Linearity, Motor Primitives and Low-Dimensionality in the Spinal Organization of Motor Control

by

Max Berniker

M.S.M.E., Massachusetts Institute of Technology, 2000
B.M.E., University of California, San Diego, 1998

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2005

Author .....
Department of Mechanical Engineering
August 23, 2005

Certified by .....
Emilio Bizzi
Institute Professor
Thesis Supervisor

Accepted by .....
Lalit Anand
Professor of Mechanical Engineering
Chairman, Committee on Graduate Students

Linearity, Motor Primitives and Low-Dimensionality in the Spinal Organization of Motor Behavior

by

Max Berniker

Submitted to the Department of Mechanical Engineering
on August 23, 2005, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Mechanical Engineering

Abstract

The typical biological system is nonlinear, high-dimensional and highly redundant, all of which are burdens on controller design. Yet despite these complications, the central nervous system is able to control motor systems with an impressive level of complexity and effectivity. One such example is the frog. Evidence suggests that in frogs, the central nervous system, and the spinal cord in particular, may adopt simplifying strategies to ease the motor control problem. For instance, despite the known nonlinear nature of muscle, it has been demonstrated experimentally that spinally induced force production in the frog limb is linear in stimulation. Spinally encoded force fields have also been implicated as building blocks for generating hind limb movements. Furthermore, muscle EMG measurements for both intact and spinalized animals, have been shown to be low-dimensional; these measurements can be reconstructed as linear combinations of fixed muscle activations, or synergies. The evidence above suggests that the central nervous system may adopt simplifying strategies for the motor control problem.

First, the thesis addresses the issue of linearity in isometric force fields. It proposes that this behavior can be explained as a result of biomechanical properties. To this end, a physiologically realistic model of the frog hind limb is analyzed. The results suggest that, due to features of the musculo-skeletal structure, forces produced by the hind limb muscles are linear in activation, in large part and within the limb's workspace. The results, therefore, support our hypothesis that muscle forces which scale linearly in activation are a natural biomechanical result.

The second portion of the thesis centers on the evidence of low-dimensional motor commands and the hypothesized motor primitives (in the form both of force fields and of muscle synergies). Many investigations have examined muscle synergies, probed motor behaviors for modular features in the form of force fields, and looked for connections between synergies and force fields. However, this work has largely been descriptive in nature, trying to explain the data without reference to the underlying control structure. We offer a principled explanation for motor primitives, for how force fields and synergies arise, and for how they are implicated in the organization of motor control. A controller that utilizes a reduced order model is proposed. Using apparatuses drawn from model order reduction theory, a method for finding a low-dimensional model that estimates a nonlinear model of the frog hind limb is examined.

A formalism for defining motor primitives is proposed and the resulting primitives are compared with experimentally derived synergies. The motor primitives are found to correspond well with several synergies, and to offer practical interpretations in terms of limb biomechanics. The reduced model is shown to be capable of generating natural motor behaviors as well as optimal control solutions. The evidence suggests that frog hind limb motor behaviors, and the spinal circuitry that coordinates these behaviors, are consistent with a control architecture that utilizes a reduced order model of the musculo-skeletal system in an effort to simplify motor control.

Thesis Supervisor: Emilio Bizzi
Title: Institute Professor

Acknowledgments

I have been at M.I.T. for what seems like a long, long time now. During this time it has been my priveledge and honor to interact with many people, many friends, many labs, and just as many advisors. And, if there are words to express my sincerest appreciation and gratitude, I don't know them; so I'll just say, *thank you all.*

for mom and dad and bill

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The goals of this thesis are relatively simple and tangible. The difficulties of motor control are extensive and formidable. The field of motor control strives to understand how animals, including humans, control their bodies; how they locomote, swim, jump, fly, reach, grasp and generally interact with their environment. This is, in a nutshell, the motor control problem. From a control engineer's perspective the motor control problem is exceedingly hard: a typical biological system is nonlinear, high-dimensional and highly redundant, and while separately each of these characteristics is a burden for controller design, when combined the motor control problem is nearly overwhelming. This leaves one wondering - how does the central nervous system do it? Techniques from many varied fields can be brought into play when examining this problem: machine learning, neuroscience, nonlinear dynamical systems theory, control engineering, and even mechanical design are fields that can contribute to the investigation of motor control. The theoretical and computational demands for such an investigation are significant. Indeed, an argument could be made that solving the motor control problem, and understanding how the central nervous system commands the body, is AI complete[1].

We will parse the complexities of motor control in three categories: nonlinearity, high-dimensionality and redundancy. While discussing these categories we will highlight the specific difficulties of biological motor control that the thesis addresses. We will then introduce methods for alleviating these difficulties, present evidence to suggest that the central nervous system of some animals may be adopting similar strategies. In doing so, we introduce the thesis objectives: to explore simplifying strategies that the central nervous system may adopt to alleviate the difficulties of biological motor control.

1.1 The Difficulties of Motor Control

Complete descriptions of nonlinear systems (that is, for arbitrary initial conditions, for all times and arbitary inputs) are difficult to find. Descriptions for linear systems, on the other hand, are relatively simple to find. As long as the system inputs are well

[1] An indication that the solution presupposes a solution to the strong AI problem.

behaved (i.e., continuous and bounded), descriptions of the system, or the solutions, are easy to obtain. And even when the inputs are not well behaved, solutions may be obtained. The solutions for linear systems can be analyzed to better understand the system and its dynamics; this ability to analyze and understand linear systems dynamics is a great asset when designing a controller.

For nonlinear systems, solutions and system analysis are much more difficult. An example of such a nonlinear system is a multi-link limb, like the models commonly used to describe an arm or leg. This system has been widely analyzed and examined, and is well understood (e.g. [22], [63], [62]). However, no solutions to the differential equations that define them are known. As a result, the path a multi-link limb takes when arbitrary torques are applied to its joints is not known; this path can easily be simulated, but it is not known in advance. Because an analytic solution for this nonlinear system is lacking, the resulting motion of the limb is difficult to assess without the aid of a computer. Similarly, what the joint torques ought to be to produce a specific motion is difficult to predict. Yet this is exactly what we aim to do when our goal is to control the system.

The muscles that drive the limb are nonlinear as well. Muscle properties have nonlinear dependencies both on their mechanical state (e.g., position and velocity) and, as the models used to describe this behavior suggest, on the motor commands that excite them (e.g., [9], [71]). Therefore, the force a muscle produces changes in a nonlinear manner as it lengthens, possibly complicating attempts to apply a constant torque about a joint while the limb is moving. The neurons of the central nervous system, including those of the motor neuron pools that command muscles, are also nonlinear in their dynamics.

When all these aspects of the motor apparatus–the limb mechanics, the dynamics of muscles and the neurons that command them–are taken into conisderation, the nonlinearities render the system considerably difficult to analyze and control.

Another source of complexity in biological motor control is the high system dimensionality. By dimensionality, we are referring to the number of variables or states required to completely describe a system. Systems with many dimensions or states can burden the issues of solution determination, analysis, and control mentioned above; analyzing a system with one state is, generally speaking, far easier than analyzing a system with twenty states. By definition, more dimensions allow a system to vary in more independent ways, complicating descriptions that characterize the system, and in turn complicating attempts to control that system in a predictable manner.

Returning to the example of the limb, the number of states, or dimensions, for this system may be managable depending on the degrees of freedom specified. Adding another link, perhaps to model a hand or foot, adds more states. When muscles are included in the model more states may be needed to describe their dynamics as well. If more states were added, for example to include the effects of motor neurons, then the system dimension further increases. At this point the number of states may be high. This need to take into account a large number of states, that can independently influence system dynamics, is a significant complication. Not only does this complicate predictions of system responses, it makes analysis of said responses arduous. It is precisely these complications that make control design difficult.

Finally, system redundancy can also be a complicating factor in motor control too. When two or more variables are redundant, they have similar effects on some characteristic of the system, therefore there may be no clear indication how one should be treated respective of the other. That is, if they both do the same thing, how do you treat them differently to achieve the same goal? With the limb, there are often redundancies. Redundant kinematics can imply multiple choices for positioning the limb while achieving the same task. For instance, while pointing at a fixed target location, the elbow and wrist can be in many positions. How does one choose amongst these positions?

Many muscles have similar effects on the limb as well. There are many ways to excite the muscles to achieve the same mechanical output. When performing a reaching movement, the trajectory the limb follows can be achieved with many commands to the same muscles. For example, consider two movements, one made with little effort, and another made while co-contracting the limb's muscles. These two movements may be indistinguishable in terms of their kinematics, yet their motor commands differ.

The neurons that drive the muscles have even more redundancy. In humans, for example, a large muscle can have a hundred or more motor nuerons commanding it [34]. A typical human muscle may have 100-500 spindle fibers [46],[52], and 10-20 Golgi tendon organs [21] relaying information back to the spinal cord. All these redundancies amount to states or variables that, although, ought to be prescribed to achieve a motor goal, may not be uniquely constrained by the motor goals; they are free variables in designing the task; choices that must be made but that do not directly influence the desired outcome. These free variables may be a benefit (for instance, if co-contraction was desired) or a deficit (they represent an increased computational complexity), but they are a complication.

Now that we have identified sources of difficulty, what can be done to alleviate them? There are strategies that specifically strive to alleviate nonlinearities. We shall briefly mention a few below. There are also techniques that alleviate the difficulties associated with high-dimensionality and redundancies. The techniques we shall present do not necesarily distinguish between the difficulties due to high-dimensionality and redundancy, but they alleviate both problems.

1.2 Simplifying Strategies for Motor Control

As we have tried to explain, nonlinearity clearly can be an unwanted complication. What can be done to simplify its influences on the control problem? One possibility is to incorporate a feedback mechanism. If system variables are "fed back" to a controller, commands can be computed to effectively cancel unwanted dynamics, or nonlinearities. Feedback controllers can bias the system dynamics in favor of the feedback path over those of the forward path. If unwanted dynamics need to be completely eliminated then there is a formalism for this as well: feedback linearization [63]. When the appropriate conditions are met, a feedback controller can be designed to linearize system dynamics. If this is done, the system behavior can more easily be

understood and analyzed.

A popular methodology for overcoming the complications of controller design for high-dimensional and redundant systems is to use an optimal strategy. By including a cost function in the motor goal, especially one that is influenced by redundant variables, unique choices for motor commands and states can be obtained by minimizing (or maximing) the cost function. For instance, in the example mentioned earlier a reaching movement can be made with or without co-contraction. By minimzing the total metabolic work in the movement, motor commands with unnecesary co-contraction, and, therefore unnecesary metablic cost, are unacceptable. Thus by finding the optimal solution, we have isolated a unique choice for motor commands amongst the many possibilites.

Optimal solutions, when found, can yield a command signal or a controller. However, for nonlinear systems optimal solutions are difficult to obtain. For high dimensional nonlinear systems, optimal solutions are even more difficult to obtain. This is the so-called curse of dimensionality [4]. Yet, when applicable, they are a great advantage in simplifying motor control.

To alleviate the problems of high-dimensionality, perhaps the easiest option is to simply neglect some of the dimensions. This is done by neglecting or ignoring the dependence on particular states. For instance, if some states do not vary much, or their variation has little impact on the system, then perhaps they can be ignored altogether. Similarly, one could neglect states that have fast dynamics. Assuming these states aren't of interest, they can be neglected and the remaining, relatively slow dynamics of the system can be studied.

Another possibility for reducing redundancy and dimensions is to group together variables that have similar effects on the motor system, and treat them as a single unit. This is the idea behind motor primitives, which are compact representations for a group of covarying motor variables. By grouping redundancies together as primitives, the number of independent variables needed for the motor control problem is reduced, and the difficulties of redundancies are simplified [58]. A similar idea is to seek a reduced set of variables that compactly represents the system of interest. If these variables can be used to make predictions of the system dynamics, then they may be useful for designing control as well. And, given the reduced dimension of these variables, they may simplify the controller design. Proposals for how to learn such a reduced order variable, in a biological system, have been suggested [65], [42]. However, these proposed methods for finding reduced order variables are based on learning algorithms, and as such, a clear interpretation of the these variables may be lacking.

There is a branch of control theory that deals formally with attempts reduce system complexity (high dimensionality and redundancy). Model order reduction is a methodology which studies properties of dynamical systems in the attempt to reduce their complexity, while preserving their input-output behavior. For linear systems, the field is well developed (e.g. [2], [3]). System complexity can be decreased by reducing the system order or number of states. If a dynamical system with a large number of states can be modeled with significantly fewer states, effectively reducing the system order, this can be a great simplification. Furthermore, if motor commands

that have similar, or redundant effects on the system dynamics can be coupled, this can be a great simplification. Model order reduction achieves both of these goals, relieving the problems of system analysis and control design.

For nonlinear systems the tools for model order reduction are not as well developed. Some theoretical underpinnings exist [60], but a general methodology is still absent. Despite this deficiency, algorithms for reducing nonlinear system dimensionality have been proposed and investigated for a select few examples [12], [39].

We have briefly mentioned some theoretical and hypothetical methods for alleviating nonlinearity, high-dimensionality and redundancy. We now examine evidence that suggests biological systems may be adopting similar strategies for the purposes of motor control.

1.3 Evidence for Biological Simplifying Strategies

There are many examples of nonlinearity, high-dimensionality and redundancy being dealt with in the controls literature, but this thesis is concerned with biological systems that may have devised strategies to do the same. We shall examine experimental evidence that suggests the central nervous system may linearize system properties, reduce dimensions, and collapse redundancies all in the application of simplifying motor control. We shall first discuss evidence that, interesting in its own right, sets the context for the motor control strategies argued to be relevant for biological systems. Then we present the evidence that has motivated this thesis, linear force production in stimulation, force fields that appear to represent motor primitives, and low-dimensional motor commands, also known as muscle synergies.

As mentioned, there are motor behaviors in biological systems that are consistent with dimension reducing strategies. For instance, the equilibrium-point (EP) hypothesis can be described as a method for neglecting dimensions for biological systems. The EP theory, first proposed by Feldman (reference), posited that posture and movement are subserved by the same process, modulation of the limb's equilibrium-point. In this theory the inherent spring-like properites of muscles, combined with spinal feedback circuits, act to drive the limb towards a stable posture. By shifting this stable posture, perhaps continuously, the central nervous system can drive the limb to a new posture, and command its motion. For example, a study that presented compelling evidence for this idea examined reaching movements made by deafferented monkeys [6]. Monkeys were trained to make targeted reaching movments, holding a manipulandum that could deliver movement disturbances. It was shown that if the monkey's hand was suddenly pushed toward the target, the limb would rebound to an intermediate position in a direction away from the target, before continuing on to the target location. This suggested that the limb could be following a slowly shifting equilibrium that drove the limb to the final target destination.

If the limb's motion could be controlled through a series of static postures, then the only computations required are based on a static map from limb posture to neural motor commands. Therefore, not only have many of the states needed to define the musculo-skeletal system been neglected, but the very dynamical nature of

the system has also been neglected. With this in mind, the EP theory can also be explained in terms of neglecting fast dynamics, those that drive the limb toward the equilibrium trajectory, in favor of slower dynamics, namely those of the modulated equilibrium. If the central nervous system were to employ some variation of this theory in the coordination of motor behaviors, it would constitute a reduction of dimensions, greatly reducing the difficulty of generating motor controls.

As mentioned, optimization strategies are a technique for constraining redundancies. There are examples in biological systems of what appear to be motor behaviors that reflect an optimal strategy. For instance, many features of motor behavior are stereotyped. In humans, reaching movements are generally made along straight paths [31]. These straight paths have a typical "bell-shaped" velocity profile as well [31]. There are also typical features when the reaching movements include intermediate targets [20]. However, it is not the geometric qualities of these movements that are interesting. What is interesting, is these geometric qualities persist over a large variety of conditions. Because these movements contain invariant patterns, despite the opportunity to do otherwise, this behavior can be framed in the context of an optimization. Given this information, it has been argued that these movements are the result of a cost optimizing strategy. Human reaching movements in particular, have been argued to be consistent with a motor plan that penalizes, end-point jerk [20], [31], joint torque rate [38], motor commands [50] or end-point variance [28]. If a cost function is being optimized, it could be a strategy the biological system has devised to help choose amongst the multitude of possible motor plans.

Regarding the system nonlinearities, there is also some intriguing evidence for simplifying strategies in biology. Descriptions of muscle behavior have a nonlinear dependence on their controlling signals [71]. Specifically, these models suggest the force that muscles produce varies nonlinearly with muscle activation. However, experimental evidence suggests that this dependence is approximately linear. Has the central nervous system adopted a linearizing strategy?

Bizzi et.al. [5] were first to report on what would begin a series of investigations on the linear production of forces in the frog hind limb. Experiments were conducted on the common bull frog, *Rana catesbeiana*. The results of microstimulating the spinal cord were examined. Typically, the spinal cord of the animal was transected, this is the so called spinalized state. Neural connections with the cortex and any other super spinal structures were therefore eliminated. Motor behaviors of an animal prepared in this manner are the exclusive result of local spinal circuits. Interestingly, in this state, frogs, as well as other animals [68], still retain basic motor skills. For example, spinalized frogs are capable of coordinated motion for reflexive removal of, or withdrawal from noxious stimuli. Kicks and leg movements that mime swimming motions can also be induced through spinal stimulation (fictive swimming and kicking? reference).

Bizzi et.al. found that stimulating the spinal cord with small amounts of electric current would induce coordinated muscle activity within the hind limb. With a transducer attached to the ankle, forces could be measured across the limb's workspace while the stimulus was active. By repeatedly stimulating the same site within the spinal cord, force measurements were found to be repeatable, invoking the notion of a spinally encoded force field. Often this activity drove the limb towards a stable

orientation, or equilbrium, within the workspace of the hind limb.

This first study also found that these spinally induced force fields, surprisingly, exhibited the quality of superposition. The force field associated with costimulating two spinal sites was approximately equal to the vectorial sum of the force fields associated with the individual sites. This finding was unexpected given the nonlinearities in muscle force generation, neuron properties and neuron interactions.

Since these initial findings were reported, further experiments have examined these spinally induced isometric force fields and their linear properties. A study conducted on approximately 100 frogs found that though the sudden onset of stimulus induced transients in the magnitude of the field's strength, the orientation remained approximately constant [25], suggesting that spinal circuitry was encoding force fields with an invariant spatial structure. This evidence further implicated isometric force fields as motor primitives.

In another study [49], the superposition properties of isometric force fields were specifically investigated. Force fields resulting from individual stimulation and simultaneously costimulated sites were measured. The costimulated force fields were compared with the vector summation of the individually stimulated force fields. The costimulated and vectorially summed force vectors were compared on a point-by-point basis, it was found that the cosine of the angle between them was >0.9 for over 87% of the cases, implying that the orientation of the costimulated force fields was similar to that of the vectorially superposed field (if they were exactly proportional the cosine would be 1.0). Further analysis of the force fields compared the relative magnitudes between the costimulated and vectorially superposed fields. This too was found to be in good agreement.

Further quantifying the linear properties of isometric force fields [41], three properties of the elicited force fields were examined: the dependence of force magnitude on stimulation parameters, the invariance of force field structure with varying stimulation strength, and the costimulated force field's similarity to a vectorial superposition of individual force fields. The study found that for a great majority of cases, the correlation coefficient between force field magnitude and stimulus strength (measured in charge delivered) was highly correlated (>0.9). The field structure, measured in terms of average angular shift with stimulus strength, was small, approximately 1 degree. And the costimulated fields closely matched those predicted by a superposition of the individual fields. The results indicated a strong tendency for force fields that increased linearly in stimulation and could be linearly superposed through linear combinations of stimuli.

These studies, and others like them, e.g. [29], indicate what could be examples of a biological system that has found strategies for alleviating nonlinearities. These isometric force fields could allow for further advantages too. Force fields could be a redundancy or dimensionality reducing strategy as well. The isometric force fields encoded within the frog's spinal interneurons were quickly recognized as a possible organizing principle for motor control. The force fields, and their associated equilibria could be a way to group muscle activities into functional units. It was speculated that these force fields could represent a mechanism whereby the central nervous system mapped motor plans to motor commands [7].

Several investigators, trying to understand how the frog's motor control was organized, proposed that these force fields could be used to construct the complex and coordinated motor behaviors produced by awake and naturally behaving animals. It was suggested that these force fields were the component pieces used by the motor circuitry of the animal to generate motor behaviors, i.e. force fields were motor primitives.

In one of the early force field studies [25], movements were compared with their inferred equilibrium trajectories. Evoking reflexive movements through cutaneous stimulation, free limb trajectories were recorded. With the limb held in fixed locations across the workspace, subsequent stimulations were used to compute a force field. The time varying equilibria of these force fields corresponded closely with the limb's trajectory for many movements. It was speculated that force fields might be used as motor primitives to generate movements, possibly with an equilibrium trjaectory.

In a study that claimed to present the first direct evidence that force field primitives combined to create motor behaviors, reflexive limb movements with and without cutaneous afferents were investigated [35]. Results indicated that both kinematics and EMG signals, before and after cutaneous deafferentation, were similar. However, when the limb encountered an obstacle there were clear differences in the measurements. The frogs made corrective movements in response to the obstacle, that were absent after deafferentation. Corrective movements consisted of sequences of stereotypical forces that varied with time of contact and force of impact with the obstacle. The authors argued that these corrective forces, were the result of a corrective force field superposed on the nominal movement generating forces. A follow up investigation [23], found further support for modular independent units, in the form of force fields, which appeared to be combined in the generation of movements.

And in another noteworthy study [24], free limb kinematics of spinalized frogs with and without afferents were compared. Movements were elicited through microstimulation of spinal gray matter. The isometric force fields for these same stimulations were computed. The isometric force fields were then used to simulate limb kinematics using a dynamic model of the limb. Limb kinematics of the actual and simulated system were found to agree well during the initial acceleration of the movement. This was offered as support for the hypothesis that the spinal cord executes movements through the use of force field primitives.

Just as force fields might represent motor primitives, the commands that induce muscle forces could also represent a motor primitive strategy. Hence, another biological example of dimension reduction may come in the form of motor command reduction. In many studies investigators collect electromyographic (EMG) data from awake and behaving frogs. Theoretically, the activity of each individual muscle can vary independent of all others. Therefore, for example, the EMG recordings of 13 muscles can constitute 13 independent signals. However, in many behaviors, such complete independencies have not been observed. Instead, the activity of muscles has been observed to covary.

It should be noted that the behaviors examined are rich enough to allow for more independent variability in the muscle activations than observed. Furthermore, the EMG signals are measured with high enough precision to detect the indepence of all

muscles. These observations argue against dismissing the low-dimensionality as an epiphenomenon.

If muscle activations covary, that is to say they are not independent, then their measurements will be contained on a hypersurface within the ambient space of their full dimension; if the activations have a linear dependence, then their measurements are contained on a linear hyperplane (a plane with dimension greater than 2). For example, if the activations of three muscles had a linear dependence, such that two activations covaried, and the third was independent of the two, then the three measurements would be contained on a two-dimensional plane within their three-dimensional ambient space. With the plane (or hyperplane) identified, the activations of all three muscles can be represented as two time varying signals, that scale two bases that represent fixed ratios of muscle activity.

Looking for regularities in muscle activity, many studies, (including some of those mentioned earlier) have found that EMG signals from as many as 13 muscles often can be reconstructed to a high degree of precision with as few as 4 or 5 ratios of activity. These fixed ratios of muscle activity have been defined as muscle synergies.

This low dimensionality of muscle activation during behavior was investigated under many different conditions. One study [67], collected EMG data of movements evoked through cutaneous stimulation with spinalized frogs. EMG recordings from 9 muscles were made. By fitting the data to fixed ratios of muscle activations, it was found that the data could be described with as few as four synergies.

In another sudy on spinalized frogs, motor behaviors were evoked through a chemical means (NMDA microstimulation) in spinalized frogs. This experimental method had the distinction of evoking motor behaviors that were not peripherally mediated relflexes, and as such it was argued that the EMG data was richer than in previous studies. The EMG data collected from 12 muscles was used to find a low-dimensional representation using a principal component analysis method. It was found that 96% of the data was accounted for with 7 synergies (a reduction in dimension of almost half, from 12 to 7). Due to the task independence of the motor behaviors evoked, and the similar synergies found amongst different frogs, the possibility that these synergies were a basis for all motor behaviors was put forth.

Extending the idea of synergies further, it has been suggested that the fundamental units being manipulated for the generation of motor behavior are not fixed ratios of muscle synergies, but rather patterns of muscle activation that vary in time. In a study using defensive kicks from awake and behaving animals, EMG data was fit to time varying synergies [15], [14]. As few as three time varying synergies could be used to fit the data from 13 muscles.

If synergies are the bases for muscle activations, both for reflexive and awake behaving motor behaviors, this would represent a large reduction in the number of controlled variables. What's more, commands produced with linear combinations of synergies could be an effort to group together redundant motor signals into functional units [11]. It has also been suggested that synergies are the bases for force fields [43], further implicating them in the use of motor primitives. All of these possiblities are an opportunity for the central nervous system to simplify the problems associated with high-dimensionality and redundancies in motor control.

1.4 Thesis Objectives

The spinal motor circuitry of the frog has revealed some auspicious artifacts. These spinalized animals exhibit forces that are linear in stimulation; they appear to scale linearly, obey vector superposition, and have a structure that is invariant with increasing stimulation. These same isometric force fields reflect characteristics consistent with motor primitives; modular force fields appear to be the basis for some spinally mediated behaviors. The motor commands of frogs are low-dimensional and may be composed of fixed ratios of muscle activity; motor commands could be restricted to combinations of muscle synergies. In part, it is these phenomena that have led investigators to conclude that spinal circuits contain a great deal of function and motor "intelligence". Whereas previously the spinal cord was considered a passive device relaying information from the cortex to the animals extremities, now it is clear there is more to the story [7]. Could these spinal circuits by working to linearize, reduce dimensions and redundancies in the interest of simplifying motor control? We shall explore these issues in the thesis.

The "linear" nature of isometric force fields is an intriguing characteristic. Although linearity is not necessary for a control structure based on force fields, linearity has distinct advantages. When speculating on spinal motor primitives in the form of force fields, linearity is often a prominent feature. Arguments typically point out that motor behaviors produced through combinations of force fields can be easily interpolated or extrapolated to new behaviors by exploiting linearity. For example, when motor goals can be framed in terms of forces, behaviors can be easily generated by stringing together appropriate force fields in linear combination. This may greatly reduce the computational load in generating motor commands.

The linear isometric force fields, and the force fields inferred from motor primitive studies, have intriguing implications for motor control. There is compelling evidence that spinal circuitry generates motor behaviors through a combination of discrete force fields. These force fields may represent the elemental units of the motor generating structure.

The ability of complex behaviors to be explained as combinations of synergies indicate that motor commands are low-dimensional and restricted to time varying combinations of specific commands. This could be an effort to reduce the high-dimension of motor commands. They could also be the basis for motor primitives. In either event, they are a simplification over control that must manage all the motor commands available independently.

Much has been made concerning the connection between the linearity in isometric force fields, motor primitives and synergies. However, there is little direct evidence to link them, either functionally or physiologically. What's more, aside from ad hoc attempts to explain possible links between these phenomena, there has been little in the way of a principled explanation for a motor control framework that would give rise to these phenomena. The main goal of this thesis is to provide such an explanation.

The thesis is divided into two parts. First the issue of linearity in isometric force fields is addressed. It has been observed that the strength of isometric force fields increases linearly with stimulation parameters. This apparently linear phenomena

in a nonlinear system is commonly observed and reported yet relatively little effort has been made to analyze and explain this seemingly counterintuitive behavior. The linearity is investigated in the context of motor control. The thesis asks if linearity is somehow a result of specific motor commands and feedback strategies, or whether it is in fact a natural result of the biomechanical properites of the musculo-skeletal system. The results found indicate that while a detailed model of the limb is fully capable of exhibiting nonlinear behavior, when restricted to parameters that accurately describe the animal in question, the model predicts nearly linear behavior. Based on these results we propose that force fields that scale linearly in activation are a natural result of the specific biomechanic properties of the limb and are not the outcome of specific linearizing motor control strategies. We also propose a corollary hypotheses, that the animal has evolved to allow for these linear behaviors is.

The second investigation centers on the connection between low-dimensional motor commands and hypothesized motor primitives. Many experimental investigations have examined muscle synergies, fixed patterns of muscle activation. It has been found that a small number of synergies may account for motor commands across different behaviors. Other experiments have probed motor behaviors for modular features in the form of force fields, that might indicate the presence of motor primitives, uniquely controlled variables that can be flexibly combined to accomplish motor tasks. And further studies have looked for connections between synergies and motor primitives, specifically questioning if the former is the basis for the latter. However, this work has largely been descriptive in nature, trying to explain the low-dimensional data without reference to the underlying control structure. For instance, it has not been addressed how synergies are specified, or how control might be enhanced or limited by their use. There is still want for, if not a causal connection between, a principled explanation for how motor primitives, low-dimensional motor commands and synergies arise, how they are connected, and how they are implicated in the organization of motor control. This thesis investigates one such explanation, using apparatus drawn from model order reduction theory.

A reduced order model is a simplified representation for estimating the input-output dynamics of a dynamical system. Many low-dimensional features result from employing a reduced order model in motor control. We propose that a reduced order model of the input-output behavior of the frog hind limb's musculo-skeletal dynamics is capable of estimating system properties relevant for controlling behavior. These estimates allow the central nervous system to generate motor commands to produce behaviorally relevant actions and predicts the existence of low-dimensional motor commands, and motor primitives in the form of synergies.

We examine a method for finding a low-dimensional model that estimates the full system. The method attempts to find a compact representation for states that most significantly impact the energy of the input-output dynamics. The reduced model is examined for its abilities to estimate system properties and for its use as a controller, generating motor commands for the full system. The reduced system is further examined to illuminate the role of motor primitives in motor control and their relationship to limb biomechanics. A formalism for defining motor primitives is defined and the resulting primitives are compared with experimentally derived

synergies.

In chapter 2 we introduce the many models used in this thesis for analysis and simulation of the musculo-skeletal system. In chapter 3 we investigate the biomechanical sources of linearity in force production. Our results indicate that the linearity experimentally observed may be explained as the natural result of the biomechanical properties of the frog's musculo-skeletal structure. In chapter 4 model order reduction theory is reviewed in the context of our eventual application, computing a reduced order model for a nonlinear system. In chapter 5 we investigate a reduced order model for use with controlling a physiologically accurate frog hind limb model. The model presented is interpreted as an existence proof for a means to simplify motor control. We use the model to demonstrate system estimation and control, as well as to explain motor primitives. Chapter 6 is a discussion of the contributions of this thesis and some connected ideas concerning biological motor control. We conclude the thesis in chapter 7 and provide some ideas for future work.

Chapter 2

System Models

In this chapter we describe the component models used to simulate the hind limb of *Rana pipiens*. Some of the models used are quite intricate and merit a detailed description. To best procede we break up the description into three parts. First we review some of the pertinent muscle physiology and mechanics to motivate and introduce the muscle models used. Building on these muscle models, a model of the frog hind limb's musculo-skeletal system is presented. This model is created using the software package, Software for Interactive Musculoskeletal Modeling, or SIMM. The SIMM model incorporates details of the inertial and geometric properties of the hind limb bones, as well as the particulars of how the muscles couple to it. The SIMM model is used to address the first portion of this dissertation, that of linear force production in stimulation.

The third section of this chapter will present a second dynamic model of the frog hind limb. This second dynamical description is based on the SIMM model but incorporates a few simplifications to allow for fast simulations in Matlab. It is this model that we use for the second half of the thesis study, to analyze the feasibility of a reduced order model to control the limb.

2.1 Muscle Modeling

In this section we briefly review the salient features of muscle, or more precisely musculo-tendon, anatomy and physiology. Our goal is not to present an exhaustive analysis of muscle, but rather enough information to motivate the model used to describe it. There are many published examples in the literature that do a thorough and compelling job of motivating these same models (e.g. [9], [70], [72] and [73]). What follows is true for mamallian, and in large part, vertebrate skeletal muscle. For further details refer to [34], [71], [46], [8].

2.1.1 Muscle Anatomy

Muscle is composed of muscle fibers and tendon. The muscle fibers constitute what is referred to as the muscle "belly", the active force generating portion of muscle. The

tendon, an elastic connective tissue, joins disparate groups of muscle fibers to each other and the skeleton.

Muscle fibers are composed of smaller functional units called myofibrils. The myofibrils are in turn composed of even smaller units called sarcomeres. The sarcomere is the smallest functional unit of muscle contraction and force generation.

The sarcomere is responsible for both active and passive force generation. Composed of long overlapping protein fibers, the sarcomere contracts as its constituent proteins slide past each other in a calcium-mediated electro-chemical process. This process of mechanical deformation produces tension across the sarcomeres and in turn the muscle fibers. The mechanical deformation and resulting force generation requires a metabolic cost. Furthermore, the process is influenced by the level of stimulation the muscle fiber receives from it's innervating neruons. The level of stimulation a fiber receives determines the strength of these protein contracting dynamics and is termed the active state. Therefore, the forces derived from the sarcomere in this way are referred to as active force generation.

In addition to producing force actively, the sarcomeres (due to further protein chains and structural tissue) generate resistance to stretch that is independent of stimulation, or active state, and does not require a metabolic cost. This is the so-called passive force.

Tendon, too, resists stretch passively, though generally with much greater force than muscle. That is, tendon is generally much less compliant than muscle fiber. The tendon acts to transmit the force muscle generates to the skeletal links muscle effects.

Each muscle fiber is innervated by a motor neuron. Each motor neuron, along with all the muscle fibers it innervates is referred to as a motor unit. The motor units are the smallest individual element the central nervous system controls.

2.1.2 Muscle Physiology and Mechanics

Force-length relation

When stimulated, muscle fiber contracts. The subsequent force generated is a function of both the length and velocity (shortening or lengthening) of the muscle fiber. From experimental studies (e.g. [46]) it is known that the force generated as a function of length has a typical profile. The force is maximal about a unique length, l_o^m, the optimal muscle fiber length, and decreases as the fiber length increases or decreases. The details of this force-length profile are beyond the scope of this thesis, suffice it to say, it can in large part be explained in terms of the sarcomere's mechanical deformation. The strength of this response increases with increasing muscle activation state. This is therefore referred to as the active force-length curve.

In addition to this, muscle fiber also generates force as a function of length that is independent of activation level. This passive tension is generated as the muscle fiber is stretched to lengths larger than l_o^m. The force increases in a nearly exponential relation to stretch [73].

Force-velocity relation

In 1938 Hill performed the quick release epxeriments for which he is famous. Af-

ter a load attatched to muscle was released, the initial velocity of the load (and the lengthening muscle) was measured. The load's mass was varied and the subsequent initial velocities were measured. After what appeared to be an instantaneous deflection of the load/muscle the velocity was approximately constant. The relationship between load's mass and initial velocity appeared linear. Based on heat and work measurements Hill derived the famous force-velocity curve that bares his name.

$$(T+a)(v+b) = (T_o+a)b \implies T = \frac{(T_o b - av)}{(v+b)} \tag{2.1}$$

In this curve, T is the muscle tension, v is the muscle velocity and a and b are constant coefficients. T_o is the tension at rest.

Tendon Complicance
Tendon has no active characteristics. It is merely connective tissue. In force-length measurements tendon appears elastic (or slightly viscoelastic). Initially the tendon is relatively compliant, this defines the so-called "toe" region. The stress a tendon experiences as the length further increases, approaches a near constant level. If, as is often assumed, the stress-strain relation of tendon is constant, then the compliance of tendon can be computed based on its rest, or slack length, l_s^T. As a consequence of the constant stress-strain assumption, the longer a tendon's slack length, the larger its compliance.

Twitch and tetanus
When muscle is stimulated with a brief electrical impulse force begins to develop. The force slowly increases to a peak level, and then decays back to zero. This transient force profile is referred to as a twitch. Once a twitch has expired, a second electrical impulse will elicit the same response. If another impulse is delivered before the twitch has decayed to zero, the response appears as the superposition of two twitch responses. As the rate at which the electrical impulses is delivered, the force profile appears as a ripple, or sinusoidal signal, with a nonzero mean value. Stimulating the muscle at higher frequencies, the strength of the force eventually saturates. This is the tetanic fusion state; muscle in this state is tetanized.

Activation Dynamics
As mentioned above, the electrical excitation of muscle elicits what is referred to as an activation state (the strength of which can be inferred from the muscle force). Experimentally, activation is approximately proportional to time integral of excitation. The activation dynamics can be approximated as a low-pass filtered version of the excitation. However, it has been widely observed ([73]) that this relation is nonlinear, as the rate for increasing activation is faster than that for deactivation.

2.1.3 2-element Nonlinear Hill model

The first muscle model we present is an effort to incorporate muscle fiber contraction dynamics (Hill's force velocity relation, and the empirical force-length relation) in a

simple fashion. This is done with a two element nonlinear model. The two elements are an active contractile element in parallel with a passive elastic element (see figure 2-1). The contractile element models the active force generating character of muscle fiber. The passive elastic element represents the muscle fiber's resistance to deflection. This may or may not represent tendon elasticity depending on modeling assumptions. If the tendon is assumed to act in the same direction or nearly parallel to, the muscle's tension (which may be a valid assumption for some muscles) then the parallel elastic element could incorporate tendon properties. However, in this thesis we shall assume that all the muscles modeled have tendons in series with the forces they generate. Under this assumption, the parallel elastic element models muscle fiber's passive elasticity only. This is a reasonable model if it is assumed the tendon's elasticity is significantly less than that of the muscle fiber's.

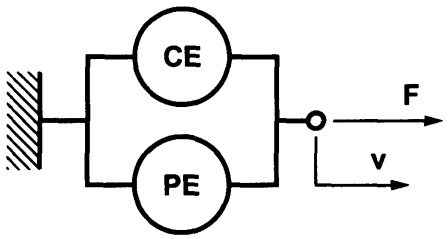

Figure 2-1: 2-element Hill model diagram. The two elements are the contractile element (CE) and the parallel passive elastic element (PE).

The contractile element models the muscle fiber's dependence on fiber length, fiber velocity, and activation state, denoted as l^m, v^m and α respectively. The length and velocity dependence have already been described. The activation state of the muscle fiber is a description of the amount of muscle fibers actively generating force; zero when the muscle is wholly passive, and one when completely excited. The functional form of this length-velocity-activation dependence is commonly represented as being multiplicative. That is, if the force-length and force-velocity functions are $f_x(l^m)$ and $f_v(v^m)$, then the active contractile element's force is,

$$F_{ce} = \alpha f_x(l^m) f_v(v^m) \tag{2.2}$$

The passive elastic element represents the muscle's resistance to displacements without regard to activation state. If, as we are modeling it here, this is the muscle fiber's elasticity and not that of the tendon, it is typically represented as a force that increases exponentially with displacement. We shall denote the passive elastic force as F_{pe}.

