

DYNAMIC DEXTEROUS MANIPULATION: BENEFITS
OF THE EDGE OF INSTABILITY IN EXPLORING
COMPLEX DYNAMICAL BEHAVIOR

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Madhusudhan Venkadesan

January 2007

© 2007 Madhusudhan Venkadesan

ALL RIGHTS RESERVED

DYNAMIC DEXTEROUS MANIPULATION: BENEFITS OF THE EDGE OF INSTABILITY IN EXPLORING COMPLEX DYNAMICAL BEHAVIOR

Madhusudhan Venkadesan, Ph.D.
Cornell University 2007

Behaviors such as fine manipulation in humans are quintessentially nonlinear, dynamic and complex. Exploring, quantifying and characterizing such complex dynamic systems is essential to expand our understanding of biological systems and to design artificial systems that can match the versatility and performance of their biological counterparts. However, we currently lack means to quantify or model nonlinear dynamical behavior for most such complex systems.

Here we present a paradigm that exploits ubiquitous low-dimensional phenomena at the edge of instability in arbitrarily complex nonlinear dynamical systems to quantify and understand their behavior. Specifically, we examine dynamic manipulation at the edge of instability by asking subjects to compress a slender spring using their thumbpad almost to a point of instability. The spatiotemporal dynamics of a one-dimensional nonlinear dynamical system based on bifurcation theory and dynamics of spring buckling was indistinguishable from experimental measurement at the edge of instability.

We then use this model to answer a neurophysiologically important question: why do we normally handle objects effortlessly without looking at them, but rely heavily on vision if our fingers were numb? We extend our simple model to incorporate feedback from noisy and time-delayed multiple sensory channels — namely, thumbpad sensors, non-digital sensors, and vision. Using numerical optimization, we find that the selective use of vision depending on finger numbness is a cumulative effect of noise and time-delays on task-optimal multisensory integration strategies.

Next we provide preliminary evidence that an upcoming treatment for thumb osteoarthritis — intra articular Hyaluronan injection — causes self-reported improvement in patients primarily due to pain relief and not due to any improvement in their dynamic manipulation ability. We are able to distinguish between the contribution of pain-relief and innate sensorimotor control ability because dynamic manipulation ability at low forces is independent of pain.

We conclude with a brief presentation of various directions for short and long term goals that build upon the foundation established by this thesis for exploring dynamic sensorimotor behavior. We observe that there are several clinical, neurophysiological, and control theoretical benefits of studying complex dynamical systems at the edge of instability. The ubiquitous nature of well-classified low-dimensional phenomena at the edge of instability in arbitrarily complex nonlinear dynamical systems makes the techniques developed in this thesis widely applicable to biological or artificial systems alike.

BIOGRAPHICAL SKETCH

Madhusudhan Venkadesan was born in Madras (currently known as Chennai), Tamilnadu, on September 13, 1978. He spent most of the first eighteen years of his life in Kalpakkam, Tamilnadu, where he studied at the central government school, Kendriya Vidyalaya II and finished high school in 1996. He joined the Indian Institute of Technology (IIT), Madras the same year for his undergraduate studies, specializing in Mechanical Engineering and in 2000, received the degree of Bachelor of Technology. In August 2000, he was admitted to the graduate program at Cornell University, Ithaca, New York, in the Sibley School of Mechanical & Aerospace Engineering and was awarded the degree of Master of Science in 2003. He then continued as a candidate in the doctoral program at Cornell University in the Sibley School of Mechanical & Aerospace Engineering.

In memory of my grandmother – *Train Paati*

ACKNOWLEDGEMENTS

This dissertation has been an educational experience, with enthralling moments of discovery, inspiring glimpses of new horizons, some dead-ends and most of all, an exciting and enjoyable journey. Right through this journey, the unquestioning support, encouragement, inspiration and relentless hard work of my companion Soumya were instrumental in making it a truly rewarding experience.

The chair of my special committee Francisco Valero-Cuevas, was an advisor to me in every sense of the word. A lunch I had with Francisco nearly six years ago launched me on the path I traveled at Cornell. I thank him for leading me by the hand when needed and letting me fly independently when I sought it. It was undoubtedly his inspiration as an advisor, company as a friend, and vision as a philosopher that have made my graduate study immeasurably enriching. My committee member John Guckenheimer, opened my eyes to the stunning world of nonlinear dynamical systems. The benefit of studying the ‘edge of instability’ under John’s guidance was indeed the most gratifying experience! I also thank my committee member Raffaello D’Andrea whose insights have helped improve this dissertation.

Negotiating my path through Cornell and a new country would have been a nightmare if not impossible without the always timely and invaluable help provided by Judy Thoroughman. I thank her also for ‘insider’ information about Ithaca, which made the stay all the more enjoyable for Soumya and I.

My mother Bhooma and my father Venkadesan have been pivotal in my life by motivating independence and creativity of thought. I thank them for their inspiration and unwavering support throughout.

For the frequent, captivating and fruitful discussions that often ran well into the night, I am indebted to Mahesh, Sridhar (Machi), Manoj and Ramu. I am also thankful for some comments by Andy Ruina and Walter Herzog early in my Ph.D. that proved instrumental in helping orient my research.

Large parts of my research would not have been possible without the support of several others. Stephanie Roach, Robert McNamara, Eric Samorodnitsky and Alex Deyle were critical to the experiments carried out at Ithaca. Lisa Mandl, Sherry Backus, Lily Ariola, Allison Swigart, and Elizabeth Gluck at the Hospital for Special Surgery, New York, were instrumental to the clinical portion of my work. I thank Russell Lloyd and Karen Grace-Martin for advising me on the statistical analyses I had to perform.

Finally, it goes without saying that Soumya and I were fortunate to have had the company of an enthusiastic group of friends and coworkers – Veronica, Kevin, Vic, Tuhin, Nishu and Pritam.

My research work was partly supported by the Whitaker Foundation, US National Science Foundation CAREER Award 0237258, US National Institutes of Health grants R21-HD048566, R01-AR050520, R01-AR052345 and also by the Sage Fellowship award, Engineering Supplementary Fellowship Award, and the Sibley School of Mechanical & Aerospace Engineering at Cornell University.

TABLE OF CONTENTS

1	Introduction	1
1.1	Theories of sensorimotor control	1
1.1.1	Predictive aspects of sensorimotor control	2
1.1.2	Optimality in sensorimotor control	3
1.1.3	Bayesian inference	8
1.2	Sensorimotor control of manual “dexterity”	13
1.2.1	Static pinch and tactile dominance	13
1.3	Nonlinear dynamics in biological sciences	15
1.3.1	Benefits of the edge of instability	18
1.4	What this thesis is about	18
1.5	What this thesis is not about	20
1.6	Why study the hand?	21
1.7	Layout of this thesis	22
2	Dynamic manipulation at the edge of instability — Experimental design and mathematical modeling	23
2.1	Introduction	23
2.1.1	Background	23
2.1.2	Past work – The “Strength-Dexterity Test”	24
2.1.3	“Manipulating the edge of instability”	27
2.1.4	Hypotheses	30
2.2	Experimental methods	31
2.2.1	Experimental setup	32
2.2.2	Experimental protocol	33
2.2.3	Metric of performance: F_s	33
2.2.4	Experiments to test for effect of training	34
2.2.5	Safety-margin	34
2.2.6	Analysis of endcap rotation	34
2.2.7	Statistical analyses	35
2.3	Mathematical model of manipulation at the edge of instability . . .	37
2.3.1	Why a subcritical pitchfork bifurcation model?	37
2.3.2	Endcap rotational buckling as a secondary bifurcation	44
2.3.3	Subcritical pitchfork bifurcation model	44
2.3.4	Parameter values for the low-order model	45
2.4	Results	47
2.4.1	Dimensional collapse at the edge of instability	47
2.4.2	Strength independent limit of sensorimotor performance . .	50
2.4.3	Lack of measurable safety-margin	51
2.4.4	No measurable effect of training	51
2.5	Discussion	52
2.6	Conclusions	54

3	Task optimal multisensory integration during dynamic manipulation	55
3.1	Introduction	55
3.1.1	Background – recapitulation of Bayesian inference	56
3.1.2	Why bother about time-delays?	56
3.1.3	Hypotheses	57
3.2	Computational model	57
3.2.1	Separating “controller” from “plant”	58
3.2.2	Multisensory feedback and linearly weighted sensor fusion	59
3.2.3	Noise and time-delay estimates	61
3.2.4	Sensory variance estimates	62
3.2.5	Sensory time-delay estimates	62
3.2.6	Metrics for stability: “Survival times” and “Success rates”	63
3.2.7	Numerical optimization: sensory weights that maximize F_s	65
3.2.8	Sensory weights that minimize the effects of noise alone	65
3.2.9	<i>Best-fit</i> vs. <i>Baseline</i> simulations	66
3.3	Experimental methods	66
3.3.1	Sensory occlusion	66
3.4	Results	67
3.4.1	Experimental results of sensory occlusion show selective use of vision	67
3.4.2	Combined effects of noise and time-delays drive the selective use of vision	68
3.5	Discussion	72
3.6	Conclusions	75
3.6.1	Do we better understand the effects of growing old or cold?	75
4	Clinical implications	77
4.1	Introduction	77
4.1.1	Why does hand function improve in CMC OA patients after treatment with intra-articular HA injection?	78
4.1.2	Intra-articular HA – a potential treatment for CMC OA	78
4.1.3	Existing outcome measures do not quantify impairment and recovery of dynamic manipulation	79
4.2	Hypotheses	80
4.3	Methods	81
4.4	Statistical Analyses	82
4.4.1	Testing Hypothesis 1: Change in metrics of hand function after treatment	82
4.4.2	Testing Hypothesis 2: Multiple linear regression	82
4.4.3	Stepwise linear regression	83
4.4.4	Pairwise correlation analysis	85
4.5	Results and Discussion	85

4.5.1	DASH and VAS are the only metrics that showed improvement after treatment	87
4.5.2	Patients with better initial sensorimotor control ability benefited more from pain-relief	87
4.5.3	Only pain relief and baseline F_s could explain self-reported improvement	88
4.5.4	Static pinch strength is not informative of any aspect of functional improvement	88
4.6	Conclusions	88
5	Future Work and Conclusions	90
5.1	Refinement of the normal form model	90
5.1.1	Inclusion of “imperfections”	90
5.1.2	Trial-specific parameter estimation	93
5.2	Proportional-derivative feedback controllers	93
5.2.1	Proportional-derivative control model	95
5.2.2	Parameter values for proportional-derivative control model	96
5.2.3	Results of <i>Best-fit</i> simulations	97
5.3	“Computational” time-delay — artifact or real?	98
5.4	Carpal tunnel syndrome vs. multiple sclerosis	99
5.5	Optimal feedback control	101
5.6	Hierarchical nearly-optimal feedback control	103
5.7	Conclusion	105

LIST OF TABLES

2.1	Statistical post-hoc planned comparisons of F_s (sustained load) and $F_{\max}$ (maximal load) for trials when the spring slipped vs. when the spring did not slip.	51
3.1	Task-optimal and noise minimizing sensory weights for all simulation results.	69
4.1	This table gives a list of potential factors that could explain any improvement in DASH scores after treatment with intra-articular HA injection.	84
4.2	Results of pairwise Spearman correlation for all the predictors from the master-set in the stepwise regression model.	86
4.3	DASH score and VAS for pain are the only metrics that showed a change improvement after 26 weeks.	87
5.1	This table presents results of the <i>Best-fit</i> simulations using a PD controller. The results reveal that a task-optimal strategy explains multisensory integration during dynamic manipulation.	97
5.2	This table presents testable predictions of our model to determine whether “computational” time-delays really exist for non-digital sensors.	99
5.3	This table summarizes testable model predictions for the functional consequences of MS vs. CTS.	101

LIST OF FIGURES

2.1	The discretized strength-dexterity plane. Adapted from Valero-Cuevas et al. (2003b) with permission.	25
2.2	Effective length for a buckling column.	26
2.3	Experimental setup for manipulation at the edge of instability. . .	29
2.4	Schematic showing a Hopf (left) and pitchfork (right) bifurcation for a buckling spring.	38
2.5	The bifurcation diagram for the normal form of a supercritical pitchfork bifurcation.	41
2.6	The bifurcation diagram for the normal form of a subcritical pitchfork bifurcation.	42
2.7	An illustration of spring buckling under a dead-load that depicts a non-ideal guided end-condition resulting in both lateral and rotational buckling.	43
2.8	Bifurcation diagram for the normal form used to model manipulation at the edge of instability.	46
2.9	Eccentricity of endcap rotation data.	48
2.10	Endcap rotation along first principal component agrees with model prediction.	49
2.11	Endcap rotation along second principal component showed no dependence on F_s	50
2.12	Box-plot showing the effect of training.	52
3.1	Block diagram of multisensory feedback control model.	60
3.2	This box plot summarizes the experimental results of sensory occlusion.	68
3.3	This figure shows the effect of time-delays in addition to noise on task-optimal multisensory integration.	70
3.4	Results of the global optimization using the <i>Baseline</i> simulation. .	73
3.5	Results of the global optimization using the <i>Best-fit</i> simulation. . .	73
5.1	The bifurcation diagram of a subcritical pitchfork bifurcation with the inclusion of a small imperfection.	92

Chapter 1

Introduction

In this work, an attempt will be made to quantify and reveal some aspects of the neuromuscular control of manual dexterity. More specifically, this work develops a mathematical-experimental method to quantify and explore the complex, noisy and nonlinear system of dynamic dexterous manipulation. Results from dynamical systems theory and the theory of elastic buckling are used to develop a method to experimentally quantify and mathematically model dynamic manipulation. We then use this method to reveal hitherto unexplored effects of time-delays on how the nervous system combines signals from (or not) and uses signals from multiple sensory channels¹. Finally, a clinical application of this paradigm is presented. We demonstrate potential use of our experimental method to predict treatment outcome for patients with osteoarthritis of the thumb — a disease with devastating effects that is very common, if not inevitable in the older population.

