UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Biologically Plausible Control of Fast Reaching Movements Using Non-Traditional Cost Functions

Permalink

Author

Gamble, Geoffrey

Publication Date

2016

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Biologically Plausible Control of Fast Reaching Movements Using Non-Traditional
Cost Functions

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

in

Computer Science

by

Geoffrey George Gamble

Committee in charge:

    Professor Garrison Cottrell, Co-Chair
    Professor Robert Hecht-Nielsen, Co-Chair
    Professor Gert Cauwenberghs
    Professor William Griswold
    Professor David Kriegman

2016

The Dissertation of Geoffrey George Gamble is approved and is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

_____

Co-Chair

_____

Co-Chair

University of California, San Diego

2016

DEDICATION

This work is dedicated to my parents, George and Judith Gamble, who have been supportive of my endeavors, scientific and otherwise. I thank you for your boundless patience, unwavering support and endless encouragement. I love you very much.

The overwhelming majority of theories are rejected because they contain bad explanations, not because they fail experimental tests.

David Deutsch

The truly privileged theories are not the ones referring to any particular scale of size or complexity, nor the ones situated at any particular level of the predictive hierarchy, but the ones that contain the deepest explanations.

David Deutsch

No great discovery was ever made without a bold guess.

Isaac Newton

TABLE OF CONTENTS

vii

LIST OF FIGURES

PREFACE

In Sir Isaac Newton's seminal work, *Philosophiae Naturalis Principia Mathematica (The Mathematical Principles of Natural Philosophy, third edition)*, he made a subtle but important point regarding his newly conceived laws of gravitation. He states:

"Hitherto we have explained the phenomena of the heavens and of our sea by the power of gravity, *but have not yet assigned the cause of this power, I have not been able to discover the cause of those properties of gravity from phenomena, and I frame no hypotheses*".

The many translations of his work from Latin to other languages vary slightly with their perceived meaning. Still, Newton, perhaps unwittingly, was emphasising a new focus in science...A type of science that seeks to discover mathematical rules, but in addition, places emphasis on discovering the underlying *explanations* of how and why those rules exist. I call this distinction the "what happens" vs. "why and how does this happen". Newton freely admitted that, although he ascertained mathematical rules that predicted the gravitational relationship between distance and mass, which helped us to predict the motion of the planets and tides, he had absolutely no understanding of *why and how* bodies of matter attract each other in this way. He wondered why he could predict the gravitational forces between particles, but could not understand why these forces exist.

My work follows in the vein of Sir Isaac Newton's scientific principles. I explore *what* happens with human movement (as described by mathematical rules) and with these rules in mind, I attempt to move closer to explaining *why and how* these movements occur. Ultimately, I demonstrate a new framework that can be used to support the connection of these two concepts with regard to some aspects of human movement. This effort places value on the *explanation* of observed

phenomena over their *prediction*, but still uses *prediction* as a justification for an *explanation*. I find that discovering an explanation for a given phenomenon to be much more satisfying than finding a formula for predicting some outcome without understanding the underlying mechanisms. If you can explain, then you can predict, in which case you have the best of both worlds.

ACKNOWLEDGEMENTS

simple control policy for achieving minimum jerk trajectories.", Neural Networks, vol. 27, pp. 74-80, 2012.

My role in this effort included:

- Acquired, parsed and deciphered a raw data set of human reaching movements

- Matlab programming

- Literature review

- Co-wrote text

Chapter 3 is a reprint of the material as it appears on Cornell University's arXiv. Gamble, G.G.; Yazdani, M. "Sparse signals for the control of human movements using the infinity norm", http://arxiv.org/abs/1601.01410, 2016.

Chapter 4 is an augmented reprint of the material that has been submitted to Neural Networks, Gamble, G.G.; Yazdani, M.; Hecht-Nielsen, R. "Suboptimal combinations of cost functions fit fast human reaching movements as well as optimal combinations", 2015.

VITA

EMPLOYMENT

- Co-founder of Augmented Data, 2016-present

- Graduate Student Researcher at the University of California, San Diego, 2010 – 2015

- Senior Software Engineer at XIFIN, San Diego, CA, Jan. 2004 – Feb. 2010

- Senior Software Engineer at Optimal Gain, Boston, MA, Sept. 2001 – Oct. 2003

- Software Engineer/technologist at Zefer, Boston, MA, Sept. 1999 – Aug. 2001

- Software Engineer at Exchange Applications, Boston, MA, Sept. 1998 – Sept. 1999

EDUCATION

2016   Doctor of Philosophy in Computer Science, University of California, San Diego

- Focus in computational neuroscience

2011   Masters in Computer Science, University of California, San Diego

- Overall GPA 3.74/4.0
- Focus in software engineering
- Winner of Frieda Daum Urey Fellowship

1998   Bachelor of Arts majoring in computer science, Boston College

- Major: Computer Science
- Graduated Cum Laude

- 3.7/4.0 GPA within computer science major
- Graduated with computer science departmental honors

PUBLICATIONS

- Gamble, G.G.; Yazdani, M. "Sparse signals for the control of human movements using the infinity norm", arXiv:http://arxiv.org/abs/1601.01410, 2016.

- Gamble, G.G.; Yazdani, M; Hecht-Nielsen, R. "Suboptimal combinations of cost functions fit fast human reaching movements as well as optimal combinations", submitted to Neural Networks, Oct. 2015.

- Yazdani, M.; Gamble, G.G.; Henderson, G.; Hecht-Nielsen, R.; "A simple control policy for achieving minimum jerk trajectories", Neural Networks, vol.27, pp 74-80, 2012.

- Gamble, G.G.; Yazdani, M.; et al. (2010) A model of spinomuscular posture attainment. Program No. 109.1. 2010 Neuroscience Meeting Planner. San Diego, CA: Society for Neuroscience, 2010. Online.

- Yazdani, M.; Gamble, G. G.; et al (2010) A neuronal network model of the static postural goal circuit in the spinal cord. Program No. 108.6. 2010 Neuroscience Meeting Planner. San Diego, CA: Society for Neuroscience, 2010. Online

- Henderson, G.; Gamble, G.G.; et al (2010) The Role of Oscillatory Networks in Motor Programs. Proceedings of the fourteenth annual International Conference on Cognitive and Neural Systems. Boston, MA: page 79.

FIELDS OF STUDY

Computational Neuroscience
Professor Robert Hecht-Nielsen

Optimal Control
Self taught with tutelage from Dr. Mehrdad Yazdani

Software Architecture
Professors Ingolf Krueger and William Griswold

ABSTRACT OF THE DISSERTATION

Biologically Plausible Control of Fast Reaching Movements Using Non-Traditional
Cost Functions

by

Geoffrey George Gamble

Doctor of Philosophy in Computer Science

University of California, San Diego, 2016

Professor Garrison Cottrell, Co-Chair
Professor Robert Hecht-Nielsen, Co-Chair

Optimal control has been used as a technique to uncover mathematical principles which are observed regularly in the dynamics of human movement. We present two new models of human reaching movements. While both are rooted in optimal control theory, the models were conceived by questioning basic tenets and typical practices used in optimal control as applied to human movement. In the first model, we use cost functions that measure various control signals via the L_∞ norm as opposed to the commonly used L_2 norm. Doing so models human

reaching movements as well as current approaches, but results in control signals that can be reasoned about in terms of neural spikes and their timing. In the second model, we change the organization of the terms within a single, multi-term cost function by transforming it into many single-term cost functions. This approach yields sub-optimal results with regard to cost, yet produces equal or better results when applied to accuracy in modeling human reaching movements. The traditional optimal control approach to modeling human movement assumes that humans have an optimal design in terms of the anatomy and physiology of their motor systems. This design is assumed to optimally minimize costs such as energy consumption, or error while attaining a goal. However, it is more likely that in changing environments/niches, humans and other animals are still evolving, and therefore have not yet arrived at an optimal design. By reorganizing the terms of a cost function in a cost-suboptimal way, while achieving high accuracy with regard to modeling the movements, we challenge the basic premise of cost-optimality that underlies optimal control based models of human movement. Additionally, this reorganization of cost function terms into multiple cost functions results in multiple, interacting control signals, making it possible to combine the these signals in ways that resemble the connectivity of the human motor system, which contains a diverse set of neural signals working in concert, each with its own character and purpose. For this reason, we introduce a framework that generalizes these concepts, which can be utilized for further modeling of human movement. The framework expands upon traditional optimal control as applied to modeling human movements by supporting multiple interacting control signals. This allows for experiments which more closely resemble the neural architecture of the motor system, thereby making it easier to reason about experimental results in terms of the construction of the human motor system.

Chapter 1

Introduction

1.1 The problem of understanding the neural basis of motor control in animals

Animals move in a myriad of complex yet graceful ways, from a monkey swinging from tree to tree, to an eagle soaring high in the sky, to a gymnast artfully flipping and tumbling across a mat. Consider a cheetah in the savannah pursuing a gazelle. Both animals are extremely nimble and extremely fast, each with maximum speeds of nearly 100 km/h. In order for the cheetah to capture its prey, and conversely for the gazelle to avoid capture, the animals must incorporate many different types of information into the decisions they make during each moment in order to achieve their respective goals. This information includes, but is not limited to, their position, velocity, acceleration, rate of change in acceleration, as well as their opponent's set of dynamic variables, their surrounding environment (e.g. obstacles in their paths). In addition, their internal states must also be incorporated into their decision making, e.g. what they have learned from similar past experiences, their sense of balance, and their sense of where their limbs currently reside. All of these variables are somehow represented within the animals via some neural substrates.

1.1.1 Multiple neural signals in the animal motor system

Figuring out how these and other animals sense and encode these variables into *neural signals* is only part of the problem. The signals from a given component of their nervous systems combine with signals from other components, and may influence many other components directly or indirectly. Many of these components are very different from one another regarding their form and function. This complicated, yet apparently well orchestrated symphony of neural signals somehow drive hundreds of muscles in the animals to produce the incredibly adept movements we observe.

1.1.2 Problem of relating biological computation to traditional computation

Given the myriad of variables and the complexity of the system, how does the nervous system perform the computations necessary to achieve these results? It is generally accepted that the brain does not compute in a way that is reflected in a modern computer based on the von Neumann architecture, and for all but the simplest of organisms, it is not known how the above mentioned motor control computations occur. As alluded to in section 1.1.1, there are a myriad of brain structures/systems involved in motor control in some capacity. In humans, this list includes (but is not limited to) the cortex, the cerebellum, various structures in the brain stem, the spinal cord, and sensors in various locations (e.g. the inner ear, muscles, and skin), that all relay state information back to the central nervous system.

We have yet to come close to building a machine that can perform a breadth of motor tasks at human level of proficiency. This fact is indicative of how far we have to go in order to fully understand the human motor system in terms of

anatomy and physiology. When we do build a machine that performs a small subset of human-like motor tasks well, its architecture typically has very little to do with the architecture of the human motor system (e.g. a mechanical arm at a car assembly plant doesn't reflect the control and architecture of the human arm). Current robotic solutions may mimic human/animal movements in some respects, and are extremely useful and cleverly designed, but typically do not inform us about how the human motor system functions.

1.2 Why explore the mechanisms of the human motor system?

1.2.1 Understanding our own construction

Hitherto, we have referred to the "human motor system" and the "animal motor system", henceforth, we may refer to the them as the "HMS" and "AMS" respectively. Because this work focuses on human movement, if not specified explicitly, we are referring to the HMS. Part of the motivation for studying the HMS comes from an innate desire to understand one's self. Understanding the HMS will help us to understand our own construction in terms of both form and function. The study of the HMS directly helps us to understand the most complicated system in the known universe, the human brain, hence helps us towards understanding the nature of thought.

More immediately, the study and understanding of the HMS will help to uncover the causes (and potential cures) of various motor diseases e.g. Parkinson's disease and amyotrophic lateral sclerosis (ALS, or Lou Gehrig's disease). In this work, we address the very real possibility that models of movement presented here are very likely to be able to detect motor irregularities leading to early detection and treatment. Additionally, such an understanding may provide the means to aid

ongoing efforts to create limb replacements which are controlled naturally by the nervous system, helping those with birth defects or catastrophic injury.

1.2.2 Human-like robots

Understanding the HMS will bring us to the doorstep of creating human-like robots. This notion does not refer to "thinking machines", but instead, well programmed machines that possess the graceful motor skills needed to navigate complex terrains that are typically accessible only to humans. The benefits of such an accomplishment are many...These robots might be capable of carrying out dangerous tasks that currently only humans are adept at, but risk their safety in order to perform (e.g. mining, or deep sea and planetary exploration). Additionally, menial yet intricate tasks could be offloaded to these robots (e.g. sewing or iPad assembly), allowing global economies to become more efficient, while the need for child labor and dangerous factory conditions might be reduced or eliminated. Finally, as developed nations face aging populations, human-like robots might possess the gentle touch needed to assist the elderly, prolonging their quality of life. Similar to the AMSs mentioned earlier in section 1.1.1, these robots would likely require multiple motor and sensory signals that interact with each other in complex ways in order to achieve the varied movements that animals are capable of. Modeling the complex interaction of multiple signals of different types is one of the motivations for this work. We discuss a unique approach to this problem in section 5.1.

1.3 What is optimal control and how is it applied to the exploration of motor systems function?

"Optimal control" has been used to find solutions to diverse problems in many fields (e.g. equity trading, [2], spacecraft control, [3], and factory control [4]). These applications of optimal control relate to the engineering of a solution to a problem that is "better", in some sense, than any other solution....Perhaps the optimal solution is more accurate, more reliable, more productive, or faster than other solutions based on standard engineering techniques, which commonly rely upon the cunning, ability, and intuition of the engineer.

In this work, we, and others that will be discussed, do not use optimal control to solve traditional engineering problems. Instead, we seek the underlying principles that govern human movement. In a sense, we use optimal control to *reverse engineer* the motor system by taking educated guesses at what principles the "engineer" (evolution) may have used to guide its "decisions" when crafting the AMS. A hallmark of this work is that we provide a path which gives the pure mathematics of optimal control a more plausible mapping to actual neural function when compared to other optimal control approaches, thus, getting a step closer to the goal of reverse engineering the AMS.

Optimal control is a broad topic with many facets, a small subset of these are used in this work. We have chosen to adapt approaches which have been successful in the past at modeling human movements, while at the same time, questioning their fundamental assumptions.

In this chapter, we will discuss the high level concepts necessary for the uninitiated reader to follow the rest of the work, and why optimality principles

may, or may not, underly the way in which the HMS has evolved. More detailed aspects of optimal control will be discussed as needed in the following chapters.

1.3.1 Optimal control concepts relevant to this work

Loosely speaking, a "controller" produces a "control signal" to drive a system to some desired state while satisfying a set of constraints. For example, the controller of a self driving car might seek to get from point A to point B. One of the constraints on the controller might limit the maximum speed of the car to 80 mph. The controller would then send a "control signal" to the accelerator to move towards point B without violating the constraint of not surpassing 80 mph.[1]

Consider a human that wants to move her hand from point A to point B. An infinite number of trajectories exist that will deliver the hand from one point to the other, see figure 1.1. There are also infinite permutations of sequences of muscle activations that could lead to the hand reaching its destination. However, when performing these movements, the (properly functioning) HMS selects a very small subset of these trajectories and muscle activation permutations. The HMS routinely selects highly predictable and stereotypical paths for reaching movements. The same is true for other related types of movements, e.g. catching a baseball [5]. How and why does the HMS select this small subset of movement trajectories from an infinite pool of choices? This is known as the "redundancy problem" [6]. This phenomenon leads us to believe that the nervous system has somehow evolved to select a set of very specific solutions from an infinite pool of possibilities. These selections are, presumably, optimally advantageous to humans. If they were not, it is reasoned, evolutionary pressures would have arrived at some other solution that *is* optimally advantageous. This redundancy problem, along with its solution, that

---

[1]A more detailed discussion of control signals will follow.
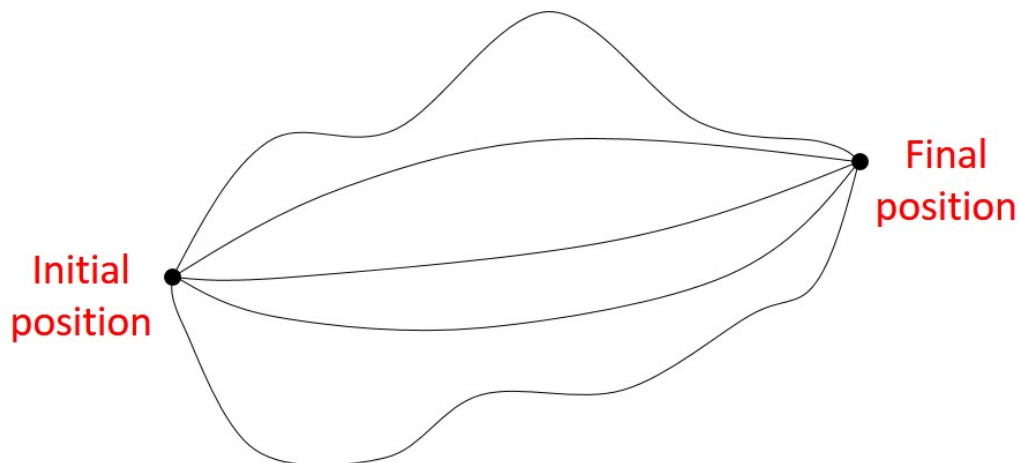
Figure 1.1. A depiction of the so called "redundancy problem" of movement. An infinite number of possible pathways exist between an starting position and desired ending position. Furthermore, an infinite number of velocity, acceleration, etc. profiles exist along the positional path.

evolution has "selected" an optimal trajectory for movements, in part, led to the application of optimal control theory for modeling human movement.

An *optimal control problem* consists of several parts...Informally defined as:

1. *System*: a mathematical model representing some real world system, e.g. the moving parts of a factory, or the dynamics of stock trading. The system dictates how a *control signal* changes system state and the outputs outputs of the system.

2. *Control signal constraints*: a set of permissible *control signals* (inputs) to the system. E.g. the set of possible control signals to a factory arm might be constrained to direct the arm within a limited range of motion.

3. *Output constraints*: a set of permissable outputs from the system. Each acceptable output in the set can specify either an exact value, or a range of values. E.g. fuel consumption for a flight from Las Vegas to San Diego must

be below x gallons of fuel. These can also be thought of as the "goals" of the optimization problem.

4. *Objective function*: a measure of cost/reward performance of the permissible *control signals*, allowing the assessment of the "optimal" control signal.

The *solution* to an optimal control problem is a *controller* that produces a control signal within the control signal constraint set (item 2), that achieves the desired output, specified in the output constraints (item 3), with the lowest cost (or highest reward) as defined by the objective function (item 4).

In the 1970's and 1980's, we began to see the application of optimal control principles to the study of the HMS [7, 8, 9, 10]. Works that followed, which will be discussed, continued to build upon the concept that the HMS conforms to optimal control principles.

The notion of optimality is (perhaps obviously) central to the field of optimal control. In this context, "optimal" means finding the single, best possible solution to a problem, given a definition of its cost (which is to be minimized) and/or reward (which is to be maximized), and its constraint sets. The concept of finding the single, best approach to achieve a goal amongst potentially infinite solutions, at least on the surface, appears to be a perfect resolution to the aforementioned redundancy problem of animal movement. Amongst a sea of possibilities, one trajectory and muscle activation pattern will be selected for some motor task and its constraints. Inherently, this approach maintains the *assumption* that this solution (chosen trajectory and muscle activation) should somehow be optimal for animals in some way, and that animal movement movement has evolved to achieve optimality in some way.

Additionally, keeping in mind the concepts of cost and reward as they

relate to evolutionary pressures, it seems natural to apply optimal control to the animal motor system, which evolved under various evolutionary pressures (costs and rewards) such as minimizing energy consumption, maximizing the attainment of scarce resources (e.g. food), and reproducing maximally. Optimal control has the potential to elucidate biological mysteries because it gives us tools that mimic the process of evolution by finding optimal solutions under the penalty of costs/rewards and failure to meet goals. Several examples of these approaches will be discussed.

In this work, we investigate how current optimal control approaches might be augmented to better reflect the physical implementation of the HMS. We will question whether or not the optimal control approach, in its modern form, is a valid approach for understanding the HMS to begin with. We present evidence that our new formulation of state of the art optimal control techniques may be more accurate and more relevant when modeling human movement.

1.3.2 Open questions regarding the application of optimal control to movement

As encouraging as the premise of using optimal control to investigate the nature of the HMS may be, the state of the art has its failures and shortcomings. The choice of cost functions used to determine an optimal solution is an example of this. Many, seemingly unrelated cost functions can be used to model human movement accurately, a litany of cost functions produce adequate results when modeling human movement [11]. This leads to the question, which cost functions were the true drivers behind the evolutionary process leading to the HMS?

Even if we knew the set of cost functions that molded the HMS via evolution, how do the resultant control signals work together at a systems level? In physiological terms, how do the different neural signals that control our motor

system work together to drive our motor output? The current literature in this field typically utilizes a single control signal, and ignores the fact that many neural control signals of different types work together to achieve motor tasks in the AMS. We address this shortcoming in this work in chapters 4 and 5.1.

At what level of abstraction should we model these systems in order to understand the HMS? Current optimal control related approaches are nearly devoid of any physiological or anatomical explanations of *how* the HMS is implemented. In order to create machines that move with animal-like grace and precision and/or to understand human motor deficits/illness, we need a model that informs us about the construction of the HMS. We also address this shortcoming in this work.

1.3.3 Select neuroscientific discoveries leading to the concept of "controllers" in the animal motor system

Despite hundreds of years of investigating the mechanisms of the AMS, we have only a rudimentary understanding of how its neural mechanisms control the skeletomuscular system to produce effective movement. In the late 1700's, Italian physician Luigi Galvani performed a series of seminal experiments which began to elucidate this problem. He was able to actuate frogs' leg muscles via electrical stimulation. Clearly, this showed that electrical signals were capable of driving muscle activation, at least in a very rudimentary manner. His experiments were at the core of a the new field of "bioelectricity", which studied electrical signals in the nervous system [12]. Still, these experiments did not encapsulate the full meaning of "control" in the modern sense, as there was no desired state for the frog's nervous system to achieve, nor were there constraints. However, it is easy to imagine that, somewhere in the (living) frog, similar signals were emitted which drove the frog's motor system to achieve some goal, within certain physical limitations (which can

be viewed as constraints).