Before proceeding to the kinetics of the model, a brief statement on impedance and admittance modeling is necessary. As the model stands, it can either produce forces in repsonse to imposed motions, or it can produce motions in response to imposed forces. In the first scenario, the muscle is modeled as an impedance, in the latter scenario the muscle is an admittance. The two scenarios require distinct conditions from the model.

When the muscle model is an impedance, the muscle's velocity (and in turn position) is specified and the model produces the corresponding forces. The total muscle length and velocity shall be denoted as l, and v. Due to the parallel configuration, $l^m = l$ and $v^m = v$, and the muscle force is,

$$F = F_{ce}(l, v, \alpha) + F_{pe}(l) \tag{2.3}$$

When the muscle model is an admittance, the muscle's force is specified and the model predicts the appropriate motions. To do this the contractile and passive element's force relations must be inverted to compute an instantaneous velocity for the muscle. However, in this thesis the two-element model is never used in this capacity.

2.1.4 3-element Nonlinear Hill model

The three-element model is an attempt to incorporate tendon mechanics into our muscle model. Now the model is more accurately referred to as a musculo-tendon model (throughout the thesis the term musculo-tendon shall be used to describe models that represent muscle). A passive elastic element is added in series with the previous two-element model. In general, the series elastic element, of which the tendon is likely the dominant contributer, will produce a force, maybe nonlinear, that's a function of displacement. A schematic of the three-element model is shown in figure 2-2.

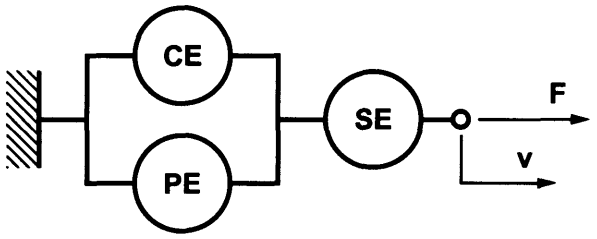

Figure 2-2: Three element Hill model diagram. The three elements are the contractile element (CE), the parallel passive elastic element (PE) and series elastic element (SE).

Just as before, this model may be an admittance or an impedance and the defining relations will be used in different ways accordingly. If motions are imposed on the muscle, the musculo-tendon must be modeled as an impedance. However, under these conditions some computational difficulties arise. The model is not well-posed; the series elastic element will be in derivative causality and to solve for the musculo-tendon force will require taking the time derivative of the contractile element's force. This is an unwanted and unnecesary complication.

To avoid this computational difficulty we shall add a small mass element between the series elastic element and the contractile element. The mass element also has the added benefit of possibly representing the mass of the muscle belly. To better

illustrate how this influences the system the bond graph notation shall be utilized. In figure 2-3 bond graphs are shown with the addition of a mass element. The left schematic is the bond diagram for the case when the model is coupled to an effort source (forces are prescribed) and the musculo-tendon is an admittance. The right schematic is the bond graph when the model is coupled to a flow source (motions are prescribed) and the musculo-tendon is an impedance.

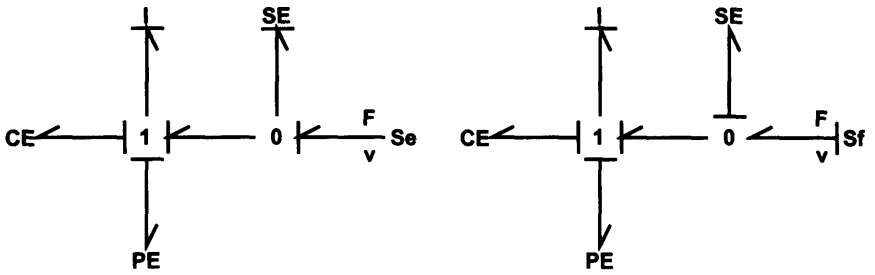

Figure 2-3: Bond graph for 3 element Hill model with added inertia element. On the left, the system is coupled to an effort source. On the right the system is coupled to a flow source.

Now, returning to the case when motions are imposed on the model, the mass element is now in integral causality and the relations for prescribing the musculo-tendon force are well-posed. The system relations are third order. Alluding to the tendon, the series elastic length shall be denoted l^T and the force as, F_{se}. The inertia of the mass element is δ.

$$\begin{aligned} \frac{d}{dt}l^T &= v - v^m \\ \frac{d}{dt}l^m &= v^m \\ \frac{d}{dt}v^m &= \frac{1}{\delta}\left[F_{se} - F_{pe}(l^m) - F_{ce}(l^m, v^m, \alpha)\right] \end{aligned} \tag{2.4}$$

The musculo-tendon force is the tension across the tendon, $F_{se}(l^T)$.

When forces are imposed on the model, the musculo-tendon is an admittance. The relations are now second order.

$$\begin{aligned} \frac{d}{dt}l^m &= v^m \\ \frac{d}{dt}v^m &= \frac{1}{\delta}\left[F(t) - F_{pe}(l^m) - F_{ce}(l^m, v^m, \alpha)\right] \end{aligned} \tag{2.5}$$

The total musculo-tendon length is $l = l^m + F_{se}^{-1}(F(t))$.

2.2 Hind Limb Model I

To investigate the biomechanical sources of linearity in force production, we need a model of the frog hind limb. First and foremost, the model will be required to reproduce musculo-skeletal behavior at a level of detail commensurate with our hypothesis.

Therefore, the transient dynamics of the system are not of great importance. For instance, while the exact temporal dynamics of twitch and tetanus are not of primary concern, the force-length curves for both muscle fiber and tendon ought to be reproduced fairly accurately. Similarly, model details like variable muscle moment arms and peak isometric tension are not key features of the hypothesis. However, whatever details these or other model features bring to our model, that make it a more accurate description of the actual frog hind limb, are a welcome advantage.

With all this in mind, the model we use was built to be simulated using a commercially available software package called Software for Interactive Musculoskeletal Modeling, or SIMM. In this section we introduce the SIMM model for the *Rana pipiens* hind limb, made available to us from the Larry Rome laboratory. The structure of the SIMM model and the capabilites of the simulations are discussed. But first we briefly introduce the capabilities of the SIMM application.

2.2.1 SIMM

The SIMM application is a graphics based software package built for the express purpose of developing and analyzing musculo-skeletal models. The software is built specifically to represent bones, muscles and their interactions. Some of the SIMM features taken advantage of are dynamic muscle models, detailed skeletal geometry and musculo-tendon geometry.

The SIMM package comes complete with several classes of muscle models. It can also be edited to include user defined models. To this end our own software was added to SIMM to model the 3-element nonlinear Hill musculo-skeletal model for the hind limb.

The bones in a model are defined with axis of rotation (possibly translating) and inertial parameters. For instance, a femur can (and does) have three axis of rotation at the hip and an inertial tensor to describe the mass distribution about each axis. With two or more bones in a skeleton, the location and degrees of freedom of each bone are defined relative to the neighboring bones. For instance, the femur's location and axis of rotation, are defined relative to the hip bone. Similarly, the tibiofibula is defined in terms relative to the femur bone. Realistic visual representations of bones can also be included in the application. Bone files, provided by the user, can be incorporated into the program for realistic images of the model.

Within SIMM the user must define attachment, or insertion points, for all muscles. These are locations on bones that the muscle is rigidly connected to. For arbitrary skeletal configurations SIMM automatically computes the length and "line of action", or direction in space each muscle applies forces to the bones. This allows for variable muscle moment arms.

If the user wishes, via points can be specified. A via point defines a location in space that the muscle will always pass through. For example, a via point can be used to enforce that a muscle never passes through a bone.

For more realistic descriptions of contact between bones and muscles, SIMM includes wrap objects. Wrap objects, defined with spheres, ellipsoids and cylinders, designate surfaces muscles "wrap" around and over. SIMM calculates the muscle

paths over these wrap objects and applies forces to bones in accordance with this muscle geometry.

When a model in SIMM is fully defined, the application creates a dynamical description of the system which can then be used in simulation. Once complied, SD/Fast, another commerically available software package, is used to compute forward or inverse dynamics for the model. The kinematics for the model can be supplied, in which case SD/Fast will compute the joint torques necessary to achieve the motion (inverse dynamics), or muscle excitations can be supplied and SD/Fast will compute the resulting motions (forward dynamics). In addition, as long as the system is not overly specified, some of the kinematics may be specified, as well as the muscle excitations, and SD/Fast will compute the resulting motions. When the simulations are run, output variables are written to a file for later use. Joint torques and kinemtics, as well as most of the intermediate variables, e.g. muscle forces, and joint accelerations, can also be written to a file when a simulation is run.

2.2.2 Model Structure

The SIMM model used in this thesis was, in large part, created in the lab of Lawrence Rome. The model, created through detailed measurements of muscle and skeletal properties of many frogs, was detailed in a series of published papers ([36], [37]). The full model includes bones for the hip, femur, tibiofibula, an astragalus-calcaneus bone, and a single unit for the metatarsophalangeals. The bones all have experimentally derived degrees of freedom and geometry. Most of the limb's muscles are in the model, including all the parameters necessary to describe them. In addition, wrap objects and insertion points for muscles have been created using experimental data.

Though the model has potential for a great deal more, for the purposes of this study, the hind limb was restrained to movements within a plane. The movements were also restricted to two degrees of freedom for the femur and tibio-fibula (hip and knee joints). For this reason we will limit our review of the model description to the hip, femur and tibio-fibula and their attendant properties. For similar reasons the muscles used in this study are limited to thirteen of the hind limb. These thirteen muscles constitute the bulk of the hindlimb's mass (approximately 90%) and also have been studied extensively by other researchers in frog motor behavior studies.

2.2.3 Parameter Determination

Multiple frogs of approximately the same size (28g) were sacrificed and the muscles, organs and connective tissue were dissected from the skeleton. All connective tissue surrounding limb joints were left intact to ensure proper location and motion of the bones. Three-dimensional images of the skeletons were then constructed. The intact skeletons, placed in a reference position on a rotating stage, were imaged with a three-dimensional laser scanner. The stage was rotated through 10^o and another scan was made. 36 such scans of the skeleton, each rotated 10^o apart were collected to construct a three dimensional image of the skeleton. The individual bones were then dissected from the skeleton and individually scanned. Three dimensional images of the

individual bone segments were then overlayed on the intact skeletal image, to locate the relative positions of individual bones within the skeleton once the connective tissue was no longer present. This data was used create both a visual representation for the animal within the SIMM enviroment and also to locate the local coordinate frames for each bone segment.

In a seperate series of experiments, excess tissue was again dissected from the skeleton and the relative motion of bone segments was measured to determine the ranges of motion for each joint. Custom made jigs were used to measure rotation and translation of one bone segment relative to another.

The hip bone is used as a grounding element; it is fixed to an inertial reference frame. The femur is connected to the hip via a ball and socket joint allowing for three rotational degrees of freedom. For this study, only one rotational degree of freedom is unconstrained, allowing for hip flexion and extension in the animals horizontal plane. The tibiofibula is connected to the distal end of the femur via a single degree of freedom pin joint. This allows for rotation in the same plane as the hip joint.

Each bone's mass, center of mass and rotational inertia were carefully measured on multiple frogs and averaged (see Appendix A for parameter values). Animals were frozen and cut into segments (e.g. upper leg and lower leg). The mass and volume of each segment was measured. The moments of inertia were then calculated based on assumptions of simplified geometric solids.

The attachment, or insertion points for thirteen muscle were found as follows. After animals were sacrificed, the muscles were partially removed, leaving small portions affixed to the bone. After drying, the skeletons were scanned using a three-dimensional laser scanner as described above. The remaining muscles were then removed from the bones and the skeleton was once more scanned. The two images of the skeleton were then overlayed to determine the location of the insertion points.

As mentioned abobe, the SIMM model allows for the inclusion of wrap objects, to simulate physical boundaries between the muscles and the bones. The model includes four wrap objects. At the hip joint there are two wrap objects, both serve to keep the musculo-tendons that cross the hip joint from passing through the femur head and the hip. The remaining two wrap objects are affixed to the distal end of the femur. One allows msuculo-tendons to slide over the rostral surface of the knee joint, and the other provides the same function for the caudal surface of the knee joint. Additional via points were added to the model to approximate the effects of connective tissue.

For each muscle, optimal muscle fiber length, l_o^m, tendon slack length, l_s^T, pennation angle, θ and peak isometric tension, F_o^m were computed. The procedure is briefly explained here, but the details can be found in [37]. Under controlled conditions, muscle sarcomere lengths were visually measured. Then, based on an assumed muscle tension and tendon strain, l_o^m and l_s^T were computed[1]. Pennation angle, θ was found with caliper measurements. F_o^m was computed as the product of physiological cross-sectional area and muscle stress, assumed to be $260kN/m^2$.

[1]We found through an independent analysis that the results are not sensitive to the assumed muscle tension.

2.2.4 Model Overview

Unfortunately, it seems many of the muscles of the frog's hind limb are referred to with multiple names. For consistency, we will use the terminology used in [37]. Shown in tabular form, are the muscles contained in the model. In the first and second columns are the abbreviated and full muscle names. For convenience the third column lists alternate names for the muscles.

Abbrev.	Muscle Name	Alternate Name
ILe	Iliacus Externus	Gluteus
ILi	Iliacus Internus	Ilio-Psoas
ILf	Ileo-Fibularis	Ilio-Fibularis (biceps)
TFL	Tensor Fasciae Latae	Rectus Femoris Anticus
GL	Gluteus Magnus	Vastus Externus
CR	Cruralis	Vastus Internus
SM	Semimembranosus	-
SA	Sartorius	-
GR	Gracilus Major	Rectus Internus Major
STd	Semitendinosus dorsal	-
STv	Semitendinosus Ventral	-
ADd	Adductor Magnus Dorsal	-
ADv	Adductor Magnus Ventral	-

Table 2.1: Muscle Names and abbreviations

An image of the SIMM model is shown in figure 2-4. The red cords are the paths of the musculo-tendons. In blue are the wrap objects, displayed as hatched cylinders and hemispheres. At the hip joint there are two wrap objects, two more are affixed to the knee joint.

The SIMM model has two kinematic degrees of freedom, the hip and knee angle. The range of both joints is $\pm 90^o$. When the hip angle is zero degrees, it is perpendicular to the long axis of the pelvis. When the knee angle, which is defined relative to the hip joint, is zero degrees, it is perpendicular to the long axis of the femur. For the remainder of the thesis, the orientation corresponding to both joints at zero degrees shall be referred to as the zero orientation. This is depicted in figure 2-4.

Positive angluar displacements, for both joints, are defined as clock-wise rotations. To better visulaize how the hip and knee angles sweep out the ankle position refer to figure 2-5. In figure 2-5 isoclines for the hip (left plot) and knee (right plot) are plotted across the ankle's workspace. Also depicted in the plot are the hip, knee and ankle location within the workspace when the leg is in the zero orientation.

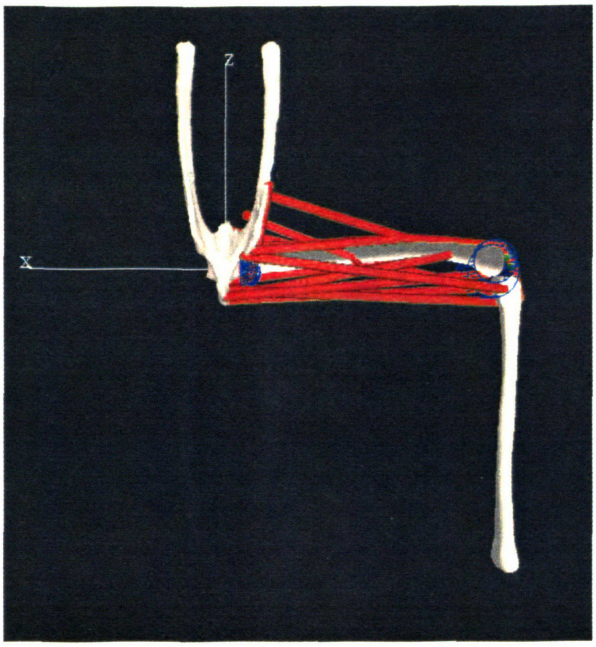

Figure 2-4: Image of SIMM model of *Rana pipiens* hind limb (provided by the lab of Larry Rome) showing the hip, femur and tibiofibula. In red are the 13 muscles. In blue are the 4 wrap objects.

2.3 Hind Limb Model II

The SIMM model, with it's precise wrap objects and muscle insertion points allows for very detailed simulations. This detail is exploited for the purposes of simulating realistic isometric force fields. However, for running dynamic simulations where the limb is free to move on its own accord, the model is computationally intense; the simulations take a significant amount of time to run. For this reason we concluded that the SIMM model would not be appropriate to run the many simulations that would be needed to investigate the use of a reduced order representation for control. Therefore a second, more computationally tractable model was sought.

To reduce the model complexity, two structural changes were made. The model's configuration dependent muscle moment arms were replaced with constant moment arms, and the 3-element Hill muscle models were replaced with 2-element models. There is no clear rationale for including constant moment arms, as many of the musculo-tendon's moment arms varied considerably, some even changing sign within the workspace. However, this simplification greatly reduces the simulation run-time. What's more, with constant moment arms the model still retains similar joint torques and muscle excursions.

The switch to the 2-element Hill model though, was justified. As will be shown in the following chapter, the musuclo-tendons in this hind limb model are relatively short and stiff. Therefore the series elastic element of the muscle model simply transmits

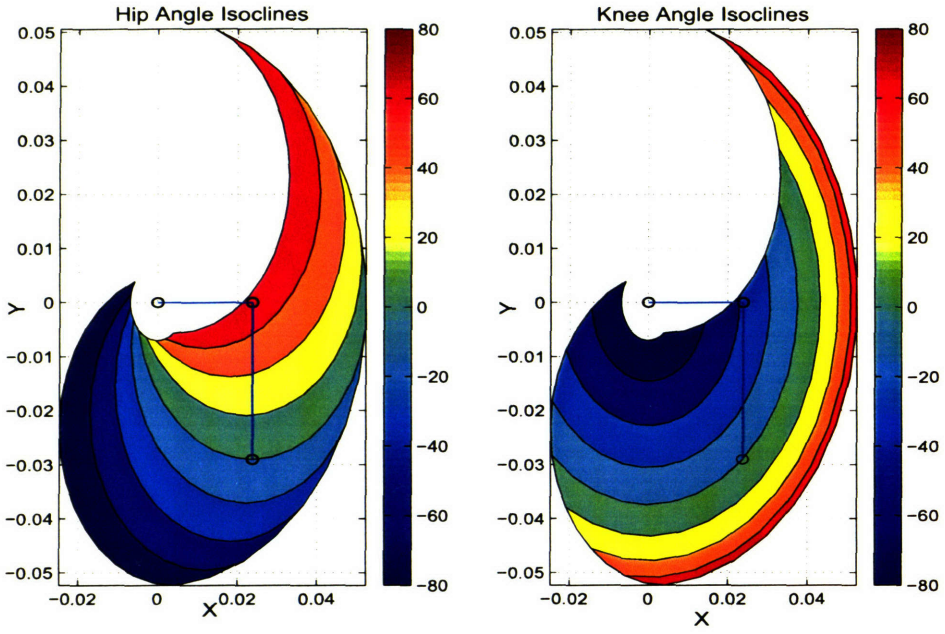

Figure 2-5: Hip and knee agnle isoclines for the ankle position. Also depicted are the locations of the hip, knee and ankle when the hip and knee joints are 0 degrees.

the contractile and parallel elastic element's forces and contributes little to the overal musculo-tendon dynamics.

2.3.1 Moment Arms

In order to define constant moment arms for each muscle, averages were computed across the workspace, $\pm 90^o$ for both hip and knee joints. For many of the biarticular muscles that passed over wrap objects, there was little variation in the moment arms. The averaged muscle moment arms can be grouped into a matrix, R for symbolic notation. The joint torque due to all the muscles can then be expressed as,

$$\tau = RF \tag{2.6}$$

where F is a diagonal matrix, with the force due to the i^{th} muscle at the (i, i) element. R, in centimeters, is,

$$R = \begin{bmatrix} 0.3 & 0.19 & 0.17 & -0.11 & -0.03 & -0.33 & -0.02 & -0.25 & 0.28 & 0.28 & 0.11 & -0.07 & -0.23 \\ 0.08 & 0.18 & 0.18 & -0.21 & -0.21 & -0.21 & 0.1 & 0.13 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

2.3.2 Muscle Models

With constant moment arms, the length of any muscle is a linear function of the joint angles. In matrix form, all muscle lengths are,

$$L = L_0 + R^T\theta \tag{2.7}$$

where L_0 is chosen such that each muscle has the same length as their SIMM counterpart when the limb is in the zero orientation, $(0^o, 0^o)$.

To model the muscle's active contractile force element, a quadratic curve, rather than the cubic spline in the SIMM model, was used. This greatly speeds up simulation time. The quadractic force profile is modeled as,

$$f_x(l) = \alpha[-3 + 8l/l_o^m - 4l^2/l_o^{m2}]^+ \tag{2.8}$$

where l is the muscle length. Note that we impose zero force at $l_o^m/2$ and $3l_o^m/2$. The contractile element's force profile is now discontinuous at the extremities, however, for most muscles these muscle lengths are never experienced within the defined ($\pm 90^o$) workspace.

To model the velocity dependence, a Hill curve was used. The parameter values were adjusted to closely approximate the SIMM cubic spline. Defining $v_m = v^m/v_{max}$ as the normalized muscle velocity, we have,

$$f_v = \frac{(b_1 - a_1 v_m)}{(b_1 + v_m)} \tag{2.9}$$

for muscle shortening, and,

$$f_v = \frac{(b_2 - a_2 v_m - 2v_m)}{(b_2 - v_m)} \tag{2.10}$$

for muscle lengthening. The overal normalized contractile element's force is, as before, $F_{ce} = \alpha f_x f_v$.

The passive parallel elastic element was also modeled with a quadratic curve rather than SIMM's cubic spline.

$$F_{pe} = \rho(l - l_o^m)^2 \tag{2.11}$$

ρ was chosen to fit the SIMM passive force length curve. Note that this curve is also normalized with respect to the muscle's maximum isometric tension, F_o^m.

The remaining variable to model is the activation state. The excitation-activation dynamics are typically modeled as a low-pass filter. Though the evidence suggests these low-pass dynamics may be nonlinear, with a faster rate for activation than deactivation, we model the dynamics as linear.

$$\frac{d\alpha_i}{dt} = (u_i - \alpha_i)/\tau_i \tag{2.12}$$

where u_i is the command to the i^{th} muscle.

As described in a previous section, for this 2-element model, the total muscle force is,

$$F = F_{oi}^m(F_{ce} + F_{pe}) \tag{2.13}$$

where F_{oi}^m is the maximum isometric tension for the i^{th} muscle, taken directly from the SIMM model.

2.3.3 Miscellaneous

The remainder of version II of the hind limb model is equivalent to the SIMM model. The inertial parameters as well as the link lengths and mass centers were taken directly from the SIMM model. However, the dynamics were simulated using Matlab rather than the SIMM application. This was done to simplify the continual need to modify and update the software written to examine and analyze the system.

One further modeling issue shall be addressed, that of joint damping. Though the contractile element does have velocity dependence, it is modulated by activation state and the musculo-tendon models do not passively produce forces to resist velocity. The mass of the limb links is very small relative to the forces the muscle's produce, so much so that the forces the muscles generate can accelerate the limb across the workspace in tens of milliseconds. To add numerical stability to the system, as well as increase the accuracy of the model, velocity dependent damping terms were added to the joint torques. The joint damping effects competing factors. Large damping forces slow down the system's natural dynamics, increasing the time required to attain steady-state conditions. This also increases the simulation run-times (it slows down Matlab). On the other hand, light damping forces make the natural response fast and speed up the time to attiain steady-state and decrease simulation run-times. However, this can also result in unnaturally fast limb movements. Through trial and error, and comparison with published results of frog limb movements [1], the damping at the hip and knee joints was set to $B = 2.5 \times 10^{-4} Nms/rad$.

Chapter 3

Weakly Nonlinear Muscle Behavior

3.1 Linearity Hypothesis

As discussed in the introduction, in experiments investigating force production in spinalized frogs the isometric force fields appear to be linear in stimulation. This is, ofcourse, only approximately true. Isometric force fields and the muscle behavior that generate them are decidedly not linear. However, there is a long-standing convention in the literature of describing the isometric force fields as "linear" in stimulation. More accurately, experimental studies have found these force fields to approximately conform to the rules of scaling and superposition, suggesting they are only weakly nonlinear. That being the case, we seek to examine the sources of this linearity, and the degree to which it holds. To do so, we will examine the biomechanical properties of a hind limb model that may allow for such behavior.

We propose the following hypothesis:

The linear force production in stimulation can be explained in terms of the natural biomechanical properties of the limb's musculo-skeletal system.

Below we present evidence to support this hypothesis. We demonstrate that although the hind limb model can exhibit force production that is nonlinear in stimulation (either at the level of individual muscles or force fields), the limb does in fact appear largely linear in stimulation for an appropriate range of parameter values and regions of operation. What's more, the parameter values that best describe the frog, and a large portion of the limb's workspace, coincide with this linear region. This seeming "coincidence" argues that the frog may have evolved to exploit this linear characteristic.

To examine this hypothesis, a realistic musculo-tendon and musculo-skeletal model of the frog hind limb is analyzed. Therefore, a musculo-tendon model will be examined. Furthermore, to limit our study to biomechanical sources, muscle behavior in the absence of feedback is examined. We seek to determine the degree to which the relationship between muscle force and activation is essentially linear.

3.2 Empirical Evidence Review

In one of the earliest studies specifically examining the linearity of spinal force fields, Mussa Ivaldi et.al. [49] reported the results of isometric force fields at the ankle and their ability to be superposed. In the study, force fields were recorded for individual stimulation sites. Then the sites were costimulated and the resulting force fields were recorded. These costimulated force fields were compared with the linear superposition of the individually stimulated sites.

For a quantitative analysis, two measures were computed. The first was a test of the superposition of the active component of the force fields (the force field that resulted from subtracting the passive contribution). Two fields were measured when individual spinal sites were stimulated. Then, the field elicited by costimulating the individual sites was measured. This costimulated field was compared with the field that would result if the individual fields were vectorially summed. The two fields were found to be in close agreement, supporting the hypothesis of linearity in superposition.

To compute the second measure, a "cosine" between two force fields was defined. This cosine measure was used to compare the costimulated field with the vectorially summed fields as described above. This measure quantified the proportionality of the two fields, that is, if the two fields were linearly scaled versions of eachother, then the cosine measurement would be one. This measure was found to be > 0.9, further supporting the linearity in superposition hypothesis.

In a follow up study by Lemay et.al. [41] a thorough examination of the the linearity in force production was conducted. The researchers examined not only the superposition property of force fields, but the scaling of force magnitude with stimulation as well. In contrast to the study by Mussa-Ivaldi et.al., the total force field and not simply the active component were examined.

The effects of varying stimulus parameters on force strength were examined. Multiple parameters effecting stimulus strength were systematically varied during electrical stimulus delivery to spinal sites. In most (16 of 17 spinal sites) cases, the average correlation coefficient between charge delivered and magnitude of resulting forces was > 0.9, implying that the relation between force strength and stimulus was linear.

To compare changes in force field strength, the ratio of force magnitudes for different stimulation levels were compared at each location. Before ratios were computed, the forces were normalized by the maximum force component. Under these conditions, if the force fields scaled linearly with stimulation the ratio would always be 1.0. The mean ratio was 1.1 ± 0.5. This result indicated the force fields scaled linearly.

The final comparison of interest was that of summation under costimulation. Just as in the study by Mussa-Ivaldi et.al., the hypothesis was that the field obtained by costimulating two sites was the vector sum of the field obtained by stimulating the individual sites. However, in this study more parameters were varied to better examine the hypothesis. The results again implied that the spinal force fields were linear in superposition.

3.3 Problem Formulation

The "linear" findings of experimental studies have been investigated through varied means. For instance, linearty has been examined through force fields whose strength increases linearly with stimulation, force fields whose orientation remains invariant with varying stimulation, and costimulation force fields that are approximately the vectorial sum of the individually stimulated force fields. Now we shall frame all this evidence under the mathmatical formalism of linearity in scaling and superposition. With this done, it becomes apparent that all the linearity results can be explained in terms of muscles that individually exhibit linear force production in activation.

For an isometric force field F, with activation vector α, (perhaps induced through a spinal stimulation) the property of linearity in scaling, is expressed as

$$F(c\alpha) = cF(\alpha) \tag{3.1}$$

where c is a scalar. Since we are considering isometric force fields, the end-point forces, at a specific location, are a function of muscle forces and joint angles only (as opposed to inertial forces of the skeletal links). Furthermore, since there are only two links in our hind limb model (experimentally, the hind limb is restrained to the same two degrees of freedom), the mapping from joint torques to end-point forces is unique, and we can equivalently examine joint torques to investigate end-point force linearity in stimulation.

Consider the torque about an individual joint due to a single muscle, $F_j^m(l_j, \alpha_j)$ $r_j(\theta)$. For the j^{th} muscle, r_j is the muscle moment arm corresponding to the joint angles θ, and l_j is the muscle length. To simplify the expression, the muscle length can be rewritten in terms of joint angles as well so that the expression for muscle force becomes, $F_j^m(\theta, \alpha_j) r_j(\theta)$. The total joint torque is the sum of the individual contributions from each muscle,

$$\tau = \sum_j F_j^m(\theta, \alpha_j) r_j(\theta) \tag{3.2}$$

Assume that the activations α_j are due to stimulation at a spinal site. Further assume that increasing the stimulation to the spinal site by c produces activations $\tilde{\alpha}_j$. Then the question of linearity in scaling is wheter the following relation holds,

$$\sum_j F_j^m(\theta, \tilde{\alpha}_j) r_j(\theta) = c \sum_j F_j^m(\theta, \alpha_j) r_j(\theta) \tag{3.3}$$

There are two possibilities under which this can fail: either increasing the stimulation does not linearly increase the activation, $\tilde{\alpha}_j \neq c\alpha_j$, or the muscle force generated is not linear in activation.

This thesis assumes that the former has not occured, that is we assume $\tilde{\alpha}_j$ is always equivalent to $c\alpha_j$. This is the assumption made in previous studies, and an examination of the robustness of this linearity between stimulation and activation is beyond the scope of this thesis.

Regarding the second possibility, that the muscle force is not linear in activation, would violate the following,

$$\sum_j F_j^m(\theta, c\alpha_j) r_j(\theta) = c \sum_j F_j^m(\theta, \alpha_j) r_j(\theta) \tag{3.4}$$

However, assuming the preceding relation does hold, for arbitrary c, α_j, θ, implies,

$$F_j^m(\theta, c\alpha_j) = cF_j^m(\theta, \alpha_j) \tag{3.5}$$

or, the force produced by each individual muscle is linear in activation. Thus, linearity in scaling can be assesed by examining the behavior of single muscles.

To examine the property of linearity in superposition, we can follow a similar line of reasoning. Again, we consider joint torques to examine end-point force field linearity in superposition. Assuming stimulation at a spinal location, site A, produces activations α_{Aj}, and stimulation at site B produces activations α_{Bj}, the individual stimulations produce joint torques,

$$\tau_A = \sum_j F_j^m(\theta, \alpha_{Aj}) r_j(\theta) \tag{3.6}$$

$$\tau_B = \sum_j F_j^m(\theta, \alpha_{Bj}) r_j(\theta) \tag{3.7}$$

Now, if costimulating at both sites A and B simultaneously produces activations α_{ABj}, then the resulting joint torque is,

$$\tau_{AB} = \sum_j F_j^m(\theta, \alpha_{ABj}) r_j(\theta) \tag{3.8}$$

Therefore, if superposition holds under these circumstances, $\tau_A + \tau_B = \tau_{AB}$. There are two possibilities for this to fail. Either costimulating at both sites, A and B doesn't produce activations that are the linear sum of the individual activations due to those site stimulations, or the msucle force generated is nonlinear in activation.

We assume the former has not occured, and that α_{AB} is always equivalent to $\alpha_A + \alpha_B$. The second possibility, that the muscle force is not linear in stimulation, would violate the following,

$$\sum_j F_j^m(\theta, \alpha_{Aj} + \alpha_{Bj}) r_j(\theta) = \sum_j [F_j^m(\theta, \alpha_{Aj}) + F_j^m(\theta, \alpha_{Bj})] r_j(\theta) \tag{3.9}$$

However, assuming the above relation holds, for arbitrary α_{Aj}, α_{Bj}, implies,

$$\begin{aligned} F_j^m(\theta, \alpha_{Aj} + \alpha_{Bj}) &= F_j^m(\theta, \alpha_{Aj}) + F_j^m(\theta, \alpha_{Bj}) \\ \Longrightarrow \quad & F_j^m(\theta, c\alpha_j) = cF_j^m(\theta, \alpha_j) \end{aligned} \tag{3.10}$$

which is the same condition for scalability described above.

Based on the implications for force fields at the limb end-point, that are linear in stimulation, under scaling and superposition, the question of linearity collapses to the level of individual muscle force production. With this in mind, the thesis will now posit a hypothesis for the observed linearity, and outline how this hypothesis will be examined.

3.4 Biomechanical Analysis

The musculo-tendon model used for this study is a nonlinear 3-element Hill model, as described in chapter 2. In this model an active force generating contractile element (representing muscle fiber force as a function of fiber position, velocity and activation level) in parallel with a passive elastic element represent the muscle fiber. Both elements are in series with another elastic element modeling the tendon. The details of the model's force length curves have been experimentally found (see chapter 2). Below a description of the model structure and behavior is presented.

3.4.1 Musculo-Tendon Model

The present analysis of force generation with varying activation levels is independent of muscle fiber or tendon dynamics. The properties we are investigating are static in nature. Therefore, just as velocity effects are irrelevant in isometric studes, the contractile element's fiber velocity dependence will not come into play for this analysis.

The active force generating, or contractile element, produces force as a function of activation state α, fiber length l^m and fiber velocity, v^m. The dependence on l^m and v^m are modeled independently by force-length and force-velocity relations. The total force is assumed to be a multiplicative relationship, $\alpha f_l(l^m) f_v(v^m)$. The force-length curve for the contractile element is based on single muscle fiber experiments. As the fiber length increases from zero, the force-length curve first ascends, reaches a peak, and then descends back to zero. The maximum tension muscle fiber can produce, maximally activated and with $v^m = 0$, is referred to as the maximum isometric tension, F_o^m. The length at which F_o^m occurs is referred to as the optimal muscle fiber length, l_o^m.

As previously mentioned, we shall make the assumption throughout, that stimulation and activation are linearly related at steady-state. The stimulus is assumed to elicit excitation-contraction dynamics that in turn elicit an activation state. These dynamics are known to have a nonlinear relation [73], however, because these nonlinearities are most accutely observed at the extrememes of muscle recruitment (for instance where saturation effects are observed), this nonlinearity is often neglected [70]. Furthermore, in the studies that have examined isometric force production, care has been taken to avoid the levels of stimulus that might induce these nonlinearities in muscle recruitment [25], [41].

In figure 3-1 the active force curve is plotted. Two active curves are depicted, one with maximal activation, and another with 50% activation (dashed lines). The active force profile begins to ascend from zero at small ($\approx 40\% l_o^m$) displacements, and descends back to zero for large ($\approx 160\% l_o^m$) displacements. As modeled, the active force component is nonlinear in length, but linear in activation state.

Both axis of figure 3-1 are normalized. The force is normalized with respect to the maximum isometric tension, F_o^m, while the length is normalized with respect to the optimal fiber length, l_o^m. Using this normalization, the force is maximal (1.0) when the normalized length is 1.0.

Also depicted in figure 3-1 is the parallel passive elastic element (dotted line). The passive force is also zero initially, but exponentially increases as the displacements increase from l_o^m. The passive and active components are in parallel, so their force contributions add. The sum of active and passive forces can be seen in the solid lines of figure 3-1 for the two activation levels. Note how the activation level does not influence the passive componenet; it is not linear in activation.

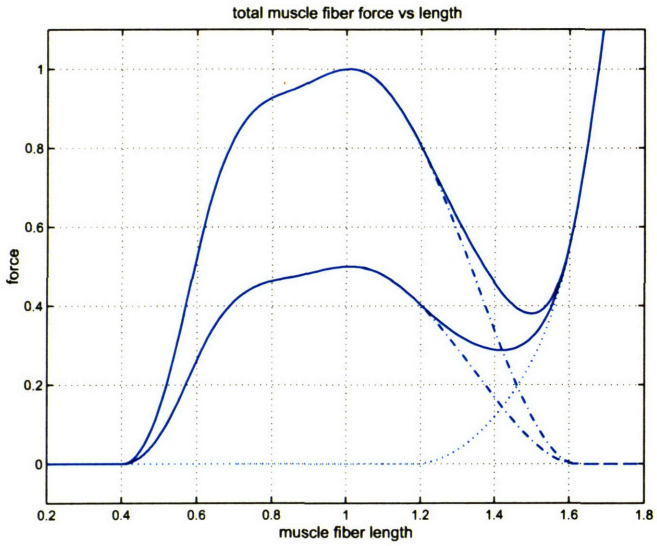

Figure 3-1: Muscle fiber force-length curves for activation levels 50% and 100% (solid lines). The dashed and dotted lines are the active and passive contributions, respectively.

In figure 3-2 is a plot of the normalized tendon force as a function of tendon strain. The tendon force is normalized with peak isometric tension, F_o^m. Tendon strain is computed as the ratio of tendon length to tendon slack length, $\varepsilon^T = l^T/l_s^T$. l_s^T is the length of the tendon when the tension across it is zero. We note here that tendon strain, when its force equals peak isometric fiber force, is assumed musculo-tendon independent. This strain value, denoted ε_o^T is often reported to be 0.033 (e.g. [73]). However, the tendon force-strain curve in our model is significantly more compliant. This tendon model predicts $\varepsilon_o^T \approx 0.05$. As will be shown later, this relatively compliant tendon has the effect of making our musculo-tendon model err on the side of nonlinear force production in activation, further supporting our eventual conclusions.