A brief overview of some necessary background material is included below to help read this thesis. This overview is by no means exhaustive or detailed. It is meant only to provide representative examples of ongoing research in various fields relevant to this work. The overview also serves the dual purpose of motivating the work reported in this thesis.

1.1 Theories of sensorimotor control

For the benefit of readers who are not familiar with contemporary theories of how the nervous system controls the body and its interaction with the environment, I will first provide a brief overview of theories relevant to the present work. Although not all of these theories have been established beyond doubt, we can nonetheless garner some useful facts and observations about the nervous system from the experimental, mathematical and computational studies underlying these theories.

It is important to first define the term *sensorimotor control*. By sensorimotor control, I refer to *sensory* (e.g. vision), *neural* (e.g. spinal cord) and *musculoskeletal* (muscles, tendons, ligaments, bone, etc.) components of the body acting as a fused dynamical system that is *controlled* by inherent properties or variables such as neural signals, muscle contraction, etc. to yield useful interactions with the environment. The mechanisms that you use to hold and read this thesis is an example of sensorimotor control. Now I present some theories of sensorimotor control that are relevant to the work presented here.

¹The process of combining and using multiple sensory channels is often called *multisensory integration* or *sensor fusion*.

1.1.1 Predictive aspects of sensorimotor control

The ability of the nervous system to predictively control the body, i.e., control the body even when it has no sensory feedback, is essential for activities of daily living. When you walk with a full glass of water without spilling it, you move your arm in sync with your foot-falls by predictively anticipating your foot-falls. You need to predictively anticipate and not reactively respond to footfalls because, time-delays in your sensors and in your nervous system make it impossible to move your arm in sync with footfalls by relying solely on sensory feedback. Contemporary theories of the predictive capability of the nervous system provide a systematic mathematical formalization that has led to a much clearer picture of *how* the nervous system uses predictive capabilities in everyday tasks.

One key aspect of this predictive capability is known as *motor prediction* (Wolpert and Flanagan, 2001). Motor prediction refers to predicting the behavior of systems that are directly influenced by the nervous system's commands to muscles in our body (motor commands). For example, motor prediction could mean predicting how your leg will swing during walking in response to a motor command and also predicting what will happen to the bag that is swinging from your shoulder as your body rises up and down during walking. One dominant thought in the field of motor control says that to predict the consequences of the motor command on our body or on the environment we interact with, there must be some mechanism in the nervous system to simulate the dynamic behavior of our body (Wolpert et al., 1995, 1998) and the environment (Dingwell et al., 2002, 2004). This hypothesized mechanism is commonly referred to as the *internal forward model* — *internal* because it has to be physically inside the nervous system and *forward* because it is used to predict the causal relationship between motor action and its consequence. One of the earliest systematic investigations into the existence of an internal forward model were carried out by a 19th century German scientist and philosopher, Hermann Ludwig Ferdinand von Helmholtz when he tried to understand how the brain identifies the spatial location of visual objects (von Helmholtz, 1867; von Helmholtz et al., 1971). The spatial location of an object relative to the head depends on where the object lies on the retina and on the orientation of the eye in its socket. Hermann von Helmholtz made an audacious (at that time) and ingenious proposition that the brain does not use any sensors to measure the orientation of the eye inside the socket, rather it uses the motor commands sent to the eye muscles to predict the eye orientation (von Helmholtz, 1867; von Helmholtz et al., 1971, original and translation, respectively). He used a simple experiment to demonstrate this idea — when the eye is moved without any motor command to the muscles (close one eye and gently press the open eye through the eyelid), the visual scenery seems to shift because the location of objects on the retina changes, but the predicted eye orientation is not updated. The corresponding neural mechanism by which the nervous system informs itself of outgoing motor commands is called the *effference copy*, the existence of which was further reinforced by more recent experiments (for instance, Sperry, 1950).

The flip-side of a forward model is the ability of the nervous system to take a desired outcome (say, pointing to an object) and determine the appropriate motor command. This is equivalent to inverting the forward model and hence referred to as the *inverse internal model*. The inverse internal model is considered an important part of sensorimotor control (Wolpert et al., 1995, 1998, 2001; Wolpert and Ghahramani, 2000; Kawato, 1999; Jordan, 1996). However, as we shall soon see (in Section 1.1.2), from the point of view of optimal feedback control, this nomenclature of an inverse model becomes somewhat redundant.

It should be noted that there is still disagreement and debate about the existence of (forward and inverse) internal models (Ostry and Feldman, 2003; Feldman and Latash, 2005). The primary objection is that there exists an alternate to internal models² — the nervous system could tune the stiffness and damping properties of the muscles appropriately so that all control is achieved solely via the viscoelastic mechanical properties of the body and not via direct control of muscle forces. This is based on the equilibrium point hypothesis proposed by Feldman (1966a,b), Bizzi et al. (1984, 1992), Hogan (1984a,b), etc. However, a vast amount of recent evidence from psychophysical (Shadmehr and Mussa-Ivaldi, 1994; Flanagan and Wing, 1997; Kawato, 1999; Flanagan and Lolley, 2001) and neurophysiological (Gribble and Scott, 2002; Li et al., 2001) studies strongly support the existence of internal models.

This discussion takes us naturally to theories about how the nervous system puts its predictive capability to use. A thorough review of various theories that attempt to explain sensorimotor control would fill several books and is not the purpose of this overview. It suffices to say that theories of sensorimotor control have evolved from the purported existence of a ‘vital force’ that drives all living things³ to optimality as a principle of sensorimotor control (Jordan, 1996; Todorov and Jordan, 2002; Todorov, 2004).

1.1.2 Optimality in sensorimotor control

The subtle, yet important connection between the predictive capability of the nervous system that we surveyed above and sensorimotor control appears to be in the form of “optimality laws”. Optimality in the context of predicting or explaining sensorimotor coordination in biological systems, refers to generating detailed predictions of how muscles are activated by the nervous system using a low-dimensional and well-defined goal that is specified in the form of a ‘cost-

²An internal model is defined as an internal neural mechanism that can simulate the dynamics of the body, the world and their interaction.

³The ‘vital force’ or ‘vitalism’ is the metaphysical notion that living things possess an inner “force” that makes them different from non-living things. Hermann von Helmholtz, the German philosopher and scientist was a pioneer of the viewpoint that the study of physiology must be founded on laws of physics and chemistry (Mackey and Santillan, 2004) and not on vital forces.

function’⁴ (in fact, almost always a scalar function). An example application of optimality principle is in predicting how humans walk by specifying only that we want to minimize the metabolic energy that is consumed as the fitness-function (Srinivasan and Ruina, 2006; Srinivasan, 2006). In order to use optimization to predict how the nervous system controls the body, it is also necessary to constrain allowable behaviors to a meaningful, yet not over-constrained set. For example, in human locomotion one could (mathematically) specify that our bipedal computer model can only produce periodic gaits and of the infinite possibilities of periodic gaits, find the gait that minimizes energy consumption. In other words, optimization provides a mathematical method for finding the best possible way of doing something given constraints on what can be done and a measure of fitness or cost, i.e. a scalar metric to compare two choices and qualify one as better or worse than the other. There are other methods that use non-scalar measures of fitness or cost. However, they are often plagued by uncertainties about how to relatively weight independent dimensions or goals since they try to achieve multiple goals. The uncertainties arise because in general, there is no unambiguous way to call one vector “bigger” than another when the each of independent axes of the vectors could be scaled differently. In all, we will henceforth focus on optimization techniques that rely on scalar cost-functions.

Of the mind-bogglingly large number of scientific attempts at explaining sensorimotor coordination, optimality holds the appeal of being the most parsimonious explanation. It could be argued that optimality principles well approximate both evolutionary pressures and task-demands in biological systems. A more thorough discussion of optimality in biological systems in general can be found in Smith (1978) and Alexander (1996, 2001). Optimality has been successfully applied in the past to biomechanical systems like maximal static force production (Valero-Cuevas et al., 1998), locomotion (Srinivasan and Ruina, 2006; Srinivasan, 2006, and references therein), etc. Valero-Cuevas et al. (1998) used linear programming (a linear convex optimization method) to predict muscle coordination for maximizing fingertip force in a specified direction with well-defined constraints on fingertip forces/torques. Srinivasan and Ruina (2006) used energy minimization as the goal of the optimization to explain transitions from ‘walking’ to ‘running’ in bipedal animals.

The two examples cited above use optimization to predict average behavior. They translate a desired goal into detailed specifications for how best to coordinate muscles / actuators on average. This form of application of optimization to predict control action is referred to as *open-loop optimal control*. Undoubtedly, given the presence of unavoidable environmental and internal disturbances (or noise), some form of feedback is essential for single trial performance. In open-loop optimal control, there is no unambiguous specification of how to implement

⁴The cost-function is also alternately referred to as ‘objective-function’ or ‘fitness-function’. In its typical usage, the cost-function is a quantity that we desire to minimize, while the fitness-function is a quantity that we desire to maximize.

the feedback portion of the controller. The superset of open-loop optimal control is then *closed-loop optimal control*, that optimizes not only the open-loop portion (from goal to average control action), but also the response to trial-to-trial perturbations according to some high-level goal or cost-function. So, closed-loop optimal control predicts both average behavior and trial-to-trial variability. We present below a more detailed overview of open-loop vs. closed-loop optimal control in the context of sensorimotor control. It is important not to confuse this terminology of open-loop vs. closed-loop with the absence or presence of feedback — both flavors of optimal control necessarily entail feedback in order to be implemented in real-world systems. ‘Open’ or ‘closed’ loop in optimal control refer only to the optimization procedure — whether only the desired goal to control action is optimized (open-loop) or the feedback loop is incorporated (closed-loop) into the optimization routine.

Open-loop optimal control

In the context of sensorimotor control, a large class of optimal control models predict average behavior (over repeated measurements of the same animal or measurements across several animals) by optimizing some preset cost-function in the same spirit as energy optimality models of locomotion. These models can be lumped under the category of *open-loop optimal control* models. Cost functions other than energy minimization were proposed because energy minimization failed to account for behaviors such as arm movements (Nelson, 1983). Several authors proposed that the nervous system might be trying to maximize “smoothness” of limb movements i.e., jerk (the time-derivative of acceleration) (minimum-jerk models such as Hogan, 1984b; Flash and Hogan, 1985; Smeets and Brenner, 1999) or the nervous system might be trying to minimize the effects of sensorimotor noise⁵ (minimum-variance models such as Harris and Wolpert, 1998). Irrespective of the debate whether these fitness-functions seem reasonable or not, the point is that control actions needed to achieve a desired task goal can be computed using some form of optimality law⁶. However, it is important to note that there is one important shortcoming of open-loop optimal control models. They focus only on average behavior and tell us nothing about how the nervous system responds to unexpected disturbances. For example, if you picked up an empty can of soda, but

⁵Noise is a general term used to describe variability in a system that cannot be captured by a deterministic model of a specified complexity. The source of ‘noise’ in biological and other physical systems is in general extremely hard to identify. For an interesting discussion of whether neuronal variability is a hindrance (noise) or useful information, see Stein et al. (2005).

⁶This is a reason why the nomenclature of an inverse internal model is rendered superfluous by optimal sensorimotor control. An inverse internal model becomes just another name for the mechanism (whatever it may be) that computes the control action needed to achieve the desired goal.

expected it to be full, your arm would move “unnaturally” before your muscular coordination adjusts appropriately to restore your arm motion to something more “natural”. Does the nervous system just use some open-loop strategy to pick up the can and try to update it online using some form of simple feedback control, thus acting like a servo to the open-loop strategy? Or does it have the capability to choose feedback control laws that are also optimal in some sense? Hence, although both open-loop and closed-loop optimal control utilize sensory feedback, the distinction between them is whether the use of sensory feedback is also optimal (i.e., closed-loop optimal control) or if sensory feedback is used in a predetermined, non-optimal, servo-like manner (i.e., open-loop optimal control).

Closed-loop optimal control

Closed-loop optimal control is a superset of open-loop optimal control. It optimizes (i) the preplanned sensorimotor control strategy for a given goal (i.e. determines optimal open-loop behavior) and, (ii) the feedback control strategy (i.e. online response to unexpected perturbations). In other words, it does what it takes to achieve the specified goal and lets the dynamics arising out of the interaction between the body and the world determine the optimal feedback strategy. A very thorough review and discussion of optimal feedback control (synonymous to closed-loop optimal control) and open-loop optimal control in the context of sensorimotor systems can be found in Todorov (2004). The main feature of optimal feedback control models that is relevant to the work presented in this thesis is the necessity for existence of optimal state estimators. Optimal state estimators are mechanisms that can use sensors with noise and time-delays to make an optimal guess⁷ of the true state of the system being sensed. For a broad range of sensory noise characteristics (i.e., even for non-Gaussian, colored noise) and definitions of optimality, optimal estimators are Bayesian⁸ in nature (Clark and Yuille, 1990), i.e. they take into account the measured sensory signals, the dynamics of the system being controlled, the context of the task being performed⁹, current and past control signals, and combines them by weighting them according to their relative reliability. Several studies have experimentally found remarkable agreement between experiments and optimal feedback control models that use Bayesian estimators. For example, if optimality for the state estimate is defined in terms of least-squared error, then maximum likelihood estimation, i.e., Bayesian estimation

⁷The guess is optimal in the sense of minimizing the error in the estimate.

