

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Fast human movements and sparse optimal control policies

Permalink

<https://escholarship.org/uc/item/8s79b77v>

Author

Yazdani, Mehrdad

Publication Date

2012

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Fast Human Movements and Sparse Optimal Control Policies

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Electrical Engineering (Intelligent Systems, Robotics, and Control)

by

Mehrdad Yazdani

Committee in charge:

Professor Robert Hecht-Nielsen, Chair
Professor Clark Guest, Co-Chair
Professor William Hodgekiss
Professor Virginia de Sa
Professor David Sworder

2012

Copyright
Mehrdad Yazdani, 2012
All rights reserved.

The dissertation of Mehrdad Yazdani is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Co-Chair

Chair

University of California, San Diego

2012

DEDICATION

Dedicated to the memory of Cezario Tebcherani.

EPIGRAPH

“It is far easier to point out the faults and errors in [a great work] than to give a clear
and full exposition of its value.”

Arthur Schopenhauer from *Criticism of the Kantian Philosophy*

TABLE OF CONTENTS

Signature Page	iii
Dedication	iv
Epigraph	v
Table of Contents	vi
List of Figures	ix
List of Tables	xi
Acknowledgements	xii
Vita	xiv
Abstract of the Dissertation	xv
Chapter 1	Introduction	1
	1.1 Why Human Movements	1
	1.2 Why Optimal Control	3
	1.3 Why Sparsity	3
	1.4 Organization of Dissertation	5
Chapter 2	Neuroanatomy and Fast Human Movements	6
	2.1 Background on Human Movement	6
	2.1.1 Fundamental Nature of Movement	6
	2.1.2 Overview of the Spinal Cord	7
	2.2 Organization of the Gray Matter in the Spinal Cord	9
	2.2.1 The Organization of the Dorsal Horn Laminae I through VI	9
	2.2.2 The Organization of the Ventral Horn Laminae VII through X	11
	2.2.3 Descending Tracts	12
	2.3 Static Postural Goals	13
	2.4 Fast Human Arm Movements Data	14
	2.4.1 Velocity Profiles	15
	2.4.2 Regularities in Velocity Profiles	19
	2.4.3 Peak Velocity to Average Velocity Ratio	22
	2.5 Notes	24

Chapter 3	A Simple Control Policy for Smooth Movement Maximization . .	25
	3.1 Background	26
	3.2 Model Description	27
	3.2.1 Minimizing jerk as a control variable	28
	3.3 Methods	31
	3.4 Results	33
	3.5 Discussion and Future Work	37
	3.6 Notes	40
Chapter 4	Sparse Optimal Control Policies	41
	4.1 Optimal Control Overview	41
	4.1.1 min time	42
	4.1.2 Minimum-Effort Control	44
	4.1.3 Some Remarks on Optimal Control and Optimization	45
	4.2 Sparse Optimal Control Policies for Straight Point-to-Point Trajectories	46
	4.2.1 Sparse Optimal Control Signals in Fast Human Movements	49
	4.2.2 Experimental Setup	51
	4.3 Sparse Optimal Control Policies for Curved Trajectories	53
	4.4 Notes	54
Chapter 5	Multiple Cost Funcions	55
	5.1 Numerical Solutions to the Minimum Effort Control Problem	55
	5.1.1 The solution to the minimum effort control problem using the l_2 norm	59
	5.1.2 The solution to the minimum effort control problem using the l_∞ norm	60
	5.1.3 Minimum Jerk as a Minimum Effort Control Problem	61
	5.2 Minimum Effort Control with a Hybrid Metric	61
	5.3 Minimum Acceleration Criterion with Constraints	65
	5.4 Combination of Multiple Control Signals	66
	5.4.1 Convex Combination Trade-off Control Model	67
	5.4.2 Optimal Confluence Model	70
	5.5 Notes	73
Chapter 6	Confluence	74
	6.1 Introduction	74
	6.1.1 What Policy or Policies?	75
	6.1.2 An Approach to Modeling the Motor System Using Optimal Control: The Optimal Confluence Control Framework	76

6.1.3	An Application of Optimal Confluence Control for Minimum Effort Policies: JCOC	78
6.2	Results	80
6.2.1	Evaluation of Model	82
6.3	Discussion	84
6.3.1	The Role of Neural Integration and the Derivatives of Position in Motor Control	84
6.3.2	Confluence of Neural Signals	86
6.3.3	Motor Primitives and the Optimal Confluence Paradigm	87
6.4	Notes	87
Appendix A	Sparse Optimal Filters	88
A.1	Overview	88
A.2	Problem definition and solution	89
A.2.1	Optimal Filter Formulation	90
A.2.2	Sparse Filter Formulation	91
A.3	Simulation Results	93
A.3.1	Detection of Switching States	93
A.3.2	Estimating Sinusoidal Signal Corrupted in White Gaus- sian Noise	95
A.3.3	Estimating Randomized Signal Corrupted in White Gaussian Noise	96
A.4	Conclusions and Future Work	97
A.5	Maximum SNR filters	97
A.6	Notes	98
References	99

LIST OF FIGURES

Figure 1.1:	Figure taken from [17]. Electronic version of the figure and reproduction permissions are freely available at www.izhikevich.com . . .	4
Figure 2.1:	A typical spinal segment.	8
Figure 2.2:	Neuronal pools span several spinal segments.	8
Figure 2.3:	Experiment setup for data collection and typical movement profile.	15
Figure 2.4:	Typical velocity profile illustrating the “ballistic” and “corrective” portions of a movement.	16
Figure 2.5:	Examples of typical trial velocity profiles. The left column shows the velocity profile in the time domain, and the right column shows the corresponding magnitude Fourier transform.	18
Figure 2.6:	Average velocity profile.	19
Figure 2.7:	Average Fourier Magnitude velocity profile.	20
Figure 2.8:	Top Six Principle Components of Movement.	21
Figure 2.9:	Projection onto the top 2 principle components.	21
Figure 2.10:	Projection onto the top 2 principle components.	22
Figure 2.11:	Top two PCA	23
Figure 2.12:	Distribution of r for each subject.	24
Figure 3.1:	The profiles corresponding to the L_2 norm control policy and L_∞ norm control policy are shown in red and blue respectively. These plots are generated according to Theorems 1 and 2 respectively. . .	29
Figure 3.2:	Histogram of the ratio of average starting and stopping velocities to peak velocity for each trial. Ideally, the trials should have start and end velocities as close to rest as possible in order to conform to both models’ assumptions.	34
Figure 3.3:	Bars A - E show the comparison of the average MSE of position profiles predicted by MJ_2 and MJ_∞ vs. human trial data along with standard error bars. “All Subjects” is an aggregate of the trials in A-E. The MJ_∞ model performs better than the MJ_2 model in all cases. The results for A, B, E are significant with $p < 0.05$ by a Wilcoxon rank-sum test. The “All Subjects” aggregate results are extremely significant with $p < 0.001$ by a Wilcoxon rank-sum test. This test was utilized due to the non-normality of the data as is typically done in this situation [42].	35
Figure 3.4:	Overlay of MJ_∞ model with human movement data of a single trial.	36
Figure 4.1:	Example of a sparse signal and its spike encoding.	46
Figure 4.2:	Examples of sparse optimal control signals $u_i(t)$ for $i = 3, \dots, 6$. Shown are movements that starts from $t = 0$ and end at $t = 1$	50
Figure 4.3:	Average Mean Squared Error for all trials for each subject.	51

Figure 4.4:	A curved trajectory specified by a via point.	54
Figure 5.1:	The solution to equation 5.3 with $n = 3$ and equality matrices as defined by equation 5.14. The blue line shows the solution to equation 5.3 with $p = \infty$, while the red line shows the solution to equation 5.3 with $p = 2$	62
Figure 5.2:	The control policies resulting from varying θ from zero to one for the optimization problem in equation 5.15.	64
Figure 5.3:	The control policies resulting from varying θ from zero to one for the optimization problem in equation 5.23.	69
Figure 5.4:	The solution to equation 5.3 with $n = 5$ and $p = \infty$	71
Figure 5.5:	The optimal convex combination (shown in cyan) of the jerk (blue) and crackle (green) control signals to be as close as possible to the desired peak velocity of a given movement dynamics (shown in red).	72
Figure 6.1:	Figure adapted from [73].	77
Figure 6.2:	The green and blue curves correspond to the minimum jerk and minimum crackle velocity profiles respectively. The red curve is the velocity profile of a give trial.	80
Figure 6.3:	Plotted in blue is the peak velocity of each trial of subjects A-E. Plotted in red is the average peak velocity of each subject. Solid green and orange lines are the peak velocities corresponding to jerk and crackle minimization respectively.	81
Figure 6.4:	Average mean squared error of jerk 2-norm, jerk inf-norm, snap, crackle and JCOC for each subject and for all subjects, all trials. JCOC on average has a smaller error than all models shown with $p < 0.001$	82
Figure 6.5:	Plots of mean squared error (MSE) of three models compared with subject E's horizontal left-to-right movements. Points plotted from the same trial are connected with straight lines.	83
Figure A.1:	Adapted from Moon and Stirling[112]	90
Figure A.2:	Original and Corrupted Signal	93
Figure A.3:	Filter Outputs. Note that an appropriate scaling amplitude should be applied after filter to recover the appropriate amplitude of the original signal. However, in this problem we are most interested in switching times, so this step was not taken.	94
Figure A.4:	Filter Coefficients	94

LIST OF TABLES

Table A.1: Estimation of Sinusoidal Signals	96
Table A.2: Estimation of Random Signal	96

ACKNOWLEDGEMENTS

If it takes a village to raise a child, it must take a planet to raise a PhD. I am deeply indebted to Professor Robert Hecht-Nielsen for his invaluable insights and perspectives. Without his guidance, this dissertation would not have been possible. I have learned a great deal from working with Professor Hecht-Nielsen's broad knowledge of the sciences and his experience in both academia and industry. I would also like to thank my co-chair, Professor Clark Guest, and the rest of my committee for supporting this research and providing valuable feedback.

I am extremely grateful to Amir Karniel for sharing his thoughts and data for our research. His openness is what helps accelerate scientific progress and form breakthroughs. From his generosity I am inspired to do the same and distribute my research data, source code, or anything else freely.

This dissertation would also not have been possible without the countless hours of discussions that I had with my colleagues, Geoffrey Gamble and Gavin Henderson. We have had many disagreements, but I think we all agree that our persistence on high quality research has finally resulted in cracking some of the mysteries of the brain. I am sure that if we keep our persistence at high quality research, we will all reach our dreams of becoming true scientists and solve the pressing questions and problems of the universe.

This dissertation owes a large part to the influence of Cezario "Cez" Tebcherani. I became interested in Artificial Intelligence and Neural computation when, on a whim, I took ECE 173 (Artificial Neural Networks and Fuzzy Logic) my senior year at UCSD, and Cez happened to be in the classroom. He ended up not taking the class, but since then we discussed AI and its related issues and attended the AI seminars during our first year in graduate school. I know that he would have read the contents of this work with great interest. While I won't be able to hear his feedback, his free spirit and hunger for life continues to inspire me.

I also would like to thank Syrus Nemat-Nasser for being an outstanding teacher and inspiring me to get into this research. If it was not for his Artificial Neural Networks and Fuzz Logic class, I would have taken a very different path in life. Great thanks to the many friends and family who have supported me through this endeavor. There are

too many to name and you know who you are!

Finally, I would also like to thank my parents, Hushmand Yazdani and Homa Massih, for supporting me all these years and keeping me on the straight path.

Chapter 2, in part, is a modified reprint with major revision from of the material as it appears in Society for Neuroscience. Yazdani, M., G. Gamble, G. Henderson, R. Hecht-Nielsen. “A neuronal network model of the static postural goal circuit in the spinal cord”. Program No. 108.6. 2010. San Diego, CA: Society for Neuroscience, 2010.

Chapter 3, in full, is a modified reprint with minor revisions of the material as it appears in *Neural Networks*. Yazdani, M., G. Gamble, G. Henderson, R. Hecht-Nielsen. “A Simple Control Policy for Achieving Minimum Jerk Trajectories.” *Neural Networks*, vol. 27, pp. 74-80, 2012.

Chapter 4, in part, has been submitted for publication of the material. The dissertation author is the primary investigator and author of this material.

Chapters 5 and 6, in part, have been submitted for publication of the material. The dissertation author is the primary investigator and author of this material.

Appendix A, in full, is a modified reprint with minor revisions of the material as it appears in SPIE. Yazdani, M. and R. Hecht-Nielsen. “Optimal filters with heuristic 1-norm sparsity constraints.” Proc. SPIE 8137, 813709, 2011.

VITA

Doctor of Philosophy, 2012

Electrical Engineering (Intelligent Systems, Robotics, and Control)

University of California, San Diego

Candidate of Philosophy, 2009

Electrical Engineering (Intelligent Systems, Robotics, and Control)

University of California, San Diego

Master of Science, 2008

Electrical Engineering (Intelligent Systems, Robotics, and Control)

University of California, San Diego

Bachelor of Science, 2005

Electrical and Computer Engineering

University of California, San Diego

PUBLICATIONS

Yazdani, M., G. Gamble, G. Henderson, R. Hecht-Nielsen. “Confluence of Neural Signals and a New Optimal Control Paradigm” *In Preparation*.

Yazdani, M., G. Gamble, G. Henderson, R. Hecht-Nielsen. “Optimal Control and Sparsity” *In Preparation*.

Yazdani, M., G. Gamble, G. Henderson, R. Hecht-Nielsen. “A Simple Control Policy for Achieving Minimum Jerk Trajectories.” *Neural Networks*, vol. 27, pp. 74-80, 2012.

Yazdani, M. and R. Hecht-Nielsen. “Optimal filters with heuristic 1-norm sparsity constraints.” Proc. SPIE 8137, 813709, 2011.

Yazdani, M., G. Gamble, G. Henderson, R. Hecht-Nielsen. “A neuronal network model of the static postural goal circuit in the spinal cord”. Program No. 108.6. 2010. San Diego, CA: Society for Neuroscience, 2010.

Luke Barrington, **Mehrdad Yazdani**, Douglas Turnbull, Gert R. G. Lanckriet. “Combining Feature Kernels for Semantic Music Retrieval.” ISMIR: 614-619, 2008.

ABSTRACT OF THE DISSERTATION

Fast Human Movements and Sparse Optimal Control Policies

by

Mehrdad Yazdani

Doctor of Philosophy in Electrical Engineering (Intelligent Systems, Robotics, and Control)

University of California San Diego, 2012

Professor Robert Hecht-Nielsen, Chair

Professor Clark Guest, Co-Chair

In this work we seek to find the underpinnings of neuronal control and movement. In doing so, we hope to gain insight into the function of the brain and also inspiration for the development of advanced robotics and neuronal prosthetics. In motor control, there is strong evidence that many different signals in the brain are combined. We propose that these various signals are different control policies and that these signals are sparse (reflecting the simplicity of neuronal computation). Additionally, neuronal time-integration plays a crucial role in motor control. Hence, we propose that biological motor control involves

1. Sparse control signals.
2. Time-integration of control signals.
3. Combination of multiple control signals.

We used these three simple ideas to outperforms previous models significantly better on human data (with $p < 0.001$).

Chapter 1

Introduction

1.1 Why Human Movements

The desire to create robotic systems endowed with motor capabilities similar to those of humans has existed for centuries. In 1495, Leonardo da Vinci created a humanoid robot that came to be called “Leonardo’s Knight”. Reproductions based on da Vinci’s designs yield a mechanical knight run by pulleys capable of performing rudimentary imitations of a few human-like movements such as opening and closing its arms, opening its visor, and moving its head [1]. Over 500 years later, roboticists have yet to design robots that can accurately mimic human-like movements or operate in environments within which humans move with ease. A notable fact that can potentially explain the failures over the years is that existing approaches in robotics have typically avoided implementing solutions that replicate the systems level architecture of the mammalian neural and biomechanical system.

The goal of humanoid robotics is to deliver human-like movement, navigation, and object manipulation capabilities in environments humans perform well in (e.g. cities, deserts, mountains, etc.). Attempts to create humanoid robots have involved programming a controller to attempt to accomplish specific tasks. Roboticists have been very successful in providing solutions to well defined engineering problems and will likely continue to engineer adequate solutions to such problems. However, despite enduring efforts, roboticists have not been able to properly mimic basic tasks that humans master at early ages, like running, adjusting to perturbations, and obstacle avoidance. Animal

motor systems are the only known set of systems excelling at the tasks described above. We believe that mimicking the high level mathematical concepts which emerge from an in-depth study of animal motor systems is the most promising path forward in the effort to create humanoid robots. However in order to find these mathematical principles we must look to the anatomy and physiology of animal motor systems as a guide. We propose a computational framework equally grounded in control theory, mathematics, and neuroscience that bridges the gaps between these fields. Once the barriers between these disciplines are broken down, this new family of models will be harnessed to develop human like robots and advanced neural prosthetics.

How the Central Nervous System (CNS) implements a control policy to achieve movements is not understood. Areas of cortex send axons to the spinal cord and generate movements when stimulated, which has led many to believe that cortical structures are responsible for most aspects of movement control [2, 3, 4]. Determining the exact parameters by which the cortex is able to alter behavior has proved to be difficult. Many experiments have correlated cortical activity with various aspects of a behavior [2, 5, 3, 4]. For example, Graziano's work [6] implies that the cortical description of a behavior may be limited to a high level goal state representation. Other evidence shows that the spinal cord plays a key role in the generation of behaviors [7, 8].

One method used to understand the nature of movements is to reduce them to simpler components. For this study we limit ourselves to examining simple, ballistic point-to-point reaching movements. Simple movements might be considered to make up a basis set of which more complicated movements are composed [9, 10, 11]. By studying these simple movements, we may be able to gain insight into the control of more complex movements such as curved movements defined by via points [12] or paths [13]. Furthermore, it has been observed that this class of movements consistently follows bell-shaped velocity profiles [14, 12]. This observation suggests that there exists a set of constraints placed on the dynamic system by a controller. Because this set of constraints exists across a range of movements, it can be said that these movements result from a common control policy. We show here that the observable invariant parameters of movement suggest a neural control policy which is corroborated by neurophysiological experiments.

1.2 Why Optimal Control

Evolution is a process that optimizes species ability to survive and reproduce within their environments. For example, a reasonable hypothesis of an evolutionary reward function that likely shaped the form and function of the retina might be maximizing the resolution of vision. Photoreceptor cells are optimal in relation to this criterion, since they are sensitive to the smallest unit of light, single photons. These cells have been completely optimized by evolutionary processes due to their unique reward function [15]. Optimization theory has great potential to elucidate biological mysteries because it gives us tools that mimic the process of evolution by finding optimal solutions given a notion of cost or reward [16].

We will apply optimal control theory in an effort to understand the vertebrate motor system. The optimal control approach has the added advantage of direct applicability to the control of machines. We seek to demonstrate computationally that understanding and reproducing the myriad of human movements requires a viewpoint broader than most models of movement currently espouse. Using optimal control theory, we can then ask which attributes of the nervous system's control policy might have adapted to provide a competitive advantage. It is reasonable to assume that a bio-mechanical system evolves to find an optimal control policy by optimizing over some cost or reward function. In this paper, we suggest a new cost function and discuss the control policy this cost function dictates.

1.3 Why Sparsity

We, and many others, can view the CNS as a machine that estimates the "state" of the world and, based on this state, forms a response. The response is typically manifested as a motor command so the organism can interact with the environment and potentially change the estimated state. It is well accepted that neurons form a complex network for information processing (for state estimation) and motor control. While the complexity of the neural network in even the simplest organisms can be astronomically high, the individual neuron itself can be viewed as a very simple computational unit. A

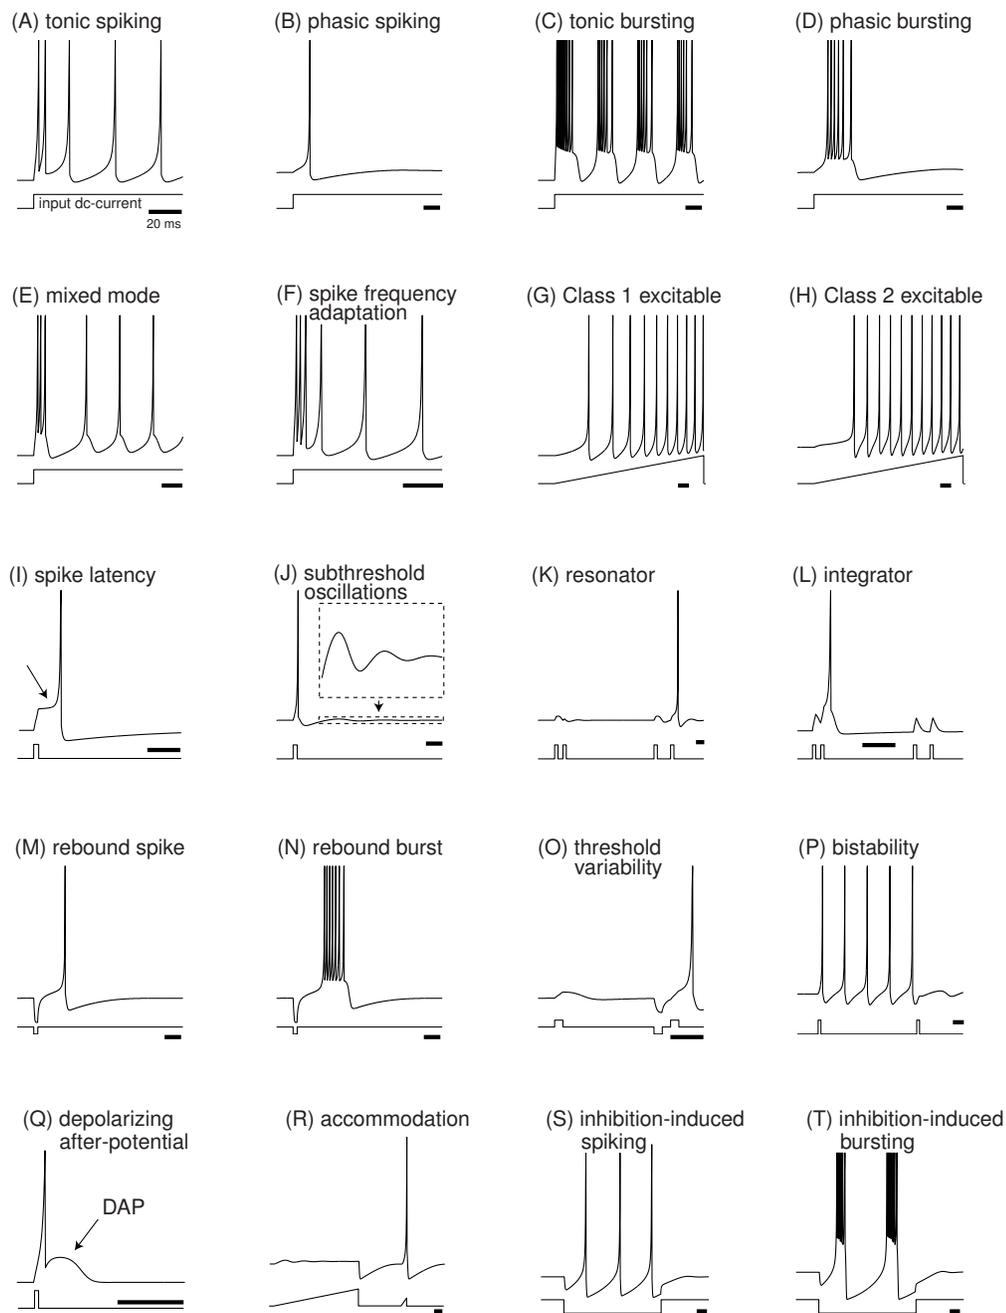


Figure 1.1: Figure taken from [17]. Electronic version of the figure and reproduction permissions are freely available at www.izhikevich.com

neuron is in a quiescent state until its membrane potential reaches a threshold and the neuron fires. Figure 1.1 shows several examples of observed neural responses. These responses can be characterized as *sparse* signals. Thus, if the typical neuronal responses are as shown in Figure 1.1, then sparsity is at the heart of neuronal signal processing. In recent years, sparsity has played a crucial role in developing computational models for information processing in the brain. In this work we propose that sparsity also plays a crucial role for motor control in the CNS. Thus, the brain, in addition to meeting optimality conditions, capitalizes on sparsity in estimation of the state *and* in the control of movement.

1.4 Organization of Dissertation

In Chapter 2 we discuss the anatomy of the spinal cord that is relevant for the control and generation of movement. In this chapter we also describe the variant and invariant characteristics of fast human arm movements by analyzing fast human movement data. In Chapter 3 we introduce a sparse control policy known as a “bang-bang” controller and show that such a control policy fits human data better than the influential model of Flash and Hogan [12]. We extend the bang-bang controller for a general n -th order controller in Chapter 4 and also compare to human data. In addition, we show how sparse optimal control policies can be extended for curved movements. In Chapter 5 we introduce the idea of multiple cost functions in the objective function of an optimization problem and also finding the optimal combination of many control policies. Finally, in Chapter 6 we apply the idea of multiply control policies to human data and show the relationship to biological evidence.

Our research into the sparsity of control signals also led to the design of sparse filters. We report these preliminary results in the Appendix.

Chapter 2

Neuroanatomy and Fast Human Movements

In this chapter we are interested in providing the biological background of sensorimotor control and movement. We focus on the movements that are accomplished by the static postural goal circuit and focus on movements created by humans. The chapter is organized to first give an overview of the neuroanatomy of the spinal cord (the material for which follows closely the studies of [18, 19].) and we conclude by quantitative studies of fast human arm movements in accomplishing point-to-point straight movements (these movements are the type of movement that a static postural goal circuit implements).

2.1 Background on Human Movement

2.1.1 Fundamental Nature of Movement

It has been observed that the trajectories of motion of animal movements follow a bell-shaped velocity profile [20, 14, 21]. We further investigate this claim and other properties of the velocity profile of humans later in this chapter. Flash and Hogan were the first to propose that animal movements can be formulated as the trajectory between an initial and final position in a set time that minimizes the third derivative of the position profile, or “jerk” [12]. In later chapters we further explore this hypothesis and make

connections with neuroscience. Specifically, we show that defining the mathematics of measuring minimization (that is defining the appropriate Lebesgue metric) is crucial in gaining insight into the control policy that the central nervous system may implement. By formulating the appropriate mathematics, we will show in Chapter 4 that the control policy the central nervous systems utilizes is sparse and simple to implement. In Chapters 5 and 6 we also present supporting evidence that the neuronal control signals the central nervous system implements represent dynamic variables of high order derivatives and that many optimal control signals (with respect to different cost functions) are combined.

2.1.2 Overview of the Spinal Cord

A typical horizontal segment of the spinal cord is shown in Figure 2.1. Bror Rexed showed that the gray matter is partitioned into 10 populations referred to as “Rexed Laminae” [22]. The dorsal laminae largely process sensory information (proprioception, nociception, etc.), while the ventral laminae are involved with regulation of movement. Various connections between the laminae exist (including feed-forward and feed-back, excitatory and inhibitory), but their exact properties and anatomy is unknown. Lamina IX is in several disparate clusters where motor neurons (α , β , and γ) reside.

Motor neurons form neuronal “pools” or “needles” that span several spinal segments (indicated as red needles in Figure 2.2). Each needle innervates a single muscle. Each motor neuron innervates several muscle fibers. These neurons innervate the muscle through the ventral root nerves. The dorsal root nerves transmit vital state information to the dorsal laminae in a retrograde fashion.

The corticospinal tract largely terminates in the medial laminae. In addition to the primary motor cortex, the premotor, supplementary, and somatosensory cortical areas also contribute significantly to the corticospinal tract. This tract provides the necessary parameters for initiating the neuronal circuit for a movement/behavior. The information sent down this tract is sparse and we hypothesize that the primary motor cortex (and possibly other cortical motor areas) provide the necessary parameters for

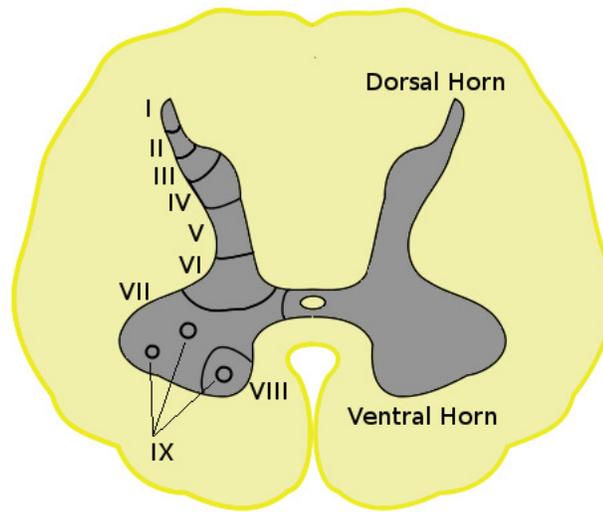


Figure 2.1: A typical spinal segment.

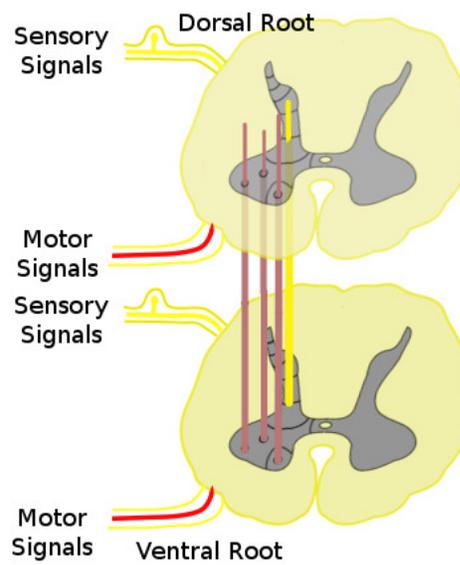


Figure 2.2: Neuronal pools span several spinal segments.

achieving a “static postural goal.” These parameters are final limb position, duration of movement, and final limb stiffness. In addition to receiving inputs from the higher centers in the cortex, the dorsal root ganglion that encodes state information from the extra-spinal system (muscles, bones, loads, etc.) also provides additional input for the spinal circuits. For example, in the case of the “static posture” circuit, the initial posture is provided via the dorsal root ganglion.