As explained in chapter 2, these three elements, the parallel passive elastic element, the contractile force generating element and the series elastic element, compose the 3-element Hill model we shall use in this analysis. When all three elements of the Hill model are coupled (the passive and contractile elements in parallel, and the series elastic element in series) the output force is no longer a simple algebraic function of the components. For a given muscle length and activation level, the tendon

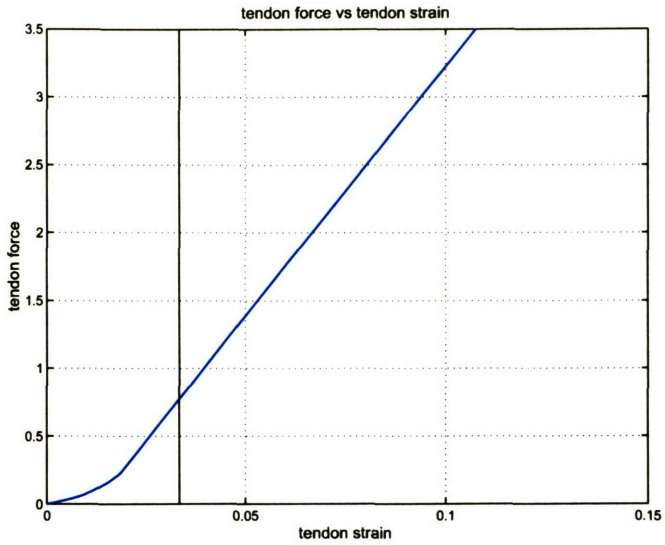

Figure 3-2: Normalized force-strain curve for tendon. Tendon force is normalized with maximal isometric tension, F_o^m, and muscle strain is normalized with optimal muscle fiber length, l_o^m.

strain and contractile and passive lengths are unknown. The muscle force for a given displacement must be found through simulation of the dynamic system (in which case an inertial element is added to the model, see chapter 2) or a numerical search algorithm. Regardless of the method, the muscle force, due to the interplay between all three elements, is not linear in activation.

The 3-element Hill model is characterized by three parameters (for now we are only concerned with static conditions so no dependence on fiber velocity is needed): the peak isometric tension, F_o^m; the optimal muscle fiber length, l_o^m; and the tendon slack length, l_s^T. By defining these three parameters we completely characterize the static properties of our musculo-tendon models.

Each musculo-tendon in the *Rana pipiens* hind limb model has distinct values for these parameters. Rather than examining each individual musculo-tendon, it would be convenient to find a stereotypical example for explanatory analysis. To do this we shall normalize the force-length curve for musculo-tendon with the three defining parameters, F_o^m, l_o^m and l_s^T. As we have done previously, force will be normalized with F_o^m. Temporarily, for illustrative purposes, displacements will be normalized with l_o^m (another more permanent choice will be made in the following section). The remaining free parameter, the tendon slack length, l_s^T, will also be normalized with l_o^m. This ratio of tendon slack length to optimal muscle fiber length will be a referred to as the "slack length ratio", denoted ε_s. By normalizing the force-length curves we have a framework with which to compare individual musculo-tendon models, without regard to the values of the independent parameters. The parameters for each musculo-tendon can now be collapsed down to one unique variable, ε_s.

3.4.2 Quantification of Linearity

We now procede to analyze the musculo-tendon model's force production as a function of activation. A few example force-length curves with different activation levels have been computed and plotted below (figures 3-3 and 3-4). Figure 3-3 displays a few normalized force-length curves for a musculo-tendon with a slack length ratio, ε_s, of .5 (the tendon slack length is half of the optimal muscle fiber length). Curves for three activation levels (100%, 50% and 10%) are shown in the solid blue lines. Overlaid on these curves in dashed red lines are linear extrapolations; these are the force profiles that would result if the musculo-tendon were linear in activation. The linear curve is created by scaling the maximally activated force length curve by the relevant activation level. For example, for the 50% activation curve, the linear profile is the maximally active curve scaled by 0.5.

Observe that for this particular model, ($\varepsilon_s = .5$), the force-length curve is very nearly linear over much of the displacement shown. It is only when the model begins to experience large displacements and the active force component can no longer contribute to the force output, that the musculo-tendon force begins to appreciably differ from the linear extrapolation. Once the muscle fiber has been strained past the point where the active force generation can no longer contribute, the musculo-tendon behaves as a purely passive element. Therefore, the force is no longer influenced by the activation state, and is not linear in activation.

In figure 3-4 we plot another set of force-length curves for a model with $\varepsilon_s = 5$. The three blue lines are the curves with activation levels 100%, 50% and 10%. In red dashed lines are the linear extrapolations. We note that relative to the previous model, with $\varepsilon_s = .5$, these curves deviate much more from linearity in activation. The tendon in this model is much more compliant (due to its length relative to the muscle fiber) and masks the linear increase in force of the active contractile element, making the overall force production more distinctly nonlinear.

A final remark should be made concerning the displacements shown in figure 3-4. Due to the normalization, l^m/l_o^m this musuclo-tendon experiences a smaller total displacement, relative to its total musculo-tendon length, than in the previous case where $\varepsilon_s = .5$. This was done for illustrative purposes, as it allows for force-length curves that are similar in appearance. However, we will choose a more appropriate normalization shortly.

The above plots demonstrate, qualitatively, how the force output can deviate from linearity. To quantify this behavior, we need to define a metric. The goal of this investigation is to analyze how musculo-tendon force scales in activation. In particular we are interested in knowing how close to linear this scaling is. Therefore, we shall define an error in linearity, E_L, to be:

$$E_L(l) = \frac{1}{F^m(1,l)}\sqrt{\sum^{n}(F^m(\alpha_i,l) - \alpha_i F^m(1,l))^2/n} \tag{3.11}$$

where $F^m(\alpha_i, l)$ is the musculotendon force with activation α_i and total musculotendon length l, and n is the number of distinct activations. E_L is the root mean

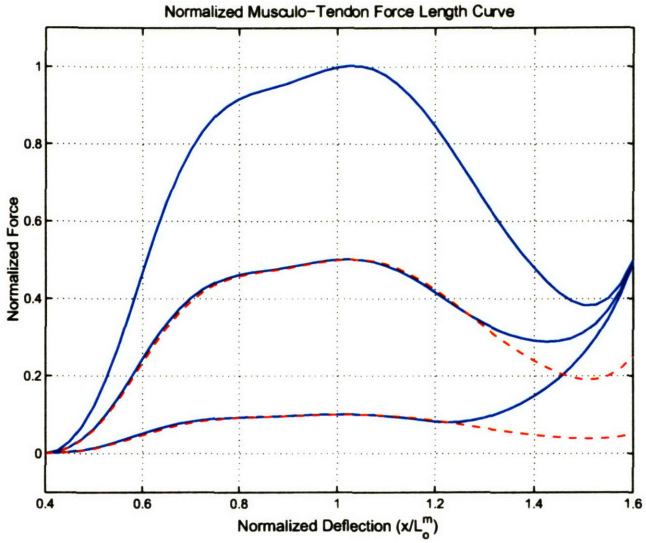

Figure 3-3: Normalized force-length curves for musculo-tendon, $\varepsilon_s = 0.5$. Overlaid, in dashed red, are the linearly extrapolated curves.

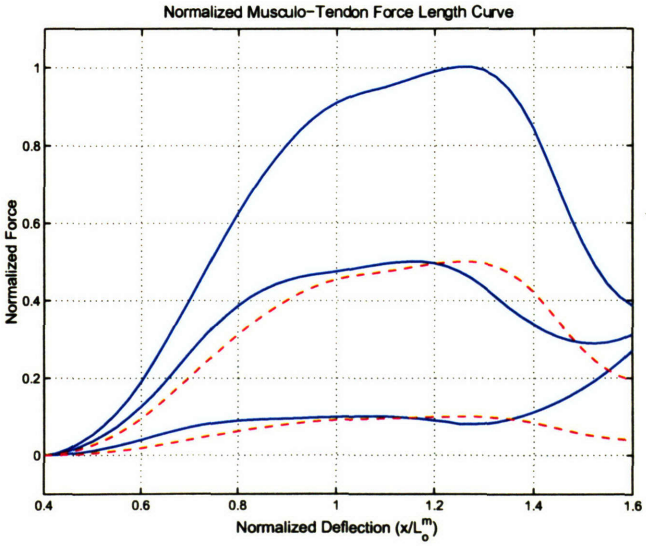

Figure 3-4: Normalized force-length curves for musculo-tendon, $\varepsilon_s = 5.0$. Overlaid, in dashed red, are the linearly extrapolated curves.

square error from a linearly scaled force, normalized by the force at maximal activation. Note that for a given ε_s, this error is a function of musculo-tendon length only.

This error in linearity, or E_L score, is computed for the musculo-tendons of figures 3-3 and 3-4 and plotted in figure 3-5. This is done using activation levels of 100%, 75%, 50%, 25%, and 10%. In the top plot we see $E_L(l)$ for the case when $\varepsilon_s = .5$. The bottom plot is for $\varepsilon_s = 5$. We note three trends in these curves, which are also consistent with the earlier plots. First, the curve computed for the musculo-tendon with a longer slack length has a greater error score; it is relatively less linear. Second, the error scores are greater for both models as the musculo-tendon experiences large displacements. Third, the error score is generally low when the deflection is in the vicinity of the opimal fiber length; i.e. the muscle force behaves most linear about the optimal muscle fiber length.

We have computed the error in linearity score for two specific tendon slack length ratios, $\varepsilon_s = .5, 5$. However, we would like to use this E_L score to predict behavior for arbitary musculo-tendon models, so we sweep through varying values of ε_s and compute the score for various displacements. Again, this is done using activation levels of 100%, 75%, 50%, 25% and 10%. The result is a surface, plotted in figure 3-6.

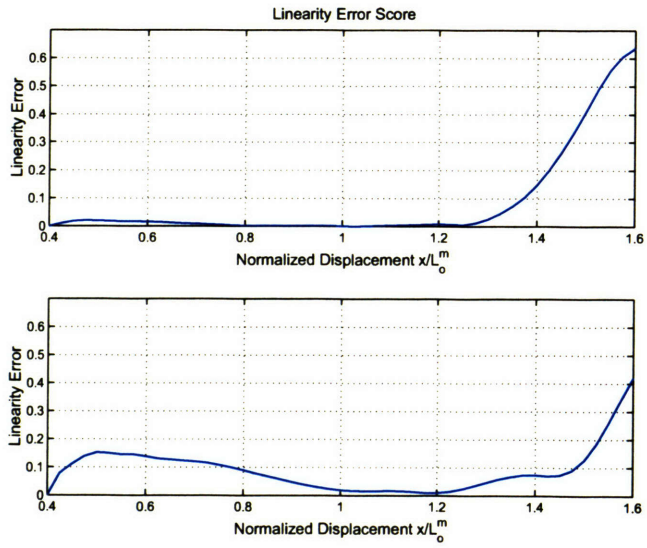

Figure 3-5: Error in Linearity score, $E_L(l)$, for musculo-tendon models depicted in figures 3-3 and 3-4.

The surface of figure 3-6 is the E_L score a musculo-tendon model will exhibit for a given set of parameters, F_o^m, l_o^m, l_s^T and a particular displacement l. The horizontal axis are the total musculo-tendon displacement normalized with respect to the sum $l_o^m + l_s^T$ and the tendon slack ratio, ε_s. Normalizing the musculo-tendon displacement with respect to the sum $l_o^m + l_s^T$, rather than l_o^m, has the effect of normalizing the displacement across musculo-tendons relative to their total length.

Note the following observations of this E_L score surface. For small slack length ratios, say $\varepsilon_s < 0.5$, the E_L score is small (the musculo-tendon is approximately linear in stimulation) for all but very large displacements. As the slack length ratio and displacement increase, the E_L score generally gets larger. As the normalized displacement increases past 1.0, there is an abrupt change and the E_L score surface quickly increases. This happens as the E_L score approaches values around 0.5. This happens relatively qucikly for large ratios, ε_s. This indicates the location where the musculo-tendon model's passive characteristics, defined by the parallel and series elastic elements, begin to dominate in force production. For all musculo-tendons, the E_L score eventually saturates at the same level, $\approx .65$; this is the E_L score value when the musculo-tendon is, within numerical accuracy, purely passive and no longer varies with activation.

There is also a region of moderately large E_L scores for large ε_s and small normalized muscle displacments. This can be seen in the left corner of figure 3-6. We note, however, that in this region the musculo-tendon's force is very close to zero and the nonlinear effects are therefore insignificant.

With this surface, we can now predict when the force a musculo-tendon model produces will, or will not behave linearly in activation. In addition, due to the E_L score's particular normalization, we can also use this surface to predict the error in linearity of the joint torques or limb end-point forces a musuclo-tendon generates. This invariance can be seen by considering the joint torque due to a musculo-tendon at a given fixed orientation, and for a single joint,

$$\tau = F^m(l, \alpha) r(\theta) \tag{3.12}$$

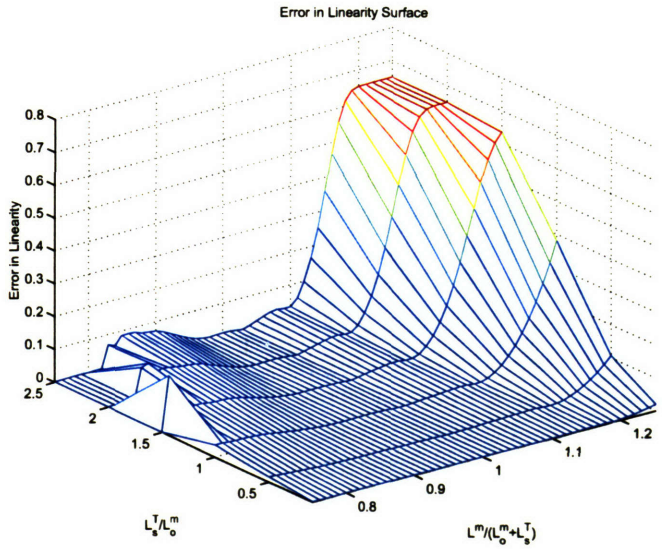

Figure 3-6: $E_L(l)$ surface plot. The horizontal axis are the normalized tendon slack length, and normalized displacement. The discontinuities observed in the left corner are due to round-off errors associated with muscle forces that are nearly zero.

where $r(\theta)$ is the muscle's joint moment arm and τ is the joint torque. The E_L score in terms of joint torque is then,

$$\begin{aligned} E_L(l) &= \frac{1}{\tau(1,\theta)}\sqrt{\sum(\tau(\alpha_i,l)-\alpha_i\tau(1,l))^2/n} \\ &= \frac{1}{f(1,\theta)r(\theta)}\sqrt{r^2(\theta)\sum(f(\alpha_i,l)-\alpha_i f(1,l))^2/n} \\ &= \frac{1}{f(1,\theta)}\sqrt{\sum(f(\alpha_i,l)-\alpha_i f(1,l))^2/n} \end{aligned} \tag{3.13}$$

The same line of reasoning demonstrates that the error in linearity for end-point forces, for example measured at the ankle, is also equivalent to this score. That is, although we have devised this linearity measure in terms of muscle force, it applies to the end-point force field as well.

Now, in order to illustrate the linearity analysis we shall perform on the *Rana pipiens* model, as well as validate the utility of the E_L surface for linearity predictions, we shall examine a "toy" model system.

3.5 Example Linearity Analysis

A toy musculo-skeletal system, simpler than, but analgous to the *Rana pipiens* model, shall be examined. The model will be created to exhibit desired linearity properties based on the previously found E_L surface. This is done for two reasons. First, it will serve as an illustrative example of how we shall analyze the linearity of musculo-skeletal system. Second, this effort will demonstrate the utility of the E_L surface in quantifying the nearly linear, or lack thereof, behavior of the 3-element Hill model without the need for computationally intensive simulations.

We shall design our system to have four muscles. Two muscles will be biarticular, with large and small ε_s, .5 and 2. The remaining two musculo-tendons will be uniarticular, also with large and small ε_s, .5 and 2. To distinguish between them, we shall refer to the musculo-tendons as Uni flexor, Uni extensor, Bi flexor and Bi extensor.

For each musculo-tendon a point on the error surface will be chosen in order to fix its linear behavior. This will define the slack length ratio's (ε_s) value, and the average musculo-tendon length in terms of l_o^m and l_s^T. Figure 3-7 is a top view of figure 3-6 with the location of the four designed muscles's linearity score superposed. Three isoclines (lines of constant value) are also depicted on figure 3-7. In blue, green and red are lines indicating a linearity score of 0.5, 1.0 and 1.5 respectively. These values were chosen for descriptive purposes as they indicate the approximate region where the error surface quickly ascends to large values, corresponding to a nonlinear operating region. As can be seen from this figure, this set of design parameters will allow for a relatively linear biarticular and uniarticular musculo-tendon, and a relatively nonlinear biarticular and uniarticular musculo-tendon. The biarticular extensor and uniarticular flexor will on average operate in relatively linear regions, while the uniarticular extensor and biarticular flexor, on the other hand, will on average operate in nonlinear regions.

As described, the musculo-tendon parameters, l_o^m and l_s^T have now been defined in terms of the average musculo-tendon length. To determine this average length, insertion points for each musculo-tendon must be chosen. The insertion points will define the mechanical actions of the musculo-tendon on the limb, knee flexor, hip extensor, etc. The model, constructed using the SIMM application, can then be used to find the musuclo-tendon length as a function of joint angles across the limb's workspace. Once this is determined, the average musculo-tendon lengths are computed and l_o^m and l_s^T are solved for. F_o^m is chosen to be unity for each muscle in this example system.

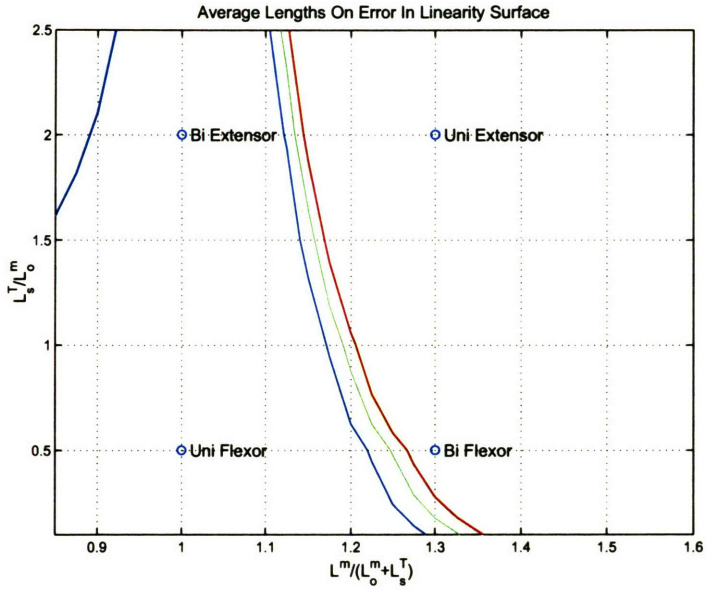

Figure 3-7: Location of average musuclo-tendon length on the E_L score surface, for the four toy muscles.

The skeletal links used for the model are those of the *Rana pipiens* model (see chapter 2). The wrap objects at the hip and knee are included in the model. The muscle insertion points are chosen by hand to be consistent with the musculo-tendon actions, i.e. extensor, flexor, uniarticular and biarticular. Figure 3-8 is an image of the SIMM model.

The SIMM application is used to compute the musculo-tendon excursions as a function of joint location. Workspace plots of the musculo-tendon lengths are shown in figure 3-9. The average excursion for each musculo-tendon is computed and used with the previously defined location on the E_L score surface plot to find l_o^m and l_s^T.

With the musculo-skeletal model completely defined, the force fields can be simulated and the linearity scores computed. To find the fields for each musculo-tendon, the isometric force experiments are performed virtually. In simulation, the limb is held in a fixed position and various activation levels are commanded to individual muscles. The limb is then moved (in simulation) to another fixed location in the

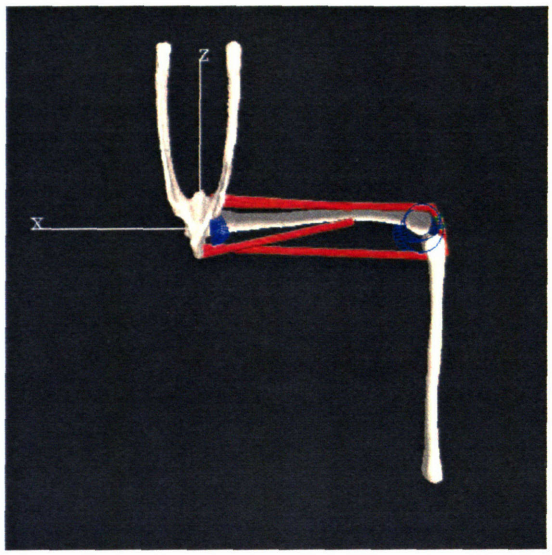

Figure 3-8: Figure of the toy musculo-skeletal system as modeled in SIMM.

workspace and the same activation levels are commanded again.

The simulations are performed using a grid over hip and knee joint angles. Each joint was discretized with 9 evenly spaced points ranging from -80^o to 80^o. In total, 81 locations are used to find isometric forces. At each location, activation levels, 0%, 10%, 25%, 50%, 75% and 100% are commanded for approximately 2 seconds. The resulting forces at the limb endpoint (the ankle) are recorded. The steady-state force is found and assigned to its respective activation level and joint location. When the force fields for each muscle and for each activation level are found, the error in linearity, E_L, is computed using equation 3.11.

The E_L score is computed and displayed in figure 3-10, in cartesian workspace coordinates. Uni Flexor, the muscle with a short tendon ($\varepsilon_s = .5$), designed to be linear in activation, has the lowest E_L scores across the worksapce of all the muscles. This can be explained by examining the musculo-tendon length across the workspace (figure 3-9) in conjunction with the E_L surface (figure 3-7). Although the E_L scores for the model were obtained by simulating the model's force fields, at any point in the workspace, the same score can be found by using figure 3-9 to find the muscle length and then finding the corresponding score on the E_L surface. Therefore, from figure 3-7 we see that Uni Flexor's average musculo-tendon length is, amongst the four muscles, the furthest distance from the contour lines indicating large E_L scores (a nonlinear operating region) for a muscle with its tendon length. Similarly, from figure 3-9 we observe that the largest (normalized) displacements Uni Flexor experiences, at the left-most region of the workspace, are not large enough to exhibit a nonlinear effect. Uni Flexor's maximum normalized displacement, approximately 1.15, is too short to enter into the nonlinear region of the E_L surface.

Bi Flexor, even though it also has a short tendon ($\varepsilon_s = .5$) does achieve large E_L

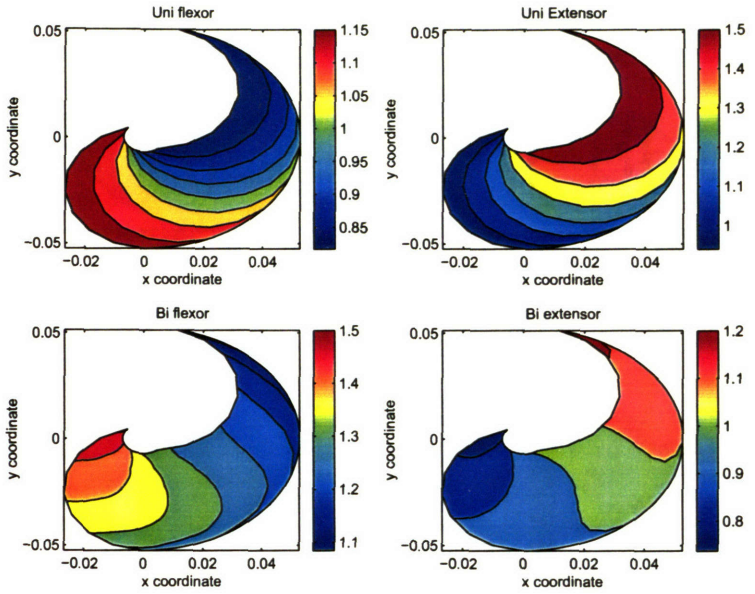

Figure 3-9: Musculo-tendon lengths of toy muscles across workspace.

scores. Just as with Uni Flexor, this can be explained in terms of the musculo-tendon lengths achieved across the workspace. First we note the average musculo-tendon length is to the right of the contours in figure 3-7. Thus, for large portions of the workspace, its length will enter nonlinear regions. Second, observe from figure 3-9 that Bi Flexor's displacments are, relative to Uni Flexor, large. In figure 3-7 we see that for $\varepsilon_s = .5$, muscles begin to enter a nonlinear working region around a normalized length of 1.2. From figure 3-9 we observe that Bi Flexors displacements reach this value in the left side of the workspace, starting with the yellow region. Comparing the musculo-tendon length and linearity score (figures 3-9 and 3-10) one can easily observe the correspondence between musculo-tendon length and error in linearity.

Uni Extensor, and Bi Extensor, are nonlinear and linear, respectively, over the majority of their workspace. Just as with Uni Flexor and Bi Flexor, this behavior of the muscles can be also be explained by observing where in the workspace each musculo-tendon achieves lengths that place it in nonlinear regions.

These results demonstrate several points. First we've outlined our method for quantifying the linear force production in simulated isometric force fields. This analysis and quantification will be done for the *Rana pipiens* model next. Furthermore, we've shown how the error in linearity is a function of musculo-tendon length across the workspace, as was suggested by our analysis in section 3.4.

The analysis has also established how the E_L score surface can be utilized to asses the linear behavior of muscles with respect to activation. The surface of figure 3-6 was found by examining muscle force generating properites. However, becuase of its invariance with respect to coordinate transformations, it can also be used to examine end-point forces. What's more, as long as the 3-element Hill model is a

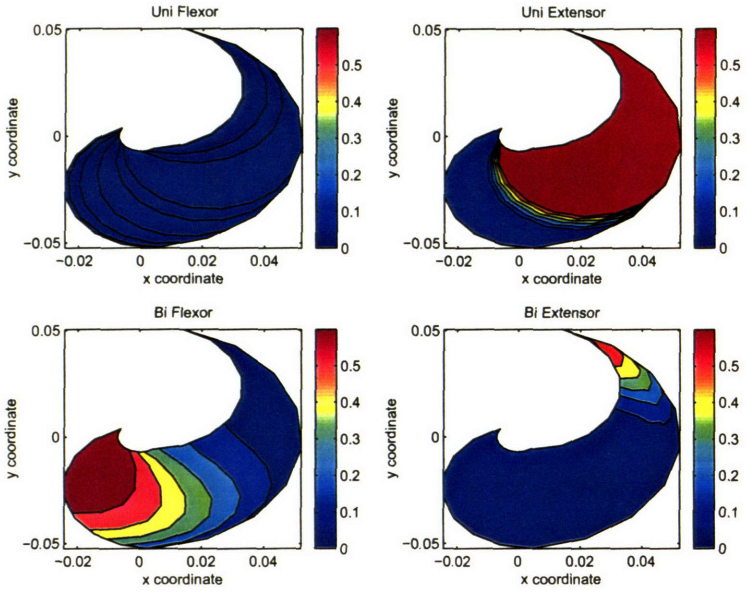

Figure 3-10: Plot of error in linearity scores across workspace, for toy model.

good characterization of the system in question, this surface can be used to predict linearity in force production in a computationally efficient manner, without resorting to simulations for investigating force fields.

We also emphasize that the linearity in force production (at the muscular level or in end-point forces) in activation, varies with location. That is, a muscle may behave linearly in one region of its workspace, and nonlinearly in another. This however, is not a statement about the variation of muscle force with length. As this chapter and the previous one have clearly presented, muscle forces are nonlinear in muscle length and velocity. With that said, a similar analysis to the one just performed may find that muscle forces are approximately linear across the workspace for muscles that experience small displacements. This analysis, however, is beyond the scope of the thesis.

3.6 Linearity predictions for *Rana pipiens*

We have demonstrated our approach to analyzing the linearity in force production on a toy musculo-skeletal system. Now we use the same method to examine the model of *Rana pipiens*. As discussed previously, the frog hind limb exhibits a significant degree of linearity in stimulation when examined experimentally. We want to know if this can be explained solely in terms of biomechanical properties rather than involving properties of efferent feedback or neural control. Therefore, just as in the previous the section, the simulated isometric force fields for each muscle shall be determined across a range of activations. These force fields will then be used to compute our defined error in linearity score, E_L. Note that the endpoint force field is always the

sum of the force fields of the individual muscles it is comprised of. Therefore the linearity, or lack thereof, for the individual musculo-tendons will help explain the overall behavior.

Without some understanding of their origin, the force field simulations are little more than another empirical phenomena. Hence, just as was done in the previous section, the E_L scores shall be examined and explained in terms of the musculo-tendon displacements. By examining each musculo-tendon's excursion across the workspace, the error score can be used to explain how the linearity, or lack thereof, arises, and where in the workspace it is expected.

We use SIMM to simulate the isometric force fields. The model we use, discussed in chapter 2, contains detailed descriptions of the musculo-skeletal geometry including insertion points for the musculo-tendons, and wrap objects to simulate the joint surfaces that the muscles come into contact with. To find the fields, the isometric force experiments are performed virtually. The analysis of linearity performed here, explained in the previous section, is meant to replicate the performed experiments. In addition to simulating the force field for each musculo-tendon, its length across the workspace is also obtained. This is used to analyze the sources of linearity.

In figure 3-11 the result of one such force field simulation for Semimembranosus is plotted. The left half of the figure is a plot of the force vectors found at each of the simulated locations when the activation level was 50%. The force vectors display the direction the ankle "pushes" when held at a particular location. Semimembranosus is a hip extensor. The right half of the figure is a plot of the magnitude of these same forces. Note that the magnitude of the force vectors does not change much except near the boundaries of the workspace. Close to the edges, where the joint angles are approaching kinematic singularities (they occur at $\pm 90^o$) the end-point forces quickly increase in magnitude.

For a given muscle, once the force field for all the activation levels is found, the linearity score can be computed. This was done just as defined in equation 3.11, using the 6 activation levels, 0%, 10%, 25%, 50%, 75% and 100%.

The E_L score for all 13 muscles in the *pipiens* model is presented in Appendix B. Most of the muscles have low error scores across most of the workspace. The exceptions are Semitendinosus Dorsal and Ventral, Tensor Facia Latae and Iliacus Internus. Both heads of Semitendinosus attain large error scores in the rostral-lateral region of the workspace, where the hip is flexed and the knee is extended. Tensor Facia Latae and Iliacus Internus attain large values in the medial most region of the workspace where the hip and knee are extended.

All 13 musculo-tendon's lengths across the workspace were also computed (see Appendix B). These are plotted by sweeping through hip and knee ankles from -80^o to 80^o. Each plot is normalized by the sum $l_o^m + l_s^T$.

To summarize the simulated isometric force field findings two plots are presented. One is a plot of the E_L score averaged across the workspace, for each musculo-tendon. This average is shown in figure 3-12. This plot illustrates the points made above. Both heads of Semitendinosus, Tensor Facia Latae and Iliacus Internus are the most nonlinear of the 13 muscles evaluated. The rest behave very nearly linear over most of the workspace.

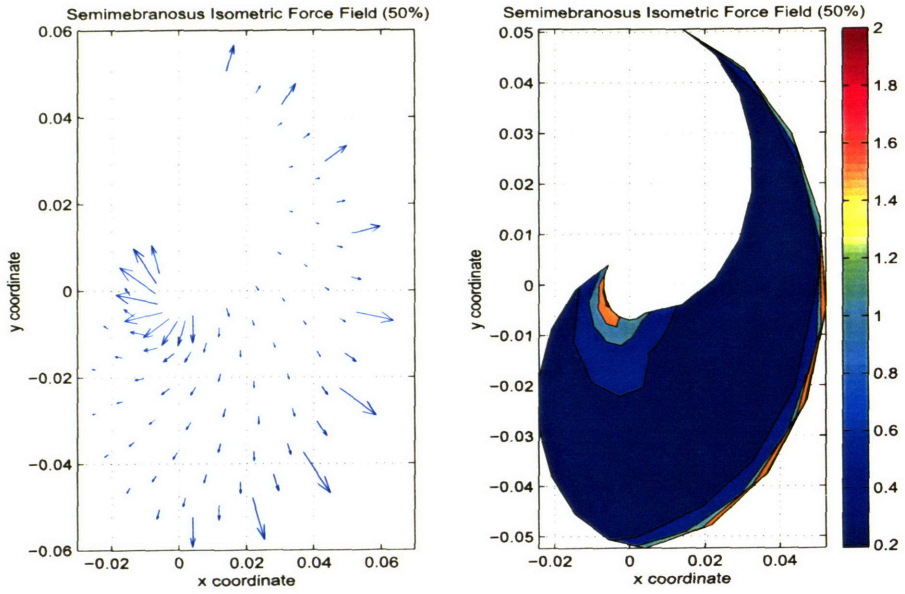

Figure 3-11: Results of isometric force field simulation for Semimembranosus with an activation level of 50%.

To further summarize the results of the linearity analysis, the location of each musculo-tendon on the E_L score surface is plotted. Using the musculo-tendon lengths found previously, an average length over the workspace is computed for each muscle. This average, along with the ratio ε_s is used to locate each musculo-tendon on the E_L score surface. The result is shown in figure 3-13. Observe that the averaged location of each musculo-tendon, with the notable exception of Semitendinosus Dorsal and Ventral, are within the low-scoring region of the surface (to the left of the contours that delineate 0.5, 1.0 and 1.5).

To explain these results and the relationship between linearity and musculo-tendon length, we shall examine two representative muscles, Gluteus and Semitendinosus Dorsal. In figure 3-14 the linearity scores of these two muscles are displayed. Gluteus and Semitendinosus Dorsal are examples of the best and worst E_L scoring muscles respectively. Note how Gluteus has a low linearity error score throughout all of the computed workspace. Semitendinosus Dorsal has a low score over the medial half of the workspace, where the hip is extended, but quickly makes a transition to high linearity error scores in areas where the hip is flexed.

In figure 3-15 are plots of the musculo-tendon lengths of Gluteus and Semitendinosus Dorsal across the workspace. Gluteus, a hip and knee flexor, has a relatively short normalized musculo-tendon length (≈ 0.8) at the lateral edge of the workspace, where the hip and knee are extended. In the medial region of the workspace, around the location of the hip ((x, y) coordinates (0,0)) the knee is fully flexed and the normalized lengths are larger (≈ 1.0). Now, we examine the E_L surface in figure 3-13 and the correspondance between these scores and Gluteus' musculo-tendon lengths across

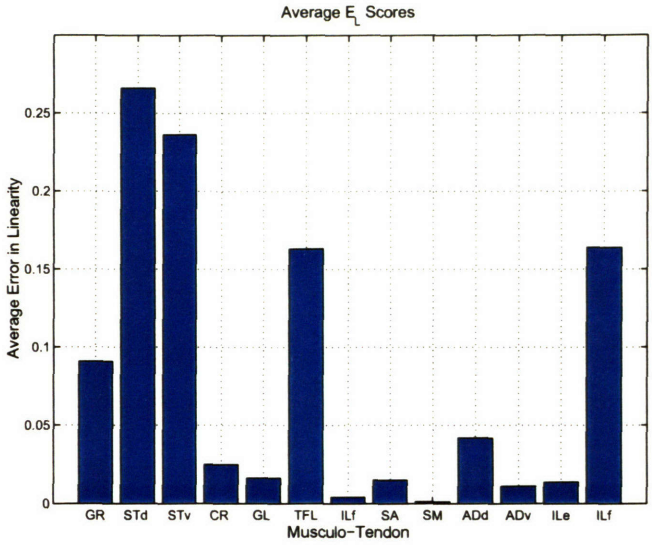

Figure 3-12: E_L scores averaged across workspace for each muscle.

the workspace. For musculo-tendon lengths between 0.8 and 1.0, and an appropriate ε_s (as indicated by the vertical location of Gluteus on the plot) the E_L scores are far to the right of the nonlinear contours. Therefore, according to this hind limb model, Gluteus never attains lengths that would allow for nonlinear force generation in activation.

Now, for Semitendinosus, a hip and knee extensor, the same analysis can be applied. In the right region of the workspace, the normalized musculo-tendon length is long, ≈ 1.2, while in the left region it is short, ≈ 0.8. When figure 3-13 is examined for these musculo-tendon displacements, we see that Semitendinosus Dorsal clearly moves to the right of the delineating contours. Therefore, according to this hind limb model, Semitendinosus Dorsal attains lengths in the medial side of the workspace, that dictate it generates forces that are nonlinear in activation.

The same analysis based on musculo-tendon length can be used to explain the nonlinearity results for all the other muscles examined.

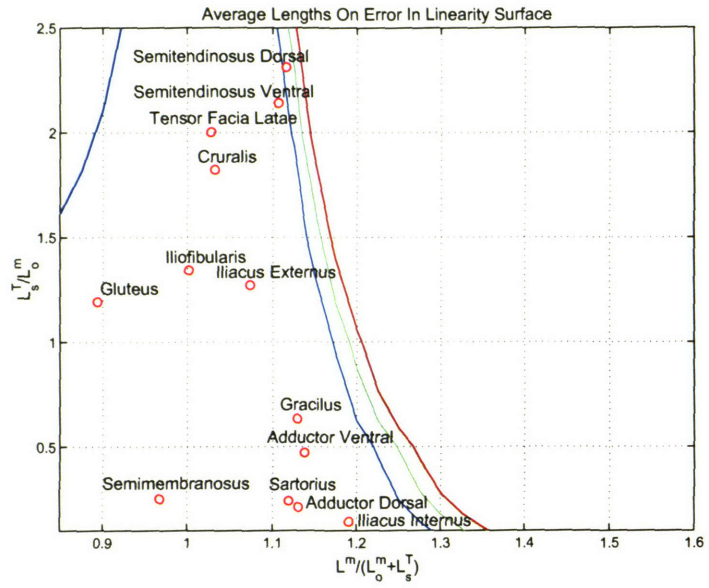

Figure 3-13: Average musculo-tendon excursions plotted on error in linearity surface.

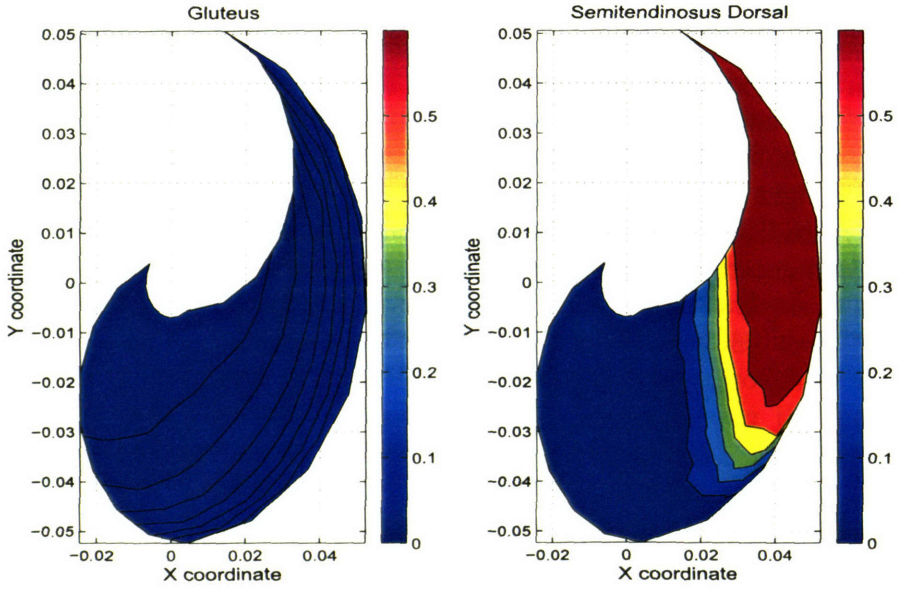

Figure 3-14: Example plots for simulated error in linearity scores of Gluteus and Semitendinosus Dorsal.