⁸The term ‘Bayesian’ refers to Bayes’ theorem that provides a way to calculate conditional probability, i.e. the probability of event B, given that event A has happened.

⁹It is not clear how exactly to define context. However, a loose definition could be based on the time-scales. Phenomena with such slow time-scales compared to other dynamics in the system so that they can be considered as static (time-invariant) can be grouped together to define ‘context’.

is the optimal thing to do. For the purpose of this thesis and for typical scenarios encountered in biological settings, we will accept the hypothesis that optimal estimators are Bayesian to be a true statement. In summary, the important message to remember is that optimal feedback control, Bayesian estimation and internal forward models go hand-in-hand.

Here is a brief overview of some studies to provide a sample of the evidence in support of optimal feedback control in the nervous system. Studies of balance control during quiet standing provide an informative comparison between optimal control models and experiments. Using linearized models of inverted pendula, several studies have demonstrated how the measured response of joint angle trajectories to small perturbations closely resembles optimal feedback control models that try to maintain balance while minimizing control effort (Loeb et al., 1990; Kuo, 1995, 2005; van der Kooij et al., 1999). The most convincing evidence in support of optimal feedback control in the nervous system (Todorov and Jordan, 2002, 2003) arise from attempts at explaining key ideas in sensorimotor control that were originated by the pioneering work of Bernstein (1967). Bernstein noted that almost all tasks we perform everyday require resolving redundancy, i.e. the same goal can be achieved in many ways (usually infinite), thus leading to the problem of how to choose one out of the many possibilities. For example, you could touch your nose using (infinitely) many trajectories of your arm, most of which appear ‘unnatural’ to you. What made you decide that a specific trajectory was ‘natural’ and others were ‘unnatural’? Although open-loop optimal control models specify how best to behave on average based on some well defined task goal, they do not specify how to use feedback control to implement the desired behavior in reality where there is noise from the surrounding and from within the neuromuscular system. In the presence of noise, even if the average behavior was identical, different feedback control strategies typically produce different behavior (Loeb et al., 1999; Todorov, 2004). Hence, understanding how noise affects behavior served as an ideal test bed for theories of feedback control in the nervous system. Specifically, one important prediction of optimal feedback control models that try to achieve the task goal using the least amount of effort is the emergence of ‘minimal intervention’, i.e. optimal feedback control models do not try to control any departure from the desired average behavior unless the departure will affect the task goal (Todorov and Jordan, 2002; Todorov, 2004). There are several phenomena observed in sensorimotor control in humans that naturally emerge from the concept of minimal intervention thus providing support for the utility of optimal feedback control in understanding human sensorimotor behavior:

1. **Selective response to external perturbations.** The nervous system responds to external perturbations only if the perturbation will affect the desired goal (Bernstein, 1967; Cole and Abbs, 1988; Gracco and Abbs, 1985; Burdet et al., 2001, 2006). This is exactly what an optimal feedback controller (and not others) would do, since perturbations that do not affect the goal, would not affect the cost-function, but extra control effort does affect

the cost-function.

2. **Structure in variability — Uncontrolled manifold.** The variability in most tasks have been repeatedly observed to be elongated along those directions that have no effect on the desired goal when the variability is represented in the space of variables that describe the task. This has been named the ‘uncontrolled manifold’ (Scholz and Schöner, 1999; Scholz et al., 2000). Once again it is clear why minimal intervention would naturally lead to an uncontrolled manifold. Because, variability due to internal perturbations that has no effect on the task outcome is best left uncontrolled, but it is beneficial to control variability in directions that affect task outcome, the observed structure in variability is a natural outcome of optimal feedback control.
3. **Muscle synergies.** Applying statistical techniques like principal components analysis to muscle EMGs often show that several muscles act in a concerted way as if they were controlled by the nervous system to represent only the relevant features of the task, leading to the idea of ‘motor synergies’ (Bernstein, 1967; Cheung et al., 2005; d’Avella and Bizzi, 2005; Ivanenko et al., 2003). In optimal feedback control models, the absence of control in certain directions implies that control signals are correlated or some control signals are completely absent. Statistical techniques such as principal components analysis find that very few principal components suffice to account for the variance in a correlated data set. The muscle synergies that are found in experimental data rely on statistical techniques like principal components analysis. Hence, the frequent observation of synergies in muscle coordination is once again an expected outcome of optimal feedback control.

1.1.3 Bayesian inference

The final piece of our overview of theories of optimal sensorimotor control is Bayesian estimation. The most commonly used name in the literature of motor control for optimal state estimation using Bayesian techniques is ‘Bayesian inference’. Studies on state estimation in humans have repeatedly found that Bayesian inference models resemble human behavior in perception (Knill and Pouget, 2004; Ernst and Bulthoff, 2004; Ernst and Banks, 2002a,b; Wolpert et al., 1995), sensorimotor learning (Körding and Wolpert, 2004; Körding et al., 2004), motor planning (Sober and Sabes, 2003, 2005; Wolpert et al., 1995), posture balance (Kuo, 1995, 2005; van der Kooij et al., 1999) and even in tasks where the same sensory channel (vision) perceives multiple signals (position and velocity of the fingertip; Saunders and Knill, 2003, 2004). Here, I will go over some of the evidence presented in support of Bayesian inference in the nervous system and more importantly, examine the specifics of the types of tasks that were used in these studies. We will see how the type of tasks restrict the results from these studies to reveal only the

effect of sensorimotor noise and not the effect of time-delays in the nervous system on behavior. The strength and weakness of these past studies both arise from the simple tasks they use, like pointing to a target, where the goal of the task is always static, or tasks where assumptions of linearity were justifiable. It is a strength of these studies, because it is astounding how much about the nervous system can be revealed with such simple tasks. It is a weakness because the static or linear nature of these tasks restricts them to conclusions that do not truly reflect everyday tasks like object manipulation that are typically nonlinear and dynamic (Bicchi, 2000; Valero-Cuevas, 2005). Past studies of multisensory integration come under one of three groups listed below.

Static task goal

Several previous studies have examined how humans combine multiple sensory signals. Most of them use a static task goal of moving the hand to a given target as accurately as possible to study sensory estimation (Ernst, 2004; Ernst and Banks, 2002a; Ernst and Bulthoff, 2004; Knill and Pouget, 2004; Körding et al., 2004; Körding and Wolpert, 2004; Saunders and Knill, 2004; Sober and Sabes, 2003, 2005; van Beers et al., 2002, 2004; Wolpert et al., 1995). In the language of optimal control models, the goal is to minimize errors in pointing to a static target. The errors during behavior like pointing to a target arise from noise (uncertainty) in: (i) sensory signals that inform the nervous system of the target’s location and the posture of the arm, (ii) the motor system, i.e., apparently random fluctuations in muscle force that could affect the accuracy of task performance and, (iii) neural computations — although noise and variability in sensorimotor control is typically attributed to either sensors or muscles, it is possible that a significant amount of noise might arise from neural computation. The review by van Beers et al. (2002, page 7) provides a good list of studies that try to elucidate the source of noise at every level in the sensorimotor system — from sensors and muscles to neuronal networks.

Ernst and Banks (2002a) used a study where human subjects had to estimate which out of two blocks held between their index finger and thumb was taller. The perceptual variance of volunteers in the experiment was estimated using a virtual reality setup to provide only one of visual or haptic feedback. Subjects then performed trials where both visual and haptic feedback were simultaneously available, but visual feedback was randomly perturbed, so that there was a random discrepancy between visual and haptic feedback. A model based on Bayesian inference found a rate of successful discrimination that was indistinguishable from experimental data. This sort of experiment is a recurring theme, where the variance of each sensory channel is first obtained experimentally and when all sensory channels are available, one of them is randomly perturbed to deliberately induce a discrepancy between the available sensory channels. This can then be compared to a model to determine how the different sensory channels are combined. When considering only the effect of noise on multisensory integration, assuming independent

noise processes for different sensors, the optimal (in the sense of minimizing effect of noise) sensory estimate is given by, $\hat{x} = \sum_i \omega_i x_i$, where x_i are noisy sensory signals from different sensory channels with noise variances σ_i^2 , $\hat{x}$ is the net sensory estimate and ω_i are sensory weights such that $\sum_i \omega_i = 1$. The net estimate $\hat{x}$ is optimal when

$$\omega_i = \frac{1/\sigma_i^2}{\sum_j (1/\sigma_j^2)} \quad (1.1)$$

Other studies (van Beers et al., 1999) have found similar results when examining the ability of the nervous system to combine vision and proprioceptive sensation from the arm to localize the hand in space. Instead of a 1D static task like height of an object (Ernst and Banks, 2002a), van Beers et al. (1999) had to contend with a 2D distribution of variance in sensory perception since they tested the ability of subjects to sense the location of their fingertip on a plane. The novelty of their result is that these 2D variances were non-uniform, i.e. they were elliptical in shape. For such distributions, if there is a discrepancy in the mean location from visual and proprioceptive sensations, then the optimal combination did not even lie on the line connecting the two mean locations — a possibly counter-intuitive result. However, they found that their humans subjects also made the same error, namely, when both vision and proprioception were available, they perceived the location of their fingertip to be off the line connecting the mean location of the visual and proprioceptive feedback signals! K rding and Wolpert (2004) found that not only do humans use Bayesian inference to combine sensory signals from multiple channels, but also use Bayesian inference to combine *measured* sensory signals with the *expected* sensory signal — called the *prior* distribution (a probabilistic internal model). The result of probabilistically combining the actual sensory signal with expected sensory signal (i.e., combining distributions) yields yet another distribution called the *posterior*. K rding and Wolpert (2004) found that via sensorimotor learning, humans could learn to represent and combine their prior distributions with measured noisy sensory signals, thus behaving as predicted by optimal estimation models. They also found similar results for fingertip force discrimination (K rding et al., 2004). In another study, Saunders and Knill (2004) studied optimal estimation in the nervous system when the same sensory channel measures different variables simultaneously. For example, human vision is known to measure both location and velocity of objects in the visual field. Using a linear Kalman filter model Basar (2001) found that during reaching movements with the arm, humans combine information about fingertip position and velocity optimally to carry out the task of reaching a static target. Finally, Sober and Sabes (2003, 2005) examined motor planning, i.e. determining arm trajectory to point to a target and found that vision and proprioception are not combined based only on the sensory variance, but based on the additional noise injected by the neural system. They argue that this additional noise exists for proprioception since these sensors respond primarily to joint angle changes and need to be converted to an spatial

reference-frame which will entail necessary noisy neural computation, thus increasing the noise-level in proprioceptive sense and thus decreasing its sensory weight. This evidence is indirect, but is one of the first studies to try to address how known noise in neural computation might be significant to sensorimotor control.

These studies provide a broad spectrum of evidence that human perception probably uses Bayesian inference to minimize deleterious effects of sensory noise during *static* tasks.

Linear or linearized dynamics

Similarly, other studies have examined whether optimal feedback control that uses optimal state estimation can explain human balance control (Kuo, 1995, 2005). These studies differ from those cited above in one key aspect — they do not rely on conscious perception (both “conscious” and “perception” are ill-defined terms; see Prochazka et al., 2000), but use some form of “subconscious” control to keep balance during quiet standing. Moreover, the systems being studied here are dynamic in the sense that time-varying trajectories of the center of pressure should stay within reasonable bounds for the model (or person) not to lose stability. Since, the full-blown nonlinear system of balance control is extremely complex, these studies restrict themselves to a linear analysis of optimal sensory estimation by studying the response to very small random external perturbations delivered to their standing human subjects. They find that their linear models that use Kalman filtering for state estimation were able to reproduce various features of experimental center-of-pressure trajectories remarkably well.

Time-delay compensation

Time-delays are ubiquitous in the nervous system. The first category of studies on Bayesian inference in the nervous system that we reviewed previously all involved static task goals, which renders time-delays irrelevant. However, time-delays are critical for stability of dynamical systems such as balance control. Unfortunately, all model-based studies to test optimality of sensory integration in the nervous system during dynamic tasks such as balance control, whether linear (Kuo, 1995, 2005) or nonlinear (van der Kooij et al., 1999) use time-delay compensation to make their models robust (in this case independent) of time-delays so long as the delays are not extremely large. Time-delay compensation is just the name for the feature of optimal controllers that use optimal state estimators such as Kalman filters. If sensory signals are time-delayed, then the optimal estimator just relies on its forward internal model to predict the current state of the system being controlled using delayed sensory feedback. The only studies that have examined sensory estimation during dynamic sensorimotor behavior (and hence are potentially influenced by time-delays) are of balance control which are modeled using linear or nonlinear models that exclude effects of sensorimotor time-delays on multisensory integration by using time-delay compensation techniques. A natural

question to ask at this point is, “why bother more about time-delays if we know that the nervous system’s behavior closely resembles optimal control models that can compensate for time-delays”? There are two reasons why this is of concern to us.

1. A very well known and big problem related to aging is the dramatic increase in how often people drop objects or fall down themselves. The study of the reasons behind this dramatic increase is a vast area of research. With respect to object manipulation, it suffices to say that one careful study (Cole et al., 1998) has revealed that the reason for loss of manual dexterity is not degraded sensors (hence greater noise), thus by a process of elimination, blames the overall slowing of cortical processing, i.e. increased time-delays with aging.
2. Although optimal controllers are capable of compensating for time-delays, this ability is not robust to noise when the closed-loop system (controller + body + world) is close to instability (Stein, 2003), which is typically the case when the sensorimotor system is at its limit of performance (with the exception of static activities like producing maximum force). It is important to note that sensorimotor control is highly regarded for its remarkable performance at the limit of performance (at least compared to robotic controllers) and disease is regarded as devastating because of the dramatic loss of ability to push the limit of performance. But, because of the unavoidable and ubiquitous nature of sensorimotor noise, any time-delay compensation techniques are rendered ineffective at the edge of instability. In fact, the importance of time-delays extends beyond just multisensory integration in biological systems to the mathematical theory of Bayesian inference. Clark and Yuille (1990, page 181) in their book on mathematical techniques of multisensory integration rightly point out that, “*Time is an important, and all too often ignored, aspect of the operation of sensory systems*”. Mathematically, they find that optimal Bayesian estimate of a measured parameter using time-delayed and noisy sensory signals should depend on both noise and time-delays. In other words, the knowledge gained so far based on the monumental volume of research on human sensorimotor behavior is only half the story of multisensory integration. Keep in mind that for the purpose of this discussion and this thesis we will restrict ourselves to noise and time-delays as the primary factors that can affect sensorimotor performance and hence, multisensory integration strategies.