2.2 Organization of the Gray Matter in the Spinal Cord

The Gray matter of the spinal cord, as shown in Figure 2.1, is organized in a “butterfly” shape and is categorized in two groups: the dorsal and the ventral horns. The dorsal horn, corresponding roughly to laminae I through VI, can be viewed as the input region for sensory afferent input from the muscle, skin, and other peripheral nerves. The ventral horn, on the other hand, can be thought of as an output region where the motor neurons (that are responsible for motor function and control) reside. The cell bodies are organized in 10 laminae and each of the neurons can be classified broadly in four groups: motor neurons, tract neurons, interneurons, and propriospinal neurons. The motor neurons have axons that leave the spinal cord and innervate muscles. The tract neurons have axons that ascend through the white matter to higher centers (such as in the brain stem or the motor cortex). The interneurons serve as intermediate processing units in the spinal segment. Propriospinal neurons have axons that ascend or descend through the white matter tracts to other spinal segments. Each lamina has at least one of these types of neurons and may also be the site of termination of descending tracts from higher centers such as the brainstem or the motor cortex. Each lamina may also receive inputs from neighboring lamina or neighboring spinal segments.

2.2.1 The Organization of the Dorsal Horn Laminae I through VI

Laminae I and II are involved with the processing of pain or nociception. Lamina I, also known as the zone of Waldeyer, has most of its primary afferent input originating from cutaneous nociceptors and thermoreceptors. Lamina II, also known as the substantia gelatinosa of Rolando, in addition to having primary afferent input from cu-

taneous nociceptors and thermoreceptors, this region also has primary afferent input from mechanoreceptors. The neurons in this lamina are almost entirely excitatory and inhibitory interneurons with dendrites that branch within the lamina and also project to other neighboring laminae.

Lamina III and IV have most of their primary afferent input originating from mechanoreceptors. Lamina III and IV also have inputs originating from Lamina II neurons that project ventrally into these laminae. Lamina V also has inputs from cutaneous mechanical nociceptors and the more dorsal laminae (such as laminae II through IV). Lamina VI is believed to be heavily involved in motor control since it is close to the motor regions in the ventral horn. Many of the descending tracts from the higher centers, such as the primary motor cortex and the brain stem, terminate in this area. Also, in this area there are numerous interneurons and propriospinal neurons. This suggests that this lamina can be viewed as a preprocessing stage for volitional motor control, where volitional commands originate from the higher centers (such as the primary motor cortex and the brain stem), and are projected to neighboring spinal segments via propriospinal neurons or to ventral motor regions via interneurons.

To summarize, the dorsal horn receives a significant amount of input from the peripheral nerves and higher centers. The inputs to the dorsal horn cause the activation of many types of interneurons that modulate ventrally to the ventral horn. The interneurons are both of the excitatory and inhibitory type and form mutually inhibitory networks. This implies that when one particular response pathway along with its corresponding interneurons is activated, all other interneuronal pathways get inhibited. It is believed that this pathway activation allows the spinal cord to respond selectively to any input, either afferent or efferent, without inappropriately activating other antagonistic pathways that might hinder the appropriate response. In addition to modulating the spinal networks, the interneurons can also modulate the output of the higher-order neurons originating from the higher centers.

2.2.2 The Organization of the Ventral Horn Laminae VII through X

The neurons in laminae VII through X are in the ventral horn of the spinal segment and are believed to be heavily involved with motor processing and control. The majority of the afferent inputs to the ventral horn are concerned with proprioception. Lamina VII is involved with the regulation of posture and has numerous tract neurons with connections to the cerebellum and the midbrain. The interneurons and propriospinal neurons in Lamina VII connect to adjacent laminae and cord segments and are involved in the processing of reflexes, posture, and autonomic movement. Again, there are both excitatory and inhibitory type interneurons identified as in the dorsal horn with mutually inhibitory networks that allow selective pathway activation.

Lamina VIII receives inputs from the descending tracts (that originate from higher centers) and also proprioceptive afferents. The interneurons and propriospinal neurons in lamina VIII project to adjacent laminae and spinal segments and also receive inputs from adjacent laminae. Thus lamina VIII forms another layer of neural network processing involving both inhibitory and excitatory type neurons that modulate motor activity based on inputs from both the higher centers (that represent volition or desired movement) and proprioceptive afferents (that represent the state of the environment). Lamina VIII, however, is not the final output layer for motor control.

Lamina IX can be thought of as the final output layer for motor control. The motor neurons, that span several spinal segments, innervate muscle fibers. The motor neurons are organized spatially, in that the more lateral the motor neurons, the more distal the muscles they innervate: the more dorsal motor neurons innervate flexor muscles and the more ventral motor neurons innervate extensor muscles. There are three types of motor neurons in lamina IX: α , β , and γ . α motor neurons exclusively innervate extrafusal muscle fibers, and γ motor neurons exclusively innervate intrafusal muscle fibers. β motor neurons innervate both extrafusal and intrafusal type muscle fibers. In addition to motor neurons, lamina IX also includes interneurons (such as inhibitory Renshaw interneurons) and propriospinal neurons.

Lastly, lamina X, despite its label, is not the final output for motor control but yet another neural processing center. The neurons are either interneurons or propriospinal

neurons with decussating axons. This region most likely serves as a relay to the contralateral spinal segment. Again, the mutually inhibitory network that is formed with the excitatory and inhibitory type neurons allow for the activation or deactivation of pathways on either side of the spinal segment. This is crucial for contralateral motor coordination.

2.2.3 Descending Tracts

The descending tracts are formed by the axons that are projected from the higher centers, such as the motor cortex and the brain stem, to the spinal segment interneurons and motor neurons. A tract that is of particular importance in volition control is the corticospinal tract. This tract originates from the motor cortical areas (the primary motor cortex, the supplementary motor cortex, and the somatosensory motor cortex). In all vertebrates the corticospinal tract terminates in the medial laminae interneurons. In addition to terminating in the medial laminae interneurons and propriospinal neurons, the corticospinal tract also terminates in the motor neurons for humans. It is important to note, however, that while it maybe tempting to deduce that in humans the primary motor cortex has direct control over the motor neurons and hence direct control over the motor control of muscle flexion and extension, this neural connection is not sufficient for the control of efficient and useful movements [23, 24]. The direct connection to the motor neurons most likely provides a level of dexterity in motor control that humans posses, but is not fundamental for the useful movements that have lead to the dominance of vertebrate life forms on Earth. In lower mammals the corticospinal tract terminate exclusively in the medial laminae via polysynaptic connections and interlaminar connections to the motor neuron output.

The corticospinal tract axons terminate on numerous neuron types in the spinal cord, including inhibitory Renshaw cells, excitatory and inhibitory interneurons, and primary afferent fibers. The medial laminae, such as laminae IV to VIII are the site of heavy corticospinal tract termination either via monosynaptic or polysynaptic connections. It is believed that the descending corticospinal tracts that terminate in medial laminae also influence several neural pools indirectly through propriospinal neurons. The propriospinal neurons in turn innervate motor neurons that directly control mus-

cles. Thus, a model emerges: the volitional desire of movement that is generated in the cerebral cortex excites, through the corticospinal tract, a network of spinal neuronal pools that consist of propriospinal neurons (that span parallel spinal segments) and interneurons. This network, with the sensory feedback, integrates to a control command to the motor neurons that then finally control the muscle fibers.

2.3 Static Postural Goals

Much research in computational neuroscience and motor control often attempts to hypothesize the function of cortical motor areas (M1, premotor cortex, etc.) (for example [25, 23, 26, 27]) and largely ignores the spinal cord [24, 28, 29, 30]. These hypotheses attempt to define the coding of the neurons in M1 and many have noted the confusion in what exactly the layer V neurons (that contribute to the corticospinal tract) encode [31, 32, 33, 34]. While there is no dispute that the cerebral cortex is central in behavior processes, the role of the spinal cord is too often misleadingly oversimplified as a mere “bundle of wires” for the cortex to exert direct control over the limbs and bodily functions [30]. We have deduced from the neuroscientific literature that the spinal cord plays an important role in controlling limbs for the generation of controlled movements, and that the cerebral cortex has little direct control over skeletal muscles. In fact, the spinal cord ought to be thought of as a “second brain” taking cues from the cerebral cortex and exerting direct control over limbs for the generation of precise movements.

Although M1 contributes a majority of fibers in the descending corticospinal tract, other cortical motor areas, such as the premotor cortex, the supplementary motor cortex and, perhaps strangely, the somatosensory cortex, also contribute significant fibers to the descending corticospinal tract [35, 36, 37]. Due to these significant projections, we conclude that M1 must not be the sole structure to study. We hypothesize that the subset of the descending corticospinal tract that originates from M1 (and possibly alongside with the somatosensory cortex) encodes “static postural goal” parameters that are received as inputs to spinal circuits. This hypothesis is strongly suggested and supported by the studies of Graziano and colleagues [6]. Here by “static postural goal” we mean the achievement of a stationary final position for a set of limbs. The spinal cir-

circuits are the neuronal gray matter within the spinal column that function independently and carry out different roles. These independent spinal circuits receive appropriate input parameters to generate and directly control precise movements. In the case of static postural goal circuits, the input from the higher centers (most likely M1) consists of three parameters: movement duration, final static posture, and the final muscle stiffness in posture (these parameters are suggested by Graziano's experiments and also computational models [38, 12] typically use these three parameters). In addition to receiving inputs from the higher centers in the cortex, the dorsal root ganglion that encodes state information from the muscles also provides additional input for the spinal circuits. Finally, the output of the spinal circuits is manifested through the motor neurons in the Rexed lamina IX of the spinal column that directly innervate muscle fibers.

2.4 Fast Human Arm Movements Data

Complex movements can be considered to be combinations of shorter, simpler motor programs or movement primitives [39]. To better understand complex movements, one needs to first have a strong understanding of the simpler components that they consist of. For this reason the work described here focuses on simple point-to-point reaching movements. The past literature has claimed several characteristics regarding fast human arm movements:

1. Fast human movements are smooth and have bell-shaped velocity profiles [20, 14, 21].
2. The peak-velocity to average-velocity ratio of fast human movements is an invariant characteristic and on average is 1.805 with a standard deviation of 0.153 [12].

These observations have been crucial in developing computational models that describe human arm movements, and therefore we spend the rest of the chapter to verify the validity of these claims. We will show that fast human movements can indeed be characterized by stereotypical bell-shaped profiles, but the data set that we have does not support

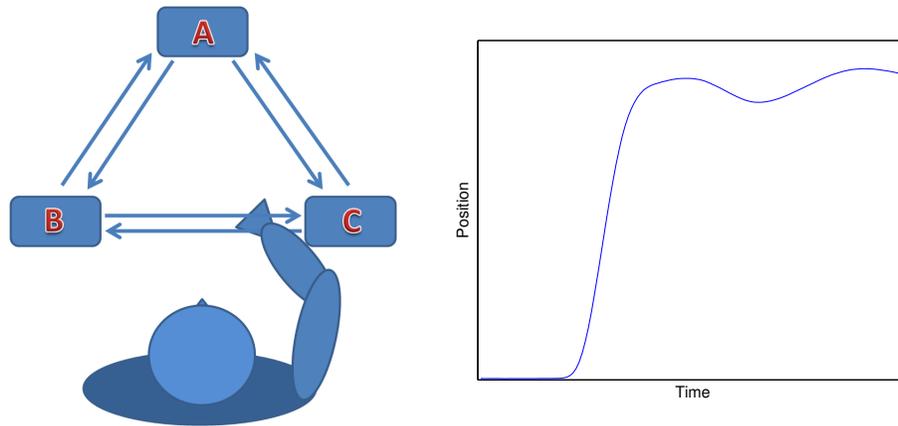


Figure 2.3: Experiment setup for data collection and typical movement profile.

the idea that the peak-velocity to average-velocity ratio of fast human movements is an invariant characteristic, but rather dependent on the subject and the movement type.

We explore the validity of these claims with a data set acquired from Amir Karniel at the Ben-Gurion University of the Negev and was originally described in [40]. The recordings include velocity and position profiles of short point-to-point movements made by 5 healthy individuals and recorded using a manipulandum. Each subject was instructed to move his or her hand between two of three targets arranged in an equilateral triangle, 10 cm apart. The arrangement of the manipulandum and the three points are shown in Figure 2.3(a). The movements lasted a third of a second ± 50 ms. A more formal description of the data collection methods can be found in [40].

2.4.1 Velocity Profiles

Figure 2.3(b) shows a one-dimensional (on the horizontal axis of Figure 2.3(a)) projection of a typical position profile. Note that in this example the subject has an “over-shoot” of the target: that is, when making the movement, the hand of the subject went beyond the target point. As a result of this overshoot, the subject had to make corrective movements to settle at the target point. Similarly, instead of an overshoot, a subject may have an under-shoot: that is, the initial burst of the movement does not reach the end-point target, and thus the subject has to make corrective movements to

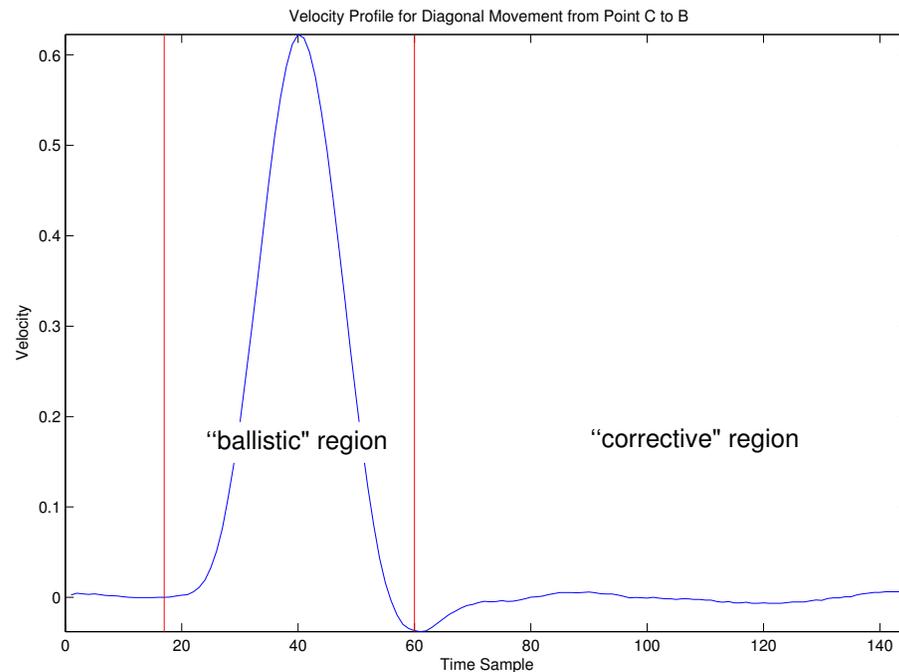


Figure 2.4: Typical velocity profile illustrating the “ballistic” and “corrective” portions of a movement.

reach the target. We can summarize the movement consisting of two portions: the “ballistic” portion of the movement that may have an under-shoot or over-shoot, and the “corrective” portion of the movement that adjusts the hand to reach the final end-point target. Figure 2.4 illustrates these portions.

A fundamental claim regarding fast human arm movements is that the velocity profiles of such movements have stereotypical bell-shaped profiles. We now attempt to verify this claim on the Karniel and Mussa-Ivaldi data set. The left column of Figure 2.5 shows the velocity profile of three trials. Note that all three trials consist of the two portions previously defined: the ballistic portion and the corrective portion. Also note that these trials have various number of time samples. That is, if we view each trial as a vector belonging to a real vector space (in the Euclidean sense), each trial belongs to a different vector space since the dimension of each vector (that is, trial) is different. Also note that in different trials the movements have different onsets, i.e., there is a phase difference between different trials. We wish to explore the statistical behavior of the

velocity profiles, but since each trial belongs to a different vector space it is not clear how to compute various statistics on data points that are in different vector spaces.

To resolve the different number of samples issue, we use the technique of up/downsampling. We find the median trial length (out of the 1,830 trials) is 146, and we up or down sample appropriately each trial to this median length. As a result, all trials now have a length of 146 samples and we can compute the average velocity profile (trials that are in the negative direction are inverted): this is shown in Figure 2.6. This figure verifies that on average the velocity profile of fast human arm movements take a bell-shaped profile followed by a small corrective portion. This figure also verifies that on average the velocity profiles are smooth with no abrupt discontinuities. In addition to a time-domain based analysis, we also consider the Fourier domain. Again, we have to resolve the issue of different trials having a different number of samples, and, additionally, we do not want the phase differences between different trials to be reflected in our statistical analyses since these phase differences do not reveal information about the dynamics of movement (instead, phase information may reveal insights about subject reaction times, which is not the focus of this work).

To resolve these issues, we zero-pad the velocity profiles so that all trials have 512 dimensions and analyze the trials in Fourier magnitude domain. We are only interested in the Fourier magnitudes since the phase information is not relevant (that is, we are not interested in the different onset times that different subjects made on different trials). The right-column of Figure 2.5 shows the 512-dimensional Fourier transforms magnitudes of the corresponding velocity profiles in the time-domain. If the velocity profiles of these trials were indeed perfect bells as [20, 14, 21] and many others have claimed, then the Fourier magnitudes of these velocity profiles would also have bell-shaped profiles. However, we can see that there are deviations from the bell-shaped profile in the Fourier magnitude of the velocity profiles. These deviations are most likely due to the corrective portions of movement that divert the Fourier magnitude from having a perfect bell-shape behavior.

Figure 2.7 shows the average Fourier magnitude transform of the velocity profiles. This is taken over all 5 subjects over 4 days and amounts to 1,830 trials. As

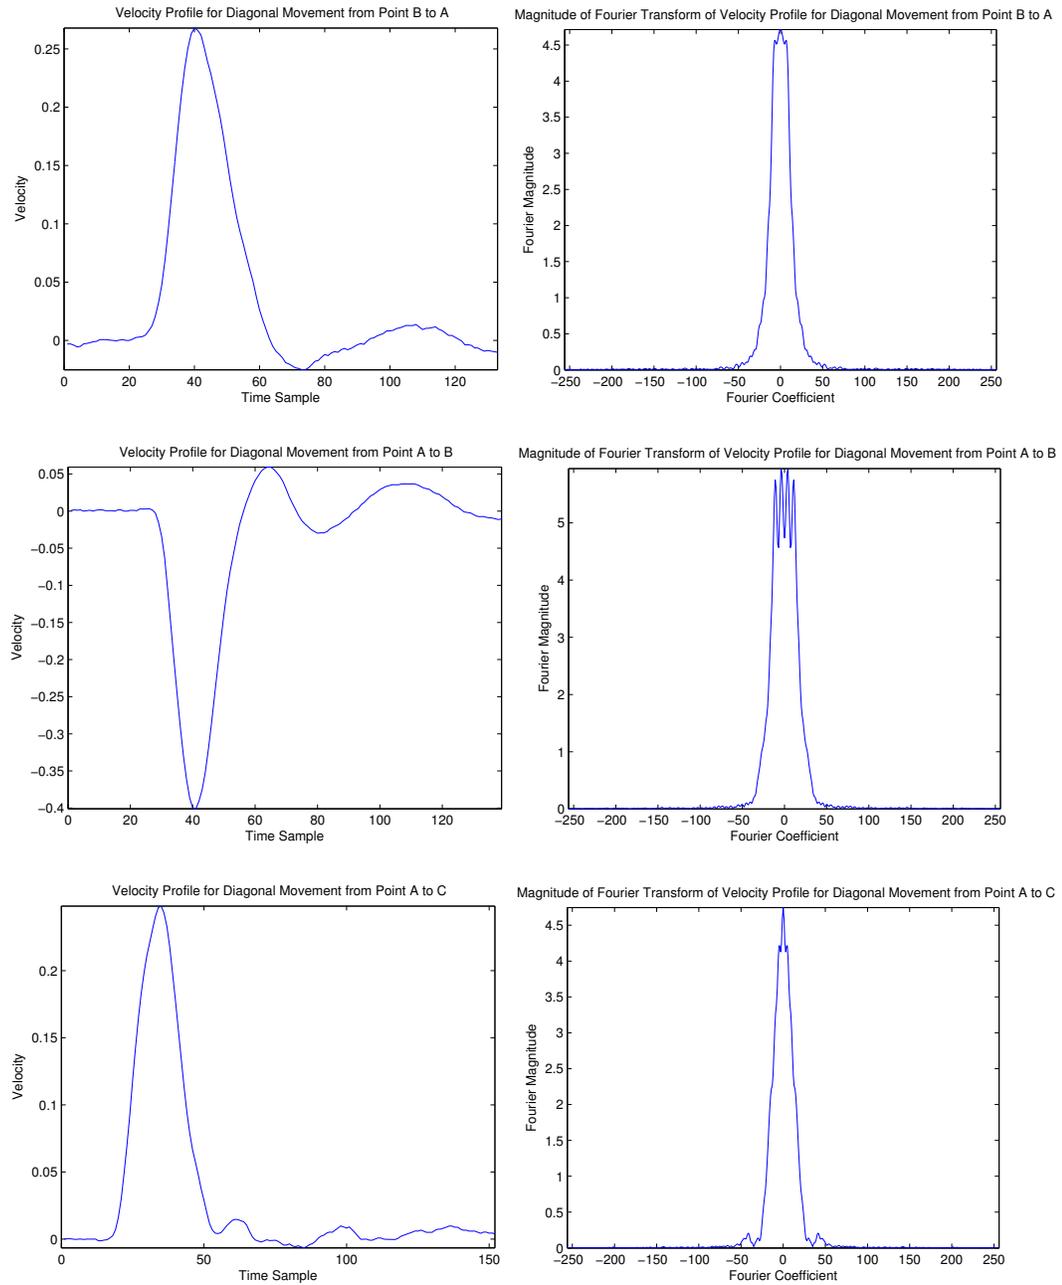


Figure 2.5: Examples of typical trial velocity profiles. The left column shows the velocity profile in the time domain, and the right column shows the corresponding magnitude Fourier transform.

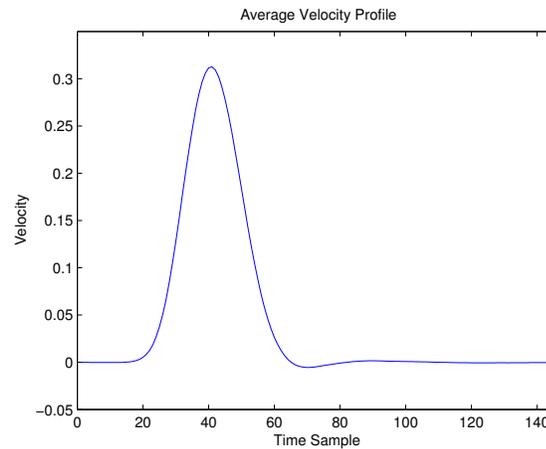


Figure 2.6: Average velocity profile.

Figure 2.7 shows, the average Fourier magnitude-transform of the velocity profiles is nearly a perfect bell. This verifies that the velocity profiles of human subjects, on average, are indeed bell-shaped. In other words, we can conclude that the corrective portions of the movement are mostly random and that when averaged, they all cancel each other out.

2.4.2 Regularities in Velocity Profiles

One of the difficulties in exploring the various velocity profiles is that each trial is a high-dimensional vector that is difficult to visualize. There are various methods of visualizing high-dimensional data, many of which involve projecting the data onto a 2-dimensional subspace for ease of visualization. For our analysis, we compute the Principle Component Analysis (PCA)[41] decomposition on this data. This was done by first computing the covariance matrix of the data set in the Fourier magnitude domain (a 512 dimensional vector space), followed by the eigenvalue decomposition on this covariance matrix. The eigenvectors then correspond to the “principle components” or the directions with the most amount of variance in the Fourier magnitude domain. The eigenvalues correspond to the variance in each principle component. In effect, PCA projects the data onto an ordered set of axes, where the first axis corresponds to direction with the largest variance and the last axis (in this case, the 512th direction) corresponds

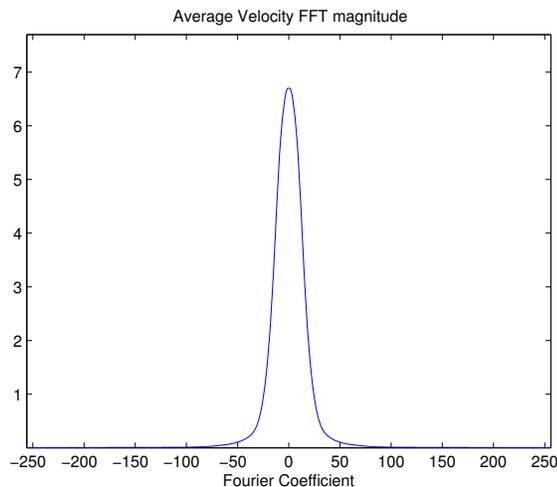


Figure 2.7: Average Fourier Magnitude velocity profile.

to the direction with the smallest variance. This ordering allows us to only consider the directions with the largest variance.

Figure 2.8 shows the 6 principle components with the largest variance. These 6 principle components account for 99.34% of the variance in this data set. Similarly, the top 2 principle components account 94.1% of the variance in this data set. To be able to visualize and compare each trial, we project each trial onto the top 2 principle components and this is shown in Figure 2.9. Clearly we see that there are two prominent clusters. Figure 2.10 shows the the same projection but highlighting each trial by subject. Figure 2.11 shows these projections for the 5 subjects and all subjects, but highlighting the movement type. Figure 2.11 clearly shows what the two prominent clusters correspond to: one cluster corresponds to horizontal type movements (movements that are between points B and C in Figure 2.3(a)) and another corresponds to movements that are diagonal (movements that are between points A and B and between points A and C in Figure 2.3(a)). Clearly, this indicates that for these 5 subjects, the dynamics of fast arm movements fall into two categories. Since the system has not changed, we infer from this result that the central nervous system utilizes (at least) two different control strategies for the control of movement. This leads us to consider different control

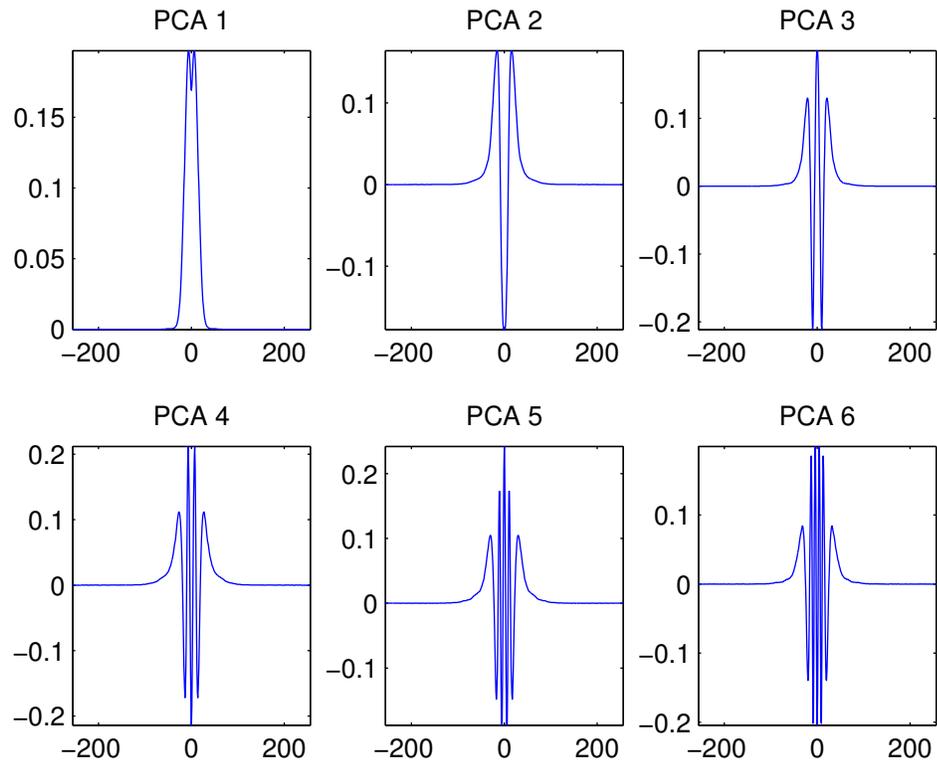


Figure 2.8: Top Six Principle Components of Movement.

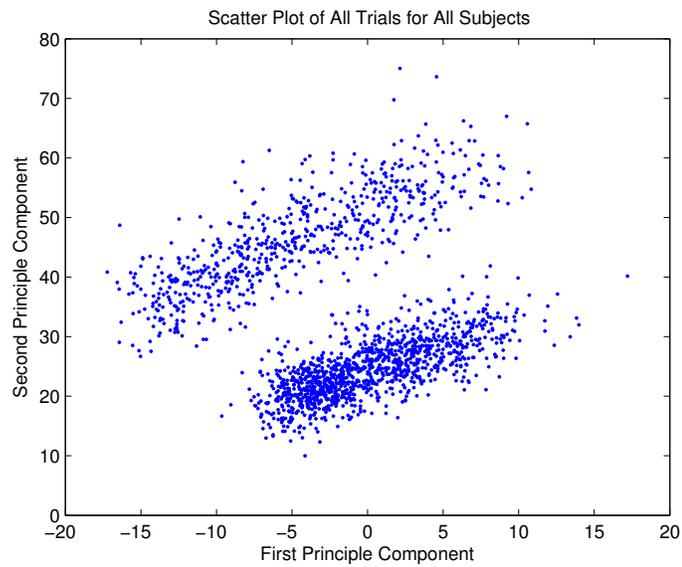


Figure 2.9: Projection onto the top 2 principle components.