Moving forward to the late 1800's, we began to understand one of the most basic tenets of mammalian motor control. Fritsch and Hitzig showed that electrical stimulation to the motor cortex of dogs could elicit movement [13] (translated in [14]). This movement was not random, instead they showed that stimulations of select areas of the cortex repeatedly activated the same muscle or muscle groups. Furthermore, these stimulations formed a mapping between certain areas of the cortex with certain muscles. In a sense, they had found a *neural controller*, albeit one that was distributed across many areas of the cortex (later to be considered the motor cortex). It is worth reiterating that, like Galvani and his frog experiments, these stimulations that drove motor functions were electrical signals. It appeared that each area of this distributed system sent *control signals* to activate different muscles. Again, we witness multiple signals from multiple sources having the ability to control motor function, as first discussed in section 1.1.1.

1.4 General philosophy and scientific approach that have guided this work

1.4.1 The importance of explanatory power

As mentioned in the preface, whether scientists are aware of their own biases of not, they often have different goals when studying natural phenomena. A scientist may choose to create a model that is predictive in nature. For example, Newton's laws of motion could be utilized to *predict* exactly when an apple will hit the ground when falling from a tree. However, Newton lamented that his laws did not explain *how or why* gravity affects the apple in the way that it does. His "law" of gravity lacked an *explanation* of the observed phenomenon, and while Newton's laws were, and still are incredibly useful, he was left unsatisfied.

Contrast Newton's dissatisfaction regarding his theories with Einstein's formulation of general relativity, which posits that gravity is an effect of curved space-time. Einstein offers a better explanation of why the apple falls from the tree at a given rate. Mass reshapes the fabric of space. As an example, light is not "bent" when it travels passed a massive star, it is simply travelling normally along a curved surface. Planets are not "drawn" to the star they orbit, they are are travelling via a trajectory through space that is warped around the mass of a star. In the same way that a planet "falls" towards the Sun (hopefully in an orbit), an apple is drawn towards the Earth due to warped space. Whether or not this is the correct conception of gravitation, it is much more satisfying than Newton's purely mathematical account as it offers an explanation of *how and why* gravity functions as it does.

An analogous situation exists in the literature of applying of optimal control to the modeling of human movement. As will be demonstrated, the majority of the applications of optimal control to the modeling of human movements are purely mathematical in nature. They make only feeble efforts to explain how these optimality principles are implemented in the nervous system (analogous to the lack of explanation of how gravity is "implemented" in the universe). They lack explanatory power. In this work, we do not claim that we can explain the HMS. However, whenever possible, we attempt to offer some, plausible relationship between the mathematics of optimal control, and the anatomy and physiology of the HMS. In doing so, we hope that the techniques introduced here can be used to model the nervous system more directly, thereby bringing us closer to an understanding and explanation of *how and why* the HMS functions as it does.

1.4.2   Bottom up vs. top down

In order to understand how optimal control purports to give us insight into the HMS, and to understand the value it may provide, in this section we frame the traditional approach from a different perspective, illustrate the mindset that has guided this work, and give intuition regarding the scientific contributions this work makes. Most techniques in engineering and science can be divided into two broad categories. We refer to these as "bottom up" (also called inductive reasoning) and "top down" (also known as deductive reasoning). Understanding the differences between these two approaches will help the reader to appreciate the contributions of this work. We give a loose definition of each approach here as they relate to scientific investigation:

The bottom up approach can be thought of as having five steps, we give an illustrative examples from a fictitious scientific problem regarding the ocean's algae population fluctuations.

1. Start with data collection on low level observations of some phenomenon, assuming nothing regarding the nature of the data: At a certain latitude, take samples of the temperatures of the world's oceans, along with their algae content, over several years.

2. Notice patterns, generalize based on observations, begin to be able to predict future behavior based on past observations: Algae content peaks at during the fall months.

3. Develop one or more hypotheses to explain the observations A) The ocean temperature is optimal at this time for algae growth in certain regions B) The temperature at this time is unfavorable for the main predator of algae, allowing algae to proliferate.

4. Test hypotheses A) Choose a different latitude which reaches this optimal temperature at a different time of year, and see if that corresponds to the algae population peak. B) Find an area of ocean where this predator doesn't exist, and see if the same pattern is observed

5. *Finish* by producing a *theory* based on which hypothesis (or hypotheses) stand(s) up to the testing.

The top down approach can be thought of as having three steps. Again we provide a illustrative examples.

1. *Begin* by proposing a general *theory* of how some process works, which has not been validated by experimentation at this point: Guess that algae population fluctuation is due to carbon emission.

2. Come up with specific disprovable hypotheses of what should happen if the theory were true: There should exist a correlation between carbon levels in the atmosphere from the pre-industrial period until now, and algae populations.

3. Observe and experiment in order to support or disprove the hypotheses and by extension support or disprove the more general theory: Use the sedimentary record to estimate the algae population dating back to pre-industrial times. See what the correlation is between algae population and the known atmospheric carbon level over the same time period.

4. Accept or reject the proposed theory based on the strength of experimental evidence.

Notice how the bottom up approach *finishes* by producing a theory while the top down approach *begins* by *postulating* a theory. Both approaches attempt to

validate a theory that should have a plausible explanation for the phenomenon it addresses. Neuroscientific endeavors often take one or the other approach. Some use the bottom up approach, starting at a low level and working "upwards". This approach may begin as low as the level of atoms, molecules, and physical forces, or arbitrarily higher, e.g. individual neurons, networks of neurons, etc. These approaches may include the modeling of spiking neuronal networks based on known neuroanatomy. This type of endeavour has led to explanations (whether correct or incorrect) of how small but relevant subsections of the nervous system may work [15, 16, 17, 18], but hasn't yet been able to work "up" to explaining the system as a whole.

In contrast, we consider the use of optimal control to be a top-down approach, as it starts by postulating a high level theory and, ideally, attempts to move toward something that may help us explain the inner workings of the motor system at a lower level (e.g. at the level of individual neurons). For example, one might start by proposing a theory that the HMS has evolved to minimize energy consumption while completing motor tasks. One then can come up with disprovable hypotheses and experiment to confirm or refute them, thereby supporting or disproving the theory. However we feel that the current state of the optimal control approach is still at the high level theory and hypothesis testing stage and lacks explanatory power with regard to the neural implementation of the motor system. Others have stated that the next step in the optimal control approach is to extend the reach of the explanations provided closer to the anatomical and physiological level (e.g. [19]). With this directive in mind, in this work, we explore, test, and justify a number of ideas which change the optimal control approach to modeling animal movement in ways that bring us closer to modeling the neural implementation of the motor system.

1.4.3 Questioning fundamental assumptions and practices in the field of optimal control as applied to the exploration of the HMS

Perhaps the most important concept guiding this work is the questioning of the assumptions and standard practices used by researchers in this field. Science typically builds upon itself, a concept wish is succinctly captured in Newton's famous quote, "If I have seen further it is by standing on the shoulders of giants". However, there are times when a few of those scientific building blocks are flawed. Reexamination of core assumptions can affect all science built upon those assumptions.

Research leading up to this work involved a continuous exploratory process in which we questioned core assumptions that scientists typically employ when modeling human movements via optimal control. Altering these assumptions sometimes led nowhere. However, in some cases, experimenting with altered assumptions led to interesting and unexpected results. Chapters 3, 4, and 5.1 all address how questioning two basic assumptions of optimal control can lead to fruitful results, which bring the discipline a step closer to providing explanations that can be reasoned about in terms of anatomy and physiology of the human motor system.

1.5 Overview of select chapters

Chapter 2: A simple control policy for achieving minimum jerk trajectories.

Clearly, evolutionary forces impose some notion of cost or reward on the organisms that survive and reproduce, as opposed to those that become extinct. Organisms incurring high costs but low rewards are unlikely to succeed in an

evolutionary sense. Optimal control directly models cost and reward. The question becomes, how does one measure cost or reward? When finding an optimal solution to a problem with some set of costs/rewards, over some time series, the 2-norm is typically used as the de facto metric (as it is in many error related calculations), and has led to cost functions which have been used to model human movement. However, using different metrics to define cost leads us to very different results. We will show the importance of using a different cost metric, the *infinity-norm*. By comparing one of the canonical models of human reaching movements dubbed "minimum jerk" ("jerk" being the rate of change of acceleration, i.e. the third derivative of position), in which error is measured by the 2-norm (as proposed by Flash and Hogen [20]), against the same cost (minimum jerk) as measured by the infinity-norm. We show that the choice of error measurement has drastic implications with regard to the two models' predictions of the control signals that drive movements.

Chapter 3: Sparse signals for the control of human movements using the infinity norm.

This chapter is an extension of our findings described in chapter 2, "A simple control policy for achieving minimum jerk trajectories". We expand our exploration to optimal control strategies which minimize higher derivatives of human movement position, namely, jerk, snap, and crackle, the third, fourth, and fifth derivatives of position respectively. We again compare the these costs as measured by the 2-norm vs the *infinity-norm*. We argue that in all cases, minimization via the infinity-norm produces a control signal reminiscent of a "bang-bang" control signal, one of the simpler control strategies used to control mechanical systems. We extend the previous work in chapter 2 by outlining a method for encoding this signal

with sparse Dirac delta functions, which more closely resemble both individual neural spikes, along with sudden bursts of neural neural activity from populations of neurons.

Chapter 4: Suboptimal combinations of cost functions fit fast human reaching movements as well as optimal combinations.

This chapter presents a generalized framework called the "Optimal Confluence Control Framework", or OCCF. It defines the OCCF and begins to justify why it can be useful in the modeling of human movements. We then demonstrate an example of how to use the OCCF framework, which is a proof of concept that the framework is effective. To date, when applying optimal control to the modeling of human movements, a single control signal is used. This control signal is based upon a single cost function, sometimes consisting of many cost terms. In contrast, the HMS utilizes scores of signals interacting to produce muscle contractions resulting in effective movement, a concept that the OCCF supports.

We question the use of a single cost function, and a single control signal based on that function. There is a reason that a single, multi-term cost function, which generates a single control signal, is used universally. It is because this formulation is mathematically guaranteed to be "optimal" in terms of costs, the notion of which is the bedrock of optimal control.

Instead, we espouse the use of multiple cost functions, each evaluated separately from the others, which result in multiple control signals. This approach is not optimal in terms of minimizing costs, but can more closely mirror anatomy and physiology of the HMS than the traditional, single function, single control signal method. We offer a simple example of this approach and compare the differences in efficacy between this and traditional approaches when modeling human reaching

movements.

Chapter 5: Conclusions, conjectures, and scientific contributions.

In this chapter, we explore anatomical and physiological reasons why our modification of traditional optimal control techniques holds promise to become a valuable tool when reverse engineering and reasoning about the HMS. We provide an overarching framework that, we argue, should be used in place of the existing paradigm. We offer examples from neuroscience which show that the traditional optimal control approach, keeping in mind its top down strategy, has no hope of explaining the HMS at a level "lower" than it already has...That is, the current paradigm has reached the end of its explanatory power. It is for these reasons we have explored the mechanisms necessary for the new paradigm offered in this work.

Chapter 2

A simple control policy for achieving minimum jerk trajectories

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 paper we propose a bang-bang optimal control policy 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.

2.1 Introduction

The process of evolution drives species to differentiate to produce some competitive advantage. Just as it is easy to observe that evolution has lead to a great many modes of locomotion [21], it can be assumed that evolution has had effects on the control of behavior by the nervous system. Under this assumption, 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.

How the 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 [22, 23, 24]. 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 [22, 25, 23, 24]. For example, Graziano's work [26] 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 [27, 28].

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 [29, 30, 31]. 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 [20] or paths [32]. Furthermore, it has been observed that this class of movements consistently follows bell-shaped velocity profiles [33, 20]. 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. * There have been several attempts to mathematically model human point-to-point movements of an end effector (the hand) [20, 32, 34]. A key study of this type is [20]. 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 [35]. 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 [36]. 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 it's 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.

2.2   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 [37, 38, 39, 40]. 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 2.1, notice that near time t=0 and t=1 second, the jerk resulting from the MJ_2 model far exceeds
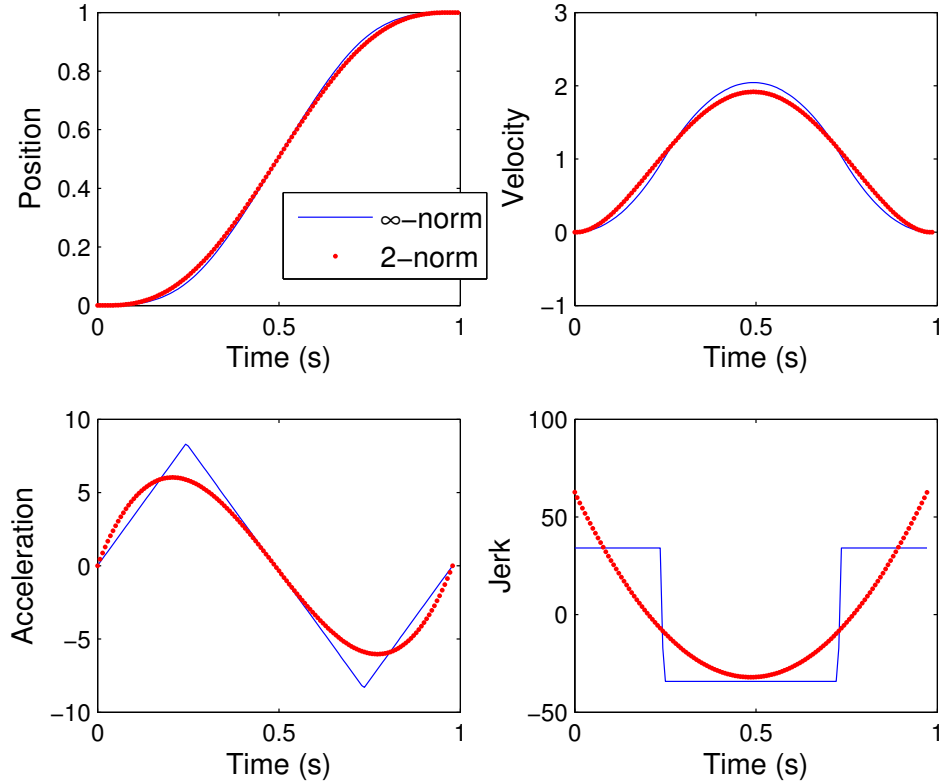
Figure 2.1. Dynamic profiles of reaching movements corresponding to the L_2 norm based control policy and L_∞ norm based control policy are shown in red and blue respectively.

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).

2.2.1   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 seeking a minimum jerk trajectory $x(t)$ is formulated as follows:

$$\begin{aligned}
\underset{u(t)}{\text{minimize}} \quad & \| u(t) \|_p \\
\text{subject to} \quad & \dot{\mathbf{x}}(t) = A\mathbf{x}(t) + Bu(t)
\end{aligned} \tag{2.1}$$

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

The solution to equation 2.1 will determine the optimal control policy $u(t)$. The selection of the L_p norm can result in vastly different control policies. For $1 \le p < 2$, the control policy will result in physiologically unrealistic movements and as a result these 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 1. *The solution to equation 2.1 with $p = 2$ is a straight line trajectory given by the following control policy:*

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

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 [20] $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ $\square$

The control policy corresponding to the above theorem has been shown to fit human data very well [20]. The next theorem shows that if the L_∞ norm of jerk is minimized (i.e, $p = \infty$) in equation 2.1 then we have a bang-bang control policy:

Theorem 2: Bang-bang Control Theorem. *The solution to equation 2.1 with $p = \infty$ is a straight line trajectory given by the following control policy:*

$$u(t) = \dddot{x}(t) = \begin{cases} J & 0 \le t < {}^T/_4 \\ -J & {}^T/_4 \le t < {}^{3T}/_4 \\ J & {}^{3T}/_4 \le t \le T \end{cases} \qquad (2.2)$$

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 [38] $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ $\square$

Ben-Itzhak and Karniel [35] 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.
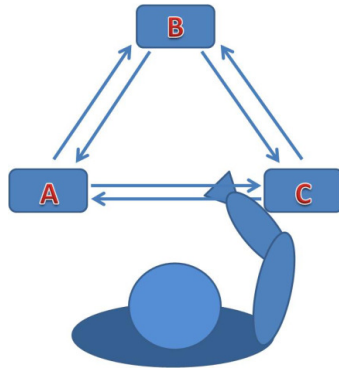
Figure 2.2. The boxes labeled A, B and C are "targets" which appeared on a screen in front of the subjects. The triangle at the end of the subject's arm represents the subject's hand and the manipulandum he is moving. The subject moves his hand between two of the points within a specified time window. The three targets allow for six distinct movement types indicated by the six arrows. The subject receives visual feedback indicating if he or she has successfully reached a target and if he or she has done so within the allotted time window (diagram adapted from [35]).

2.3 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 [35]. The data originated from a 2002 paper by Karniel and Mussa-Ivaldi [41] 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 [41].

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 2.2). The subjects were instructed to move
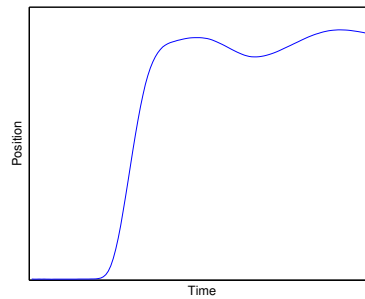
Figure 2.3. A typical recording of the position profile. The oscillatory behavior at the end of the movement corresponds to the overshoot and correctional effects described in the text and are discarded since they are not part of the ballistic portion of the movement. The units of time and position are not relevant and are not shown.

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, ±50ms. 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 4.2 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 [42, 43, 44]. 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:

$$E = \sum_{i=1}^{5} v^2[i] \tag{2.3}$$

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 ($^T/_2$) and end (T) times of a movement in the same way as done in [35]. End effector velocity profiles for ballistic point-to-point movements are known to have a symmetric bell shape [33, 20]. In order to determine the end time of the movement, we first define the middle position of the movement ($^T/_2$) 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:

$$c = \frac{v_s + v_e}{2v_p} \tag{2.4}$$

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 2.4). 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.
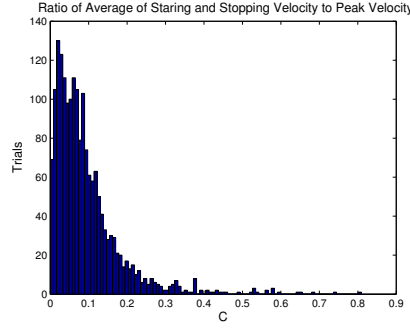
Figure 2.4. Histogram of the ratio of average starting and stopping velocities to peak velocity for each trial. Ideally, the trials should have start and end velocities as close to rest as possible in order to conform to both models' assumptions.

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.

2.4   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 subject's 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:

$$MSE_k = \frac{1}{t_k} \sum_{i=1}^{t_k} (x_k[i] - p_k[i])^2 \tag{2.5}$$

where k refers to the trial, t_k the number of samples for the kth trial, x_k is the model position profile for the kth trial, and p_k is the recorded data position data
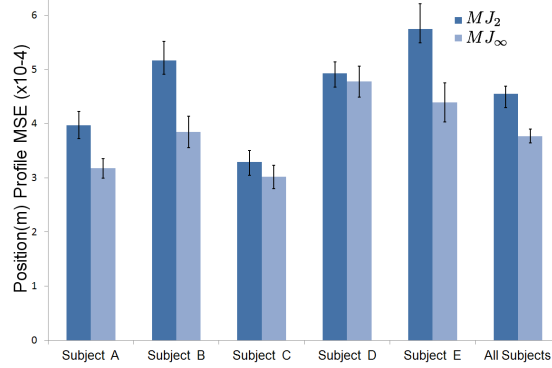
Figure 2.5. Bars A - E show the comparison of the average MSE of position profiles predicted by MJ_2 and MJ_∞ vs. human trial data along with standard error bars. Error units are $m \times 10^{-4}$. "All Subjects" is an aggregate of the trials in A-E. The MJ_∞ model performs better than the MJ_2 model 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 [35].

for the kth trial.

Figure 2.5 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 by Ben-Itzhak and Karniel [35] 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
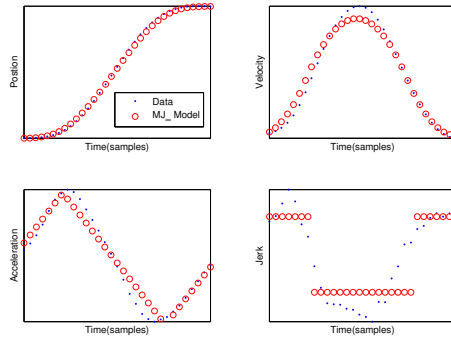
Figure 2.6. Overlay of MJ_∞ model with human movement data of a single trial.

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 minimizing 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 2.6).

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.

2.5   Discussion and future work

In this work, 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 [45] 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 [20]). 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 [46, 47]. Other models have shown that central pattern generators can be driven via step inputs [48].

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 [28] 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 [49] 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 exists in supraspinal centers [26, 46], within the spinal cord [27, 28], or both in the brain stem and spinal cord [49]. Furthermore, it is plausible that populations of bursting neurons could implement a bang-bang control signal [50]. 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 [51]. Integration of a neural signal encoding velocity would lead to a signal encoding position as happens in the occular-motor system [52, 53]. 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 [32, 54]. 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.

2.6   Acknowledgements

This chapter, in full, is a modified reprint of the material as it appears in

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

Chapter 3

Sparse signals for the control of human movements using the infinity norm

This chapter is an extension of the previous chapter, building upon its conclusions and extending its explanatory power. We expand the argument that minimizing the L_∞ norm of the effort term when optimizing a control problem results in control signals that have a physiological interpretation.