Because of the above cited reasons, despite their novelty and insight into the effects of sensorimotor noise, past studies on Bayesian inference cannot address how unavoidable time-delays in the nervous system affect the ability of the nervous system to use multiple sensory channels. A whole chapter of this thesis presents novel results that provide evidence for the effects of time-delays in addition to noise on multisensory integration and provides a model-based explanation for the increased reliance on vision for manipulation when you get old or cold.

1.2 Sensorimotor control of manual “dexterity”

So far we have covered what is known about the sensorimotor control of a wide range tasks, but primarily simple ones such as finger pointing. A quintessential feature of humans is our ability for dexterous manipulation using our fingers. The research into dexterous manipulation (in humans or robots) is still in its infancy, with much that is not yet known or understood. This thesis delves primarily into aspects related to dexterous manipulation in the context of previous research on optimality of sensorimotor behavior. As Sir Isaac Newton famously said¹⁰, “If I have been able to see further, it was only because I stood on the shoulders of giants”. If this thesis goes any further at all in our present understanding of sensorimotor control and of dexterous manipulation it is by building on the monumental effort of past researchers who have revealed features of sensorimotor control (reviewed above) and others who are beginning to unravel the horrendously complex system of the human hand (e.g., Valero-Cuevas, 2005; Bicchi, 2000; Johansson and Birznieks, 2004; Cole and Abbs, 1988; Johansson and Cole, 1992; Zatsiorsky et al., 2002). Here, I present a brief overview of some features of sensorimotor control using the hand relevant to the work presented in this thesis.

1.2.1 Static pinch and tactile dominance

In the neuroscience literature, holding or manipulation with the fingertips is frequently referred to as *precision pinch* or *precision manipulation*. Using the whole hand to hold an object like a hammer is called a *power grip*. *Grasp* is a term most often used to describe the process of acquiring a pinch or a grip by appropriately bringing our fingers into contact with an object. A very large number of studies on sensorimotor control of dexterous manipulation have focused exclusively on static precision pinch using the thumb and index finger to hold an object. In these past studies, the objects that experimental volunteers grasped or manipulated were of two types, *passive* or *active*. Passive objects are static objects held between the index finger and thumb, and active objects are similar to passive objects but for unexpected external perturbations (Johansson and Cole, 1994).

The first systematic and quantitative study to examine how frictional properties of the object surface and weight of the object affect the coordination of fingertip force exerted on the object was carried out by Westling and Johansson (1984); Johansson and Westling (1984). These two papers found several key aspects of fingertip force coordination during static pinch that have been repeatedly validated by several other studies:

1. The grip force (fingertip force) scaled extremely consistently and linearly with load force (weight of object supported by the fingers). For a given load, the

¹⁰This quote appeared in a letter to Robert Hooke who discovered “Hooke’s law of elasticity” for springs.

slope of the linear relationship between load force and grip force also systematically varied with frictional properties of the object surface. Importantly, the grip force was always just marginally greater than the force at which the object would slip (slip force). The results of this study indicate that the nervous system predictively modulates grip force so that it is appropriate to the weight and frictional properties of the object to be lifted.

2. The ratio of the slip force vs. load force was defined as the slip ratio and the difference between the ratio of grip force vs. load force and the slip ratio as the *safety-margin* (Westling and Johansson, 1984; Johansson and Westling, 1984). Although the safety-margin was found to be variable across individuals, it was typically very small. Hence, these researchers suggested that the nervous system uses predictive sensorimotor control to hold objects hard enough so that they don't slip, but not too much harder than necessary. Interestingly, the safety-margin increased dramatically with any tactile impairment (using a local anaesthetic; Johansson and Westling, 1984; Cole and Abbs, 1988; Monzee et al., 2003; Augurelle et al., 2003). This has led to the proposal that the safety-margin is a hallmark of sensorimotor control of manual dexterity.
3. Perhaps the most salient find from several studies on static pinch is the dominance of finger tactile sensation to static pinch. For a passive object, tactile impairment dramatically increased the safety-margin, and completely abolished the stereotypical modulation of grip force with changes in frictional properties. Cole and Abbs (1988) found that with all sensation intact, the grip force responded as fast as 60ms to any unexpected external perturbations. But, once tactile sensation was impaired the latency of the grip force response dramatically increased to as much as 160ms and also became extremely variable. This strong dependence of fingertip force coordination on tactile sensation from the digits has been repeatedly established by several studies on static pinch (e.g., Johansson and Westling, 1987; Johansson and Cole, 1992, 1994; Johansson, 1996; Monzee et al., 2003; Augurelle et al., 2003). In fact Johansson and Birznieks (2004) used the paradigm of static pinch to find that the rapid response to perturbations using tactile sensation is because information about object slip is encoded using only the first neuronal spike from an ensemble of tactile sensors in the fingerpad.

A series of publications also examined how these stereotypical patterns of static pinch developed in children (Eliasson et al., 1995; Forssberg et al., 1991, 1995, 1992; Gordon et al., 1992) and revealed that the predictive features of precision pinch took several years to develop and mature.

Cole and Abbs (1988) observed that, after losing digit sensation because of an administered local anesthetic, "subjects found it difficult to grasp and lift the object without visual monitoring of the hand and object". Subsequently, other studies (Häger-Ross and Johansson, 1996; Macefield and Johansson, 1996; Monzee

et al., 2003; Augurelle et al., 2003; Cole et al., 1998) have begun to explore how sensations other than tactile sensation contribute to static pinch. For example Häger-Ross and Johansson (1996) found that if the hand was not restrained at the wrist, even subjects with a tactile nerve-block showed grip force modulation in response to external perturbations. But, the latency of response in grip force modulation was greatly prolonged after a digital anesthesia. Therefore, they concluded that digital tactile sensors are dominant in feedback control and non-digital receptors such as non-digital skin, muscle spindles, etc. are used only when digital tactile sensation is absent. Augurelle et al. (2003) find that even in the absence of digital tactile sensation, subjects could respond appropriately to external perturbations and speculate that it could be due to subjects using visual or proprioceptive sensory feedback. However, Monzee et al. (2003) found that even with vision the force vectors from the thumb and index finger were not properly aligned causing the held object to rotate and slip for subjects with digital anesthesia. They speculate that the lack of sufficient visual acuity could be to blame for this.

In summary, although the theories of optimal control that we reviewed should be applicable to object manipulation, there are no studies that examine this quintessential human behavior of manual dexterity using the framework of optimal estimation. Instead, current studies on object manipulation are limited to heuristic speculation about the relative merits of various sensory channels. Moreover, as evident from the literature survey presented in this section, almost all studies on dexterous manipulation are still restricted to static tasks although in the very first paper that used this paradigm of static pinch Westling and Johansson (1984) observed that static precision pinch “serves [only] as a basis for further precision handling which may involve rotation of the object about one of its own axes ...”. Hence, there is undoubtedly a vital need for methods to quantify and explore the nonlinear dynamic aspects of dexterous manipulation.

1.3 Nonlinear dynamics in biological sciences

Not only do we lack studies of dynamic manipulation in humans, but, existing studies of nonlinear dynamic manipulation in robotic mechanisms are also highly limited despite the use of extremely complicated mathematical methods (Bicchi, 2000). However, there exist some studies that have explored nonlinear dynamical behavior in other sensorimotor systems and various other dynamic systems (such as disease dynamics, weather patterns etc.) that could provide insight into how we could study the nonlinear dynamics of manipulation (Longtin et al., 1990; Mackey and Glass, 1977; Mackey and Santillan, 2004; Collins and DeLuca, 1994; Cabrera and Milton, 2002, 2004a; Beuter, 2003; Poon and Merrill, 1997; Wiesenfeld and Moss, 1995; Robertson et al., 2004).

At this juncture, it is useful to briefly define what we mean by a “dynamical system” and “bifurcations” which are terms that will be frequently used in this section. A dynamical system in continuous time is usually described by a set of

ordinary differential equations such as $\dot{x} = f(x; \lambda)$, where x is a set of n variables that suffice to describe the state of the dynamic system under consideration, $\dot{x}$ is the time-derivative of x determined by x and p parameters λ as specified by the function $f : \mathbb{R}^{n \times p} \rightarrow \mathbb{R}^n$. An equivalent geometric picture of this dynamical system for specific values of λ is the n -dimensional picture that traces how x evolves over time for every possible initial condition — called the phase portrait of this dynamical system. Sometimes, when one or more components of λ change, the phase portrait of the system undergoes a sharp, qualitative change that is called a bifurcation. For example, consider a *pitchfork bifurcation* (Guckenheimer and Holmes, 1983) for a dynamical system described by $\dot{x} = \lambda x - x^3$, where x is a scalar; $x = 0$ is the only equilibrium point and is stable for $\lambda < 0$. However, at $\lambda = 0$, $x = 0$ is no longer stable, but just marginally stable, i.e. the trajectories of this dynamical system are no longer attracted towards $x = 0$. Once $\lambda > 0$, the point $x = 0$ becomes an unstable equilibrium point and two new stable equilibria arise at $x = \pm\sqrt{\lambda}$. This qualitative change that happens at $\lambda = 0$ is called a bifurcation (of the pitchfork type in this example).

Longtin et al. (1990) studied the human pupil light reflex where oscillatory phenomena of varying complexity have been observed. They choose that system as a paradigm for neurological control systems because, the “advantage of studying the pupil light reflex... is the ease with which the feedback loop can be opened (i.e., the effect of the output on the input can be removed). This has led to extensive studies of the linear and nonlinear properties of the reflex components... closed-loop oscillations can be studied using an experimental setup in which the normal feedback is replaced by controllable external electronic feedback.” Hence, they could model the dynamic behavior of this system using first-principles¹¹ as a delayed differential equation with noise. They explore and reveal several interesting interactions between noise and bifurcations (for example, they find that a Hopf bifurcation occurs in this system that is postponed by the presence of additive or multiplicative noise). Although the exact details of the interactions are not directly relevant to us, the message to take away from their studies is the power of a combined experimental-mathematical approach that uses nonlinear dynamical systems to reveal the workings of a neurological system. The edited book by Beuter (2003) contains a very good collection of such studies that have examined nonlinear dynamical aspects of various biological systems ranging from a single cell to a whole human.

Not all of the above cited studies started with systems that could be modeled from first-principles. For example, Robertson et al. (2004) sought to explain how infants choose to focus attention when presented with a visual scenery. Specifically, they sought to determine if the frequent shifts (saccades) that are experimentally observed in infants is driven by some visual processing or by some intrinsic noisy process. They found that a nonlinear dynamic system that had multiple stable

¹¹See Section 1.4 for our chosen definition of first-principles vs. phenomenological models.

points (like $\lambda > 0$ in the example presented earlier) representing objects of interest in the visual scenery, together with hysteresis and additive Gaussian white noise acted quantitatively like experimental measurements on 4-week-old infants. Thus they concluded that noise and hysteresis could explain the ‘visual foraging’ that is observed in infants. Yet another series of studies that are well summarized in the review by Wiesenfeld and Moss (1995) found that the presence of noise in nonlinear dynamical systems could be beneficial. The common theme in the studies reviewed by Wiesenfeld and Moss (1995) is the development of nonlinear dynamic models either from first-principles or from the phenomenology of the system of interest (that ranged from biological systems, climate systems to electromagnetic systems) and a model-based examination of the effects of noise.

Another class of studies that use nonlinear dynamic techniques such as Collins and Deluca (1994); Cabrera and Milton (2002, 2004a) rely on nonlinear time-series analysis to reveal aspects of interest about the systems they study. Collins and Deluca (1994) found that the postural sway observed during quiet standing was indistinguishable from a random walk so long as deviations of the center of pressure under the feet were not too large. For large deviations, there seemed to be rapid corrections that returned the center of pressure toward the center of contact of the feet. This result is strong evidence for the presence of hybrid control (continuous-time + event-driven) in human balance control. Similarly, Cabrera and Milton (2002, 2004a) find that the human ability to balance sticks reveals something known as ‘on-off intermittency’ that points to the benefit of multiplicative noise in feedback control and of deliberately pushing the system to the edge of instability.

It is hard to miss that most of these studies are interested in bifurcations in one way or another. There are two main reasons for this that arise from the mathematical properties of almost all nonlinear dynamic systems, irrespective of their complexity.

1. The center-manifold theorem states that at the point of bifurcation of almost all dynamic systems, all non-trivial dynamics are restricted to a low-dimensional (compared to the dimension of the full system) manifold. This result is of course governed by several caveats about properties that the dynamic system under consideration should satisfy. However, from the application point of view, this creates the possibility of modeling otherwise intractably complex systems such as sensorimotor control of the hand, global climate, etc.
2. There are only a finite number of possible bifurcations of a fixed point when one or two parameters (e.g., components of λ in the example discussed above) are changing. These bifurcations can then be described by ‘normal form’ equations that closely approximate the dynamics of the full system when restricted to the center-manifold. We restrict our attention to fixed points and only one or two changing parameters, since that is the only focus of this

thesis and permits us to make a stronger statement (finiteness) about the number of possible bifurcations.