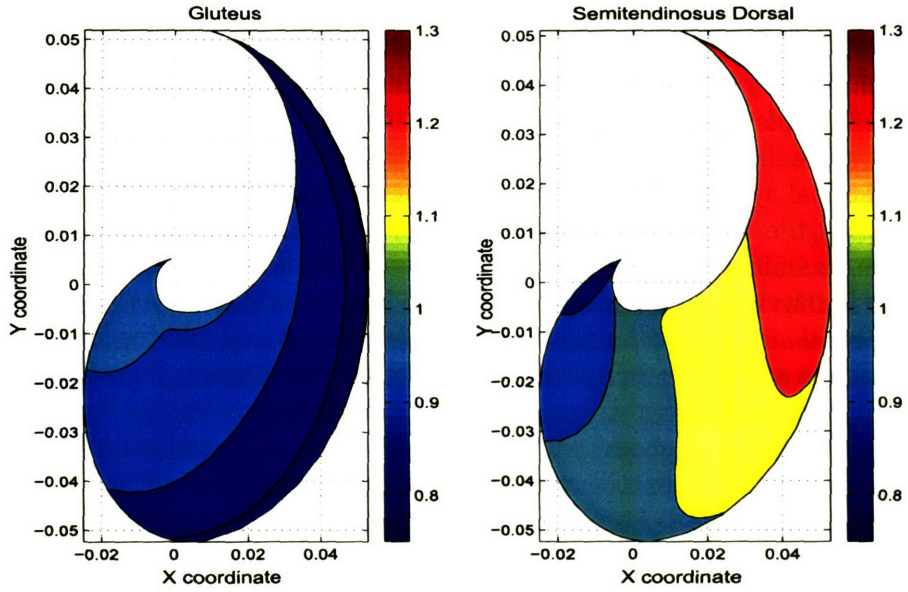

Figure 3-15: Example plots for musculo-tendon lengths across workspace for Gluteus and Semitendinosus Dorsal.

3.7 Summary

In this chapter we have analyzed the reported ability of isometric force fields to scale linearly with stimulation. We have shown how this linear behavior can be explained in terms of the linearity of force production at the level of individual muscles and we offered the following hypothesis: linear force production in stimulation can be explained in terms of the natural biomechanical properties of the limb's musculo-skeletal system.

To quantify this linearity, which we have couched in terms of activation, we defined a measure, or score, of linear force production with respect to activation. We have termed this score an error in linearity, or E_L. Through analysis of a normalized musculo-tendon model, we found a relation between muscle length and the error in linearity. This was used to compute an error in linearity surface, which could then be used to predict the E_L score for any musculo-tendon model.

Through examination of a toy musculo-skeletal system, we analyzed the musculo-tendon's ability to generate forces that scaled linearly in activation, and demonstrated by counterexample, the ability to deviate from this linearity. We thus concluded that the model's predictions for linear force production depend on the structure of the musculo-skeletal system. The linear behavior is a biomechanical result of the length of tendon and the operating region of the muscle.

We then examined the model for the frog hind limb of *Rana pipiens*. The isometric force fields for the individual muscles of this model were simulated and the results were interpreted and analyzed using the E_L score. Many of the musculo-tendons were found to predict a low error in linearity score (meaning they were nearly linear in activation) throughout much of the workspace. Some muscles, though attaining large E_L scores, usually did so in areas of the worksapce that were not examined experimentally, suggesting their results are still consistent with experimental evidence.

Our results imply that isometric force fields are, to a large extent, linear in activation over much of the limb's workspace, and most of the the area examined experimentally. From this we conclude that the linearity observed experimentally does not require feedback linearizing signals or other forms of motor controllers. Indeed, the results support our hypothesis, that the limb's linear force production is a natural result of biomechanics.

Chapter 4

Model Order Reduction

The dynamical system being studied in this thesis, the frog hind limb, is high-dimensional and highly-redundant. The limb has 13 muscles, many of which have very similar influences on the dynamics. As has been previously motivated, we would like to investigate a controller for this system that utilizes a simplified, low-dimensional representation, or reduced order model.

The process of obtaining a low-dimensional approximation to a high-dimensional dynamical system is known as model reduction, or model order reduction. Ideally, the reduced representation will be helpful for both estimating and controlling of the full-dimesional system. There are many methods for linear (time invariant) model order reduction. Under the heading, Krylov-based, there are many widely used algorithms (references). These are so called moment matching methods, designed to solve the coefficients of an expansion of the system's impulse response. Some algorithms, for instance the Arnoldi and Lanzcos procedures, are numerically stable and efficient. However, these approaches have no gurantees for estimation properties or stability of the resulting reduced model (references).

Another class of techniques known as projection methods, attempt to find subspaces of the original phase space upon which to project the system dynamics. The Karhunen-Loeve decomposition and balanced truncation are two popular methods within this class. The former is a popular method, in part because it is numerically feasible for large dimensional systems (e.g. 10,000 states or more, references). The later has the distinct advantage of carrying guarantees for both estimation and stability properties of linear systems (references). We will focus on these two methods, both of which have been applied for use with nonliear system model reduction.

Two more methods closely related to balanced truncation are Hankel norm and singular perturbation approximations. However, these methods are not as ammenable to applications with nonlinear systems.

With few exceptions, there is little theory for nonlinear model order reduction ([60]), and few examples of nonlinear solutions ([51]). Despite this short-coming, there have been many succesful examples of linear techniques in the application to problems of nonlinear model order reduction (e.g. [12], [39], [53]). We will briefly review some of the applicable nonlinear theory and techniques from the literature.

4.1 Proper Orthogonal Decomposition

Proper orthogonal decomposition (POD), or more formally, Karhunen-Loeve decomposition, (but also known as principal component analysis, factor analysis, and total-least-squares-estimation) is a method for finding the "best" approximating subspace, in terms of data variance. POD (in contrast to balanced truncation) is data-driven, e.g. the results are inherently a function of a specific data set. The POD method minimizes the distance between a data set and it's projection onto a k-dimensional subspace, $\mathcal{S}_k \subset \mathbb{R}^n$. A linear basis for this subspace defines k low-dimensional variables. There are at least two ways to motivate this, through a singular value decomposition, or an eigen decomposition.

4.1.1 The Singular Value Decomposition Approach

Assume we have a set of data, or "snapshots" of our dynamical system,

$$\mathcal{X} = [x_1, x_2, ...x_N] \in \mathbb{R}^{n \times N} \tag{4.1}$$

where $N \gg n$ and we assume $\mathcal{X}$ has full row rank. Let

$$x_i = x'(t_i) - \bar{x}, \;\; \bar{x} = \frac{1}{N} \sum x'(t_j), \;\; x' \in \mathbb{R}^n \tag{4.2}$$

where $x'(t_i)$ are state values of our dynamical system of interest at discrete times t_i, $i = 1, ..., N$. If we want to find a "best" approximation to this data, of a lower rank, k, then we can look for the approximation that minimizes the 2-norm,

$$\|\mathcal{X} - \hat{\mathcal{X}}_k\|_2 \tag{4.3}$$

where $\hat{\mathcal{X}}_k$ is our rank k approximation to $\mathcal{X}$. The optimal approximation, in the 2-norm sense, has the following cost,

$$\min \|\mathcal{X} - \hat{\mathcal{X}}_k\|_2 = \sigma_{k+1}(\mathcal{X}) \tag{4.4}$$

where σ_{k+1} is found from the singular value decomposition,

$$\mathcal{X} = U\Sigma V^*, \;\; \text{where} \;\; UU^* = I, \;\; VV^* = I, \;\; \sigma_1 \geq ... \geq \sigma_k \geq ... \geq \sigma_n \tag{4.5}$$

and a *non-unique* optimal approximation is found as,

$$\hat{\mathcal{X}}_k = U_k \Sigma_k V_k^* = \sum^{k} \sigma_i u_i v_i^* \;\; \text{where} \;\; U_k \in \mathbb{R}^{n \times k}, \;\; \Sigma_k \in \mathbb{R}^{k \times k}, \;\; V_k^* \in \mathbb{R}^{k \times N} \tag{4.6}$$

the orthogonal vectors u_i are referred to as the POD modes. U_k is a basis for the k-dimensional subspace we are projecting the data onto, $\mathcal{S}_k$.

4.1.2 The Eigen Decomposition Approach

Assume we have the same data set, $\mathcal{X}$, but now we wish to minimize the Euclidean distance between each data point and its projection onto some k-dimensional subspace,

$$\min \sum \|x_i - P_k x_i\|^2, \quad P_k : \mathbb{R}^n \longmapsto \mathbb{R}^n \tag{4.7}$$

where the projection matrix is constructed, as the method's name suggests, with an orthogonal basis (or modes) set, $\{\varphi_1, \varphi_2, ...\varphi_k\}$,

$$P_k = \sum_k \varphi_i \varphi^*, \quad \text{and} \quad \varphi_i^T \varphi_j = \delta_{ij} \tag{4.8}$$

the optimal modes can be found once we construct the (approximate) correlation matrix

$$R = \mathcal{X}\mathcal{X}^* = \sum_i^N x_i x_i^* \tag{4.9}$$

now the optimal modes (in the sense above) are the eigenvectors of the correlation matrix,

$$R\varphi_i = \varphi\lambda_i \tag{4.10}$$

P_k constructed this way uses the eigenvectors of the correlation matrix for a basis of the k-dimensional subspace. $P_k x_i$ is the inner product of x_i with each of the eigenvectors, projected onto these k bases.

$$P_k x_i = \sum_j^k \varphi_j(\varphi_j^* x_i) = \sum_j^k \varphi_j z_j \tag{4.11}$$

4.1.3 The Connections

First we note that $\hat{\mathcal{X}}_k$ and $P_k\mathcal{X}$ are equivalent. Since both $\hat{\mathcal{X}}_k$ and $P_k\mathcal{X}$ have the same singular values, they both optimize the cost given in equation 4.3, which as mentioned, is not unique. Because they are equivalent, they both minimize the second cost given in equation 4.7. Also note that the eigenvectors of the correlation matrix, or the modes, are the same modes that were found with the SVD approach,

$$\mathcal{X} = U\Sigma V^*, \quad R = \mathcal{X}\mathcal{X}^* = U\Sigma\Sigma^* U^* \tag{4.12}$$

Both approaches yield a reduced basis (or modes) for representing the data. In this reduced basis space, the approximation error is minimized, and the "energy" in the approximation, as measured by the correlation matrix, (the singular values retained) is maximized. We can map from the ambient $\mathbb{R}^n$ to the k-dimensional subspace, $\mathcal{S}_k$

with P_k (note that $U_k U_k^* \neq I_n$, but $U_k^* U_k = I_k$). The data elements have effectively been reduced to k variables where previously there were n,

$$\hat{x}_i = U_k(U_k^* x_i) = U_k z_i \text{ where } z_i \in \mathbb{R}^k \tag{4.13}$$

We will refer to z as our low dimensional variable.

We can summarize the state reduction as follows: we map our state variable to a low-dimensional variable, z with the relation, $z = U_k^* x$. This low-dimensional variable can then be used to map "up" to an (low-dimensional) estimate of the state variable, as follows: $\hat{x} = U_k z$. We can identify an exact relationship between these variables as,

$$x = \hat{x} + x_r = U_k z + x_r \tag{4.14}$$

where x_r is the residual, an error in the estimate of x that is in the null space of U_k^*.

With the low-dimensional subspace identified, the next step is to model the dynamics of the low-dimensional variables. For linear systems this is straight forward, but for arbitrary systems this is not the case. The method most commonly used is Galerkin projection.

4.2 Galerkin projection

All of the above was to justify our low dimensional representation of data, but now we would like to find a low dimensional representation for a dynamical system,

$$\dot{x} = f(x, u), \quad y = g(x) \tag{4.15}$$

or in the LTI case,

$$\dot{x} = Ax + Bu, \quad y = Cx \tag{4.16}$$

The Galerkin projection replaces the dynamics of the system with the approximate dynanics of the low-dimensional variable by projecting the vector field onto the subspace $\mathcal{S}_k$. Consider the relation,

$$x - x_r = U_k z \tag{4.17}$$

where x_r is the residual error associated with our projection. Taking the derivative we find,

$$\dot{x} - \dot{x}_r = f(x, u) - \dot{x}_r = U_k \dot{z} \tag{4.18}$$

or, rearranging,

$$U_k \dot{z} = f(U_k z + x_r, u) - \dot{x}_r \tag{4.19}$$

At this point, the relation above is exact. The approximation results from assuming that the residual's effects upon the dynamics are negligible, and neglecting any dependence on it,

$$\dot{z} \approx U_k^* f(U_k z, u) \tag{4.20}$$

In the linear case, to find our reduced POD model we start with the following relation,

$$\begin{aligned} x - \bar{x} - x_r &= U_k z \\ x - \bar{x} &\approx U_k z \end{aligned} \tag{4.21}$$

where $\bar{x}$ is the stationary mean of the data set, which we assume may not be negligible. Continuing,

$$\begin{aligned} \dot{x} = U_k \dot{z} &= Ax + Bu \\ \dot{z} &= U_k^* A x + U_k^* B u \\ \dot{z} &\approx U_k^* A U_k z + U_k^* A \bar{x} + U_k^* B u, \quad y = C U_k z + C\bar{x} \end{aligned} \tag{4.22}$$

With this relation we can either define a quadruplet for the reduced system,

$$\{A_k, B_k, C_k, D_k\} = \{U_k^* A U_k, \ [U_k^* B, U_k^* A], \ C U_k, \ [0, C]\} \tag{4.23}$$

with motor command $u_z = [u^T, \bar{x}^T]^T$, or we can carry an extra state, $\bar{x}$, whose derivative is zero. In this case our system is unstable (has a free integrator), but the command is equivalent to the original, $u_z = u$.

4.3 Balanced Truncation

For this method we begin by finding a balanced realization for our LTI system. A balanced state-space representation has equivalent and diagonal controllability and observability gramians. The controlability gramian for an LTI system may be defined as follows,

$$\mathcal{P} = \int_o^\infty e^{At} B B^T e^{A^T t} dt \tag{4.24}$$

The controllabilty gramian is a measure of the amount each state is excited by the command. For instance, for two states, w, v, if $\|w\| = \|v\|$ and $w^T \mathcal{P} w > v^T \mathcal{P} v$ then w is more "controllable" than v; it requires a smaller input to drive the system from rest to w. The controllability gramian can also be used to find the minimum command required to drive the system from some initial state, x_o, to zero,

$$\min_{u(t), x(0)=x_o} \int \|u(t)\|^2 dt = \frac{1}{2} x_o^T \mathcal{P}^{-1} x_o \tag{4.25}$$

In addition to the controllability gramian, for LTI systems we can define the observability gramian,

$$\mathcal{Q} = \int_0^\infty e^{A^T t} C^T C e^{At} dt \tag{4.26}$$

The observability gramian is a measure of how much each state excites the output. For instance, if $x(0) = x_o$ and $u(t) = 0$, $\|y\|_2^2 = x_o^T \mathcal{Q} x_o$; states which excite larger output are more observable.

The gramians are also the symmetric positive definite solutions to the Lyapunov equations,

$$\begin{aligned} A\mathcal{P} + \mathcal{P}A^T + BB^T &= 0 \\ A^T\mathcal{Q} + \mathcal{Q}A + C^TC &= 0 \end{aligned} \tag{4.27}$$

Now suppose the system coordinates are such that the both grammians are diagonal and equivalent,

$$\tilde{\mathcal{P}} = \tilde{\mathcal{Q}} = \Sigma = \text{diag}(\sigma_1, \sigma_2, ...\sigma_n) \tag{4.28}$$

Systems with this property are said to be balanced. The entries on the diagonal of Σ are defined as the Hankel singular values, they are a property of the system dynamics and independent of coordinates. The Hankel singular values are the square roots of the eignevalues of the product $\mathcal{PQ}$,

$$\sigma_i = \sqrt{\lambda(\mathcal{PQ})} = \lambda(\hat{\mathcal{P}}) = \lambda(\hat{\mathcal{Q}}) \tag{4.29}$$

When the Hankel singular values are ordered, $\sigma_1 \geq \sigma_2 \geq ... \geq \sigma_n$ then the i^{th} state is more easily influenced by the command, and its influence on the output is more easily observed, than the $(i+1)^{th}$ state. The balanced realization is, in effect, an ordering of the states in terms of energy in the output, as measured by the Hakel singular values.

Any system that is both observable and controllable can be balanced. We proceed by finding an invertible coordinate transform $\tilde{x} = Tx$, yielding a new description of the dynamics,

$$\{\tilde{A}, \tilde{B}, \tilde{C}\} = \{TAT^{-1}, TB, CT^{-1}\} \tag{4.30}$$

The new gramians are then,

$$\tilde{\mathcal{P}} = T\mathcal{P}T^* \quad \tilde{\mathcal{Q}} = (T^{-1})^*\mathcal{Q}T^{-1} \tag{4.31}$$

we can use the above to find a relation for T,

$$T\mathcal{P}\mathcal{Q}T^{-1} = \Sigma^2 \tag{4.32}$$

which is the eigenvalue problem. We observe that the matrices $\mathcal{P}$ and $\hat{\mathcal{P}}$ are not similar. $\lambda(\mathcal{P}) \neq \lambda(\hat{\mathcal{P}})$. The coordinate transformation has altered the eignevalues.

The same holds for $\mathcal{Q}$ and $\hat{\mathcal{Q}}$. In addition, recall that eigenvectors are not unique. Therefore the above relation is not sufficient to find the coordinate transform T (many matrices T will satisfy the above relation, but not all of them will diagonalize the individual matrices, $\mathcal{P}$ and $\mathcal{Q}$).

The gramians, being symmetric positive definite, can be factored as $\mathcal{P} = UU^*$ and $\mathcal{Q} = LL^*$, where U and L are upper and lower traingular matrices respectively. Denoting the singular value decomposition of U^*L as $Z\Sigma Y^*$, then $T = \Sigma^{1/2}Z^*U^{-1} = \Sigma^{-1/2}Y^*L^*$. With T sovled for, the system can be balanced. We now express the balanced system as,

$$\tilde{A} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}, \quad \tilde{B} = \begin{bmatrix} B_1 \\ B_2 \end{bmatrix}, \quad \tilde{C} = \begin{bmatrix} C_1 & C_2 \end{bmatrix} \tag{4.33}$$

where the state is ranked from largest to smallest Hankel singular values. To find a reduced order model, we use the Galerkin method with the projection matrix $[I_k \; 0]$, that is, we simply truncate after the k^{th} state, where the σ_{k+1} is an appropriately small Hankel singular value. The modes of our projection are the states with the appropriately large Hankel singular values. The low-dimensional variable is the fist k states, $z = [\tilde{x}_1, \tilde{x}_2, ..., \tilde{x}_k]^T$. The new system triplet is,

$$\{A_k, B_k, C_k\} = \{A_{11}, B_1, C_1\} \tag{4.34}$$

4.3.1 Balancing Coordinates

A brief description of the coordinates of balanced systems wil help to clarify how they are related to the controllability and observability gramians. Consider a point mass with two degrees of freedom, acted upon by two force sources. We will choose the velocity in the x and y direction as our state, assume unity damping to ground and unity mass, and define the dynamics as,

$$\dot{x} = \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix} x + Bu, \quad y = Cx \tag{4.35}$$

For this system the gramians are easily calculable. Expressing B and C as follows,

$$B = \begin{bmatrix} b_1 & b_2 \\ b_3 & b_4 \end{bmatrix}, \quad C = \begin{bmatrix} c_1 & c_2 \\ c_3 & c_4 \end{bmatrix} \tag{4.36}$$

the controllabity gramians is,

$$\mathcal{P} = \int_0^\infty e^{-2t} BB' dt \;=\; \frac{1}{2} BB' \;=\; \frac{1}{2} \begin{bmatrix} b_1^2 + b_2^2 & b_1 b_3 + b_2 b_4 \\ b_1 b_3 + b_2 b_4 & b_3^2 + b_4^2 \end{bmatrix} \tag{4.37}$$

and the observability gramian is,

$$\mathcal{Q} = \int_0^\infty e^{-2t} C'C dt \;=\; \frac{1}{2} C'C \;=\; \frac{1}{2} \begin{bmatrix} c_1^2 + c_3^2 & c_1 c_2 + c_3 c_4 \\ c_1 c_2 + c_3 c_4 & c_2^2 + c_4^2 \end{bmatrix} \tag{4.38}$$

Now, if B and C are the identity matrices, then the Gramians are equivalent and diagonal ($\frac{1}{2}I$); the system is balanced and the balanced coordinates cooincide with our original ones. If $b_1 = 2, b_2 = b_3 = 0, b_4 = 1$ and C is the identity matrix, then the gramians are,

$$\mathcal{P} = \begin{bmatrix} 2 & 0 \\ 0 & .5 \end{bmatrix}, \quad \mathcal{Q} = \begin{bmatrix} .5 & 0 \\ 0 & .5 \end{bmatrix} \tag{4.39}$$

they are still diagonal, but not equal, so the system is not balanced, and a change of coordinates is needed. In figure 4-1 are pictorial representations for both gramians and the transformed coordinates. The dotted ellipse is the "shape" of the controllability grammain; those are the points we get when we map the unit circle through the Gramian. The ellipse denoted by x's is the shape of the observability gramian. The arrows indicate how the standard basis vectors (unit vectors along the horizontal and vertical axis) are mapped to the balanced coordinates. We make the following observations. The balanced coordinates are orthogonal and colinear with the original coordinates (the original gramians are diagonal so no rotation is needed). The coordinate along the vertical axis has not changed; both gramians have the same eigenvalue/singular value along that coordinate so no change is needed. The balanced coordinate along the horizontal axis is scaled down, to compensate for the magnification in that direction the controllabilty gramian is causing; the controllability gramian's eigenvalue/singular value along that coordinate is larger than the observability's.

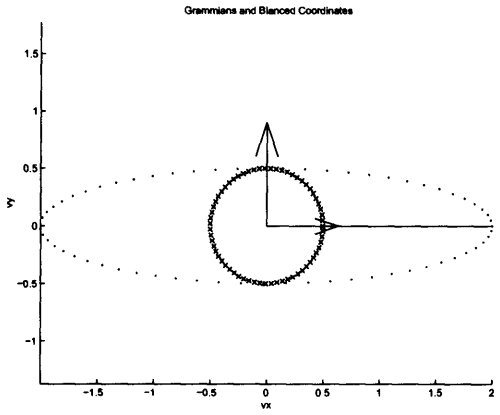

Figure 4-1: Controllability and observability gramians graphically depicted along with balanced coordinates.

Now consider the case where $b_1 = \alpha, b_2 = b_3 = 0, b_4 = 1$ and we let α attain large values,

$$\mathcal{P} = \begin{bmatrix} \alpha^2/2 & 0 \\ 0 & .5 \end{bmatrix}, \quad \mathcal{Q} = \begin{bmatrix} .5 & 0 \\ 0 & .5 \end{bmatrix} \tag{4.40}$$

the horizontal axis of the balanced coordinate will continue to shrink as α increases, while the vertical axis remains constant.

Now let $B_{ij} = C_{ij} = \delta_{ij} + \alpha$, both gramians have their axis rotated through 45^o and are equivalent (figure 4-2). The balancing transformation is simply a rotation that brings the coordinates in line with the major and minor axis of the gramians. The columns of the balancing transformation, T are orthonormal.

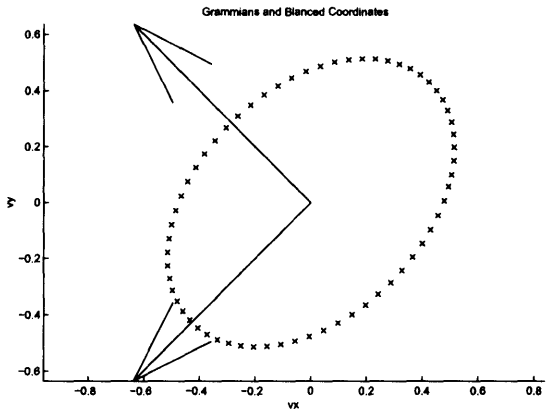

Figure 4-2: Equivalent gramians rotated through 45 degrees and their balaned coordinates.

However, if C is the identity matrix, but B is as above, a rotation of the original axis won't balance the system. A scaling of the coordiantes is again needed. But, the balanced coordinates, pictorially, do not appear to be scaled along the major and minor axis of $\mathcal{P}$, figure 4-3.

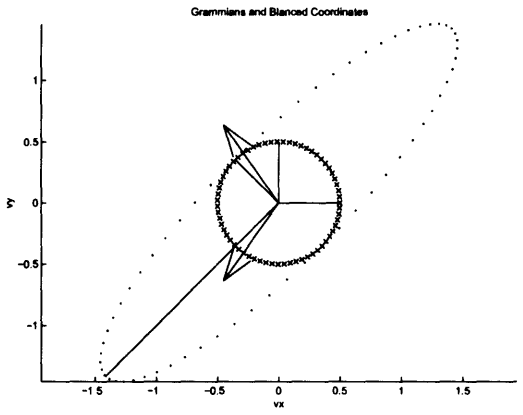

Figure 4-3: Gramians, graphically depicted, and their balanced coordinates.

The balanced coordinates *are* rotated, and *then* scaled along their respective dimensions. The balancing transform can be viewed as follows,

$$T = \begin{bmatrix} \alpha_1 & 0 \\ 0 & \alpha_2 \end{bmatrix} \begin{bmatrix} cos(\theta) & sin(\theta) \\ -sin(\theta) & cos(\theta) \end{bmatrix} \tag{4.41}$$

When both gramians are non-diagonal and unequal, the new coordinate frame is rotated and scaled appropriately.

4.4 Singular Perturbation Approximation

We now briefly review singular perturbation approximation[1]. The method assumes the existence of a balanced realization (the system must be observable and controllable) but instead of truncating those states with low Hankel singular values, they are fixed at their steady-state values. This makes intuitive sense considering these states have relatively little energy in the output. The dynamics can be found as follows,

$$\begin{aligned} \dot{\tilde{x}} &= \begin{bmatrix} \dot{x}_k \\ 0 \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} x_k \\ x_{k-n} \end{bmatrix} + \begin{bmatrix} B_1 \\ B_2 \end{bmatrix} u \\ y &= \begin{bmatrix} C_1 & C_2 \end{bmatrix} \tilde{x} \end{aligned} \tag{4.42}$$

solving for the dynamics of x_k we find the quadruplet,

$$\{A_k, B_k, C_k, D_k\} = \{A_{11} - A_{12}A_{22}^{-1}A_{21},\ B_1 - A_{12}A_{22}^{-1}B_2,\ C_1 - C_2A_{22}^{-1}A_{21},\ -C_2A_{22}^{-1}B_2\} \tag{4.43}$$

We note that in contrast to the balanced truncation method, the above equations will predict the correct steady-state values.

4.5 Reduced Model Properties

Assuming a low-dimensional representation for the full-dimensional dynamical system has been found, we would like to know what kind of behavior we can expect; how good are the estimates our reduced system creates? is the reduced system even stable? For example, POD is optimal at approximating our data set, but the POD model is not necessarily the best approximation of the dynamics responsible for that data. Infact, we don't know if the POD model will even be stable. It could be that low "energy" dynamics that were left out of the reconstruction are crucial to the system behavior. Here we shall review the guarantees that exist for the methods described above in application for the reduction of linear time invariant systems.

Denoting the system transfer function as $G(s) = C(sI - A)^{-1}B$, we have the following guarantee ([2]),

$$\sigma_{k+1} < \|G - G_k\|_\infty^2 \tag{4.44}$$

where σ_{k+1} is the $k + 1^{th}$ Hankel singular value. This lower bound, holds for *any* reduced order model of dimension k.

[1]This method is distinct from attempts to design controllers that drive nonlinear dynamics onto stable manifolds, see [22]. Though the names are similar, the singularly perturbed manifold theory is not a model order reduction process.

For the model obtained by the POD method the above lower bound holds. However, we have no further guarantees. Even if the original system is stable, the reduced model could be unstable. What's more, there is no guarantee that increasing the dimension of the reduced model will lead to better system properties. An increase in dimension will lead to better approximations of the *data set*, however increasing the dimension may drive the reduced dynamic model unstable. The estimation and stability properties of the reduced model obtained through the POD method are a function of the particular data set used to create the model and have to be scrutinized on a case by case basis.

System properties for the reduced model obtained through the balanced truncation model are not as bleak. In addition to the lower bound, we also have the following guarantee ([2]),

$$\|G - G_k\|_\infty^2 < 2\left(\sigma_{k+1} + ... + \sigma_n\right) \tag{4.45}$$

And as can be inferred from above, with some minor restrictions, stability is preserved[2]. This can be proved by showing that the reduced system equations also satisfy the Lyapunov equations. The reduced system may produce errors at steady state. This short-coming can be avoided by enforcing that steady-state solutions are in the range of the balancing transformation, $A^{-1}B \in \mathcal{R}(T)$.

One last note, a balanced truncation is not optimal (in the sense give above). The reduced model has the upper error bound given above, but other better approximations may exist. However, when the Hankel values drop off quickly in magnitude, then σ_k may be much greater than σ_{k+1}, and the guarantees ensure the reduced model will have small infinity norm bounds.

The reduced model obtained through singular perturbation approximation has the same guarantees as the balanced and trucated model. In addition, the reduced model will predict the correct steady state solutions. However, because the system matrices have been manipulated to retain the correct steady-state (low frequency) behavior, the reduced model tends not to approximate the full order dynamics at high frequencies as well as the balanced and truncated model.

4.6 Nonlinear MOR

Up to this point the information presented has focused on model order reduction for linear systems. There is precious little theory for nonlinear model order reduction. Some of the ideas presented above have natural extensions to nonlinear systems, but lack the mathmatical formalism to merit a generally applicable technique (ref's, [60]). This is not to suggest nonlinear model order reduction is not feasible. The literature contains examples of reduced nonlinear models, but the approaches are tailored to the particular models.

It should be noted that there are examples in the literature of nonlinear balancing by way of linearization: nonlinear systems are linearized about operating points and

[2]if $\sigma_k \neq \sigma_{k+1}$ then stability is assured.

used to create linear time varying systems with balanced realizations (references). However, this method is not strictly a nonlinear method and the information presented above suffices.

As discussed earlier, the POD reduction method is inherently a data-driven method. This being the case, a transformation to a low-dimensional variable (and subspace) can be found regardless of the linear, or nonlinear character of the system dynamics. A Galerkin projection can be computed as well. Aside from the uncertainty of the resulting system dynamics (estimation and stability properties) this procedure may actually increase the computations required to simulate the system dynamics. Consider again the resulting dynamical relation,

$$\dot{x} = f(x,u) \longrightarrow \dot{z} = Vf(Wz,u) \tag{4.46}$$

where V and W are the appropriate maps to and from spaces. Though the state has been reduced, we have incurred extra computations, namely, Wz and $Vf(\cdot)$, to compute the rate of this reduced state. If these equations cannot be collapsed (as in the case of working with a large simulation where the rates are computed via software) then the advantage of a state reduction is called into question. Consequently, it would be advantageous to further reduce the dynamics as well. An appropriate function approximation technique could be employed (e.g. an artificial neural network).

Given the advantages a balanced truncation has over a POD model, any line of reasoning that will balance a nonlinear system is tempting. There are formal extensions of balancing to nonlinear systems. Unfortunately, the theory does not readily give rise to solutions. We will briefly review some of the theory here.

For nonlinear systems there are no gramians to be computed, however, we can define controllability and observability functions,

$$L_c(x_o) = \min_{u \in L2} \int_{-\infty}^{0} \|u\|^2 \, dt, \quad \text{with } x(-\infty) = 0, \quad x(0) = x_o \tag{4.47}$$

$$L_o(x_o) = \frac{1}{2} \int_{0}^{\infty} \|y\|^2 \, dt, \quad \text{with } x(0) = x_o, \quad u(t) = 0 \tag{4.48}$$

$L_c(x_o)$ is the minimum amount of control "energy" required to reach state x_o, and $L_o(x_o)$ is the amount of output "energy" generated by x_o. These functions do not necessarily exist. If x_o cannot be reached, L_c is infinite, and if the system is unstable, then L_o too is infinite.

For nonlinear systems of the form,

$$\dot{x} = f(x) + b(x)u, \quad y = g(x) \tag{4.49}$$

if the controllability and observability functions exist, and are smooth functions of x, then in a neighborhood of an asymptotic equilibria, say $x = 0$, these functions are solutions to the partial differential equations,

$$\frac{\partial L_c}{\partial x} f(x) + \frac{1}{2} \frac{\partial L_c}{\partial x} b(x) b^T(x) \frac{\partial^T L_c}{\partial x} = 0, \quad L_c(0) = 0 \tag{4.50}$$

$$\frac{\partial L_o}{\partial x} f(x) + \frac{1}{2} g^T(x) g(x), \quad L_o(0) = 0 \tag{4.51}$$

and in this neighborhood, a coordinate transformation exists, such that in the new coordinates the controllability and observability functions are diagonal in their arguments. With this coordinate transform one can define singular value functions.

The above equations, are deceptively compact. Equation 4.50 is a Hamilton-Jacobi-Bellman equation. Solving it is equivalent to solving a family of optimal control solutions for the system. With rare exceptions this solution is intractable.

Despite the difficulties associated with the formal approach to nonlinear balancing, a method for approximate balancing of nonlinear systems has been proposed. This method, known as empirical balancing, attempts to find approximate gramians about some operating point of the nonlinear system, and find a transformation that locally balances the nonlinear system. Empirical gramians will be presented below.

4.6.1 Empirical Gramians

In this section we will first review methods for obtaining empirical gramians for linear systems. We will then present a method that is applicable to both linear and nonlinear systems.

As mentioned earlier, for the LTI case, the controllability gramian is defined as,

$$\mathcal{P} = \int_o^\infty e^{At} B B^T e^{A^T t} \, dt = \int_0^\infty \left(e^{At} B\right) \left(e^{At} B\right)^T dt \tag{4.52}$$

where $e^{At}B$ is simply the impulse response of the system. We can express the integral in an alternate but equivalent form as,

$$\int_0^\infty \left(x_1(t) x_1^T(t) + ... + x_r(t) x_r^T(t)\right) dt \tag{4.53}$$

where $x_i(t)$ is the impulse response from input i^{th} input,

$$x_i(t) = e^{At} b_i \tag{4.54}$$

Now the integral is recognizable as the correlation matrix of the impulse responses.

If instead, we have "snapshots" of the system, at N points in time, we can organize the data as such,

$$\mathcal{X} = [x_1(t_1), x_1(t_2), ..., x_1(t_N), x_2(t_1), ..., x_r(t_N)] \quad \in \mathbb{R}^{n \times rN} \tag{4.55}$$

Then the approximate (modulo a constant) correlation matrix or equivalently, the controllability Gramian is,

$$\mathcal{X}\mathcal{X}^* \approx \mathcal{P} \tag{4.56}$$

Returning to the observability gramian,

$$\mathcal{Q} = \int_0^\infty e^{A^T t} C^T C e^{At} \, dt = \int_0^\infty \left(e^{A^T t} C^T\right) \left(e^{A^T t} C^T\right)^T \, dt \tag{4.57}$$

which we observe is the correlation of the impulse responses of the adjoint system,

$$\dot{z} = A^T z + C^T v \tag{4.58}$$

Therefore, following the strategy above, we construct a data set,

$$\mathcal{Y} = [z_1(t_1), z_1(t_2), ..., z_1(t_N), z_2(t_1), ..., z_m(t_N)] \quad \in \mathbb{R}^{n \times mN} \tag{4.59}$$

where z_i is the impulse response from the i^{th} input,

$$z_i(t) = e^{A^T t} c_i \tag{4.60}$$

The approximate correlation matrix/observability gramian is,

$$\mathcal{Y}\mathcal{Y}^* \approx \mathcal{Q} \tag{4.61}$$

If we have followed the procedures outlined above, we can use the matrices, $\mathcal{X}$ and $\mathcal{Y}$ to find a balancing transform. Instead of computing the transformation from the observability and controllability gramians (which may not be full rank) as presented in the balanced truncation section, we find a SVD of the product of these matrices,

$$\mathcal{Y}^T\mathcal{X} = \begin{bmatrix} U_1 & U_2 \end{bmatrix} \begin{bmatrix} \Sigma_r & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} V_1^T \\ V_2^T \end{bmatrix} = U_1 \Sigma_r V_1^T \tag{4.62}$$

where Σ_r has rank r (no relation to the dimension of our control signal). The balancing transformation and its inverse are,

$$T_r = \mathcal{X} V_1 \Sigma_r^{-1/2}, \quad T_r \in \mathbb{R}^{n \times r} \tag{4.63}$$

$$S_r = \Sigma_r^{-1/2} U_1^T \mathcal{Y}^t, \quad S_r \in \mathbb{R}^{r \times n} \tag{4.64}$$

The approaches above assume we can manipulate the state dynamics. Specifically, they assume we can find solutions to the adjoint system's impulse responses. However, if we cannot do this (if the system is not available to us) or if the system is nonlinear, then there is another method. This method has been suggested for use in balancing nonlinear systems, but it does reduce to the correct gramians when applied to linear systems (references).

Let $\mathcal{T}^n$ be a set of q orthogonal $n \times n$ matrics $\{\mathcal{T}_1, \mathcal{T}_2, ..., \mathcal{T}_q\}$, and let $\mathcal{M}$ be a set of positive constants, $\{c_1, c_2, ..., c_s\}$, and let $\mathcal{E}^n$ be the set of standard unit vectors in $\mathbb{R}^n$, $\{e_1, e_2, ..., e_n\}$.

With the previous definitions, and the sets $\mathcal{T}^r, \mathcal{E}^r$, we can define an approximate, or empirical controllability gramian by,

$$\hat{\mathcal{P}} = \sum_{l}^{q} \sum_{h}^{s} \sum_{i}^{r} \frac{1}{q s c_h^2} \int_0^\infty \Phi^{ilh}(t) dt \tag{4.65}$$

where

$$\Phi^{ilh} = (x^{ilh}(t) - \bar{x}^{ilh}) \ (x^{ilh}(t) - \bar{x}^{ilh})^T \tag{4.66}$$

and x^{ilh} is the state corresponding to the impulse command $u(t) = c_h \mathcal{T}_l e_i \delta(t)$.