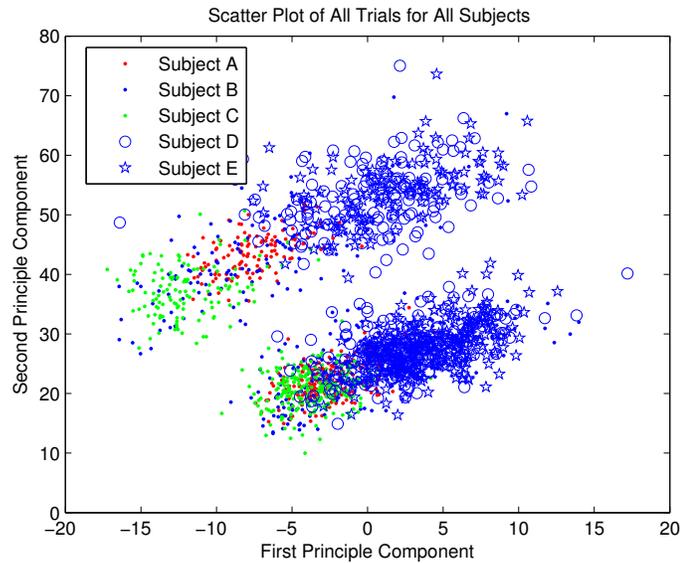


Figure 2.10: Projection onto the top 2 principle components.

policies that we discuss in later chapters.

2.4.3 Peak Velocity to Average Velocity Ratio

Another assumption regarding fast human arm movements is that the peak-velocity to average-velocity ratio of the “ballistic” portion of the movement is an invariant characteristic and on average is 1.805 with a standard deviation of 0.153 [12]. Here we clarify the validity of this claim on the Karniel and Mussa-Ivaldi data set. In the rest of the document we refer to this ratio as r . We compute r for the “ballistic” portion of each trial and show the distribution of r for each subject in Figure 2.12. The “ballistic” portion of each trial was found by the technique described in Chapter 4. A Chi-squared goodness-of-fit shows that the r for each subject is normally distributed with $p < 0.001$. Figure 2.12 also shows that the mean r for each subject is different and can be very different from the purported 1.805. A one-way ANOVA test shows that the mean r for each subject is indeed different (with $p < 0.001$), hence we cannot verify that r is an invariant characteristic of fast human arm movements. This further supports the idea that there are various dynamics in fast arm movements. Again, as in the case of the PCA discussion above, since the system has not changed, we infer that there are

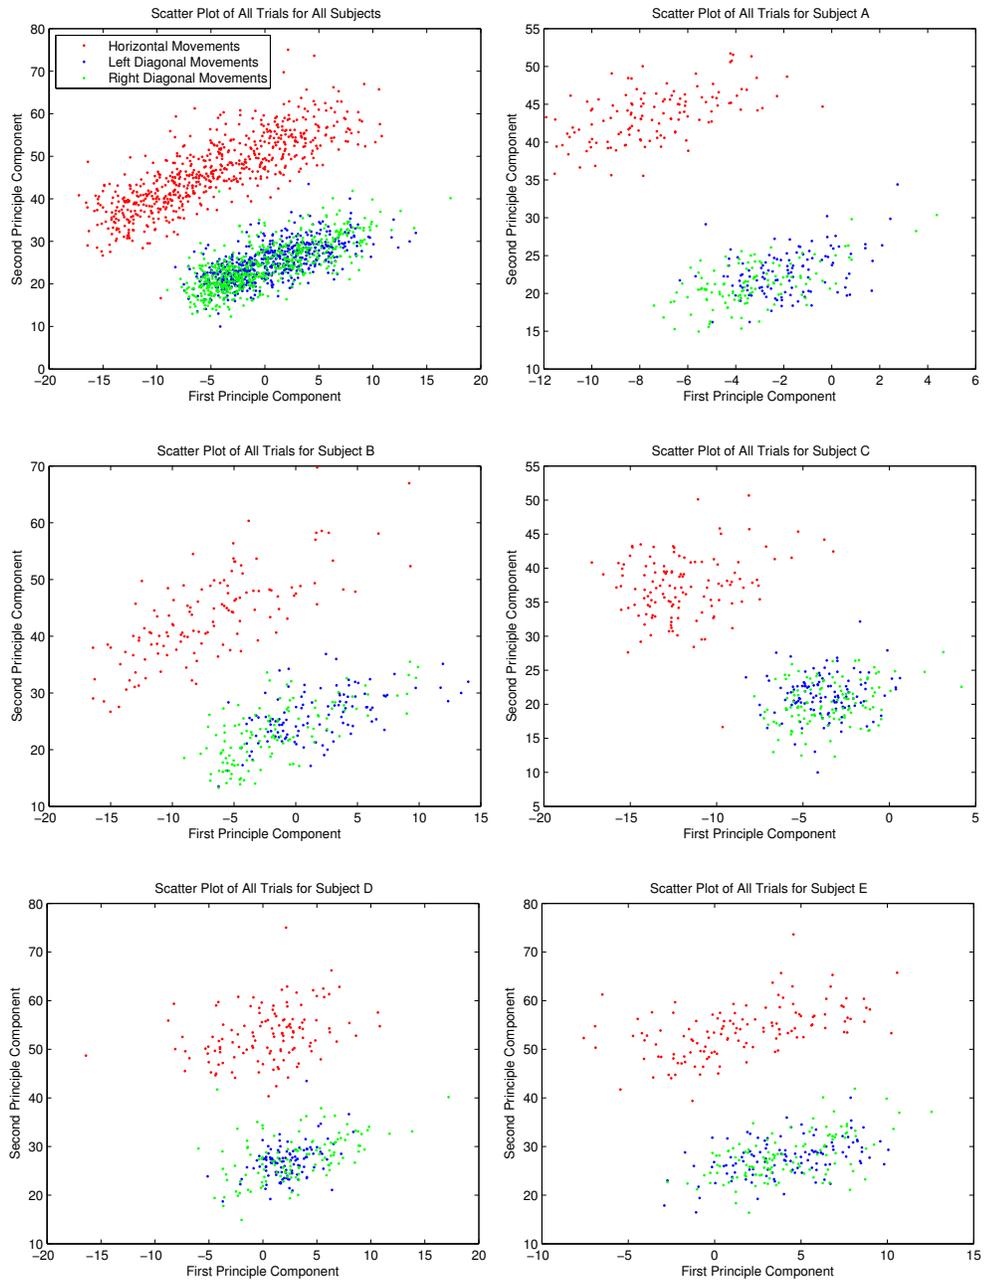


Figure 2.11: Top two PCA

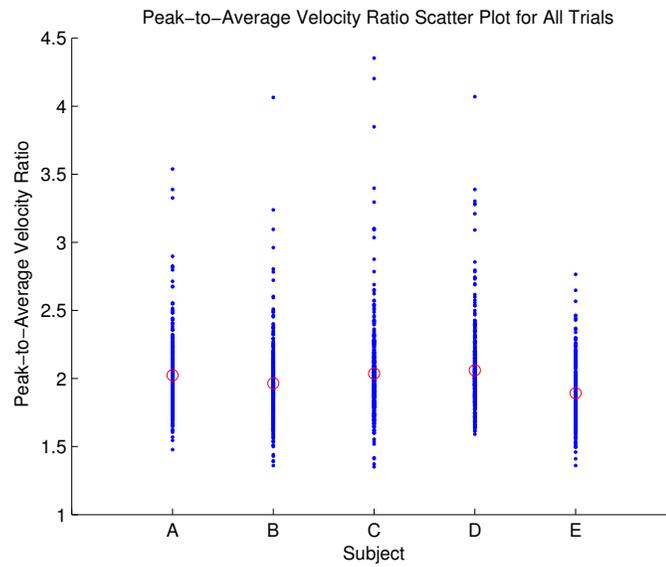


Figure 2.12: Distribution of r for each subject.

a multitude of control policies used by the central nervous system hence resulting in a myriad of movement dynamics. These ideas are discussed in much greater detail in later chapters.

2.5 Notes

Chapter 2, in part, is a modified reprint with major revision from of the material as it appears in Society for Neuroscience. Yazdani, M., G. Gamble, G. Henderson, R. Hecht-Nielsen. “A neuronal network model of the static postural goal circuit in the spinal cord”. Program No. 108.6. 2010. San Diego, CA: Society for Neuroscience, 2010.

Chapter 3

A Simple Control Policy for Smooth Movement Maximization

Point-to-point fast hand movements, often referred to as ballistic movements, are a class of movements characterized by straight paths and bell-shaped velocity profiles. In this chapter we propose a bang-bang optimal control policy¹ that can achieve such movements. This optimal control policy is accomplished by minimizing the L_∞ norm of the jerk profile of ballistic movements with known initial position, final position, and duration of movement. We compare the results of this control policy with human motion data recorded with a manipulandum. We propose that such bang-bang control policies are inherently simple for the central nervous system to implement and also minimize wear and tear on the bio-mechanical system. Physiological experiments support the possibility that some parts of the central nervous system use bang-bang control policies. Furthermore, while many computational neural models of movement control have used a bang-bang control policy without justification, our study shows that the use of such policies is not only convenient, but optimal.

¹In Chapter 4 we generalize these control signals as “sparse” signals.

3.1 Background

There have been several attempts to mathematically model human point-to-point movements of an end effector (the hand) [12, 13, 38]. A key study of this type is [12]. Flash and Hogan began their work on modeling movement by observing that short duration, straight line reaching movements (ballistic movements) exhibit a stereotypical bell-shaped velocity profile. Their work resulted in a model that was exceptionally good at reproducing the trajectory of a movement given the limited information of initial position, final position, and movement duration. This model was achieved by finding the trajectory that minimized the L_2 norm of the 3rd derivative of the position trajectory. The 3rd derivative of position is known primarily as “jerk”, but is also known as “shock”, “jolt”, “surge” or “lurch”. We henceforth refer to their model minimizing the L_2 norm of jerk as MJ_2 .

Although the MJ_2 is exceptionally good at reproducing trajectories from limited constraints, it remains unclear how the central nervous system would generate these trajectories. Various other models have been proposed that minimize other derivatives of position such as acceleration [42]. This recent work by Ben-Itzhak and Karniel has produced a model that not only generates accurate point-to-point movements but also suggests a control policy by which the CNS could be generating such movements. The work presented here expands upon those findings and suggests an alternative model that presents a simple control policy the CNS may implement.

We propose an optimal control policy for achieving ballistic movements based on minimizing the jerk of the trajectory of the end effector. We formulate the problem as an optimal control problem wherein the jerk is treated as the control signal. Our model minimizes the L_∞ norm of the jerk and shows that the optimal control policy is of a “bang-bang” type controller, a policy which simply switches a system between two states [43]. The appeal of such controllers is that they are inherently simple to implement. Furthermore, minimizing the L_∞ norm minimizes the maximum allowable jerk for the system, which can reduce wear and tear. Henceforth, we refer to our proposed model as MJ_∞ .

Flash and Hogan also solved for the trajectory that minimizes jerk, however their cost function utilizes the L_2 norm. While such a cost function yields bell-shaped velocity

profiles and position profiles as observed in humans, the jerk profiles given by the model are not a simple bang-bang type controller. Ben-Itzhak and Karniel developed a model (MACC) that also yields accurate trajectories. In contrast to MJ_2 model, their model implies a simple bang-bang controller. They arrive at this model by minimizing the L_2 norm of the acceleration with a free parameter constraining the maximum allowable jerk. While they show that their model improves error significantly, we argue that their usage of a free parameter is not needed and adds unnecessary complexity. The model presented here still yields a bang-bang control policy but is less complex in that no free parameter is required.

This paper is divided into the following sections. In section 2, we describe our model used to describe realistic human movements as well as provide justification for its biological importance. Next, in section 3, we explain the process by which the human movement data were collected and parsed for proper analysis and comparison. In Section 4 we discuss the results of the data comparison with our model, MJ_∞ , and the MJ_2 model. The last section discusses the importance of this work and draws conclusions about the insights provided by modeling human movement.

3.2 Model Description

The reasons the central nervous system minimizes the jerk of movements are not immediately apparent. Mechanical systems have maximum tolerances related to various dynamic variables (velocity, acceleration, jerk, etc.). Beyond these tolerance levels, components of the system may begin to fail. Biological systems are mechanical systems and therefore also have thresholds that, when exceeded, may lead to damage such as ligaments and muscles tearing or bones breaking. Jerk is one of the dynamic variables that bears directly on the well-being of a mechanical system. Mechanical engineers and roboticists have recognized the benefits of minimizing jerk and have incorporated this concept into their systems [44, 45, 46, 47]. Optimizing animal movement by minimizing jerk is beneficial in that it can reduce stress on the mechanical components of the body.

It is not obvious what function of the instantaneous jerk should be minimized to match biological observation. The L_2 norm (as used in the MJ_2 model) measures

the summation of squared jerk over the course of the movement while the L_∞ norm (as used in the MJ_∞ model) minimizes the maximum jerk value over the course of the movement. While the L_2 norm metric penalizes high jerk values, it does not explicitly force the system to keep the maximum instantaneous jerk as low as possible. In the jerk profile figure shown in Figure 3.1, notice that near time $t=0$ and $t=1$ second, the jerk resulting from the MJ_2 model far exceeds the maximum jerk over the entire movement by the MJ_∞ model. In contrast to the L_2 norm metric, limiting the maximum instantaneous magnitude of jerk via an L_∞ norm cost function reduces the possibility of the movement passing some critical jerk threshold, after which damage to the body may occur. This intuitive rationale helps justify why evolution may have minimized the maximum magnitude of instantaneous jerk (L_∞ norm) during a movement rather than the sum of squared jerk over the course of a movement (L_2 norm).

3.2.1 Minimizing jerk as a control variable

In this section we formulate the problem of minimizing the jerk of a ballistic point-to-point movement as a control problem where the control signal is the jerk, the initial and final positions are known, and the duration of movement is also known. For ease of notation, here we restrict our problem formulation to one dimension and note that extensions to higher dimensions are straight-forward. The control signal that we seek to achieve a minimum jerk position trajectory $x(t)$ is formulated as follows:

$$\begin{aligned} & \underset{u(t)}{\text{minimize}} && \| u(t) \|_p \\ & \text{subject to} && \dot{\mathbf{x}}(t) = A\mathbf{x}(t) + Bu(t) \end{aligned} \tag{3.1}$$

where $A = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$, $B = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$, $\mathbf{x}(t) = \begin{bmatrix} x(t) \\ \dot{x}(t) \\ \ddot{x}(t) \end{bmatrix}$, $u(t) = \ddot{x}(t)$, and where $\| \cdot \|_p$ denotes the L_p norm.

The solution to equation A.3 will determine the optimal control policy $u(t)$. The selection of the L_p norm can result in vastly different control policies. For $1 \leq p < 2$, the control policy will result in physiologically unrealistic movements and as a result these

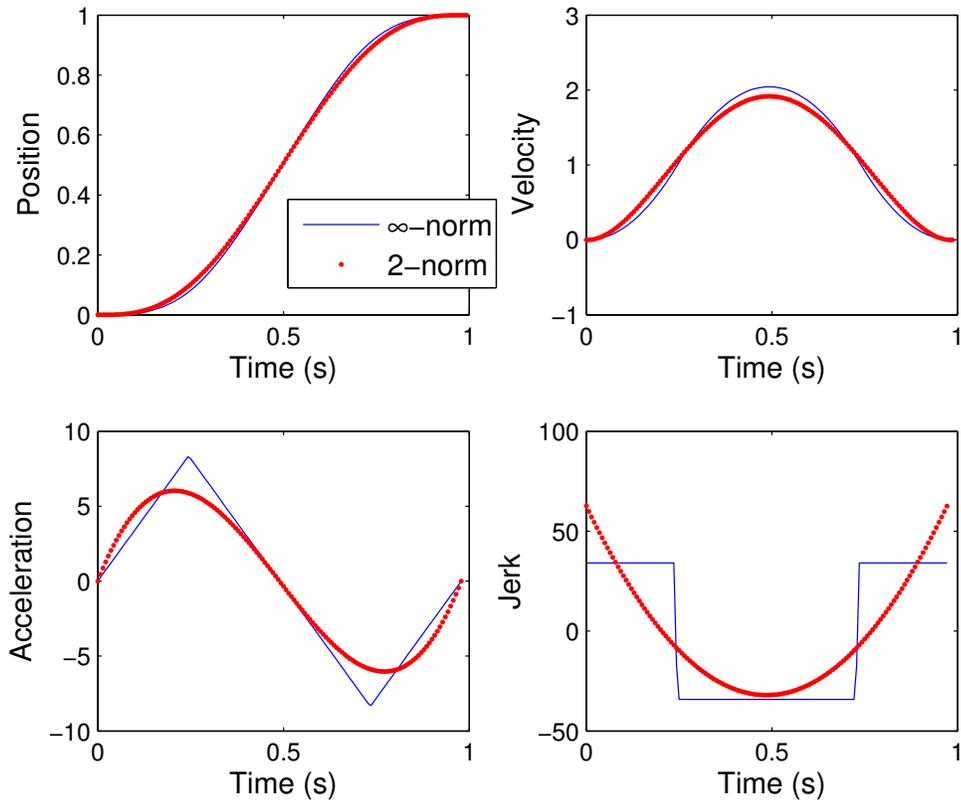


Figure 3.1: The profiles corresponding to the L_2 norm control policy and L_∞ norm control policy are shown in red and blue respectively. These plots are generated according to Theorems 1 and 2 respectively.

types of control policies are not discussed here. Instead we will pay close attention to cases where $p = 2$ and $p = \infty$. For $p = 2$, we have the following policy as described in the following theorem:

Theorem 1. *The solution to equation A.3 with $p = 2$ is a straight line trajectory given by the following control policy:*

$$u(t) = \ddot{x}(t) = (x_f - x_i) \left(\frac{360}{T^5} t^2 - \frac{360}{T^4} t + \frac{60}{T^3} \right)$$

where x_i is the initial hand position at time $t = 0$ and x_f is the final hand position at time $t = T$.

Proof. This was originally shown by [12] □

The control policy corresponding to the above theorem has been shown to fit human data very well [12]. The next theorem shows that if the L_∞ norm of jerk is minimized (i.e, $p = \infty$) in equation A.3 then we have a bang-bang control policy:

Theorem 2: Bang-bang Control Theorem. *The solution to equation A.3 with $p = \infty$ is a straight line trajectory given by the following control policy:*

$$u(t) = \ddot{x}(t) = \begin{cases} J & 0 \leq t < \frac{T}{4} \\ -J & \frac{T}{4} \leq t < \frac{3T}{4} \\ J & \frac{3T}{4} \leq t \leq T \end{cases} \quad (3.2)$$

with $J = 32 \frac{x_f - x_i}{T^3}$ where x_i is the initial hand position at time $t = 0$ and x_f is the final hand position at time $t = T$.

Proof. This was originally shown by [45] □

Ben-Itzhak and Karniel [42] proposed a similar bang-bang control policy for achieving ballistic point-to-point movements. Their control policy minimizes acceleration and also places a threshold on the jerk of the trajectory. This threshold is a free parameter in their model that controls the amount of allowable jerk. Here we show that achieving a bang-bang control policy can be done without introducing any free parameters simply by minimizing jerk as measured by the infinity norm.

3.3 Methods

The human arm movement data used in this study for comparison with our model were provided by Amir Karniel at the Ben Gurion University of the Negev, the same data he and Ben-Itzhak used in their 2008 paper outlining their model of fast arm reaching movements [42]. The data originated from a 2002 paper by Karniel and Mussa-Ivaldi [40] investigating the nervous system's ability to adapt to perturbations. An abbreviated description of the data collection techniques is given below. For a complete description, see [40].

Seated subjects held a robotic manipulandum which was restricted to two dimensional movements corresponding to the horizontal plane of the subjects. They watched a screen which displayed the position of their hand (and the manipulandum) in relation to three positional markers A, B and C. The markers were positioned to form an equilateral triangle (see Figure 2.3(a)). The subjects were instructed to move the on-screen representation of the manipulandum from one target to another. The distance between the targets was 10cm. This motion was to occur within one third of a second, ± 50 ms. Feedback was given to the subjects indicating if they had reached the target and if they did so within the appropriate time window. Position profiles were recorded for all six possible movement types for five subjects over the course of four days. The original data included a subset of trials in which the arm was perturbed during movement from one marker to another. This subset was excluded from our analysis. Only unperturbed movements were analyzed. See Figure 2.3(b) for an example of a typical movement.

The data included uninteresting aspects such as near stationary positional information before a subject began moving and after a subject reached his or her goal and stopped moving. Various methods have been used for movement onset detection [48, 49, 50]. Unfortunately, there is no consensus regarding which technique is best for choosing the relevant portion of a movement as the definition of what is relevant may change from study to study or from one movement type to another. We employ a simple method to determine the start and end times of each movement. We start by finding the onset of the movement. To do so, we compute the energy of a moving window of five

time steps over the velocity profile from a given trial:

$$E = \sum_{i=1}^5 v^2[i] \quad (3.3)$$

where $v[i]$ is the velocity of the manipulandum at time step i of the current window. The window starts at the beginning of the recorded data of a trial (time steps 1-5) and moves forward in one time step increments (e.g. 2-6, 3-7, etc). If E is not greater than a threshold, δ , the window continues moving forward and the test is repeated.

When the window moves over a portion of the velocity profile where the manipulandum is both stationary and close to the starting position (i.e. before the movement begins), E is low. As the window moves over a portion of the velocity profile which is increasing, E becomes greater. We define the starting time of the movement to be the beginning time step of the window at the first window position where $E \geq \delta$.

Finding the time at which the ballistic portion of the movement ends is difficult due to corrective movements made by the subjects after they reach their target, e.g. the correction of perceived target overshoot. Recall that we are only interested in ballistic movements. The corrective portions of the movement fall outside the ballistic portion of the movement. We define the middle ($\frac{T}{2}$) and end (T) times of a movement in the same way as done in [42]. End effector velocity profiles for ballistic point-to-point movements are known to have a symmetric bell shape [14, 12]. In order to determine the end time of the movement, we first define the middle position of the movement ($\frac{T}{2}$) to be the point of maximum velocity, i.e. the top of the symmetric bell. We then simply double this value to find the end time T .

Like other methods, this heuristic technique is not guaranteed to find the ideal onset and end of the recorded movements. Both the MJ_2 and MJ_∞ models assume an initial and final position at which velocity is zero. For this reason, it is appropriate to filter the trials, keeping those for which our start/end detection algorithm has chosen points which most closely meet the zero velocity start and end point assumptions of the models. By definition these are the only trials that are relevant to the models.

Since none of the trials have exactly zero velocity start and end points, some degree of tolerance must be allowed. Furthermore, some metric must be employed to define the degree of closeness to zero velocity for a given start/end point. We define

close to zero velocity for start and end points on a trial by trial basis by computing a ratio between the velocities in question and the peak velocity:

$$c = \frac{v_s + v_e}{2v_p} \quad (3.4)$$

Where c is a unitless indicator of closeness to zero, v_s is starting velocity, v_e is the ending velocity and v_p is the peak velocity of the movement. Applying this metric to all trials yields a distribution between 0 and 1 (see Figure 3.2). As the c value gets closer to zero, the corresponding trial increasingly conforms to the zero start and end assumptions of the MJ_2 and MJ_∞ models and is therefore more appropriate for comparison against the models. We included all trials with a c value less than the harmonic mean of the entire distribution. That is, we favored trials which most closely conform to the assumptions of the models while still leaving enough trials to properly gauge statistical significance. The use of the harmonic mean as a level of tolerance of deviation from zero velocity is somewhat arbitrary. What is important is that the trials with large c values (that do not conform to the models assumptions) are discarded while keeping enough trials to maintain statistical significance. In all, our filtered data set included 406 movement trials.

3.4 Results

To assess the performance of the minimum L_2 norm jerk model (MJ_2) in comparison to the minimum L_∞ norm jerk model (MJ_∞), we compute the time-series mean-squared error between the model's predicted position trajectory with the human subjects trajectory. Our intention is to show that MJ_∞ 's much simpler control policy can fit data at a high accuracy. In fact, we show that our model significantly exceeds the accuracy of the MJ_2 model for this data set. Formally, the mean-squared error is computed as follows:

$$MSE_k = \frac{1}{t_k} \sum_{i=1}^{t_k} (x_k[i] - p_k[i])^2 \quad (3.5)$$

where k refers to the trial, t_k the number of samples for the k th trial, x_k is the model position profile for the k th trial, and p_k is the recorded data position data for the k th trial.

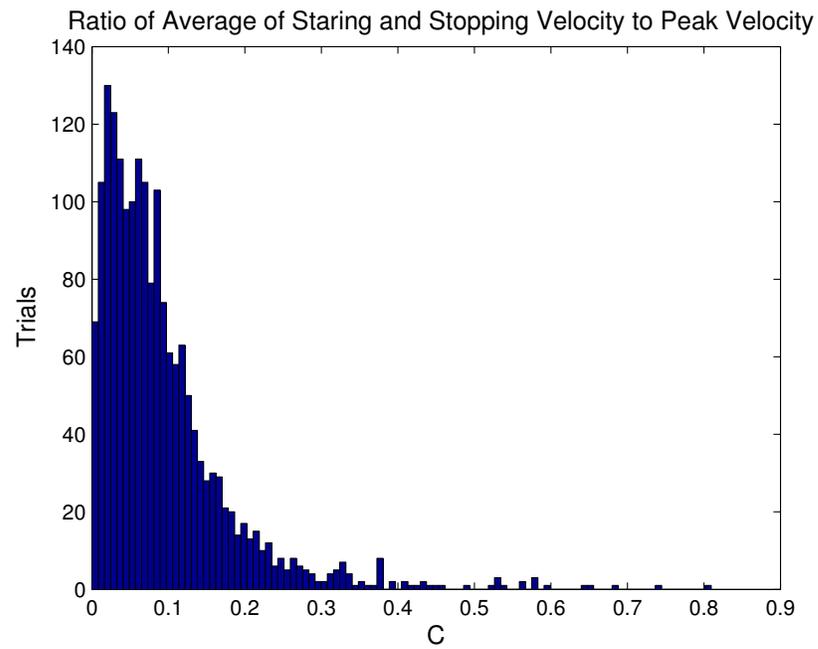


Figure 3.2: Histogram of the ratio of average starting and stopping velocities to peak velocity for each trial. Ideally, the trials should have start and end velocities as close to rest as possible in order to conform to both models' assumptions.

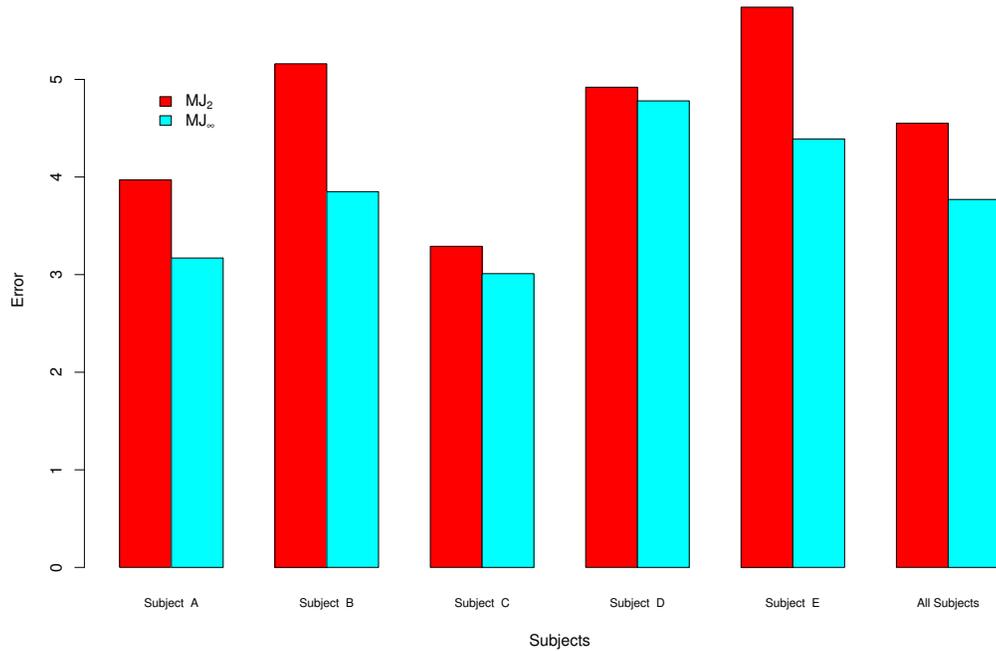


Figure 3.3: Bars A - E show the comparison of the average MSE of position profiles predicted by MJ_2 and MJ_∞ vs. human trial data along with standard error bars. “All Subjects” is an aggregate of the trials in A-E. The MJ_∞ model performs better than the MJ_2 model in all cases. The results for A, B, E are significant with $p < 0.05$ by a Wilcoxon rank-sum test. The “All Subjects” aggregate results are extremely significant with $p < 0.001$ by a Wilcoxon rank-sum test. This test was utilized due to the non-normality of the data as is typically done in this situation [42].

Figure 3.3 shows the comparison of the average MSE trajectories between MJ_2 and MJ_∞ . MJ_∞ has a smaller MSE in all cases. We therefore can conclude that the trajectory estimates of MJ_∞ are just as good or significantly better than those of MJ_2 . Furthermore, their control policies differ significantly and we suspect that the simple control policy (the bang-bang controller of MJ_∞) would be favored by a biological system.