1.3.1 Benefits of the edge of instability

Based on these two results stated above, it seems that the benefits of pushing a system to an edge on instability, where a bifurcation might ensue are tremendous. That is indeed the case and the principal motivation for the work presented in this thesis. It is especially important to examine the dynamics at the edge of instability of sensorimotor control systems for reasons other than just the convenience of mathematical modeling. A very illuminating lecture by Stein (2003) called ‘Respect the unstable’ goes over the challenges posed by systems that are nearly unstable to the design of robust controllers from the engineering perspective. The realization that skilled behavior such as dexterous manipulation or ice skating are often at the sensorimotor limit of performance makes the importance of studying sensorimotor control near an instability especially pertinent.

I now present examples of studies that used the edge of instability as a tool to examine complex nonlinear dynamics of various systems. Based on observed time-series of occurrence of measles, Earn et al. (2000) find that bifurcations of a simple nonlinear model can explain various cyclic occurrences of measles outbreaks. Alvarado et al. (1994) observe that voltage collapse and blackouts in electric power systems due to increase in demand is because the power grid loses stability through a saddle-node bifurcation. They use this observation to design algorithms to detect such voltage collapses before they occur. Milton et al. (2004) find that epileptic attacks can be modeled as a dynamic system at the edge of instability and propose along the lines of theories of self-organized criticality (Bak et al., 1987, 1988) that most neural systems are naturally poised at the edge of instability.

In a study that was the direct precursor to the work in this thesis, Valero-Cuevas et al. (2003b) used the edge of instability as a measure of the limit of sensorimotor performance to quantify the effects of impairment to the thumb due to osteoarthritis.

1.4 What this thesis is about

The different ways to model a system could be roughly categorized into two blocks — first-principles vs. phenomenological. A ‘first-principle’ is a statement in a formal logic system that cannot be deduced from any other. In that strict sense, for physical systems, only a molecular or smaller-scale model whose physical laws derive purely from empirical observations and cannot be broken down further can be regarded as a first-principles model. But, we will stick to a less strict definition in this thesis — if the system can be modeled using empirical “laws” that describe its constituent components, then we will refer to it as a first-principles model. On the other hand, if we use empirical observations of the entire system of

interest to develop a model with no knowledge about its constituents, we will refer to the model as a phenomenological model. In this chosen grammar for classifying models, a first-principles model is just an assembly of several phenomenological models! First-principles models are easy to justify based on known physics of the system being modeled, but more often than not, such detailed models cannot be developed because of unknown or unmeasurable properties of very complex systems like the human sensorimotor system. On the other hand, phenomenological models are often not restricted by our ability to measure properties of the original system, but typically suffer from lack of generalizability of results they find because of their specificity. Almost independent of the modeling approach chosen, “understanding” can be achieved in one of two ways — dynamical/emergent behavior vs. goal-oriented behavior (e.g. optimality, external forcing or somehow explicitly specifying a goal, i.e., fitness function approach). I say almost independent of modeling approach because there is no obvious coupling that is necessary between any of the modeling approaches, vs. any of the understanding approaches, but one typically finds that most studies that use first-principle models will resort to the emergent behavior flavor of understanding and most studies that use a phenomenological approach will use a goal-oriented behavior flavor of understanding. The reason for this coupling is in some ways a purely technological limitation. Goal-oriented approaches are often optimality based and suffer from computational and convergence issues if the system being studied is very high-dimensional, so they resort to simpler models that are often either derived phenomenologically or are in some way a cartoon of the real system being studied. However, this coupling is often trespassed in relatively contemporary studies either by using sheer computational power or by ideas such as emergent low-order dynamics near a bifurcation that may be considered as valid cartoons of the real system with as few assumptions as possible. The body of work presented in this thesis sits firmly in the territory of a “cartoon of reality” approach, and uses bifurcation theory of nonlinear dynamics to build the cartoon (i.e., phenomenological) model. In some ways, a purpose of this thesis could be considered as trying to highlight the strength of bifurcation theory in helping generate cartoons of reality that are very close and valid approximations of the system being studied while still remaining mathematically and computationally tractable. In terms of “understanding”, this thesis lies somewhere between the emergent behavior and goal-oriented flavors. The development of the cartoon model relies on mathematical theories that classify the types of emergent behavior when a system undergoes a qualitative change in its dynamics. But, once the cartoon model is developed, the work presented here resorts to numerical optimization that tries to achieve an externally specified goal (maximize a well defined fitness function), which is an interpretation of the nervous system’s functioning in the context of our experimental design.

In summary, this thesis tries to address three intertwined questions:

1. How can we design an experimental technique and the mathematical tools necessary to extend our understanding of dexterous manipulation into the

realistic regime of nonlinear dynamic manipulation?

- We address this issue by designing an experimental task of dexterous manipulation at the edge of instability. This enables us to model the otherwise intractable behavior of dynamic manipulation using a simple low-order model that is mathematically and computationally tractable, but because it naturally emerges due to the property of dynamic systems at the edge of instability, it remains extremely generalizable to a large class of systems (biological and others).
2. During dynamic sensorimotor control, does the nervous system perform sensory estimation according optimal estimation theories, namely, does multisensory integration emerge from the combined effect of noise and time-delays?
 - We perform an equivalent of previous studies on sensory occlusion (see Section 1.2.1), but use our novel paradigm and provide an explanation in the framework of Bayesian inference for the tactile dominance in intact hands and visual dominance once tactile sensation is impaired.
 3. Can we quantify and understand the clinical consequences of treatments and diagnose impairments of the hand using the neurophysiological understanding of sensorimotor control that we have obtained?
 - We use the techniques developed in this study to understand why a specific treatment for osteoarthritis of the carpometacarpal joint of the thumb caused improvement in everyday function for the patients who underwent treatment.
 - We use our model of multisensory integration to provide explicit predictions and direction for future work to design a measurement technique for detecting specific neurological diseases such as carpal tunnel syndrome or multiplesclerosis in their early stages.

1.5 What this thesis is not about

It is important to clearly outline what this thesis is not about, especially since there is a body of literature in sensorimotor control that uses similar language to this thesis. Schöner and Kelso (1988) performed experiments of periodic bimanual finger motions that they modeled as coupled oscillators. They found that increasing the frequency of finger oscillation automatically led to mirror symmetric bimanual oscillations and any other form of oscillation could not be sustained by their subjects. They related this phenomena of stereotypical behavior beyond a critical frequency to the work of Haken (1975) and claimed that the nervous system has a tendency to spontaneously self-organize sensorimotor control. Kelso (1995) provides a collection of several such studies and goes further to propose that the nervous system spontaneously organizes itself near some form of criticality. To

state this in terms of the dynamical systems language we have been using so far, Schöner and Kelso (1988); Kelso (1995); Sternad (2000) postulate that the nervous system and sensorimotor control in specific exhibits self-organized criticality (Bak et al., 1987, 1988). They go further to postulate (Kelso, 1995) that the emergent self-organized state is often at the edge of instability, i.e. near a bifurcation, which is not necessarily true of self-organized criticality. Although these studies use similar language to the work in this thesis, the important distinction lies in the fact that those studies (Schöner and Kelso, 1988; Kelso, 1995; Sternad, 2000) do not use results of bifurcation theory to develop generalizable models, but develop specific phenomenological models of coupled oscillators that best fit their experimental data. Moreover, these studies are along the lines of “emergent behavior” route of “understanding” sensorimotor control, which is a distinctly different flavor from the “goal-oriented” approach of this thesis. In fact, it is not even clear if the postulated self-organization of sensorimotor control is a valid claim since more recent studies have shown experimental evidence to the contrary (Mechsner et al., 2001).

1.6 Why study the hand?

The hand is considered by some to be what sets humans apart from most other species (Valero-Cuevas, 2005; Bicchi, 2000; Clack, 2004). Not surprisingly, robust and versatile hands have proven to be the hallmark of a large volume of contemporary research in robotics (Bicchi, 2000; Okamura et al., 2000). Above all, any impairment to the hand often has devastating effects on a person’s quality of life (Valero-Cuevas et al., 2003b; Valero-Cuevas and Hentz, 2002). Yet, despite this apparent gravity of the hand, very little is known or understood about the hand and its quintessential, yet incredible versatility (Valero-Cuevas, 2005; Bicchi, 2000).

What made the human hand what it is today? How can one model or understand the complex interaction of the nervous system, the hand and the world that yields such rich and versatile behavior? Is the human hand distinguished because of or despite of its nonlinearities, noisiness, time-delays and such aspects that are extremely hard to mathematically describe? How do the interactions of the nervous system at the level of individual neurons and the muscles at the level of motor-units (if not smaller) result in the emergence of such coherent and robust behavior? In all, the hand appears to be the *Kurukshetra*¹² of science, where biology, physics and mathematics converge in a battle of many unsolved problems, several of which promise to deliver deep insights into the working of not just our

¹²The purported site of a legendary battle in ancient India around 3067 BCE that is depicted in the Hindu epic *Mahabharata*. This battle served as the setting for the *Bhagavad Gita* that summarizes a number of philosophical explanations in Hinduism for the meaning and purpose of life.

body, but the physical world in general. This work aims to establish the foundation upon which some of these questions can be explored in the future.

1.7 Layout of this thesis

The following is the specific layout of this thesis based on the goals and hypotheses listed in Section 1.4. In Chapter 2, we introduce an experimental task of compressing a slender spring that pushes the closed-loop thumb + object + nervous system to the edge of instability. We use the normal form of a subcritical pitchfork bifurcation to model this system based on the mechanics of the object and measured spatiotemporal dynamics. We provide evidence that the spatiotemporal dynamics of our model are indistinguishable from experimentally measured dynamics. We also provide experimental evidence to demonstrate that this task truly takes people to their limit of sensorimotor performance that is independent of muscular strength. We present evidence from our experiments that is contrary to the notion of ‘safety-margins’ that is prevalent in the literature of static precision pinch and provide a possible explanation for our result. In Chapter 3, we extend our model to incorporate multisensory feedback from thumbpad tactile sensation, vision and non-digital mechanoreceptors. Importantly, we incorporate noise and time-delays in this model based on known neurophysiological data. We then use numerical optimization on this model to explain the effects of sensory occlusion in the framework of Bayesian inference, where both time-delays and noise influence sensory weights. We also provide experimental evidence that experimental results of sensory occlusion closely resembled model behavior. Chapter 4 provides preliminary evidence that explains improvement in manual dexterity after treatment using Hylan G-F 20 (Hyaluronic acid) for osteoarthritis of the carpometacarpal joint of the thumb. Specifically, we propose that self-reported improvement in dexterity was primarily due to alleviation of pain, thus permitting the patients to better use their inherent sensorimotor ability. Then, in Chapter 5 we present a brief list of short, intermediate and long term goals for future work based on the results and methods developed in this thesis. We conclude the thesis with a discussion of implications of this work to optimal feedback control theory in general and sensorimotor control in specific.

Chapter 2

Dynamic manipulation at the edge of instability — Experimental design and mathematical modeling

2.1 Introduction

In this chapter we provide evidence that experimentally measured spatiotemporal dynamics of manipulation at the edge of instability are indistinguishable from that of a low-order nonlinear dynamical model that we developed using bifurcation theory. The work presented in this chapter lays the necessary foundation for subsequent chapters of this thesis where we test neurological and clinical hypotheses using the combination of our experimental task and low-order model.

2.1.1 Background

In Chapter 1 we reviewed existing literature about sensorimotor control of dexterous manipulation. We saw that the study of manipulation in humans has been restricted to static force production and force coordination in the past. However, a paradigm used in some of these past studies that is relevant to the work presented in this thesis is that of pushing the neuromechanical¹ system of the hand to some point of uniqueness. To clarify this further, let us look at an example. Valero-Cuevas et al. (1998) imposed a goal of trying to maximize static force production from the index fingertip. They found subject-independent muscle coordination patterns and a computer model predicted muscle coordination patterns for maximizing force output (using linear programming — an optimization method) that agreed remarkably well with experimental measurements. Other examples are studies that have examined muscle coordination during maximum height jumping (Zajac et al., 1981, 1984; Pandy et al., 1990). Once again, they found consistent muscle coordination patterns that agreed well with optimal control models. In both these examples, the researchers pushed the neuromechanical system to a limit of performance, thereby pushing the system to near-uniqueness, i.e., only very few motor control strategies could yield task-optimal performance (be it maximum force or

¹This phrase is used to represent the fact that the nervous system, muscles of the hand and the kinematic mechanism of the digits are best treated as a coupled dynamic system (Valero-Cuevas, 2000). This approach is in contrast to the strictly hierarchical picture of sensorimotor control with the nervous system at the top as adopted by several researchers. However, the phrases *neuro* and *mechanical* do reflect the fact that we still distinguish the neural and musculoskeletal components of this system as distinct entities. This distinction helps place any analysis of this system in the realm of goal-oriented behavior as discussed in Section 1.4.

maximum jump height). We sought to design a similar paradigm, where the neuromechanical system of the hand manipulating an object was pushed to some limit of performance. However, in contrast to the past studies on manipulation that relied on static tasks, we developed a dynamic manipulation task that required very low forces.