Optimal control models have been successful tools in describing many aspects of human movements. While such models have a sound theoretical foundation, the interpretation of such models with regard to the neuronal implementation of the human motor system, or how robotic systems might be implemented to mimic human movement, is not clear. One of the most important aspects of optimal control policies is the notion of *cost*...We offer a mathematical method to transform the current methodologies found in the literature from their traditional form by changing the method by which cost is assessed. In doing so we show how sparsity can be introduced into current approaches that currently use continuous control signals. We assess cost using the L_∞ norm. This influences the optimization process to produce optimal control signals which can be represented by a small amount of Dirac

delta functions, instead of continuous control signals. In recent years sparsity has played an important role in theoretical neuroscience for information processing (such as vision). Typically, sparsity is imposed by introducing a cardinality constraint or penalty (measured or approximated by the one-norm). In this work, however, to obtain sparse control signals, the L_∞ norm is used as a penalty on the control signal, which is then encoded with Dirac delta functions. We show that, for a basic physical system, a point mass can be moved between two points in a way that resembles human fast reaching movements. Despite the sparse nature of the control signal, the movements that result are continuous and smooth. These control signals are simpler than their non-sparse counterparts, yet yield comparable if not better results when applied towards modeling human fast reaching movements. In addition, such sparse control signals, consisting of Dirac delta functions have a neuronal interpretation as a sequence of spikes, giving this approach a biological interpretation. Actual neuronal implementations are clearly more complex, as they may consist of large numbers of neurons. However, this work shows, in principle, that sparsely encoded control signals are a plausible implementation for the control of fast reaching movements. The presented method could easily be scaled up to arbitrarily large numbers of duplicates, thus representing larger numbers of spikes. We show how leading techniques for modeling human movements can easily be adjusted in order to introduce sparsity, and thus the biological interpretation and the simplified information content of the control signal.

3.1   Introduction

Optimal control theory has provided a great deal of insight with regard to developing mathematical models that describe human movements (for example, [20, 55, 56, 57, 58]). These works, amongst many others, have shown that humans

move with strategies that can be described/driven by various control signals[1] and related cost functions to model movement. However, as [19] points out, while the development of optimal control models has given mathematical insights into the properties of human movements, and perhaps the costs that forged our motor system via evolution, the connection to the neuronal implementation of the motor system is not clear. In contrast to these models, we show that a novel penalty on a control signal results in signals which can be represented more simply, and that have more plausible biological interpretations, while maintaining the ability to model human movements accurately.

To demonstrate the utility of sparse optimal control signals for human movements, we will compare two versions of a class of problems called "minimum effort" control problems, which attempt to minimize the "size" or "effort" of the control signal when modeling human movements, as explained in [59]. In the first, and more traditional version, the effort to be minimized is defined as the L_2 norm of the control signal over some time course within which a movement is completed. The goal is to minimize that signal. One of the first, and most famous of this family of models is the "minimum jerk" control policy, proposed originally by Flash and Hogan [20]. The Flash and Hogan control strategy models human reaching movements, and uses jerk (the third derivative of position) as a control signal, and minimizes that signal to the extent allowed by a well defined reaching task. Intuitively, most are more familiar with thinking of acceleration as a control signal (e.g. pressing the accelerator in a car controls the speed). In the Flash and Hogan case, jerk is the derivative of acceleration, so it "controls" the acceleration in the same way acceleration "controls" velocity. Because jerk is the minimized control

---

[1]Control signals are to be described in more depth later, but for now it is a signal that controls elements of a system.

signal, we refer to jerk as the "effort term". Because they attempt to minimize their control signal, this as a minimum effort problem, and is referred to as "minimum jerk".

Work since Flash and Hogan has considered different cost functions as the effort term as defined above, such as minimizing torque over the course of a movement [60], or minimizing torque change (derivative of torque) over the movement [55, 61]. Other works have added more terms to the a cost function but have maintained an effort term. Additions to the effort term include end-point stability (how much adjustment is needed once the target area of a reaching movement is breached) or end-point accuracy (how close to the target when the reaching movement ends) [58]. Such extensions, however, have not lead to any insights into the neuronal implementation of control signals in the CNS, nor do they simplify the nature of the control signal.

We will show that using the L_∞ norm instead of the L_2 norm for measuring and penalizing the "effort" of the control signal results in signals that can be encoded sparsely via Dirac delta functions. There exists a family of models where this technique is applicable, specifically, because they all employ an "effort" term. Practically, the sparsification of signals generated by this family of optimal control models might be useful in a robotic system in order to achieve human like-movement. Due to the simplicity (sparsity) of the resulting signal, implementations of human-like robotic control may be easier to comprehend and construct.

Models of reaching movements leading to this work were considered in [35, 62, 63]. These works referred to their signals as "bang-bang" or "intermittent" (see [64, 65] for more on intermittent control). However, the control signals were not sparse, rather, they were *square wave continuous*. Other types of motor control such as standing and keeping balance have been modeled via intermittent

control, notably, two such models are compared in [66]. Here, we demonstrate a mathematical relationship that can convert non-sparsity to sparsity with regard to the control signal. This relationship is related to the metric used to measure the control signal (i.e. how is the effort measured?), but more importantly a sparse encoding of the signal via Dirac delta functions. We also show that sparse optimal control signals model real human arm movements with high accuracy, thus supporting sparse optimal control signals as a plausible control strategy used by a human's biological system. We emphasize that, using sparse signals in the cases shown here has no downside in terms of model performance, but has the benefits of a simpler encoding of the control signal, a biological interpretation in terms of neural spike timing, and potentially, a simpler control strategy which may be useful in robotics...All of these qualities are absent from non-sparse signals.

3.1.1   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 [36, 19]. Optimal control theory is an application of optimization theory to the control of a dynamic system. In optimization theory, we seek to find an element in a domain that minimizes (or maximizes) a criterion (also referred to as an objective), while satisfying a constraint set. When the elements in the intersection of the domain and constraint sets 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:

$$\dot{\mathbf{x}}(t) = \mathbf{a}(\mathbf{x}(t), \mathbf{u}(t), t). \tag{3.1}$$

$\mathbf{x}(t)$ 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 time-invariant (LTI) and can be expressed as

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t). \tag{3.2}$$

Given the dynamic system of equation 3.1 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 theory, we seek an optimal control signal $\mathbf{u}^*(t)$ that, in addition to transferring the system to a desired state, also minimizes a cost functional $\mathbf{J}(\mathbf{u}(t))$. The cost functional is application dependent and the optimal solution $\mathbf{u}^*(t)$ depends on what we consider to be "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.

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 possibly. Thus, the cost functional penalizes the total time it takes to transfer the initial state to a final state and can be expressed as

$$\mathbf{J}(\mathbf{u}(t)) = t_f - t_0 \tag{3.3}$$

where the initial state $\mathbf{x}(t_0)$, initial time t_0, and final state $\mathbf{x}(t_f)$ are known, while t_f is unknown, and the system dynamics are described by equation 3.1. We furthermore constrain the control signals to be bounded

$$|\mathbf{u}(t)| \leq B. \tag{3.4}$$

We now show the solution to the minimum-time control problem for an LTI system as described by equation 3.2. 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

$$
\begin{aligned}
\underset{\mathbf{u}(t)}{\text{minimize}} \quad & T \\
\text{subject to} \quad & \dot{\mathbf{x}}_n(t) = \mathbf{A}\mathbf{x}_n(t) + \mathbf{B}\mathbf{u}(t) \\
& \mathbf{x}_n(0) = \mathbf{x}_i \\
& \mathbf{x}_n(T) = \mathbf{x}_f \\
& |\mathbf{u}(t)| \leq B
\end{aligned}
\tag{3.5}
$$

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. Their conclusions lead to several important

points upon which this work builds, as they guarantee a control signal which switches a finite number of times between two possible values. These points are summarized below (see [67] for more details).

P.1 For a given LTI system, there is one, and only one optimal signal to drive the system from an initial state to a desired state.

P.2 Because the goal of a minimum-time problem is to move the system to the desired state in the least amount of time, a control signal representing a dynamic variable (such as acceleration) is always at one of two extremes, $+B$ or $-B$. These extremes are defined in the constraints of equation 3.5. Intuitively, if you want to get from point X to point Y as fast as possible, you would change from zero acceleration to maximum positive acceleration to speed up initially, and then to maximum deceleration to slow down to reach point Y, and then to zero acceleration to maintain your starting position. The two extreme values are the only values that yield to a minimum time (optimal) result.

P.3 The control signal will switch between these two extremes at most $n + 1$ times, where n is the derivative of position we choose to be the control signal. For example, for velocity, acceleration, and jerk, n is 1, 2 and 3 respectively, and thus has a maximum of 2, 3, and 4 (respectively) switches between the extremes of the control signal.

This type of control signal is sometimes referred to as a "bang-bang" control signal since the signal switches between the lower bound and upper bound of the inequality constraint of equation 3.4 (see [36] for more). It can also 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 bounded

number of switches between the lower and upper bound. These switches can be encoded by a series of Dirac delta functions, which resemble neural bursts or spikes.

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

$$\mathbf{J}(\mathbf{u}(t)) = \left(\int u(t)^p dt \right)^{1/p} \tag{3.6}$$

which denotes the L_p norm (and is typically the L_2 norm), and the system dynamics are described by equation 3.1. 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 3.3 and 3.6

$$\mathbf{J}(\mathbf{u}(t)) = \gamma \left(\int u(t)^p dt \right)^{1/p} + t_f - t_0 \tag{3.7}$$

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

As an example of a minimum-effort problem, Flash and Hogan considered the following minimum-effort control problem which introduced constraints on

initial and final state in order to describe human movements:

$$\underset{u(t)}{\text{minimize}} \quad \left(\int u(t)^2 dt \right)^{1/2}$$

$$\text{subject to} \quad \dot{\mathbf{x}}(t) = A\mathbf{x}(t) + Bu(t)$$

$$\mathbf{x}(0) = \mathbf{x}_i$$

$$\mathbf{x}(T) = \mathbf{x}_f$$

(3.8)

where $\mathbf{x}(t) = \begin{bmatrix} x(t) & \dot{x}(t) & \ddot{x}(t) \end{bmatrix}^\top$ 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 $(u(t) = \dddot{x}(t))$, and furthermore used a third-order integrator model for the linear time-invariant dynamic equation parameters:

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

(3.9)

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 these extensions. For example, [57] 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 3.8, which specify an exact final state.

3.1.2   Sparse optimal control policies for straight point-to-point trajectories

In a previous work, we show how square wave control signals, with abrupt switches between two states (first alluded to in P.2, with an example given in figure 3.1), can effectively model smooth human reaching movements [62]. This work extends that notion by developing a method to represent neural signals via sparse usage of the Dirac delta function. These sparse signals can be thought of as encoding a series of positive (excitatory) or negative (inhibitory) neural spikes, or, more plausibly, groups of neurons spiking for brief periods. In summary, we take the square wave control signals described in [62], and encode them as a sparse series of Dirac delta functions, each of which signifies one of the abrupt switching points for the control signal.

We define a *sparse optimal control policy* as a control policy that meets optimality constraints with the lowest cost, as defined by the chosen cost function, that can be encoded by a finite number of discontinuous changes in the signal. An example of a signal that can be encoded as a sparse series is a rectangular pulse function, much like the control signals explored in [62], and shown in figure 3.1. Functions like these can be encoded by impulse functions (see Figure 3.2). The control signals in the minimum-time control problem discussed on in section 3.1.1 is an example of sparse control signals that are optimal in terms of state transfer time. *Henseforth, we will refer to the control signal (square wave) as "sparse", as it is easily transformed into a sparse signal consisting of Dirac delta functions.*

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 $x(t)$,
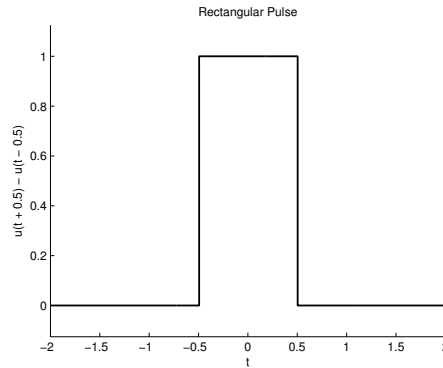
Figure 3.1. An example of the types of sparse signals generated by minimizing the control signal as measured by the L_∞ norm.
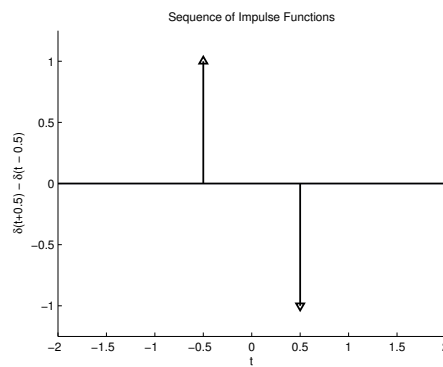
Figure 3.2. Rate of change of sparse signal shown in Figure 3.1. These impulse functions are all that are needed to characterize the control signal in Figure 3.1. They resemble neuronal spike trains, giving them a plausible and mapping to neural systems and making their neural coding predictions testable.

$$u_n(t) = \frac{d^n}{dt^n}x(t). \tag{3.10}$$

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

$$\begin{aligned}
\underset{u_n(t)}{\text{minimize}} \quad & \sup_{0 \le t \le T} |u(t)| \\
\text{subject to} \quad & \dot{\mathbf{x}}_n(t) = \mathbf{A}_n\mathbf{x}_n(t) + \mathbf{B}_n u_n(t) \\
& \mathbf{x}_n(0) = \mathbf{x}_i \\
& \mathbf{x}_n(T) = \mathbf{x}_f
\end{aligned} \tag{3.11}$$

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, T is the duration of the movement (with movement starting at time $t = 0$), and $\sup_{0 \le t \le T} |.|$ is the L_∞ norm. Here we consider a system that is an n-th order integrator, thus

$$\mathbf{A}_n = \begin{bmatrix} \mathbf{0}_{(n-1)\times 1} & \mathbf{I}_{(n-1)\times(n-1)} \\ 0 & \mathbf{0}_{1\times(n-1)} \end{bmatrix} \text{ and } \mathbf{B}_n = \begin{bmatrix} \mathbf{0}_{(n-1)\times 1} \\ 1 \end{bmatrix}. \tag{3.12}$$

The authors of [62] considered the special case $u_3(t) = \dddot{x}(t)$, $(n = 3)$, and showed that this particular sparse control signal explains the trajectories of human movements better than the traditional Flash and Hogan model of equation 3.8. Perhaps more importantly, sparse control signals are biologically more realistic than non-sparse signals (See the Discussion section for more on this). As we will demonstrate, most formulations of minimum effort problems can be easily converted to generate sparse control signals, complete with the afore mentioned benefits. By simply measuring the effort term in the objective function via the L_∞ norm (as opposed

to the 2-norm), we can frame the control signal solution to these problems in terms of discontinuous switching states.

We now show the general analytic solution for equation 3.11. To derive the general solution, we assume that the boundary conditions are

$$\mathbf{x}_i = \begin{bmatrix} x_i \\ \mathbf{0}_{(n-1)\times 1} \end{bmatrix} \text{ and } \mathbf{x}_f = \begin{bmatrix} x_f \\ \mathbf{0}_{(n-1)\times 1} \end{bmatrix}. \tag{3.13}$$

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 3.11 to a form that has been previously solved. Namely, note that every optimization problem can be written equivalently as an optimization problem with a linear objective by introducing an auxiliary variable K and we can equivalently express equation 3.11 as follows:

$$
\begin{aligned}
\underset{u_n(t), K}{\text{minimize}} \quad & K \\
\text{subject to} \quad & \dot{\mathbf{x}}_n(t) = \mathbf{A}_n \mathbf{x}_n(t) + \mathbf{B}_n u_n(t) \\
& \mathbf{x}_n(0) = \mathbf{x}_i \\
& \mathbf{x}_n(T) = \mathbf{x}_f \\
& |u_n(t)| \le K.
\end{aligned}
\tag{3.14}
$$

where $u_n(t)$, $\mathbf{A}_n$, $\mathbf{B}_n$, $\mathbf{x}_i$, and $\mathbf{x}_f$ are defined as before in equations 3.10, 3.12, and 3.13 respectively, and we have used the fact that $||u_n(t)||_\infty \le K \implies |u_n(t)| \le K$. The equivalency between equations 3.11 and 3.14 is due to the fact that every objective can be bounded, and this bound is expressed as an additional constraint in equation 3.14.

The optimization problem of equation 3.14 has the same form as equation

3.5. 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 3.5 the unknown is time T, whereas in equation 3.14 the unknown is the bound K on the control signal $u_n(t)$. Since the dynamic system in equation 3.14 is an n-th order integrator, we can use the result from [68] 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 3.5 or 3.14 where the system dynamic equations are an n-th order integrator (as in equation 3.12), then the number of switchings in the control signal is exactly $n + 1$ and the control signal is symmetric.*

In other words, as initially discussed in P.3, as the order of the control signal increases (as n increases), the number of switches in the control signal increases by the same amount. [69] solved the general n-th order minimum-time control problem of equation 3.5 for an n-th order integrator. We adapt their results for the general n-th order minimum-effort control problem of equation 3.14 and summarize the solution as follows:

$$K_n^* = \frac{2^{2(n-1)}(n-1)!(x_f - x_i)}{T^n} \tag{3.15}$$

$$t_i^* = T\sin^2\left(\frac{\pi i}{2n}\right), \quad i = 0, \ldots, n \tag{3.16}$$

where K^* denotes the optimal amplitude or bound on the control signal, and t_i^* denotes the optimal switching times. Figures 3.3, 3.4, 3.5, and 3.6 show examples
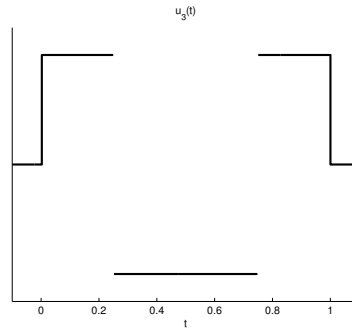
Figure 3.3. Control signal for $u_i(t)$ where $i = 3$, and the movement starts at $t = 0$ and ends at $t = 1$.
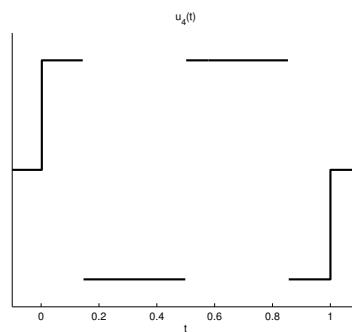
Figure 3.4. Control signal for $u_i(t)$ where $i = 4$, and the movement starts at $t = 0$ and ends at $t = 1$.

of several optimal control signals that can be encoded sparsely.

3.1.3   Sparse Optimal Control Signals in Fast Human Movements

The sparse optimal control signals introduced in section 3.1.2 are not only optimal with respect to a minimum-effort objective, but are also more biologically plausible when compared with non-sparse signals. Sparse optimal control signals can be are those that can be efficiently represented with Dirac delta functions, which resemble neuronal bursts or spikes. If we treat each spike as an idealized
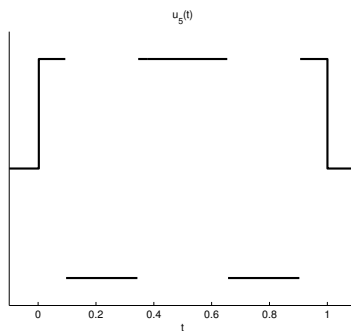
Figure 3.5. Control signal for $u_i(t)$ where $i = 5$, and the movement starts at $t = 0$ and ends at $t = 1$.
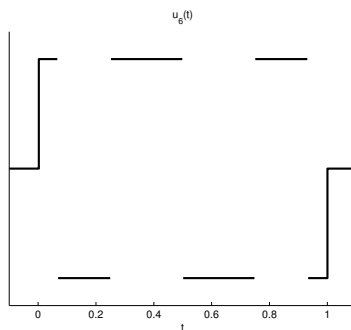
Figure 3.6. Control signal for $u_i(t)$ where $i = 6$, and the movement starts at $t = 0$ and ends at $t = 1$.

Dirac delta function, as done in [70], and as visualized in Figure 3.2, then a spike sequence that represents an n-th order optimal control signal can be expressed as

$$\rho_n(t) = K_n^* \sum_{i=0}^{n} (-1)^i \delta(t - t_i^*) \tag{3.17}$$

where K_n^* and t_i^* can be found from equations 3.15 and 3.16 respectively. The spike train represented by equation 3.17 is not postulated to be from a single neuron, but rather a population of excitatory and inhibitory neurons forming a network.

The work of [35, 62] showed that a sparse optimal control signal that corresponds to jerk (expressed as $\rho_3(t)$ in the notation of equation 3.17) can model fast human movements with greater accuracy than the smooth control signal that results from using the L_2 norm. We now also propose that the control signals the nervous system uses are not limited to the jerk control signal. There is nothing preventing the nervous system from using a higher-order control signal (see Figure 3.9 for comparisons higher-order derivative control signals). With each increase in the order of the control signal, the number of spikes increases and the timing of those spikes changes (as shown in Figures 3.3, 3.4, 3.5, and 3.6). Because each wave form is different, these high-order derivative control signals can form a basis set, the elements of which are combined to form a subspace of control signals. The neuroscience motor control literature commonly refers to the elements in such basis sets as "motor primitives". These primitives are combined to control a variety of animal movements[71].

3.1.4   Application of Sparsity to Extensions of Minimum Effort Control

The minimum effort control problem identified in section 3.1.1 can be extended in numerous ways, e.g. [55, 72, 62, 35, 73, 20, 58]. These works and

others, account for various aspects of movement and draw different conclusions regarding the nature of the motor system. For example, in [58], several types of reaching movements under various conditions are analyzed and modeled via an extension of the minimum effort problem. The types of movements included both two and three dimensional reaching, with and without target perturbation, with and without obstacle avoidance, and under various instructions to the subject regarding how the target should be impacted. The following is a simplified version of the model in [58] that maintains its core concepts: a term for effort and a term for final state error. Equations 3.18 and 3.19 give an example of how minimum effort control problems can easily be adapted to our method of generating sparse signals.

$$
\begin{aligned}
&\underset{u(t)}{\text{minimize}} \quad ||\mathbf{x}(T) - \mathbf{x}_f||_2^2 + w_{\text{effort}} \int_0^T u(t)^2 dt \\
&\text{subject to} \quad \dot{\mathbf{x}}(t) = A\mathbf{x}(t) + Bu(t) \\
&\qquad\qquad \mathbf{x}(0) = \mathbf{x}_i
\end{aligned}
\tag{3.18}
$$

Equation 3.18 is the same as the minimum effort control problem discussed earlier, with the exception that hard equality constraints (the end-point boundary conditions) are now "soft" constraints and are penalized as a cost. The w_{effort} term dictates a trade off between minimizing effort and meeting the final boundary conditions. To have a sparse implementation of the above, we use the infinity-norm (sup) as before:

$$
\begin{aligned}
&\underset{u(t)}{\text{minimize}} \quad ||\mathbf{x}(T) - \mathbf{x}_f||_2^2 + w_{\text{effort}} \sup_{0 \leq t \leq T} |u(t)| \\
&\text{subject to} \quad \dot{\mathbf{x}}(t) = A\mathbf{x}(t) + Bu(t) \\
&\qquad\qquad \mathbf{x}(0) = \mathbf{x}_i
\end{aligned}
\tag{3.19}
$$

This optimization problem, similar to the sparse minimum effort problem above in equation 3.11 can be written as follows:

$$
\begin{aligned}
\underset{u(t),K,K_1,K_2}{\text{minimize}} \quad & K \\
\text{subject to} \quad & \dot{\mathbf{x}}(t) = A\mathbf{x}(t) + Bu(t) \\
& \mathbf{x}(0) = \mathbf{x}_i \\
& K = K_1 + K_2 \\
& ||\mathbf{x}(T) - \mathbf{x}_f||_2^2 \leq K_1 \\
& |u(t)| \leq K_2/w_{\text{effort}} \\
& K_1 \geq 0, K_2 \geq 0
\end{aligned}
\tag{3.20}
$$

As before, we are using the property that every optimization function can be written equivalently as an optimization problem with a linear objective by introducing an auxiliary variable. Again, we have used the property that $w_{\text{effort}} \sup_{0 \leq t \leq T} |u(t)| \leq K_2 \Rightarrow w_{\text{effort}} |u(t)| \leq K_2$. Therefore, we also have a "bang-bang" solution, since the control signal is hard bounded.