The MACC (Minimum Acceleration Criterion with Constraints) model proposed

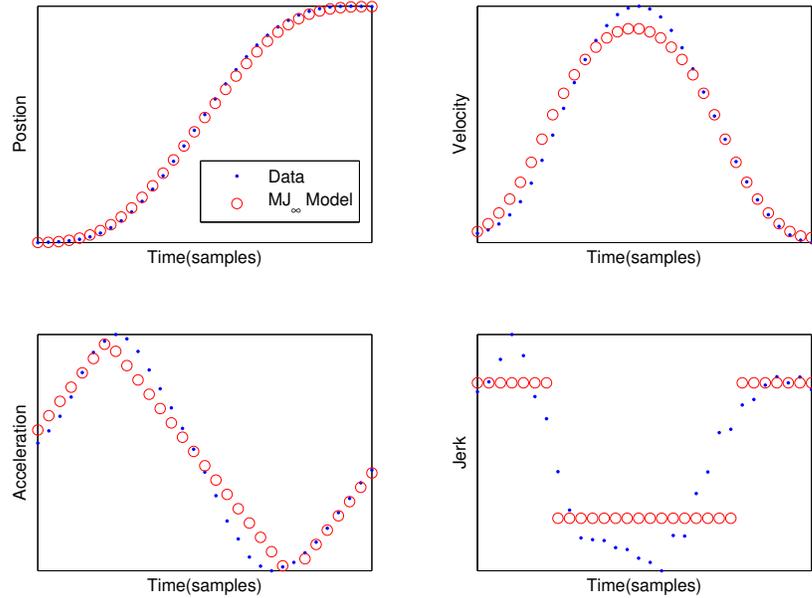


Figure 3.4: Overlay of MJ_∞ model with human movement data of a single trial.

by Ben-Itzhak and Karniel [42] also implies a simple bang-bang controller, however, their approach requires a free parameter that places a cap on the maximum allowable jerk (manifested as a constraint on the control signal). Although changing this parameter can change the switching times of the bang-bang controller, it is not clear how to select an appropriate value for this parameter. Ben-Itzhak and Karniel choose the value for this parameter by performing a grid search and selecting the best value that fit the data for each trial. Even though they showed that using this method of selecting the parameter value results in trajectories that have MSE significantly smaller than the MJ_2 , these results depend upon the model's ability to tune this parameter on a trial by trial basis. This may explain the improvement over other models they have used for comparison. While it is feasible for the CNS to implement additional parameters, free parameters add unnecessary complexity for achieving bang-bang control. It is worth noting that while our model does not require any free parameters, it is a special, but important, case of the MACC model where the free parameter is chosen such that the infinity-norm of the jerk profile is minimized and hence is equivalent to the MJ_∞ jerk profile.

Our suggestion of a bang-bang control policy is based upon the fact that mini-

mizing jerk with an L_∞ norm measure results in a two state jerk profile. Since the jerk is an observable characteristic of the movement, experimental data should be able to confirm the true nature of the jerk profile. Unfortunately, numerical computations of jerk by derivative approximation (as done in this paper) amplify noise inherent in the original recording. This makes drawing direct conclusions about the nature of the jerk profile difficult (see jerk approximation in Figure 3.4). To overcome this difficulty we decided to evaluate the error of the models with respect to the positional data. Future work on this hypothesis should include recordings of acceleration, which should reduce noise amplification and possibly make the true jerk profile apparent.

3.5 Discussion and Future Work

In this paper we studied an optimal control policy for achieving point-to-point ballistic movements using the minimum-jerk criterion. We focused on minimizing the L_∞ norm of jerk (MJ_∞) to achieve a simple bang-bang control policy as opposed to using the L_2 norm (MJ_2). We compared the two policies with human motion data recorded with a manipulandum and showed that the MJ_∞ outperforms the MJ_2 at predicting ballistic human arm movements.

Determining the precise contributions of the various components of the CNS to the control of movement is difficult, since observations of the motor system's neural activity in behaving animals are hard to obtain. However, measurements of external motor behavior are much easier to record. It is natural then to attempt to leverage the movement data we have to explain what the control policy used by the CNS may be. Since the movements were performed by well trained individuals in an unperturbed workspace it is reasonable to assume that feedback due to movement errors would be minimal. Furthermore, experimental studies on deafferented animals have demonstrated that trajectory planning for fast point-to-point movements is not disrupted [21] and that proprioceptive or cutaneous feedback is not necessary for the execution of such movements. Because of the lack of feedback involved with these movements, we can model these movements as a feed forward control problem.

With this in mind, our proposed optimal control problem was solved once, and

the solution was used for each trial (as done in [12]). We remind the reader that we are not proposing that the biological system is performing an optimization computation each time a movement occurs. Instead, we are suggesting that evolution has already performed the optimization process via some cost function and arrived at a neuromechanical system (the human body) with a construction intrinsically built to minimize jerk.

The selection of a cost function to optimize is crucial and can result in vastly different control policies. If we assume a cost function that evolution has used to optimize animal movement, we can extrapolate a corresponding control policy. In addition, we can draw inferences regarding the nature of the biological mechanisms that might implement such a system. Since minimizing the L_∞ norm of jerk results in a bang-bang control policy, we can hypothesize that simple two-state step functions are utilized to control a biomechanical system. These two-state step functions are desirable because binary control is simple.

In addition, utilizing two-state control policies have been shown to be effective in computational models of movement. Recent computational models of spinomuscular control require only step functions representing supraspinal inputs in order to drive a network to achieve human like movements [51, 52]. Other models have shown that central pattern generators can be driven via step inputs [53].

Similarly, on/off control policies have been observed both *in vivo* and *in vitro* in multiple vertebrates. Complex movements such as walking can be activated by gross on/off stimulation of groups of neurons in the brain stem or the spinal cord. The experiments reported in [8] induced various patterns of locomotion in a spinally transected cat by administering a simple step-like electrical stimulation of the lower region of the cat's spinal cord. The experiments in [54] showed that fictive locomotor patterns could be induced with the use of step-like excitation to either the brain stem or spinal cord of mice. Classes of neurons in the brain stem and lumbar regions of the spinal cord of the mice were genetically engineered to contain channelrhodopsin light gated ion channels. Using this technique, light stimulation (or lack thereof) served as an on/off switch for the genetically modified motor system neurons. Gross "on" stimulation to either the brain stem or the lumbar region of the spinal cord activated a class of neurons in those regions

and induced fictive locomotion patterns. The gross “on” stimulation was a control signal driving the generation of locomotor patterns.

The computational and animal experiments explained above indicate that one or more bang-bang type controllers may exist somewhere in the nervous system. It is still an open question as to how and where their neural implementations exist. The evidence cited here suggests that these controllers may exist in supraspinal centers [6, 51], within the spinal cord [7, 8], or both in the brain stem and spinal cord [54]. Furthermore, it is plausible that populations of bursting neurons could implement a bang-bang control signal [55]. This minimum jerk-based bang-bang control signal could then be converted to a signal representing any lower order derivative (acceleration, velocity, etc) via integration of the signal. Certain populations of neurons are known to perform functions akin to time integration [56]. Integration of a neural signal encoding velocity would lead to a signal encoding position as happens in the ocular-motor system [57, 58]. We stress that these suggestions for our control policy’s neural implementation are merely hypotheses backed by neuro-scientific evidence.

We suggest several important extensions to this work. As discussed earlier, to gauge a more accurate jerk profile than that attained by numerical approximations of higher derivatives, we propose using a manipulandum device where the jerk profile can be recorded directly with high bit precision. As shown in Figure 1 the jerk profile that we propose has step-like features and discontinuities. If the jerk profile of human arm reaching movements has such features, then care must be taken to acquire and digitize the jerk signal appropriately. Only then can we effectively compare the MJ_∞ jerk profile with the acquired jerk profile from human subjects.

Another interesting extension would be to investigate how well MJ_∞ models curved movements as done by MJ_2 and other models [13, 59]. Although the MJ_∞ performs very well at modeling straight point-to-point ballistic movements, and it is likely to perform well with curved movements, it is possible that more exotic control policies might be needed to explain more complex movements. With this in mind, we plan on making simultaneous use of multiple control policies by switching between them or blending them depending on the nature of the task at hand. In addition, to achieve a model that replicates a wider array of human movements (such as perturbed

movements), future research directions should extend this feed-forward model to include feedback (environmental and proprioceptive). Finally, to acquire a deeper understanding of the central nervous system, we propose modeling ballistic movements with neuronal networks in order to test the control signal schemes outlined in this work and their potential neural implementations.

3.6 Notes

Chapter 3, in full, is a modified reprint with minor revisions of the material as it appears in *Neural Networks*. Yazdani, M., G. Gamble, G. Henderson, R. Hecht-Nielsen. "A Simple Control Policy for Achieving Minimum Jerk Trajectories." *Neural Networks*, vol. 27, pp. 74-80, 2012

Chapter 4

Sparse Optimal Control Policies

In this chapter we generalize the bang-bang minimum jerk model of Chapter 3 to a sparse control policy that results from solving a minimum “effort” control problem where the “effort” is measured by the L_∞ norm, and the control signal is any high-order derivative. We call this generalization the sparse minimum effort control problem. We test minimum snap and minimum crackle models with human data and show that such high-order control signals are just as good, if not better, than the minimum jerk control signal. We argue that such high-order derivative control signals can be characterized by high frequency pulse trains (on the millisecond order) which can be easily encoded by spiking neurons. Lastly, we extend the sparse control signals for curved movements.

4.1 Optimal Control Overview

In this section we give an overview of optimal control theory and highlight two optimal control problems: the minimum-time and the minimum-effort control problems. Our overview is meant as means to establish common notation and terminology. For a more in depth overview, see the many excellent texts available [43, 60]. Optimal Control theory is an application of optimization theory to the control of a dynamic “plant” or system. In optimization theory, we seek to find an element in a domain and a constraint set that we wish to minimize (or maximize) a criterion (also referred to as an objective). When the elements in the intersection of the domain and constraint set are functions, the criterion is typically referred to as an objective functional or a cost functional, whereas

when the elements are points in a vector space, the criterion is referred to as an objective function or a cost function. In optimal control, we seek an optimal controller (typically a function of time if the system is continuous or a vector if the system is discrete) that has certain constraints (for example, the controller is limited by a specified amount of power or resources) that minimizes a criterion.

We describe dynamic systems as a set of first-order differential equations:

$$\dot{\mathbf{x}}(t) = \mathbf{a}(\mathbf{x}(t), \mathbf{u}(t), t). \quad (4.1)$$

$\mathbf{x}(t)$ is referred to as the “state” of the system, $\mathbf{u}(t)$ is the controller of the system, and $\mathbf{a}(\cdot, \cdot, t)$ is, in general, a non-linear time dependent function describing the dynamics of the state as determined by the state and controller at time t . We assume that the initial state $\mathbf{x}(t_0)$ and initial time is known. Often the dynamic system is assumed to be linear (or locally linear). In that case, the linear system can be either time variant and expressed as

$$\dot{\mathbf{x}}(t) = \mathbf{A}(t)\mathbf{x}(t) + \mathbf{B}(t)\mathbf{u}(t)$$

or time-invariant and expressed as

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t). \quad (4.2)$$

Given the dynamic system of equation 4.1 and an initial state $\mathbf{x}(t_0)$, we seek a control signal $\mathbf{u}(t)$ to transfer the system to a desired state in a finite time. In practice, the control signal $\mathbf{u}(t)$ is not unconstrained, but rather bounded by the available resources (such as fuel, energy, or supply). In optimal control, we seek an optimal control signal $\mathbf{u}^*(t)$ that, in addition to transferring the system to a desired state, the control signal also minimizes a cost functional $\mathbf{J}(\mathbf{u}(t))$. The cost functional is application dependent and, clearly, the optimal solution $\mathbf{u}^*(t)$ depends on what we consider cost. For example, in the cost functional we may penalize large control signals or penalize deviations from a desired trajectory. Subsequently we will discuss two important cost functionals.

4.1.1 Minimum-Time Control

In the minimum-time control problem, the objective is to transfer a system to a final state with a constrained control signal as quickly as possible. Thus, the cost

functional penalizes the total time it takes to transfer the initial state to a final state and can be expressed as

$$\mathbf{J}(\mathbf{u}(t)) = t_f - t_0 \quad (4.3)$$

where t_f is unknown, the initial state $\mathbf{x}(t_0)$ and initial time t_0 are known, and the system dynamics are described by equation 4.1. We furthermore constrain the control signals to be bounded

$$|\mathbf{u}(t)| \leq B. \quad (4.4)$$

We now show the solution to the minimum-time control problem for linear and time invariant system as described by equation 4.2. We consider \mathbf{A} and \mathbf{B} to be constant $n \times n$ and $n \times m$ matrices respectively. Thus the minimum-time control problem can be expressed as

$$\begin{aligned} & \underset{\mathbf{u}(t), T}{\text{minimize}} && T \\ & \text{subject to} && \dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) \\ & && \mathbf{x}(0) = \mathbf{x}_i \\ & && \mathbf{x}(T) = \mathbf{x}_f \\ & && |\mathbf{u}(t)| \leq B \end{aligned} \quad (4.5)$$

where we have defined $T \equiv t_f - t_0$ and assumed $t_0 = 0$. This special case has been solved by Pontryagin and colleagues and their main results can be summarized in the following two theorems (see [61] for more details):

Existence and Uniqueness Theorem. *If all the eigenvalues of \mathbf{A} have non-positive real parts, then an optimal control policy exists that can transfer a linear time-invariant system from any initial state $\mathbf{x}(t_0)$ to any desired state. Furthermore, if an optimal control signal exists, then it is unique.*

Number of Switches Theorem. *If the eigenvalues of \mathbf{A} are all real and a unique optimal control signal exists, then the control signal can switch at most $n + 1$ times.*

This type of control signal is typically referred to as a “bang-bang” control signal since the signal is step-like and switches between the lower bound ($-B$ in the case of the constraint in equation 4.4) and the upper bound (positive B in the case of the constraint in equation 4.4). This type of control signal can be regarded as a “sparse” control signal

since the number of changes in the signal's values is small. In other words, the changes in the control signal can be described by a small number of Dirac delta functions.

4.1.2 Minimum-Effort Control

In the minimum-effort control problem, the objective is to transfer a system from an initial state to a final state with a control signal that is as “small” as possible (hence, minimum “effort”). Typically, the “size” of a control signal is measured with a penalty function. In this work we consider the L_p norm penalty function and can express the cost functional as

$$\mathbf{J}(\mathbf{u}(t)) = \|\mathbf{u}(t)\|_p \quad (4.6)$$

where $\|\cdot\|_p$ denotes the L_p norm (and is typically the L_2 norm), and the system dynamics are described by equation 4.1. We can also have additional constraints in the minimum-effort control problem, and just as in the minimum-time control problem, there can be many variations by introducing additional constraints or additional costs to the objective. For example, a simple extension would be to consider a control problem where the cost functional trades-off between “effort” and the transfer time and can be expressed as a combination of equations 4.3 and 4.6

$$\mathbf{J}(\mathbf{u}(t)) = \gamma \|\mathbf{u}(t)\|_p + t_f - t_0 \quad (4.7)$$

where $\gamma \geq 0$ is a trade-off parameter between “effort” and the state transfer time and can be varied depending on the application.

A minimum-effort control problem that we have considered in describing human movements is the model originally introduced by Flash and Hogan:

$$\begin{aligned} & \underset{u(t)}{\text{minimize}} && \|u(t)\|_2 \\ & \text{subject to} && \dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}u(t) \\ & && \mathbf{x}(0) = \mathbf{x}_i \\ & && \mathbf{x}(T) = \mathbf{x}_f \end{aligned} \quad (4.8)$$

where $\mathbf{x}(t) = \begin{bmatrix} x(t) & \dot{x}(t) & \ddot{x}(t) \end{bmatrix}^T$ is the state vector, \mathbf{x}_i and \mathbf{x}_f are the initial and final boundary conditions, and T is the duration of the movement (with movement starting

at time $t = 0$). Flash and Hogan used a jerk control signal ($u(t) = \ddot{x}(t)$), and furthermore used a third-order integrator model for the linear time-invariant dynamic equation parameters:

$$\mathbf{A} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} \text{ and } \mathbf{B} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}. \quad (4.9)$$

This simple model yields trajectories that are remarkably similar to those of humans. Naturally, simple extensions of this optimization problem can yield results that are even more realistic and many researches have begun exploring extensions. For example, [62] has noted that when humans make movements to a target, the target that is reached is a not a specific point, but rather a distribution of points. Hence, in their optimization procedure they relaxed the constraints of equation 4.8.

4.1.3 Some Remarks on Optimal Control and Optimization

Optimal control theory and optimization have been applied to a diverse set of fields. Caution must be taken when it is claimed or suggested that a control policy is “optimal.” A control policy is only optimal with respect to the objective and the constraint set of the optimization problem. Therefore, in some sense, every optimal control policy comes with the caveat that it is dependent on the objective and constraint in the design of the problem. If an optimal control policy does not meet a desired specification, this is because the optimization problem has not been appropriately defined. For example, some have criticized the Flash and Hogan model of equation 4.8 that, for example, the control policy does not allow any variations in the final target (when it has been shown that humans demonstrate reaching a distribution of targets as opposed to a single precise target). This does not imply that the control policy that Flash and Hogan suggested is “sub-optimal,” but rather that characteristics that are sought after have not been appropriately defined.

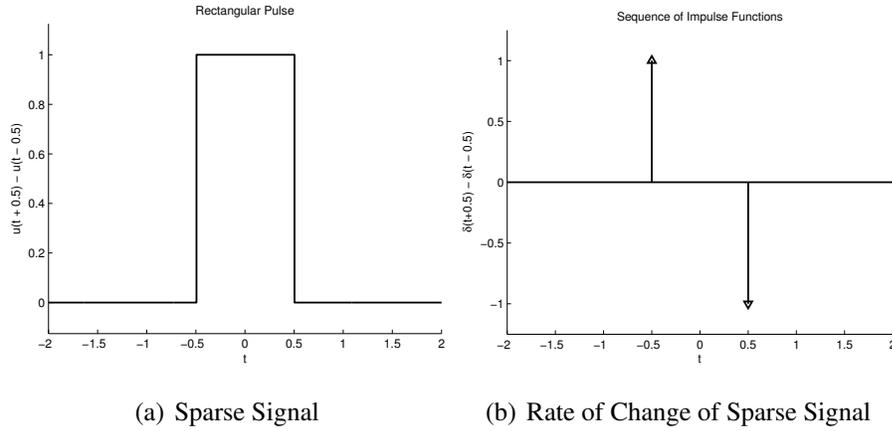


Figure 4.1: Example of a sparse signal and its spike encoding.

4.2 Sparse Optimal Control Policies for Straight Point-to-Point Trajectories

We define a sparse optimal control policy as a control policy that meets an optimality condition and is also sparse. Here we define a sparse signal as a signal where the number of changes in the signal is small. For example, a Heaviside step function is a very sparse signal since the changes in this signal is described by a unit impulse function. Similarly, a rectangular pulse is characterized by two impulse functions (see Figure 4.1). The “bang-bang” control signals in the minimum-time control problem discussed above are examples of a sparse control signal that are also optimal in that they minimize the time of state transfer for a system.

Here we discuss sparse optimal control signals that solve the minimum-effort problem. The control signal is defined as the n -th order derivative in terms of position $x(t)$,

$$u_n(t) = \frac{d^n}{dt^n} x(t). \quad (4.10)$$

The minimum effort control problem that results in sparse control signals uses the L_∞ norm and is written as

$$\begin{aligned}
& \underset{u_n(t)}{\text{minimize}} && \|u_n(t)\|_\infty \\
& \text{subject to} && \dot{\mathbf{x}}_n(t) = \mathbf{A}_n \mathbf{x}_n(t) + \mathbf{B}_n u_n(t) \\
& && \mathbf{x}_n(0) = \mathbf{x}_i \\
& && \mathbf{x}_n(T) = \mathbf{x}_f
\end{aligned} \tag{4.11}$$

where $\mathbf{x}_n(t) = \left[x(t) \quad \frac{d}{dt}x(t) \quad \frac{d^2}{dt^2}x(t) \quad \dots \quad \frac{d^{n-1}}{dt^{n-1}}x(t) \right]^T$ is the state vector, \mathbf{x}_i and \mathbf{x}_f are the initial and final boundary conditions, and T is the duration of the movement (with movement starting at time $t = 0$). Here we consider a system that is an n -th order integrator, thus

$$\mathbf{A}_n = \begin{bmatrix} \mathbf{0}_{(n-1) \times 1} & \mathbf{I}_{(n-1) \times (n-1)} \\ 0 & \mathbf{0}_{1 \times (n-1)} \end{bmatrix} \text{ and } \mathbf{B}_n = \begin{bmatrix} \mathbf{0}_{(n-1) \times 1} \\ 1 \end{bmatrix}. \tag{4.12}$$

The L_∞ norm can be computed by solving $\underset{0 \leq t \leq T}{\text{maximize}} |u(t)|$. In Chapter 3 we showed that when $u(t) = \ddot{x}(t)$ and $n = 3$, the sparse control signal explains the trajectories of human movements better than the traditional Flash and Hogan model of equation 4.8. Furthermore, and perhaps more important, sparse control signals are biologically more realistic than non-sparse signals.

We now show the general analytic solution for equation 4.11. To derive the general solution, we assume that the boundary conditions are

$$\mathbf{x}_i = \begin{bmatrix} x_i \\ \mathbf{0}_{(n-1) \times 1} \end{bmatrix} \text{ and } \mathbf{x}_f = \begin{bmatrix} x_f \\ \mathbf{0}_{(n-1) \times 1} \end{bmatrix}. \tag{4.13}$$

That is, we assume that the movement starts at rest and ends at rest. We solve the general sparse minimum effort control problem by manipulating equation 4.11 to a form that has been previously solved. Namely, note that every optimization problem can be written equivalently as a linear optimization problems by introducing an auxiliary variable K

and we can equivalently express equation 4.11 as follows:

$$\begin{aligned}
& \underset{u_n(t), K}{\text{minimize}} && K \\
& \text{subject to} && \dot{\mathbf{x}}_n(t) = \mathbf{A}_n \mathbf{x}_n(t) + \mathbf{B}_n u_n(t) \\
& && \mathbf{x}_n(0) = \mathbf{x}_i \\
& && \mathbf{x}_n(T) = \mathbf{x}_f \\
& && \|u_n(t)\|_\infty \leq K
\end{aligned} \tag{4.14}$$

where $u_n(t)$, \mathbf{A}_n , \mathbf{B}_n , \mathbf{x}_i , and \mathbf{x}_f are defined as before in equations 4.10, 4.12, and 4.13 respectively. The equivalency between equations 4.11 and 4.16 is due to the fact that every objective can be bounded, and this bound is expressed as an additional constraint in equation 4.16. We can simplify this bound as follows:

$$\begin{aligned}
& \|u_n(t)\|_\infty \leq K \\
\implies & \underset{0 \leq t \leq T}{\text{maximize}} |u_n(t)| \leq K \\
\implies & |u_n(t)| \leq K.
\end{aligned} \tag{4.15}$$

In words, the result of equations 4.15 is due to the fact that if the largest element of $|u_n(t)|$ for $0 \leq t \leq T$ is less than or equal to K , then every element in $|u_n(t)|$ for $0 \leq t \leq T$ is less than or equal to K , that is $|u_n(t)| \leq K$ for $0 \leq t \leq T$. We can therefore write equation 4.16 as

$$\begin{aligned}
& \underset{u_n(t), K}{\text{minimize}} && K \\
& \text{subject to} && \dot{\mathbf{x}}_n(t) = \mathbf{A}_n \mathbf{x}_n(t) + \mathbf{B}_n u_n(t) \\
& && \mathbf{x}_n(0) = \mathbf{x}_i \\
& && \mathbf{x}_n(T) = \mathbf{x}_f \\
& && |u_n(t)| \leq K.
\end{aligned} \tag{4.16}$$

The optimization problem of equation 4.16 has the same form as equation 4.5. We can therefore use the results from the minimum-time control problem and apply them here (namely that the results of Pontryagin and colleagues still hold). The difference is that in equation 4.5 the unknown is time T , whereas in equation 4.16 the unknown is the bound K on the control signal $u_n(t)$. Since the dynamic system in equation 4.16 is an n -th order integrator, we can use the result from [63] and write the following Theorem:

Number of Switches for an N-th Order Integrator Theorem 1. *For a control problem of the type in equations 4.5 or 4.16 where the system dynamic equations are an n -th order integrator (as in equation 4.12), then the number of switchings in the control signal is exactly $n + 1$ and the control signal is symmetric.*

In other words, as the order of the control signal goes up (that is, as n increases) then the number of switchings in the control signal increases by the same amount. [64] solved the general n -th order minimum-time control problem of equation 4.5 for an n -th order integrator. We adapt their results for the general n -th order minimum-effort control problem of equation 4.16 and summarize the solution as follows:

$$K_n^* = \frac{2^{2(n-1)}(n-1)!(x_f - x_i)}{T^n} \quad (4.17)$$

$$t_i^* = T \sin^2 \left(\frac{\pi i}{2n} \right), \quad i = 0, \dots, n \quad (4.18)$$

where t_i^* denote the optimal switching times. Figure 4.2 shows examples of several sparse optimal control signals.

4.2.1 Sparse Optimal Control Signals in Fast Human Movements

The sparse optimal control signals introduced in the previous section are not only optimal with respect to a minimum-effort objective, but are also biologically more realistic. These sparse optimal control signals can be efficiently represented with neuronal spikes. If we treat each spike as an idealized Dirac delta function [65], then a spike sequence that represents an n -th order optimal control signal can be expressed as

$$\rho_n(t) = K_n^* \sum_{i=0}^n (-1)^i \delta(t - t_i^*) \quad (4.19)$$

where K_n^* and t_i^* can be found from equations 4.17 and 4.18 respectively. The spike train represented by equation 4.19 is not from a single neuron, but rather a population of excitatory and inhibitory neural network. The exact architecture of such a neural network is unknown, but the construction of such a network to generate spikes as expressed in equation 4.19 is trivial.

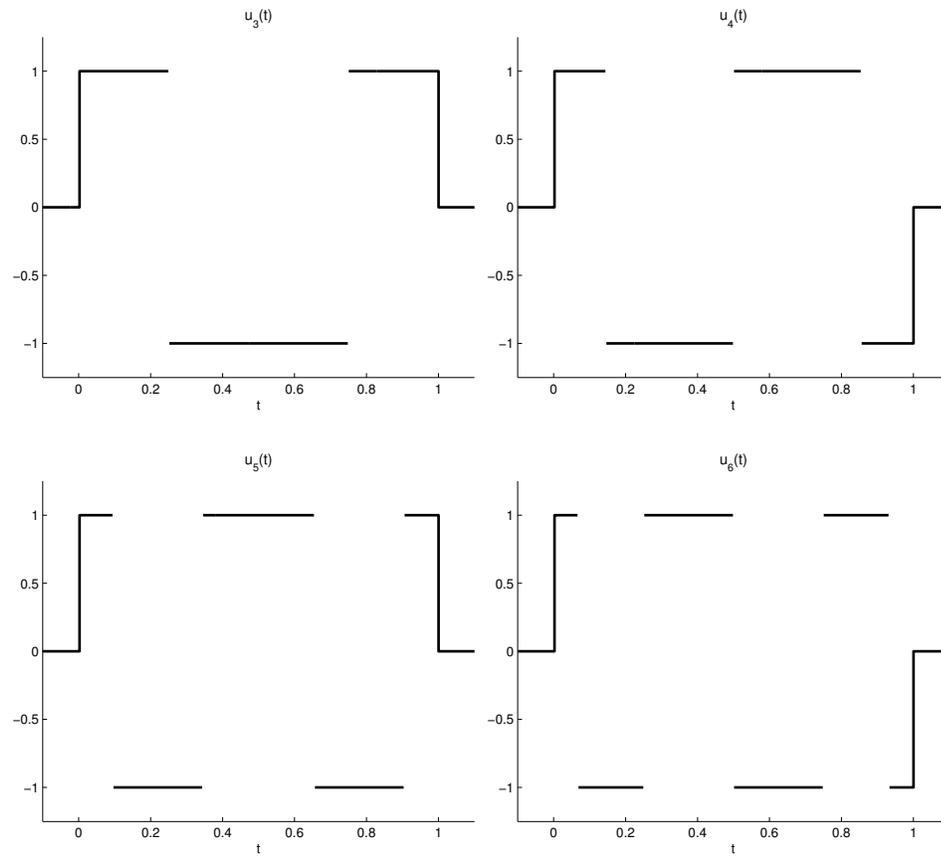


Figure 4.2: Examples of sparse optimal control signals $u_i(t)$ for $i = 3, \dots, 6$. Shown are movements that starts from $t = 0$ and end at $t = 1$.

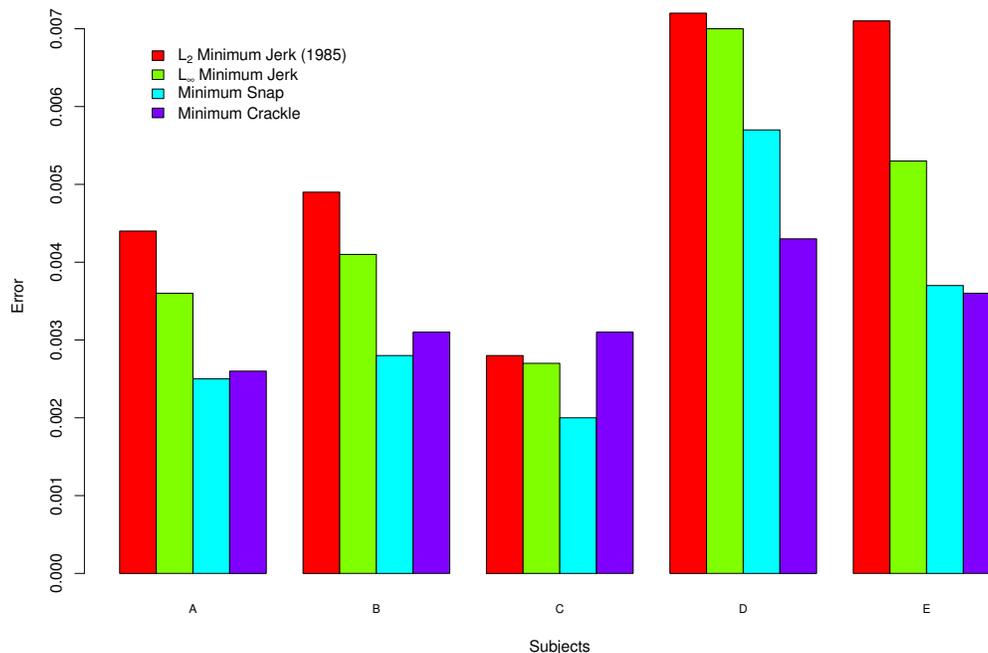


Figure 4.3: Average Mean Squared Error for all trials for each subject.