To find the empirical observability gramian, using sets $\mathcal{T}^n, \mathcal{E}^n$, we have

$$\hat{\mathcal{Q}} = \sum_{l}^{q} \sum_{h}^{s} \frac{1}{qsc_h^2} \int_0^\infty \mathcal{T}_l \Psi^{lh}(t) \mathcal{T}_l^T dt \tag{4.67}$$

where

$$\Psi_{ij}^{lh}(t) = (y^{ilh}(t) - \bar{y}^{ilh})^T \ (y^{jlh}(t) - \bar{y}^{jlh}) \tag{4.68}$$

and y^{ilh} is the output corresponding to the inital condition response with $x_0 = c_h \mathcal{T}_l e_i$. Keep in mind, if this is applied to a nonlinear system about equilibrium, then the mean values are simply the values at equilibrium. Similarly, when applied to linear systems, the mean values are simply zero.

With this method, we require that we can indpendently define initial conditions on the system, and that we can input impulsive commands.

4.6.2 Reduced Order Dynamics Modeling

To find a reduced order model, for a nonlinear dynamical system, there are several approaches. We shall break the process into two steps; finding a transformation to a balanced and reduced state, and modeling the dynamics of the resulting reduced state.

For a nonlinear system, there are multiple ways to compute a reduced order state through a balanced trunctation. The nonlinear system could be linearized about one, or many, operating points. The linearized system/s could then be exactly balanced and then truncated.

If only a single operating point were utilized, the balancing transformation would only be valid over the region the linearization was valid. This could limit the effectiveness of the balancing process if the system dynamics were not amenable to a local linearization. However, this approach greatly simplifies the necessary computations and results in a single, linear transformation between the full and reduced state.

If multiple operating points were used to obtain the balancing, the resulting transformations to a reduced state would be linear but time varying. A means of switching between the different reduced states according to the current region of state space would be required. Extra memory is required to retain all the multiple transformations as well. The resulting reduced order state would certainly be superior to the previous example. Since the time varying transformation can be utilized to discretize state space at arbitary levels of precision, the balancing, in theory, could approach an exact solution. However, the added computational complexity of a means to switch between transformations and accompanying reduced states, as well as the extra memory required, may overshadow the advantages.

An alternative is to bypass any linearizations and utilize the empirical gramians. The procedure for obtaining empirical gramians is essentially a locally linear approximation to a nonlinear phenomena, so there is every reason to believe it will work as well as obtaining exact gramians for a linearized version of the system. In addition,

just as with the approach outlined above, empirical gramians may be obtained for multiple operating points.

The same benefits and drawbacks to finding balanced truncations about single or multiple operating points apply. In this thesis, we shall advocate finding empirical gramians about a single operating point. We recognize the short-comings of this approach, but view it as a worthy starting point for examining reduced order models for nonlinear systems.

Just as there are multiple approaches to finding a balanced transformation for a nonlinear system, there are multiple ways to model the reduced state dynamics once a transformation has been found. If the system is linearized, then exact balancing transformations for the linear system can be obtained. A linear representation for the reduced order dynamics can be found as well. The system can be linearized about many operating points, and multiple linear reduced order models can be found as well. Just as with the transformations, this does allow for arbitrary accuracy in the reduced order models, but at the expense of memory and computational power.

The approach we shall employ is to learn a representation for the reduced order state, linear or nonlinear, once the empirical gramians and their corresponding balancing transformations have been found. This method has manageable memory requirements and does not require a method of switching between representations (it is not time varying).

Consider dynamics of the form (the form of dynamics for our hind limb model, sections 2.3),

$$\dot{x} = f(x) + Bu \tag{4.69}$$

after finding empirical observability and controlability grammains a transformation that allows for the following relations, $z = Vx$ and $\hat{x} = Wz$ can be found. From the system dynamics a data set of the follwing can be found.

$$\begin{matrix}\{x(t_o), x(t_1), ..., x(t_N)\} \\ \{u(t_o), u(t_1), ..., u(t_N)\}\end{matrix} \Longrightarrow \{z(t_o), z(t_1), ..., z(t_N)\} \tag{4.70}$$

We desire a dynamical relation for the evolution of the low-dimensional state, $z(t)$. Note however, there is no true function $f_z(z)$ that will account for the data. The data only has a underlying function in x, namely $f(x)$. Since the data has been compressed, in general, there is no functional relationship between $\dot{z}$ and z. The best that can be hoped for is an accurate estimate.

To learn $f_z(z)$ we could use the relation, $Vf(WVx) = Vf(Wz)$. This is an appropriate function in that it maps z to temporal derivative of z. However, this mapping may have unintended consequences, like going unstable, or entering regions of state-space never normally seen. For instance, Wz can, (and in practice does) allow for negative muscle activations.

An alternative is to learn a best estimate of the dynamics that would best generate the *data*. That is, find $f_z(\cdot)$, such that integrating

$$\dot{z} = f_z(z) + v \tag{4.71}$$

accurately generates the empirically obtained data, $z(t_o), z(t_1), ...$ This is the approach we shall take in this thesis.

Chapter 5

Reduced Order Control

In this chapter we present a proposal for control of a biological system based on a reduced order model. The reduced order model presented here is meant to be an existence proof for a simplifying motor control strategy. We argue that control based on a reduced order representation of the hind limb elucidates the existence and role of low-dimensional motor commands, commonly observed as synergies, while allowing for simplifications in both estimating the influence of motor commands on the system, and generating said motor commands.

Model order reduction (see chapter 4) is a technique for reducing a dynamical system's complexity (often considered equivalent to system order) while preserving the input-output behavior as accurately as possible. Balanced truncation is a method for model order reduction that has desirable guarantees for stability and estimation errors when applied to linear systems. We employ a method of approximate balancing for nonlinear systems by computing approximate controllability and observability gramians and finding a transformation to a reduced order state variable.

The original system dynamics of the frog limb model are of the form,

$$\begin{aligned} \dot{x} &= f(x) + Bu \\ y &= Cx \end{aligned} \tag{5.1}$$

We shall find a transformation between the reduced state and the complete state, $z = Vx$, $\hat{x} = Wz$ $(\hat{y} = CWz)$. Following constructing this transformation, we find the dynamics for the low-dimensional variable. As discussed in chapter 4, the dynamical relations for the reduced variable are not known and must be approximated. We investigate two representations,

$$\dot{z} = A_k z + B_k u \tag{5.2}$$

$$\dot{z} = f_k(z) + B_k u \tag{5.3}$$

The first is linear in the state variable, the latter is not. Based on the known nonlinearities of the original system, a naive conjecture might be that the latter will perform better. However, we will demonstrate that the trade-offs between the two are not clear.

Furthermore, though we do consider nonlinearities in the reduced state, we restrict ourselves to models that are linear in the motor command, u. This assumption allows for straightforward interpretations of motor primitives.

5.1 Motor Primitives

In this section we focus on motor primitives; what they are, how we define them, and how we will use them. First, a few notes on notation, we assume the following,

$x \in \mathbb{R}^n$	state variable for full-dimensional system
$u \in \mathbb{R}^r$	motor command for full-dimensional system
$z \in \mathbb{R}^k$	reduced order state variable
$y, \hat{y} \in \mathbb{R}^m$	output of full and reduced dimension models

The low-dimensional system is a compact approximation of how the dynamics in x influence the output, y. We would like to use the low-dimensional system to both estimate the output of the full-dimensional dynamics and to generate control signals for the full-dimensional system. These two efforts place restrictions on the command signals. For the purposes of estimation, we ought to restrict our motor commands such that their influence on the full-dimensional dynamics can be best interpreted by the low-dimensional system. For the purposes of using the low-dimensional system to derive control signals for the full-dimesnional system, the commands need to be non-negative values; the motor commands correspond to muscle activations and can therefore not be negative.

The reduced order models we investigate are linear in the command signals, as witnessed by the input matrix, B_k. This being the case, much of the following motivation for motor primitives relies on the existence of a linear vector space, and the operations defined on such a space. If instead B_k were nonlinear, then the following analysis would not, in general, hold. We will examine the implications of a nonlinear input matrix in a later chapter.

The low-dimensional input matrtix, B_k, describes how the r-dimensional motor command, u, influences the dynamics in the k-dimensional state, z. For the reduced model, $k < r$, and B_k has rank k. Because B_k maps motor commands to a lower dimension, we lose information contained in arbitrary motor commands, i.e. B_k has a null space.

Knowing that B_k has rank k, we discuss some implications of using a reduced order model for control. First we consider the implications for estimation. When obtaining an estimate of the full-dimensional system, motor commands are supplied to the reduced order model and the resulting approximation, $\hat{y}$ is observed. Yet if a motor command partially or wholly overlapped with the null space of B_k, the corresponding portion of the motor command would not influence the low-dimensional dynamics and a poor estimate of the full-dimesnional output may result. If, however, the motor commands do not coincide with the null space of B_k, all the information contained in the motor command will be transmitted to the low-dimensional dynamics and accurate estimates of the full-dimensional system will be possible.

Now, consider the implications when the reduced order model is used to generate commands. Since $k < r$, the commands computed with a reduced order model will always lie within k-dimensional subspace of $\mathbb{R}^r$. Therefore, we can represent any possible command generated by the reduced order model as a linear combination of k 13-dimensional vectors. Furthermore, for the purposes of controlling the hind limb, these commands will have to be restricted to producing positive values. To do this we can enforce that the low-dimensional system only produce commands that are a linear combination of k non-negative vectors.

Based on the above observations, motor commands, either for use with estimations or control, should be restricted as,

$$u = \sum_{i}^{k} u_i^* c_i \tag{5.4}$$

where the sum, u, never allows negative muscle activations, and u_i is not in the null space of B_k. We shall refer to u_i^* as a motor primitive and denote the set of k primitives as $\{u^*\}$.

There are several qualitites we desire from the primitives. First we reiterate, motor commands ought to be non-negative. We define our motor commands as those contained in the positive orthant of the r-dimensional space, $\mathbb{R}^{r+}$. That is, u is admissable if $u \succeq 0$ (element-wise greater than or equal to zero). We impose this constraint by restricting motor primitives to the positive orthant as well, $u_i^* \in \mathbb{R}^{r+}$.

As explained above, we also want to restrict our motor commands such that they avoid the null space of B_k, denoted $\mathcal{N}(B_k)$. Any motor commands contained within the null space are lost on the low-dimensional dynamics. The low-dimensional model will not be able to predict the outcome of these commands and will produce poor estimates. On the other hand, commands that are orthogonal to the null space uniquely map to the low-dimensional dynamics. The influence of these commands on the system dynamics can be estimated. Also mentioned above, we know there should be k primitives. If we have more than this, we are retaining redundant information, or information that is lost in the null space.

Identifying the null space of B_k is relatively trivial. We find a singular value decomposition of B_k as $U\Sigma V^T$. The last $r - k$ columns of V are an orthonormal basis for the null space, $\mathcal{N}(B_k)$. The first k columns of V, which we shall denote as $V_{1,k}$, are an orthonormal basis for the orthogonal complement of the null space, $\mathcal{N}(B_k)^\perp$. It would seem we should choose our motor primitives to coincide with the spanning set for $\mathcal{N}(B_k)^\perp$, e.g. $\{u^*\} = V_{1,k}$. Indeed, this would be our best choice for primitives (based on the arguments presented), however, using these vectors to define primitives does not necessarily meet our first constraint for maintaining admissible, non-negative commands.

Taking the positive constraint into consideration, it is tempting to define primitives in terms of the set of all possible commands that are both in the orthogonal complement, and in the positive orthant. This union is a cone, where C is a cone if $\forall x \in C, \alpha \in \mathbb{R},\ \alpha x \in C$. The primitives could be defined by a p-dimensional cone, $C_p^+ = \mathbb{R}^{r+} \cap \mathcal{N}(B_k)^\perp$. However, this definition of primitives is too restrictive, as the

dimension, p is often much less than k.

The implicit requirement in choosing our primitives as a spanning set for $\mathcal{N}(B_k)^\perp$ is that our primitives should be capable of controlling the low-dimensional model. If any primitive is orthogonal to $\mathcal{N}(B_k)^\perp$ the low-dimensional model will have an uncontrollable mode. We could require that each have a nonzero projection onto $\mathcal{N}(B_k)^\perp$, yet this constraint is lax. Even if this requirement was met, the low-dimensional model could be weakly controllable if two or more primitives were nearly parallel. Therefore we can strengthen our demands on the primitives by requiring that they are nearly, mutually orthogonal. To this end, we present the following two objectives which will be used in defining a set of primitives.

O-1. $\|V_{1,k}^T u_i^*\| \approx 1,\ i = 1, 2, ..., k$

O-2. $\|u_i^{*T} u_j^*\| \approx \delta_{ij},\ i = 1, 2, ..., k$

Objective 1 is assures us that each primitive projects onto the orthogonal complement, as much as possible. Objective 2 assures us that the primitives are nearly orthogonal, and therefore independent from the other primitives. Meeting these objectives ensures that the primitives are a "good" basis, guarding against a weakly controllable low-dimensional model. In essence, these objectives, combined with the admissible constraint, are an attempt to identify a non-negative basis for $\mathcal{N}(B_k)^\perp$.

We can also add another objective if needed. There may be cases where certain commands are desired, for instance based on the behavioral results they produce. For instance, if we would like to estimate and control withdrawl reflexes, then commands that excite the flexor muscles are desirable. We shall define these commands as behaviorally relevant, $\{u\}_b$. Based on this we present one last possible objective,

O-3. $span\{u\}_b \subseteq span\{u^*\}$

This objective encourages primitives which will account for the desired behaviorally relevant commands.

In summary, the spanning set, $V_{1,k}$ is an ideal candidate for motor primitives, but in general cannot generate admissible commands. We propose that motor primitives should be chosen as to maximize objectives O-1 and O-2 while meeting the positivity constraint. When necessary, O-3 can be included amongst the objectives to search for primitives that are behaviorally relevant.

Motor commands will be composed of linear combinations of primitives, as in equation 5.4. Commands defined this way will ensure the low-dimensional system is controllable (and not weakly so). In addition, when the low-dimensional model is employed to generate control signals the problem is reduced from that of determing the r independent variables of the command signal, u, to determining the k variables, $c_1, c_2, ..., c_k$. Furthermore, while $c_i \geq 0$, these commands will be admissible.

5.2 Example Analysis

In order to illustrate the use of a reduced model and the identification of motor primitives as described in the previous section, we shall examine a "toy" problem. Our toy problem is of a similar dimension to the nonlinear system we will eventually investigate (section 5.3), but linear, to make the computations simpler and the interpretations straight forward. The details of how the balanced realization and this low-dimensional model are found are skipped (for linear systems the process is relatively trivial, see section 4.3). We want to focus on the determination of primitives once a reduced model is in hand and on the subsequent utility of system estimation and control with the reduced model.

5.2.1 Model Order Reduction

The system we examine will consist of a point mass with two degrees of freedom, which we'll describe with a horizontal component, x_1, and an orthogonal vertical component x_2. The mass is lightly damped and sprung to ground (in analogy with our frog model which has very small passive forces). The mass is acted upon by forces, the magnitude of which are a low-pass filtered version of command signals. Each force acts in a fixed direction with a maximum magnitude, described by the vector $\mathbf{f}_i$.

The dynamical relations for the system are as follows. For the two degrees of freedom of the mass, the acceleration is defined by,

$$\begin{aligned} \ddot{x}_1 &= \sum^{r} f_{1i} - b_1\dot{x}_1 - k_1 x_1 \\ \ddot{x}_2 &= \sum^{r} f_{2i} - b_2\dot{x}_2 - k_2 x_2 \end{aligned} \tag{5.5}$$

where f_{1i} is the i^{th} force acting along the x_1 (horizontal) direction, and b_1 and k_1 are the damping and stiffness acting along this direction as well. Each force is computed as,

$$(f_1, f_2)_i = \mathbf{f}_i \alpha_i \tag{5.6}$$

where α_i is the activation state of the i^{th} muscle, whose dynamics are of the form,

$$\dot{\alpha}_i = (u_i - \alpha_i)/\tau_i \tag{5.7}$$

Combining all the dynamical relations into state-space form, we have,

$$\dot{x} = Ax + Bu \tag{5.8}$$

With r inputs, the system has $r + 4$ states ($n = r + 4$), two positions, two velocities, and r forces. For the output, we'll choose the positions and velocities of the mass, $y = [x_1,\ x_2,\ \dot{x}_1,\ \dot{x}_2]^T$.

13 forces will be modeled, analagous to the thirteen muscles in the frog hind limb model. The maximum value of each force is equivalent to the averaged muscle force of the nonlinear SIMM model, at the limb end-point. The force vectors are,

$$\mathbf{f} = \begin{bmatrix} 0 & -.12 & -.10 & 1.06 & .26 & .16 & 0 & 0 & 0 & 0 & 0 & .07 & .14 \\ -1.00 & 0 & 0 & -1.06 & -0.26 & 0 & .09 & .19 & -.75 & -.22 & -.12 & .07 & .14 \end{bmatrix}^{\mathbf{T}} \tag{5.9}$$

With the system defined, we proceed to find a balanced realization. Recall that the balancing procedure finds the states that most significantly impact the energy within the input-output dynamics. To do this, the observability and controllability gramians are computed. These gramians are then used to find an invertible transformation to a new state variable. With this new balanced state representation, the observability and controllability gramians are equivalent and diagonal. Each new state is aligned with an entry on the diagonal of the gramians. These diagonal entries are the system Hankel singular values. By truncating those states associated with small Hankel singular values we have, with a linear system, guaranteed stability and error limits.

In figure 5-1 is a plot of the Hankel singular values for the system. We see that that singular values drop off quickly in magnitude. The first four Hankel singular values represent approximately 99.9% perent of the total sum. Therefore we reduce our system to 4^{th} order ($k = 4$). The low-dimensional dynamics are now of the form,

$$\begin{aligned} \dot{z} &= A_k z + B_k u \\ \hat{y} &= C_k z \end{aligned} \tag{5.10}$$

We make no numerical comparisons of the full and reduced system (the numbers are too cumbersome to present here) but we note that the system eigenvalues of both systems are distinct; the reduced order model is not simply the slow dynamics of the full system. We also observe that based on the system's Hankel singular values, the reduced model has the following error bounds,

$$0.3319 < \|G - G_k\|_\infty < 1.1213 \tag{5.11}$$

The balanced and truncated system is a simplified representation for the input-output dynamics. Now we proceed to define the resulting primitives.

5.2.2 Defining Primitives

For illustrative purposes we present the positive cone, C_p^+ for the case where $k = 4$. We numerically search for the intersection of the positive orthant and the orthogonal complement, $\mathbb{R}^{r+} \cap \mathcal{N}(B_k)^\perp$. We find that the cone is a two dimensional subspace, approximately,

$$\begin{aligned} C_2^+ = \; & span\{[0\ .13\ .06\ 0\ .15\ .19\ 0\ 0\ 0\ 0\ 0\ .38\ .88]', \\ & [0\ .05\ 0\ .46\ .24\ .22\ 0\ 0\ 0\ 0\ 0\ .33\ .75]'\} \end{aligned} \tag{5.12}$$

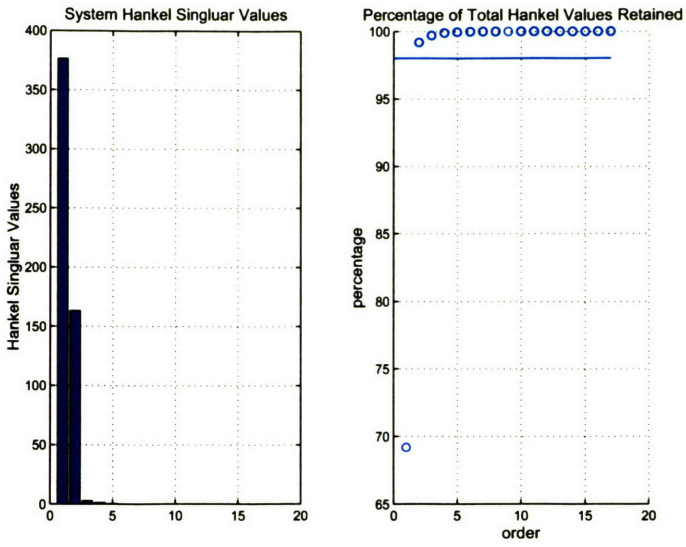

Figure 5-1: In the left plot are the Hankel Singular Values. The right plot displays the ratio of Hankel values retained for reduced models of increasing order. The horizontal line indicates 98%.

Therefore, if we limited ourselves to the positive cone, we would have at most 2 primitives, rendering our low-dimensional system uncontrollable. The low-dimensional model would be capable of generating estimates of the full system, but inadequately limited for generating control.

To further motivate our reasoning for primitives that maximize the stated objective costs, we shall explore two cases: 1) motor primitives chosen naively and 2) motor primitives chosen to maximize costs O-1 and O-2.

Case 1:
We shall first choose primitives based on some intuition of their influence on the system output. For instance, we could define them as below,

$$\begin{aligned}\{u^*\} \;=\; &\{[0\,1\,1\,0\,0\,0\,0\,0\,0\,0\,0\,0\,0]', \\ &[0\,0\,0\,1\,1\,1\,0\,0\,0\,0\,0\,1\,1]', \\ &[0\,0\,0\,0\,0\,0\,1\,1\,0\,0\,0\,0\,0]', \\ &[1\,0\,0\,0\,0\,0\,0\,0\,1\,1\,1\,0\,0]\}\end{aligned} \tag{5.13}$$

By examining the force directions (equation 5.9) we see that with these commands the mass can be driven in all directions (up, down, left and right) while obeying the positivity constraint. These commands, therefore, would seem to be good candidates for motor primitives.

We note that these primitives are mutually orthogonal and none is an element of the null space of B_k. But, we would like to know how strongly each primitive projects onto the orthogonal complement, and in turn, how strongly it influences

the low-dimensional dynamics. To do this, we look at the sum of inner products of each command and the spanning set for the null space of B_k (found from a singular value decomposition). This value must be between one and zero, or 0% and 100% (since the primitives have unitary norm and are orthogonal). If it is 0%, then the primitive is in the null space, and if it is 100%, the primitive is an element of the orthogonal complement. Figure 5-2 is a plot of the percentage projected of each of the above proposed primitives. We observe that the second and fourth primitives have strong projections within the orthogonal complement. The first and third primitives, however, only weakly project to the complement; they are closely aligned with the null space of B_k.

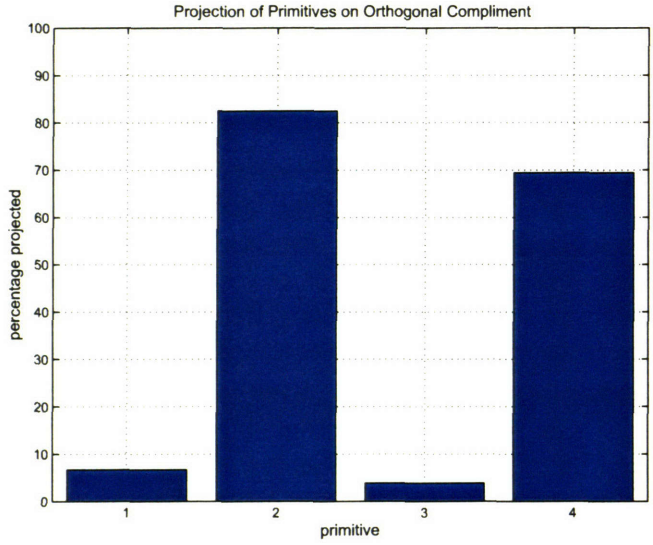

Figure 5-2: Percentage of each naively proposed primitive projected onto orthongonal complement.

We can further quantify how strongly these primitives project onto the low-dimensional space by examining the singular values of $V_{1,k}\{u^*\}$, where $V_{1,k}$ is the matrix formed from the first k columns of V (found from a singular value decomposition of B_k). For example, if $\{u^*\} = V_{1,k}$, the singular values of $V_{1,k}\{u^*\}$ would all be one. If a primitive projected poorly (was nearly in the null space) or was nearly parallel to another primitive, it would result in a small singular value. This would in turn imply that the low-dimensional model was weakly controllable. For the proposed primitives, the singular values of this matrix are 0.9079, 0.8656, 0.2102, 0.0738. The two small values are due to the first and third primitive, which weakly project onto the orthogonal complement.

From figure 5-2 and the singular values of $V_{1,k}\{u^*\}$ we gather that every primitive projects onto the complement; two do so strongly and the other two weakly. The primitives are mutually orthogonal, so they project in all directions of the orthogonal complement. Based on these observations, we can control the low-dimensional sys-

tem, though perhaps weakly so.

Case 2:
Now we use the proposed objectives to define primitives. We will find primitives by formulating the search problem as that of optimizing a cost function that maximizes objectives O-1 and O-2, subject to the positivity constraint.

$$\begin{aligned} L &= \frac{1}{2}\sum_{i}^{k}\left(\|V_{1,k}^{T}u_i^*\|^2 + \|\{u^*\}^T u_i^*\|^2\right) \\ &= \frac{1}{2}\sum_{i}^{k}\left(u_i^{*T}V_{1,k}V_{1,k}^{T}u_i^{*T} + u_i^{*T}\{u^*\}\{u^*\}^T u_i^*\right) \end{aligned} \tag{5.14}$$

We can find an analytical gradient for this expression and try to maximize the cost. Starting from random initial $\{u^*\}$ the solutions always converge on the same primitives. The result is shown below,

$$\begin{aligned} \{u^*\} = \{&[0, 0.12, 0.05, 0, 0.13, 0.19, 0.0, 0.0, 0, 0, 0, 0.39, 0.88]' && (5.15)\\ &[0.01, 0, 0, 0.97, 0.21, 0.10, 0, 0, 0, 0, 0, 0, 0]', && (5.16)\\ &[0.99, 0.08, 0.06, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0.0]', && (5.17)\\ &[0, 0.10, 0.06, 0, 0.09, 0, 0, 0, 0.92, 0.31, 0.19, 0, 0]\} && (5.18) \end{aligned}$$

Figure 5-3 is a plot of the projection of each of these primitives onto the orthogonal complement of B_k. The bottom half of the figure displays the inner products between the primitives. As can be seen, these primitives have large projections in the complement, and are nearly mutually orthogonal. The singular values of $V_{1,k}'\{u^*\}$ using these new primitives are $1.0251, 0.9901, 0.9723, 0.9689$. In essence, what we have found with these new primitives is a best approximation to a non-negative spanning set for the orthogonal complement, $\mathcal{N}(B_k)^{\perp}$. The low-dimensional model will be controlable, and minimal information from the motor commands will be lost in the null space of B_k.

Case 1 vs 2:
Now we shall address how the objective maximizing primitives of case 2 are in fact better than the naively chosen primitives of case 1, by using the primitives to command the system. Both sets of primitives were used to command step inputs (for one second) to both the full and low-dimensional systems. Figure 5-4 is a plot of the output (positions and velocities) and the low-dimensional estimate, for each of the case 1 primitives. By design, these primitives allow for the mass to be driven in any direction. To see this clearly figure 5-5 is a plot of the workspace (x-y coordinates) and the location of the mass after 15 seconds. We see that although the commands that drive the mass in the leftward and upward directions are weak (primitives 1 and 3, which have weak projections in the complement) it is still possibile to generate movement in these directions. We also note that the low-dimensional model has generated qualitatively good estimates.

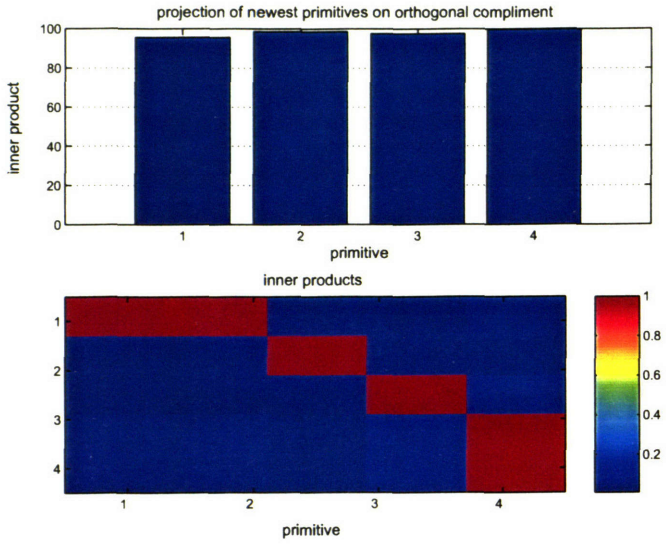

Figure 5-3: percentage of each new and impoved primitive projected onto the orthogonal complement (top), inner product betweem primitives (bottom).

Moving on to the case 2 primitives, we get figure 5-6 and figure 5-7. We call the reader's attention to three aspects of these figures. First, as in the previous results, the estimates generated with the low-dimensional model are relatively accurate (we can't compare the two results directly because they are movements with different trajectories). The second thing to note is the distance this set of primitives has moved the mass. On average, the case 2, cost maximizing primitives have moved the mass much further than the naive case 1 primitives. And lastly, we note that for all practical purposes, restricted to control with these primitives we have lost the ability to move the mass in the leftward direction.

Conisdering these observations, we note that the two sets of primitives cannot be distinguished on the basis of the ability to generate good estimates of our desired output. The Hankel singular values for this system dropped off quickly, so we were assured good estimates given the order of our reduced system, despite the fact that some of the primitives examined were closely aligned with the null space of B_k.

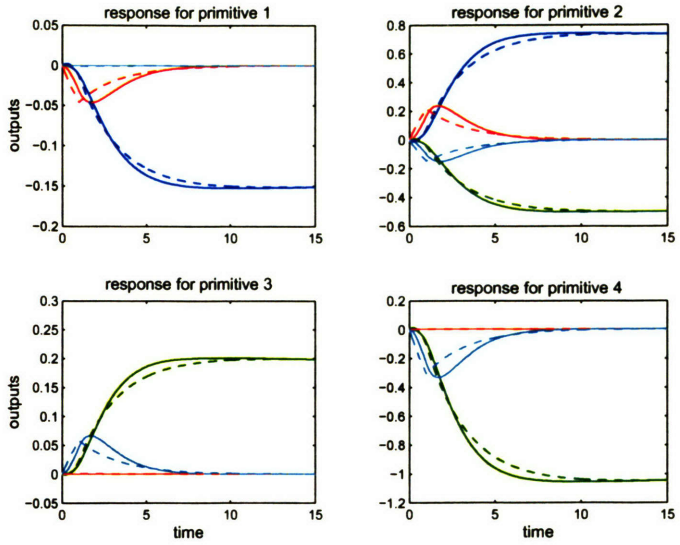

Figure 5-4: Estimated and actual output of system (positions and velocities) for the four naive primitives. Solid lines are the full system, dashed lines are the low-dimensional estimate.

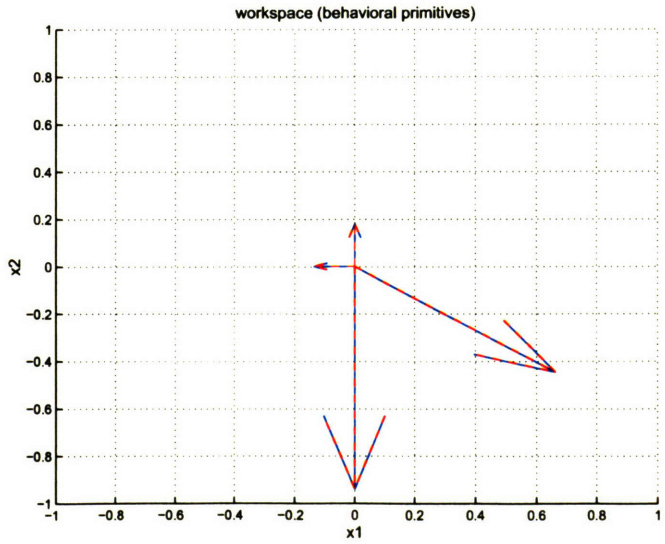

Figure 5-5: Actual (solid line) and estimated (dashed line)location of the mass in the workspace after 15 seconds, for the four naive primitives.

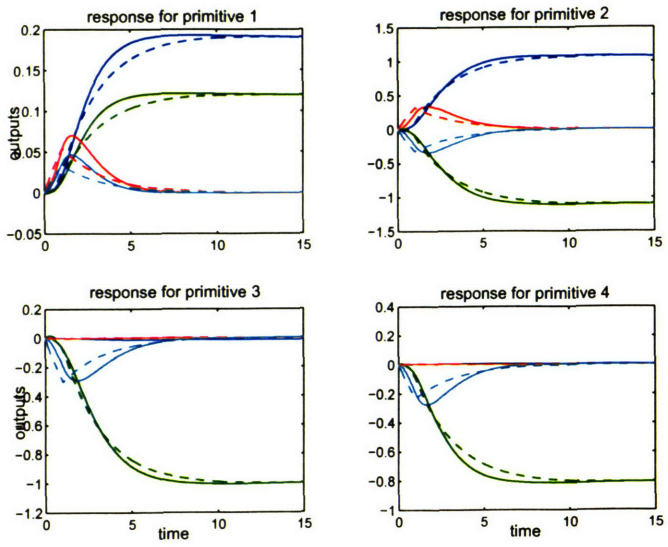

Figure 5-6: Estimated and actual output of system (positions and velocities) for the four newest (cost minimizing) primitives. Solid linear are the full system, dashed lines are the low-dimensional estimate.

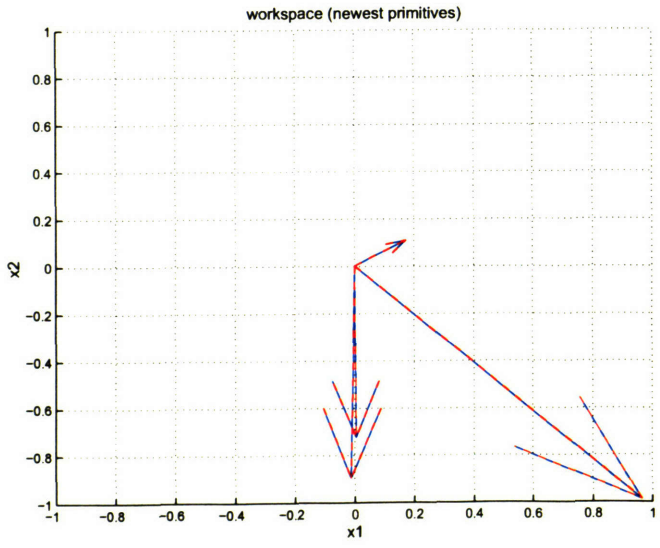

Figure 5-7: Actual (solid line) and estimated (dashed line) location of the mass in the workspace after 15 seconds, for the new (cost minimizing) primitives.

Recall that the balancing procedure aligns states with Hankel singular values, and the truncation has retained only those states with large Hankel values. Therefore commands that are in the null space, or very nearly so, do not influence these states, and are not capable of generating as much relative energy in the output (of the full or reduced system). The output energy is computed as,

$$E_y = \int_o^\infty \|y(t)\|^2 dt \tag{5.19}$$

assuming the command is zero from $t = 0$ to infinity. The toy system, modeled after the frog model, takes a substantial amount of time to return to equilibrium (the frog model has very small passive forces to drive the limb back to equilibrium as well), and the velocities are so small, the output energy is essentially a function of the displacement the command induces. Therefore, we can estimate the relative amount of output energy each primitive induces by considering the displacment of the system after it comes to quasi-steady-state (as the velocities change sign).

Comparing the primitives by their ability to control the system, we note that a significant difference between the primitives of case 1 and 2 is the energy they impart into the system output. We quantify this difference by comparing the energy in the outputs from each of the primitive responses. The norm of the output at quasi steady-state for the case 1 primitives is 0.1513, 0.8888, 0.1983, 1.0466. The norm of the output at quasi steady-state for the case 2 primitives is 0.2235, 0.9977, 1.5331, 0.8047, considerably greater.

The above comparison illustrates that by choosing primitives to maximize objectives O-1 and O-2, motor commands are maximally effective at influencing the system output. Though perhaps not obvious, this implies that the commands derived from the low-dimensional model will require minimal control effort for a given output.

A final note on objective 3:

We noted how with the case 2 objective cost maximizing primitives, limited to non-negative commands, we have lost the ability to move in the left-ward direction. This loss is due to the small output energy associated with the commands that move the mass in that direction, as seen directly in the numbers presented above, and indirectly by these commands' poor projection in the orthogonal complement. However, if this result was unacceptable and movement in this or another direction was desired, we can accomodate this within the framework of the balanced and reduced modeling.

We can include the third objective, O-3 in the cost function. Adding this objective to the cost would constitute a compromise between the desire to have certain specified commands as primitives, and primitives that are a mututally orthogonal spanning set of the complement.

A second approach to accomodate objective O-3 is to append the system output to reflect these desired behaviorally relevant commands. For instance, by adding an extra output to our full system that is the sum of the activations for the second, third, seventh and eighth forces (all the forces that drive the mass left-wards and upwards) we can find a new (4^{th} order) balanced and reduced system. Maximizing objectives

O-1 and O-2, we find a new set of primitives. The response for each of these new primitives is in figure 5-8. Now we have the ability to move in any direction in the workspace while restricting our commands to positive values.

Although both of these approaches can be used in case a particular direction of movement is not represented in a particular synergy set, we note that in our analysis of the realistic frog hind limb model (section 5.3) this is not a problem, and every direction of movement is represented.

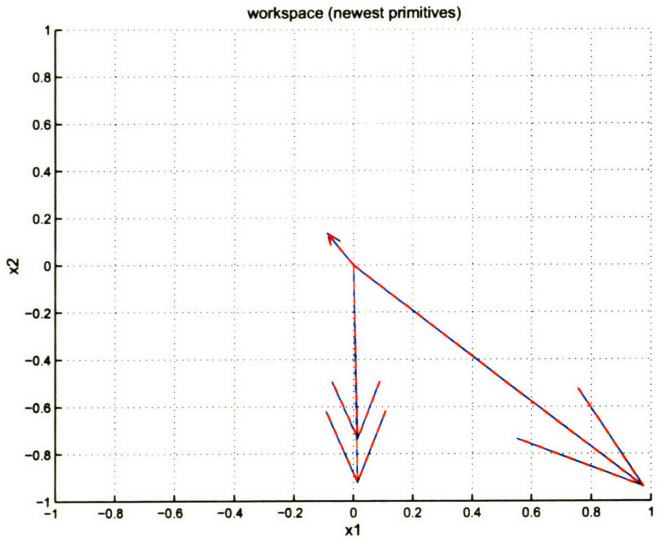

Figure 5-8: Actual (solid line) and estimated (dashed line) location of the mass in the workspace after 15 seconds, for the new behaviorally relevant primitives.