3.2 Materials and methods

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 ability 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 [41] for a more complete description.

Five subjects participated in an experiment involving a manipulandum that restricted their movements to a the horizontal plane in front of the seated
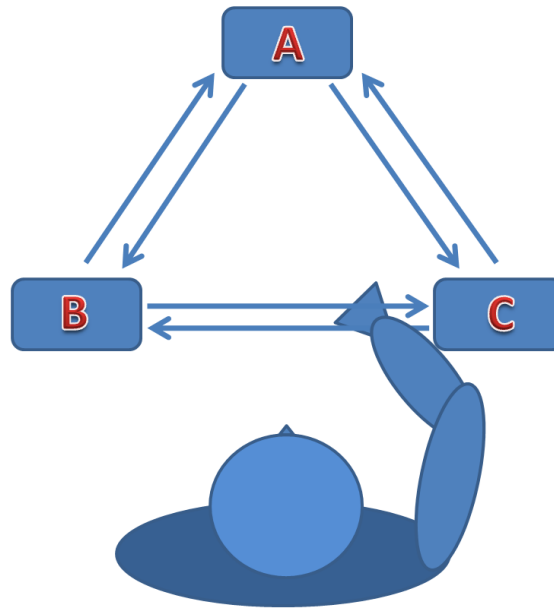
Figure 3.7. Experimental setup for collection of fast reaching movement data. Subjects sat down and held a manipulandum with their hands which they could maneuver about a 2D plane positioned in front of them perpendicular to their torsos. Regions 'A', 'B' and 'C' (Figure adapted from [62] with permission.)

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 (see Figure 3.7). 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 ±50ms. 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. In all, there are 366 trials for the 5 subjects.
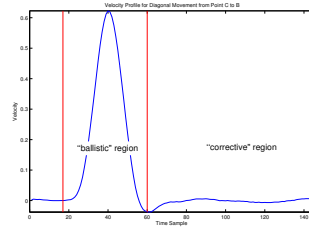
Figure 3.8. A velocity profile for a typical hand movement trial. We attempt to identify the "ballistic" (fast) portion of the movement by using an onset and offset detection algorithm to automatically detect the beginning and end of the fast reaching portion of the movement. The "onset" is indicated by the leftmost red line and the "offset" is indicated by the rightmost red line. The "corrective region" refers to the time during which the subject attempts to correct any over or undershoot of the target, this region is not modeled. Data between the red lines is modeled, the rest is discarded.

For each trial, we only select the so-called "ballistic" portion of the movement (see Figure 3.8). That is, we select the portion of the trial where movement had started and the movement had completed its feedforward portion, the portion of the movement that was preplanned and not affected by peripheral feedback. Part of the justification for this understanding of fast reaching movements is that they are happening too quickly for a corrective proprioceptive signal to make a meaningful difference, as discussed in [74, 75]. Thus we do not model the "corrective" portion of the movement that likely involves additional feedback information from the subjects' visual system and limbs in order to fix any error when attempting to reach to the targets. For this reason, we frame our approach as a feed forward control problem, because feedback is not involved in the movements. This is the same approach used in several studies including [62, 76, 35].

There are various methods for movement onset and offset detection [42, 43, 44], 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 profile has reached it's peak velocity. Fast movements 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 find the sample that falls below a pre-determined threshold. This methods extracts the ballistic region of the movement.

The optimization procedures were implemented using CVX and Matlab. CVX is package for specifying and solving convex optimization problems [77].

3.3 Comparison of sparse and non-sparse model predictions for human reaching movements

Figure 3.9 shows the average mean squared error (MSE) between the human subjects' trial velocity profiles and the velocity profiles generated by four computational models. We computed the MSE between the models' velocity profiles and those of the human subjects across time steps of the recorded movement. For all models, portions of the velocity profile which were forced to be accurate due to setting boundary conditions were not included in the MSE calculation. The figure highlights the canonical minimum jerk (as measured by the L_2 norm) model which results in a continuous, non-sparse control signal. We compare this model with one that minimizes jerk as measured by the L_∞ norm. In all cases the sparse signal generated by this model has a lower error than the minimum jerk L_2 norm model. This is not to say that the L_∞ norm model is decidedly better in all cases, but it is in most, and at the very least it performs similarly to the L_2 norm model while
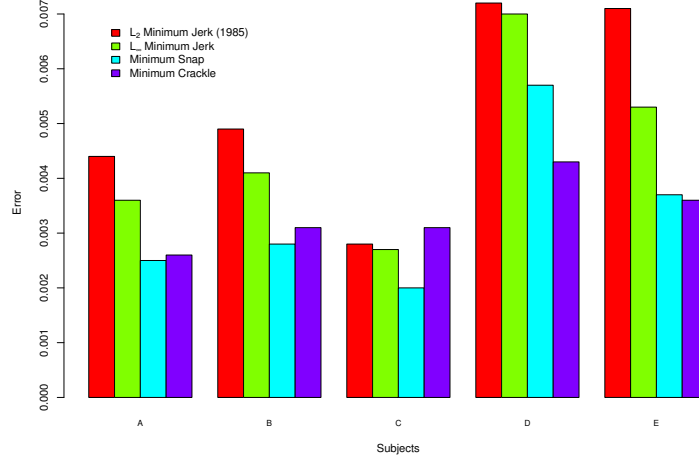
Figure 3.9. Comparison of average mean squared error between the velocity profiles of reaching trials for all five subjects and predictions made on those trials by four models. Error units are in m/s. Red bars indicate MSE for the L_2 norm based minimum jerk model (*continuous*). MSEs shown in green, cyan, purple are all based on the L_∞ norm, and thus result from *sparse* control signals. They represent minimum jerk (*sparse*), minimum snap, and minimum crackle, respectively. The sparse control signals on average have smaller error than the continuous control signal (resulting from use of the L_2 norm). In all comparisons but one, the sparse control signals result in lower MSE than the continuous control signal. The one exception is in the case of subject C where we see the the minimum crackle sparse control signal produce a greater error that the L_2 norm minimum jerk (continuous) control signal, however, the difference between the errors is not statistically significant (as determined by a Wilcoxon rank-sum test).

retaining the characteristic of an intuitive mapping to a spike train representation.

We also contrast the results of the L_2 norm model with three additional models that minimize snap, crackle, and pop (fourth, fifth, and sixth derivatives of position) as measured by the L_∞ norm (resulting in sparse control signals). Errors of all models employing sparse control signals are smaller than those generated by the minimum jerk model employing a non-sparse control signal in all cases but one. In this case, for subject C, the sparse crackle (L_∞) based model has a higher error than the non-sparse jerk (L_2 norm) based model, however, the difference between

the non-sparse minimum jerk error and sparse crackle error is not statistically significant (by a Wilcoxon rank-sum test). This demonstrates that regardless of the chosen derivative of position used as the control signal (jerk, snap, or crackle), sparse signals are effective control strategies.

3.4   Discussion

3.4.1   Sparse Signals and Their Biological Plausibility

Converting traditional minimum effort models to their sparse counterparts via methods outlined in this work can bring these models closer to a plausible biological interpretation in several ways. At the level of the observation of human movement, studies have indicated that human subjects use sparse (intermittent) control strategies for ballistic movements to control activities that are continuous in nature [78, 35].

At the neural system level, there is evidence that various neural structures exhibit intermittent behavior. For example, the basal ganglia have been shown to be key components in the control of movement. Inputs to the basal ganglia, arriving from a large portion of the cerebral cortex, exhibit intermittent behavior [79]. More generally, the basal ganglia as a whole are thought by most to be a network that switches between well defined states, in an intermittent fashion [79]. Another example that is directly applicable to this work examines a primate tasked with making reaching movements towards two *possible* targets that the animal is accustomed to. For a given trial, the "correct" target is not initially known to the animal. At this point, two sustained signals (sustained neural activity), representing each potential target, are present in its pre-motor cortex. Once the "correct" target (for a given trial) is revealed, an abrupt switch occurs where the

neural signal representing the "incorrect target" is suppressed and the "correct" neural signal remains [80]. This is is clearly an abrupt switch between two states, as suggested by the sparse model of reaching tasks outlined in this work.

At the neuronal level, it has been shown in [81] that Purkinje cells in the cerebellum (well known to be involved with motor control) exhibit bistability. That is, they have two modes of operation, each of which persist until a switching event occurs. This event consists of *brief neural pulses*, which switch the Purkinje cell back and forth from a highly active state to an inactive state. The modeled spike trains suggested in this work (e.g. Figure 3.2) may be interpreted as single neurons or they may be interpreted as ensembles of neurons working in concert. In either case, the the pulses controlling the bistable state of the Purkinje cell can be represented by Dirac delta functions, which is an accepted technique to mathematically represent spike trains [70]. In this way, these signals more closely mirror the physiology of neurons when compared to their non-sparse counterparts (as shown in Figure 3.10). Conceptually, there is a mapping from sparse switches (Dirac deltas) to neuronal spikes, or groups of neurons spiking. Furthermore, spike timing, and its relevance in the neural coding of control information is directly represented, which is not the case with the continuous control signal method. In addition, the sparsity of these signals simplify the necessary output of a neural circuit used to drive motor function. For example, it would require only four spikes to encode the signal shown in Figure 3.3. This concept lends itself to a hierarchical control structure employing "higher level" neural motor control structures that focus on learning and producing simple switching times, which drive and offload more complicated tasks and signal processing to lower level structures (as discussed in [72]) which exist in the brain stem and spinal cord. To reiterate, we are not stating that the models in this work imply that a single neuron is driving any kind
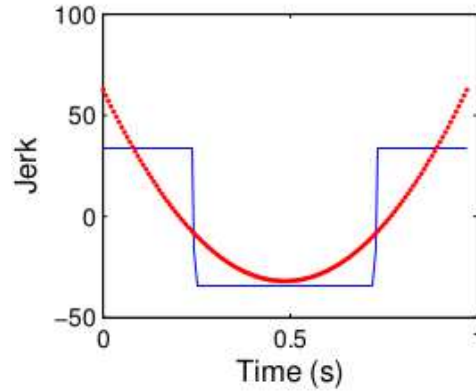
Figure 3.10. The red plot shows a minimized jerk control signal as measured by the L_2 norm. Note that it is parabolic and continuous (non-sparse). The blue plot shows a minimized jerk control signal as measured by the L_∞ norm. Note the distinct qualitative difference between the two. The L_∞ based signal is encoded by a finite number of abrupt changes. Such a signal can be characterized by Dirac delta functions as show in Figures 3.1 and 3.2. These spikes resemble neuronal spikes. In contrast, the L_2 norm based signal and has no obvious mapping to the physiology of neurons. (figure originally used in [62])

of motor function, or that these signals *directly* drive muscles. Rather, that, the concepts here are an abstraction which indicates that at some level of a hierarchy of neural motor control hierarchy, this sparse control approach is plausible.

It is clear that the timing of spike trains plays an important role in their meaning and information content [82, 83]. Unlike continuous control signals, the sparse signals discussed here have very specific switching times, each switching time being integral to the character of the resulting modeled movement. Because of this, the sparse approach can model spike timing as it pertains to neural coding of information, while the continuous models have no explicit representation of spike timing or abrupt system level switches.

It has been proposed that signal dependent noise contributes to the variability of observed movements [57]. Both [84, 85] provide empirical support for this proposal. This viewpoint correlates neural signal magnitude with the level of noise in the

system, and therefore the accuracy of the movement. In other words, signals that are extremely strong (i.e. all available neurons for a given task are firing), inherently create more noise in their own system, lessening the accuracy of the movement. Using the L_∞ norm to generate sparse control signals has the additional benefit of setting a upper bound on the magnitude of the control signal, thus upper bounding the signal-dependent noise, and increasing accuracy. Thinking of the two ideas (upper bounding a control signal, and signal dependent noise) in this way may reconcile any perceived discrepancy between them. Figure 3.10 shows the continuous signal exceeding the absolute magnitude of the sparse signal in two locations, thereby creating more noise in the system at those points. There is nothing explicitly preventing the continuous signal from becoming arbitrarily large at any point along the way, allowing the noise associated with large signals to reach arbitrarily large levels. Using L_∞ norm avoids this problem, as it sets a cap on the absolute magnitude of the control signal, and therefore, the signal dependent noise. In addition, preventing arbitrarily large signal magnitudes is in line with the physiology of (populations of) neurons, which can only fire with a maximum strength and frequency.

3.4.2 Future Work

Future work should consider extensions of the simple minimum effort cost functions used in this work in order to describe a richer set of movements. This can be accomplished by simply penalizing the norm of the effort term in any cost function (such as those outlined in [11]) with an L_∞ norm instead of the L_2 norm that is typically used. Additionally, the possibility be should explored that some type of *combination* of signals of the type shown in figures 3.3, 3.4, 3.5, and 3.6 may be advantageous when modeling a richer set of movements. Combining select

signals might form a basis set which would provide a larger subspace of signals that, by extension, controls a larger subspace of human movements than have been explored in this work. We envision a set of sparse signals generated by separate populations of neurons which produce sparse control signals that are combined (linearly or otherwise). This "combination of signals" concept is inline with the idea that the CNS achieves control through relative activation of motor primitives (as described by Giszter [71]). He empirically illustrates how spinal motor primitives can be thought of as a basis set that can be combined in varying degrees to achieve a desired movement. However, in addition to controlling the relative magnitude of activation between primitives, changing the *timing* of the activation of the primitives is equally, if not more important, as only a proper sequence of motor primitive activation will provide the desired motor output. A basis set of signals as described in this work would be advantageous as it would require a minimum amount of neural structures to control human movement. This combination of "on-off" sparse signals might be thought of as controlling motor primitives by varying the activation times between them. Future work should consider combining such sparse signals to describe a richer set of movements. Furthermore, since this work is an abstraction of neural control signals, and it models reaching movements well, the abstraction may be a useful platform to apply to robotic motor control.

3.5   Acknowledgments

This chapter is a reprint of the material as it appears on Cornell University's

arXiv. Gamble, G.G., M. Yazdani. "Sparse signals for the control of human movements using the infinity norm", http://arxiv.org/abs/1601.01410, 2016.

Chapter 4

Suboptimal combinations of cost functions fit fast human reaching movements as well as optimal combinations

Central to the approach of using optimal control to model human movements is the minimization of costs and/or the maximization of rewards in order to generate a movement trajectory. When more than one type of cost/reward is to be considered (e.g. movement accuracy or energy usage), an objective function with multiple weighted terms is used, and is typically referred to as a *composite* function. This approach yields a movement trajectory that minimizes/maximizes the weighted costs/rewards while meeting some set of constraints (e.g. the movement should be completed within a certain amount of time). Inherent in this approach is the assumption that human movements are optimal with respect to some costs and/or rewards. However, this notion has been recently challenged. In this work we compare a cost-optimal approach to modeling fast human reaching movements with a cost-suboptimal approach. We analyze two different models of fast, point-to-point human reaching movements. One model, referred to as the composite model, is optimal in terms of minimizing costs, while the other model which we

call "confluence", is not. Both models trade off between two weighted costs which minimize derivatives of position and are known to produce the bell shaped velocity profiles and smooth positional trajectories observed in human reaching movements. Specifically, both models trade off between the costs of "jerk" and "crackle", the third and fifth derivatives of position respectively. The composite (cost-optimal) model employs a two term cost function of weighted minimum jerk and minimum crackle. In contrast, the confluence model uses a cost-suboptimal approach which employs two *single* term cost functions, minimum jerk and minimum crackle, and combines the results from two separate optimizations over each cost function. We quantitatively compare the reaching trajectories generated by each model with fast point-to-point human reaching movements recorded from five subjects. Our results indicate that, for the purposes of modeling human reaching movements, although the cost is strictly higher for the the confluence (cost-suboptimal) approach presented here, it is statistically equivalent to the composite (cost-optimal) approach in terms of error when comparing the models' dynamic profiles to human reaching movements. Additionally, on a per subject basis, when root mean squared error is averaged across all reaching trials, the confluence approach always results in lower error when compared with the composite approach. Together, these results indicate that confluence should be considered as a viable alternative for modeling human movements, while calling into question the assumption that human movements are governed by minimum cost optimality principles.

4.1   Introduction

The method by which animals achieve graceful and effective movement is not understood. Optimal control theory has the potential to elucidate such biological mysteries. The optimal control approach to understanding animal motor systems

assumes that animal motor systems have evolved to be "optimal" in some sense. When optimal control is applied in an effort to understand human movement, the validity of the approach rests upon this assumption of optimality. However, it is plausible that animal motor systems have evolved to simply be "good enough" to survive in their niche without achieving optimality. The assumption of optimality in animal motor systems has been recently questioned, as has the use of optimal control as the best paradigm by which to model and understand animal motor systems [86, 87]. In this work we introduce a suboptimal model of human reaching movements that compares favorably with the optimal approach.

When applying optimization theory to the modeling of biological movement, the investigator must consider what cost function (or functions) best model the movements in question. An optimal control policy is a strategy that is "optimal" in the control of a system (biomechanical system, robotic arm, etc.) with respect to minimizing (or maximizing) a cost (or reward) function (from this point on, we refer only to "cost", with the understanding that analogous concepts apply to "reward"). Techniques from optimal control theory which utilize a single cost function to arrive at a control policy for human movements vary in nature. Some focus on minimizing torque change around a joint [55], some concern themselves with the state of the hand during reaching and grasping maneuvers [56], others focus on the relationship between signal strength and the amount of noise in the motor system [57], and still others minimize higher derivatives of position in order to produce smooth trajectories [10, 32, 35, 62]. These policies are optimal with respect to the cost function that each work has proposed.[1] But, are human movements truly optimal with respect to all of these (or other) cost functions?

---

[1]In finding an optimal control policy, in addition to the cost function, the set of constraints in the optimization problem also plays a role in determining the characteristics of an optimal control policy. Details of the constraint set used in this work can be found in section 4.2.1.

It is accepted that a single-term cost function is not likely to capture the breadth of human movement over a multitude of tasks. A clear improvement to this is the multi-term cost function [11, 88], which trades off amongst multiple costs and therefore is more flexible. These weighted, multi-term functions are called *composite* cost functions. Models that employ composite cost functions seek to find the single unique solution to minimizing those collective costs while still meeting a set of constraints. It is proven that for a well defined class of problems, this organization of cost terms will yield the solution that is optimal, in that it has the lowest cost of any solution [36]. Deviation from the composite cost function organization (shown in equation 4.1) is suboptimal, again, in terms of *cost*, but not necessarily in terms of *accuracy* when the model's dynamic profiles are compared to real world data.

While the dynamic profiles generated by a model of movement depend upon defining the notion of cost, specifically in this work, they also depend upon the *organization* of the functions chosen to represent that cost. In this work, we offer an alternate way to express the notion of cost. Composite cost functions weight the importance of several terms within a single function. These terms may represent, for example, the amount of energy expended in a movement, or the degree by which the target position of a movement is missed. The technique introduced here, called "confluence", instead uses *multiple single-term cost functions* that are used to optimize control signals *individually*.

Although the confluence approach outlined here is rooted in optimal control, its formulation is cost-suboptimal in terms of the mathematical organization of its cost functions. Despite this, for the modeling of human reaching movements, the confluence based model performs as well as the the traditional composite (cost-optimal) approach. This indicates that optimal composite cost function should

not be considered the de facto technique for modeling animal movement, and agrees with the suggestion by Loeb that the vertebrate motor system may employ a suboptimal "good-enough" control strategy [86].

In order to simplify the comparison between the composite (cost-optimal) and confluence (cost-suboptimal) approaches, we choose two simple functions from a well known family of kinematic cost functions which have proven effective in modeling human reaching movements, and apply them within each of the two paradigms. Minimizing various derivatives of position has proven effective to model various aspects of human movement. This approach yields the smooth position and velocity profiles we observe in human reaching movements. The two cost functions we employ in both models are the minimization of "jerk", the third derivative of position, and "crackle", the fifth derivative of position. We choose these because they are well understood, and serve as a straight forward way to compare the composite approach with confluence. For further discussion on the choice of cost functions, see 4.1.1 We run experiments using these costs organized in the manner prescribed by both confluence and composite and then compare the results with human reaching movements.