In Chapter 3 we showed that a sparse optimal control signal that corresponds to jerk (expressed as $\rho_3(t)$ in the notation of equation 4.19) explains fast human movements better than the smooth control signal that results when using the L_2 norm. We now also propose that the control signals the brain uses are not limited to the jerk control signal. There is nothing preventing the brain from having a high-order control signal (see Figure 4.3 for the error of using a higher-order derivative control signal). With each increase in the order of the control signal, the number of spikes increases. Furthermore, as shown in Figure 4.2, these high-order derivative control signals can form a basis set for a subspace of fast human movements. In the literature such basis sets are often referred to as motor primitives [39].

4.2.2 Experimental Setup

The human arm movement data for this work was originally collected by Karniel and Mussa-Ivaldi and used in their 2002 paper to investigate the nervous system's abil-

ity to adapt to perturbations. We used a subset of this data that was relevant for our study of fast movements (the baseline unperturbed movements) and summarize their experimental setup below and refer the reader to [40] for a more complete description.

Five subjects participated in an experiment involving a manipulandum that restricted their movements to a two-dimensional horizontal plane of the seated subjects (subjects participated separately in these experiments). During each trial, the subject watched a screen that displayed the position of their hand and the manipulandum in relation to three positional markers A, B and C. Each marker was separated by 10 cm and formed an equilateral triangle as shown in Figure 2.3(a). For each trial, the subject was instructed to move the on-screen representation of the manipulandum from one target to another in about one third of a second with a tolerance of ± 50 ms. At the end of each trial feedback was given indicating if the subject had reached the target and also if the execution of their movement was within the allowed time window. The trajectories were recorded for all six possible movement types for all subjects over the course of four days. The experiments also included trials where the arm was perturbed during movement. We excluded these trials from our study and only considered unperturbed trials. For each trial, we only select the so-called “ballistic” portion of the movement. That is, we select the portion of the trial where movement had started and the movement had completed its “feedforward” portion. There are various methods for movement onset and offset detection [48, 49, 50], and there is no standard technique for choosing the relevant portion of a movement since the definition of what is relevant may change from study to study or from one movement type to another.

We approach finding the start and end of movement by finding the point in time when the velocity has reached its peak velocity. Fast movement always have a unique global maximum in the trial (unless the trial is an outlier) so finding the time at which this maximum occurs is unique. Once this point in time is found, we proceed to consider velocity samples before and after the peak velocity and check to see if they fall below a pre-determined threshold. Once the velocity sample is below this threshold, we refer to this sample as the “start” or “end” depending on whether the sample is before or after the peak velocity. We have considered different thresholds and even considered using the onset and offset detection method of Chapter 3, however the qualitative results

do not change hence our overall results in this paper are not sensitive to the method used in finding the onset and offsets of movement. Regardless of what onset and offset detection method used, real trials do not neatly satisfy the boundary conditions assumed in theorems 1 and 2. That is, it is very rare for the trial to start and end at complete rest (with velocity and higher derivatives equal to zero). As a result we apply theorems 1 and 2 by not arbitrarily set the boundary conditions to zero, but rather set the boundary conditions to the velocities and higher derivatives of the given trial. Again, this change does not have an impact on the qualitative results, but we make this change since it fits the assumptions of data more accurately.

4.3 Sparse Optimal Control Policies for Curved Trajectories

The sparse optimal control policy for straight trajectories that we showed in the previous section can also be extended to curved trajectories. We introduce a via-point into the constraint set of equation 4.11:

$$\begin{aligned}
 & \underset{u_n(t)}{\text{minimize}} && \|u_n(t)\|_\infty \\
 & \text{subject to} && \dot{\mathbf{x}}_n(t) = \mathbf{A}_n \mathbf{x}_n(t) + \mathbf{B}_n u_n(t) \\
 & && \mathbf{x}_n(0) = \mathbf{x}_i \\
 & && \mathbf{x}_n(t_v) = \mathbf{x}_v \\
 & && \mathbf{x}_n(T) = \mathbf{x}_f
 \end{aligned} \tag{4.20}$$

where $\mathbf{x}_n(t_v) = \mathbf{x}_v$ specifies the state (denoted \mathbf{x}_v) at time t_v . The time t_v can be inferred from the data. This is the same formulation that Flash and Hogan proposed for curved movements in [12], but here we use the L_∞ norm instead. By using the L_∞ norm, we have sparse optimal control signals. The optimization problem of 4.20 can be solved numerically with the techniques discussed in Chapter 5. Figure 4.4 shows an example simulation of a curved trajectory specified by a via-point. Note that even though the movement has a curved trajectory, the control signal (jerk in this case) is still sparse. Future work should test this sparse control signal for curved movements.

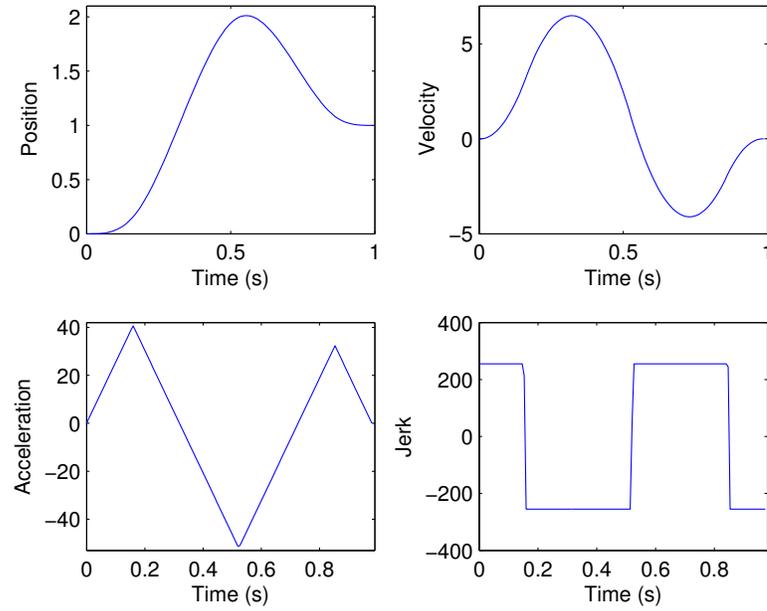


Figure 4.4: A curved trajectory specified by a via point.

4.4 Notes

Chapter 4, in part, has been submitted for publication of the material. The dissertation author is the primary investigator and author of this material.

Chapter 5

Multiple Cost Funcions

In this chapter we extend the sparse minimum effort control problem of Chapter 4 to multiple composite cost functions. We first begin by giving a numerical solution to the minimum effort control problem. This is so we can use numerical methods and optimization techniques and no longer need to rely on analytic solutions¹. Using numerical techniques, we solve the minimum effort control problem with multiple control signals that are combined.

5.1 Numerical Solutions to the Minimum Effort Control Problem

In this section we formulate numerical solutions to solving the minimum effort control problem. We define the minimum effort control problem as a control signal that corresponds to a high order derivative of the position of the end effector and define the cost function as minimizing the norm of the control signal. Formally, we define the control signal $u(t)$ in terms of the position $x(t)$ as

$$u_n(t) = \frac{d^n}{dt^n}x(t). \quad (5.1)$$

The minimum effort control problem then becomes

¹In general, closed-form analytic solutions do not exist for a general optimal control problem.

$$\begin{aligned}
& \underset{u_n(t)}{\text{minimize}} && \| u_n(t) \|_p \\
& \text{subject to} && \dot{\mathbf{x}}_n(t) = A\mathbf{x}_n(t) + Bu(t) \\
& && \mathbf{x}_n(0) = \mathbf{x}_i \\
& && \mathbf{x}_n(T) = \mathbf{x}_f
\end{aligned} \tag{5.2}$$

where $\mathbf{x}_n(t) = \left[x(t) \quad \frac{d}{dt}x(t) \quad \frac{d^2}{dt^2}x(t) \quad \dots \quad \frac{d^{n-1}}{dt^{n-1}}x(t) \right]^T$ is the state vector, \mathbf{x}_i and \mathbf{x}_f are the initial and final boundary conditions, and T is the duration of the movement (with movement starting at time $t = 0$). Here we consider a system that is an n -th order integrator, thus $A = \begin{bmatrix} \mathbf{0}_{(n-1) \times 1} & \mathbf{I}_{(n-1) \times (n-1)} \\ 0 & \mathbf{0}_{1 \times (n-1)} \end{bmatrix}$ and $B = \begin{bmatrix} \mathbf{0}_{(n-1) \times 1} \\ 1 \end{bmatrix}$.

To reformulate the problem in discrete time, we define the vector \mathbf{x} as follows:

$$\begin{aligned}
x[k] &= x(k\Delta T), \quad k = 0, \dots, N \\
\mathbf{x} &= \begin{bmatrix} x[0] & x[1] & \dots & x[N] \end{bmatrix}
\end{aligned}$$

That is each element of the vector \mathbf{x} is a sample from the continuous time position trajectory $x(t)$ sampled at intervals of ΔT . Clearly, to have $N + 1$ samples, we must have $\Delta T = \frac{T}{N}$. We can approximate the derivative of $x(t)$ using \mathbf{x} with the difference equation $(\Delta T)^{-1}(x[i] - x[i - 1])$ for $i = 1, \dots, N$. If we define the $N \times (N + 1)$ matrix D_1 as:

$$D_1 \equiv (\Delta T)^{-1} \begin{bmatrix} -1 & 1 & 0 & \dots & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & \dots & 0 & 0 \\ 0 & 0 & -1 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & \dots & \dots & \dots & 0 & -1 & 1 \end{bmatrix}$$

then we can write the velocity as $D_1\mathbf{x}$. We can also write the velocity using the second difference $(\Delta T)^{-2}(x[i + 1] - 2x[i] + x[i - 1])$ and we can write the acceleration matrix

D_2 as

$$D_2 \equiv (\Delta T)^{-2} \begin{bmatrix} 1 & -2 & 1 & 0 & \cdots & 0 & 0 & 0 & 0 \\ 0 & 1 & -2 & 1 & \cdots & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 2 & \cdots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & 1 & -2 & 1 \end{bmatrix}$$

We can similarly write the jerk as the difference equation $(\Delta T)^{-3}(x[i+2] - 3x[i+1] + 3x[i] - x[i-1])$ for $i = 1, \dots, N-2$. If we then define the $(N-2) \times (N+1)$ matrix D_3 as:

$$D_3 \equiv (\Delta T)^{-3} \begin{bmatrix} -1 & 3 & -3 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & -1 & 3 & -3 & 1 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & -1 & 3 & -3 & 1 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & -1 & 3 & -3 & 1 & 0 \\ 0 & 0 & 0 & 0 & \cdots & 0 & -1 & 3 & -3 & 1 \end{bmatrix}$$

Thus the jerk can be computed as $D_3 \mathbf{x}$.

It can be shown that the general n -th order derivative matrix D_n is an $(N-n+1) \times (N+1)$ Toeplitz matrix where the non-zero elements of the matrix are found from the binomial coefficients. To write this more generally, we introduce some notation. We define $p_i \in \mathbb{R}^{1 \times (i+1)}$ as a row-vector with elements belonging to the i th row of a modified Pascal triangle:

$$\begin{array}{l} i = 0: \qquad \qquad \qquad 1 \\ i = 1: \qquad \qquad -1 \quad 1 \\ i = 2: \qquad \qquad 1 \quad -2 \quad 1 \\ i = 3: \qquad -1 \quad 3 \quad -3 \quad 1 \\ i = 4: \quad 1 \quad -4 \quad 6 \quad -4 \quad 1 \end{array}$$

For example, $p_1 = [-1 \ 1]$, $p_2 = [1 \ -2 \ 1]$, and $p_3 = [-1 \ 3 \ -3 \ 1]$. D_1 then, as shown above, is the $N \times (N+1)$ Toeplitz matrix where the first row is $[p_1 \ \mathbf{0}_{1 \times (N-1)}]$, D_2 is the $(N-1) \times (N+1)$ Toeplitz matrix where the first row is

$\begin{bmatrix} p_2 & \mathbf{0}_{1 \times (N-2)} \end{bmatrix}$, and D_3 is the $(N-2) \times (N+1)$ Toeplitz matrix where the first row is $\begin{bmatrix} p_3 & \mathbf{0}_{1 \times (N-3)} \end{bmatrix}$. Therefore, we can compactly write D_n as the $(N-n) \times (N+1)$ Toeplitz matrix where the first row is $\begin{bmatrix} p_n & \mathbf{0}_{1 \times (N-n)} \end{bmatrix}$. Clearly the derivative order must satisfy $n < N$, otherwise the notion of discrete-time derivative is meaningless here.

We can now formulate the discrete-time counterpart to equation 5.2 as follows:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \| D_n \mathbf{x} \|_p \\ & \text{subject to} && A_{eq_n} \mathbf{x} = b_{eq_n} \end{aligned} \quad (5.3)$$

where A_{eq_n} and b_{eq_n} are defined respectively as:

$$A_{eq_n} \equiv \begin{bmatrix} 1 & \mathbf{0}_{1 \times (N)} \\ \mathbf{0}_{1 \times (N)} & 1 \\ (\Delta T)^{-1} p_1 & \mathbf{0}_{1 \times (N-1)} \\ \mathbf{0}_{1 \times (N-1)} & (\Delta T)^{-1} p_1 \\ (\Delta T)^{-2} p_2 & \mathbf{0}_{1 \times (N-2)} \\ \mathbf{0}_{1 \times (N-2)} & (\Delta T)^{-2} p_2 \\ \vdots & \vdots \\ (\Delta T)^{-(n-1)} p_{n-1} & \mathbf{0}_{1 \times (N-(n-1))} \\ \mathbf{0}_{1 \times (N-(n-1))} & (\Delta T)^{-(n-1)} p_{n-1} \end{bmatrix} \quad \text{and } b_{eq_n} \equiv \begin{bmatrix} x_i \\ x_f \\ \dot{x}_i \\ \dot{x}_f \\ \ddot{x}_i \\ \ddot{x}_f \\ \vdots \\ x_i^{(n-1)} \\ x_f^{(n-1)} \end{bmatrix}.$$

Note that A_{eq_n} is a $2n \times (N+1)$ matrix and b_{eq_n} is a $2n \times 1$ column vector. The equality constraint of equation 5.3 essentially constrains the solution set to the boundary conditions of the trajectory that are determined by the initial and final positions, and the initial and final $(n-1)$ derivatives. Note that when the minimum effort control problem was solved previously, we only considered the case where the initial and final $(n-1)$ derivatives had the value of zero (the movement starts and ends at ‘‘rest’’). Here, however, we relax this requirement since we are not computationally restricted from considering the more general (and realistic) case of non-zero boundary conditions.

5.1.1 The solution to the minimum effort control problem using the l_2 norm

We now proceed to solve equation 5.3 for the case that $p = 2$ (that is, the case that the l_2 norm is used)

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} \quad \| D_n \mathbf{x} \|_2 \\ & \text{subject to} \quad A_{eqn} \mathbf{x} = b_{eqn} \end{aligned} \quad (5.4)$$

using the method of Lagrange multipliers. We formulate the Lagrangian for equation 5.3 as follows:

$$L(\mathbf{x}, \lambda) = \mathbf{x}^T D_n^T D_n \mathbf{x} + \lambda^T (A_{eqn} \mathbf{x} - b_{eqn}). \quad (5.5)$$

We formulate the dual function $g(\lambda)$ by solving the following problem:

$$g(\lambda) = \underset{\mathbf{x}}{\text{minimize}} L(\mathbf{x}, \lambda). \quad (5.6)$$

The solution to the optimization problem of equation 5.6 (that is, the dual function $g(\lambda)$) can be found by setting $\frac{\partial}{\partial \mathbf{x}} L(\mathbf{x}, \lambda)$ to zero and solving for \mathbf{x} :

$$\begin{aligned} \frac{\partial}{\partial \mathbf{x}} L(\mathbf{x}, \lambda) &= 2D_n^T D_n \mathbf{x} + A_{eqn}^T \lambda = 0 \\ \implies \mathbf{x} &= -\frac{1}{2} (D_n^T D_n)^{-1} A_{eqn}^T \lambda \end{aligned} \quad (5.7)$$

Note that since equation 5.5 is convex (due the fact that $D_n^T D_n \succeq 0$), this solution is the true minimum. Thus, we can find the dual function by evaluating equation 5.5 at this minimum:

$$\begin{aligned} g(\lambda) &= \left. \frac{\partial}{\partial \mathbf{x}} L(\mathbf{x}, \lambda) \right|_{\mathbf{x} = -\frac{1}{2} (D_n^T D_n)^{-1} A_{eqn}^T \lambda} \\ &= -\frac{1}{2} \lambda^T A_{eqn} (D_n^T D_n)^{-1} A_{eqn}^T \lambda - \lambda^T b_{eqn} \end{aligned} \quad (5.8)$$

We now evaluate the maximum of the dual function as follows:

$$\begin{aligned} \frac{\partial}{\partial \lambda} g(\lambda) &= -\frac{1}{2} A_{eqn} (D_n^T D_n)^{-1} A_{eqn}^T \lambda - b_{eqn} = 0 \\ \implies \lambda &= -2 (A_{eqn} (D_n^T D_n)^{-1} A_{eqn}^T)^{-1} b_{eqn} \end{aligned} \quad (5.9)$$

Again note that this is the true maximum of the dual function since $g(\lambda)$ is a concave function (due to the fact that $-\frac{1}{2} A_{eqn} (D_n^T D_n)^{-1} A_{eqn}^T \leq 0$). Thus, we can now evaluate

equation 5.5 at this solution and solve for \mathbf{x} to find the solution to equation 5.3:

$$\begin{aligned} \frac{\partial}{\partial \mathbf{x}} L(\mathbf{x}, \lambda) \Big|_{\lambda = -2(A_{eqn}(D_n^T D_n)^{-1} A_{eqn}^T)^{-1} b_{eqn}} &= 0 \\ \implies \mathbf{x}^* &= (D_n^T D_n)^{-1} A_{eqn}^T (A_{eqn}(D_n^T D_n)^{-1} A_{eqn}^T)^{-1} b_{eqn} \end{aligned} \quad (5.10)$$

Since equation 5.4 is a convex optimization problem (the objective function is the norm of a linear function, thus a convex function, and the constraints are linear), the solution \mathbf{x}^* as expressed in equation 5.10 is the true global minimum.

5.1.2 The solution to the minimum effort control problem using the l_∞ norm

We now proceed to solve equation 5.3 for the case that $p = \infty$ (that is, the cost of the control signal is measured using the l_∞ norm):

$$\begin{aligned} \underset{\mathbf{x}}{\text{minimize}} \quad & \| D_n \mathbf{x} \|_\infty \\ \text{subject to} \quad & A_{eqn} \mathbf{x} = b_{eqn} \end{aligned} \quad (5.11)$$

We solve this optimization problem by formulating equation 5.11 as a Linear Program that can then be solved with an off-the-shelf software package. To pose equation 5.11 as a Linear Program, we introduce a scalar auxiliary variable $s \in \mathbb{R}$ and solve the following optimization problem that is equivalent to equation 5.11:

$$\begin{aligned} \underset{\mathbf{x}, s}{\text{minimize}} \quad & s \\ \text{subject to} \quad & A_{eqn} \mathbf{x} = b_{eqn} \\ & \| D_n \mathbf{x} \|_\infty \leq s \end{aligned} \quad (5.12)$$

Similarly, since $\| D_n \mathbf{x} \|_\infty \leq s \implies -s\mathbf{1} \preceq D_n \mathbf{x} \preceq s\mathbf{1}$ (here \preceq denotes an element-by-element inequality and $\mathbf{1}$ denotes a vector of all 1's), we can write equation 5.13 equivalently as

$$\begin{aligned} \underset{\mathbf{x}, s}{\text{minimize}} \quad & s \\ \text{subject to} \quad & A_{eqn} \mathbf{x} = b_{eqn} \\ & -s\mathbf{1} \preceq D_n \mathbf{x} \preceq s\mathbf{1} \end{aligned} \quad (5.13)$$

which is also equivalent to equation 5.11. Equation 5.13, however, is a Linear Program since the objective function, the equality constraint, and the inequality constraints are all linear. This can be solved with any off-the-shelf Linear Program solver (such as [66, 67]).

5.1.3 Minimum Jerk as a Minimum Effort Control Problem

Here we give a specific example of a control signal and solve the control problem of equation 5.3. We consider the control signal to be the third derivative of position, or jerk, and consider an initial position of 0 and a final position of 1 units. The duration of the movement is 1 sec and the number of samples is 50 (thus, $i = 0, 1, \dots, 49$). The movement starts at rests, and ends at rest. Hence, the objective function and constraints of equation 5.3 are

$$\| D_3 \mathbf{x} \|_p$$

and

$$A_{eq3} = \begin{bmatrix} 1 & \mathbf{0}_{1 \times 49} \\ \mathbf{0}_{1 \times 49} & 1 \\ (\Delta T)^{-1} p_1 & \mathbf{0}_{1 \times 48} \\ \mathbf{0}_{1 \times 48} & (\Delta T)^{-1} p_1 \\ (\Delta T)^{-2} p_2 & \mathbf{0}_{1 \times 47} \\ \mathbf{0}_{1 \times 47} & (\Delta T)^{-2} p_2 \end{bmatrix} \text{ and } b_{eq3} = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \quad (5.14)$$

where $\Delta T = 0.02$. Having defined the objective function and the constraint matrices, we solve equation 5.3 for $p = 2$ and $p = \infty$ with the methods described above. The solved solution is shown in Figure 5.1 where the blue shows the solution to equation 5.3 with $p = \infty$ and red shows the solution with $p = 2$.

5.2 Minimum Effort Control with a Hybrid Metric

The previous section illustrated the numerical solutions to the minimum effort control problem for the $p = 2$ and $p = \infty$ case. The solutions show that when $p = \infty$, the minimum effort control problem of equation 5.2 results in a sparse ‘‘bang-bang’’

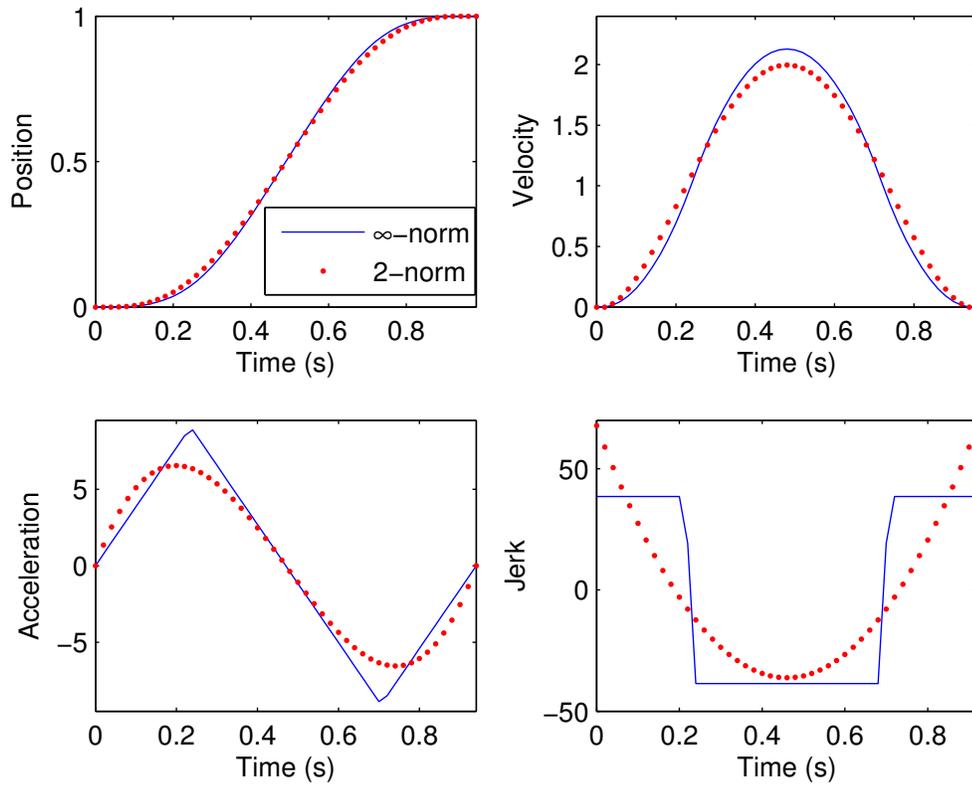


Figure 5.1: The solution to equation 5.3 with $n = 3$ and equality matrices as defined by equation 5.14. The blue line shows the solution to equation 5.3 with $p = \infty$, while the red line shows the solution to equation 5.3 with $p = 2$.

control policy, while when $p = 2$ the control signal is smooth and non-sparse. [68] have suggested a control signal that appears to be a trade-off between the smooth and sparse control policies. They advocate that such a control policy has certain advantages for biological control. For example, such a control policy has the benefit that it is partially hard-limited, but has fewer discontinuities than the sparse optimal control signal derived from the control problem in equation 5.11. Here we show that the proposed control signal of [68] can be derived by solving an optimization problem that trades-off between the cost function measuring the “effort” of the control signal with the l_2 norm and the l_∞ norm. This can be written as follows:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && (1 - \theta) \| D_n \mathbf{x} \|_\infty + \theta \| D_n \mathbf{x} \|_2 \\ & \text{subject to} && A_{eqn} \mathbf{x} = b_{eqn} \end{aligned} \quad (5.15)$$

where $0 \leq \theta \leq 1$ is a “free” trade-off parameter to be selected.

When θ is close to one, the $\| D_n \mathbf{x} \|_2$ term in the cost function of equation 5.15 dominates and the solution to the optimization problem is close to the solution of equation 5.4. Similarly, when θ is close to zero, the $\| D_n \mathbf{x} \|_\infty$ term in the cost function of equation 5.15 dominates and the solution is close to the solution of equation 5.11. The optimization problem of equation 5.15 can be solved by introducing the auxiliary variable s and solving the following equivalent problem:

$$\begin{aligned} & \underset{\mathbf{x}, s}{\text{minimize}} && (1 - \theta)s + \theta \| D_n \mathbf{x} \|_2 \\ & \text{subject to} && A_{eqn} \mathbf{x} = b_{eqn} \\ & && \| D_n \mathbf{x} \|_\infty \leq s \end{aligned} \quad (5.16)$$

As before, since $\| D_n \mathbf{x} \|_\infty \leq s \implies -s\mathbf{1} \preceq D_n \mathbf{x} \preceq s\mathbf{1}$ and $\| D_n \mathbf{x} \|_2 = \mathbf{x}^T D_n^T D_n \mathbf{x}$, we can write equation 5.16 equivalently as

$$\begin{aligned} & \underset{\mathbf{x}, s}{\text{minimize}} && (1 - \theta)s + \theta \mathbf{x}^T D_n^T D_n \mathbf{x} \\ & \text{subject to} && A_{eqn} \mathbf{x} = b_{eqn} \\ & && -s\mathbf{1} \preceq D_n \mathbf{x} \preceq s\mathbf{1} \end{aligned} \quad (5.17)$$

Equation 5.17 is a Quadratic Program (QP) since the objective function is quadratic, and the equality and inequality constraints are linear. There are numerous off-the-shelf

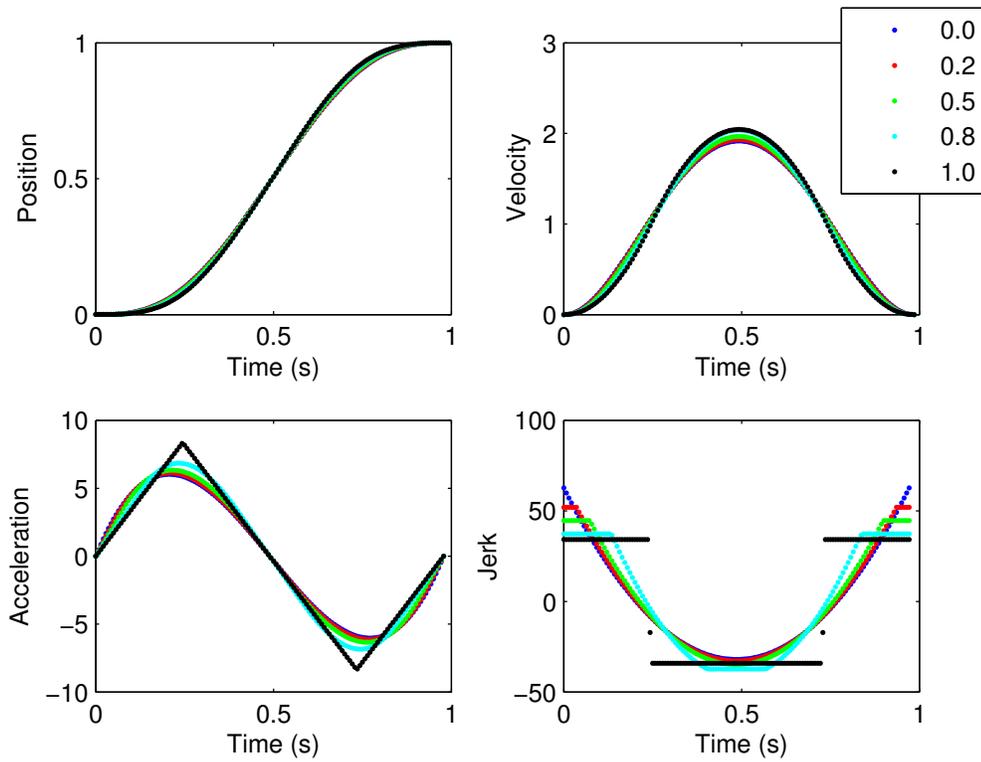


Figure 5.2: The control policies resulting from varying θ from zero to one for the optimization problem in equation 5.15.

software packages that solve QPs. Figure 5.2 shows equation 5.17 solved many times as θ is varied from zero to one. Clearly, we can see that as θ increases the control signal ranges from the sparse policy of equation 5.11 to the smooth policy of equation 5.4. The disadvantage of such a model is that selecting the appropriate θ is not clear.

