and Richard B Stein. 1990. "Phase-dependent reflex reversal in human leg muscles during walking." *Journal of Neurophysiology* no. 63 (5):1109-1117.
- Yang, Jaynie F, Marilee J Stephens, and Rosie Vishram. 1998. "Transient disturbances to one limb produce coordinated, bilateral responses during infant stepping." *Journal of neurophysiology* no. 79 (5):2329-2337.
- Yang, JF, RB Stein, and KB James. 1991. "Contribution of peripheral afferents to the activation of the soleus muscle during walking in humans." *Experimental Brain Research* no. 87 (3):679-687.
- Young, R. R. 1994. "Spasticity: a review." *Neurology* no. 44 (11 Suppl 9):S12-20.
- Zehr, E Paul, Jaclyn E Balter, Daniel P Ferris, Sandra R Hundza, Pamela M Loadman, and Rebecca H Stoloff. 2007. "Neural regulation of rhythmic arm and leg movement is conserved across human locomotor tasks." *The Journal of physiology* no. 582 (1):209-227.
- Zehr, E Paul, and Jacques Duysens. 2004. "Regulation of arm and leg movement during human locomotion." *The Neuroscientist* no. 10 (4):347-361.
- Zehr, E Paul, Kathryn L Hesketh, and Romeo Chua. 2001. "Differential regulation of cutaneous and H-reflexes during leg cycling in humans." *Journal of Neurophysiology* no. 85 (3):1178-1184.
- Zehr, E Paul, and Richard B Stein. 1999. "What functions do reflexes serve during human locomotion?" *Progress in neurobiology* no. 58 (2):185-205.
- Zehr, EP, K Fujita, and RB Stein. 1998. "Reflexes from the superficial peroneal nerve during walking in stroke subjects." *Journal of neurophysiology* no. 79 (2):848-858.
- Zehr, EP, T Komiyama, and RB Stein. 1997. "Cutaneous reflexes during human gait: electromyographic and kinematic responses to electrical stimulation." *Journal of neurophysiology* no. 77 (6):3311-3325.
- Zhao, Y, S Inayat, DA Dikin, JH Singer, RS Ruoff, and JB Troy. 2008. "Patch clamp technique: review of the current state of the art and potential contributions from nanoengineering." *Proceedings of the Institution of Mechanical Engineers, Part N: Journal of Nanoengineering and Nanosystems* no. 222 (1):1-11.

Žilinskas, Antanas. 1992. "A review of statistical models for global optimization." *Journal of Global Optimization* no. 2 (2):145-153.