4.1.1 Background

We compare the composite and confluence approaches by implementing two versions of a "minimum effort" control problem, which is an optimization problem that seeks a control policy that meets a desired set of constraints and uses the smallest possible "size" or "effort" of the control signal. For example, Flash and Hogan authored one of the seminal works describing fast reaching movements [10]. Their model is a minimum effort model which uses jerk as the control signal, and therefore uses the smallest possible jerk signal to achieve its goals. For this reason

it is called a "minimum jerk" model. They (and others) chose to quantify the dynamics of a fast reaching movement by finding the ratio of the movement's peak velocity to its average velocity. We use r to refer to this value:

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

Minimum effort models for fast reaching movements based on single term cost functions have rigid values for r that do not change for a given set of initial and final conditions. As seen in figure 4.1 and Table ??, the r value is different for each of the derivatives of position that each model minimizes. As indicated by [89], experimental measurements of r can be used to determine the best derivative of position to minimize. Flash and Hogan's data was based on one subject's movements over thirty trials and r was estimated to be 1.805 with a standard deviation of 0.153. One justification for Flash and Hogan's choice of a minimum jerk model is that it predicts an r value of 1.875, very close to their experimental data [10]. These results were corroborated in subsequent work [90]. However, there is confusion in the literature regarding whether the minimum jerk model fits human reaching movements best. Some studies with more extensive data show that models based on minimum snap (fourth derivative of position) fit the data better [91, 92, 93]. Other studies show that crackle fits best for hand movements that control a non-rigid object [94]. Even still, other studies continued to use jerk minimization but noted that snap minimization is a better fit in many cases [95]. Modern scientific works continue to accept the Flash and Hogan minimum jerk model as a baseline assumption and use this to build their arguments [96, 97, 98].

All of these models fit the data only on average, and do not account for the

Table 4.1. r values predicted by several minimum effort control policies.

Predicted r Values	
Derivative Order	r
3	1.875
4	2.188
5	2.461
6	2.707
10	3.523
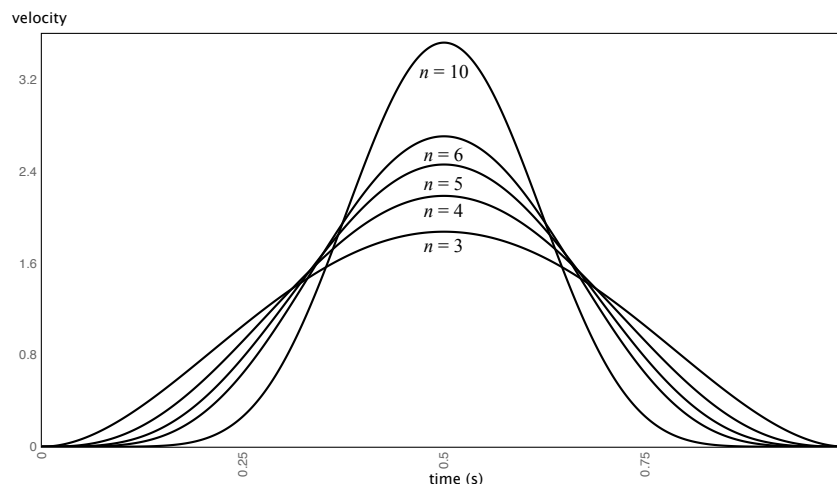
Figure 4.1. Different velocity profiles resulting from minimizing various derivatives of position. n represents the order of the positional derivative that is minimized.

variability of r observed across subjects and trials. An inspection of our data set of fast, point-to-point hand reaching movements (explained fully in section 4.2.1) confirms that movement dynamics vary from trial to trial, and that minimization of derivatives higher than jerk fit best in many cases. Often, no single model in the minimum effort family provides a good fit because a trial's r value falls in between the r values generated by two of the minimum effort models. This leads us to suggest that combinations of multiple minimum effort models are required to more accurately capture the variation in movement from trial to trial and subject to subject.

4.2 Materials and methods

Here we describe the collection of human data for fast reaching movements. We then define the "Optimal Confluence Control Framework" (OCCF), a general framework for modeling human movement. We then demonstrate an instance of this framework called the "Multi-Derivative Optimal Confluence" (MDOC) model, which is a basic implementation of the more general OCCF. The MDOC models fast reaching movements by combining two optimal control signals related to two cost functions each based on a single term, and will be compared to a model based on a traditional composite cost function, with the two terms combined into one function. The methods of comparison between the two are discussed.

4.2.1 Experimental setup

The human hand movement data used in this work was originally collected by Karniel and Mussa-Ivaldi [41]. We used a subset of this data that was relevant for our study of fast movements (baseline unperturbed movements). We summarize their experimental setup below and refer the reader to [41] for a more complete

description.

Five subjects (four male and one female) participated in an experiment involving a manipulandum (see [99] for a detailed description of the 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 10cm and formed an equilateral triangle as shown in figure 2.2. 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 ($\pm$50ms). Many trials fell outside this window. These trials were kept, but at the end of each trial, visual feedback was given indicating if the execution of their movement was within the desired 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. Figure 4.2 shows the profile of a typical movement. Since these movements are restricted to the two-dimensional plane and tend to be straight from the starting to ending position, we only consider the component of the movement that corresponds to the straight line direction between the start and end points A, B, or C.

Evidence suggests that only feedforward signals are necessary to achieve fast "ballistic" reaching movements. Studies on deafferented animals show that trajectory planning for these fast point-to-point movements is not disrupted and that proprioceptive or cutaneous feedback is not necessary for the execution of such movements [45]. Because of this, we can model these movements as a feedforward
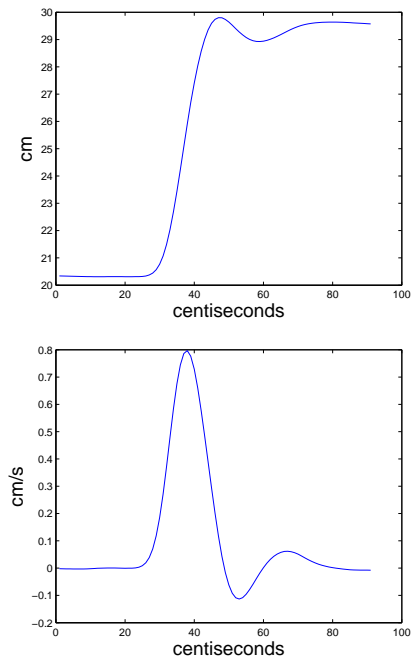
Figure 4.2. A recording of position and velocity profiles from a typical reaching movement. The oscillatory behavior at the end of the movement corresponds to the overshoot and correctional effects described in the text and are discarded since they are not part of the ballistic portion of the movement. The units of time and position are not relevant and are not shown.

control problem, as done in [62, 35, 100]. For each trial, we identify the "ballistic" portion of the movement, that is, we select the portion of the trial between the points where the movement has started and completed its "feedforward" portion (see figure 4.3). There are various methods for movement onset and offset detection [42, 43, 44], 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 the ballistic portion of a velocity profile algorithmically in two steps. The first step is intended to find a point near the beginning or the end of the relevant portion where the velocity profile is close to zero. We begin by finding the point in time when the velocity profile has reached its peak. Fast movements always have a unique global maximum (unless the trial is an outlier). Once this maximum point is found, we proceed to consider sequential velocity samples before and after the peak velocity and check to see if they fall below a threshold of .05 mm/s. I.e., We descend on both sides of the bell shape curve looking for a velocity value that is lower than the stated threshold of .05mm/s. Once the sample is found, we refer to this sample as the "start" or "end" depending on whether the sample is before or after the peak velocity. This threshold may first be crossed on either the start half of the bell curve, or the end half of the bell curve.

The second step, which we call the"corrective step" starts at the near zero point on the velocity curve. Its purpose is to move the entire range of the chosen velocity profile (both the start and end points in lock step) to find a start and end point that differ by no more than .2mm/s. Let us consider the case of the initial threshold (in step one) being first crossed on the start (left) side of the velocity bell curve. That velocity is guaranteed to be close to zero ($<= .05mm/s$). This means
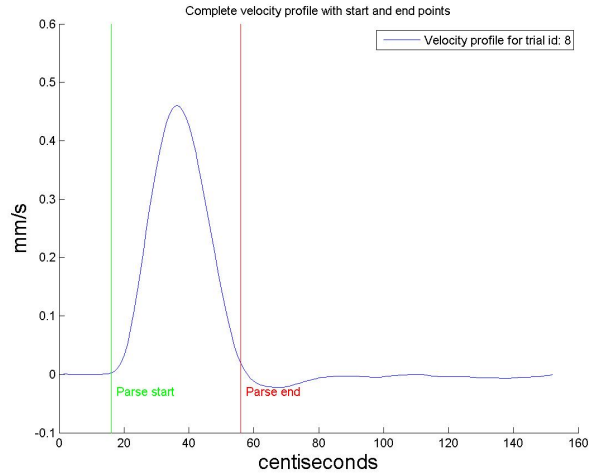
Figure 4.3. A typical result for the onset and offset detection algorithm which identifies the "ballistic" portion of the velocity profile.

that the chosen starting point will be less than or equal to the end point velocity. The corrective step then moves both the start and end point forward one step at a time, searching for a point at which velocity values both are within $.2mm/s$ of each other. This means the start point is climbing back up the left side of the bell curve, and the end point is descending down right side of the bell curve. When $End_{vel} - Start_{vel} < .2mm/s$, the final start and end velocities are decided upon. A similar approach is used in the case where the initial $.05mm/s$ is crossed on the end (right) side of the curve.

Regarding the parameters $.05mm/s$ and $.2mm/s$, we visualized many possible selections. Qualitatively, any choice within a given range these captured the ballistic portion of the fast reaching movement very well. Figure 4.3 is indicative of the onset and offset detection of almost all non-outlier trials.

Regardless of what onset and offset detection method used, it is very rare for the trial to start and end at complete rest (i.e. with velocity equal to zero). As a result we do not set the boundary conditions to zero, but rather to the starting

and ending velocity and acceleration of the given trial. We posit that the higher order dynamics in the boundaries must be zero at some point. As a result, we set the final boundary condition of jerk to zero (similar to the "traditional" constraints mentioned in [58]). Later, we compare this approach to the composite approach, and we apply the boundary conditions in the same way.

4.2.2 The Optimal Confluence Control Framework: A generalized framework for modeling human movements

There is abundant evidence that heterogeneous neural control signals originating in disparate parts of the motor system are combined (see chapter 5.1 and figures 4.4 and 4.6a) in order to produce final motor output. We propose an approach that consists of a confluence of multiple control signals that are each optimal with respect to different cost functions (Fig. 4.5). We suggest that our proposed framework, The "Optimal Confluence Control Framework" (OCCF) is a step towards applying optimal control theory as a means to represent and understand biological motor systems more directly.

We choose the term "confluence" to represent the concept of individually optimized neural signals working together, being transformed and combined, to control the motor system. Our OCCF 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. Instead, our approach posits that the optimization of individual neural control signals was accomplished over millennia via evolutionary processes, and that the only real-time neural computations performed during a movement resembles a weighted combination (linear or
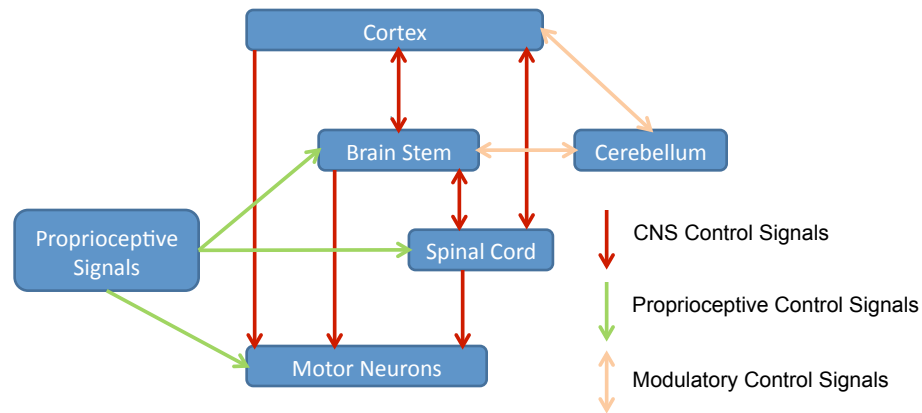
Figure 4.4. A high level view of some important components of the vertebrate motor system and the the connectivity between them. Despite the high level abstraction and simplicity of the diagram, note that signals from multiple areas of the motor system converge simultaneously on other areas of the motor system. Evidently, they all work together to achieve desirable motor control.
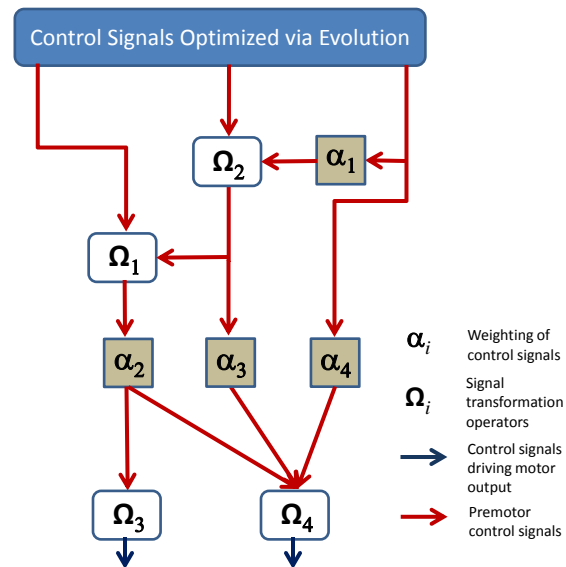
Figure 4.5. High level schematic of one example instance of the Optimal Confluence Control Framework. The figure is not intended to mirror any particular neural system. Instead, it shows that any desired network architecture can be achieved in order to mirror any known connectivity found in the HMS, or elsewhere.

otherwise) of neural signals, similar to a suggestion made in [11], along with simple transformations, e.g. integration of a signal over time.

There are key differences between OCCF and traditional optimal control strategies. The OCCF can abstractly represent the nervous system's network architecture. Its method for trading-off between various properties of a movement (e.g. speed vs. accuracy) represents signals from one population of neurons becoming more or less dominant relative to others. The OCCF allows the investigator to map the network architecture of the model to network architectures found in the nervous system, allowing scientists to use the model to draw inferences related to anatomy and physiology (See Fig. 4.6). By contrast, the traditional multi-term approach intertwines all optimization goals together into a single cost function. This leads to each term influencing the other with regard to final output, with no notion left of what the individual terms would have represented by themselves. An example of this distinction is illustrated in the following equations. The first is the generalized equation used to arrive at a cost employed by traditional, multi-term composite cost functions[11]:

$$\underset{u(t) \in S}{\arg\min} \sum_{j=1}^{k} \alpha_j c_j(\mathbf{x}(t), u(t)). \tag{4.1}$$

The second equation represents the method employed by the OCCF:

$$\sum_{j=1}^{k} \alpha_j \underset{u(t) \in S_j}{\arg\min} c_j(\mathbf{x}(t), u(t)). \tag{4.2}$$

In both equations, $\mathbf{x}(t)$ is a function of state over time, $u(t)$ is the control signal, k is the number of cost functions being considered, c_j is some set of cost functions,

(a)   **(b)**
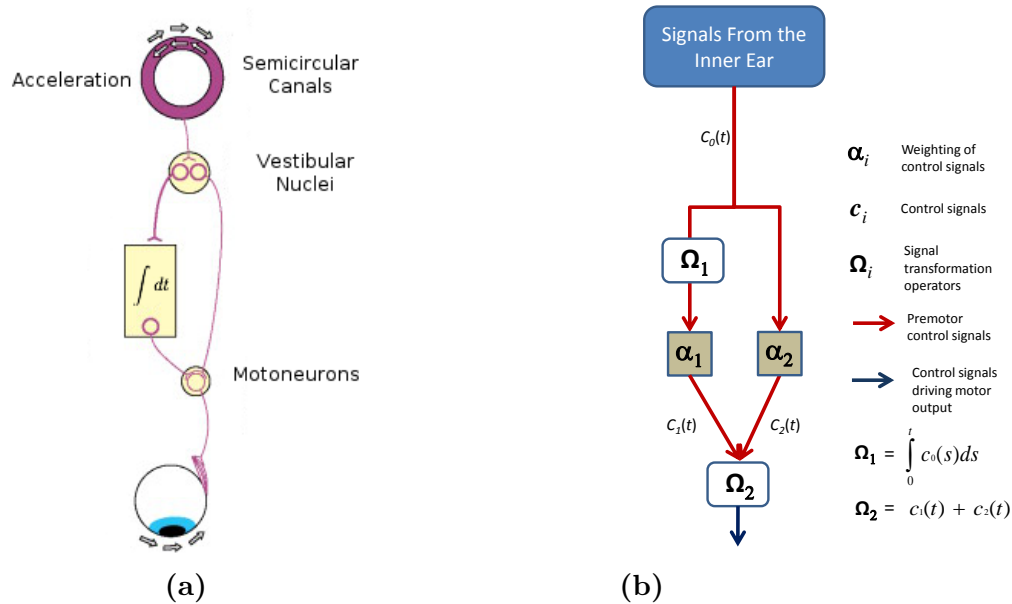
Figure 4.6. Fig. 4.6a depicts the anatomy of the vestibulo-ocular system (adapted from [1] and reprinted with permission). The VOR system utilizes measurements of head movement to control eye movements. Information related to the angular velocity of the head flows from top (semicircular canals) to bottom (oculomotor neurons). Some pathways are integrated over time and all pathways are *combined* to produce the final control signal to the muscle. Fig. 4.6b shows an example of how the Optimal Confluence Control Framework can be used to mirror a biological system. At the top of the figure, a pre-optimized control signal $(c_0(t))$ represents the ability of the inner ear to detect acceleration and output a signal representing velocity. This ability was optimized over time via evolutionary pressures. On the left branch, Ω_1 is an operator integrating the velocity signal over time to produce a positional signal. The left and the right signal branches are combined with various weights (α_i) before being delivered to actuators in the system. Note the combination of signals at the bottom of both diagrams. This transformed, weighted combination of optimized signals is what we call *confluence*.

α_j are the weights paired with each cost function, and S is a set of constraints that $u(t)$ must meet for the composite formulation (eq. 4.1). However, S_j are the constraints that each individual $u(t)$ must meet in the confluence formulation (eq. 4.2), allowing confluence a separate set of constraints over each optimization. Equation 4.1 yields the optimal $u(t)$ based on the sum of weighted costs while equation 4.2 sums weighted $u(t)$ functions which have been optimized over each cost and constraint set individually.

In our example implementation (MDOC), the resulting $u(t)$ signal then controls an end effector (a representation of a hand) to arrive at a desired final position. We demonstrate how these signals can work together to model dynamic profiles of reaching movements effectively.

In order to offer a direct comparison to equation 4.1, and for simplicity, equation 4.2 (OCCF) uses a summation to combine the control signals, but any kind of combination strategy may be used under the OCCF. These concepts are further justified in chapter 5.1: "Anatomical and physiological justification for cost and control signal reorganization".

4.3 Definition of the OCCF

As indicated in Figure 4.5, an instance of the OCCF consists of five sets of elements (see the section 4.4 section for a concrete example):

1. *A set of pre-optimized control signal producers crafted via selection of cost or reward*: "Pre-optimized" means that, for all chosen cost functions, the optimization of their corresponding signals are carried out prior to being combined (confluence) with each other. An optimized signal producer is crafted once, based on some optimality condition, then held static, henceforth

only differing in its relative influence compared to other signals with which it interacts (see item 3). Elements of this set are analogous to evolutionary process forging structures in the nervous system and the character of the neural signals they produce.

2. *A set of signal transformation operators which can be defined in any way such that they take one or more signals as input and transform them to produce an output signal (See section "A simple application of the OCCF (confluence) to model human movements: MDOC" for an example of how signals might be combined/transformed)* . Elements of this set are analogous to any structure in the nervous system where neural signals are transformed and/or combined (e.g. Fig. 4.6a).

3. *A set of weights used to adjust the relative amplitude of the signals mentioned in item 1.* Differing ratios between cooperating neural signals produce different movement dynamics[2].

4. *A set of network connections.* These outline the ways in which signals connect to signal transformation operators, signal weighting points, and output points. They are analogous to the overall network architecture of the neuronal system being modeled. [3]

5. *A set of strategies for choosing the values for the weights (item 3) based on the desired characteristics of movement.* We hypothesize that changing

---

[2]This is not the first framework of motor control to blend multiple signals, e.g. [101]. However, such models are not suited to be a framework that mirrors the known connectivity of the nervous system. In addition, they do not pair this reflection of the nervous system's network architecture with optimization techniques that mimic evolutionary processes as a means to craft a basis set of optimized control signals.

[3]In practice, the set of weights can be combined with the set of network connections, where there is no connection, the weight becomes zero, but here we separate them for clarity.

the weighting of signals is an important real-time computation occurring during the control of movement. The weightings are analogous to excitatory neurons increasing the activity of their targets, and to inhibitory neurons either inhibiting or disinhibiting the activity their targets (e.g. Renshaw cells which down modulate populations of alpha motor neurons in the spinal cord, or the disinhibitory systems involving pallidal neurons in the basal ganglia). See section 4.4 for an example of a basic strategy for weight selection.

4.4   A simple application of the OCCF (confluence) to model human movements: MDOC

Here, we implement a simple model using the generalized OCCF called the Multi-Derivative Optimal Confluence model (MDOC). The key difference between the MDOC and traditional composite approach to a minimum effort problem is its method for trading-off between various minimum effort strategies. The MDOC separates each goal into a separate optimization problem, then later combines the results. By contrast, the traditional multi-term approach intertwines all optimization goals together into a single cost function. This leads to each term influencing the other with regard to final output, with no notion left of what the signal related to the individual terms would have represented by themselves. This distinction is shown in equations 4.1 and 4.2.

We define two generalized control problems that can vary based on derivatives of position, boundary conditions and norm used to measure the control signal. These control problems are formulated in discrete time and will form the basis of the MDOC.

$$\underset{\mathbf{x}}{\text{minimize}} \quad \| D_m \mathbf{x} \|_p$$

$$\text{subject to} \quad A_{eq_s} \mathbf{x} = b_{eq_s}$$

$$(4.3)$$

and

$$\underset{\mathbf{x}}{\text{minimize}} \quad \| D_n \mathbf{x} \|_p$$

$$\text{subject to} \quad A_{eq_s} \mathbf{x} = b_{eq_s}$$

$$(4.4)$$

Where $\mathbf{x}$ is a vector containing discrete samples of hand position from continuous time recordings of reaching movements, n and m are derivatives of position to be minimized, $m \neq n$, D_n and D_m are matrices used to approximate the n-th and m-th order derivatives of the position vector $\mathbf{x}$, A_{eq_s} and b_{eq_s} represent the constraints of the control problem, such as starting and ending boundary conditions. (see [102] for details).