5.3 Minimum Acceleration Criterion with Constraints

Karniel and colleagues proposed a model they called the Minimum Acceleration Criterion with Constraints (MACC) that they defined as follows:

$$\begin{aligned}
 & \underset{u_3(t)}{\text{minimize}} && \| \ddot{x}(t) \|_2 \\
 & \text{subject to} && \dot{\mathbf{x}}_3(t) = A\mathbf{x}_3(t) + Bu(t) \\
 & && \mathbf{x}_3(0) = \mathbf{x}_i \\
 & && \mathbf{x}_3(T) = \mathbf{x}_f \\
 & && \| u_3(t) \| \leq \beta
 \end{aligned}$$

where β is a free parameter that hard-limits the control signal $u_3(t) = \ddot{x}(t)$. Karniel and colleagues showed that the solution to this optimization problem is a bang-bang (that is, “sparse”) control policy, where the switching times are controlled by the β parameter. Here we will show that the MACC is yet another minimum effort control problem with a multiple criterion cost function. We can move the constraint on the control signal that is set by β to the objective function and write the following optimization problem:

$$\begin{aligned}
 & \underset{u_3(t)}{\text{minimize}} && \| \ddot{x}(t) \|_2 + \gamma \| u_3(t) \|_\infty \\
 & \text{subject to} && \dot{\mathbf{x}}_3(t) = A\mathbf{x}_3(t) + Bu(t) \\
 & && \mathbf{x}_3(0) = \mathbf{x}_i \\
 & && \mathbf{x}_3(T) = \mathbf{x}_f
 \end{aligned}$$

Clearly, this is a minimum effort control problem with an additional cost associated with the acceleration of the end-effector. While the γ parameter does not have a one-to-one relationship with the β parameter of the previous optimization problem, they both represent the trade-off with hard-limiting the control signal (that is, keeping the cost of the effort small) and the “size” of the acceleration. This optimization can also be solved in discrete time as a Quadratic Program. First, we write the MACC in discrete time as follows:

$$\begin{aligned}
 & \underset{\mathbf{x}}{\text{minimize}} && \| D_2\mathbf{x} \|_2 + \gamma \| D_3\mathbf{x} \|_\infty \\
 & \text{subject to} && A_{eqn}\mathbf{x} = b_{eqn}
 \end{aligned} \tag{5.18}$$

where again γ is the trade-off parameter. In a manner similar to equation 5.15, the optimization problem of equation 5.18 can be posed as a Quadratic Program and solved with a QP solver.

5.4 Combination of Multiple Control Signals

We propose a control paradigm where multiple control signals (that correspond to different dynamic variables) are involved in controlling a biomechanical system. This paradigm is inspired by evidence from the central nervous system encoding different dynamic variables (such as velocity and acceleration) as control signals. The multi-control policy model hypothesizes that segregated populations of neurons throughout the nervous system (such as the brain stem or the motor cortex) represent control policies that reflect the various trade-offs that are defined by a cost function. For example, one neural population may represent a control signal that creates “peaked” velocity profiles, while another population may represent a control signal that creates a velocity trajectory with a more “broad” profile. While the full biological details of such a hypothesis are unknown, we propose a simplified mathematical model to illustrate our proposed hypothesis. Here we propose that a population of neurons represent a control policy with a trade-off in one extreme (“peakiness” in the velocity profile for example), while another population represents a control policy operating in an opposite extreme (“broadness” in the velocity profile). In addition, there may be several neural populations where each population represents different control policies that reflect different costs. These neural populations are then combined to produce the rich set of movements we observe.

Our goal here is to illustrate the feasibility of combining control signals to achieve a rich set of movements, and we do not focus on the specific control policies that neural populations encode. The specific control policy that neurons represent is unknown, but since sparse control policies are biologically realistic, we propose that each population represents a sparse control policy for a control signal that represents a high-order derivative of the end-effector dynamics (such as “jerk” corresponding to the third derivative of position). The model that we propose here is for a one-dimensional end-effector point mass (that we treat as the hand of the subject) to simplify the math-

ematics and the model. Extensions to higher dimensions can be analogously made and more sophisticated and accurate models of the musculo-skeletal system can be included, but we defer this work for future studies since the simplifications we have made here are sufficient for illustrating the concept and architecture of combining multiple control policies.

While such evidence makes it certain that the central nervous system uses multiple (and different) control signals, how such control signals are combined is unknown. We propose that the control policy the CNS follows is a simple linear combination of various control signals (each optimal with respect to a particular cost function: the l_∞ norm in our case) that result in a rich set of movements. It is possible to consider other optimal control signals (such as ones proposed by [69, 62, 70]) and to combine these signals in other ways that are perhaps even non-linear, but these possibilities are not explored here and are deferred for future work. We propose two models: the Convex Combination Trade-off Control model and the Optimal Confluence model. The Convex Combination Trade-off Control model considers finding a control signal that is a convex combination of two different control signals with two different sets of constraints. While the Convex Combination Trade-off Control model can be extended to any control system, tuning the trade-off parameters that are introduced makes its application difficult in practice. The Optimal Confluence model, on the hand, is tailored towards the specific case where the combined control signal should match a desired peak velocity. Clearly, this model is specialized for a specific criterion, but for our application of fast human movements, it is sufficient.

5.4.1 Convex Combination Trade-off Control Model

We define the Convex Combination Trade-off Control model as an optimization problem where we seek a control signal that is a convex combination of two different control signals and also a “convex” combination of the two different constraints. That is, when we design a particular control signal according to equation 5.3, we have a particular set of constraints associated with the particular optimization problem. Solving a different control signal optimization problem according to equation 5.3 involves a different set of constraints. Therefore, if we wish to design a convex combination of

two different control signals, then we also need to consider a “convex combination” of the two different constraints. More concretely consider the following two optimization problems:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \| D_{n-1} \mathbf{x} \|_p \\ & \text{subject to} && A_{eq_{n-1}} \mathbf{x} = b_{eq_{n-1}} \end{aligned} \quad (5.19)$$

and

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \| D_n \mathbf{x} \|_p \\ & \text{subject to} && A_{eq_n} \mathbf{x} = b_{eq_n} \end{aligned} \quad (5.20)$$

where $A_{eq_{n-1}}$, $b_{eq_{n-1}}$, A_{eq_n} , and b_{eq_n} are defined as before. We wish to now design an optimization problem where the cost function reflects the convex combination of the objective functions of equations 5.19 and 5.20 and also reflects a convex combination of their respect constraints $A_{eq_{n-1}} \mathbf{x} = b_{eq_{n-1}}$ and $A_{eq_n} \mathbf{x} = b_{eq_n}$. Note that we can write A_{eq_n} and b_{eq_n} recursively as:

$$A_{eq_n} = \begin{bmatrix} A_{eq_{n-1}} \\ \mathbf{a}_n \end{bmatrix} \text{ and } b_{eq_n} = \begin{bmatrix} b_{eq_{n-1}} \\ \mathbf{b}_{a_n} \end{bmatrix} \quad (5.21)$$

$$\text{where } \mathbf{a}_n = \begin{bmatrix} (\Delta T)^{-(n-1)} p_{n-1} & \mathbf{0}_{1 \times (N-(n-1))} \\ \mathbf{0}_{1 \times (N-(n-1))} & (\Delta T)^{-(n-1)} p_{n-1} \end{bmatrix} \text{ and } \mathbf{b}_{a_n} = \begin{bmatrix} x_i^{(n-1)} \\ x_f^{(n-1)} \end{bmatrix}.$$

Having introduced this notation, we can define the Convex Combination Trade-off Control optimization problem as follows:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \theta \| D_{n-1} \mathbf{x} \|_p + (1 - \theta) \| D_n \mathbf{x} \|_p + (1 - \theta) \gamma \| \mathbf{a}_n \mathbf{x} - \mathbf{b}_{a_n} \|_2 \\ & \text{subject to} && A_{eq_{n-1}} \mathbf{x} = b_{eq_{n-1}} \end{aligned} \quad (5.22)$$

where θ is the parameter that control the convex combination of the two different control signals, and γ determines the sensitivity of the augmented constraints. When θ is close to one, the optimization problem in equation 5.22 approaches the solution of equation 5.19, while when θ approaches zero, the optimization problem in equation 5.22 approaches the solution of equation 5.20. The γ parameter must be selected beforehand for the convex combination of the two different control signals to work. We now give a more concrete example where a jerk control signal and a snap control signal are combined

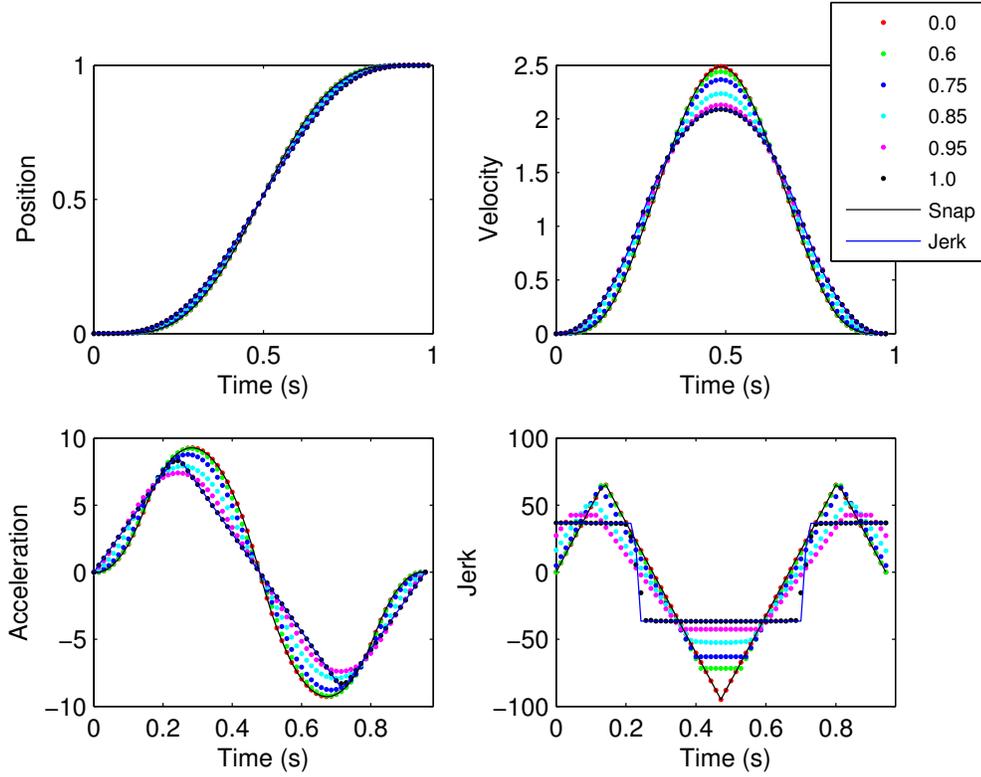


Figure 5.3: The control policies resulting from varying θ from zero to one for the optimization problem in equation 5.23.

with the model of 5.22. This problem can be written as:

$$\begin{aligned}
 & \underset{\mathbf{x}}{\text{minimize}} && \theta \| D_3 \mathbf{x} \|_{\infty} + (1 - \theta) \| D_4 \mathbf{x} \|_{\infty} + (1 - \theta) \gamma \| \mathbf{a}_4 \mathbf{x} - \mathbf{b}_{a_4} \|_2 \\
 & \text{subject to} && A_{eq_3} \mathbf{x} = b_{eq_3}
 \end{aligned} \tag{5.23}$$

where we have considered the case $p = \infty$ case for measuring the effort of the jerk ($D_3 \mathbf{x}$) and snap ($D_4 \mathbf{x}$) control signals. Figure 5.3 shows that as θ is varied from zero to one, the control signal (along with the constraints) are varied from a minimum snap control model to a minimum jerk. Note that the γ parameter makes varying θ be very sensitive to switching from a minimum snap to a minimum jerk control signal.

5.4.2 Optimal Confluence Model

Since studies (as indicated in Chapter 2) have shown that the peak velocity of a fast movement is an important parameter in determining the characteristics of the dynamics, we propose that in addition to using the boundary conditions of the initial and final positions to also use the desired peak velocity as an additional parameter in determining the characteristic of the movement. We define the Optimal Confluence model as an optimization problem where we seek a control signal that is a convex combination of two different optimal control signals, where the combined control signal has a peak velocity as close as possible to a desired peak velocity. More concretely, consider the following two optimization problems:

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \| D_m \mathbf{x} \|_p \\ & \text{subject to} && A_{eq_m} \mathbf{x} = b_{eq_m} \end{aligned} \quad (5.24)$$

and

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \| D_n \mathbf{x} \|_p \\ & \text{subject to} && A_{eq_n} \mathbf{x} = b_{eq_n} \end{aligned} \quad (5.25)$$

where $A_{eq_{n-1}}$, $b_{eq_{n-1}}$, A_{eq_n} , and b_{eq_n} are defined as before, and $m \neq n$. Let \mathbf{x}_m^* and \mathbf{x}_n^* be the optimal solutions to equations 5.24 and 5.25 respectively. Define $v_m^{Max} \equiv \| D_2 \mathbf{x}_m^* \|_\infty$ and $v_n^{Max} \equiv \| D_2 \mathbf{x}_n^* \|_\infty$ (that is, v_m^{Max} and v_n^{Max} are the peak velocities for \mathbf{x}_m^* and \mathbf{x}_n^* respectively). The Optimal Confluence model then finds the optimal convex combination between the control policies of equations 5.24 and 5.25 so that the combination is as close as possible to a desired peak velocity. This is done by solving the following optimization problem:

$$\begin{aligned} & \underset{w_1, w_2}{\text{minimize}} && (w_1 v_m^{Max} + w_2 v_n^{Max} - v^{peak})^2 \\ & \text{subject to} && w_1 \geq 0 \\ & && w_2 \geq 0 \\ & && w_1 + w_2 = 1 \end{aligned} \quad (5.26)$$

where v^{peak} denotes the desired peak velocity for the combined control signal, and w_1 and w_2 are the convex combination weights.

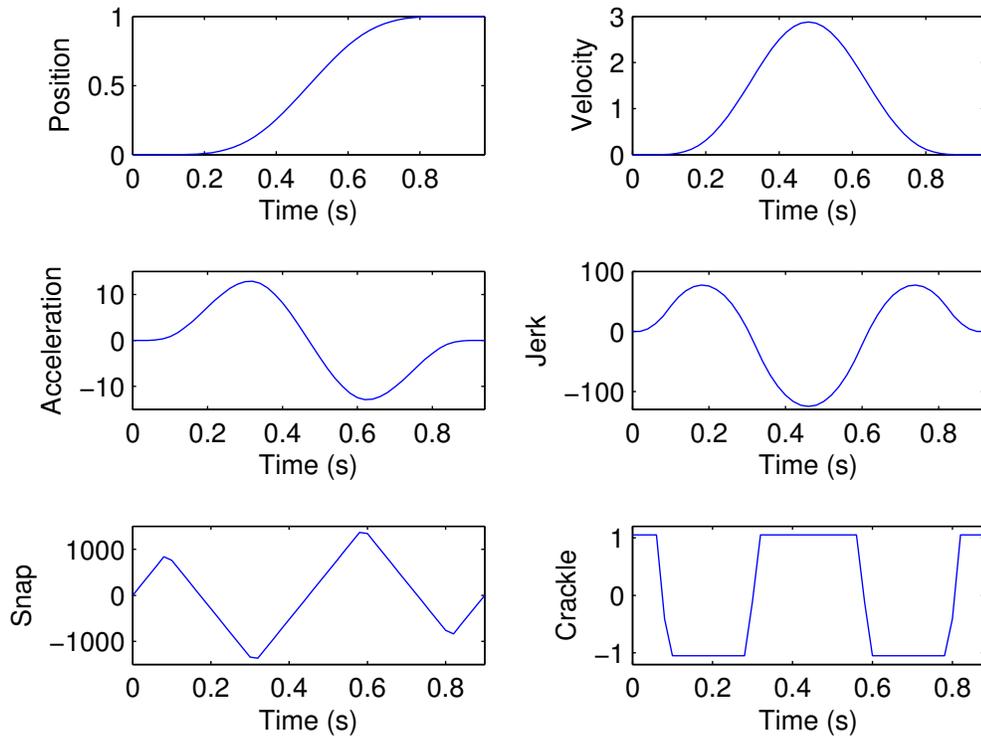


Figure 5.4: The solution to equation 5.3 with $n = 5$ and $p = \infty$.

To illustrate the Optimal Confluence model, consider the example where we wish to combine two control signals that correspond to jerk and crackle (the 5th derivative of position). We consider an initial position of 0 and a final position of 1 units, and that the movement starts at rests, and ends at rest. The duration of the movement is 1 second and the number of samples is 50 (thus, $i = 0, 1, \dots, 49$). The optimal jerk control policy for this example was solved before and is shown as the blue curve in Figure 5.1. The optimal crackle control policy for this example has the solution shown in Figure 5.4 (note that A_{eq_5} and b_{eq_5} equality matrices can be similarly defined as equation 5.14). The Optimal Confluence model finds the optimal combination of the jerk control signal shown in blue in Figure 5.1 and the crackle control signal shown in Figure 5.4 to a desired peak velocity. Figure 5.5 shows the jerk control signal in blue, the crackle control signal in green, and the desired trial in red. The Optimal Confluence model (shown in cyan) is the then the optimal convex combination of the jerk and crackle

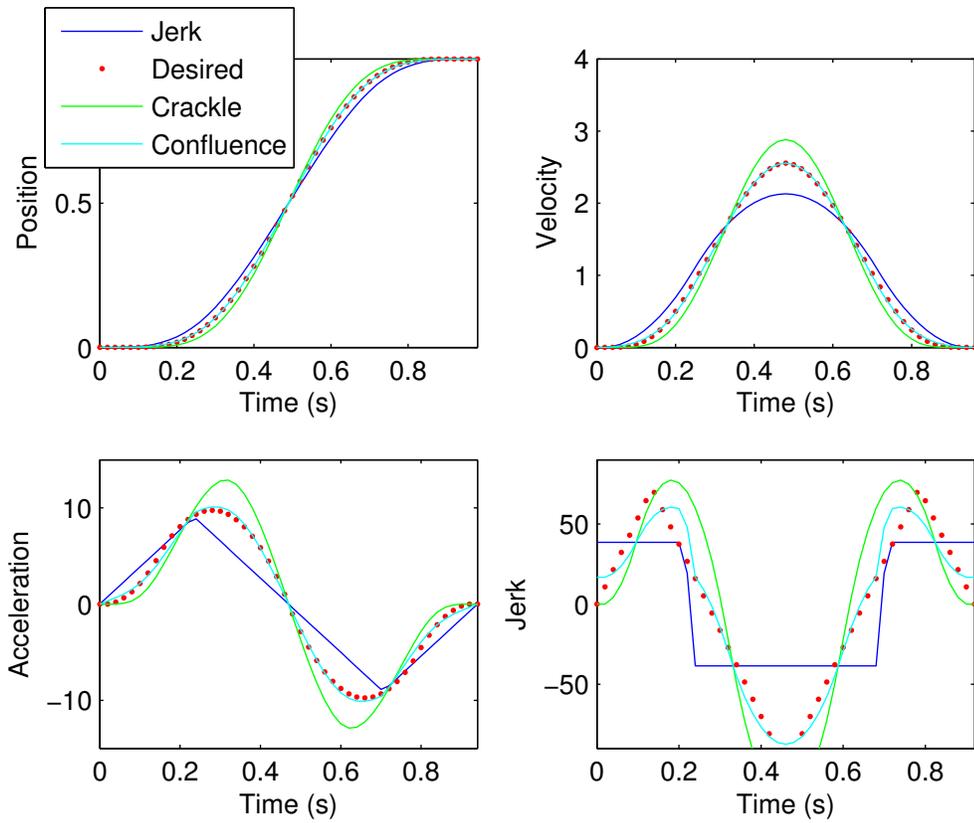


Figure 5.5: The optimal convex combination (shown in cyan) of the jerk (blue) and crackle (green) control signals to be as close as possible to the desired peak velocity of a given movement dynamics (shown in red).

control signals to be as close as possible to the desired velocity profile.

5.5 Notes

Chapters 5, in part, have been submitted for publication of the material. The dissertation author is the primary investigator and author of this material.

Chapter 6

Confluence

We propose an optimal control model of human movement that combines multiple control policies and has direct ties to neuroscience. We model optimal control signals generated by disparate populations of neurons that together drive the CNS to generate movement. These signals originating from different parts of the nervous system are combined in varying ratios to achieve varied movements. Each control signal is optimized with respect to a cost function. Combinations of control signals based upon different cost functions result in different control strategies. The confluence of these control signals allows for greater variability in the combined signal and leads to the variability observed in human movement. We cite examples of how elements of the vertebrate motor system map directly to our mathematical model. By relating the model to anatomy and physiology, we narrow the gap that typically exists between optimal control theory and the functioning of the motor system.

6.1 Introduction

How humans achieve graceful and effective movement is largely unknown. The neuroscientific literature suggests that the control signals emanating from disparate parts of the nervous system are combined in varying ratios to achieve a vast array of movements. We propose that different weightings of these signals lead to different movement characteristics. Each control signal emanating from a given population of neurons has been optimized via evolutionary processes with respect to a cost or reward function or

functions. Our new optimal control paradigm utilizes multiple cost functions (that are not necessarily composite) that reflect the diversity of tasks and environments biological systems encounter over the course of their evolution. When applying optimization theory to biological motor control, two questions arise:

1. What control policy (or policies) does the motor system use to achieve movements?
2. How are control policies implemented by the nervous system?

6.1.1 What Policy or Policies?

An optimal control policy is a strategy that is optimal in the control of a system (such as a biomechanical system, a robotic arm, etc.) with respect to minimizing (or maximizing) a specific cost (or reward) function. Techniques from optimal control theory that utilize a single cost function to arrive at a control policy include: [12, 69, 71, 62, 13, 42]. These policies are only optimal with respect to a the single cost function each respective author has proposed.¹

Here.

Optimal control cost functions that are typically considered involve either a single criterion or trade-offs between several criteria (e.g. find a policy to get to a destination as fast as possible while consuming the least amount of fuel). One of the earliest applications of optimal control to human movement is the work of Flash and Hogan [12] where the cost function has the effect of minimizing “effort.”² The control signal which Flash and Hogan proposed is a signal that corresponds to the third derivative of position of the hand (also known as jerk) during a point-to-point arm reaching task. Other control signals and minimum effort control problems have been considered. For example, [69] considered torque change as a control signal while Ben-Itzhak and Karniel [42] considered acceleration. While the mathematical formulations of more recent optimal control

¹In finding an optimal control policy, in addition to the cost function, the set of constraints in the optimization problem also plays a major role in determining the characteristics of an optimal control policy. Details of the constraint set used in this work are discussed in the methods section.

²A “minimum effort” control problem is an optimization problem that seeks a control policy that meets the desired set of constraints of the problem and uses the smallest “size” or “effort” of the control signal [43]. The mathematical definition of measuring the size of a control signal can lead to insights in biological control as proposed in Chapter 4.

models (such as [62, 13, 72]) may appear fundamentally different from the approach of Flash and Hogan, they are in fact extensions of the minimum effort type control problem that Flash and Hogan proposed. These works employ a single cost function that trades off between minimizing effort (the control signal) and other criteria such as end-point accuracy or stability.

To the best of our knowledge, all applications of optimal control to human movements have consisted of a single proposed cost function³ producing a single control signal. In contrast, biological systems have many interacting subsystems composed of disparate structures that coevolved together, each system forged by a different cost or reward function. Likewise, it is reasonable to assume that most elements of the motor system (e.g. spinal cord, cerebellum, brain stem motor nuclei, motor cortex) did not evolve based on a single cost function, but instead each coevolved utilizing different cost functions.

6.1.2 An Approach to Modeling the Motor System Using Optimal Control: The Optimal Confluence Control Framework

Despite optimal control theory's great promise to shed light on biological systems that have been forged by evolutionary processes, researchers that utilize this technique typically do not attempt to connect the mathematics of their models with neuronal mechanisms [60]. The main contribution of this work is the explanation we offer regarding the highly plausible implementation of optimal control in the vertebrate motor system. In doing so, we address the second question regarding applying optimization theory to biological motor control, "How are control policies implemented in the nervous system?"

There is abundant evidence that heterogeneous neural control signals originating in disparate parts of the motor system are combined (see Discussion section for further details and see Figure 6.1). This combination implies that the animal motor system can trade-off between various control policies. We propose a novel framework that consists

³Typically, a cost function that trades off between various criterion is referred to as a composite cost function. While a composite cost function contains several criteria, they are still encapsulated by a single function.

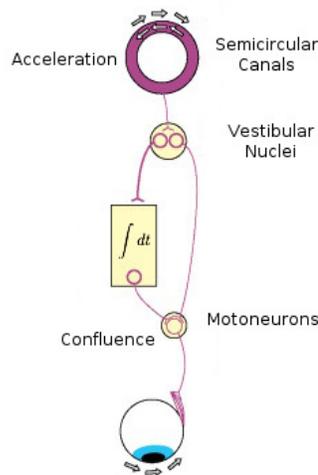


Figure 6.1: Figure adapted from [73].

of multiple control policies that are each optimal with respect to different cost functions. We call this formulation an “Optimal Confluence” control framework and assert that it is a step towards applying optimal control theory to understanding biological systems.

Although this formulation is mathematically and fundamentally different from previous optimal control models, any set of optimal control models can be adapted to be implemented within the Optimal Confluence framework. The key difference between our framework and traditional optimal control policies is the way in which trade-offs between various cost/reward goals are achieved (e.g. [70] trades-off between end-point position accuracy, end-point stability, and effort). Any set of cost functions can be chosen, and a separate control signal/policy is pre-optimized for each. Trade-offs are achieved by combining the resulting optimal signals in different ratios to produce different types of movements. The Optimal Confluence control problem is to find the optimal combination of these different control signals for the particular movement considered.

Our model draws its inspiration from the vertebrate nervous system and the evolutionary forces that presumably shaped its sub-components and neural control signals. Real-time optimization performed by the nervous system during movements seems computationally implausible. With this in mind, our approach posits that the optimization of neural control signals was accomplished over millennia via evolutionary processes, and that the only real-time neural computation performed during a movement resembles a

linear combination of neural signals.

6.1.3 An Application of Optimal Confluence Control for Minimum Effort Policies: JCOC

The Optimal Confluence paradigm provides a universal framework for exploring the control of movements by the central nervous system by allowing for multiple control policies of any type to be combined. The remainder of this paper will discuss the particular control problem of making fast point-to-point reaching movements. In particular, we utilize the Optimal Confluence framework to implement a combination of minimum effort control policies that minimize jerk and crackle (5th derivative of position) and combine the resultant signals to drive a system. The Jerk-Crackle Optimal Confluence model (JCOC) demonstrates the ability of the Optimal Confluence paradigm to model a basic human movement.

One of the characteristics of fast human arm movements that Flash and Hogan and others have defined is the ratio of peak velocity to average velocity of the hands movement. We use r to refer to this value:

$$r = \frac{\text{peak velocity of movement}}{\text{average velocity of movement}}$$

Minimum effort models typically have a rigid value for r that does not change. An inspection of our human fast hand movement data reveals that r varies from trial to trial in ways that are not captured by minimum effort control policies. Flash and Hogan's data was based on one subject's movements over thirty trials and yielded an r value of 1.805 with a standard deviation of 0.153. They chose jerk as the derivative of position to minimize in part because their minimum jerk model predicts an r value of 1.875, very close to their experimental data [12]. However, studies with more extensive data show that models based on minimum snap fit the data better than those based on minimum jerk [74, 75, 76]. Yet, other studies continued to find that minimum jerk models were a better fit to the data [77]. Other studies still use jerk minimization but admit that snap minimization is a better fit in many cases [78]. A comparison of a minimum snap model vs a minimum jerk model against our data supports the notion that minimum snap is a better fit (see Figure 6.4). Scientific studies that are tangentially related to models of

human movement accept the Flash and Hogan minimum jerk model as baseline truth, and use this to bolster the arguments they are making [79, 80].

There is a reason for this confusion. All of these models fit the data only on average, and do not account for the variability of r observed across trials. Often, no model in the minimum effort family provides a good fit because a trial's r value falls in between that generated by two of the minimum effort models (see Figure 6.2). The application of our framework to minimum effort control problems, JCOC, does not suffer from these deficiencies. The confluence of pre-optimized control signals based upon the minimization of multiple derivatives of position allows for a flexible r value that fits data better not just on average but on a trial to trial, subject to subject basis (see Figure 6.3). We propose that in this way, the nervous system combines many policies to achieve the variability observed in human movement.

We use sparse control policies⁴ in this work as an example to showcase the application of the confluence of multiple optimal control signals. We consider two control signals that are optimal with respect to a cost function based on minimizing two high-order derivatives of the position of the end-effector (hand). One control signal corresponds to jerk (this is the same control signal that Flash and Hogan used) and another corresponds to the 5th derivative of position, also known as “crackle.” Each of these control signals is minimized, as measured by the infinity-norm, to result in a sparse control signal for an example of minimizing a control signal with the infinity-norm). The minimum jerk control signal has an r of 1.81 while the minimum crackle control signal has an r of 2.09. The combination of these control signals yields an r between these two extremes, a range in which most fast reaching movements lie. Our Optimal Confluence framework combines the control signals in such a way that the combined control signal has an r as close as possible to the trial's r . See Chapter 5. We remind the reader that any set of control policies can be plugged into our proposed framework, such as those proposed in [62, 13]. We choose cost functions related to the derivatives of position because there is neuroscientific evidence that such signals exist in many parts of the motor system.

⁴A sparse control policy is a control policy that uses a sparse signal (that is, the signal is characterized by a small number of spikes). See Chapter 4 for more details.

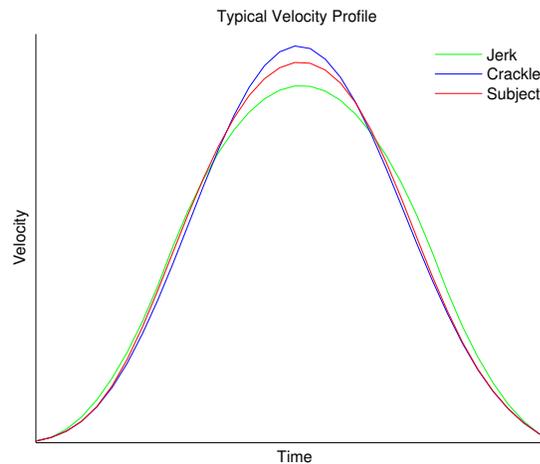


Figure 6.2: The green and blue curves correspond to the minimum jerk and minimum crackle velocity profiles respectively. The red curve is the velocity profile of a give trial.