2.1.2 Past work – The “Strength-Dexterity Test”

The immediate precursor to the work presented in this study, referred to as the “Strength-Dexterity Test” or simply the “S-D Test”, has been previously reported (Valero-Cuevas et al., 2003b; Valero-Cuevas, 2003). In this past study (Valero-Cuevas et al., 2003b), we developed a method to decouple pinch “strength” from pinch “dexterity”. We chose to define dexterity as the ability to regulate fingertip force vector (i.e., magnitude and direction) as well as fingertip position in order to control an inherently unstable object. We graded “dexterity” by asking subjects to handle objects of different dexterity requirements. The objects were helical compression springs and the manipulation task was to compress these spring until they reach their solid length, i.e., the coils come in contact with one another. For a given spring, the slenderness ratio (length/diameter) represents the dexterity requirement and the stiffness (or force magnitude required to compress it to solid length) represents the strength requirement. This ability to decouple strength and dexterity enabled us to define a Strength-Dexterity plane that we discretized using springs of various slenderness ratios and stiffnesses (see Figure 2.1).

Dexterity Index

We defined “Dexterity Index” as a function of the geometry of the spring given by,

$$\text{Dexterity Index} = \frac{1}{DC_1} \sqrt{\frac{2y_{max}C_1L_0 - y_{max}^2}{C_2}} \quad (2.1)$$

where D is the mean spring diameter, C_1 and C_2 are constants that depend on the spring material, and L_0 is the free length of the spring. For a given spring, y_{max} is the maximal distance the spring can be compressed before reaching its solid length. This formula is obtained directly from the well-known equation to calculate the critical spring compression at which it buckles under a static load (Timoshenko, 1961; Wahl, 1963; Shigley and Mischke, 1989) for given end-conditions (say, fixed, pinned, free, etc...). The vertical compression (distance) at which the spring elastically buckles is related to the imposed end-condition. The end-condition is quantified by the ratio of the ‘effective-length’ vs. ‘free-length’ of the spring. The effective length for a given end-condition is the distance between zero-moment points of the spring (or column) just before it is about to buckle. The end-condition parameter is then defined as the ratio: $\alpha = L_{\text{eff}}/L_0$. For the example of a column buckling with one end ‘fixed’ and the other end ‘free’ that is shown in Figure 2.2,




Figure 2.1: The discretized strength-dexterity plane. Adapted from Valero-Cuevas et al. (2003b) with permission. We found that every unimpaired young adult (< 40 years) could compress every spring inside the rectangle bounded by dashed lines. The same could not be said for older adults, impaired adults (Valero-Cuevas et al., 2003b) or children (unpublished data).




Figure 2.2: A column buckling example with fixed-free end conditions to illustrate the definition of “effective length”. The curvature of the column and the lateral deflection of the free end are exaggerated for clarity. The definition for end-condition parameter α holds precise meaning at the load when the column just begins to buckle, i.e., in this example, the free end just begins to elastically deflect away from the previously stable centerline.

the effective length is twice the length of the column, hence, $\alpha = 2$. With this definition of α and the same definitions of D , C_1 , C_2 and L_0 as listed above, the critical spring compression is given by,

$$y_{cr} = L_0 C_1 \left(1 - \sqrt{1 - \frac{C_2 D^2}{(\alpha L_0)^2}} \right) \quad (2.2)$$

Using this formula, a natural definition for dexterity arises in the context of using the fingertips to compress a spring. If a person successfully compressed a spring to its solid length, i.e. $y_{max} = L_0 - L_s$, where L_s is the solid length of the spring, then we can replace y_{cr} in Equation (2.2) with y_{max} , the compression at solid length, to calculate a lower bound on the effective end-condition (α_{min}) they could emulate using their digits. The Dexterity Index (Equation (2.1)) was then defined as the reciprocal of α_{min} for each spring in the discretized grid of the S-D plane. Say, for a specific spring, the critical compression for fixed-free end-conditions was at the compression just before it reached solid length (e.g., springs in row ‘D’ of Figure 2.1), then if the person could emulate a fixed-free end condition using their digits, they would be able to successfully compress the spring. However, this person may also be able to emulate stricter end conditions like fixed-pinned (where the top end is allowed to rotate, but not deflect, e.g., row ‘F’). That is why, if a person could successfully compress a given spring, the corresponding ‘Dexterity Index’ is only a lower bound.

Dexterity – the ability to control an instability

By making subjects compress every spring from the S-D grid (Figure 2.1) and quantifying the area of the spring-grid that each subject could successfully compress, we found that unimpaired young adults could cover a larger area than unimpaired older adults, who were in turn better (i.e., covered a larger area) than older adults with osteoarthritis of the carpometacarpal joint. The main message to take away from this study is that defining “dexterity” as the ability to control an instability seems to be informative of manipulation ability, at least consistent with the qualitative notion that age degrades dexterity and disease degrades it further.

2.1.3 “Manipulating the edge of instability”

Dynamic sensorimotor behaviors such as manipulation are quintessentially complex, non-linear and high-dimensional. Due to the mathematical intractability of such systems, past attempts at a nonlinear dynamic description of sensorimotor control have remained limited in one of two ways.

1. The systems studied were simple enough so that they could be modeled by assembling empirically well-known properties for their constituent components. For example, Longtin et al. (1990) sought to understand various rhythms observed in the pupil light reflex. Although there are no known studies that

have modeled sensorimotor control systems using first-principles in the strict sense², the study by Longtin et al. (1990) is an example that comes close to a first-principles approach. In their own words, the “advantage of studying the pupil light reflex... is the ease with which the feedback loop can be opened (i.e., the effect of the output on the input can be removed). This has led to extensive studies of the linear and nonlinear properties of the reflex components... closed-loop oscillations can be studied using an experimental setup in which the normal feedback is replaced by controllable external electronic feedback.” This quote is illuminating because it reveals an important facet of the system they chose to study, namely, the dynamical behavior of the underlying components of the pupil light reflex were well studied and modeled, hence, they could model the complete system of the reflex starting from first-principles, i.e., empirically observed, well-defined constituent properties. However, that is unfortunately not the case for more complex sensorimotor behavior like object manipulation, where almost none of the underlying components have been well understood!

2. A second category of studies on nonlinear dynamics of sensorimotor control examine complex systems with little knowledge about its constituents, but rely on phenomenological models that are often not generalizable beyond the specific time-series they model. For example, several studies have examined apparent qualitative transitions in behavior during rhythmic bimanual finger oscillations as the frequency of oscillation is increased (Schöner and Kelso, 1988; Kelso, 1995; Sternad, 2000). Any system with stable limit-cycles will exhibit oscillatory behavior in some region of its phase space. They chose to model their experimental observation using coupled van der Pol - like oscillators. These studies found that their model exhibited the same qualitative transitions as the real systems they were studying and posited that their model was in fact a true representation of the underlying phenomena. Such approaches to phenomenological modeling do not yield uncontested models or explanations of the real system being studied since the models lack any form of uniqueness or genericity, thus limiting the scope of their applicability. In fact, for the example of bimanual rhythmic motor control, there is more recent experimental evidence that challenges (Mechsner et al., 2001) the inferences drawn from the phenomenological models that were developed for bimanual coordination.

However, for the complex system of dynamic manipulation, we overcome this impasse between first-principles models (not necessarily in the strict sense) vs. phe-

²A ‘first-principle’ is a statement in a formal logic system that cannot be deduced from any other. In that strict sense, only a molecular or smaller-scale model whose physical laws derive purely from empirical observations and cannot be broken down further can be regarded as a first-principles model for most physical phenomena.

nomenological models by defining dexterity as the ability to control an instability. This definition leads to a natural choice for modeling and analyzing the dynamics of manipulation at the edge of instability as we shall see in this chapter. As discussed in Sections 1.3 and 1.3.1, the mathematical theory of bifurcations provides necessary tools for identifying and modeling ubiquitous low-dimensional dynamics in arbitrarily complex nonlinear dynamic systems at the edge of instability. Thus, studying manipulation at the edge of instability allows us to model the otherwise intractably complex nonlinear dynamics of dexterous manipulation.

Based on these fundamental properties of non-linear dynamical systems, we designed the novel paradigm of asking our subjects to compress a slender spring to near-instability using their thumb-pad (Figure 2.3), thus pushing the fused thumb + spring + nervous system to the edge of instability. This is a nonlinear




Figure 2.3: A schematic showing the experimental setup used in this study. A representative sample compression is shown on the right, where a subject slowly compressed the spring to minimize the volume of the audio feedback provided (not shown), without letting the spring slip and maintained that compression for 7 seconds before slowly releasing the spring. For the sake of clarity of the figure, we have not shown the hand and arm restraints.

dynamic task using noisy and time-delayed multisensory information from thumb-pad sensors, non-digital sensors (cutaneous sensors and joint sensors outside the thumb, muscle spindles, and Golgi tendon organs) and vision. The dynamics of spring buckling is known to undergo either a pitchfork (one-dimensional buckling dynamics) or Hopf (two-dimensional buckling dynamics with sustained oscillations)

bifurcation (Croll and Walker, 1972; El Naschie, 1990; Guckenheimer and Holmes, 1983; Lobas et al., 2002; Lobas and Lobas, 2004). Since we never observed any oscillatory behavior in our experiments – the signature of a Hopf bifurcation – we modeled our task with the normal form of a subcritical pitchfork bifurcation (a one-dimensional nonlinear dynamical system; Figure 2.8). We found model parameters using the frictional properties of the thumbpad-object interface and the measured highest compressive load reached by subjects.

In summary, we propose to extend our prior experimental study (Valero-Cuevas et al., 2003b) by forcing the closed loop thumb + spring + nervous system to an edge of instability. This enables us to exploit ubiquitous low-dimensional phenomena in nonlinear dynamical systems to quantify and understand dynamic manipulation.

2.1.4 Hypotheses

Based on the above discussion, we proposed and tested the following hypotheses.

1. A low-order model based on a subcritical pitchfork bifurcation normal form can capture the dominant dynamics of manipulation at the edge of instability. We expect this to be true since the chosen equation is not an arbitrary phenomenological model, but a center manifold reduction of the dynamics of the fused thumb + spring + nervous system at the edge of instability.
2. The edge of instability reached in this task of compressing a slender spring is a strength independent limit of performance. Based on past work (Valero-Cuevas et al., 2003b) we know that the area of the S-D plane covered by unimpaired subjects was independent of pinch strength, i.e. the ability to regulate fingertip force and motion to control an instability was independent of strength. Hence, we expected the ability to push the edge of instability to the furthest possible (compress the spring to the furthest possible) to be independent of strength as well.
3. There is no safety-margin during dynamic manipulation like observed in static precision pinch. This hypothesis is motivated by the qualitative observation in everyday activities that holding an object “hard-enough” is not the only determinant of typical object manipulation. An object may slip even if held hard enough or conversely, an object may be manipulated successfully despite very low contact forces (e.g. twirling a pen). So, we hypothesized that the notion of a “safety-margin” is not applicable to dynamic manipulation tasks.

2.2 Experimental methods

After giving written informed consent, nine males and three females (19-40 years of age, mean=23 years) participated in this study approved by the Cornell University Committee on Human Subjects. All subjects were right-handed, healthy young adults with no known impairments or recent injuries to their hand and had normal visual acuity or normal with correction. None of the subjects had any prior experience with this experimental task, namely compressing slender springs.

We tested our hypotheses by asking our subjects to compress a slender spring prone to buckling as far as possible (i.e., maximize compressive force) without letting the spring's endcap slip from under their thumbpad (see Figure 2.3). The performance metric was the maximal compressive spring force that they could sustain (F_s) for 7 seconds, which is a measure of the greatest instability they could control for a prolonged period of time. To ensure that our measurements were not during a transient learning phase, we performed a rigorous training routine for all subjects over 2 days (see below). Note that all data that we report in the remainder of the paper refer only to day 2 (post-training) unless otherwise explicitly mentioned. We also measured subjects' performance using a 2x2 randomized factorial design of 4 sensory occlusion conditions: with or without vision (via blindfolding) x with or without thumbpad sensation (via Lidocaine nerve block of the radial and ulnar branches of digital nerves of the thumb; see Section 3.3.1 on Page 66).

The data from the sensory occlusion experiments are presented in Chapter 3 where we use the experimental task and model developed in this chapter to test a hypothesis about optimal state estimation in the nervous system. However, because the data are all derived from the same experiment (same subject pool), we included all data in our statistical analyses to control for overall Type-I error rate (the probability that we reject our null-hypothesis incorrectly). Type-I errors are a predominant source of incorrect scientific conclusions, since incorrectly rejecting the null-hypothesis amounts to acceptance of the alternate hypothesis, in contrast to not rejecting the null-hypothesis, when no inference is made. For example, say we wanted to test a statistical (null-) hypothesis that the mean age of all readers of this thesis is 50 years and surveyed the age of a randomly selected, sufficiently large set of people who have read this thesis to test our null-hypothesis. We need to propose an alternate hypothesis, say, the mean age is not equal to 50. If our statistical test does not reject the null hypothesis (say using a t-test, the 'p-value' was 0.1, greater than our threshold for rejecting the null-hypothesis), that does not imply that the null hypothesis is true. The only conclusion that can be drawn is that our data are insufficient to reject the null-hypothesis, in contrast to saying that we accept our null-hypothesis. Though this seems like a trivial subtlety, the core of statistical hypothesis testing relies on this subtlety and is by no means trivial!