We now discuss the control problem of making fast point-to-point reaching movements. We implement a minimum effort control policy consisting of two parts corresponding to equations 4.3 and 4.4, where $m = 3$, $n = 5$ and $p = 2$. That is, two minimization problems that share the same constraint set, one that minimizes jerk and the other that minimizes crackle. These two costs correspond to c_1 and c_2 in equations 4.1 and 4.2, where $k = 2$. We then combine the resultant signals from each minimization process. The method by which the signals are weighted and combined is discussed in section 4.5.1. The signal resulting from this combination is used to control a representation of an end effector (hand). This is the "MDOC", of which figure 4.7 provides a visual representation. At its core, the MDOC is the linear combination of two optimal control signals, one that is optimal with respect to a minimum jerk cost function and a separate signal that is optimal with respect to a minimum crackle cost function.

The MDOC trades off between two cost functions (and their corresponding

optimal signals) related to derivatives of position (as done with the traditional single function, multi-term approach in [35]). We choose these costs because they are relatively simple, well understood, and serve to illustrate the implications of reorganizing cost functions as defined by the contrast between equations 4.1 and 4.2. Specifically, we choose minimized jerk as it is has been accepted as a standard choice for smoothness maximization (see section 4.1.1). We choose crackle as the second minimized cost. We considered higher order derivative experiments (beyond crackle) for completeness and to fully establish the relationships between given minimized derivatives and the boundary conditions used. Generally, RMSE decreases for both composite and confluence as the order of the minimized derivative increases and as the number of boundary conditions increase. However, the drawback to minimizing higher order derivatives when modeling fast reaching movements is increased cost, as shown in figure 4.8. This reflects a tradeoff between cost and accuracy of the velocity profiles modeled. In addition, with each increase in order of minimized derivative, the velocity profile's sensitivity to very small changes in the higher derivative control signal increases. This means that small changes in the modeled control signal result in disproportionately large changes in the velocity profile. For this reason, the more complex models (higher derivatives minimized along with more boundary conditions) run the risk of modeling noise. The jerk/crackle formulation was selected to balance these concerns. That said, for comparing composite with confluence, the most important aspect of the experiment is that both approaches use the same higher minimized derivative and the same boundary conditions.

Additional model complexity is not necessary due to the constrained nature of the experimental setup (e.g. the movements have constrained distances and times, have no reward/punishment involved, etc.). By eliminating these variables, we can compare critical difference between the two approaches, i.e. the effects of

generating control signal(s) via the approach outlined in equation 4.1 (composite) vs. that in equation 4.2 (confluence). The MDOC model is intended to be a simple example that demonstrates the efficacy and plausibility of this type of cost function/optimization organization for modeling human movements.

Formal definition of the MDOC in terms of the OCCF framework

The MDOC is formally defined in terms of the five sets of elements any OCCF implementation has, as denoted in section 4.3. See figure 4.7 for a graphical representation of this definition:

1. A set of pre-optimized control signal producers:

   $\{C_1, C_2\}$, where C_1 produces a minimum effort signal defined by equation 4.3, where $m = 3$ and $p = 2$, and C_2 produces a minimum effort signal defined by equation 4.4, where $n = 5$ and $p = 2$

2. A set of signal transformation operators which can be defined in any way such that they take one or more signals as input and transform them to produce an output signal:

   $\{\Omega_{12}\}$, where $\Omega_{12} \equiv \Omega(x, y) = x + y$

3. *A set of weights used to adjust the relative amplitude of the signals generated by members of item 1:*

   $\{w_1, w_2\}$, where w_1 modulates the signal generated by C_1, and w_2 modulates

the signal generated by C_2

4. *A set of network connections:*

$V = \{C_1, C_2, \Omega_{12}\}$

$E = \{(C_1, \Omega_{12}), (C_2, \Omega_{12})\}$

Where V is the set of vertices on a graph, E is the set of edges between vertices, and Ω_1 transforms the dual input of C_1 and C_2 and also provides the output from the network.

5. *A set of strategies for choosing the weights (item 3) based on the desired characteristics of movement:*

w_1 and w_2 are chosen such that $w_1 = 1 - \theta$ and $w_2 = \theta$. A grid search is performed by iterating from $i = .00....1$, with step size 0.05. All i values become potential θ values. Of this set of potential θ values, the value that yields the lowest root mean squared error is chosen.
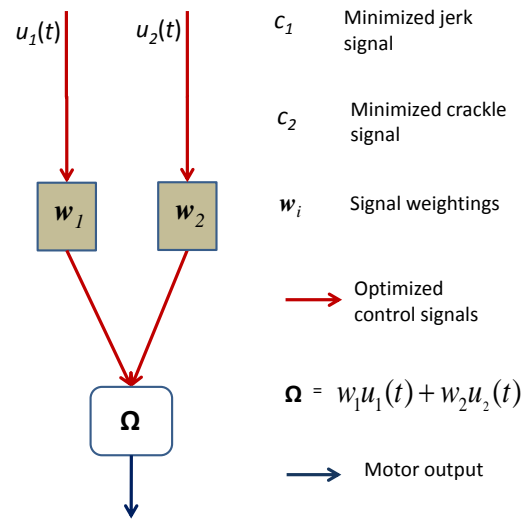
Figure 4.7. A representation of the MDOC. Separate, optimal signals are linearly combined to control motor output.
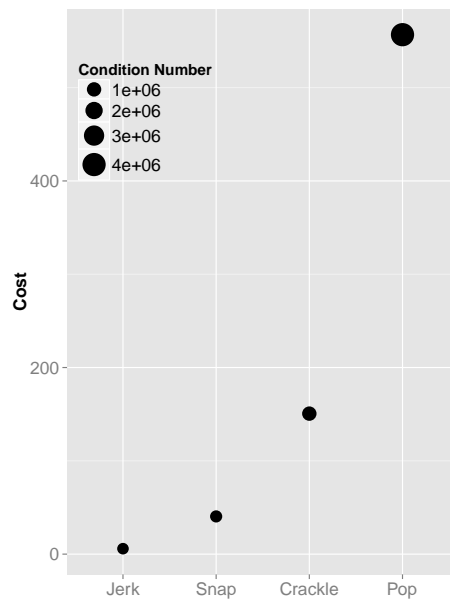
Figure 4.8. A representation of the *cost* when modeling an idealized reaching movement that is representative of a typical human trial by minimizing different derivatives of position. The height of each circle represents the relative amount of cost (unitless) across four scenarios: minimized jerk, snap, crackle and pop. The condition number for each approach is also indicated by the size of the circles.

4.5 Methods of comparison between the MDOC and the composite models

4.5.1 Selecting cost weighings for the MDOC and composite models

The proper ratio with which to weight two cost functions for a given trial and model is unknown. Therefore, for both confluence and composite, for each trial, we first selected the ratio that yielded the lowest RMSE. This approach is similar to the free parameter selection technique used in [35]. We refer to this ratio as θ. For the confluence method, for each trial, we grid searched for the θ which yielded the lowest RMSE as described in part 5 of section 4.4. We grid searched in the same way for θ to weight the two terms of the composite approach.

4.5.2 RMSE calculation

We compared the root mean squared error (RMSE) of the MDOC with its composite counterpart. We computed the RMSE between the models' velocity profiles and those of the human subjects by summing the squared error at each time step of the movement, and taking the square root of the sum. For both models, the portions of the velocity profiles which were forced to be accurate due to setting boundary conditions were not included in the RMSE calculation.

4.6 Results

As mentioned in section 4.1.1, velocity profiles vary from movement to movement and from subject to subject. One way to model this phenomenon is to combine the minimization of various derivatives of position, each with different peak velocities and thus different bell shaped velocity profiles (see figure 4.1). There

is more than one way to accomplish this. The traditional method would be to include both effort terms (e.g. jerk and crackle) in the same cost function. Each of the terms may be weighted to emphasize or deemphasize as needed. We refer to this technique as using a composite cost function.

One of the key observations in this work is that the MDOC's method of separating two or more cost terms into two separate optimizations, then combining the results of those optimizations, is fundamentally different than the traditional composite approach of adding the same terms into a single cost function. Despite the fact that the MDOC is suboptimal in terms of cost, it models fast human hand movements as well as the composite method.

We compare the overall average RMSE for both the composite and confluence approaches on a per subject basis in figure 4.9. The RMSEs of the stand alone minimum jerk, minimum snap, and minimum crackle models are also displayed as reference points, but we remind the reader that those models are less complex, and therefore are not expected to perform as well. The comparison relevant to our conclusions is that between composite and confluence, as these two models stand on equal footing with the same number of parameters. We also provide visual examples demonstrating how the two models are qualitatively similar to each other and to human velocity profiles in figures 4.10 and 4.11.

Additionally, for each subject, we report the average RMSE across all trials for the selected θ values (θ yielding the lowest RMSE per trial) for both models under consideration. Figure 4.12 shows the 95% confidence intervals of the confluence and composite average RMSEs per subject. We observe that there is not a significant difference between the confluence and composite approaches. We use the nonparametric Kolomogrov-Smirnov test and perform a two-sided pairwise test for each trial, and we define the null hypothesis that there is no difference between
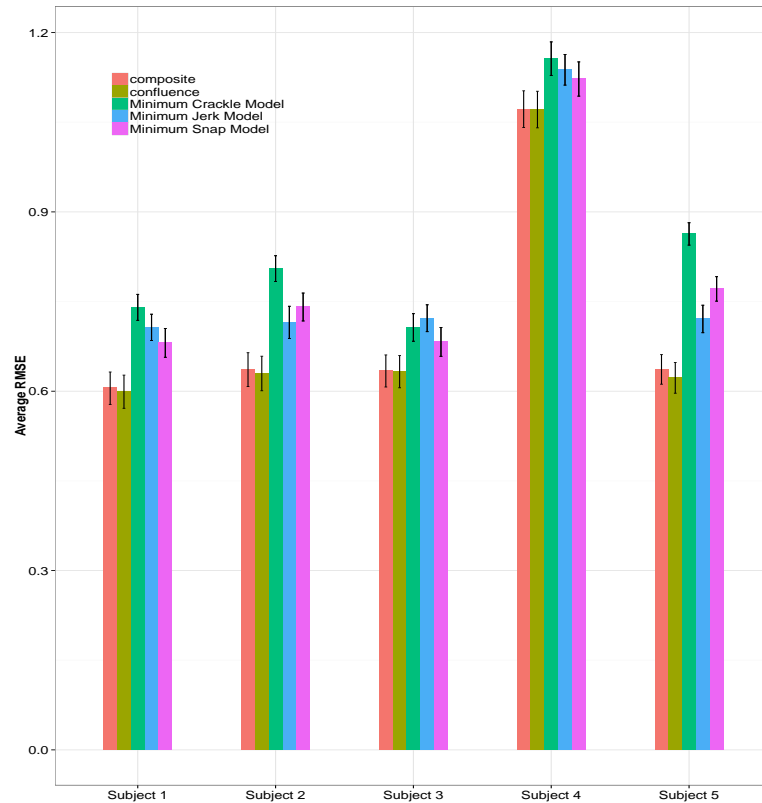
Figure 4.9. Root mean squared error (RMSE) comparing five models for all trials for all five subjects. Error units are in mm/s. The bars depict the RMSE of the minimum jerk, snap, crackle models, as well as the RMSE for the composite and the MDOC. Standard error bars are shown.

the two models in terms of RMSE. Table 4.2 shows that in all cases the p value is far too large for us to reject the null hypothesis. We therefore conclude that overall the confluence model and composite model have no significant differences in error on this data set. Both figures (4.9 and 4.12) indicate this as well.

Of additional interest is the distribution of the θ values chosen to yield the minimum RMSE. See figure 4.13 for the distribution of θ values chosen for both models for a single subject. See figure 4.14 for the θ distributions of subjects 2-5.
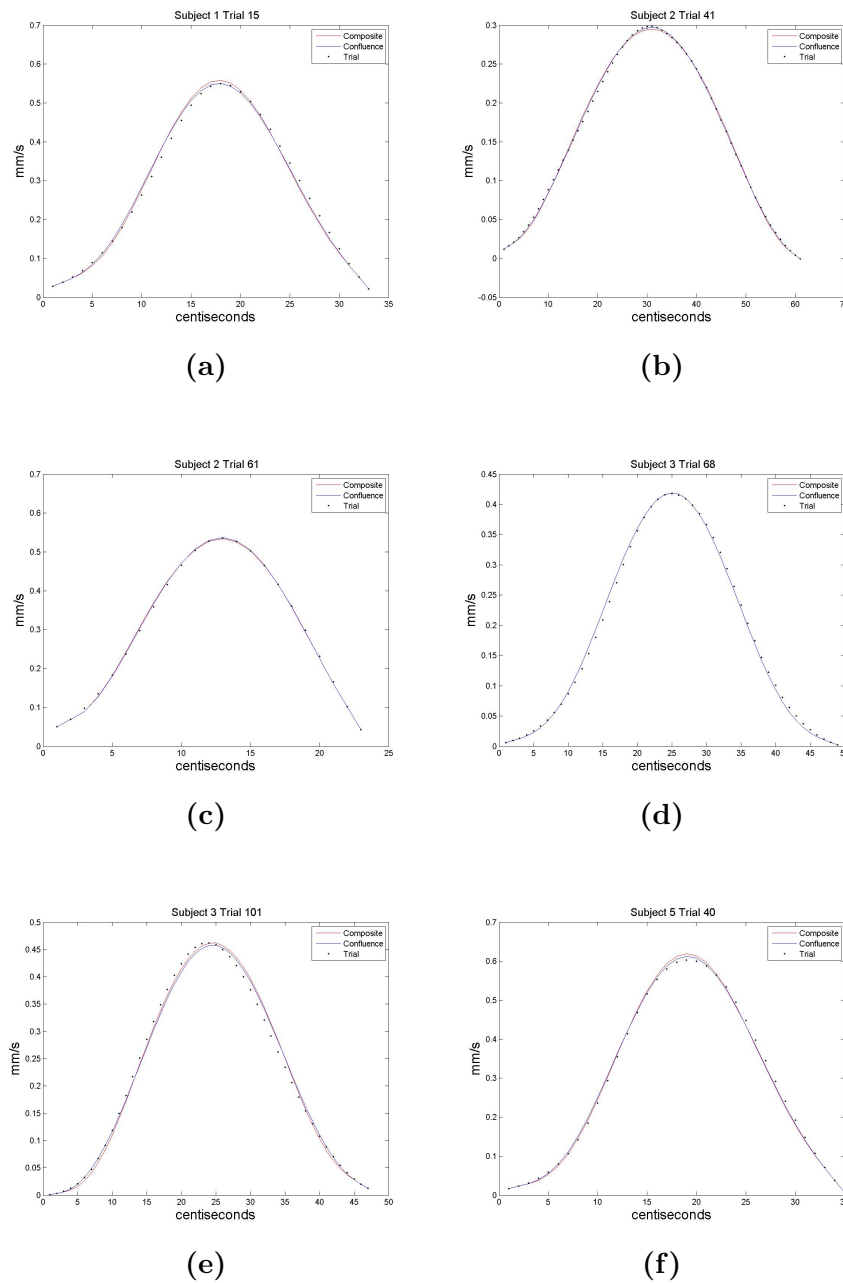
Figure 4.10. Typical subject velocity profiles compared with the profiles generated by both models. Note that, in these examples, the peak velocity varies from approximately .3 mm/s to .6 mm/s. No examples are given from subject 4 because this subject clearly has a different mode of movement than the other four. The other four subjects are comparable in terms of RMSE and qualitative inspection of their velocity profiles. Subject 4's profiles are qualitatively different for many trials, as shown in figure 4.11.

(a. thetas = 1)

(b. thetas = 0)

(c. thetas = 1)

(d. thetas = 1)

(e. thetas = 0)

(f. thetas = 0)
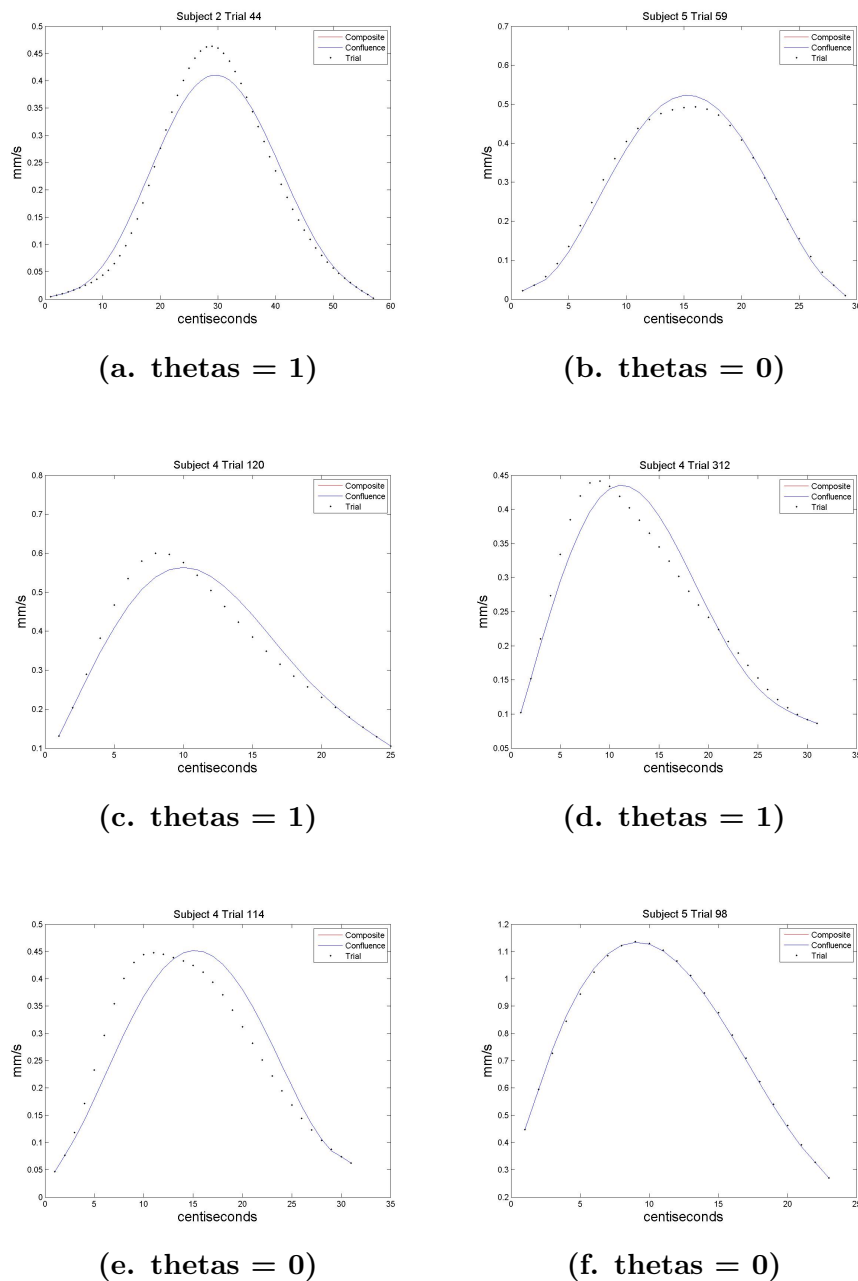
Figure 4.11. Trial and model velocity profiles for which the $\theta = 0$ or $\theta = 1$ is chosen by both models. The θ chosen by both models is listed under each subfigure. Subject 4 is represented here disproportionately because this subject's movements often differ greatly from the other four subjects, and also from the models' profiles, leading to a more frequent selection of $\theta = 0$ or $\theta = 1$. Note that only a blue profile can be seen, due to the fact that at $\theta = 0$ or $\theta = 1$, the models are either pure minimum jerk or pure minimum crackle, repsectively. Therefore, the confluence and composite plots are directly on top of each other.
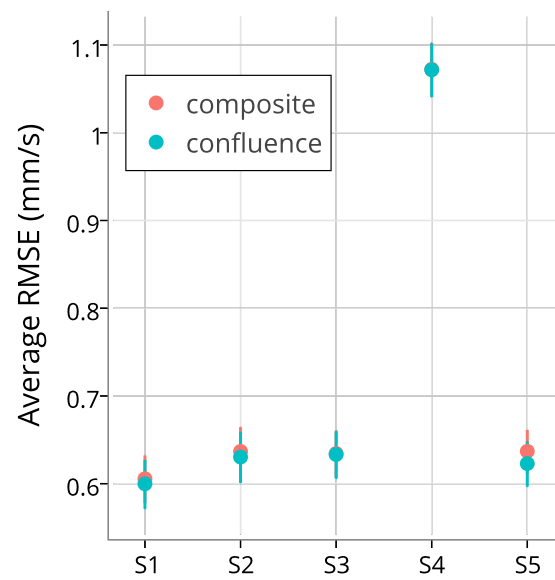
Figure 4.12. A comparison of the root mean squared error, over all trials, of confluence (error shown in blue) compared to composite (error show in red). SX represents subjects 1-5. For each subject we show the 95% confidence interval.

(Composite)   (Confluence)
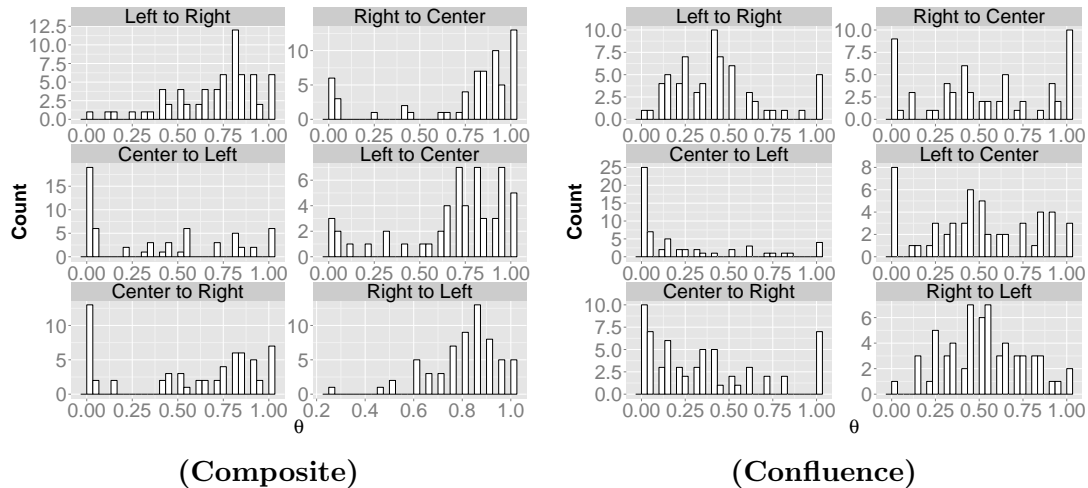
Figure 4.13. Histogram of the distribution of the selection of θ (as defined in section 4.4 and section 4.6) for all trials for subject 1. The left subfigure shows the θ distribution for the composite model, the right displays the θ distribution for the MDOC. The x axis of each histogram represents the bins of possible θ selections. The y axis represents the number of trials for which a corresponding θ value has been selected (the θ that yields the lowest RMSE for a given trial). Each subfigure has six subcomponents, each representing one of six movement types, which are illustrated in figure 2.2. The text above each histogram indicates the movement type related to that histogram, e.g. "Left to Right" corresponds to "A" to "C" in figure 2.2, "Center to Left" is "C" to "A", etc. See figure 4.14 for θ distributions for subjects 2-5.