5.2.3 Low-dimensional controller

Once we have the low-dimensional representation and primitives, the next step is to investigate how the reduced model can be used for designing control. We will examine optimal solutions for the terminal state problem. The cost to be minimized is,

$$L = \frac{1}{2}\int \xi^T Q \xi + u^T R u \, dt \tag{5.20}$$

where the state, ξ, starts from a specified initial state, zero in these examples, and ends at a specified terminal state at a fixed terminal time. To find the solutions we solve the Euler-Lagrange equations using a neighboring extremal method. Because the system in quesiton here is linear, the solutions are relatively easy to compute when the control is not restricted to positive values. When the positive restriction is enforced the solutions are harder to obtain. The solutions shown below do not enforce this positivity constraint. In all examples Q is zero, R is the identity matrix and the terminal state is one unit to the right.

The first comparison we make is between system responses using optimal commands, the commands that minimze the cost, equation 5.20, without restricting the command to a combination of primitives. The first system response is found by minimizing 5.20 subject to the full system dynamics ($\xi = x$). The second system response is found by using the cost minimizing command, subject to the reduced order dynamics ($\xi = z$). The two output trajectories are shown in the top plot of figure 5-9. In red is the optimal solution found from the full system dynamics allowing all 13 control variables to vary. The blue trajectory results from using the low-dimensional model to compute an optimal control signal allowing all 13 control variables to vary. We see that there is good qualitative agreement between the two output trajectories, though the response found with the low-dimensional command does not achieve the desired terminal state exactly.

Next we examine the system responses when we restrict the commands with the primitives found by maximizing the objectives, as described earlier. These responses are shown in the bottom plot of figure 5-9. In red is the optimal solution for the full-dimensional system when the command is a linear combination of the primitives, as in equation 5.4. This solution is quite distinct from the previous results. Restricting the command for the full-dimensional optimization has resulted in a larger cost movement with a different output trajectory. In blue is the system's response using the low-dimensional system's optimal control solution. This response is very similar to the previous result.

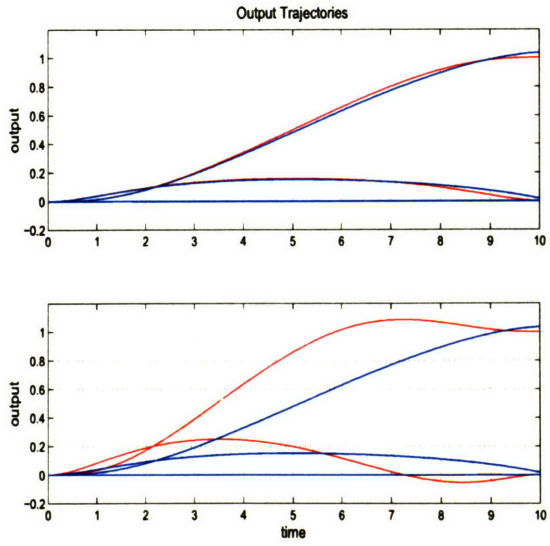

Figure 5-9: Red trajectory is full-dimensional system's optimal soluion. In blue is the full dimensional response when driven with the low-dimensional optimal solution. Top plot: solutions obtained without restricting motor commands. Bottom plot: solutions obtained when motor commands are restrained to linear combinations of motor primtives.

From these results we can infer the following: Based on the first comparison, without restricting the command to primitives, we observed that the reduced order dynamics are similar enough to the full system to compute very similar command solutions. Furthermore, though we don't explicitly restrict the command, the low dimensional system is restricted by the rank of B_k and only independently varies $k = 4$, dimensions of $\mathcal{N}(B_k)^{\perp}$. This would suggest that the low-dimensional system, by approximating the full-dimensional dynamics well, has implictly found the appropriate subspace of u to best influence the dynamics, and therefore minimize the cost funciton.

We observe from the second comparison that commands restricted to primitives, do impact the optimal commands for the full system dynamics, but not the reduced order dynamics. The full-dimensional optimal solution with primitive-restricted commands, has lost control signals that would otherwise allow for a lower cost solution (as witnessed by this response's higher cost). The commands derived with the reduced order dynamics however, have not lost any perceptible information, and the primitive-restricted commands are nearly identical to the unrestricted commands (not shown). This is due to the choice of primitives that are a best non-negative approximation to the spanning set for $\mathcal{N}(B_k)^{\perp}$.

These results suggest that a reduced order model can generate controllers in a reduced space, that are comparable to solutions found using the full system dynamics when not restricted to primitives. Though not fully explored in this document, the implication is that a low-dimensional controller does not significantly limit the motor commands obtainable by the full-dimensional system, when similar optimal solutions are sought; optimal solutions found using the low-dimensional controller are comparable to optimal solutions found using the full-dimensional dynamics.

5.3 Frog Model Results

We now examine a low-dimensional representation for the nonlinear model of the frog's hind limb (described in chapter 2). First, we describe the process of empirical balancing. This balancing yields a change of state coordinates that is used to find a reduced order state, z. Then two representations for the low-dimensional dynamics in z, one linear, one nonlinear, are examined for their respective abilities to estimate the full system responses. We then present the motor primitives found by maximizing the objective costs O-1 and O-2. The low-dimensional system's ability to generate control for the full dimensional nonlinear system is then examined.

5.3.1 Empirical Balancing

As explained in chapter 4, the empirical gramians are found by integrating a series of system responses from inputs and initial conditions. In the linear case the result is equivalent to computing the system gramians. The method for nonlinear systems is simply a linear approximation to the gramians about some operating point. For this analysis the empirical observability and controllability gramians are computed about the zero orientation (zero degrees for the hip and knee angles) for the frog hind limb model. This orientation is displayed in figure 2-4.

The empirical controllability gramian is found as,

$$\hat{\mathcal{P}} = \sum_{l}^{q} \sum_{h}^{s} \sum_{i}^{r} \frac{1}{qsc_h^2} \int_0^\infty \Phi^{ilh}(t)dt \tag{5.21}$$

where

$$\Phi^{ilh} = (x^{ilh}(t) - \bar{x}^{ilh}) \ (x^{ilh}(t) - \bar{x}^{ilh})^T \tag{5.22}$$

where q, s and c_h are scalar constants described in section 4.6 and x^{ilh} is the state corresponding to the impulse command $u(t) = c_h \mathcal{T}_l e_i \delta(t)$. However, we let $c_h = 1$ and $\mathcal{T}_l = I$, therefore $u(t) = e_i \delta(t)$ is simply the impulse response for the i^{th} muscle.

The empirical observability gramian is found as,

$$\hat{\mathcal{Q}} = \sum_{l}^{q} \sum_{h}^{s} \frac{1}{qsc_h^2} \int_0^\infty \mathcal{T}_l \Psi^{lh}(t) \mathcal{T}_l^T dt \tag{5.23}$$

where

$$\Psi_{ij}^{lh}(t) = (y^{ilh}(t) - \bar{y}^{ilh})^T \ (y^{jlh}(t) - \bar{y}^{jlh}) \tag{5.24}$$

and y^{ilh} is the output corresponding to the inital condition response with $x_0 = c_h \mathcal{T}_l e_i$. Again, we simply chose $c_h = \pm 1$ and $\mathcal{T}_l = I$. With the hind limb model, the zero orientation is approximately at rest when the command is zero, therefore this gramian is found about a stable equilibrium.

In theory the integrals above need to be integrated indefinitely, yet numerically the integration can end once the integrand is zero or sufficiently close to it. For the

hind limb system, this is achieved when the state has returned to its stable, passive equilibrium, the zero orientation. The passive stability of this equilibrium is weak, and the limb takes a considerable amount of time to achieve this steady-state; 200 seconds of simulated time were needed.

Once the gramians were computed, a balancing transformation was found. This was done as outlined in chapter 3 (We remark that though the computations outlined here are straightforward, the time required to simulate these system responses and obtain the subsequent results was substantial).

5.3.2 Reduced Order Variable

With the balanced and truncated transformations we have a map to and from the low-dimensional variable, $z = Vx$, $\hat{x} = Wz$ $(\hat{y} = CWz)$. We also have Hankel singular values obtained from the empirical gramians. Figure 5-10 is a plot of these values, along with a plot of the sum retained as the order increases. The Hankel values drop off quickly in magnitude. For a linear system, this plot would indicate an upper bound on the error limits for the reduced order model. And, as done with the previous sections toy example, a small order would suffice to retain good estimation properties. From the left plot of figure 5-10 observe that a third order reduced model would retain approximately 98% of the Hankel energy. However, the system of interest is not linear, and these gramians and Hankel values are only approximations about an operating point. Therefore more care must be given to choosing an order for the reduced model.

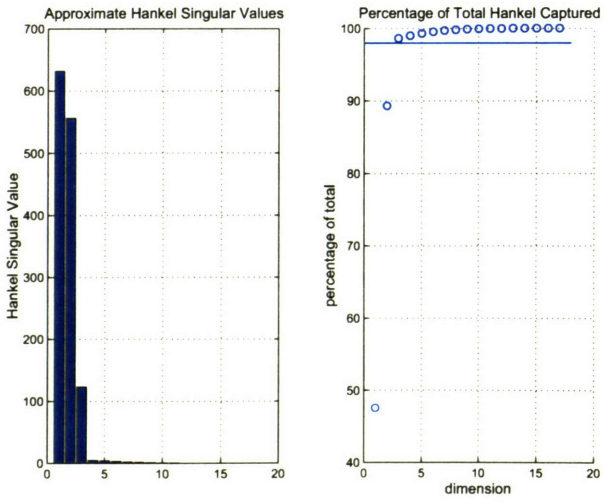

Figure 5-10: Hankel values and percent retained vs order based on empirical balancing. The horizontal line indicates 98% (left plot).

Since the Hankel singular values don't necessarily correspond to the correct error limits, we can consider the influence of system order on the error inherent in the

transformations. Consider again the transformations to the low-dimensional variable, z and back to the full-dimensional state.

$$z = Vx \tag{5.25}$$
$$\hat{x} = Wz \tag{5.26}$$

For a linear system, the above tranformations would be used to project the full system dynamics to a reduced space (see chapter 4). However, because the full-dimensional system is nonlinear, we will estimate the low-dimensional dynamics with a choice of function approximators. In order to make this estimation we will use the transformation V to compute training data in z. Consider then, that under these circumstances, in the absence of any error in the low-dimensional dynamics, there is still an error associated with the transformations, that is $\hat{x} = WVx \neq x$. Therefore, we examine how this inherent error in the mapping between low and full dimensional variables is influenced by system order. This error will better answer how well a given order captures the properties of the full system state.

In order to examine the error of this transformation (and later to learn the low-dimensional dynamics) some data is required. Step responses are individually commanded to each of the thirteen muscles. The amplitude of each muscle's command is adjusted by hand so that the resulting motions have approximately the same magnitude of angular displacements. Each command is held for 0.1 seconds. A workspace plot of the resulting data is shown in figure 5-11.

This data set is used to compute the error between the system output, $y = Cx$ and the estimate based on the balancing transformation's map to the low-dimensional variable, $\hat{y} = CWVx$. The R^2 value for $\hat{y}$ as the order of the transformation increases is displayed in figure 5-12. Note that despite the fact that the computed Hankel singular values drop off quickly after the third order, the R^2 value for third order is poor (less than 0.8). Based on this inherent error in the mapping, and the trade-off between desiring a low-order representation and good estimates of the system, the R^2 values suggests we choose an order of at least 4, since we need to ascend to a 9^{th} order model to get a significant further increase.

Another consideration when choosing the system order are primitives. The order of the reduced system will determine how many primitives there are. This being the case, too few primitives, even if the Hankel singular values and inherent error in the mapping between dimensions indicated it was viable, would be undesirable; this might limit the motor behaviors achievable when restricted to primitives, as well as limit the reduced system's ability to estimate the consequences of motor actions (by increasing the space of motor commands that reside in the null space of B_k).

Based on these considerations, and because the reduced order model is simply meant to represent an existence proof, we shall choose $k = 5$. This appears to be large enough to allow for good estimates of the state (based on the correlation coefficient) and the possibility for a variety of motor primitives, while still small enough to create a manageable reduced order model. Also worth noting, most studies of synergies have found between 5 and 7 are needed to account for the experimentally derived motor behaviors. Therefore we argue that a 5^{th} order system, with its 5 motor primitives is

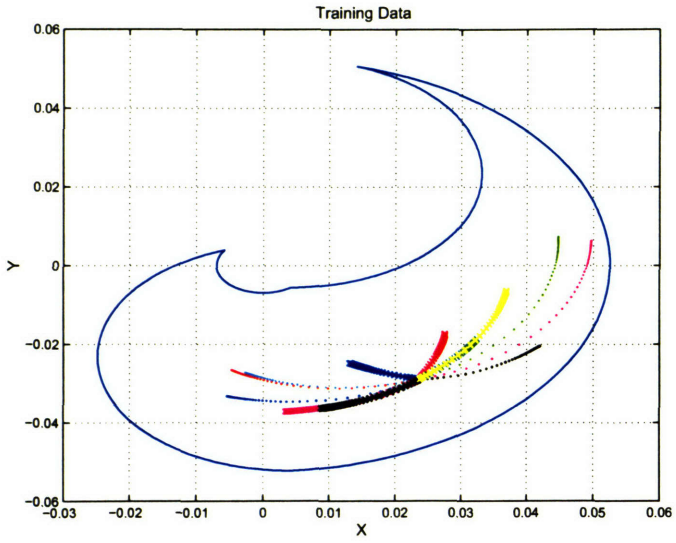

Figure 5-11: Step-commanded movements used to get training data.

a simple attempt to determine if model order reduction is feasible for motor control with the frog's hind limb. No further attempts to examine the trade-offs between models of various dimensions is examined here.

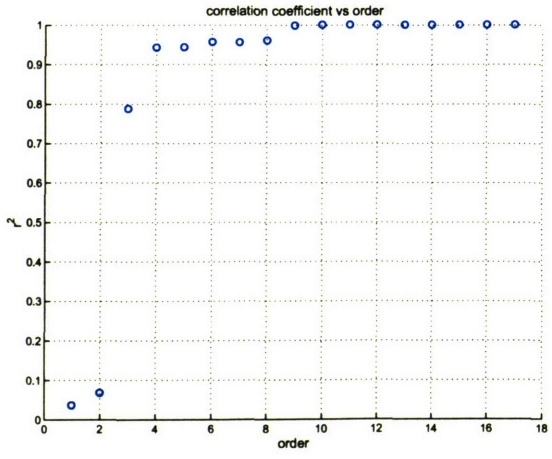

Figure 5-12: R^2 value for $\hat{y}$ vs. order/rank of transformation to the low-dimensional varibale, z.

5.3.3 Linear Reduced Order Model and System Estimation

In the last section we found a low order representation for the state, x. Now we shall model the dynamics in this reduced order state, z. We shall first approximate the low-dimensional dynamics with a linear representation. That is, we'll approximate the dynamics in z as,

$$\dot{z} = A_k z + B_k u \tag{5.27}$$

To fit this model, we shall use the data obtained above, from step inputs to each of the thirteen muscles.

With this data, we map to the low-dim variable and calculate an approximate derivative, $\dot{z}$, using an euler back-difference approximation. In the absence of any round-off noise, the smaller the time step, the better; about ten times faster than the fastest dynamics we hope to capture is more than adequate. For this model, that is very fast. This model has extremely light links and large passive forces. Some natural dynamics are on the order of a hundred hertz. The time step used, Δt was 10^{-4} seconds.

The system matrices are computed as the least-squares solution of

$$\begin{bmatrix} \dot{z}_1 & \dot{z}_2 & \dots & \dot{z}_N \end{bmatrix} = [A_k, B_k] \begin{bmatrix} z_1 & z_2 & \dots & z_N \\ u_1 & u_2 & \dots & u_N \end{bmatrix} \tag{5.28}$$

This is done for $k = 5$. With the matrices solved for, we then validate them over these same commands. The results are shown in figure 5-13. The dotted lines are the low-dimensional estimate, the solid ones are the full-dimensional output. We note that the low-dimenional model found minimizes the error in the forward dynamics (from state to state derivative), so the data shown is a validation and not simply the models' performance over the training set.

In order to evaluate this model's performance, we shall compute an error measure for the low-dimensional state. An RMS error between the low-dimensional variable, z, statically computed from the training data ($z = Vx$) and the low-dimensional variable found from the low-dimensional dynamics is computed for each movement. We shall refer to this as the low-dimensional state error. Figure 5-14 is a plot of this low-dimensional error for each movement. For now we simply note that the linear model has computed qualitatively good estimates of the full system responses.

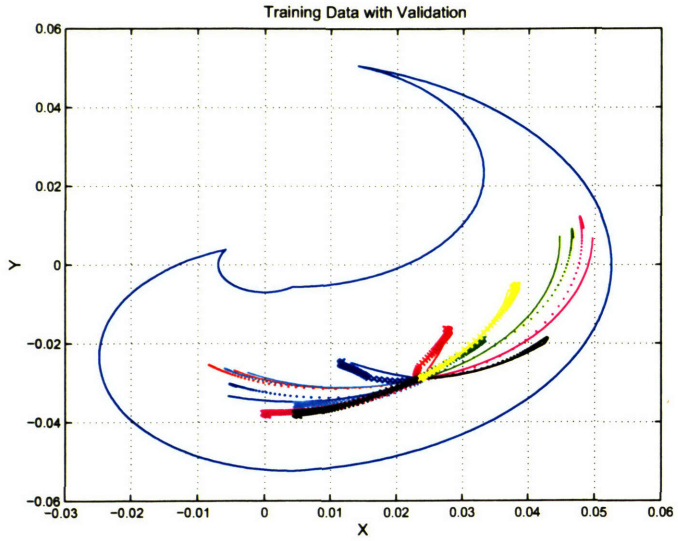

Figure 5-13: Workspace plot of training data and validation. the solid lines are the system response, and the dotted lines are the low-dimensional estimate, given the same command.

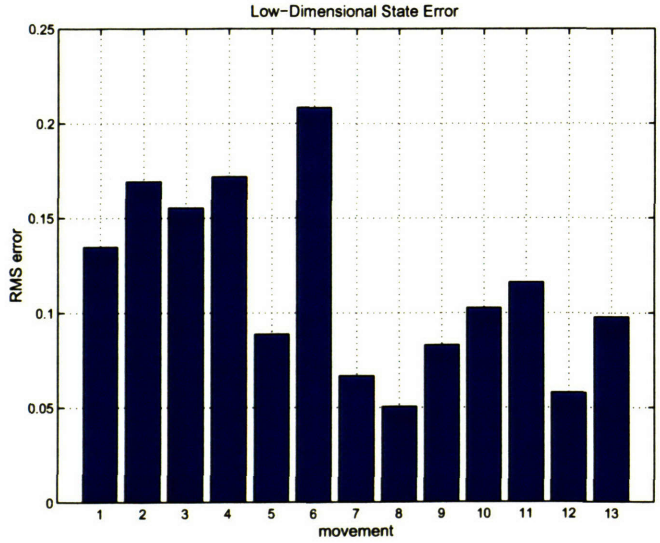

Figure 5-14: The low-dimensional output RMS error for each movement (linear system).

5.3.4 Nonlinear Reduced Order Model and System Estimation

We now compare the preceding results with a nonlinear model of the low-dimensional dynamics in z. We want to know how much better the low-dimensional model can estimate the system responses when the dynamics are modeled as,

$$\dot{z} = f_k(z) + B_k u \tag{5.29}$$

where f_k is computed using a neural network but we assume the dynamics are still linear in the command signal. We use the same data set as above. Because the dynamics are linear in the command, we use the least squares solution found earlier for B_k.

The neural network used is a multi-layer feedforward network with a single hidden layer. The hidden layer has 20 neurons, and the ouput layer has 4. Each neuron has a sigmoid activation function, $\sigma()$, bound between -1 and 1. The defining equations are as follows,

$$\hat{\eta}_1 = W_1 z + b_1 \tag{5.30}$$
$$\eta_1 = \sigma(\hat{\eta}_1) \tag{5.31}$$
$$\hat{\eta}_2 = W_2 \eta_1 + b_2 \tag{5.32}$$
$$f_k(z) = \sigma(\hat{\eta}_2) \tag{5.33}$$

where W_1 and b_1 are the weights associated with the hidden layer, and W_2 and b_2 are the weights for the output layer.

To train the network a standard gradient descent approach is used. An error metric is defined,

$$E = \sum_i \|\dot{z}_i - f_k(z_i)\| \tag{5.34}$$

the gradient with respect to the network weights is computed and updated in batch.

After training the network over this data, we validate the model just as described for the linear model (in the previous section). Figure 5-15 is a workspace plot of resulting estimates of the system output for the 13 step responses. The low-dimensional state error is computed and shown in figure 5-16. Again, there is good qualitative agreement between the estimates and the actual system output.

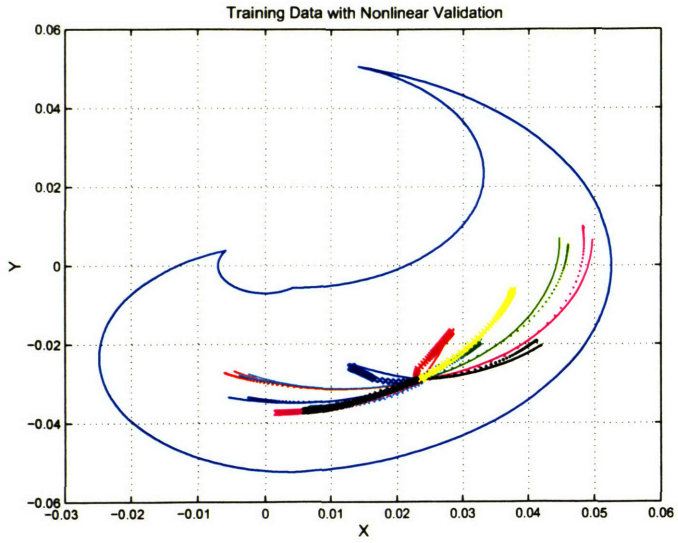

Figure 5-15: Workspace plot of training data and nonlinear validation. the solid lines are the system response, and the dotted lines are the low-dimensional estimate, given the same command.

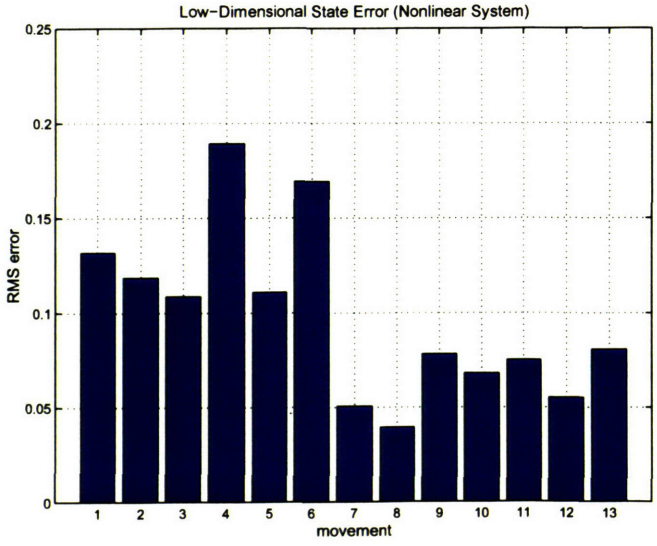

Figure 5-16: The low-dimensional output RMS error for each movement (nonlinear system).

5.3.5 Reduced Model Comparison

Qualitatively, both the linear and nonlinear reduced order models produce good estimates of the full system response. Comparing figures 5-14 and 5-16 we see that as expected, the nonlinear model does produce a superior low-dimensional estimate, albeit only slightly so. However, as discussed eariler even if the estimate of the low-dimesnional dynamics were exact, there is still an error associated with the transformation between spaces.

In figure 5-17 we plot the error inherent in the mapping between dimensions for $k = 5$. This is computed as the root mean squared error of $y - CWVx$ for each step response. This is the estimation error the low-dimensional system would produce if its dynamics exactly matched $z(t)$ from the training set, regardless of the representation for the low-dimensional dynamics.

From this figure we observe that the estimated system output error due to the mapping between dimensions is significantly greater than the error either reduced order model produces. What's more, the reduced order model will eventually be used for control, therefore a linear representation of the low-dimensional dynamics will offer a simplification over a nonlinear model. Given this information, it is not clear that the nonlinear reduced order model offers an advantage over the linear model. Therefore we assume that the linear representation for the reduced order dynamics will suffice in our efforts to estimate system dynamics and generate control. We believe the bottleneck for system estimation (when $k = 5$) is the error inherent in transformation from z to $\hat{x}$, rather than the in the estimated dynamics of $\dot{z} = f_k(z, u)$.

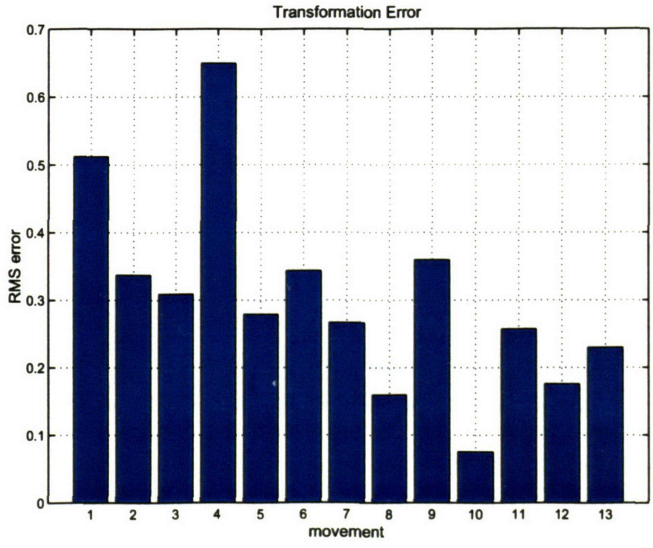

Figure 5-17: Inherent RMS error, due to transformation, for each movement.

5.3.6 Primitives

With the system matrix, B_k identified (equivalently in both the linear and nonlinear case), we proceed to find primitives. Taking our stated approach, primitives are found that maximize the objective cost function, equation 5.14, maximizing their projection onto the orthogonal complement, $\mathcal{N}(B_k)^{\perp}$, and maximizing their orthogonality, while remaining non-negative. Figure 5-18 is a plot of the percentage of each primitive projected, and the inner products for these cost maximizing primitives. Each primitive has a large projection onto the orthogonal complement, at least 90%, and they are close to mutually orthogonal. In figure 5-19 the primitives are displayed.

Before making comparisons with experimentally derived synergies and the motor primitives obtained from the reduced order model, we note their biomechanical significance. The first primitive is essentially a command to Gracilus and Semimembranosus. Both muscles are strong hip flexors. The second primitive is mostly a command to Cruralis, a hip flexor and knee extensor. The third primitive is mostly Gluteus and Tensor Faciae Latae, both hip and knee flexors. The fourth primitive is mostly commands to the dorsal and ventral head of Semitendinosus, hip and knee extensors, and the dorsal and ventral heads of Adductor, also hip extensors. The fifth primitive is combination of Tensor Facia Latae, Iliofibularis, Sartorius, and Iliacus internus and externus, all hip flexors. Thus, these primitives appear to identify common components acting upon the joints.

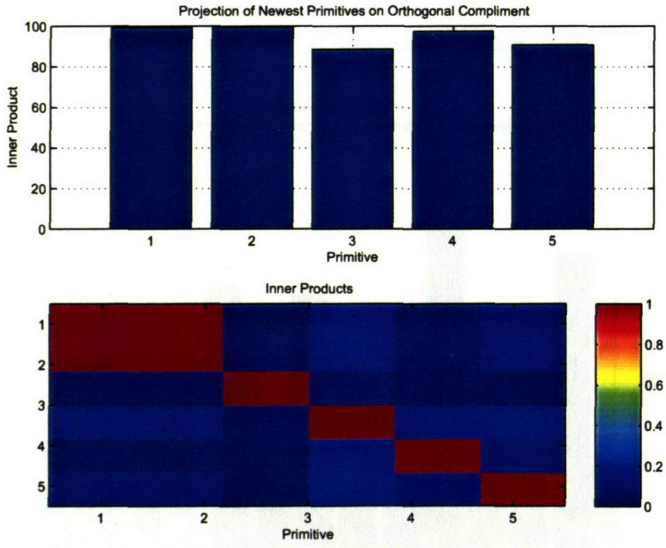

Figure 5-18: Percentage projected onto the orthogonal complement for each motor primitive. Inner products for the primitives are shown in the bottom plot.

Muscle synergies are found by extracting low-dimensional basis from experimentally determined EMG signals. Numerically, this has been done by finding linear combinations of basis, or synergies, that best reconstruct the data. The basis can be found through varying means, for instance using a gradient descent technique

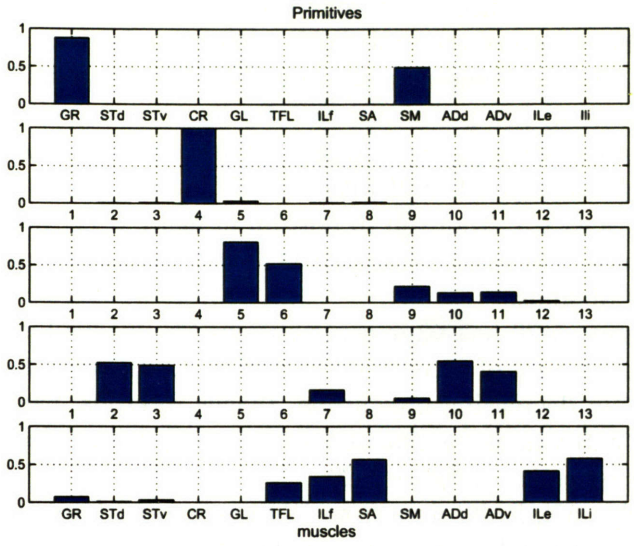

Figure 5-19: Objective maximizing primitives.

([67], [57]) or non-negative matrix factorization ([11]) or through principal component analysis. The number of synergies required to explain a given data set is a flexible quantity, typically given in terms of variance explained.

The resulting synergies are ratios of relative muscle activity. Due to experimental limitations, the magnitude of EMG activity for any muscle is normalized relative to the measurements taken, and not the maximal activaiton for that muscle. Therefore, the scalings of muscles within a synergy are relative muscle activations and not values of absolute muscle activity. This complicates a direct comparison of synergies and our motor primitives.

In addition to this incongruity, there are good reasons to suspect the motor primitives and synergies may differ: the hind limb model. The hind limb model used in this study has many simplifications and modeling assumptions that make it distinct from the real hind limb. For example, muscle properties, such as maximum isometric tension and insertion points may be inaccurate. Furthermore, the hind limb model, and the forces the musculo-tendon models produce, are in a plane, whereas the real frog limb and it's muscle forces are not constrained to a plane. These assumptions, could be sources of disparity between the model's motor primitives and the muscle synergies found experimentally.

Given the above information, the comparisons we make between our motor primitives and muscle synergies will be based only on the muscle prominently contained therein. Without reviewing the details of the experimental methods, the synergies found from three studies will be used for comparison. Each study obtained the EMG measurements from differnet motor behaviors, and extracted the synergies using different numerical methods, so the comparisons between synergies and motor primitives will be made for each study seperately.

Amongst four synergies used to reconstruct the EMG of ten muscles collected from cutaneous stimulation of the limb ([67]), two synergies bare resemblance to the motor primitives. The two synergies are reproduced in figure 5-20. Synergy 3 (the second from the bottom) has prominent activity for Gluteus, with smaller activities for Gracilus, Adductor Magnus and Cruralis (Tensor Facia Latae was not recorded from). This synergy corresponds well with our third motor primitive (we will refer to the motor primitives as 1-5, from top to bottom in figure 5-19), which commands, in decreasing magnitude, Gluteus, Tensor Facia Latae, Semimebransosus and the two heads of Adductor Magnus.

Synergy 4 in figure 5-20 has large contributions for Illiacus Internus and Iliofibularis, with a smaller amount of activity for Sartorius. Compare this with our primitive 5, which activates, in decreasing magnitude, Iliacus Internus, Sartorius, and Ilifibularis, with smaller activations for Iliacus Externus and Tensor Facia Latae, both not represented in the synergies of figure 5-20.

In another study, [57], EMG data from 12 muscles, and from 10 frogs was pooled. The data, collected through chemical stimulation of spinalized animals, was used to generate seven synergies. Three sets of seven synergies were found. Amongst these, two sets of synergies contain a synergy that bares close resemblance to our motor second primitive. The synergies, reproduced in figure 5-21 correspond to large activations of Peroneus (a muscle that is not included in our model) and Cruralis, with smaller activations for Gluteus, Sartorius, and Gracilus. Our second motor primitive is a command to Cruralis, with a very small activation to Gluteus.

In the same study, EMG activity evoked from cutaneous stimulation of one frog was also collected and used to find synergies. Of this set of seven synergies, two are similar to our motor primitives and are reproduced in figure 5-22. One prominently represents activations of Gluteus and Tensor Facia Latae, similar to our third primitive. The second synergy, similar to the chemically stimulated synergies, represents activity of Cruralis Gluteus and Sartorius, similar to our motor second primitive.

Also from the same study, in a set of seven synergies constructed from the EMG recordings of a single frog, three synergies (shown in figure 5-23) are similar to our motor primitives. The synergy labeld D, with activations for Gluteus, Tensor Facia Latae and Semimembranosus, is similar to motor primitive 3. The synergy labeled E, with activations for Semitendinosus, Adductor Magnus, Iliofibularis, and Sartorius, is similar to our fourth motor primitive. Finally, synergy labeled F, which represent activities for Sartorius, Tensor Facia Latae and Iliacus Internus, is similar to our primitive 5.

In the last study we use for comparison, [11], EMG was collected from 13 muscles while awake and behaving frogs swam in a laboratory setting. This EMG activity was used to construct six synergies using a non-negative matrix factorization algorithm. From these six synergies, we see that one in particular bares resemblance to our motor primitive 3. The synergy, shown in figure 5-24 prominently represents activity for Tensor Facia Latae, Gluteus and Semimebransosus, (as well as Iliacus Internus, Cruralis, Iliofibularis and Sartorius).

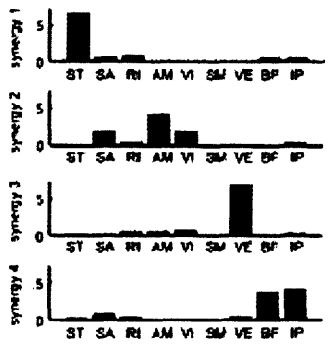

Figure 5-20: Synergies from [67]. Muscle synergies 3 and 4, are similar with motor primitives 3 and 5, respectively.

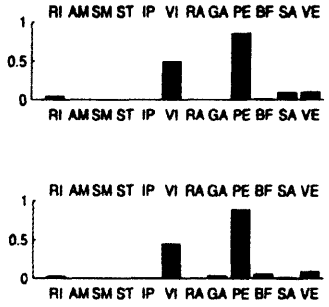

Figure 5-21: Synergies from [57]. Both muscle synergies are similar to motor primitive 2.

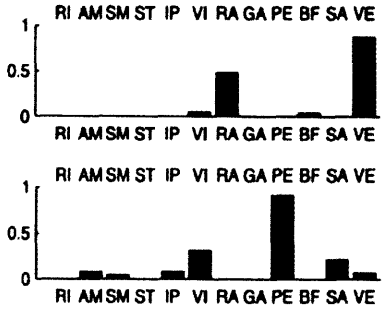

Figure 5-22: Synergies from [57]. The top muscle synergys is similar to motor primitive 3, while the bottom muscle synergy is similar to motor primitive 2.

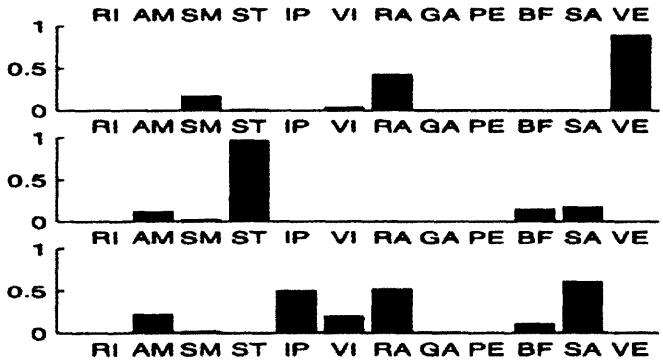

Figure 5-23: Synergies from [57]. Muscle synergy, D, is similar to motor primitive 3, synergy E is similar to motor primitive 4, and synergy F is similar to motor primitive 5.

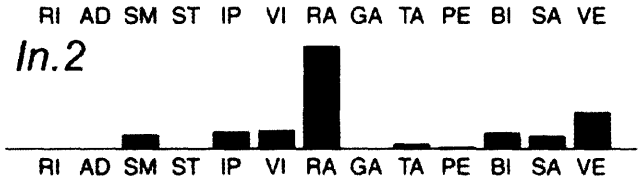

Figure 5-24: Similar synergy from [11], that is similar to motor primitive 3.

5.3.7 Low-dimensional controller

We now investigate control with the linear low-dimensional model and its attendant motor primitives. This will serve two similar purposes. First, it will establish that a controller, based on a reduced order representation can in fact be used to control the hind limb. Second, though synergies can accurately reconstruct EMG signals of motor behaviors, they have never been shown to be capable of generating the commands necessary to recreate these behaviors. There is the possibility that slight inaccuracies in the ability to reconstruct EMG records give rise to inabilities to generate a class of movements. For instance, it could be that motor commands restricted to primitives cannot produce natural reflexive movements. Therefore, by devising motor control restricted to motor primitives, we demonstrate the possibility of control using a low-dimensional model, and some of the motor behaviors possible with motor primitives.

Our first attempt to control the hind limb is to produce commands for natural kicking behavior. We merely seek a non-negative linear combination of motor primitives that can drive the along the approximate trajectory of a kick. To this end, motor primitives (see figure 5-19) were used to generate commands for the hind limb model. The desired limb kinematics for a kick could be defined and used to invert the limb's equations of motion to find the appropriate joint torques. However, to produce activations for each primitive in this manner is a complicated procedure requiring the characterizaiton of how each primitive produces torque as a function of joint position, velocity and activation level. Therefore, primitives were commanded with steps in activation. The time and magnitude of these steps was "hand-tuned" through trial and error to produce limb motion that approximated the kinematics of awake and behaving animal's kicks [17], (unpublished data, Andrea d'Avella).

In figure 5-25 three kicking movements are displayed. In the top plot is a caudally directed kick. It begins with the hip and knee flexed (the red circle at [0.01, -0.01]) and moves downward in the figure (the dotted line). After a brief pause, the limb then retracts approximately back to the initial limb position (the red x). The bottom plot displays a laterally directed kick, while the middle plot displays a kick directed at a location between these two other kicks. All kicks start in the same flexed limb location, and end in approximately the same location.

The command for each kick is composed of two portions, a brief step in command to drive the limb towards the perimeter of the workspace, and a second brief step in command to bring the limb back to the starting orientation. Each portion (the outward going, and returning) of the command was made constructed with only two of the five motor primitives. Thus, the requirement for controling the 13 independent muscles of the limb has been collapsed to specifying five primitives weightings (three of which were not used).