2.2.1 Experimental setup

Subjects tried to compress a slender helical spring prone to buckling using just their thumbpad (Figure 2.3). The design specifications for the spring used are — free length = 76.2 mm, mean diameter = 8.7 mm, wire diameter = 0.79 mm, total coils = 24, material: music wire (part no. 12201, Century Spring Corp., Los Angeles, CA). We chose this specific spring based on buckling load calculations for a dead-load. Specifically, we chose the free length so that it would be a little longer than the width of the palm when the hand was held in the orientation shown in Figure 2.3. Given the free length, we picked a spring that was sufficiently slender so that it buckled at very low loads ($< 10\text{N}$) for fixed-pinned end-conditions. This procedure was guided by Equation (2.2) and trial - and - error. We mounted this spring in polymer (ABS P400, Stratasys Inc., Eden Prairie, MN) endcaps on both ends, taking care to use only the final, closed coil of the spring so that the mounting does not change the free length of the spring. We manufactured the endcaps using a rapid prototyping 3D printer (model FDM3000, Stratasys Inc., Eden Prairie, MN) with hollow internal structure to avoid unnecessarily increasing the inertia of the endcap and affecting the unstable dynamics of the system. The endcap surface was smooth and flat (thumbpad-endcap coefficient of friction ~ 0.5) with a small (0.1 mm tall) conical projection at its center, precisely coincident with the cylindrical axis of the spring. This projection provides a tactile cue for the cross-sectional center of the spring. We then mounted the spring on a uniaxial load cell (model SML-25, Interface Inc., Scottsdale, AZ) and logged data using a 16-bit digital-to-analog data acquisition system (Vicon Peak, Lake Forest, CA), sampled at 1000 Hz. We also recorded 3D location and orientation of the spring’s endcap using a 4-camera motion capture system (Vicon Peak, Lake Forest, CA). We rigidly attached the entire setup to a table and adjusted its height so that the elbow was at approximately 90° flexion during the experiment.

As shown in Figure 2.3, the thumb rests on the endcap and the unused fingers were curled around a vertical post. We fixed the forearm using a vacuum pillow (model Versa Form, Sammons Preston Roylan, Bolingbrook, IL) with the wrist placed in neutral flexion-extension and ad-abduction. However, we did not fix the base of the thumb or the wrist since non-digital mechanoreceptors may contribute to object manipulation as seen in automatic slip-grip responses (Häger-Ross and Johansson, 1996; Macefield and Johansson, 1996). We adjusted the height of the whole arm using foam pads beneath the vacuum pillow so that the distal phalanx of the thumb was horizontal before subjects compressed the spring and they could move their thumb freely over their voluntary range of motion comfortably while still maintaining contact with the endcap. We also ensured that the palm of the hand never touched the spring during the experiment or over the range of normal thumb motion. During trials with vision, subjects could move their head so that they could see the entire spring and their thumb from a self-selected comfortable viewing angle much as in everyday object manipulation.

We provided subjects with audio feedback about the vertical compressive spring

force. The audio feedback was in the form of a loud and clearly audible 500 Hz tone that decreased in volume as the vertical compressive spring force increased. We calibrated this inverse linear relationship between the tone volume and vertical compressive spring force so that none of the subjects ever succeeded in making the tone faint enough to be completely inaudible.

2.2.2 Experimental protocol

We instructed subjects to,

“Slowly compress the spring only using your thumbpad to make the tone volume as faint as possible (i.e., maximize vertical compressive spring force) without causing the spring to slip. Once you have reached the point where you cannot decrease the tone volume without letting the spring slip, hold the compressive load so that the tone volume, although now faint, remains constant and slowly release the spring after 10 seconds. It does not matter if the spring bends or oscillates, it only matters that the volume stay constant once you reach the minimum attainable volume and that the spring does not slip.”

We accepted only those trials where the loading/unloading rate was less than 5 N/s as successful trials to prevent any postponement in the onset of instability (Berglund and Gentz, 2002) and to differentiate between trials where the spring slipped from successful control of the spring’s instability.

2.2.3 Metric of performance: F_s

We defined the performance of subjects for each compression as the mean compressive spring force that we measured during the sustained hold phase. We called the hold phase as “sustained” when the coefficient of variance ($\text{COV} = \text{standard deviation} / \text{mean}$) of the compressive spring force was less than 5% over the entire hold of 7 seconds. F_s was the largest mean compressive spring force during the sustained hold phase of 7 seconds during trials when there was no slip. For example, if the subject held the spring for 10 s, there are 10,000 data points (sampled at 1000 Hz). First, we calculated the COV for every continuous 7,000 sample segment (7 seconds, there will be 3001 such segments) and discarded all but those with $\text{COV} \leq 5\%$. We then calculated the mean compressive spring force for the remaining segments (with $\text{COV} \leq 5\%$) and defined the maximum of the calculated mean values as F_s . We used the three largest values of F_s measured from ten attempts at compressing the spring per treatment condition per subject as repeated measures for our statistical analyses. We provided rest for at least 1 minute after every five compressions to ensure that fatigue did not affect our measurement. Keep in mind that slow compression and slow release at the end of the hold phase are essential to prevent postponing the spring buckling (e.g., Berglund and Gentz, 2002) and to help distinguish a slip from a well-controlled release, respectively.

2.2.4 Experiments to test for effect of training

We conducted the experiment over two days for each subject. On the first day, subjects underwent a training phase of 100 compressions to learn to minimize the tone volume and hold for 10 seconds. The subjects had their eyes open and normal thumbpad sensibility throughout the training phase. We measured their performance before training with normal visual and thumbpad sensibility.

On the second day, we measured subjects' performance with normal thumbpad sensation, both with and without vision. We also carried out trials after a digital nerve-block where subjects lost all sensation from their thumbpad. See Section 3.3.1 in Chapter 3 on Page 66 for more details on how we blocked vision and thumbpad sensation. Since, the results presented in this chapter do not test any hypotheses related to sensory occlusion, we will not go over the details of how sensory occlusions were achieved. However, the statistical analysis of variance (ANOVA) performed will include all available data so that we control for overall Type-I error rate as discussed previously in Section 2.2.

We measured the maximum isometric force that subjects could produce in two postures, namely key and opposition pinch postures (same procedure as in Valero-Cuevas et al., 2003b) using a pinch meter before the start of the spring compression trials on both the first and second day. We picked the largest reading of three attempts as the pinch strength of the subjects. We gave enough rest (at least 2 minutes) at the end of the strength measurement to ensure that fatigue induced due to maximal voluntary force production did not affect our experiment. Subjects were always allowed extra rest, more than what we provided by default, if they asked for it.

2.2.5 Safety-margin

To see if there was a consistent safety-margin used by the subjects, i.e., if they reached a compressive spring force consistently shy of the force at which the spring slips when asked to maximize compressive spring force, we compared the values of F_s and $F_{\max}$ (maximum compressive load) between trials when the spring slipped vs. when it did not. We defined F_s for the trials where the spring slipped as the supremum of the mean value of compressive spring load with $\text{COV} \leq 5\%$ for a duration of at least 3s, since the spring often slipped well before 7s elapsed.

2.2.6 Analysis of endcap rotation

To analyze endcap rotation, we used the three markers on the top endcap and a motion capture system to determine the unit normal vector to the endcap as subjects held the spring at a sustained load for 7s (see Figure 2.3). The 3D Cartesian coordinates of motion of the tip of this unit vector unambiguously describe endcap rotation. We performed principal components analysis on the 3D motion data of the tip of the endcap's unit normal vector to test the validity of modeling our

task using a one-dimensional dynamic equation for endcap rotation dynamics. We transformed the motion along each principal component to an equivalent rotation of the endcap about a constant axis (i.e., “principal” axis corresponding to each principal component) for comparing experimental data to model predictions (see Section 2.3 for details about the model). We used the ratio of the range of endcap rotation along the first two principal components to quantify how well a single fixed axis fits the endcap rotation. For the subsequent discussion, we will refer to this ratio as the “eccentricity” of endcap rotation. We also tested to see if the range of rotation along the second and first principal components systematically changed with compressive spring force as predicted by the subcritical pitchfork bifurcation normal form equation.

2.2.7 Statistical analyses

We manipulated 6 independent variables, namely,

1. Training (measurement before or after training)
2. Available sensory modalities (absence or presence of vision/thumbpad sensation).

The dependent variable for all our statistical analyses was F_s , but for the safety-margin analyses, where $F_{\max}$ was an additional dependent variable. In all, we had six treatment levels, namely: (i) Day 1, before training, (ii) Day 1, after training, and (iii)-(vi) Day 2, all four combinations of presence or absence of vision/thumbpad sensation. We set $\alpha = 0.025$ as our threshold for significance for all our statistical analyses, since we performed two ANOVAs using this data set: (i) effect of the above stated six treatments, and (ii) effect of slip vs. no-slip for all six treatment levels. The numerical value of α defines the acceptable probability of Type-I errors. Although we present only part of the results in this chapter, by including all our data in a single ANOVA, we can more effectively reduce the overall probability of reaching incorrect statistical conclusions. We also performed a multiple regression analysis to test whether F_s is dependent on pinch strength using the independent data set of pinch strengths for the same subjects.

First, we performed a one-way repeated measures ANOVA with subjects as random factors, the six treatment levels as fixed factors and F_s as the dependent variable. We carried out three planned comparisons as post-hoc tests if the ANOVA was significant for finding the effect of training.

Second, we performed a one-way repeated measures ANOVA with subjects as random factors, the six treatment levels as fixed factors and the paired difference in F_s or $F_{\max}$ between no-slip compressions and slipped compressions as the dependent variable. If the ANOVA was significant, we carried out hypothesis tests to determine if the difference in F_s or $F_{\max}$ was significantly different from 0 for each treatment level. This way of carrying out the analysis once again ensures that we

control for overall Type-I error rate, which would not be controlled if we were to do a series of paired t-tests.

Third, we used an ANCOVA as an equivalent family of regressions to test dependence of F_s on both key and opposition pinch strength for at least one of the six treatment levels. We framed this analysis as an ANCOVA with subjects as random factors, the six treatment levels as fixed factors, key and opposition pinch strengths as covariates and F_s as the dependent variable. We obtained three separate p-values from this ANCOVA, one each for the significance of (i) the effect of treatment level (a test of intercept of the ANCOVA), (ii) slope of the key pinch strength as a covariate, and (iii) slope of the opposition pinch strength as a covariate. Since a mixed model such as the one we used does not provide R^2 for the model, we used a regular fixed factor ANCOVA to determine approximate model R^2 .

We tested to see if the mean eccentricity of the endcap rotation data was statistically different from 1 (i.e., equal range of motion along each principal component) using an 2-sided t-test.

We then found the best-fit nonlinear regression curve for the relationship between range of endcap rotation along the first principal component vs. compressive spring force (F_s) and compared this regression curve to our normal form model's prediction. Specifically, we estimated the parameters $\hat{\alpha}$, $\hat{\beta}$ and K in Equation (2.8) that is shown on Page 49. This equation is our normal form model's prediction of the relationship between the domain of attraction (i.e., range of endcap rotation) and F_s .

Finally, we used a linear regression between range of endcap rotation along the second principal component and F_s to test whether it is valid to approximate the endcap's rotational dynamics using a single fixed axis. A slope that is significantly different from 0 would invalidate our claim of one-dimensional dynamics at the edge of instability.

We verified necessary assumptions for the validity of each ANOVA/t-test we performed, namely normality and identical distribution of the residuals. We verified normality by generating normal probability plots for the residuals and found them to lie very close to the normal line. We had to log-transform the eccentricity data to make them normally distributed and carry out the t-test on the transformed data. By plotting the residuals against predicted values, we ensured that the residuals were identically distributed (in the statistical sense of identical distributions). Hence, the parametric statistical analyses we carried out are justified. We used SAS (SAS 9.1 for Windows, SAS, Cary, NC) to perform all statistical analyses.

2.3 Mathematical model of manipulation at the edge of instability

To test our primary hypothesis that a low-order nonlinear dynamic model can capture the dominant spatiotemporal dynamics of manipulation at the edge of instability, we modeled the task using a subcritical pitchfork bifurcation normal form. We calculated model parameters using frictional properties of the thumbpad-endcap interface, known mechanics of a buckling spring, and experimentally measured ‘best-performance’ (maximal value of F_s).

2.3.1 Why a subcritical pitchfork bifurcation model?

We will attempt here to provide detailed rationale for the specific choice of a subcritical pitchfork bifurcation normal form. We will then test the primary hypothesis of this chapter by comparing some predictions of our model with experimental data.

It is important to note that we used a subcritical pitchfork bifurcation normal form to describe the spring endcap’s rotational (torsional) dynamics and not the lateral deflection of the endcap. Typically, most treatises on spring (or column) buckling choose to model the dynamics of the lateral deflection of the spring’s top endcap and not the rotation of the top plane like we do. Hence, it is only meaningful to treat our subcritical pitchfork bifurcation model as a working hypothesis that we explicitly validated using experimental data (see Section 2.2.6 on Page 34). In that spirit, the rationale below is really the motivation for selecting a subcritical pitchfork bifurcation as a working hypothesis and not a justification in of itself for our model choice.

Codimension-1 bifurcations of a buckling spring

With our experimental design involving a buckling spring, we are concerned only with codimension-1 bifurcations of a fixed point³ of the spring buckling dynamics. We will provide here the rationale for why we are interested in bifurcations of fixed points and also a heuristic definition of codimension of a bifurcation.

We restricted our mathematical analysis (and hence our model choice) to bifurcations of fixed points, i.e., how stable and unstable equilibria change in the system as a function of one or more independent parameters. The rationale behind this lies in experimental results and mechanics of spring buckling when loaded at its end (Croll and Walker, 1972; El Naschie, 1990; Lobas et al., 2002; Lobas and Lobas, 2004). A slender spring or column buckles either through a pitchfork or a Hopf bifurcation. A pitchfork bifurcation occurs when a specific configuration of the spring (typically, the straight configuration) becomes unstable beyond a

³Fixed point is just another name for an equilibrium point of the system, i.e., a stationary point in the phase space of the system.

critical load and the spring bends or deflects away from the previously stable configuration. A Hopf bifurcation induces the spring or column to exhibit sustained oscillations. In our experiments where subjects compressed the spring, we never observed sustained periodic oscillatory behavior. Thus, we selected a pitchfork bifurcation normal form model.