These θ distributions tend to vary in shape for different movement types for a given subject, yet no discernible pattern emerges across subjects for a given movement type.

Some of the distributions have the characteristic of being heavy at both $\theta = 0$ and $\theta = 1$, others are heavy only at one of the two extremes. This is partially due to the phenomenon illustrated in section 4.1.1 regarding the variation of peak *velocity* and *shape* of velocity profiles across various minimized derivatives. A shift from $\theta = 0$ to $\theta = 1$ (jerk to crackle) changes the shape of the velocity curve and makes it more "skinny" and "peaky" as $\theta = 1$ is approached. Skews in θ distribution in either
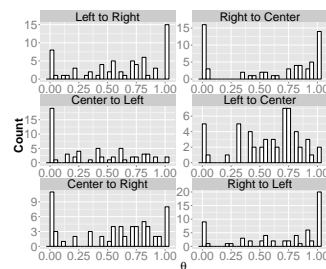
(Subj 2 Composite)
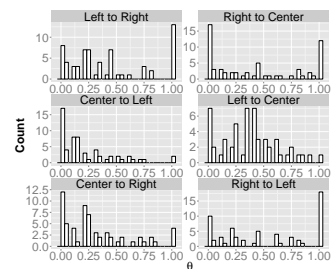
(Subj 2 Confluence)
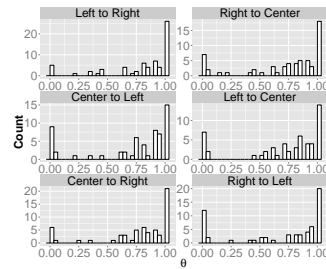
(Subj 3 Composite)
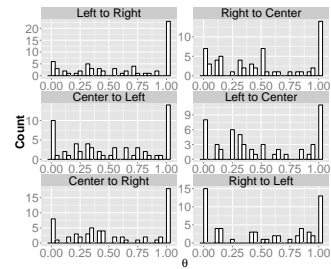
(Subj 3 Confluence)
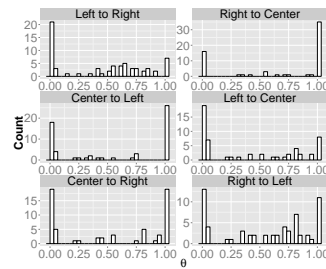
(Subj 4 Composite)
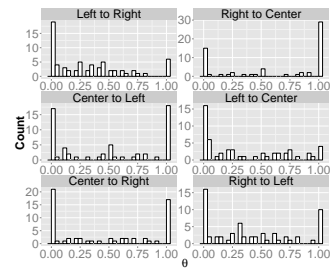
(Subj 4 Confluence)
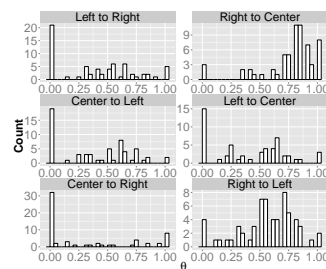
(Subj 5 Composite)
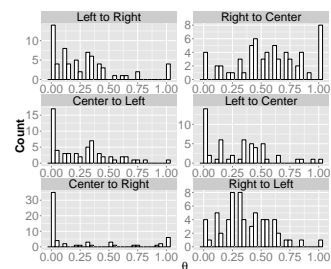
(Subj 5 Confluence)

Figure 4.14. Histogram of the distribution of the selection of θ (as defined in part 5 of section 4.4 and section 4.6) for all subjects and all trials. The left and right columns show composite and confluence (MDOC) respectively. The x axis of each histogram represents the bins of possible θ selections. The y axis represents the number of trials for which a corresponding θ value has been selected (the θ that yields the lowest RMSE for a given trial). This diagram follows the same pattern as figure 4.13

direction (or both) indicate that some trials for a particular subject/movement combination fall outside of the "bell shape range" of the jerk/crackle model tradeoff.

There are three main reasons for this to occur. First, for normative velocity profiles (standard bell shaped profiles), the trial is either too "skinny" and/or "peaky", even for a pure crackle model (i.e. $\theta = 1$). Subfigure "4.11a. thetas = 1", shows a trial in which $\theta = 1$ for both models, but the models are not skinny or peaky enough to match it. In the second case, a trial is too "shallow" and/or "fat" for a pure jerk model (i.e. $\theta = 0$). In subfigure "4.11b. thetas = 0", $\theta = 0$ for both models, but the trial is more shallow than the model output (this is not a common case).

The third main contributor to θ distributions that are heavy at the extremes is that nonnormative trials (those that deviate greatly from the standard bell curve shape) almost always gravitate towards the extreme ends of θ. The rest of the subfigures are examples of this. We highlight Subject 4 because modeling his/her movements yields the highest RMSE of all subjects (see figure 4.12). This indicates that this subject's velocity curves are nonnormative with respect to other subjects' profiles. Subject 4's trials have θ distributions heavily represented at 0 and 1, indicating that most of this subjects velocity curves fall outside of the "bell shape range" of the models. Additionally, subfigure "4.11f. thetas = 0" is particularly interesting, as it is a nonnormative velocity profile, and $\theta = 1$ for both models, but still the trial is perfectly fitted. This is an unusual case.

The other classes of θ distributions are either more heavy in the center, or more evenly spread across all θ. The distributions associated with subject 1, and to a lesser extent subjects 2 and 5 fit often fit these profile.

In order to consider the results absent the θ selection process (i.e. choosing one θ and keeping it for all trials, regardless of whether or not it represents the

lowest RMSE for that trial), figure 4.15 compares how RMSE changes with the confluence vs the composite approaches each θ values *individually*. We show the average RMSE across all trials for each θ value between 0 and 1 in increments of .05. The intervals within which the standard error bars overlap support the notion that, regardless of what θ value is selected, the models are not statistically different in terms of RMSE. However, two points are worthwhile to note. The first is that subjects 1 and 3 show large ranges within which there is no overlap between standard error bars. In these ranges, confluence's RMSE is lower than that of the composite model. The second point is that, when considering the point on the x axis (θ) that yields the lowest point on the y axis (RMSE) for each model (i.e. when choosing the optimal θ for each model), the lowest RMSE for confluence is always less than than the lowest RMSE for composite for all subjects. That is to say, when selecting the θ that yields the lowest average RMSE, the confluence approach yields an average RMSE minimum that is lower than the composite method's. This is also reflected in a more compact form in figure 4.12, in which average RMSE for confluence is always lower than that of composite.

Table 4.2. A comparison of the statistical differences in RMSE between the confluence and composite based models compared to human movent data for all subjects. The values represented in the tables correspond to the those reported in figure 4.12. The p values are arrived at via a nonparametric Kolomogrov-Smirnov test. In all cases, the two approaches are not significantly different in terms of RMSE.

Subject	p
Subject 1	0.3141994
Subject 2	0.5229139
Subject 3	0.9998817
Subject 4	0.9999999
Subject 5	0.1447257
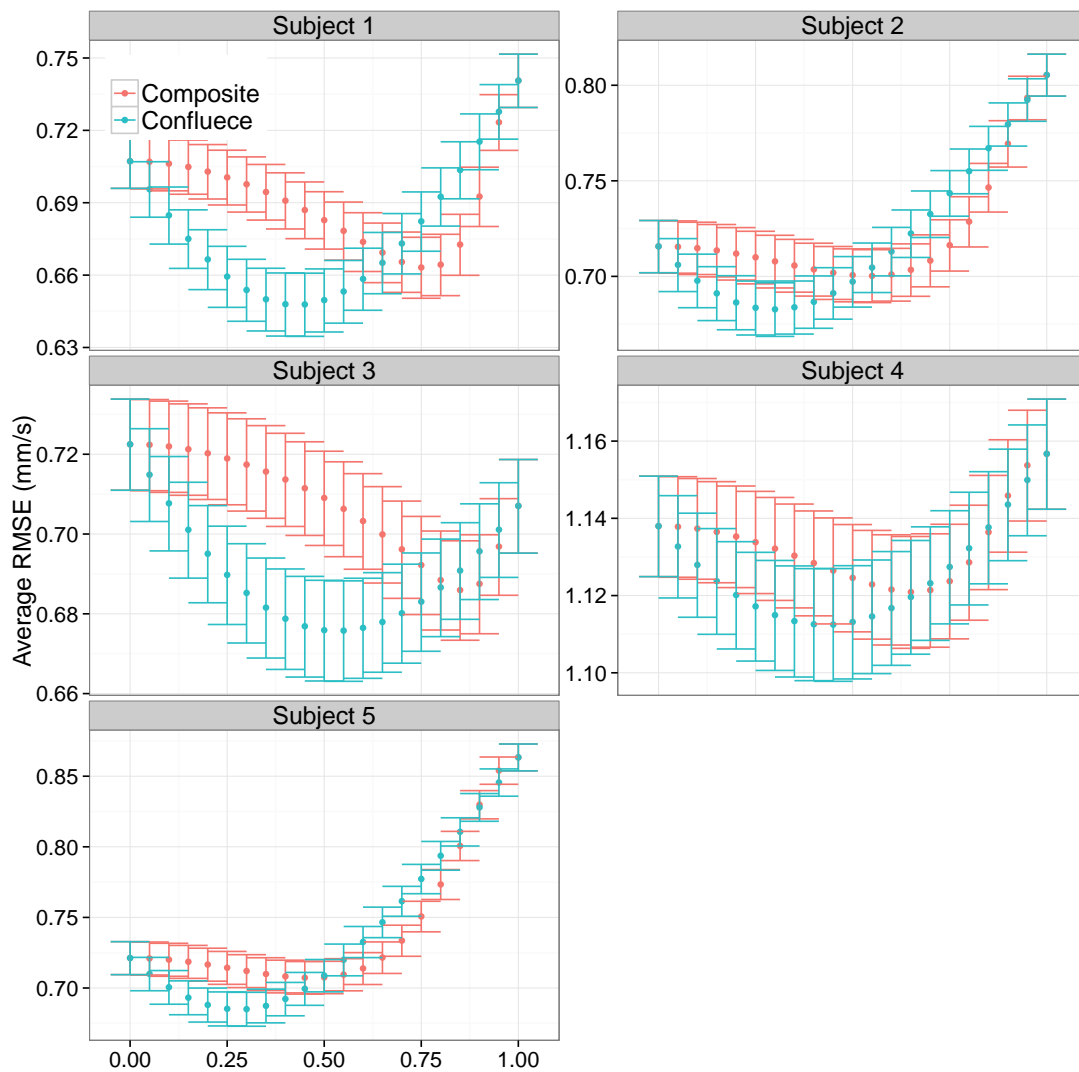
Figure 4.15. A plot of average RMSE across all trials of both MDOC and composite as θ progresses from 0 to 1 in increments of .05 (x axis). Each point represents the average RMSE across all trials for a given θ. Confluence is shown in blue, composite is shown in red. Standard error bars are shown.

4.7 Discussion

4.7.1 Anatomical Modeling Implications of the Confluence Approach

Let us consider the general confluence approach (not necessarily the MDOC) in relation to the anatomy. We first state that we do not claim the existence of signals representing jerk or crackle in the nervous system, therefore the MDOC does not have a direct mapping to the anatomy of the motor system. However, we note that the nervous system utilizes multiple signals for motor control that are up and down regulated separately. As seen in equation 4.2, the confluence approach sums multiple separate control signals, where the composite approach generates only one control signal, as shown in equation 4.1. Relative weightings of the separate confluence signals may arise from various amounts of inhibition imposed upon either the signal producer, or may result from presynaptic or postsynaptic inhibition on or near the target neural circuitry.

For example, the maintenance of balance requires a confluence of signals, including those representing disturbances in inertia and acceleration, along with vision, and signals from receptors in the muscles, tendons and skin. Additionally, it appears that the efficacy of each of these signals is up or down regulated depending on their relative congruence. If one or more of the signals involved in balance control becomes incongruent with the vestibular system's output, that incongruent signal is suppressed [103]. Hence, balance control, as well as many other neuromuscular functions, require the manipulation of many signals individually, as opposed to a single signal (as produced by a composite cost function). These concepts are not explored in this work, but future work might attempt to utilize the multi-signal approach of confluence as a way to more closely represent the nervous system.

4.7.2 Potential application to robotics

Controlling a robot in a way that mimics human movement is extremely difficult. There exists some set of algorithms to effectively mimic the breadth of human movement. For reasons mentioned in section 4.7.1, the confluence approach may be more effective than traditional control strategies in solving this problem. The confluence approach has the potential to directly relate to animal motor control signals. Robots, with multiple sensory and control signals, may move with human-like capabilities based upon the confluence approach.

4.7.3 The assumption of cost-optimality in the motor system

Optimal control theory is a well accepted technique used to model human movements. At the core of this approach is the assumption that human movements are in some way directly tied to cost-optimality principles. If this supposed cost-optimality was achieved by evolution, the assumption indicates that the human motor system has reached some type of zenith beyond which there is no room for improvement, at least within our current niche. In this work, the notion of suboptimality means that, in terms of minimizing cost, the confluence technique is suboptimal as compared to the composite technique. At least for the case of this feed forward control problem for fast reaching movements, the composite and confluence approaches are comparable in terms of RMSE when compared with actual reaching movements. These results indicate that human reaching movements do not necessarily adhere to cost-optimality principles.

Although the accuracy of each approach when compared to human movement is not statistically different, a clear trend is present. Figure 4.12 shows that for all subjects, the average RMSE across all trials for the MDOC is always lower

than that of the composite method's. This study was limited to five subjects. In order to validate this trend, more subjects and trials would be necessary. If this trend were to continue, it would offer strong evidence that the the assumption of cost-optimality of the motor system should be reexamined. Furthermore, for this set of subjects, the MDOC is more robust to an inexact (suboptimal) θ selection. Figure 4.15 shows that when spanning all θ values, the RMSE "bowl" plots are qualitatively deeper and wider for the MDOC. This means that, for a larger portion of the θ range, the MDOC has lower RMSE. This allows for a wider range of (in the parlance of Loeb [86]) "good enough" θ selections. Both of these RMSE related results indicate that more extensive models using the confluence approach may also perform as well as, or better than the composite approach.

Two cost functions were used in this work, minimized jerk and crackle. This choice is not arbitrary, as these cost functions are related to smoothness of movement, a family of costs shown to be more useful than other choices [11]. The confluence approach outlined in equation 4.2 is general enough to utilize any set of cost functions. However, it is unclear whether or not the confluence technique would yield comparable results for other types of cost functions (e.g. minimum torque change combined with an end point accuracy cost). It is also unclear how well this technique would generalize to other, more complex movement types. Future work should include an application of confluence to different movement types and/or utilizing different cost functions. Future work may also involve the inclusion of feedback.

We conclude that, for modeling reaching movements, the acceptance of the traditional organization of composite cost functions should be reexamined. The implications of this work and others that follow may lead us to question the assumption that the human motor system is "optimal" with respect to a set of

costs.

4.8   Acknowledgments

This chapter is an augmented reprint of the material that has been submitted to Neural Networks, Gamble, G.G.; Yazdani, M.; Hecht-Nielsen R. "Suboptimal combinations of cost functions fit fast human reaching movements as well as optimal combinations", Oct. 2015.

Chapter 5

Conclusions, conjectures, and scientific contributions

5.1 Anatomical and physiological justification for cost and control signal reorganization

The OCCF attempts to create a bridge between mathematical models of movement and biological motor systems. It does so by emphasizing control signals related to many cost functions, by allowing for signal operators to transform a signal or signals to mirror known neural operators (e.g. integration), and allowing for a flexible network architecture in which signals combine and cooperate, which is a phenomenon we call confluence and is witnessed throughout the nervous system. While any cost or reward function can fit into this framework, the instance of the framework demonstrated in this work (MDOC), emphasizes the role of the derivatives of position as neural motor control signals. This choice is not arbitrary, as these cost functions are related to smoothness of movement, a family of costs shown to be more useful than other choices [11]. Additionally, neurophysiological examples of such signals driving motor control are plentiful [1, 104, 105, 106, 107, 108, 109]. However, we do not assert that signals representing derivatives as high as those explored here are present in the nervous system.

5.1.1 Neural integration and the derivatives of position in motor control, a link between biology and mathematics

It is evident that the majority of animal species which have a sense of position, velocity, acceleration, jerk, 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 [110]. 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 cheetah's 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. In order to be effective, these neural representations must influence the motor systems of these animals. Animals without these abilities would have more difficulty navigating their environments in order to find food, escape predators, etc. Thus, the majority of surviving animal species likely have neural representations of dynamic variables related to position.

Neural signals related to body position *and its derivatives* are encoded in many areas of the vertebrate nervous system. Head direction cells encode for the current positional orientation of the head [111]. The vestibular system, the semicircular canals and the otolith organs all produce signals coding for derivatives of the position of the head [1, 104, 105, 106, 107], which initiate changes in eye position and body stability. Other examples include signals from muscle spindles which encode position and velocity of the limbs and body, and signals from Golgi tendon organs which encode force exerted by muscles (with an obvious relationship

to acceleration) [108, 109]. All of these signals converge upon spinal circuits and influence motor control.

An example of a neuronal system that can be reasoned about mathematically and demonstrates a wealth of cooperating control signals, thus an excellent candidate system to be modeled by using the OCCF, is the oculomotor system (control of the eyes). It consists of several subsystems, some of which use neural integration mechanisms [112, 113, 114, 115]. Here we choose to examine the vestibulo-ocular reflex (VOR), as all of its elements can be represented using the OCCF. The VOR adjusts the eye position in order that they remain locked onto a target when the head is moving, see figure 4.6a for an illustration of the circuit. The semicircular canals of the inner ear detect angular *acceleration* of the head, and the signals measured in the nerve emanating from these organs code for *velocity*, not acceleration [107]. 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 [1]. Robinson was one of the first to realize that integration (in the mathematical sense) must be occurring [116], as this is how the positional signal is calculated from a velocity signal and maintained to hold the new eye position. The existence of neural integration mechanisms in the oculomotor system is now well accepted [1, 117, 118]. 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 [112, 113, 115, 114]. This system uses signals representing position 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 [10, 57, 35, 55, 58, 62]. Again we see a connection between optimal control's mathematics and anatomy and physiology. However, models to date typically choose *one* of these variables to focus on without addressing the fact that signals of various types routinely work in concert.

As shown in figure 4.6a, in the vestibulo ocular system, signals representing velocity and position remain separate until finally combined at the point of the actuator (motor neurons/muscles). Separate signals are needed for proper motor functioning of the eye. Furthermore, the acceleration signals must remain separate, as they are sent via multiple pathways to other areas of the nervous system which may interpret or process them differently. It is an acceleration/inertial signal that motor structures in the brainstem and spinal cord expect in order to maintain balance [103], not a signal based on the amalgam of several costs in a multi-term cost function. Unlike the composite formulation found in equation 4.1, confluence's equation 4.2 supports this separation of control signals. This is not to say that mult-term cost functions are not useful, but instead that keeping multiple cost functions separate to be combined post-optimization may be more useful when envisioning or building a system or model inspired by the nervous system.

The example of the OCCF paradigm demonstrated in this work (MDOC) treats higher order derivatives of position as control signals. Any type of control signal can be used, however, if one chooses to use signals based on dynamic variables, then there is a direct linkage between the coding of known anatomical signals (mentioned above) and the model, and thus between physiology and mathematics. The nervous system has the capability to convert neural signals from one derivative order to another. 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. 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 [119], hypothesized neural integration in head direction cells [120, 121], integration in the oculomotor system [1] and both differentiation and integration in V2 [122]. In this way, the physiology of motor control signals can be mathematically expressed and reasoned about, and it is evident that representations of control signals similar to those used in this work exist. In the OCCF, *signal transformation operators* represent such neural structures that perform known mathematics, again linking mathematics with biology. In the MDOC example, the operation is a simple linear combination of two signals.

5.1.2 Confluence of neural signals

In addition to neural integration, the oculomotor system exhibits a confluence of neural signals working together to achieve a myriad of complex eye movements. For example, the otolith organs produce signals based on linear acceleration of the head. Similar to the semicircular canals in the VOR, the otolith organs also play a role in eye control [1, 123]. The VOR and the otolith organs, along with other oculomotor subsystems all act in concert 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 "bang-bang" like which originate from premotor neurons in the rostral midbrain to produce saccades [124]. 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 [124]. A single control signal, associated with a single cost function, is inadequate to control the myriad of eye movement behaviors which exist.

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 [125, 126]. Additionally, alpha motor neurons and other spinal circuitry receive signals from muscle spindles and Golgi tendon organs which directly and indirectly influence motor control [109].

The OCCF is designed to mirror neural systems such as these. If all of the cost function terms and resultant signals are combined into a larger, single cost function, all meaning and functionality contained in the *connectivity* between elements would be lost. If we wish to more closely model the complexity of the motor system, a paradigm akin to the OCCF would be helpful. Architecting such a system pushes the "top down" approach of optimal control modeling "down" a level closer to actual construction of the nervous system, as initially discussed in section 1.4.2 of the introduction. This approach can models individual components and connectivity. Observing the behaviour of the constructed system can potentially reveal insights into the inner workings of the nervous system.

5.2   Representation of evolutionary pressures

Optimal control cost/reward functions sometimes involve trade-offs between several (composite) criteria (e.g. find a policy to get to a destination as fast as possible while consuming the least amount of fuel). While a composite cost function

contains several criteria, it is still encapsulated by a single function. To the best of our knowledge, all applications of optimal control to human movements have consisted of a single proposed cost function (composite or otherwise) producing a single control signal. In contrast, biological systems have many interacting subsystems composed of disparate structures that may have coevolved together, each system forged by different cost or reward functions. It is reasonable to assume that elements of the motor system which evolved during different points in evolutionary history (e.g. spinal cord, cerebellum, brain stem motor nuclei, motor cortex, etc.) did so under different evolutionary pressures (cost functions). Our approach reflects this notion of neural systems that evolved under different evolutionary pressures separately to produce neural control signals of different types.