The kicks shown are a first pass at investigating the motor behaviors possible with a low-dimensional model. Unfortunately, the process of designing the commands by hand is lengthy and tedious, so an exhaustive investigation of similar behaviors is not practical. Instead, we turn our attention to designing a controller, rather than commands.

We now demonstrate control design with the linear low-dimensional model. To do

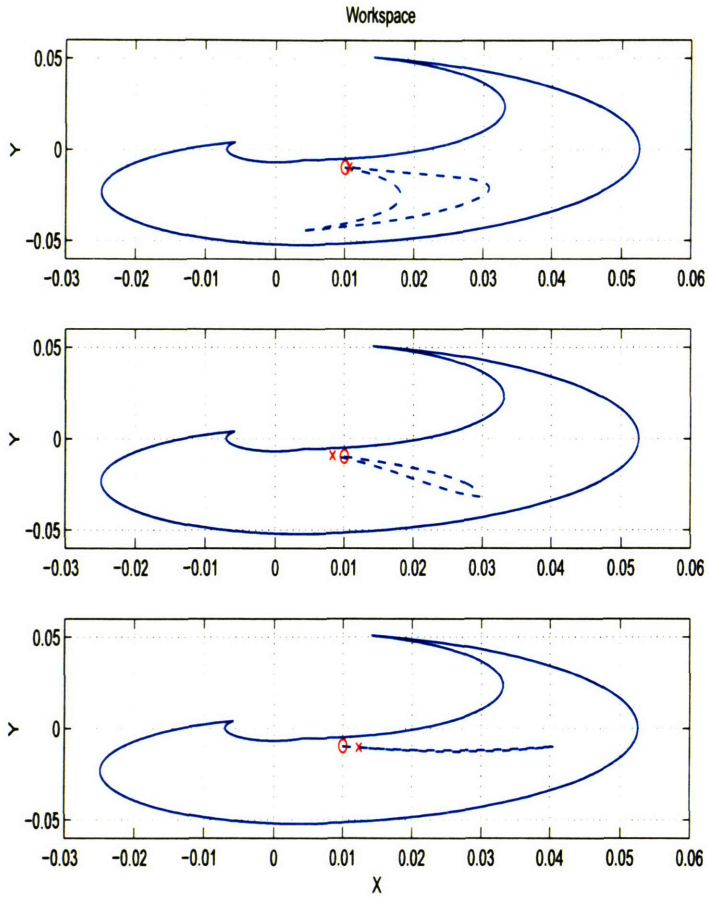

Figure 5-25: Three kicking movements built with hand-tuned primitive step-commands. Each kick starts at the red circle, and ends with the red x.

this we find the optimal solutions for the terminal state problem. As briefly mentioned in the toy problem analsys (section 5.2.3), by introducing the non-negative command constraint into the optimal solution the problem is significantly harder. Therefore, we again remove the positivity constraint on the command signals (some non-negative solutions have been found, but more work remains to complete this task). Again, we solve the Euler Lagrange equations for the reduced system, with a neighboring extremal algorithm. The high order and nonlinear features of the full nonlinear system are such that no solutions have been found, illustrating the difficulty involved in controlling a high-dimensional nonlinear system. Therefore the only data we present are the solutions for the linear low-dimensional system and the result of using this low-dimensional optimal control signal to drive the full nonlinear system. This was done for six target movements to the joint orientations, {(.5 .5), (0 .5), (-.5 .5), (-.5 -.5), (-.5 0), (-.5 -.5)} (units in radians). Figure 5-27 is a plot of the resulting trajectories. The solid lines are the full system's trajectory. The dotted lines are the

low-dimensional system's estimate. The red x's are the desired final position for the system. There is good agreement between the estimated and actual trajectories.

In figure 5-27 is a workspace plot of the same trajectories. We observe that the full-dimensional response (solid lines), again, using the motor command derived from the low-dimensional system, does come close to reaching the desired targets (denoted by x's). The low-dimensional estimate of the system response (dotted line) closely follows the trajectory of the full system.

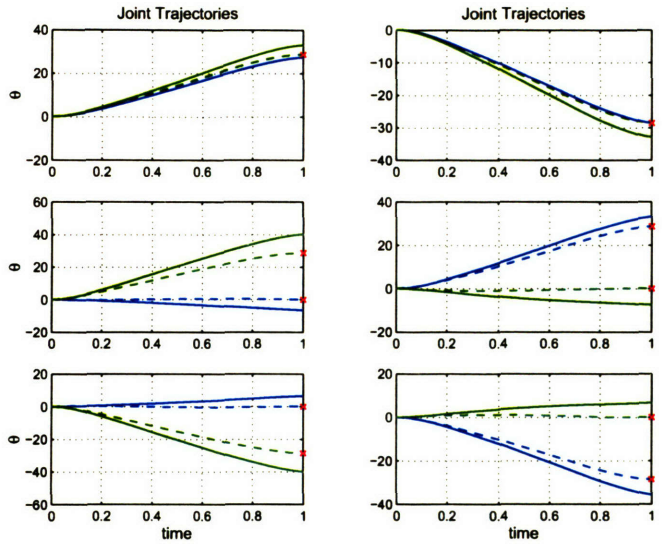

Figure 5-26: Joint trajectories for movements using low-dimensional optimal command solutions. The solid lines are the actual system response, the dashed lines the estimated response.

5.4 Summary

In this chapter we have demonstrated the utility of a reduced order model of the hind limb's musculo-skeletal system. A linear transformation based on empirical gramians was used to compute a reduced order variable, z. It was shown that when z is five (out of 17) dimensional, good transformations/estimations between the full and low-state were possible. Using this low-dimensional transformation and training data, a linear low-dimensional model for estimating the full nonlinear system output (the joint angles and velocities) was computed. This low-dimensional (5^{th} order) model was shown to produce good estimates of the full system output.

We have also presented a procedure for defining motor primitives when the dynamics are linear in the command. We argue that motor primitives, for our biological system, ought to be a best approximation to a non-negative spanning set of the orthogonal compliment of the null space of B_k, $\mathcal{N}(B_k)^{\perp}$. Through this choice

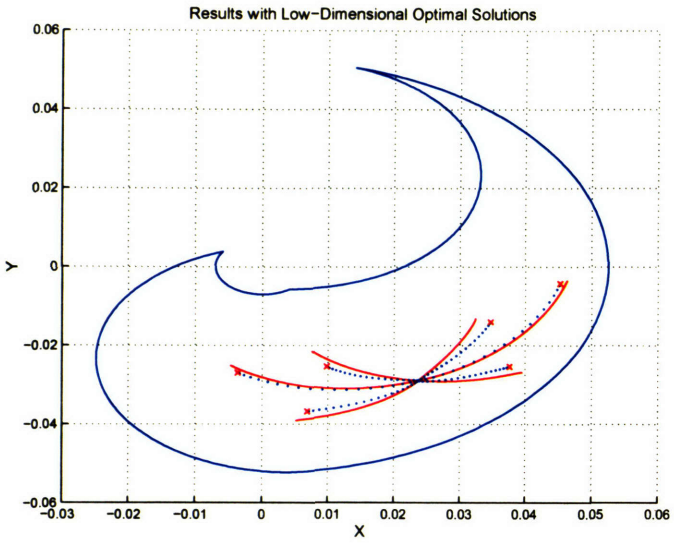

Figure 5-27: Workspace trajectories for movements using low-dimensional optimal command solutions. The blue dashed lines are the estimated system response of with the optimal command. The red solid lines show the movement of the full system when driven with the same optimal commands.

of primitives we are ensured maximal controllability of the low-dimensional system, while still maintaing non-negative motor commands.

It has been demonstrated that the motor primitives are, in part, an attempt to alleviate the redundancies of muscles, as witnessed by their groupings in primitives with consistent effects on limb dynamics. The favorable comparisons with experimentally derived synergies present evidence that our method for determining primitives is promising, if not accurate. This can also be offered as further support that muscle synergies are not an epiphenomenon due to numerical algorithms, and do in fact reflect aspects of the motor system. Furthermore, the correspondence between motor primitives and muscle synergies offers evidence that the frog's control may reflect a similar coordination of muscle activities based on biomechanical properties. This is also evidence for a stronger assertion, that the frog does in fact use a similar control system based on a reduced order representation.

The motor primitives derived from the low-dimensional model have been shown capable of generating commands for natural kicking motions. The kicks are compelling evidence that our motor primitives are capable of generating a large class of natural limb behaviors, though this has yet to be fully explored.

A further demonstration of control employed the low-dimensional model to find optimal control solutions for the full-dimensional system. It was shown that these optimal commands were sufficient to bring the nonlinear limb within close proximity of the desired target locations. Currently, there is no data yet available for optimal solutions obtained using the full-dimensional system, or the reduced nonlinear system,

so no quantitative comparisons with these optimal solutions can be made. However, this is suggestive of the difficulties associated with controlling this large dimensional nonlinear system, and the simplifications afforded with the reduced order model.

Chapter 6

Discussion

Comparison of Linearity Results with Experimental Data

Now we perform a series of computational exercises on the simulated force fields. This is done to demonstrate that the musculo-skeletal model examined in this thesis can produce results similar to the experiments reviewed in the opening chapter, and to then use the linearity analysis of chapter 3 to explain these results.

Mussa Ivaldi et.al. [49] reported the results of isometric force fields at the ankle and their ability to be superposed. Force fields of individual stimulation sites were recorded. Then these sites were costimulated and the resulting force fields were recorded. These costimulated force fields were compared with the linear superposition of the individually stimulated sites. To quantify the similarity between the costimulated force field, and the vector superposition of the individual fields, a "similarity" measure was defined. To do this, an inner product was first defined,

$$< F_A, F_B >= \sum^{N} F_A(x_i) \cdot F_B(x_i) \tag{6.1}$$

where the sum was over all sampled locations of the force fields measurements, $x_1, x_2, ...x_N$, and F_A and F_B are force fields produced by stimulating at sites A and B, respectively. With this they then defined a "cosine" between two force fields.

$$\cos(F_A, F_B) = \frac{< F_A, F_B >}{\|F_A\|\|F_B\|} \tag{6.2}$$

This cosine was used to quantify the similarity of two force fields. Note that $-1 \leq \cos(F_A, F_B) \leq 1$ and if the two fields are proportional to eachother, then the measure is 1. This measure was used to compare the costimulated force fields, F_{AB} with the superposed force fields, $F_A + F_B$.

Of 41 costimulated force fields, 87.8% were found to have a cosine of > 0.9. The average for the entire data set was 0.938 ± 0.045. The results implyed that the force fields did, in a statistically significant way, linearly superpose when simultaneously costimulated.

In a follow up study by Lemay et.al. [41] the results of a similar examiniation were reported. The scaling of force magnitude with stimulation as well as the superposition

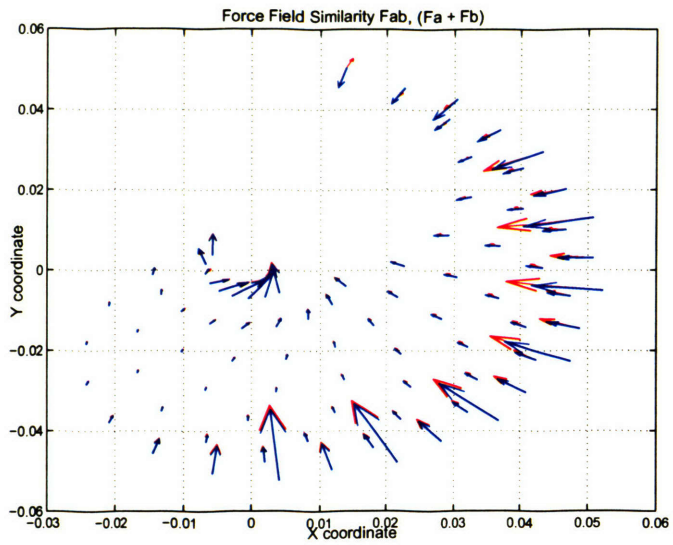

Figure 6-1: Force fields used to compute similarity measurement comprised of Gluteus and Semitendinosus Dorsal. In red is the force field found when the individual fields, F_A and F_B are summed. In blue is the field, F_{AB}.

property were reported. In contrast to the study by Giszter et.al., the total force fields and not simply the active component were examined.

The effects of varying stimulus parameters on force strength were examined under the heading of single site recruitment modulation. Four parameters were systematically varied during electrical stimulus delivery to spinal sites, pulse amplitude, pulse duration, stimulation frequency and train duration. To compare force magnitude with stimulation parameters consistently, each of the stimuli were converted into units of charge (Coulombs). In 16 of 17 spinal sites, the average correlation coefficient between charge delivered and magnitude of resulting forces was > 0.9, implying that the relation between force strength and stimulus was linear.

Field structure was also investigated in terms of changes in force direction with varying stimulus levels. Angular deviations of the force directions at each location measured were computed for varying levels of stimulation at the same spinal location. The average change in angle was $1.3^o \pm 11.8^o$.

The findings of these two papers shall be used for comparison with the thesis model. To compute the same measures with our model of the hind limb, force fields will have to be created. As noted previously, Gluteus and Semitendinosus Dorsal are examples of the most and least linear muscles in our model. Therefore, we shall examine a simulated force field that is composed of these two muscles.

Two force fields were compared. First, a force field that comprised Gluteus and Semitendinosus Dorsal at activation levels of 0.25 and 0.5, respectively, was computed. This is meant to represent stimulation at, say, site A. A second force field with activations levels of 0.5 for Gluteus and 0.25 for Semitendinosus Dorsal was also

computed for stimulation at some site, B. Then the field representing costimulation of sites A and B was computed, by simulating a field with activaiton levels of 0.75 for both Gluteus and Semitendinosus Dorsal. In this way, $F_A + F_B$ could be compared with field, F_{AB}.

The cosine, similarity measurement for these fields was found to be 0.976, within one standard deviation of the experimental results (0.938). For illustrative purposes, the two fields are shown in figure 6-1. The cosine is nearly 1.0 for most of the workspace. In the rostral lateral portion of the workspace, the length of Semitendinosus Dorsal is such that it's muscle force is nonlinear in activation. The forces it produces in this location do not scale linarly with activation. The forces produced by Gluteus do, however, scale with activation, and the resulting force at the ankle rotates. The cosine measure drops to values less than one accordingly.

Continuing on with the comparisons, from the second study, [41], we wish to compute a single site recruitment modulation score for our model. We again use force fields comprised of Gluteus and Semitendinosus Dorsal. Each muscle was activated with levels, 0.1, 0.25, 0.5, 0.75 and 1.0, and the correlation coefficient for the resulting force, at each point in the workspace, was computed. A plot of the correlation coefficient is show in figure 6-2. The correlation coefficient is greater than 0.9 for most of the workspace. The average of all the coefficients computed was 0.9656. Again, in the rostral lateral portion of the workspace Semitendinosus Dorsal reaches lengths that make it behave nonlinearly in activation. This can be seen in the figure. The linearity of Gluteus, however, masks the nonlinear increase in force from Semitendinosus Dorsal, allowing for relatively high ($> .9$) correlation coeficients.

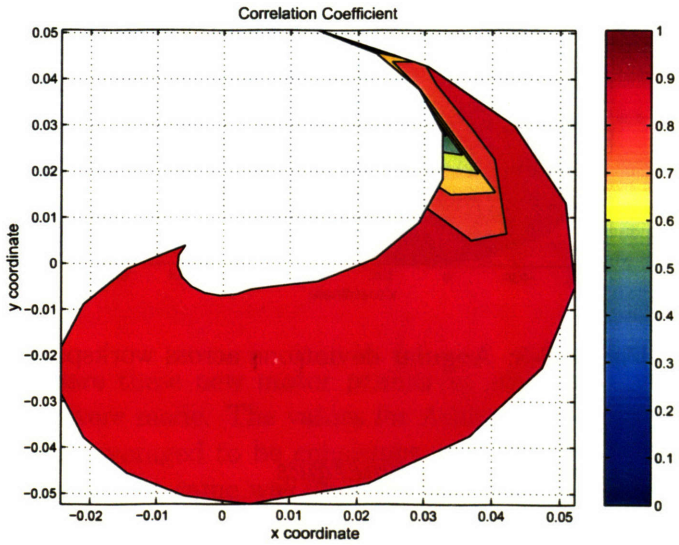

Figure 6-2: Correlation coefficients across workspace for increasing activation.

Using the field structure measure (angular deviations) from the Lemay study, two more force fields were computed for comparisons. The first field was composed of Gluteus and Semitendinosus Dorsal, both with activation levels of 0.1. The second field was the same muscles but now activated at 0.75. At each point measured in the workspace, the change in angle is computed. A plot of the angular deviation across the worksapce is shown in figure 6-3. We note how in the upper right region of the workspace, corresponding to large hip flexion and knee extension, the angluar deviations are large, $> 120^o$. This area corresponds with the nonlinear operating region of Semitendinosus Dorsal, and can be explained as the sole result of this muscle's nonlinear behavior (Gluteus is still very nearly linear in this region). The average acorss the entire workspace is 35.3^o. However, as noted several times now, this is computed over a much larger region of the limb's workspace than done in the study by Lemay et.al. In figure 6-3 we observe that over the approximate region forces were measured in the experimental study, the angular deviations are very small (close to 0^o). This, again, is due to the lengths of Glutues and Semitendinosus Dorsal in this region of the workspace.

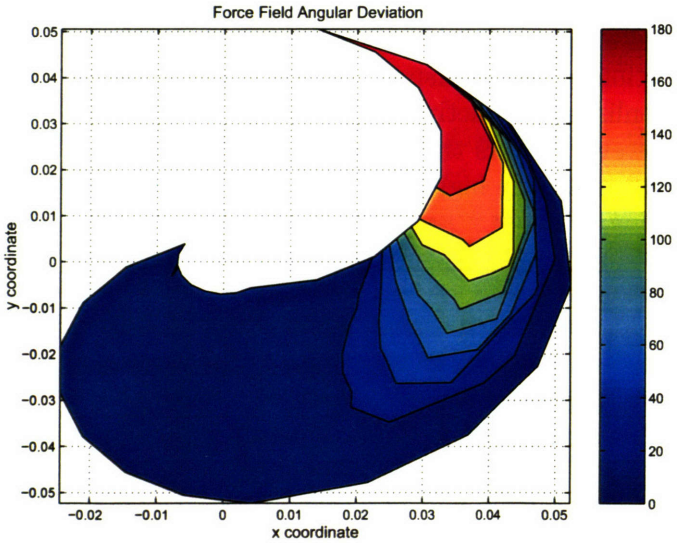

Figure 6-3: Angular deviations across workspace.

Motor primitives and muscle synergies

We have argued (section 5.3.6) that our derived motor primitives have some clear correspondences with muscle synergies reported in the literature. However, throughout the thesis we have distinguished between motor primitives and muscle synergies, in name (and spirit). This is done for a number of reasons. As was discussed in section 5.3.6, if motor primitives, similar to the ones we have defined in this thesis, are used as a basis for generating motor commands, then the muscle synergies found

experimentally would not correspond exactly to the primitives, since they are not normalized with respect to maximal EMG activity. Furthermore, various extraction algorithms for finding low-dimensional features in data might result in different muscle synergies, another source of possible discrepency between motor primitives and synergies.

That being said, we have a unique opportunity to compare synergies and primitives quantitatively, regardless of how strong or tenuous their connection might be. For comparitive purposes, we have acquired muscle synergies from the study by Cheung et.al. [11]. A set of six synergies, obtained from EMG data collected while a frog swam, were found using a non-negative matrix factorization algorithm. These muscle synergies are compared with a set of six motor primitives. The motor primitives were found using the procedures outlined in section 5.3, only with a low dimensional variable of dimension 6 ($k = 6$). Figure 6-4 is a plot of the six motor primitives. They are similar to the motor primitives presented previously for the case $k = 5$ (figure 5-19), with each primitive grouping muscles that have similar biomechanical influences on the system dynamics.

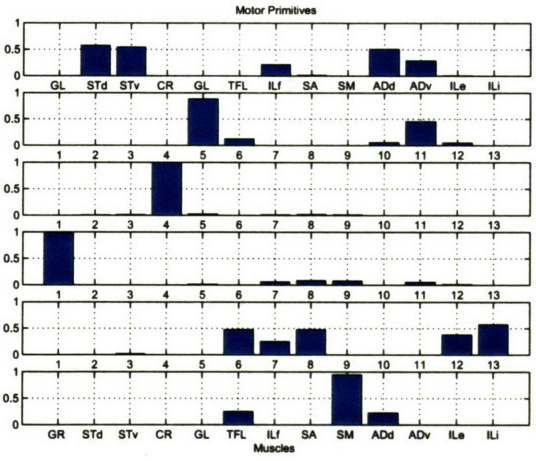

Figure 6-4: New motor primitives derived from a low-dimensional model for the case $k = 6$.

In order to compare these new motor primitives with the muscle synergies, the following alterations were made. The values for Adductor Dorsal and Ventral in the motor primitives were averaged to be consistent with the single value for Adductor in the muscle synergies. The same was done for Semitendinosus Dorsal and Ventral. Iliacus externus is neglected from our motor primitives as there is no measurement for this muscle in the synergies. Finally, Gastrocnemius, Tibialus Anticus and Peroneus are neglected from the muscle synergies as those muscles are not in our hind limb model and therefore, not in our motor primitives. Having made these changes, each motor primitive now has ten elements, consistent with the muscle synergies. The resulting altered motor primitives and the set of six muscle synergies are shown in figure 6-5.

The first comparison we make between muscle synergies and motor primitives is their "orientation", or direction in the vector space of motor commands. We seek to quantify how similar the primitives and synergies are in terms of an inner product. If the inner product between a motor primitive and a muscle synergy is one (they both have unitary norm) then they produce equivalent commands; if it is zero, then they produce commands that are orthogonal. In figure 6-6 is a matrix of inner product values between our six motor primitives and the muscle synergies. The vertical axis corresponds to the primitives and the horizontal axis to the synergies. Therefore, for instance, we see that motor primitive 6 and muscle synergy 1 have a large inner product, 0.93. Similarly, motor primitive 1 and muscle synergy 2 have a large inner product, 0.81.

For random vectors with unitary norm (chosen from a uniform distribution between zero and one) the inner product is approximately 0.76. We find that the first three motor primitives (the top three in figure 6-5) have inner products with muscle synergies of 0.77 or larger. Therefore, the first three motor primitives have a correspondence with the muscle synergies that is greater than chance.

The inner products compare the motor primitives and muscle synergies individually, without explicit reference to the commands they can generate. However, as mentioned earlier (and discussed shortly in further detail), there may be good reasons to suspect that synergies and primitives may differ. Muscle synergies, extracted using different algorithms, might differ even when found from the same data set. Yet, within error tolerances, regardless of the synergy extraction algorithm, the space the various sets of synergies span (of which the EMG data belongs to) should be invariant. That is, no matter how muscle synergies are computed, the set of motor commands they can produce should be the same. The next comparison we make will be an attempt to address this fact. We compare the spaces of commands spanned by the synergies and the primitives. To this end, we shall look at the portion of each motor primitive that is orthogonal to the span of the muscle synergies.

Figure 6-7 is a plot of the component of each motor primitive (from our set of six altered primitives, figure 6-5) that is orthogonal to the space spanned by the six muscle synergies. If the entries for each primitive are zero, then that primitive is entirely within the space of motor commands spanned by the muscle synergies. If, on the other hand, elements in a primitive's orthogonal components have large nonzero values, then that primitive commands combinations of muscles that are not available to the muscle synergies. From the figure we see that there are some clear differences between the commands motor primitives can generate and those of muscle synergies. For instance, it appears that (based on the second, third, and fourth primitives) our motor primitives are a better set of commands for generating nearly individual commands to Gluteus, Cruralis, and Gracilus (the same could be said for Iliacus Internus and Adductor as well, only these commands are not as strong). It is difficult to interpret these results physiologically as this set of muscle synergies is based on EMG data during swimming, and our motor primitives are based on a model that does not take into account interaction with the environment.

Furthermore, we can conclude that the first, fifth and sixth motor primitives are almost entirely within the span of the muscle synergies. This indicates that any

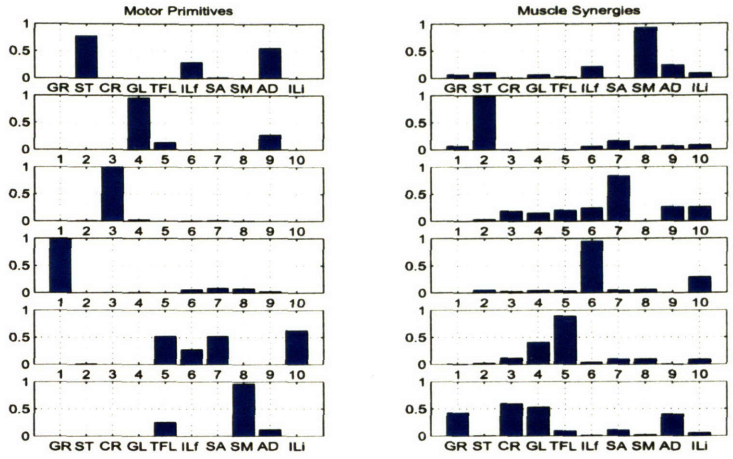

Figure 6-5: Altered motor primitives and muscle synergies for $k = 6$. The left plot displays the set of 6 motor primitives based on a 6^{th} order low-dimensional model. The right plot is the set of six muscle synergies from Cheung et.al.

motor commands generated using these three primitives can be explained in terms of linear combinations of the muscle synergies. This overlap between three of six of our primitives (1, 3 and 6) and the span of the synergies, as well as the large inner product between another three of our primitives (1, 2 and 3) and the synergies, indicates a correspondence between some of our motor primitives and the experimentally derived muscle synergies.

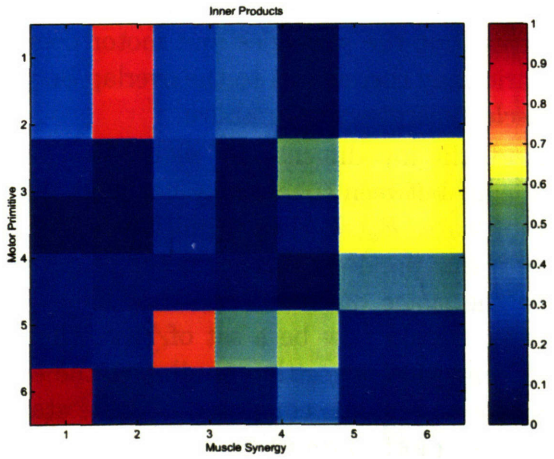

Figure 6-6: Inner products between motor primitives and muscle synergies. The vertical axis corresponds to the six motor primitives, while the horizontal axis refers to the six muscle synergies.

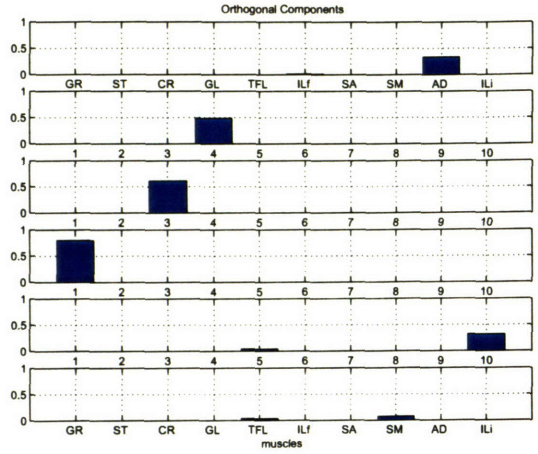

Figure 6-7: Components of each motor primitive that are orthogonal to the span of the muscle synergies.

The correspondences between motor primitives and muscle synergies are an interesting outcome not to be dismissed. Yet, there are still good reasons to distinguish between primitives and synergies. In addition to the reasons mentioned earlier, another reason to distinguish between motor primitives and muscle synergies is the "space" they occupy. For instance, in our reduced order model, the dynamics for the low-dimensional variable, z, were linear in the command signals. This was done on purpose to allow for easy definitions, and interpretations, of motor primitives; they were chosen to be a non-negative linear basis for the space of commands that influenced the low-dimensional dynamics. A numerical extraction algorithm operating on motor commands derived with our low-dimensional model would most likely find a basis for the same space that our primitives span. The synergies themselves might differ from the primitives, but the synergies and motor commands would span the same space. This scenario may correspond to the overlap between three of our motor primitives and the muscle synergies shown above.

If, on the other hand, the low-dimensional model was not linear in motor commands, this correspondence between spaces may not hold. For instance, if the input matrix was a function of state, $B_k(z)$, then primitives would have to be defined differently. In keeping with our previous theme, motor primitives could be defined as non-negative bases for the space of commands that influence the low-dimensional dynamics. This, however, would now be a set of bases that were state dependent, and varied in time. A numerical algorithm similar to the ones used for extracting synergies, acting on motor commands composed of these state-varying bases, would extract more synergies than motor primitives.

To understand why a numerical algorithm would uncover more muscle synergies than there are motor primitives, consider a simple case and assume there is only one, state-varying motor primitive (in this case $z \in \mathbb{R}$). If this single primitive were constant in value, all motor commands produced with the low-dimensional model

would lie on a line, and a single synergy could explain the resulting data. However, for the state-varying case, as z evolves in time, the motor primitive changes, and the commands it produces sweep out an arc in the space of motor commands. These commands would not lie on a line, and two or more synergies would be required to explain the data.

An example of this scenario is schematically depicted in figure 6-8. In the left plot is an image of motor commands produced when the motor primitives are constant. Two constant motor primitives define a plane, in a three dimensional space of motor commands. All motor commands will lie in this plane, and when an extraction algorithm is used to find synergies in these data, labeled u, two synergies will be found. The span of these synergies will be coincident with the plane. In the right plot is an image of the same motor commands, produced by a single state-varying motor primitive. As this primitive changes in orientation, it sweeps out a plane (in general it will not be a plane and can be a nonlinear surface) upon which the motor commands are contained. If the same extraction algorithms were used on these motor commands, the same synergies would be found (again, we assume the motor command, u is the same in both examples). It is clear in this case that there is no good correspondence between the motor primitives and the muscle synergies. Indeed, in this case, muscle synergies are merely a best approximation to a linear basis for the space that contains motor commands. This could be the case for the experimentally found synergies of Cheung et.al., and would explain why three of our motor primitives did not overlap with the span of the muscle synergies.

Finally, we consider a further reason for making a distinction between motor primitives and muscle synergies: the effects of proprioceptive feedback. Proprioceptive feedback has not been included in any of this investigation, either at the level of modeling, or in the design of control. In the frog, feedback most certainly does play a role, as motor behaviors with and without feedback have markedly distinct charac-

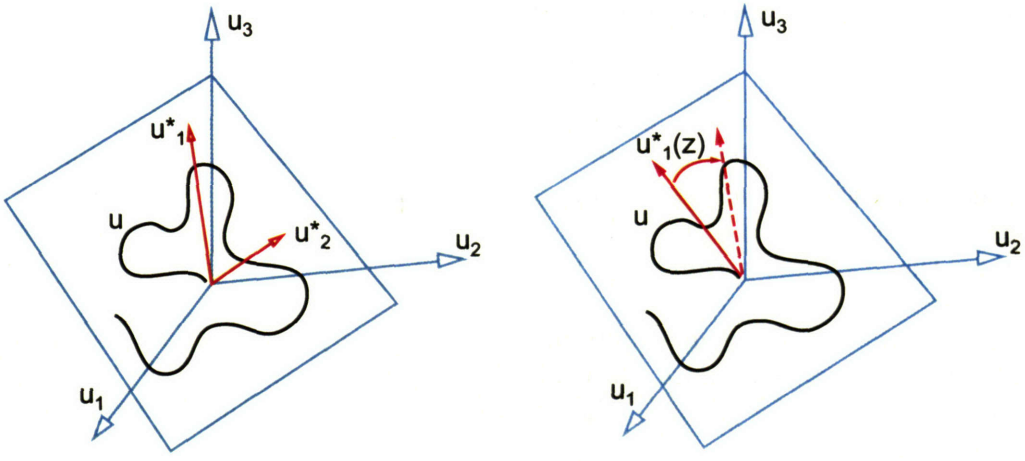

Figure 6-8: The space of motor commands for the cases of linear, constant, and state-varying, motor primitives.

terisitics (e.g. [35]). The motor primitives in this thesis are a feedforward basis for motor commands, as they are not influenced by system state. However, experimental evidence demonstrates that feedback does play a role in the muscle synergies found [11]. This further emphasizes the need to distinguish between motor primitives and muscle synergies.

In summary, there is good reason to distinguish between primitives and synergies. The strong correspondences between our motor primitives and muscle synergies, both in terms of the inner products and the spaces of motor commands spanned, are interesting and compelling evidence, but not strictly necessary outcomes of a motor control system that uses a reduced order representation.

Motor primitives and force fields

Based on our findings, we can offer an explanation for the small number of linear isometric force fields observed experimentally, and their apparent use as modular primitives in motor behaviors of spinalized animals. If a reduced order representation for the hind limb resides in the circuits of the spinal cord, then exciting the spinal cord may trigger motor primitive commands. If the excitation is strengthenend, the motor primitive commands are further increased. If two or more sites within the spine are costimulated, they may produce commands that are superpositions of motor primitive commands. Due to the biomechanic properties of the limb, as examined in chapter 3, the resulting forces will be linear in stimulation, consistent with the empirical evidence.

Reflexive motor behaviors, mediated by spinal circuits, could be based on fixed temporal sequences of motor primitive commands. If, as is done in experimental conditions, the hind limb is restrained and one such reflexive behavior is evoked, then the resulting forces measured would appear to be modular force fields. This too, is consistent with experimental evidence.

Equilbrium point control and biomechanical regularities

Several studies, e.g. [5], [25], have reported isometric force fields that are termed convergent, that is they have stable equilibria, within the area the limb's reach measured. These same studies, and others like them [43], have shown that the likelihood of random muscle activations producing convergent force fields, is small. Results like these have been viewed as an indication that spinal motor circuits encode information for the hind limb's equilibria. This could be relevant for an a equilibrium point controller.

To reap the benefits of an equilibrium point controller, over those of say a controller that utilized inverse dynamics, the controlling commands ought to have behaviorally relevant equilibria. For instance, allowing for stable equilibria that lie along a typical limb trajectory.

The motor primitives in this thesis (see Appendix), however, do not contain equilbria within the limb's workspace. The only stable equilbria for these motor primitive induced force fields are at the edge of the limb's workspace, where the joints reach

their kinematic limits. Yet our primitive force fields, unlike force fields examined experimentally, do not have proprioceptive feedback. The effects of feedback, combined with our motor primitives, could certainly allow for the introduction of different equilibria, but this has not been investigated. Instead, what has been observed is that our motor primitives reflect groupings of potentially redundant motor activity. The primitives in this thesis reflect biomechanical regularities that ease the computations of motor control, as opposed to equilibria that are relevant for a posture-based controller.

Low-Dimensional optimal controllers

A low-dimensional model, based on an approximate balanced truncation has demonstrated the capacity to generate control signals relevant for use in controlling the full-dimensional system, when trying to minimize motor commands (as defined through equation 5.20). The same low-dimensional model is capable of predicting the outcome of these control signals as well.

We have not been able to use the low-dimensional system to compute optimal commands restricted to non-negative values. Similarly, optimal solutions for the full-dimensional system, with or without the non-negative restriction, were not found either. Therefore, comparisons between the low-dimensional and full-dimensional optimal commands could not be made. Nevertheless, low-dimensional control reduces computational burden, simplifying motor control. Moreover, based on the results from linear examples (section 5.2.3), and the low-dimensional model's ability to predict the full system's output well, we are optimistic that optimal commands derived from the reduced order model are comparable to those that would be obtained using the full system dynamics. Further analysis is required for conclusive results.

In the event that commands computed using the low and full-dimensional systems differed, this would raise the possibility for testable predictions. If it was found that the hind limb trajectories predicted from a low-dimensional controller, and a controller that utilizes the full-dimensional system dynamics, differed in distinct ways, experiments could be designed to test these predictions. This would help further our understanding of the organization of motor control.

A possible connection with sub-movements

Human movements appear to be composed of discrete units of motor activity [16], [61], [40]. These units are fairly typical in their temporal character. It is as if the motor commands are "chunked" in time. This intermittent feature of motor control, sometimes termed sub-movements, has led to assertions of a sampled feedback controller [47], perhaps to avoid destabilizing feedback delays [32].

In analogy with this temporal discretization, control signals with a reduced order model could be described as "chunking" commands in space. It is tempting to speculate that just as a reduced order model simplifies control, intermittency could be a signature of further attempts to reduce computational burdens.

A reduced order representation could produce estimates of the system outputs for use with a feedback controller, allowing motor commands to be computed while

avoiding destabilizing feedback delays. Or the reduced order representation could produce feedforward estimates of optimal motor commands. Given the fidelity of the reduced representation, the information it estimates may need to be updated and recomputed, perhaps many times, before a motor goal is achieved. The results of such an episodic use of a reduced order representation would most certainly be consistent with the empirical evidence. Although pure speculation, we feel this line of reasoning deserves further investigation.

Chapter 7

Conclusions

7.1 Thesis Contributions

In this thesis we set out to examine two questions: 1) can the linearity observed in isometric force production be explained solely in terms of biomechanical properties? 2) can a compact description of the system dynamics be used for motor control, and offer an explanation for the low-dimensional phenomena associated with frog motor behaviors? Both questions were examined in the context of the simplifying strategies that biological systems might employ for the motor control problem.

To answer the first question, regarding forces that scale linearly in activation, a physiologically realistic model of the frog hind limb was examined. A measure was defined to quantify the muscle's deviation from linear behavior, and the measure was then used to explain this behavior in terms of a normalized muscle length. This analysis indicated that the biomechanical system does allow for linearity in activation, when the musculo-tendons are operating in low-scoring regions of their workspace. Due to the transduction of muscle forces to forces at the limb's endpoint, forces measured at the animal's ankle will exhibit this same linearity.

Through counterexample, we have shown that the linearity in force production is not an inevitable result of our musculo-tendon model, but a result of properly chosen parameter values. Furthermore, the parameter values that accurately describe the animal in question do allow for the linearity. This was verified through simulations meant to reproduce the isometric force field experiments that have exhibited this linear phenomenon. Based on these results, we argued that the central nervous system need not employ linearizing controllers to exploit this property, as it may be a natural biomechanical property. We also propose that the linearity is an evolved property, allowing the animal to manipulate force production through linear modulation of the controlling signals and possibly simplifying control efforts.

Regarding a low-dimensional controller, we demonstrated that a reduced order model based on an approximate balanced truncation is capable of generating estimates of the input-output dynamics of a musculo-skeletal model of the hind limb of *Rana pipiens*. We have shown that for the purposes of the model presented in this thesis, a linear reduced order model of the limb would suffice as well as a nonlinear model.