6.2 Results

The data set used in this study was acquired from Amir Karniel at the Ben-Gurion University of the Negev. The recordings include velocity and position profiles of short point-to-point movements made by five individuals and recorded using a manipulandum. Each subject was instructed to make a fast movement between two of three targets arranged in an equilateral triangle, 10 cm apart. The movements lasted a third of a second ± 50 ms. A more formal description of the data collection methods can be found in Mussa-Ivaldi and Karniel's work [40] or see Chapter 3.

The green and blue curves in Figure 6.2 result from minimization of jerk and crackle (respectively) with the same and norm. As can be seen in the figure the maximum velocity achieved by the crackle model is higher than that of the jerk model. The red line above is a typical profile from the data set of actual human movements and lies in between the two models. This typical case demonstrates the need for a more advanced model since neither jerk nor crackle minimization (nor any of the other derivatives) can accurately describe the human movement data. Jerk and crackle minimization were chosen for optimal combination because their velocity profiles lie near the lower and upper bounds of peak velocity of the data set.

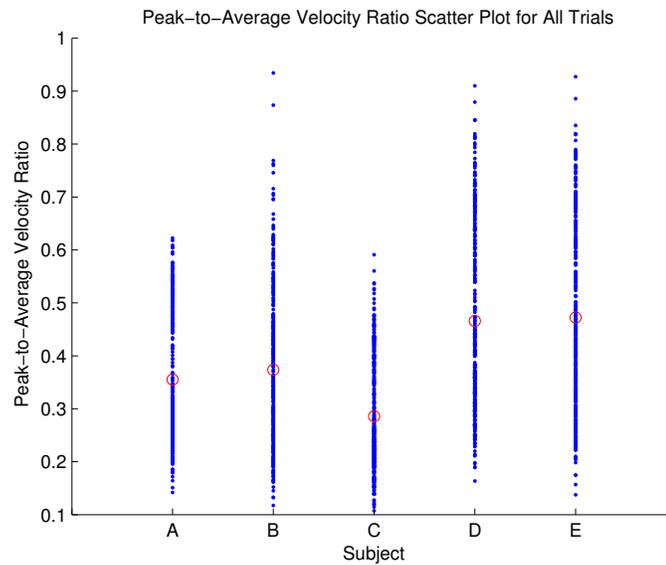


Figure 6.3: Plotted in blue is the peak velocity of each trial of subjects A-E. Plotted in red is the average peak velocity of each subject. Solid green and orange lines are the peak velocities corresponding to jerk and crackle minimization respectively.

The ratio r of a given model is sensitive to several factors. Both the minimized derivative and norm of minimization can influence r see table of r vs derivative of min. A period of time lies before and after each trial movement that are unimportant for this investigation. An automated process of selecting the ballistic portion of the movement was developed (see methods). This process can introduce error into the calculation of trial r values by way of altering the average velocity of a trial. To avoid this error, and since the distance and duration of the movements in this data set are fairly consistent, we use peak velocity for the the optimization procedure in this study (see methods section). Figure 6.3 shows that the mean peak velocity for each subject is different and does not necessarily correspond to an r value of 1.805 as reported by Flash and Hogan. A one-way ANOVA test shows that the mean peak velocity for each subject is indeed different (with $p < 0.001$), hence we can conclude that r is not an invariant characteristic of fast human arm movements. Furthermore, a principle component analysis of velocity profiles for single subjects yields two clusters. One cluster corresponds to diagonal movements while the other corresponds to horizontal movements. This further supports the idea that different control signals are needed to control movements with different

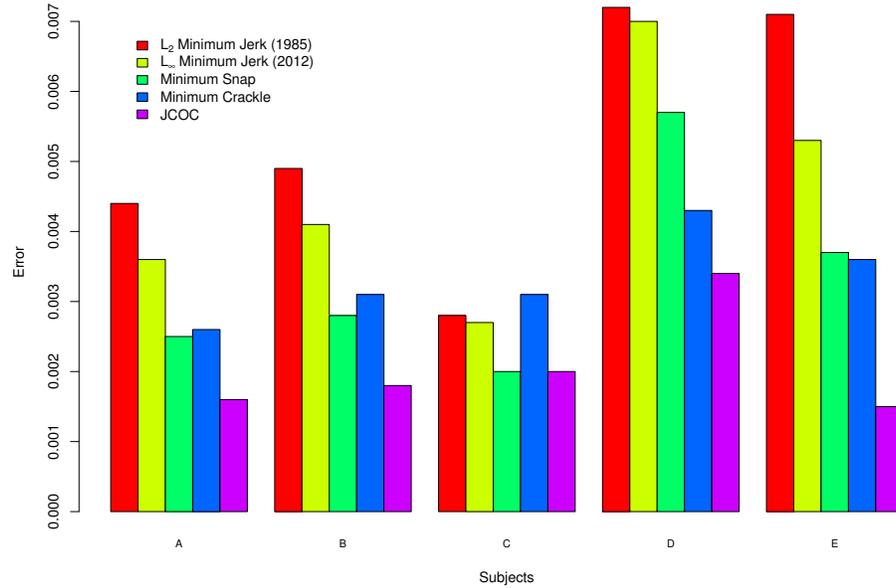


Figure 6.4: Average mean squared error of jerk 2-norm, jerk inf-norm, snap, crackle and JCOC for each subject and for all subjects, all trials. JCOC on average has a smaller error than all models shown with $p < 0.001$.

dynamics.

6.2.1 Evaluation of Model

Mean squared error (MSE) of the velocity profile of the model versus the trial's was used to compare the accuracy of the models discussed here. In Chapter 3 we demonstrated that using the infinity-norm instead of the 2-norm improves accuracy of the model, as can be seen in Figure 6.4, and for this reason we chose to continue to use the infinity-norm (see also methods). While the accuracy improvements between Flash and Hogan's model and our previous work are significant, the JCOC model based on jerk and snap significantly outperforms both models. As seen in Figure 6.5 the JCOC model has a significantly lower MSE than one or both of its constituent models in almost all cases.

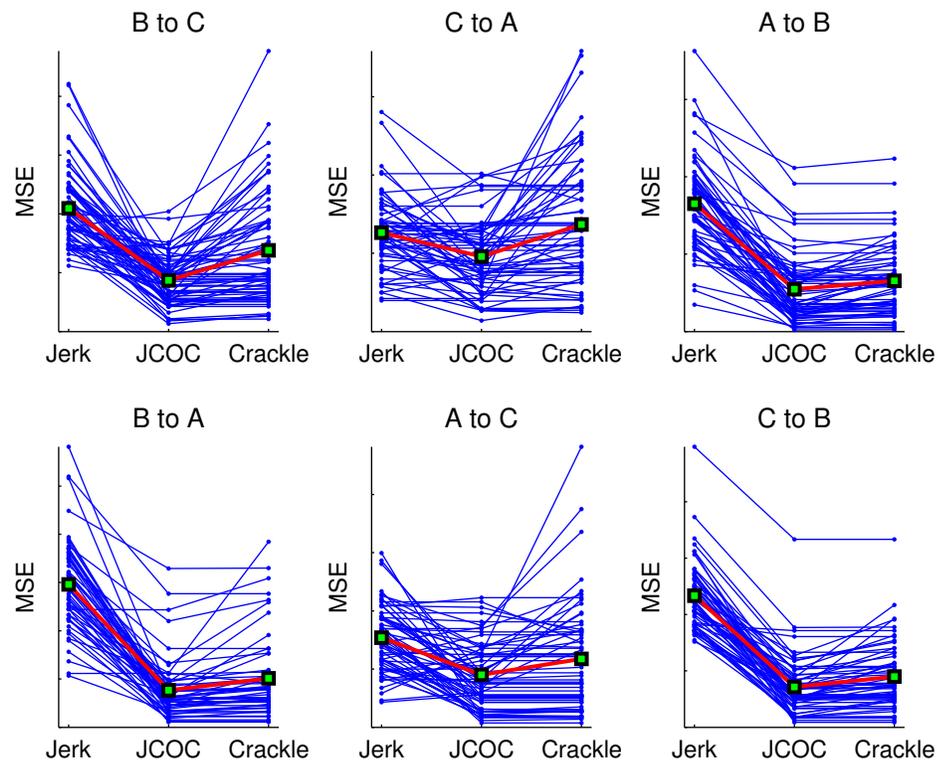


Figure 6.5: Plots of mean squared error (MSE) of three models compared with subject E's horizontal left-to-right movements. Points plotted from the same trial are connected with straight lines.

6.3 Discussion

The Optimal Confluence paradigm outlined in this work creates a bridge between mathematical models of movement and known biological systems. It emphasizes the role of position and its derivatives as neural motor control signals. Neuroscientific examples of such signals driving motor control are plentiful. This new paradigm also allows for the combining or confluence of such neural signals, as we will show, a phenomenon witnessed throughout the nervous system.

6.3.1 The Role of Neural Integration and the Derivatives of Position in Motor Control

We posit that the majority of animal species that have a sense of position, velocity, acceleration, etc., have an evolutionary advantage. A drosophila can sense the direction from which a threat is approaching (change in position). The positional information directly influences the fly's motor system as it positions its body optimally for a quick escape in the opposite direction [81]. A cheetah would not be able to change its velocity in a manner needed to catch its prey if it had no notion of acceleration. The cheetahs intended magnitude of velocity change must be encoded in some way within the its nervous system.

More generally, if an animal species is to avoid extinction, possessing some neural representations of basic Newtonian variables is evidently advantageous and necessary in order to navigate effectively. Thus, these neural representations must influence the motor systems of these animals. Animals without these abilities have more difficulty navigating their environments in order to find food, escape predators, etc. In contrast, the majority of surviving species have these capabilities.

We have proposed that neural signals representing higher derivatives of position are central to vertebrate motor control. These signals appear in numerous contexts. For example, in the vestibular system, the semicircular canals and the otolith organs both produce signals coding for derivatives of the position of the head [73, 82, 83, 84, 85, 86, 87] which initiate changes in eye position and the muscles that control body stability. Other examples of position and its derivatives contributing to the control of

motor function include signals from muscle spindles which encode relative position and velocity of the limbs and body, or signals from Golgi tendon organs which encode force exerted by muscles (with an obvious relationship to acceleration) [88, 89]. All of these signals converge upon spinal circuits and influence motor control.

The implementation of the Optimal Confluence paradigm demonstrated in this work treats higher order derivatives of position as control signals.⁵ As outlined above, it is advantageous for animal nervous systems to have access to information representing position and its derivatives. In order to convert a control signal representing a higher order derivative to a signal representing a lower derivative, all that is necessary is the operation of integration (see Figure 6.1). There is evidence for the existence of multiple, disparate classes of neural mechanisms capable of approximating mathematical integration in humans and other animals. Examples include integrators in entorhinal cortex [90], hypothesized neural integration in head direction cells [91, 92], and integration in the oculomotor system [73].

For example, the oculomotor system⁶ consists of several subsystems, many of which use neural integration mechanisms [94, 95, 96, 97]. Here we choose to examine the vestibulo-ocular reflex. The semicircular canals of the inner ear detect angular acceleration of the head [87]. Yet, the signals measured in the nerve emanating from these organs code for velocity, not acceleration. Further downstream in the system, neural signals coding for eye position emerge. These signals drive the alpha motor neurons controlling the eye muscles to move in the opposite direction of head movement in proportion to the velocity of the head movement, and then hold the new eye position, keeping the eyes on target [73]. Robinson was one of the first to realize that integration (in the mathematical sense) must be occurring [57]. The existence of neural integration mechanisms in the oculomotor system is now well accepted [73, 98, 99]. Two brain stem nuclei, the interstitial nucleus of Cajal (INC) and the nucleus prepositus hypoglossi (NPH), are key in the vertical and horizontal (respectively) oculomotor neural integration systems [94, 95, 97, 96]. This system uses signals representing position

⁵Any type of control signal can be used within the framework, this is simply one instance of the framework.

⁶The oculomotor system is a system widely studied in order to gain insight into the functioning of the motor system in general [93].

and its derivatives to drive oculomotor function. All the systems mentioned earlier provide a biological substrate for the computation of the dynamic variables found in many mathematical models of animal movement [12, 62, 42, 69, 70].

6.3.2 Confluence of Neural Signals

Aside from neural integration, the oculomotor system exhibits a confluence of neural signals working together to achieve precision motor control. For example, the otolith organs produce signals representing linear acceleration. Similar to the role of the semicircular canals in the VOR, the otolith organs also play a role in eye control [73, 100]. The VOR and the otolith organs, along with other oculomotor subsystems all act in concert in order to control ocular motor neurons. The same neural integration nuclei involved in the VOR system (INC and NPH) are also influenced by signals from the otolith organs, as well as signals representing volitional control of the eyes, and by signals that are step-like which originate from premotor neurons in the rostral midbrain to produce saccades [93]. When taken together, we see a confluence of neural signals from different areas of the brain converging, combining and cooperating in order to produce numerous types of eye movements such as smooth tracking of moving objects, saccades of varying magnitude, vergence, target fixation during head movement, etc. The oculomotor system is not controlled by a single control signal, but rather combines different signals in varying degrees depending on the desired behavior [93].

Limb movement is also achieved via a confluence of neural signals. Betz cells originating in motor cortex converge directly on alpha motor neurons in the spinal cord while other projections from motor cortex and brainstem influence spinal alpha and gamma motor neurons only indirectly, e.g. through various layers of the Rexed lamina, before the final signals reaches alpha and gamma motor neurons in the spinal cord [101, 102]. Additionally, spinal circuitry and alpha motor neurons receive signals from muscle spindles and Golgi tendon organs which directly and indirectly influences motor control [89]. All of this evidence for the combination of multiple cooperating signals supports the need for a new optimal control paradigm that more directly reflects these observations of the vertebrate motor system.

6.3.3 Motor Primitives and the Optimal Confluence Paradigm

Motor primitives are loosely defined as a group of simple movements that can be combined to create more complex movements. Evidence points to the existence of “motor primitives” located in the spinal cord [7, 103, 39]. The most well known are so called pattern generators that produce oscillatory patterns in animals which are central to locomotion. The neural machinery that accomplishes this appears to be located in the spinal cord as it does not require supraspinal input in order to produce locomotion [8].

An experiment by Giszter demonstrated the existence of limb position motor primitives in the spinal cord of frogs [7]. A frog’s limb was moved to various starting positions. The frog’s spinal cord is then stimulated. Stimulation of the same location in the spinal cord caused the frogs limb to produce force toward the same point in space, regardless of the limb’s starting position. To accomplish this, the same spinal stimulation necessarily activated muscles in different orders and in varying degrees depending on the initial position of the limb. This result indicates that locations in the spinal cord are “tuned” for locations in space regardless of the organisms initial limb position. A variation on the experiment showed that simultaneous excitation of multiple areas of the spinal cord (multiple motor primitives) produce some combination of effects not witnessed when either one area or the other was stimulated in isolation. The force (represented by vectors) created by the combination of the motor primitives was the vector sum of each of the forces generated by each primitive on its own. The two primitives were combined to create a force in a new direction with a new magnitude. This concept of combining “motor primitives” (also called muscle synergies) has been discussed frequently in the literature [104, 105, 106, 10].

6.4 Notes

Chapters 6, in part, have been submitted for publication of the material. The dissertation author is the primary investigator and author of this material.

Appendix A

Sparse Optimal Filters

A.1 Overview

Finite Impulse Response (FIR) digital filters are often used in digital signal processing due to their guarantee of stability, the ease of ensuring linear phase delays, and simplicity in implementation[107]. Implementing FIR filters involves a series of delays, multiplications, and additions of the input signal. The order of the filter is determined by the highest delayed element that has a non-zero coefficient. For example, we refer to the following filter

$$H(z) = \sum_{k=0}^{k=N} a_k z^{-k}$$

with $a_N \neq 0$ as an N -th order FIR filter. Since a major set of computations involved in implementing FIR filters are additions and multiplications, here we define the FIR filter complexity as the number of non-zero coefficients in the filter (the coefficients in an FIR filter represent a weight assigned to a delay element of the input). Clearly, for the N -th order FIR filter above, if all filter coefficients are not zero, then the complexity would be $N + 1$.

In practice, the coefficients of the FIR filter must be chosen appropriately to achieve a set of desired goals that depend on the application. The selection of these coefficients is often done algorithmically. For example, the Parks-McClellan algorithm is popularly used to select FIR filter coefficients that are close to a desired frequency response (such as lowpass, bandpass, etc.)[107]. As another example, the FIR filter may

be designed to maximize the signal-to-noise ratio (SNR), that is we wish to improve the visibility of a possibly random signal corrupted with noise[108, 109, 110]. In the context of detection theory such filters are often called “Matched Filters.”

Here we are interested in designing filters that not only meet a certain set of specifications (desired frequency response or maximizing SNR), but are also sparse. That is, we are interested in designing filters that are of high order N but exhibit low complexity M where $M \ll N$. While in general these sparse filters may be sub-optimal, their significant reduction in complexity offer an attractive advantage in ease of implementation. In our discussion we focus on optimal filters that maximize SNR. It will be shown that the solution to maximizing the SNR when the desired signal is corrupted with white Gaussian noise is to find the eigenvector of the signal autocorrelation matrix with the largest eigenvalue (the classic “Matched Filter” solution). We thus seek to find an eigenvector of the autocorrelation matrix that is also sparse.

Recently, Oppenheim and co-authors proposed linear programming algorithms for the design of sparse filters with a desired filter specification[111]. Thus their algorithm is limited to filters whose desired filter specifications is known a priori. While the algorithm we propose is a quadratic program, our algorithm is general enough to “sparsify” any set of FIR filter coefficients. This is first done by computing the filter coefficients using the appropriate algorithm (the eigenvector corresponding to the largest eigenvalue of the autocorrelation matrix in the case of maximizing SNR), followed by finding a sparse approximation to the designed filter. We will describe this method in detail in section 2 and show simulation results in section 3. We show future directions in the conclusion section.

A.2 Problem definition and solution

We wish to design a sparse FIR filter that improves the visibility (that is, the SNR) of a signal $s[n]$ corrupted with white Gaussian noise. Furthermore, we assume that the signal $s[n]$ is stationary, zero-mean, and its autocorrelation matrix \mathbf{R} is known. This signal is corrupted with white Gaussian noise $w[n]$ with variance σ_w^2 . This resultant signal we denote as $x[n]$, that is $x[n] = s[n] + \sigma_w w[n]$. We then proceed to filter $x[n]$

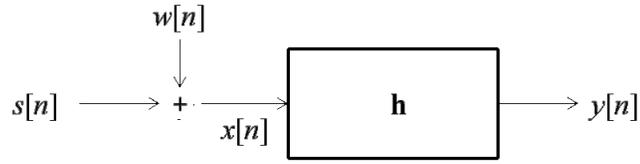


Figure A.1: Adapted from Moon and Stirling[112]

with an FIR filter of length N denoted by $\mathbf{h} \in \mathbb{R}^N$ that produces the output $y[n]$. This process is illustrated graphically in Figure 1. We select the FIR filter coefficients \mathbf{h} such that the output signal $y[n]$ recovers $x[n]$ (with a delay introduced by the filter).

A.2.1 Optimal Filter Formulation

The SNR at the output of the FIR filter is as follows (the details can be found in the end of this section):

$$SNR = \frac{\mathbf{h}^H \mathbf{R} \mathbf{h}}{\sigma_w^2 \mathbf{h}^H \mathbf{h}} \quad (\text{A.1})$$

where $(\cdot)^H$ denotes the Hermitian operation. The SNR in Equation (A.1) is also referred to as the Rayleigh quotient[113]. The FIR filter coefficients \mathbf{h} are unknown, and we would like to select them in such a way that they maximize the SNR in Equation (A.1). For ease of notation, we consider the real case and so the Hermitian operator will be denoted with the transpose operation $(\cdot)^T$. To avoid a trivial or non-volatile solution, in addition to maximizing Equation (A.1), a constraint is introduced on the filter coefficients \mathbf{h} as follows:

$$\mathbf{h}^* = \arg \max_{\mathbf{h}} \{\mathbf{h}^T \mathbf{R} \mathbf{h} : \|\mathbf{h}\|_2 = 1\} \quad (\text{A.2})$$

where $\|\cdot\|_2$ denotes the l_2 norm of a vector (which can be computed as $\|\mathbf{h}\|_2 = \sqrt{\mathbf{h}^T \mathbf{h}}$) and we can drop σ_w^2 since it is a constant.

By the Courant-Fischer Theorem[113] the solution \mathbf{h}^* to Equation (A.2) is the eigenvector with the largest eigenvalue of the autocorrelation matrix \mathbf{R} . The FIR filter \mathbf{h}^* is optimal in the sense that it maximizes the SNR in Equation (A.1) from all family of filters \mathbf{h} that are from the set $\{\mathbf{h} : \|\mathbf{h}\|_2 = 1\}$. This filter, however, is in general not sparse. In what follows we propose a method of finding a filter that is close to the optimal filter \mathbf{h}^* but is also sparse (many coefficients are zero). While the sparse filter may

be sub-optimal in the sense that it is not the solution to Equation (A.2), the low computational cost of the filter is what we are interested in. Our method introduces a parameter that can be used to trade-off between the optimality as defined in Equation (A.2) and sparsity.

A.2.2 Sparse Filter Formulation

In this section we propose a method of approximating the optimal filter \mathbf{h}^* as found by Equation (A.2) with a sparse representation. This problem is akin to finding a sparse vector representation[114, 115, 116] and we propose the following program to approximate a sparse eigenvector:

$$\begin{aligned} & \underset{\mathbf{h}}{\text{minimize}} && \|\mathbf{h}^* - \mathbf{h}\|_2 \\ & \text{subject to} && \mathbf{Card}(\mathbf{h}) = L \end{aligned} \tag{A.3}$$

where $\mathbf{Card}(\cdot)$ is the cardinality function (the function that determines the number of non-zero elements in a vector), and $L \in \mathbb{N}$ with $L \leq N$ (with the case $L = N$ yielding the same solution in Equation (A.2)). We can interpret the objective function $\|\mathbf{h}^* - \mathbf{h}\|_2$ in Equation (A.3) as a measure of the closeness of the sparse filter \mathbf{h} to the optimal FIR filter \mathbf{h}^* while the single constraint function $\mathbf{Card}(\mathbf{h}) = L$ restricts the number of non-zero coefficients the sparse FIR filter \mathbf{h} contains. Clearly, if $L = N$ the sparse filter will equal the optimal FIR filter (that is, $\mathbf{h} = \mathbf{h}^*$), and the optimal value of the objective function in Equation (A.3) will be zero.

A different but similar way to finding the sparse filter is to consider the following nearly-equivalent program:

$$\underset{\mathbf{h}}{\text{minimize}} \quad \|\mathbf{h}^* - \mathbf{h}\|_2 + \gamma \mathbf{Card}(\mathbf{h}) \tag{A.4}$$

where $\gamma \geq 0$ is a tradeoff parameter that the user selects. The interpretation of solving Equation (A.4) is that we trade off between the closeness to the optimal filter \mathbf{h}^* as measured by $\|\mathbf{h}^* - \mathbf{h}\|_2$ with the sparsity of the filter \mathbf{h} as measured by $\mathbf{Card}(\mathbf{h})$. The parameter γ is our tradeoff parameter with $\gamma = 0$ yielding the same solution as the optimal filter (but resulting in no sparsity) and larger γ placing more emphasis on improving sparsity. In practice one should vary γ to find the appropriate tradeoff for the specific

application. The equivalent tradeoff can be accomplished in Equation (A.3) by varying the sparsity constraint value L . While solving either Equation (A.3) or Equation (A.4) will yield sparse approximations, finding the optimal sparse solution is not trivial since the computations involve a combinatorial search space which can be prohibitively large. We hence revert to approximate solutions that can be solved efficiently.

Recent sparse methods propose to approximate the cardinality function with the l_1 norm, since the l_1 norm is the convex hull of the cardinality function (that is, the l_1 norm is the closest convex function to the non-convex cardinality function). While this may seem to be a poor approximation and it is suboptimal to the cardinality function, it often yields solutions that are sufficiently close to the desired design parameters and the gain in computability is an attractive advantage to the small loss in optimality. Thus, in this paper we find sparse eigenvector approximations to the optimal FIR filter \mathbf{h}^* by solving the following convex optimization program:

$$\underset{\mathbf{h}}{\text{minimize}} \quad \|\mathbf{h}^* - \mathbf{h}\|_2 + \gamma \|\mathbf{h}\|_1 \quad (\text{A.5})$$

where again $\gamma \geq 0$ is the tradeoff parameter. As in Equation (A.4), the parameter γ trades off between the closeness of the filter to the optimal FIR filter \mathbf{h}^* and the “sparsity” of the filter as approximated with the l_1 norm $\|\mathbf{h}\|_1$. This method is very similar to LASSO regression as proposed by Tibshirani[117].

The objective function in Equation (A.5) is convex since l_2 and l_1 are both convex functions (any norm is convex) and the positive linear combination of two convex functions is also convex[118]. The appeal of a convex optimization problem is the guarantee in finding globally optimal solutions. This makes the issue of developing solvers to be simple. In the case of Equation (A.5), we can actually solve this using an off-the-shelf quadratic program (QP) solver that solves the equivalent program:

$$\begin{aligned} \underset{\mathbf{h}}{\text{minimize}} \quad & \mathbf{h}^T \mathbf{h} - 2\mathbf{h}^{*T} \mathbf{h} \\ \text{subject to} \quad & -\delta \leq h_i \leq \delta, \quad i = 1, \dots, N \end{aligned} \quad (\text{A.6})$$

where here the role of $\delta \geq 0$ is very similar to the tradeoff parameter γ . Equation (A.6) is referred to as a QP, since the objective function is of quadratic form in the unknown variable \mathbf{h} and the constraints are linear inequalities.

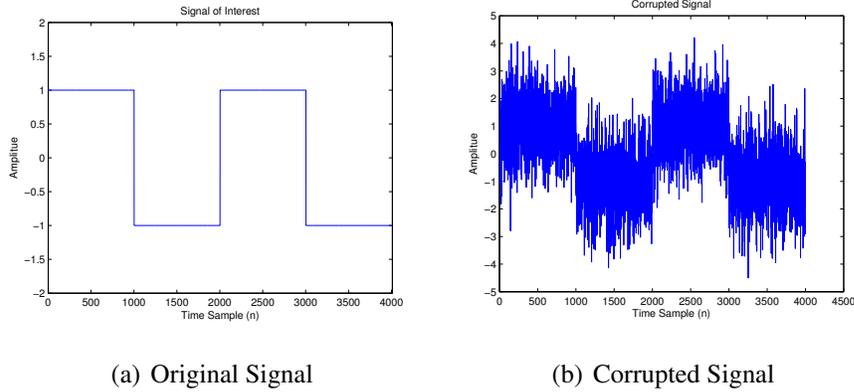


Figure A.2: Original and Corrupted Signal

A.3 Simulation Results

In this section we report results on three different experiments. In the first experiment we are interested in improving the detection time of switching states in a binary signal. In the second experiment we are interested in improving the visibility (that is the SNR) of a known sinusoid corrupted with white Gaussian noise, and in the last experiment we consider improving the visibility of a random signal realization corrupted with white Gaussian noise. We compare the performance of the traditional optimal filter (the Matched Filter) with the sparse filter we proposed in the previous section. We also compare the complexity of the traditional filter with the sparse filter (recall that we define complexity as the number of non-zero coefficients in the FIR filter).

A.3.1 Detection of Switching States

A given realization of the signal of interest is shown in Figure A.2(a): a binary signal with two states in $\{+1, -1\}$. In this problem we are interested in detecting accurate switching times between the two states. The corrupted signal that we are given, however, is given in Figure A.2(b). Clearly, estimating the switching times in this corrupted signal is difficult. Estimating the original signal from this corrupted signal using the traditional optimal (matched) filter is shown in Figure A.3(a), the same estimation

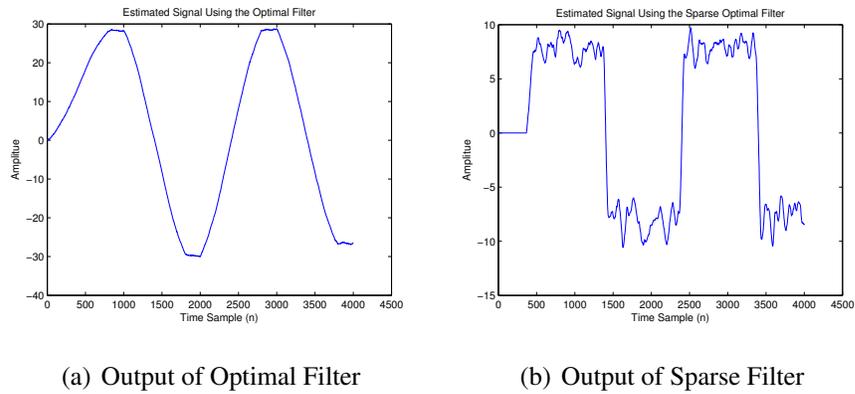


Figure A.3: Filter Outputs. Note that an appropriate scaling amplitude should be applied after filter to recover the appropriate amplitude of the original signal. However, in this problem we are most interested in switching times, so this step was not taken.

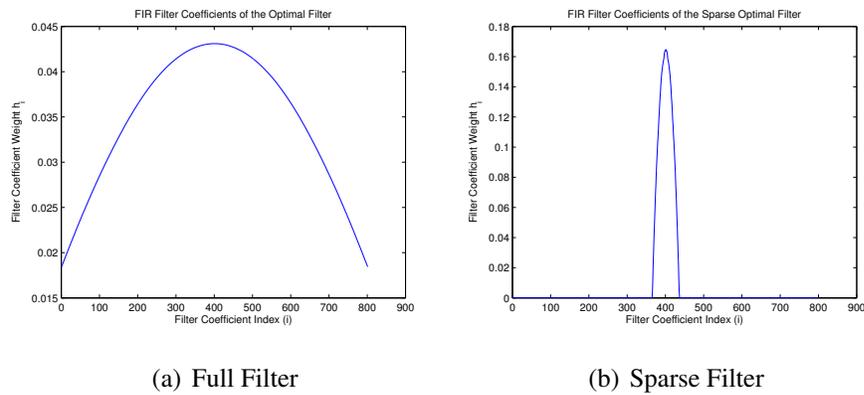


Figure A.4: Filter Coefficients

done using the sparse filter is shown in Figure A.3(b). While the output of the traditional optimal filter shown in Figure A.3(a) may be less noisy, we argue that detecting the switching times between the states is more accurate using the sparse filter in Figure A.3(b).