It is interesting to note that by its very formulation, in terms of cost, the confluence technique is sub-optimal as compared to the composite technique. Yet, at least for the limited case of this feed forward control problem of fast reaching movements, the composite and confluence approaches are comparable in every scenario. This indicates that more extensive models using the OCCF may also perform as well as the composite approach, while maintaining the advantages discussed in this work. In addition, this also raises the possibility that evolution may have arrived at a sub-optimal solution for a motor control paradigm. The OCCF allows the practitioner to model evolutionary processes related to motor control.

5.3 Conclusion

In this work, we altered some basic assumptions and practices typically used in optimal control for the modeling of human movements. Doing so gave us results

comparable to or better than the standard approach, yet the construction and nature of the resulting models have neuronal interpretations, where as with the standard approaches, the neural interpretation is lacking.

We changed the norm used to measure and penalize the control signal in a minimum effort control problem, using the L_∞ norm as opposed to the L_2 norm. We applied this technique to the classic minimum jerk model for reaching movements. This resulted in a highly accurate model, which utilizes a very simple control signal as compared to the traditional approach. We then expanded upon this concept by encoding this signal in terms of Dirac delta functions which resemble neural spikes, thereby bolstering the neural implementation of the model and making the case that, in principle, very simple signals can drive reaching movements with smooth velocity profiles. We argue that evolution would likely "favor" such a simple control strategy over a complicated one.

Additionally, we questioned the validity and usefulness of using a single, mulit-term cost function for modeling reaching movements. We introduced a general framework, the OCCF, that we believe has the power to more directly model the motor system in terms of anatomy and physiology. This framework introduces the idea of a network architecture containing multiple control signals based on multiple cost functions, whereas the traditional approach has no notion of a network and typically employs a single control signal related to a single cost function. We then demonstrate the efficacy of the framework with a model of human reaching movements which combines two separately optimized control signals, and produces highly accurate results. Finally, we give examples of the anatomy of the nervous system that show how modeling with the traditional approach is inadequate to properly represent the network architecture of the HMS, bolstering the case for using the OCCF to model components of the motor system.

Bibliography

[1] D. E. Angelaki, "Vestibulo-ocular reflex.," in *Encyclopedia of Neuroscience*, pp. 139 – 146, Oxford: Academic Press, 2009.

[2] D. Bertsimas and A. W. Lo, "Optimal control of execution costs," *Journal of Financial Markets*, vol. 1, no. 1, pp. 1–50, 1998.

[3] H. Shen and P. Tsiotras, "Time-optimal control of axisymmetric rigid spacecraft using two controls," *Journal of Guidance, Control, and Dynamics*, vol. 22, no. 5, pp. 682–694, 1999.

[4] Q. Chai, C. Yang, K. L. Teo, and W. Gui, "Optimal control of an industrial-scale evaporation process: sodium aluminate solution," *Control Engineering Practice*, vol. 20, no. 6, pp. 618–628, 2012.

[5] N. Fligge, J. McIntyre, and P. van der Smagt, "Minimum jerk for human catching movements in 3d," in *Biomedical Robotics and Biomechatronics (BioRob), 2012 4th IEEE RAS & EMBS International Conference on*, pp. 581–586, IEEE, 2012.

[6] N. A. Bernstein, "The co-ordination and regulation of movements," 1967.

[7] C. Chow and D. Jacobson, "Studies of human locomotion via optimal programming," *Mathematical Biosciences*, vol. 10, no. 3, pp. 239–306, 1971.

[8] H. Hatze and J. Buys, "Energy-optimal controls in the mammalian neuromuscular system," *Biological cybernetics*, vol. 27, no. 1, pp. 9–20, 1977.

[9] D. Davy and M. Audu, "A dynamic optimization technique for predicting muscle forces in the swing phase of gait," *Journal of biomechanics*, vol. 20, no. 2, pp. 187–201, 1987.

[10] T. Flash and N. Hogan, "The coordination of arm movements - an experimentally confirmed mathematical-model.," J. Neurosci., vol. 5, no. 7, pp. 1688–1703, 1985.

[11] 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.

[12] E. T. Whittaker, "Galvanim, from galvani to ohm," in *A history of the theories of aether and electricity. Vol 1*, pp. 67–71, London ; New York : Longmans, Green: Nelson, 1951.

[13] G. Fritsch and E. Hitzig, "Ueber die elektrische erregbarkeit des grosshirns.," Archivfur Anatomie Physiologie, und wissenschaftliche Medicin, pp. 308–314, 1870.

[14] G. Fritsch and E. Hitzig, "The electrical excitability of the cerebrum," *Neurosurgical classics*, pp. 15–27, 1870.

[15] W. Lennon, R. Hecht-Nielsen, and T. Yamazaki, "A spiking network model of cerebellar purkinje cells and molecular layer interneurons exhibiting irregular firing," *Frontiers in computational neuroscience*, vol. 8, 2014.

[16] A. L. Hodgkin and A. F. Huxley, "A quantitative description of membrane current and its application to conduction and excitation in nerve," *The Journal of physiology*, vol. 117, no. 4, pp. 500–544, 1952.

[17] X. Konstantoudaki, A. Papoutsi, K. Chalkiadaki, P. Poirazi, and K. Sidiropoulou, "Modulatory effects of inhibition on persistent activity in a cortical microcircuit model," *Frontiers in neural circuits*, vol. 8, 2014.

[18] T. Yamazaki and S. Nagao, "A computational mechanism for unified gain and timing control in the cerebellum," *PloS one*, vol. 7, no. 3, p. e33319, 2012.

[19] E. Todorov, "Optimal control theory," in *Bayesian Brain, probabilistic approaches to neural coding.* (K. Doya, S. Ishii, A. Pouget, and R. P. N. Rao, eds.), Computational Neuroscience, ch. 12.7, Cambridge, Mass.: MIT Press, 2007.

[20] 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.

[21] S. Grillner and T. M. Jessell, "Measured motion: searching for simplicity in spinal locomotor networks," *Current Opinion in Neurobiology*, vol. 19, no. 6, pp. 572–586, 2009.

[22] 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.

[23] R. Lemon, "The output map of the primate motor cortex," *Trends in Neurosciences*, vol. 11, no. 11, pp. 501–506, 1988.

[24] 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.

[25] 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.

[26] 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.

[27] 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.

[28] K. Pearson, "Control of walking.," Sci. Am., 1976.

[29] 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.

[30] C. B. Hart and S. F. Giszter, "A neural basis for motor primitives in the spinal cord.," J. Neurosci., vol. 30, no. 4, pp. 1322–1336, 2010.

[31] 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.

[32] E. Todorov and M. I. Jordan, "Smoothness maximization along a predefined path accurately predicts the speed profiles of complex arm movements.," J. Neurophysiol., vol. 80, no. 2, pp. 696–714, 1998.

[33] W. Abend, E. Bizzi, and P. Morasso, "Human arm trajectory formation," *Brain*, vol. 105, no. Pt 2, pp. 331–48, 1982.

[34] 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.

[35] 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 Comput., vol. 20, no. 3, pp. 779–812, 2008.

[36] D. E. Kirk, *Optimal control theory : an introduction.* Mineola, N.Y.: Dover Publications, 2004.

[37] 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.

[38] K. Kyriakopoulos and G. Saridis, "Minimum jerk path generation," in *Robotics and Automation, 1988. Proceedings., 1988 IEEE International Conference on*, pp. 364 –369 vol.1, apr 1988.

[39] 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*, pp. 1668 –1674, oct. 2010.

[40] A. Piazzi and A. Visioli, "Global minimum-jerk trajectory planning of robot manipulators," *Industrial Electronics, IEEE Transactions on*, vol. 47, pp. 140 –149, feb 2000.

[41] A. Karniel and F. A. Mussa-Ivaldi, "Does the motor control system use multiple models and context switching to cope with a variable environment?," Exp. Brain Res., vol. 143, no. 4, pp. 520–524, 2002.

[42] L. Botzer and A. Karniel, "A simple and accurate onset detection method for a measured bell-shaped speed profile," Front. Syst. Neurosci., vol. 3, no. 61, pp. 1–8, 2009.

[43] G. H. Staude, "Precise onset detection of human motor responses using a whitening filter and the log-likelihood-ratio test," IEEE Trans. Biomed. Eng., vol. 48, no. 11, pp. 1292–1305, 2001.

[44] G. H. Staude, W. M. Wolf, U. Appel, and R. Dengler, "Methods for onset detection of voluntary motor responses in tremor patients," IEEE Trans. Biomed. Eng., vol. 43, no. 2, pp. 177–188, 1996.

[45] 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.

[46] 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.

[47] 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.

[48] 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.

[49] 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.

[50] E. Izhikevich, *Dynamical Systems in Neuroscience: Chapter 9*. Cambridge, M.A.: MIT Press, 2007.

[51] M. Goldman, A. Compte, and X. Wang, "Neural integrator models," in *Encyclopedia of Neuroscience* (E. in Chief: Larry R. Squire, ed.), pp. 165–178, Oxford: Academic Press, 2009.

[52] D. A. Robinson, "Integrating with neurons," *Annual Review of Neuroscience*, vol. 12, pp. 33–45, 1989.

[53] 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.

[54] S. Edelman and T. Flash, "A model of handwriting," *Biological Cybernetics*, vol. 57, pp. 25–36, 1987.

[55] Y. Uno and M. Kawato, "Formation and control of optimal trajectory in human multijoint arm movement.," Biol. Cybern., vol. 61, no. 2, pp. 89–101, 1989.

[56] B. Hoff and M. Arbib, "Models of trajectory formation and temporal interaction of reach and grasp.," J. Mot. Behav, vol. 25, pp. 175–192, 1993.

[57] C. Harris and D. Wolpert, "Signal-dependent noise determines motor planning.," Nature, vol. 394, no. 6695, pp. 790–784, 1998.

[58] D. Liu and E. Todorov, "Evidence for the flexible sensorimotor strategies predicted by optimal feedback control.," J. Neurosci., vol. 27, pp. 9354–9368, 2007.

[59] L. W. Neustadt, "Minimum effort control systems," *Journal of the Society for Industrial & Applied Mathematics, Series A: Control*, vol. 1, no. 1, pp. 16–31, 1962.

[60] B. Lim, S. Ra, and F. C. Park, "Movement primitives, principal component analysis, and the efficient generation of natural motions," in *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, pp. 4630–4635, IEEE, 2005.

[61] M. Kawato, Y. Maeda, Y. Uno, and R. Suzuki, "Trajectory formation of arm movement by cascade neural network model based on minimum torque-change criterion," *Biological Cybernetics*, vol. 62, pp. 275–288, 1990.

[62] M. Yazdani, G. Gamble, G. Henderson, and R. Hecht-Nielsen, "A simple control policy for achieving minimum jerk trajectories.," Neural Netw., vol. 27, pp. 74–80, 2012.

[63] R. Leib and A. Karniel, "Minimum acceleration with constraints of center of mass: a unified model for arm movements and object manipulation.," *Journal of Neurophysiology*, vol. 108, no. 6, pp. 1646–55, 2012.

[64] A. Karniel, "Open questions in motor control," *Journal of Integrative Neuroscience*, vol. 10, no. 3, p. 385411, 2011.

[65] P. Gawthrop, I. Loram, M. Lakie, and H. Gollee, "Intermittent control: A computational theory of human control.," *Biological Cybernetics*, vol. 104, no. 1-2, pp. 31–51, 2011.

[66] P. Gawthrop, I. Loram, H. Gollee, and M. Lakie, "Intermittent control models of human standing: similarities and differences," *Biological cybernetics*, vol. 108, no. 2, pp. 159–168, 2014.

[67] L. Pontryagin, V. Boltyanskii, R. Gamkrelidze, and E. Mishchenko, *The Mathematical Processes of Optimal Processes*. New York: Interscience Publishers, 1962.

[68] A. Feld'baum, *Mathematics in Science and Engineering*. New York: Academic Press, 1965.

[69] 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.

[70] P. Dayan and L. Abbott, *Theoretical Neuroscience, chapter 1.2.*

[71] S. Giszter, "Motor primitives," in *Encyclopedia of Neuroscience* (E. in Chief: Larry R. Squire, ed.), pp. 1023–1040, Oxford: Academic Press, 2009.

[72] A. Karniel, "The minimum transition hypothesis for intermittent hierarchical motor control," *Frontiers in Computational Neuroscience*, vol. 7, no. 12, 2013.

[73] J. Emken, R. Benitez, A. Sideris, J. Bobrow, and R. DJ., "Motor adaptation as a greedy optimization of error and effort.," *J Neurophysiol.*, vol. 97, no. 6, pp. 3997–4006, 2007.

[74] S. W. Keele, "Movement control in skilled motor performance," *Psychological bulletin*, vol. 70, no. 6p1, p. 387, 1968.

[75] V. Gerdes and R. Happee, "The use of an internal representation in fast goal-directed movements, a modelling approach," *Biological cybernetics*, vol. 70, no. 6, pp. 513–524, 1994.

[76] A. Karniel and G. F. Inbar, "A model for learning human reaching movements," *Biological cybernetics*, vol. 77, no. 3, pp. 173–183, 1997.

[77] M. Grant and S. Boyd, "Cvx: Matlab software for disciplined convex programming, version 1.21 (2011)," *Available: cvxr. com/cvx*, 2011.

[78] I. D. Loram, H. Gollee, M. Lakie, and P. J. Gawthrop, "Human control of an inverted pendulum: Is continuous control necessary? is intermittent control effective? is intermittent control physiological?," *The Journal of Phisiology*, vol. 589, no. 2, p. 307324, 2010.

[79] P. Redgrave, T. J. Prescott, and K. Gurney, "The basal ganglia: a vertebrate solution to the selection problem?," *Neuroscience*, vol. 89, no. 4, pp. 1009–1023, 1999.

[80] P. Cisek and J. F. Kalaska, "Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action," *Neuron*, vol. 45, no. 5, pp. 801–814, 2005.

[81] Y. Loewenstein, S. Mahon, P. Chadderton, K. Kitamura, H. Sompolinsky, Y. Yarom, and M. Husser, "Bistability of cerebellar purkinje cells modulated by sensory stimulation.," *Nature Neuroscience*, vol. 8, no. 2, pp. 202–11, 2005.

[82] P. M. DiLorenzo and J. D. Victor, *Spike Timing: Mechanisms and Function.* 6000 Broken Sound Parkway NW, Suite 300: CRC Press, Taylor and Francis Group, 2013.

[83] W. Gerstner, A. K. Kreiter, H. Markram, and A. V. Herz, "Neural codes: Firing rates and beyond," *Proceedings of the National Academy of Sciences*, vol. 94, no. 24, pp. 12740–12741, 1997.

[84] K. E. Jones, A. F. Hamilton, and D. M. Wolpert, "Sources of signal-dependent noise during isometric force production.," *Journal of Neurophysiology*, vol. 88, no. 3, p. 15331544, 2002.

[85] A. F. Hamilton, K. E. Jones, and D. M. Wolpert, "The scaling of motor noise with muscle strength and motor unit number in humans.," *Experimental Brain Research*, vol. 157, no. 4, p. 417430, 2004.

[86] L. GE., "Optimal isn't good enough," Biol Cyber, vol. 106, no. (11-12, pp. 757–765, 2012.

[87] K. Friston, "What is optimal about motor control?," *Neuron*, vol. 72, no. 3, pp. 488–498, 2011.

[88] Y. Ueyama, "Effects of cost structure in optimal control on biological arm movement: A simulation study," in *Neural Information Processing*, pp. 241–248, Springer, 2013.

[89] R. Shadmehr and S. Wise, *The Computational Neurobiology of Reaching and Pointing: A Foundation for Motor Learning.* Computational Neuroscience, *MIT Press*, 2005. Found in online appendix.

[90] M. E. Richardson and T. Flash, "Comparing smooth arm movements with the two-thirds power law and the related segmented-control hypothesis.," J. Neurosci., vol. 15, no. 2, pp. 8201–8211, 2002.

[91] A. Wiegner and M. Wierzbicka, "Kinematic models and human elbow flexion movements: quantitative analysis.," Exp. Brain Res., vol. 88, no. 3, pp. 665–673, 1992.

[92] R. Plamondon, A. Alimi, P. Yergeau, and F. Leclerc, "Modelling velocity profiles of rapid movements: a comparative study.," Biol. Cybern., vol. 69, no. 2, pp. 119–128, 1993.

[93] S. Edelman and T. Flash, "A model of handwriting.," Biol. Cybern., vol. 57, pp. 25–36, 1987.

[94] F. A. M.-I. Jonathan B. Dingwell, Christopher D. Mah, "Experimentally confirmed mathematical model for human control of a non-rigid object," Journal of Neurophysiology, vol. 91, no. 3, pp. 1158–1170, 2004.

[95] A. Biess, D. G. Liebermann, and T. Flash, "Comparing smooth arm movements with the two-thirds power law and the related segmented-control hypothesis.," J. Neurosci., vol. 27, no. 48, pp. 13045–13064, 2007.

[96] T. Krasovsky, S. Berman, and D. Liebermann, "Kinematic features of continuous hand reaching movements under simple and complex rhythmical constraints.," J. Electromyogr. Kinesiol., vol. 20, no. 4, pp. 636–641, 2010.

[97] B. S. Liebermann D.G., Krasovsky T., "Planning maximally smooth hand movements constrained to nonplanar workspaces.," J. Mot. Behav., vol. 40, no. 6, pp. 516–31, 2008.

[98] M. Berniker and K. Kording, "Estimating the sources of motor errors for adaptation and generalization.," Nat. Neurosci., vol. 11, pp. 1454 – 1461, 2008.

[99] R. Shadmehr and F. A. Mussa-Ivaldi, "Adaptive representation of dynamics during learning of a motor task," *The Journal of Neuroscience*, vol. 14, no. 5, pp. 3208–3224, 1994.

[100] Y. Ueyama, "Mini-max feedback control as a computational theory of sensorimotor control in the presence of structural uncertainty," *Frontiers in computational neuroscience*, vol. 8, 2014.

[101] D. M. Wolpert and M. Kawato, "Multiple paired forward and inverse models for motor control.," Neural Netw., vol. 11, no. special issue, p. 13171329, 1998.

[102] M. Yazdani, *Fast human movements and sparse optimal control policies*. PhD thesis, University of California, San Diego, 2012.

[103] S. Slobounov and K. Newell, "Balance and posture control: Human," in *Encyclopedia of Neuroscience*, p. 3135, Oxford: Academic Press, 2009.

[104] 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. Neurophysio.l, vol. 39, p. 970984, 1976.

[105] C. Fernandez and J. Goldberg, "Physiology of peripheral neurons innervating otolith organs of the squirrel monkey. iii. response dynamics.," J. Neurophysiol., vol. 39, p. 9961008, 1976.

[106] C. Fernandez and J. Goldberg, "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.

[107] C. Fernandez and J. Goldberg, "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.

[108] M. P. Mileusnic and G. E. Loeb, "Force estimation from ensembles of golgi tendon organs.," J. Neural Eng., vol. 6, no. 3, 2009.

[109] T. R. Nichols, "Reflex circuits.," in *Encyclopedia of Neuroscience* (E. in Chief: Larry R. Squire, ed.), pp. 73–79, Oxford: Academic Press, 2009.

[110] T. Matheson, "Motor planning: Insects do it on the hop.," Curr. Biol., vol. 18, pp. 742–743, 2008.

[111] J. Taube, "Head direction cells.," Scholarpedia, vol. 4, no. 12, p. 1787, 2009.

[112] E. Aksay, I. Olasagasti, B. D. Mensh, R. Baker, M. S. Goldman, and D. W. Tank, "Functional dissection of circuitry in a neural integrator.," Nat. Neurosci., vol. 10, no. 4, pp. 494–504, 2007.

[113] H. Seung, "How the brain keeps the eyes still.," Proc. Natl. Acad. Sci. U.S.A., vol. 93, p. 1333913344, 1996.

[114] K. Fukushima, C. Kaneko, and A. Fuchs, "The neuronal substrate of integration in the oculomotor system.," Prog. Neurobiol., vol. 39, no. 6, pp. 609 – 639, 1992.

[115] K. Fukushima and C. Kaneko, "Vestibular integrators in the oculomotor system.," Neurosci. Res., vol. 22, no. 3, pp. 249 – 258, 1995.

[116] D. A. Robinson, "Integrating with neurons.," Annu. Rev. Neurosci., vol. 12, pp. 33–45, 1989.

[117] L. DellOsso, *Chapter 2: Neural Integration In Ocular Motility in: Neurological organization of ocular movement.* Berkeley, CA, U.S.A: Kugler Publications, 1990.

[118] M. Goldman, A. Compte, and X. Wang, "Neural integrator models.," in *Encyclopedia of Neuroscience* (E. in Chief: Larry R. Squire, ed.), pp. 165–178, Oxford: Academic Press, 2009.

[119] 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.

[120] H. Blair and P. Sharp, "Anticipatory head direction signals in anterior thalamus: evidence for a thalamocortical circuit that integrates angular head motion to compute head direction.," J. Neurosci., vol. 15, p. 62606270, 1995.

[121] J. S. Taube and J. P. Bassett, "Persistent neural activity in head direction cells.," Cereb. Cortex , vol. 13, no. 11, pp. 1162–1172, 2003.

[122] A. M. Schmid, K. P. Purpura, I. E. Ohiorhenuan, F. V. Mechler, and D. Jonathan, "Subpopulations of neurons in visual area v2 perform differentiation and integration operations in space and time.," Front. Syst. Neurosci., vol. 3, no. 15, 2009.

[123] 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.

[124] D. L. Sparks, "The brainstem control of saccadic eye movements.," Nat. Rev. Neurosci., vol. 3, no. 12, pp. 952–964, 2002.

[125] R. Burke, "Motor units: Anatomy, physiology and functional organization.," in The Nervous System (V. Brooks, ed.), vol. 1, ch. 1, pp. 345–422, 1981.

[126] R. Lemon, "Descending pathways in motor control.," Annu. Rev. Neurosci., vol. 31, pp. 195–218, July 2008.