Based on the assumption of low-dimensional dynamics that are linear in the command, we devised a procedure for defining motor primitives. We propse that motor primitives should be chosen as a best approximation to a non-negative spanning set for the orthogonal compliment of the low-dimensional input matrix's null space, $\mathcal{N}(B_k)^{\perp}$. This choice of motor primitives assures that the low-dimensional system is maximally controllable, given the non-negative constraints on motor commands.

We have shown that control with a reduced order model allows for an explanation of the low-dimensional motor commands experimentally observed in muscle synergies, and the appearance of force field-like primitives in reflexive motor behaviors. Muscle commands derived using the low-dimensional model are necessarily constrained to a low-dimensional subspace of the available motor commands. Therefore, the activity of muscles during behaviors commanded with the low-dimensional model will exhibit low-dimensionality in the form of synergy-like commands. Furthermore, if reflexive motor behaviors are based on stereotyped motor primitive based commands, the restrained limb will exhibit what appear to be modular force fields when said reflexes are elicited. These modular force fields are a natural outcome of the limb's linear force production and motor primitive commands.

Validating the utility of the reduced order model as a means to control the full-dimensional nonlinear limb, we generated motor behaviors using the low-dimensional model. Kicks, which have similar kinematics to natural behaviors, were generated by tailoring motor primitive commands to drive the limb towards the extermities of the workspace. This was done relatively simply with step commands, and confirmed the ability of motor commands restricted to primitives to arrive at natural motor behaviors. The solutions of optimal control problems using the low-dimensional model were also used to command the full system. It was found that these commands were sufficient to move the limb to the approximate location of the desired target locations. This suggested the low-dimensional model may approximate the optimal solutions for the full-dimensional dynamics. These optimal control solutions did not enforce the non-negative constraint on the control signals; however, we feel this is not a deficiency of the controller, but simply a practical limitation of the computational methods employed to find optimal solutions, it is also compelling evidence as to how difficult motor control for the full-dimensional system is, and how the low-dimensional representation may help.

Despite all the known discrepencies between our model and the real animal, several of the motor primitives found in the thesis agree well with experimentally derived muscle synergies. This alone is compelling evidence that the biological system may employ a reduced order representation and use motor primitives derived accordingly. We also argue that due to the small order of the reduced model, its use in estimation and control is a great simplification over computations that would otherwise require a complete description of the full system dynamics.

7.2 Future Work

There are still a great number of issues left to explore with this work. The need to examine more examples of low-dimensional controllers is foremost on our agenda. A few typical movements (kicks) were generated using primitive commands. However, if we find the reduced order model and its attendant motor primitives are not sufficient to produce the range of motor behaviors desired, options are available. Multiple reduced models, based on approximate balancings at multiple locations in the workspace, could be found. This would allow for better approximations to the full nonlinear dynamics, and perhaps motor primitives that varied as a function of workspace location.

The low-dimensional optimal controllers can be examined further as well. Though optimal command solutions that are restricted to non-negative values have not been obtained, the underlying motivation for explaining the low-dimensional features of motor control, and a connection to synergies, has been fruitful. If optimal non-negative commands found using the low-dimensional system were eventually found to approximate the same optimizations for the full-dimensional dynamics, this would further support the thesis work.

Appendix A

SIMM model parameters

Muscle	l_o^m (m)	l_s^T (m)	l_s^T/l_o^m	F_o^m (N)	mono-articular?
ADd	15.66 E-3	3.27 E-3	0.21	1.89	Yes
ADv	14.02 E-3	6.54 E-3	0.47	2.24	Yes
CR	10.85 E-3	19.77 E-3	1.82	14.19	No
GL	15.95 E-3	19.01 E-3	1.19	3.9	No
GR	14.82 E-3	9.31 E-3	0.63	6.46	No
ILf	11.05 E-3	14.77 E-3	1.34	1.53	No
ILe	4.31 E-3	5.49 E-3	1.27	2.7	Yes
ILi	9.43 E-3	1.31 E-3	0.14	1.63	Yes
SA	19.88 E-3	4.81 E-3	0.24	1.01	No
SM	21.55 E-3	5.45 E-3	0.25	4.81	Yes
STd	7.46 E-3	17.23 E-3	2.31	1.51	No
STv	7.80 E-3	16.72 E-3	2.14	1.46	No
TFL	9.96 E-3	19.92 E-3	2.00	1.19	No

Table A.1: Musculo-Tendon Properites

Segment	Mass (Kg)	I_x (Kgm^2)	I_y (Kgm^2)	I_z (Kgm^2)
Pelvis	3.6 E-3	2.5 E-7	3.1 E-7	4.2 E-7
Femur	2.0 E-3	0.6 E-7	1.9 E-7	2.3 E-7
Tibiofibular	1.4 E-3	0.2 E-7	1.2 E-7	1.4 E-7
Tarsals	0.4 E-3	0.1 E-7	0.8 E-7	0.8 E-7
"Foot"	0.5 E-3	0.1 E-7	0.2 E-7	0.2 E -7

Table A.2: Skeletal Properties

Appendix B

Isometric Force Field Error in Linearity Plots

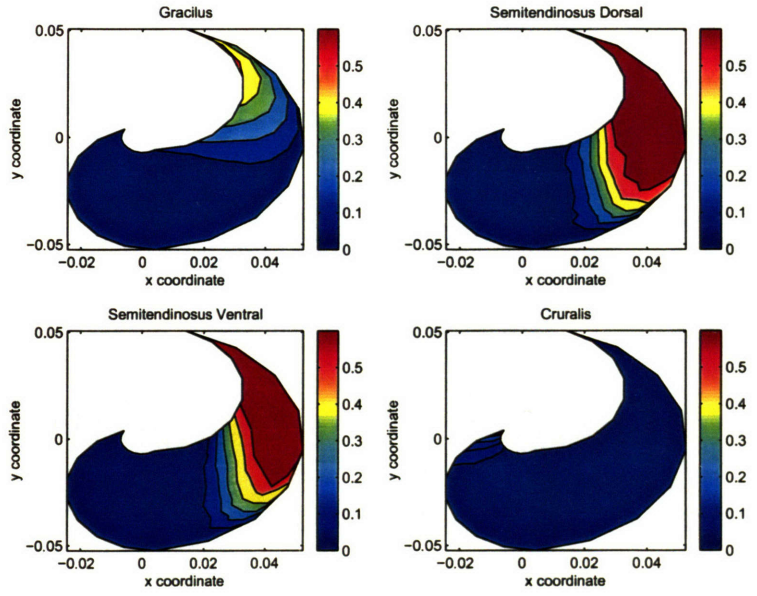

Figure B-1: Plot of error in linearity found through simulation of SIMM model for Gracilus, Semitendinosus Dorsal and Ventral, and Cruralis.

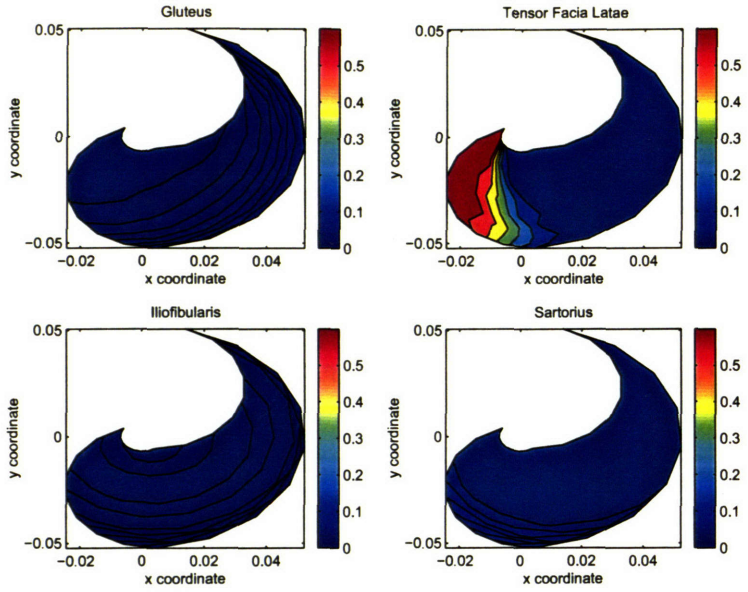

Figure B-2: Plot of error in linearity found through simulation of SIMM model for Gluteus, Tensor Facia Latae, Iliofibularis and Sartorius.

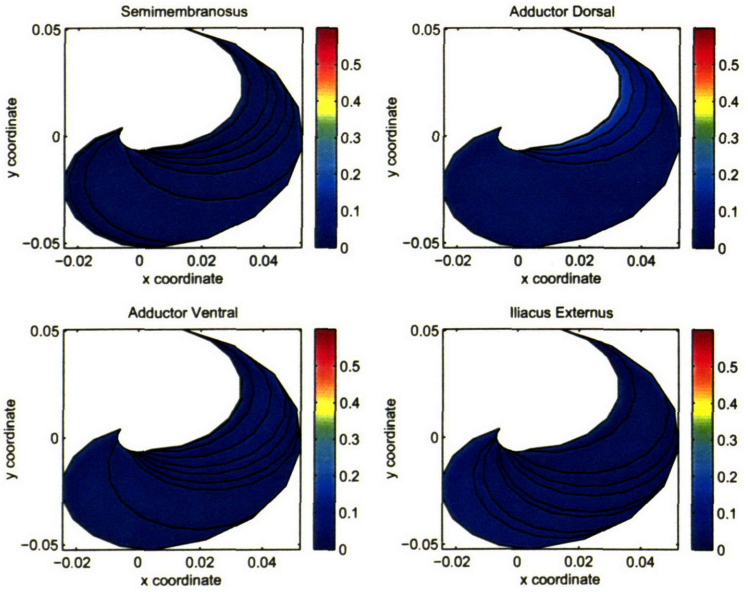

Figure B-3: Plot of error in linearity found through simulation of SIMM model for Semimembranosus, Adductor Dorsal and Ventral, and Iliacus Externus.

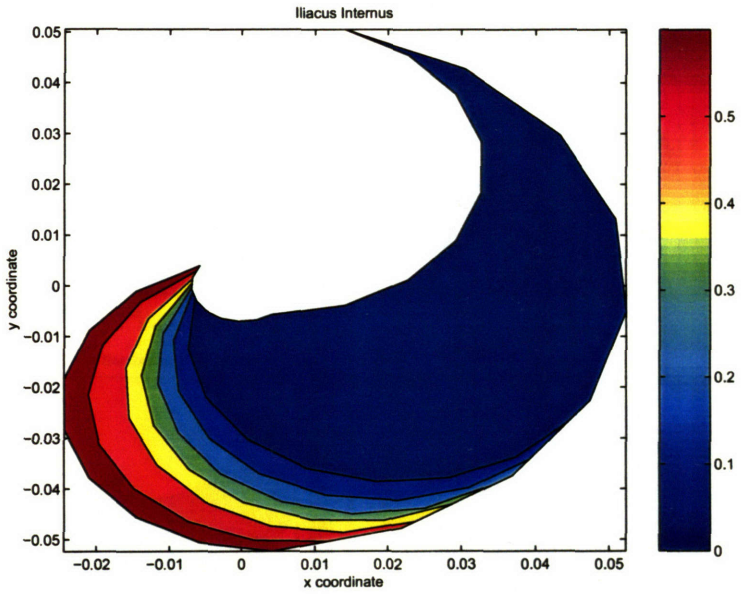

Figure B-4: Plot of error in linearity found through simulation of SIMM model for Iliacus Internus.

Appendix C

Musculo-Tendon Length Plots

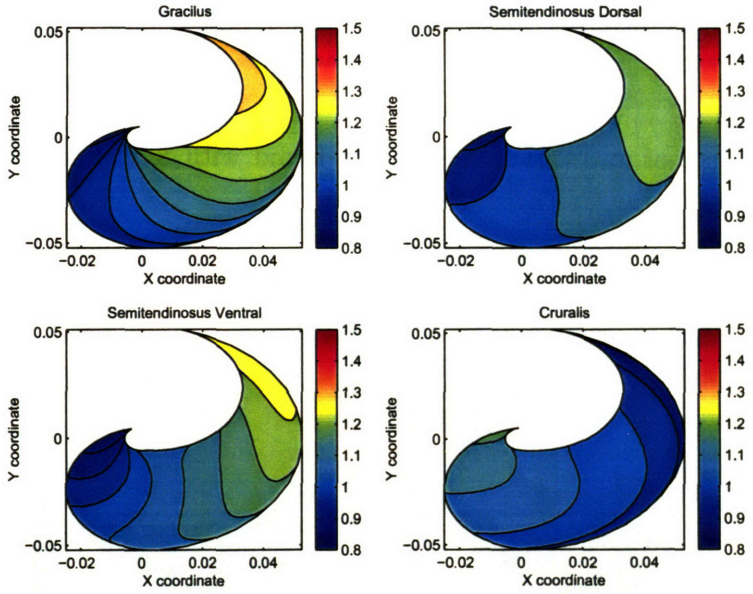

Figure C-1: Plots of musculo-tendon length found with SIMM model for Gracilus, Semitendinosus Dorsal and Ventral, and Cruralis.

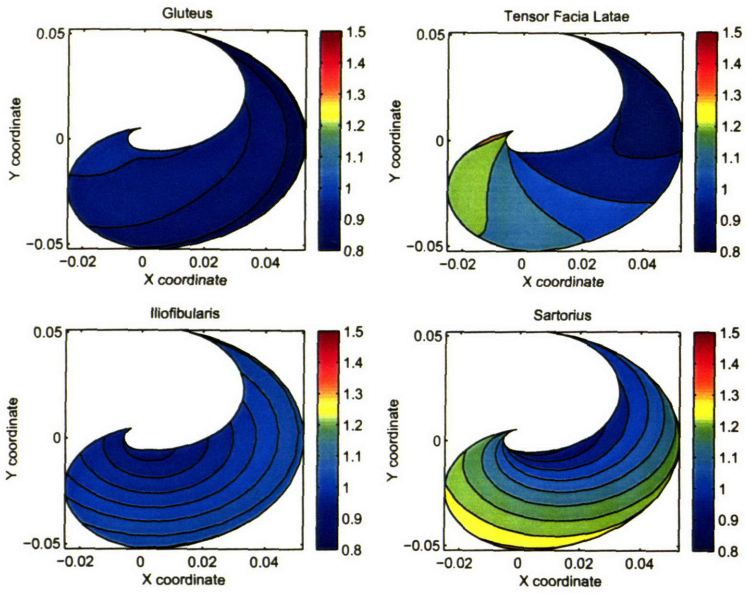

Figure C-2: Plots of musculo-tendon length found with SIMM model for Gluteus, Tensor Facia Latae, Iliofibularis and Sartorius.

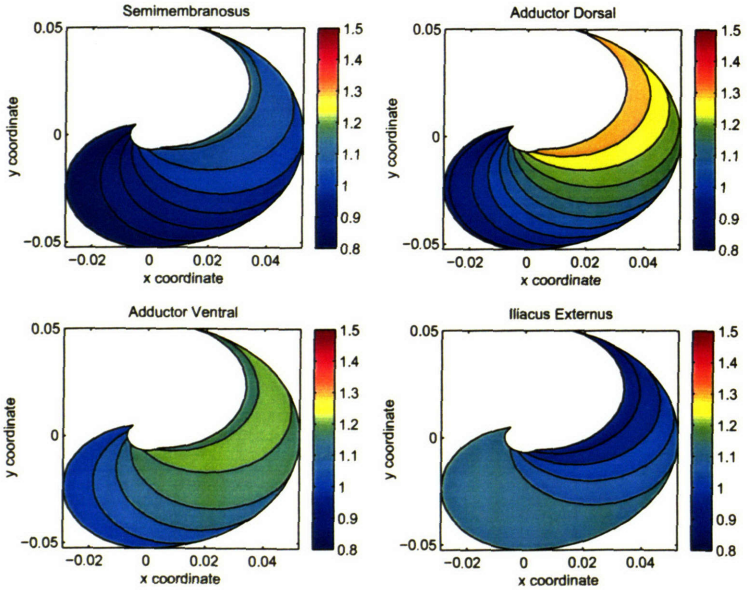

Figure C-3: Plots of musculo-tendon length found with SIMM model for Semimembranosus, Adductor Dorsal and Ventral and Iliacus Externus.

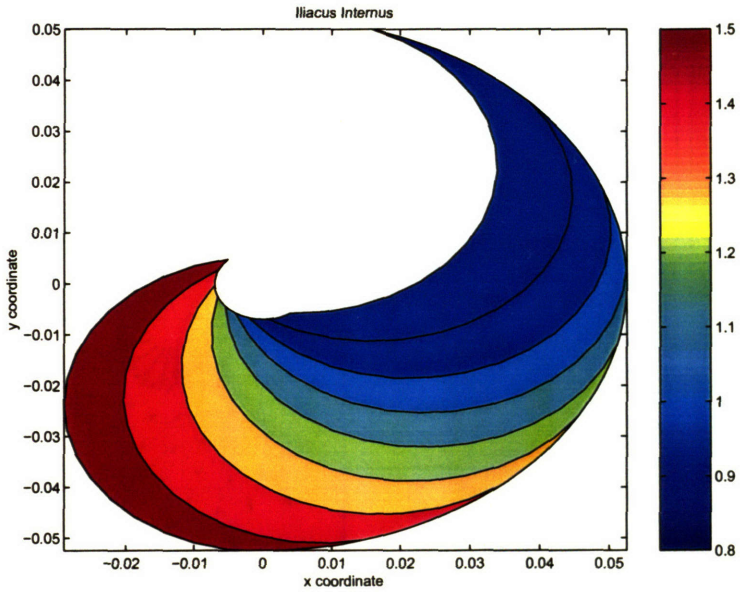

Figure C-4: Plots of musculo-tendon length found with SIMM model for Iliacus Internus.

Appendix D

Motor Primitive Force Fields

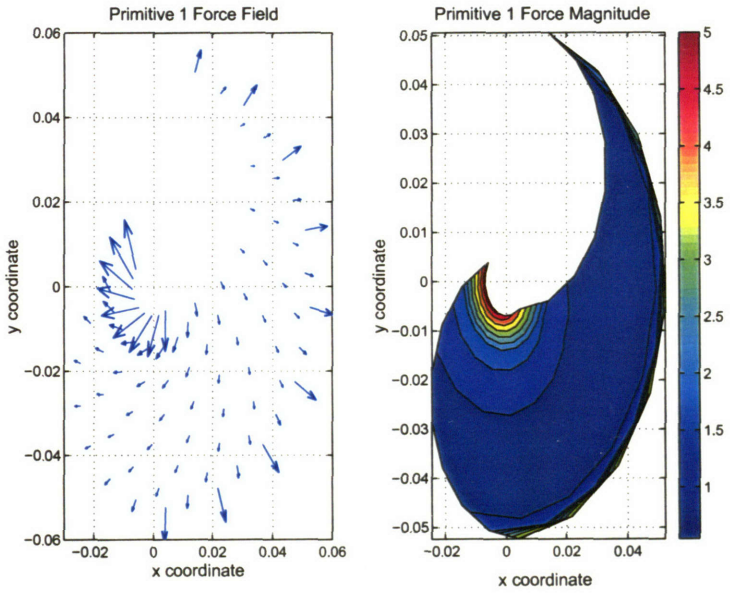

Figure D-1: Force field for motor primitive 1 (maximally activated).

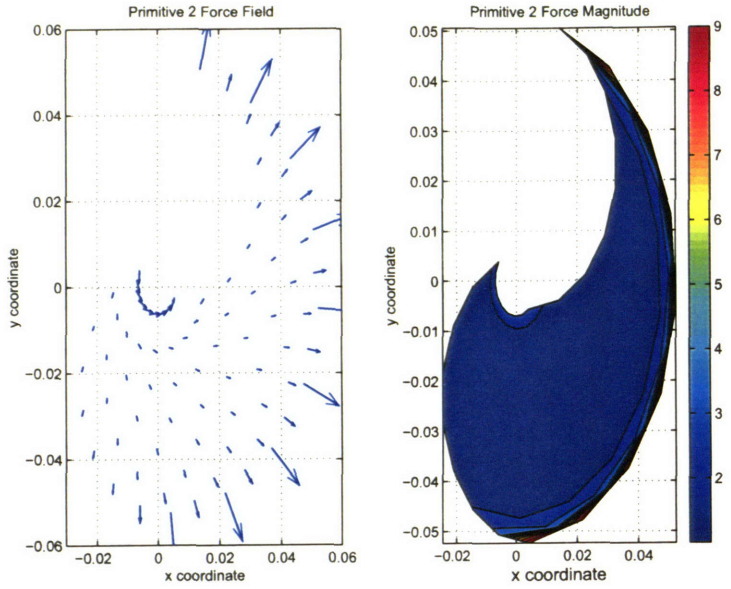

Figure D-2: Force field for motor primitive 2 (maximally activated).

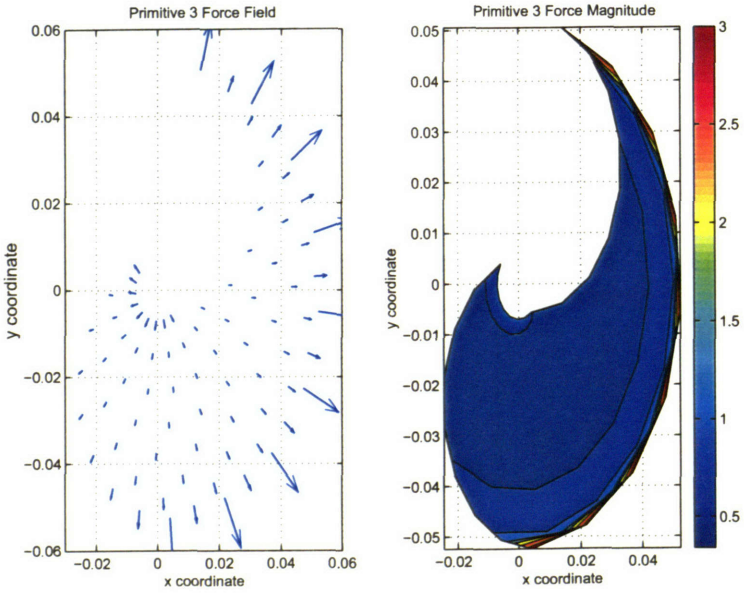

Figure D-3: Force field for motor primitive 3 (maximally activated).

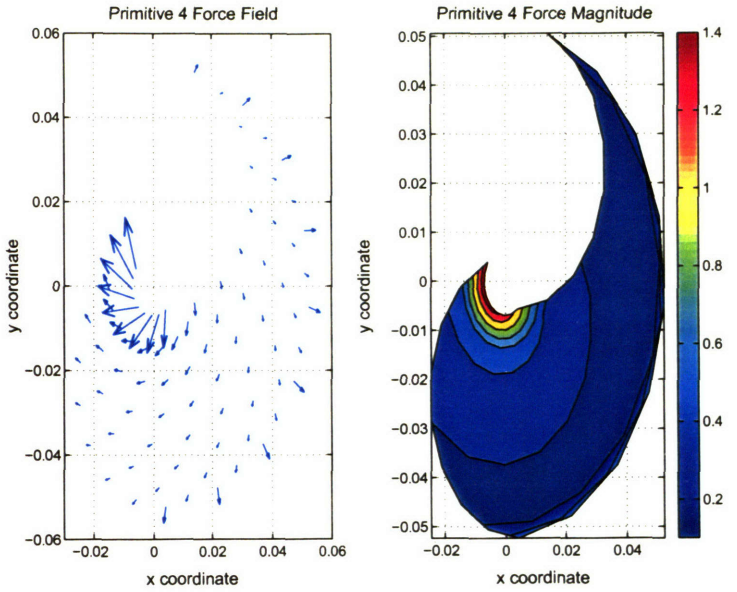

Figure D-4: Force field for motor primitive 4 (maximally activated).

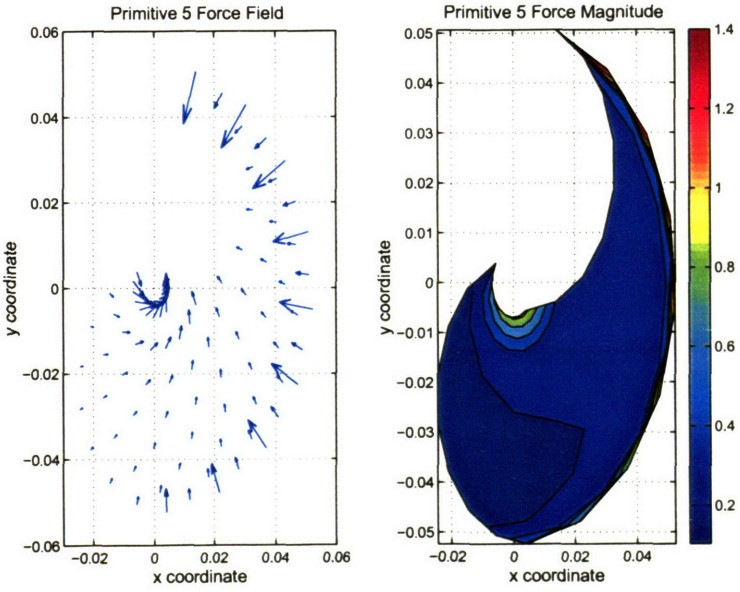

Figure D-5: Force field for motor primitive 5 (maximally activated).

Bibliography

[1] A.N. Ahn, R.J. Monti, and A.A. Biewener. *In vivo* and *in vitro* heterogeneity of segment length changes in the semimembranosus muscle of the toad. *J. Physiology*, 549(3):877–888, 2003.

[2] A.C. Antoulas and D.C. Sorensen. Approximation of large-scale dynamical systems: An overview. *MTNS*, pages 19–23, June 2000.

[3] A.C. Antoulas, D.C. Sorensen, and S. Gugercin. A survey of model reduction methods for large-scale systems. *AMS-ISM-SIAM Summer Research Conference on Structured Matrices*, June 1999. june 27-July 1.

[4] R. Bellman. *Adaptive Control Processes: A guided Tour.* Princeton University Press, 1961.

[5] E. Bizzi, S. Giszter, and F.A. Mussa-Ivaldi. Computations underlying the execution of movement: a novel biological perspective. *Science*, 253:287–291, 1991.

[6] E. Bizzi, N. Hogan, F.A. Mussa-Ivaldi, and S. Giszter. Does the nervous system use equilibrium-point control to guide single and multiple joint movements? *Behavioral and Brain Sciences*, 15:603–613, 1992.

[7] E. Bizzi, M.C. Tresch, P. Saltiel, and A. d'Avella. New perspectives on spinal motor systems. *Nat. Rev. Neurosci.*, 1:101–108, 2000.

[8] Vernon B. Brooks. *The Neural Basis of Motor Control.* Oxford University Press, 1986.

[9] I.E. Brown, S.H. Scott, and G.E. Loeb. Mechanics of feline soleus: Ii design and validation of a mathematical model. *Journal of Muscle Research and Cell Motility*, 17:221–233, 1995.

[10] John Burkardt, Qiang Du, Max Gunzburger, and Hyung-Chun Lee. Reduced order modeling of complex systems. In *Proceedings of the 20th Biennieal Conference on Numerical Analysis*, pages 29–38, 2003.

[11] Vincent C.K. Cheung, Andrea d'Avella, Matthew C. Tresch, and Emilio Bizzi. Central and sensory contributions to the activation and organization of muscle synergies during natural motor behaviors. *Journal of Neuroscience*, 25(27):6419–6434, 2005.

[12] M. Condon and R. Ivanov. Empirical balanced truncation of nonlinear systems. *Journal of Nonlinear Science*, 14:405–414, 2004.

[13] A. d'Avella and E. Bizzi. Low dimensionality of supraspinally induced force fields. *Proc. Natl. Acad. Sci.*, 95:7711–7714, 1998.

[14] A. d'Avella, P. Saltiel, and E. Bizzi. Combinations of muscle synergies in the construction of a natural motor behavior. *Nature Neuroscience*, 6(3):300–308, 2003.

[15] A. d'Avella and M.C. Tresch. Modularity in the motor system: decomposition of muscle patterns as combinations of time-varying synergies. In T.G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems*. MIT Press, 2002.

[16] Joseph A. Doeringer and Neville Hogan. Intermittency of preplanned elbow movements persists in the absence of visual feedback. *Journal Neurophysiology*, 80:1787–1799, 1998.

[17] Timothee Doutriaux. Bio-mechanical characterisation of synergistic muscle control in goal directed movements. Master's thesis, Massachusetts Institute of Technology, 2004.

[18] A.G. Feldman. Functional tuning of the nervous system during control of movement or maintenance of a steady posture. iii. mechanographic analysis of the execution by man of the simplest motor tasks. *Biophsyics*, pages 766–775, 1966.

[19] E.C. Field and P.S.G. Stein. Spinal cord coordination of hindlimb movements in the turtle: intralimb temporal relationships during scratching and swimming. *J. Neurophysiology*, 78:1394–1403, 1997.

[20] T. Flash and N. Hogan. The coordination of arm movements: An experimentally confirmed mathematical model. *The Journal of Neuroscience*, 5(7):1688–1703, July 1985.

[21] Simon C. Gandevia. *Kinesthesia: roles for afferent signals and motor commands*, pages 128–172. Oxford University Press, 1996.

[22] F. Ghorbel and M.W. Spong. Integral manifold of singularly perturbed systems with application to rigid-link flexible-joint multibody systems. *International Journal of Non-linear Mechanics*, 34:135–155, 2000.

[23] Simon F. Giszter and William J. Kargo. Conserved temporal dynamics and vector superposition of primitives in frog wiping reflexes during spontaneous extensor deletions. *Neurocomputing*, 32-33:775–783, 2000.

[24] Simon F. Giszter and William J. Kargo. Modeling of dynamic controls in the frog wiping reflex: Force-field level controls. *Neurocomputing*, 38-40:1239–1247, 2001.

[25] Simon F. Giszter, F.A. Mussa-Ivaldi, and E. Bizzi. Convergent force fields organized in the frog's spinal cord. *J. Neuroscience*, 13:467–491, 1993.

[26] Paul L. Gribble, D. Ostry, V. Sanguineti, and R. Laboissiere. Are complex control signals required for human arm movements? *J. Neurophysiology*, 79:1409–1424, 1998.

[27] Juergen Hahn and Thomas F. Edgar. Reduction of nonlinear models using balancing of empirical gramians and galerkin projections. In *Proceedings of the American Control Conference*, Chicago, Illinois, June 2000.

[28] C. M Harris and D. Wolpert. Signal-dependent noise determines motor planning. *Nature*, 394, Aug 1998.

[29] Corey B. Hart and Simon F. Giszter. Modular premotor drives and unit bursts as primitives for frog motor behaviors. *Journal of Neuroscience*, 24(22):5269–5282, June 2004.

[30] Z. Hasan. A model of spindle afferent response to muscle stretch. *J. of Neurophysiology*, 49:989–1006, 1983.

[31] Neville Hogan. An organizing principle for a class of voluntary movements. *The Journal of Neuroscience*, 4(11):2745–2754, Nov 1984.

[32] Neville Hogan, Joseph Doeringer, and Hermano Igo Krebs. Arm movement control is both continuous and discrete. *Cognitive Studies*, 6(3):254–273, 1999.

[33] M. Hollerbach, Johm and Tamar Flash. Dynamic interactions between limb segments during planar arm movements. *Biological Cybernetics*, 44:67–77, 1982.

[34] Eric R. Kandel, James H. Schwartz, and Thomas M. Jessell. *Principles of Neural Science*. McGraw-Hill Medical, 2005.

[35] W.J. Kargo and S. Giszter. Rapid correction of aimed movements by summation of force-field primitives. *J. of Neurosciences*, 20(1):409–426, 2000.

[36] W.J. Kargo, F. Nelson, and L. Rome. Jumping in frogs: assesing the design of the skeletal system by anatomically realistic modeling and forward dynamic simulation. *J. of Experimental Biology*, 205:1683–1702, 2002.

[37] W.J. Kargo and L. Rome. Functional morphology of proximal hindlimb muscles in the frog rana pipiens. *J. of Experimental Biology*, 205:1987–2004, 2002.

[38] M. Kawato. Trajectory formation in arm movements: Minimization principles and procedures. In H.N. Zelaznik, editor, *Advances in Motor Learning and Control*, pages 225–259. Human Kinetics Publishers, Champaign IL, 1996.

[39] Sanjay Lall, Jerrold E. Marsden, and Sonja Glavaski. A subspace approach to balanced truncation for model reduction of nonlinear control systems. *International Journal on Robust and Nonlinear Control*, 12(5):519–535, 2002.

[40] Daeyeol Lee, Nicholas L. Port, and Apostolos P. Georgopolous. Manual interception of moving targets. *Experimental Brain Research*, 116:421–433, 1997.

[41] M.A. Lemay, J.E. Galagan, N. Hogan, and E. Bizzi. Modulation and vectorial summation of the spinalized frog's hinbdlimb end-point force produced by intraspinal electrical stimulation of the cord. In *IEEE Trans. Neural Syst. Rehabil. Eng.*, number 9, pages 12–23, 2001.

[42] Weiwei Li, Emanuel Todorov, and Xiuchuan Pan. Hierarchichal optimal control of redundant biomechanical systems. In *The 26th Annual International Conference of the IEEE Engineering in Biology and Medicine Society*, 2004.

[43] E.P. Loeb, S.F. Giszter, P. Saltiel, E. Bizzi, and F.A. Mussa-Ivaldi. Output units of motor behavior: An experimental and modeling study. *J. Cog. Neuroscience*, 12:78–97, 2000.

[44] Gordon J. Lutz and Lawrence C. Rome. Built for jumping: The design of the frog muscular system. *Science*, 263:370–372, Jan. 1994. New Series?

[45] Joe McIntyre and Emilio Bizzi. Servo hypothesis for the biological control of movement. *J. of Motor Behavior*, 25(3):193–202, 1993.

[46] Thomas A. McMahon. *Muscles, Reflexes and Locomotion.* Princeton University Press, 1984.

[47] R.C. Miall, D.J. Weir, and J.F. Stein. Intermittency in human manual tracking tasks. *Journal of Motor Behavior*, 25(1):53–63, 1993.

[48] F.A. Mussa-Ivaldi and E. Bizzi. Motor learning through the combination of primitives. *The Royal Society*, 355:1755–1769, 2000.

[49] F.A. Mussa-Ivaldi, S.F. Giszter, and E. Bizzi. Linear combinations of primitives in vertebrate motor control. *Proc. Nat. Academy Sci.*, 91:7534–7538, 1994.

[50] E Nakano, H. Imamizu, R. Osu, Y. Uno, H. Gomi, T. Yoshioka, and M. Kawato. Quantitative examinations of internal representation for arm trajectory planning: Minimum commanded torque change model. *Journal of Nuerophysiology*, 81:2140–2155, 1999.

[51] J. Andrew Newman and P.S. Krishnaprasad. Computing balanced realizations for nonlinear systems. In *Proceedings of the 14th International Symposium on Mathematical Theory of Networks and Systems*, Peripignan, France, June 2000.

[52] Arthur Prochazka. *Proprioceptive feedback and movement regulation*, pages 89–127. Oxford University Press, 1996.

[53] Michal Jerzy Rewienski. *A trajectory Peicewise-Linear Approach to Model Order Reduction of Nonlinear Dynamical Systems.* PhD dissertation, Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, 1998.

[54] D.A. Robinson. Is the oculomotor system a cartoon of motor control? *Progress in Brain Research*, 64, 1986.

[55] C.W. Rowley. Model reduction for fluids, using balanced proper orthogonal decomposition. *Int. J. on Bifurcation and Chaos*, Mar 2005.

[56] P. Saltiel, M. Tresch, and E. Bizzi. Spinal cord modular organization and rhythm generation: An nmda iontophoretic study in the frog. *J. Neurophysiol*, 80:2323–2339, 1998.

[57] P. Saltiel, K. Wyler-Duda, A. d'Avella, M.C. Tresch, and E. Bizzi. Muscle synergies encoded within the spinal cord: evidence from focal intraspinal nmda iontophoresis in the frog. *J. Neurophysiology*, 85:605–619, 2001.

[58] Terence D. Sanger. Optimal unsupervised motor learning for dimensionality reduction of nonlinear control systems. *IEEE Transactions on Neural Networks*, 5(6):965–973, 1994.

[59] Terence D. Sanger. Human arm movements described by a low-dimensional superposition of principal components. *Journal of Neuroscience*, 20(3):1066–1072, 2000.

[60] Jacquelien M.A. Scherpen. Balancing for nonlinear systems. *Proceedings European Control Conference*, pages 1838–1843, 1993.

[61] Andrew B. Slifkin, David E. Vaillancourt, and Karl M. Newell. Intermittency in the control of continuous force production. *Journal of Neurophysiology*, 84:1708–1718, 2000.

[62] Jean-Jacques E. Slotine, Harry Asada, and Haruhiko Asada. *Robot Analysis and Control.* John Wiley & Sons, Inc., 1992.

[63] Jean-Jacques E. Slotine and Weiping Li. *Applied Nonlinear Control.* Prentice Hall, 1991.

[64] Mark W. Spong and M. Vidyasagar. *Robot Dynamics and Control.* John Wiley & Sons, Inc., 1989.

[65] E. Todorov and Z. Ghahramani. Unsupervised learning of sensory-motor primitives. In *In Proceedings of the 25th Annual International Conference of the IEEE Engineering in Biology and Medicine Society*, 2003.

[66] E. Todorov and M. I. Jordan. Optimal feedback controls as a theory of motor coordination. *Nature Neuroscience*, 5(11), Nov 2002.

[67] M. Tresch, P. Saltiel, and E. Bizzi. The construction of movement by the spinal cord. *Nature Neruoscience*, 2:162–167, 1999.

[68] M.C. Tresch, P. Saltiel, A. d'Avella, and E. Bizzi. Coordination and localization in spinal motor systems. *Brain Res. Rev.*, 1, 2002.

[69] Y. Uno, M. Kawato, and R. Suzuki. Formation and controlf of optimal trajectory in human multijoint arm movement. *Biological Cybernetics*, 61:89–101, 1989.

[70] J. M. Winters and L. Stark. Muscle models: What is gained and what is lost by varying model complexity. *Biol. Cybern.*, 55:403–420, 1987.

[71] J.M. Winters. *Hill-Based Muscle Models: A Systems Engineering Perspective*, pages 69–93. Springer-Verlag, 1990.

[72] George I. Zahalak. *Modeling Muscle Mechanics (and Energetics)*, pages 1–23. Springer-Verlag, 1990.

[73] F.E. Zajac. Muscle and tendon: properties, models, scaling, and application to biomechanics and motor control. *Critical Reviews in Biomedical Engineering*, 17(4), 1989.