The filter coefficients for the traditional optimal filter and the filter coefficients for the sparse filter are shown in Figures A.4(a) and A.4(b) respectively. Clearly, the sparse filter has a much lower complexity (by a factor of 10) than the traditional filter (the sparse filter has a complexity of 77 while the full filter a complexity of 800). From these figures we can conclude that the reason the sparse filter is not only simpler but performs better is because the traditional filter does too much averaging. In the case of slowly varying signals, this would be beneficial, but when signals have sharp instantaneous changes (as is in the case of the binary signal considered here), too much averaging results in losing temporal resolution in the state changes.

A.3.2 Estimating Sinusoidal Signal Corrupted in White Gaussian Noise

In this section we consider a sinusoid buried in white Gaussian noise. Table A.1 compares the performance of three different filters for estimating the sinusoid corrupted in white Gaussian noise. The “full” filter is the optimal filter with complexity 200. The “sparse” filter is the sparse representation of the “full” filter but with a complexity of 16. The “simple” filter is the optimal filter but with the same complexity as the “sparse” filter. Clearly, the optimal “full” filter with complexity 200 results in the best estimation (smallest Mean Squared Error or MSE). But it is interesting to note that while the “sparse filter” has the same complexity as the “simple” filter, it has a much smaller MSE than the “simple” filter. This is highly desirable, since the “sparse” filter has a much lower complexity than the “full” filter, but also much better at estimating than the “simple” filter with the same complexity. Note that all these results are *extremely significant* at $p < 0.0001$ using a Wilcoxon rank sum test.

Sinusoid Estimation		
Filter Type	MSE \pm SD	Complexity
Full	$1.78 \times 10^{-4} \pm 3.03 \times 10^{-6}$	200
Sparse	$3.97 \times 10^{-4} \pm 2.24 \times 10^{-6}$	16
Simple	$4.55 \times 10^{-4} \pm 2.89 \times 10^{-6}$	16

Table A.1: Estimation of Sinusoidal Signals

A.3.3 Estimating Randomized Signal Corrupted in White Gaussian Noise

Table A.2 similarly compares the performance of the three different filters for estimating a random signal (a realization of a Gaussian process) corrupted with white Gaussian noise. Again, the “full” filter is the optimal filter with complexity 200, the “sparse” filter is the sparse representation of the “full” filter but with complexity 4, and the “simple” filter is the optimal filter but with the same complexity as as the “sparse” filter. As shown in the table, in this case the “sparse” filter has the *smallest* MSE, even better than the “full” filter with complexity 200. The reason for this may be similar to those reasons discussed in section 3.1. That is, the “full” filter performs too much averaging, and over-averaging for signals that change values quickly (as is the case of the random signal), the resolution of the signal is lost. Again, the results reported in this table are *extremely significant* at $p < 0.0001$ using a Wilcoxon rank sum test.

Random Signal Estimation		
Filter Type	MSE \pm SD	Complexity
Full	$2.13 \times 10^{-3} \pm 3.93 \times 10^{-6}$	200
Sparse	$1.8 \times 10^{-3} \pm 1.73 \times 10^{-6}$	4
Simple	$2.3 \times 10^{-3} \pm 4.04 \times 10^{-6}$	4

Table A.2: Estimation of Random Signal

A.4 Conclusions and Future Work

In this paper we proposed developing FIR filters that are optimal in improving SNR and are also sparse. We proposed developing a method by first designing the optimal filter (by maximizing the SNR resulting in finding the eigenvector with the largest eigenvalue of the signal data correlation matrix), and then to “sparsify” this filter by finding a set of FIR filter coefficients that are close to the optimal filter but are also sparse. We measured sparsity by using an l_1 norm approximation to the cardinality function. While such a method is only an approximation to the cardinality function, this optimization procedure is convex (guaranteeing a global optimal solution) and can also be solved efficiently with off-the-shelf solvers (when the problem is posed as a quadratic program such as in Equation (A.6)).

We demonstrated three experiments, one involving improving the detection of switching times of states in a binary signals, and two experiments in estimating signal corrupted with white Gaussian noise. These experiments illustrate that the sparse filter is an attractive option at significantly improving the estimation with a low complexity FIR filter. In certain cases, the sparse filter will even outperform the high-order FIR optimal filter with high complexity. While in this paper we focused on developing optimal FIR filters for improving SNR, we propose for future work in developing sparse FIR filters that are close to a given desired filter specification. We propose that our method of finding a sparse eigenvector for optimal SNR filters can be also extended to finding a sparse eigenvector for a given filter specification using the eigenfilter approach proposed by Nguyen[119].

A.5 Maximum SNR filters

In this section we show the derivation for the optimal filter that maximizes SNR (also known as the “Matched Filter”). We closely follow the formulation taken in Moon and Stirling[112]. Consider the signals as described in Figure 1. Let $\mathbf{s}[t] = (s[t], s[t - 1], \dots, s[t - N + 1])^T$ be the signal input, and let $\mathbf{x}[t] = (x[t], x[t - 1], \dots, x[t - N + 1])^T$ be the input to the filter $\mathbf{h} \in \mathbb{R}^N$. The output of the filter at time sample t is $y[t] = \mathbf{h}^T \mathbf{x}[t]$. The output power of the filter is thus

$$E\{y[t]^2\} = E\{\mathbf{h}^T \mathbf{x}[t] \mathbf{x}[t]^T \mathbf{h}\} = \mathbf{h}^T E\{\mathbf{x}[t] \mathbf{x}[t]^T\} \mathbf{h} = \mathbf{h}^T \mathbf{R} \mathbf{h} + \sigma_w^2 \mathbf{h}^T \mathbf{h} \quad (\text{A.7})$$

where $E\{\cdot\}$ denotes the expectation operator, \mathbf{R} is the data autocorrelation matrix computed by $E\{\mathbf{s}[t] \mathbf{s}[t]^T\}$, and σ_w^2 is the noise variance. The first term in Equation (A.7) is the power due the signal alone, and the second term is the power due to the noise alone. Thus, the SNR (signal power to noise power ratio) can be written as

$$SNR = \frac{\mathbf{h}^H \mathbf{R} \mathbf{h}}{\sigma_w^2 \mathbf{h}^H \mathbf{h}}$$

A.6 Notes

Appendix A, in full, is a modified reprint with minor revisions of the material as it appears in SPIE. Yazdani, M. and R. Hecht-Nielsen. “Optimal filters with heuristic 1-norm sparsity constraints.” Proc. SPIE 8137, 813709, 2011.

References

- [1] M. E. Rosheim, *Leonardo's Lost Robots*. Berlin, Germany: Springer, 2006.
- [2] A. P. Georgopoulos, A. B. Schwartz, and R. E. Kettner, "Neuronal population coding of movement direction," *Science*, vol. 233, no. 4771, pp. 1416–1419, 1986.
- [3] R. Lemon, "The output map of the primate motor cortex," *Trends in Neurosciences*, vol. 11, no. 11, pp. 501–506, 1988.
- [4] W. Penfield and E. Boldrey, "Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical stimulation," *Brain*, vol. 60, pp. 389–443, 1937.
- [5] P. S. Hammon, S. Makeig, H. Poizner, E. Todorov, and V. R. de Sa, "Predicting reaching targets from human eeg," *Ieee Signal Processing Magazine*, vol. 25, no. 1, pp. 69–77, 2008.
- [6] M. S. A. Graziano, C. S. R. Taylor, and T. Moore, "Complex movements evoked by microstimulation of precentral cortex," *Neuron*, vol. 34, no. 5, pp. 841–851, 2002.
- [7] S. F. Giszter, F. A. Mussa-Ivaldi, and E. Bizzi, "Convergent force fields organized in the frog's spinal cord," *Journal of Neuroscience*, vol. 13, no. 2, pp. 467–91, 1993.
- [8] K. Pearson, "Control of walking," *Scientific American*, 1976.
- [9] S. Giszter, "Spinal movement primitives and motor programs - a necessary concept for motor control," *Behavioral and Brain Sciences*, vol. 15, no. 4, pp. 744–745, 1992.
- [10] C. B. Hart and S. F. Giszter, "A neural basis for motor primitives in the spinal cord," *Journal of Neuroscience*, vol. 30, no. 4, pp. 1322–1336, 2010.
- [11] F. A. Mussa-Ivaldi, S. F. Giszter, and E. Bizzi, "Linear-combinations of primitives in vertebrate motor control," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 91, no. 16, pp. 7534–7538, 1994.

- [12] T. Flash and N. Hogan, “The coordination of arm movements - an experimentally confirmed mathematical-model,” *Journal of Neuroscience*, vol. 5, no. 7, pp. 1688–1703, 1985.
- [13] E. Todorov and M. I. Jordan, “Smoothness maximization along a predefined path accurately predicts the speed profiles of complex arm movements,” *Journal of Neurophysiology*, vol. 80, no. 2, pp. 696–714, 1998.
- [14] W. Abend, E. Bizzi, and P. Morasso, “Human arm trajectory formation,” *Brain*, vol. 105, no. Pt 2, pp. 331–48, 1982.
- [15] W. Bialek, “Thinking about the brain,” *Physics of Biomolecules and Cells: Les Houches Session LXXV, EDP Sciences, Les Ulis, Springer-Verlag, Berlin*, p. 485577, 2002.
- [16] N. Angier, “Seeing the natural world with a physicists lens.” *New York Times*, p. D2, November 1 2010.
- [17] E. Izhikevich, “Which model to use for cortical spiking neurons,” *IEEE Transactions on Neural Networks*, vol. 15, pp. 1063–1070, 2004.
- [18] G. Cramer and S. Darby, *Basic and Clinical Anatomy of the Spine, Spinal Cord, and ANS*. St. Louis, Missouri: Elsevier Mosby, 2005.
- [19] D. Clark, N. Boutros, and M. Mendez, *The Brain and Behavior*. Cambridge, United Kingdom: Cambridge University Press, 2008.
- [20] P. Morasso, “Spatial control of arm movements,” *Experimental Brain Research*, vol. 42, pp. 223–227, 1981.
- [21] E. Bizzi, N. Accornero, W. Chapple, and N. Hogan, “Posture control and trajectory formation during arm movement,” *Journal of Neuroscience*, vol. 4, no. 11, pp. 2738–2744–48, 1984.
- [22] B. Rexed, “The cytoarchitectonic organization of the spinal cord in the cat,” *J Comp Neurol*, vol. 96, pp. 414–495, 1952.
- [23] R. C.B., S. C.C., B. C., and H. P.R., “Stereologic characterization and spatial distribution patterns of betz cells in the human primary motor cortex,” *The Anatomical Record*, vol. 270A, pp. 137–151, 2003.
- [24] R. Lemon and J. Griffiths, “Comparing the function of the corticospinal system in different species : Organizational differences for motor specialization?” *Muscle and Nerve*, vol. 32, pp. 261–279, 2005.
- [25] E. Todorov, “Direct cortical control of muscle activation in voluntary arm movements: a model,” *Nature Neuroscience*, vol. 3, pp. 391 – 398, 2000.

- [26] S. Kakei, D. Hoffman, and P. Strick, "Muscle and movement representations in the primary motor cortex," *Science*, vol. 285, pp. 2136 – 2139.
- [27] D. Moran and A. Schwartz, "Motor cortical representation of speed and direction during reaching," *Journal of Neurophysiology*, vol. 82, pp. 2676 – 2692, 1999.
- [28] E. Eldred, R. Granit, and P. Merton, "Supraspinal control of the muscle spindles and its significance," *Journal of Physiology*, vol. 122, pp. 498–523, 1953.
- [29] F. Bretzner and T. Drew, "Contribution of the motor cortex to the structure and the timing of hindlimb locomotion in the cat: A microstimulation study," *Journal of Physiology*, vol. 94, pp. 657 – 672, 2005.
- [30] G. Loeb, "Learning from the spinal cord," *Journal of Physiology*, vol. 533, pp. 111–117, 2001.
- [31] D. Ramanathan, J. Conner, and M. Tuszynski, "A form of motor cortical plasticity that correlates with recovery of function after brain injury," *Proc. Natl. Acad. Sci.*, vol. 103, pp. 11 370 – 11 375, 2006.
- [32] E. Evarts, "Relation of pyramidal tract activity to force exerted during voluntary movement," *Journal of Neurophysiology*, vol. 31, pp. 14–27, 1968.
- [33] S. Scott and J. Kalaska, "Changes in motor cortex activity during reaching movements with similar hand paths but different arm postures," *Journal of Neurophysiology*.
- [34] G. A.P., A. Schwartz, and R. Kettner, "Neuronal population coding of movement direction," *Science*, vol. 233, pp. 1416 – 1419, 1986.
- [35] M.-H. Boudrias, M. Park, and P. Cheney, "Contrasting properties of motor output from the supplementary motor area and primary motor cortex in rhesus macaques," *Cerebral Cortex*, vol. 16, pp. 632–638, 2006.
- [36] J. Maier, M.A. Armand, P. Kirkwood, H.-W. Yang, J. Davis, and R. Lemon, "Differences in the corticospinal projection from primary motor cortex and supplementary motor area to macaque upper limb motoneurons: An anatomical and electrophysiological study," *Cerebral Cortex*, vol. 12, pp. 281–296, 2002.
- [37] H. Freund and H. Hummelsheim, "Lesions of premotor cortex in man," *Brain*, vol. 108, pp. 697–733.
- [38] P. Viviani and T. Flash, "Minimum-jerk, 2/3-power law, and isochrony - converging approaches to movement planning," *Journal of Experimental Psychology-Human Perception and Performance*, vol. 21, no. 1, pp. 32–53, 1995.
- [39] S. Giszter, "Motor primitives," in *Encyclopedia of Neuroscience*, E. in Chief: Larry R. Squire, Ed. Oxford: Academic Press, 2009, pp. 1023–1040.

- [40] A. Karniel and F. A. Mussa-Ivaldi, “Does the motor control system use multiple models and context switching to cope with a variable environment?” *Experimental Brain Research*, vol. 143, no. 4, pp. 520–524, 2002.
- [41] C. Bishop, *Pattern Recognition and Machine Learning*. Springer, 2007.
- [42] S. Ben-Itzhak and A. Karniel, “Minimum acceleration criterion with constraints implies bang-bang control as an underlying principle for optimal trajectories of arm reaching movements,” *Neural Computation*, vol. 20, no. 3, pp. 779–812, 2008.
- [43] D. E. Kirk, *Optimal control theory : an introduction*. Mineola, N.Y.: Dover Publications, 2004.
- [44] A. Gasparetto and V. Zanotto, “A technique for time-jerk optimal planning of robot trajectories,” *Robot. Comput.-Integr. Manuf.*, vol. 24, pp. 415–426, June 2008.
- [45] K. Kyriakopoulos and G. Saridis, “Minimum jerk path generation,” in *Robotics and Automation, 1988. Proceedings., 1988 IEEE International Conference on*, apr 1988, pp. 364 –369 vol.1.
- [46] U. Pattacini, F. Nori, L. Natale, G. Metta, and G. Sandini, “An experimental evaluation of a novel minimum-jerk cartesian controller for humanoid robots,” in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, oct. 2010, pp. 1668 –1674.
- [47] A. Piazzzi and A. Visioli, “Global minimum-jerk trajectory planning of robot manipulators,” *Industrial Electronics, IEEE Transactions on*, vol. 47, no. 1, pp. 140–149, feb 2000.
- [48] L. Botzer and A. Karniel, “A simple and accurate onset detection method for a measured bell-shaped speed profile,” *Frontiers in Neuroscience*, vol. 3, no. 61, pp. 1–8, 2009.
- [49] G. H. Staude, “Precise onset detection of human motor responses using a whitening filter and the log-likelihood-ratio test,” *Ieee Transactions on Biomedical Engineering*, vol. 48, no. 11, pp. 1292–1305, 2001.
- [50] G. H. Staude, W. M. Wolf, U. Appel, and R. Dengler, “Methods for onset detection of voluntary motor responses in tremor patients,” *Ieee Transactions on Biomedical Engineering*, vol. 43, no. 2, pp. 177–188, 1996.
- [51] G. Raphael, G. A. Tsianos, and G. E. Loeb, “Spinal-like regulator facilitates control of a two-degree-of-freedom wrist,” *Journal of Neuroscience*, vol. 30, no. 28, pp. 9431–9444, 2010.

- [52] D. Bullock and S. Grossberg, “Neural dynamics of planned arm movements - emergent invariants and speed accuracy properties during trajectory formation,” *Psychological Review*, vol. 95, no. 1, pp. 49–90, 1988.
- [53] J. T. Buchanan and D. R. McPherson, “The neuronal network for locomotion in the lamprey spinal cord: Evidence for the involvement of commissural interneurons,” *Journal of Physiology-Paris*, vol. 89, no. 4-6, pp. 221–233, 1995.
- [54] M. Hagglund, L. Borgius, K. J. Dougherty, and O. Kiehn, “Activation of groups of excitatory neurons in the mammalian spinal cord or hindbrain evokes locomotion,” *Nature Neuroscience*, vol. 13, no. 2, pp. 246–52, 2010.
- [55] E. Izhikevich, *Dynamical Systems in Neuroscience: Chapter 9*. Cambridge, M.A.: MIT Press, 2007.
- [56] M. Goldman, A. Compte, and X. Wang, “Neural integrator models,” in *Encyclopedia of Neuroscience*, E. in Chief: Larry R. Squire, Ed. Oxford: Academic Press, 2009, pp. 165–178.
- [57] D. A. Robinson, “Integrating with neurons,” *Annual Review of Neuroscience*, vol. 12, pp. 33–45, 1989.
- [58] H. Seung, D. Lee, B. Reis, and D. Tank, “Stability of the memory for eye position in a recurrent network of conductance-based model neurons,” *Neuron*, vol. 26, pp. 259–271, 2000.
- [59] S. Edelman and T. Flash, “A model of handwriting,” *Biological Cybernetics*, vol. 57, pp. 25–36, 1987.
- [60] E. Todorov, *Chapter 12.7: Optimal Control Theory in Bayesian Brain, probabilistic approaches to neural coding*, ser. Computational neuroscience, K. Doya, S. Ishii, A. Pouget, and R. P. N. Rao, Eds. Cambridge, Mass.: MIT Press, 2007, ch. 12.7 authored by Todorov, E.
- [61] L. Pontryagin, V. Boltyanskii, R. Gamkrelidze, and E. Mishchenko, *The Mathematical Processes of Optimal Processes*. New York: Interscience Publishers, 1962.
- [62] H. C.M. and W. D.M, “Signal-dependent noise determines motor planning,” *Nature*, vol. 394, no. 6695, pp. 790–784, 1998.
- [63] A. Feld’baum, *Mathematics in Science and Engineering*. New York: Academic Press, 1965.
- [64] M. Svinin, M. Yamamoto, and I. Goncharenko, “Simple models in trajectory planning of human-like reaching movements,” *The 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1662–1667, 2010.

- [65] P. Daya and L. Abbott, *Theoretical Neuroscience*. Cambridge, M.A.: MIT Press, 2001.
- [66] M. Grant and S. Boyd, “Cvx: Matlab software for disciplined convex programming, version 1.21,” 2011.
- [67] ———, “Graph implementations for nonsmooth convex programs,” in *Recent Advances in Learning and Control*, ser. Lecture Notes in Control and Information Sciences, V. Blondel, S. Boyd, and H. Kimura, Eds. Springer-Verlag Limited, 2008, pp. 95–110.
- [68] H. Nagasaki, “Asymmetric velocity and acceleration profiles of human arm movements,” *Experimental Brain Research*, vol. 74, pp. 319–326, 1989.
- [69] Y. Uno and M. Kawato, “Formation and control of optimal trajectory in human multijoint arm movement,” *Biological Cybernetics*, vol. 61, no. 2, pp. 89–101, 1989.
- [70] D. Liu and E. Todorov, “Evidence for the flexible sensorimotor strategies predicted by optimal feedback control,” *Journal of Neuroscience*, vol. 27, pp. 9354–9368, 2007.
- [71] B. Hoff and M. Arbib, “Models of trajectory formation and temporal interaction of reach and grasp,” *Journal of Motor Behavior*, vol. 25, pp. 175–192, 1993.
- [72] B. Berret, E. Chiovetto, F. Nori, and T. Pozzo, “Evidence for composite cost functions in arm movement planning: An inverse optimal control approach,” *PLoS Comput Biol*, vol. 7, no. 10, p. e1002183, 2011. [Online]. Available: <http://dx.doi.org/10.1371/journal.pcbi.1002183>
- [73] D. Angelaki, “Vestibulo-ocular reflex,” in *Encyclopedia of Neuroscience*, E. in Chief: Larry R. Squire, Ed. Oxford: Academic Press, 2009, pp. 139 – 146. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/B9780080450469011074>
- [74] A. Wiegner and M. Wierzbicka, “Kinematic models and human elbow flexion movements: quantitative analysis,” *Experimental Brain Research*, vol. 88, no. 3, pp. 665–673, 1992.
- [75] R. Plamondon, A. Alimi, P. Yergeau, and F. Leclerc, “Modelling velocity profiles of rapid movements: a comparative study,” *Biological Cybernetics*, vol. 69, no. 2, pp. 119–128, 1993.
- [76] S. Edelman and T. Flash, “A model of handwriting,” *Biological Cybernetics*, vol. 57, pp. 25–36, 1987.

- [77] M. E. Richardson and T. Flash, “Comparing smooth arm movements with the two-thirds power law and the related segmented-control hypothesis,” *Journal of Neuroscience*, vol. 15, no. 2, pp. 8201–8211, 2002.
- [78] A. Biess, D. G. Liebermann, and T. Flash, “Comparing smooth arm movements with the two-thirds power law and the related segmented-control hypothesis,” *Journal of Neuroscience*, vol. 27, no. 48, pp. 13 045–13 064, 2007.
- [79] T. Krasovsky, S. Berman, and D. Liebermann, “Kinematic models and human elbow flexion movements: quantitative analysis,” *Journal of Electromyography and Kinesiology*, vol. 20, no. 4, pp. 636–641, 2010.
- [80] M. Berniker and K. Kording, “Estimating the sources of motor errors for adaptation and generalization,” *Nature Neuroscience*, vol. 11, pp. 1454 – 1461, 2008.
- [81] T. Matheson, “Motor planning: Insects do it on the hop,” *Curr Biol.*, vol. 18, pp. 742–743, 2008.
- [82] C. Fernandez and J. Goldberg, “Physiology of peripheral neurons innervating otolith organs of the squirrel monkey. i. response to static tilts and to long-duration centrifugal force.” *J Neurophysiol*, vol. 39, p. 970984, 1976.
- [83] ———, “Physiology of peripheral neurons innervating otolith organs of the squirrel monkey. ii. directional selectivity and force-response relations.” *J Neurophysiol*, vol. 39, p. 985995, 1976.
- [84] ———, “Physiology of peripheral neurons innervating otolith organs of the squirrel monkey. iii. response dynamics.” *J Neurophysiol*, vol. 39, p. 9961008, 1976.
- [85] ———, “Physiology of peripheral neurons innervating semicircular canals of the squirrel monkey. ii response to sinusoidal stimulation and dynamics of peripheral vestibular system response dynamics.” *J Neurophysiol*, vol. 34, p. 661 675, 1971.
- [86] ———, “Physiology of peripheral neurons innervating semicircular canals of the squirrel monkey. iii. variations among units in their discharge properties.” *J Neurophysiol*, vol. 34, p. 676 684, 1971.
- [87] ———, “Physiology of peripheral neurons innervating semicircular canals of the squirrel monkey. i. resting discharge and response to constant angular accelerations.” *J Neurophysiol*, vol. 34, p. 635 660, 1971.
- [88] M. P. Mileusnic and G. E. Loeb, “Force estimation from ensembles of golgi tendon organs.” *J. Neural Eng.*, vol. 6, no. 3, 2009.
- [89] T. R. Nichols, “Reflex circuits,” in *Encyclopedia of Neuroscience*, E. in Chief: Larry R. Squire, Ed. Oxford: Academic Press, 2009, pp. 73–79.

- [90] E. Fransn, B. Tahvildari, A. V. Egorov, M. E. Hasselmo, and A. A. Alonso, "Mechanism of graded persistent cellular activity of entorhinal cortex layer v neuron," *Neuron*, vol. 49, no. 5, pp. 735 – 746, 2006.
- [91] S. P. Blair, H.T., "Anticipatory head direction signals in anterior thalamus: evidence for a thalamocortical circuit that integrates angular head motion to compute head direction." *Journal of Neuroscience*, vol. 15, p. 62606270, 1995.
- [92] J. S. Taube and J. P. Bassett, "Persistent neural activity in head direction cells," *Cerebral Cortex*, vol. 13, no. 11, pp. 1162–1172, 2003.
- [93] D. L. Sparks, "The brainstem control of saccadic eye movements," *Nature Reviews Neuroscience*, vol. 3, no. 12, pp. 952–964, 2002.
- [94] E. Aksay, I. Olasagasti, B. D. Mensh, R. Baker, M. S. Goldman, and D. W. Tank, "Functional dissection of circuitry in a neural integrator," *Nature Neuroscience*, vol. 10, no. 4, pp. 494–504, 2007.
- [95] H. Seung, "How the brain keeps the eyes still." *Proceedings of the National Academy of Sciences of the United States of America*, vol. 93, p. 1333913344, 1996.
- [96] K. Fukushima, C. R. Kaneko, and A. F. Fuchs, "The neuronal substrate of integration in the oculomotor system," *Progress in Neurobiology*, vol. 39, no. 6, pp. 609 – 639, 1992. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0301008292900168>
- [97] K. Fukushima and C. R. Kaneko, "Vestibular integrators in the oculomotor system," *Neuroscience Research*, vol. 22, no. 3, pp. 249 – 258, 1995. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/0168010295009048>
- [98] L. DellOsso, *Chapter 2: Neural Integration In Ocular Motility in: Neurological organization of ocular movement*, R. Daroff, K. A. Neetens, and Ghedini, Eds. Berkeley, CA, U.S.A: Kugler Publications, 1990.
- [99] M. Goldman, A. Compte, and X. Wang, "Neural integrator models," in *Encyclopedia of Neuroscience*, E. in Chief: Larry R. Squire, Ed. Oxford: Academic Press, 2009, pp. 165–178.
- [100] D. Anastasopoulos, T. Haslwanter, M. Fetter, and J. Dichgans, "Smooth pursuit eye movements and otolith-ocular responses are differently impaired in cerebellar ataxia," *Brain*, vol. 121, no. 8, p. 14971505, 1998.
- [101] R. Burke, "Motor units: Anatomy, physiology and functional organization," in *The Nervous System*, V. Brooks, Ed., 1981, vol. 1, ch. 1, pp. 345–422.

- [102] R. Lemon, “Descending pathways in motor control,” *Annual Review of Neuroscience*, vol. 31, pp. 195–218, July 2008.
- [103] C. B. Hart and S. F. Giszter, “A neural basis for motor primitives in the spinal cord,” *Journal of Neuroscience*, vol. 30, no. 4, pp. 1322–1336, 2010.
- [104] D. M. Wolpert, J. Diedrichsen, and J. R. Flanagan, “Principles of sensorimotor learning,” *Nat Rev Neuroscience*, vol. 12, no. 12, pp. 739–51, 2011.
- [105] A. dAvella, A. Portone, L. Fernandez, and F. Lacquaniti, “Control of fast-reaching movements by muscle synergy combinations.” *J. Neurosci.*, vol. 26, p. 77917810, 2006.
- [106] G. Torres-Oviedo, J. M. Macpherson, and L. H. Ting, “Muscle synergy organization is robust across a variety of postural perturbations.” *Journal of Neurophysiology*, vol. 96, no. 3, pp. 1530–1546, 2006.
- [107] A. Oppenheim and R. Schafer, *Discrete-Time Signal Processing*. Prentice Hall, 2009.
- [108] S. Kay, *Fundamentals of Statistical Signal Processing, Volume 2: Detection Theory*. Prentice Hall, 1998.
- [109] I. Lakkis and D. McLernon, “Optimum eigenfilters and matched filters,” *Electronics Letters*, 1996.
- [110] J. Makhoul, “On the eigenvectors of symmetric toeplitz matrices,” *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 29, pp. 868 – 872, 1981.
- [111] T. Baran, D. Wei, and A. Oppenheim, “Linear programming algorithms for sparse filter design,” *IEEE Transactions on Signal Processing*, vol. 58, pp. 1605–1617, 2010.
- [112] T. Moon and W. Stirling, *Mathematical Methods and Algorithms for Signal Processing*. Prentice Hall, 1999.
- [113] C. Meyer, *Matrix Analysis and Applied Linear Algebra*. SIAM, 2001.
- [114] H. Zou, T. Hastie, and R. Tibshirani, “Sparse principle component analysis,” *Technical Report. Statistics Department, Stanford*, 2004.
- [115] A. dAspermont, L. El Ghaoui, M. Jordan, and G. Lanckriet, “A direct formulation for sparse pca using semidefinite programming,” *Advances in Neural Information Processing Systems*, vol. 17, pp. 41–48, 2005.
- [116] B. Sriperumbudur, D. Torres, and G. Lanckriet, “Sparse eigenmethods by d.c. programming,” *Proceedings of the 24th International Conference on Machine Learning*, 2007.

- [117] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society, Series B*, vol. 58, pp. 267–288, 1996.
- [118] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [119] T. Nguyen, “The design of arbitrary fir digital filters using the eigenfilter method,” *IEEE Transactions on Signal Processing*, vol. 41, pp. 1128–1139, 1993.