Figure 2.4: Schematic showing a Hopf (left) and pitchfork (right) bifurcation for a buckling spring. The figure on the left depicts spring buckling under a follower load leading to self-sustained oscillations. The bottom end-condition is ‘fixed’ in both cases, i.e., the bottom is not allowed to laterally or torsionally deflect. The top end-condition is ‘free’ ($\alpha = 2$), i.e., the top can laterally and torsionally deflect freely. On the right, the spring undergoes a pitchfork bifurcation under a dead-load. The top end-condition is ‘guided’ ($\alpha = 1$), i.e., the top can laterally deflect freely, but torsional deflections are not permitted.

The distinction between these two qualitatively and patently different types of buckling are depicted in the schematic shown in Figure 2.4. On the left the spring is shown to lose stability and buckle via a Hopf bifurcation, where beyond a critical load, the spring shows sustained oscillations. One simple loading method that will result in a Hopf bifurcation is when the load on the top of the spring is always directed perpendicular to its top surface, i.e., tangential to the spring’s axis at its

top end. This type of load is called a ‘follower load’ or a ‘non-conservative load’ and this form of oscillatory buckling is also called as ‘flutter’. On the right, is an example of the spring laterally buckling by a pitchfork bifurcation beyond some critical load. Note that in the illustration, the horizontal orientation of the endcap remains stable and does not undergo any bifurcation. Here, the top of the spring has deviated from the centerline (the desired, pre-buckled stable configuration) to a new stable configuration. This form of buckling happens when an increase in the load on the spring causes a supercritical pitchfork bifurcation (see below for an explanation of ‘supercritical’). One possible loading method that will result in a pitchfork bifurcation is a load pointed in a constant direction, say, vertically down as depicted in Figure 2.4. This type of load is also referred to in the literature as a ‘dead-load’ or ‘conservative load’ and this type of buckling is sometimes known as ‘static buckling’.

We will now attempt to relate this discussion to our earlier discussion on end conditions in Section 2.1.2 on Page 24. As described above, the figure on the right in Figure 2.4 depicts a pitchfork bifurcation for the spring buckling by lateral deflection. In the lateral deflection pitchfork bifurcation of a spring shown above, the endcap remains horizontal. So, if we examine the dynamics of the endcap rotation, the endcap rotation is still “stable” or unbuckled according to our definition of torsional/rotational buckling. In other words, the end condition corresponding to the purely lateral buckling described above (spring on the right in Figure 2.4) is fixed-guided — the base is ‘fixed’, i.e., neither lateral deflection nor rotation is permitted and the top is ‘guided’, i.e., lateral deflection is permitted, but not rotation. The spring profile for this mode of buckling is described by,

$$\delta(x) = 1 - \cos\left(\frac{\pi x}{L}\right) \quad (2.3)$$

where, L is the length of the spring (or column) and x is the axial location. The end condition parameter for this mode of buckling is $\alpha = 1$.

The papers by Lobas et al. (2002) and Lobas and Lobas (2004) provide a more detailed exposition of pitchfork vs. Hopf bifurcation in a buckling column (or spring) using a low-order approximation. They use an inverted double pendulum with stiff pivot joints as a discretization of a slender column and show how under a dead-load, it undergoes a pitchfork bifurcation and under a follower load, it undergoes a Hopf bifurcation. Note that the discussion in the papers by Lobas et al. (2002) and Lobas and Lobas (2004) focus only on the lateral deflection and not rotation of the top of the inverted double pendulum.

After this discussion on stability and transitions in stability for fixed points, i.e., non-oscillatory equilibrium configurations of a spring, I now attempt to provide a brief definition of ‘codimension’ of a bifurcation. The codimension of a bifurcation can be roughly understood as the number of parameters that need to be varied simultaneously in order for the bifurcation to occur. For a more detailed discussion and rigorous definition of codimension of a bifurcation, see Guckenheimer and Holmes (1983). For the case of the buckling spring, a bifurcation, say, loss of

stability of the vertical spring configuration happens by fine-tuning the vertical compressive load on the spring — one parameter. Assuming axisymmetry of the spring the only types of codimension-1 bifurcations for lateral deflection dynamics that are possible are either a pitchfork or a Hopf bifurcation. This is a well known fact about elastic buckling of slender columns under end-point loads (Croll and Walker, 1972; El Naschie, 1990). In terms of the endcap rotation dynamics, a Hopf bifurcation is characterized by self-sustained oscillatory behavior, i.e. the horizontal endcap orientation ($\theta = 0$) becomes unstable and self-sustained oscillations ensue. On the other hand, a pitchfork bifurcation is characterized by loss of stability of $\theta = 0$ accompanied by the appearance of two other symmetric stable equilibria.

Supercritical vs. subcritical bifurcations

Both pitchfork and Hopf bifurcations can be of two types — supercritical or subcritical. We will for now use the endcap rotational dynamics (θ) as an example to help elucidate the differences between these two types of bifurcations. Consider first the example of a pitchfork bifurcation. An approximate way to differentiate between the two types is in terms of the abruptness of change in the dynamics of the system once $\theta = 0$ becomes unstable. A supercritical pitchfork bifurcation is “continuous” in the sense that once $\theta = 0$ becomes unstable, two new symmetric stable fixed points first appear with no separation between them and grow apart continuously (see the bifurcation diagram in Figure 2.5). On the other hand, a subcritical pitchfork bifurcation is “abrupt”. Two stable fixed points with a finite (non-zero) separation appear even before $\theta = 0$ becomes unstable. Once $\theta = 0$ becomes unstable, the two already existing stable fixed points with a large separation between them cause an abrupt jump in the corresponding bifurcation diagram (see Figure 2.6). For the Hopf bifurcation, just replace the symmetric stable fixed points by a single stable limit cycle in the above description to understand the difference between the supercritical and subcritical types.

Rationale for selecting a subcritical pitchfork bifurcation for endcap rotational dynamics

We now turn our attention to the rotation of the top plane of the spring in an attempt to better define our working hypothesis of a subcritical pitchfork bifurcation model for the spring’s endcap rotational dynamics. For convenience, and without loss of generality, we define the unbuckled, desired rotational orientation as $\theta = 0$ rad.

From visual examination of the endcap’s motion data it was clear that there were no sustained oscillations of the spring. This rules out the possibility of a Hopf bifurcation (supercritical or subcritical). Another qualitative feature of the loss of stability, i.e., how the spring slips, was the abrupt and drastic rotation of the endcap just before it slips. Finally, visual examination of the spring’s profile during the experiments resembled ‘fixed-guided’ end-conditions, much as shown




Figure 2.5: The bifurcation diagram for the normal form of a supercritical pitchfork bifurcation. The normal form equation is $\dot{\theta} = \lambda\theta - \theta^3$, where λ is the parameter that is tuned, leading to a supercritical pitchfork bifurcation at $\lambda = 0$. As evident from the bifurcation diagram, the two new symmetric stable fixed points that appear for $\lambda > 0$, make their first appearance with no separation at all between them, leading to a “continuous” change in fixed points of this system as λ is changed.




Figure 2.6: The bifurcation diagram for the normal form of a subcritical pitchfork bifurcation. The normal form equation is $\dot{\theta} = \lambda\theta + \theta^3 - \theta^5$, where λ is once again the parameter that is tuned, leading to a subcritical pitchfork bifurcation at $\lambda = 0$. The quintic term is often not included in the normal form, however, its inclusion provides more realistic bounded behavior of trajectories in the system, since without the quintic term, the outer stable fixed points (not $\theta = 0$) would be absent. As evident from the bifurcation diagram, for $\lambda > 0$ the loss of stability of the fixed point at $\theta = 0$ is not accompanied by the creation of any new stable fixed points near $\theta = 0$. This creates an “abrupt jump” in the location of stable fixed points.

on the right in Figure 2.4 on Page 38. An ideal ‘guided’ end-condition at the top endcap would prevent any rotational buckling. However, in a real-world system, the ideal ‘guided’ end-condition would probably resemble a stiff torsional spring (at best) that is attached to the top endcap, leading to both lateral and rotational buckling as shown in Figure 2.7.

Lateral vs. rotational buckling




Figure 2.7: An illustration of spring buckling under a dead-load that depicts a non-ideal guided end-condition resulting in both lateral and rotational buckling. Importantly, during experimental trials with human subjects the lateral buckling happened at a lower load than the rotational buckling.

Importantly, during experimental trials with human subjects, the lateral buckling had already happened before the rotational buckling as the subjects compressed the spring. Hence, the rotational buckling is a ‘secondary’ buckling phenomena and it is this secondary buckling that we chose to model. This observed spring profile, together with the former two observations on the endcap rotational dynamics, strongly suggest that it is appropriate to model the dynamics of the closed loop thumb + spring + nervous system as a subcritical pitchfork bifurcation. This empirical observation leads to a revealing proposal. A spring (or column) undergoes a Hopf bifurcation only when the endpoint force is always directed close to the endcap normal (Croll and Walker, 1972; El Naschie, 1990). Since we never observed any sustained oscillations, thus ruling out a Hopf bifurcation, we can

postulate that the strategy used by our subjects to maximize compressive force in the spring was to somehow stabilize $\theta = 0$. Keep in mind that $\theta = 0$ is only a normative model and that in the real experiment, subjects could have stabilized $\theta = \theta_0$ (where $\theta_0 \neq 0$) and the above rationale will be identical in all other regards.

2.3.2 Endcap rotational buckling as a secondary bifurcation

A comment is called for at this point about the relationship between the lateral buckling and the rotational buckling that happen in the spring. As discussed above, the rotational buckling is a secondary buckling that happens at higher loads than the lateral buckling. Hence, as the spring is compressed, the spring would already be laterally post-buckled when the endcap rotational dynamics reach the edge of instability. The laterally buckled spring would clearly introduce an asymmetry in the rotational dynamics. This asymmetry can be captured using an additional ‘imperfection parameter’ in the subcritical pitchfork bifurcation normal form equation, called as a versal unfolding of the pitchfork bifurcation normal form (Guckenheimer and Holmes, 1983):

$$\dot{\theta} = \lambda\theta + \theta^3 - \theta^5 + \epsilon \quad (2.4)$$

However, this adds an additional parameter to our model that we cannot estimate without a more detailed model of spring buckling that can capture both the lateral and rotational buckling. Hence, for the present study, we will stick to the symmetric (without imperfection) subcritical pitchfork bifurcation normal form and leave the inclusion of this additional parameter to a future study, which is beyond the intended scope of this thesis. Nevertheless, we present a more detailed exposition on the effect of including an imperfection parameter in Section 5.1.1 on Page 90.

2.3.3 Subcritical pitchfork bifurcation model

Based on the rationale stated above, we modeled the dynamics of the entire system as a subcritical pitchfork bifurcation in the dynamics of θ , the rotation of the endcap with respect to the horizontal orientation. The subcritical pitchfork bifurcation model for the dynamics of the endcap rotation is,

$$\dot{\theta} = \alpha(F_s - K)\theta + \beta\theta^3 - \gamma\theta^5 \quad (2.5)$$

where θ represents the rotation angle of the endcap from the desired orientation. As discussed above, we will henceforth assume that the orientation of the endcap that subjects tried to stabilize was $\theta = 0$, without loss of generality. The parameter F_s is the compressive force. Note that it has been modeled as a constant parameter with no temporal dynamics of its own. This reflects the experimental protocol that the experimentally measured F_s is the sustained (near-constant) compressive spring force. K is the maximum value of F_s for which $\theta = 0$ is stable in this model, i.e., for

$F_s > K$, the orientation $\theta = 0$ becomes an unstable fixed point. The parameters α , β and γ are positive parameters that “stretch” or scale the normal form equation appropriately to reflect the specific spatiotemporal instantiation of the subcritical bifurcation we are interested in. In other words, these three parameters depend on the physics of the specific problem we are examining.

The stable and unstable equilibria⁴ for the dynamical system governed by Equation (2.5) are given by setting the right-hand side of Equation (2.5) to zero, i.e.,

$$\theta (\alpha(F_s - K) + \beta\theta^2 - \gamma\theta^4) = 0 \quad (2.6)$$

to obtain,

$$\therefore \theta^* = \begin{cases} 0 : \text{stable if } F_s < K \text{ and unstable if } F_s > K \\ \pm \sqrt{\frac{\hat{\beta}}{2} + \sqrt{\left(\frac{\hat{\beta}}{2}\right)^2 + \hat{\alpha}(F_s - K)}} : \text{stable; } K - \frac{1}{\hat{\alpha}} \left(\frac{\hat{\beta}}{2}\right)^2 \leq F_s \\ \pm \sqrt{\frac{\hat{\beta}}{2} - \sqrt{\left(\frac{\hat{\beta}}{2}\right)^2 + \hat{\alpha}(F_s - K)}} : \text{unstable; } K - \frac{1}{\hat{\alpha}} \left(\frac{\hat{\beta}}{2}\right)^2 \leq F_s \leq K \end{cases} \quad (2.7)$$

where, θ^* are the fixed points of Equation (2.5), $\hat{\alpha} = \alpha/\gamma$ and $\hat{\beta} = \beta/\gamma$. For specific parameter values of K , $\hat{\alpha}$ and $\hat{\beta}$, the dependence of θ^* on F_s and their stability is shown in Figure 2.8 on Page 46. The figure shows that a finite ‘domain of attraction’ exists around $\theta = 0$ that shrinks as F_s increases until the point of extinction at $F_s = K$. This is the predicted scaling with F_s of allowable range of rotation of the endcap so that the spring does not slip. We tested this using our experimental data as already outlined in Section 2.2.6.

2.3.4 Parameter values for the low-order model

We chose $K = 3.3\text{N}$, equal to the experimentally measured maximum value of F_s that our subjects attained. The spring slips only when the angle of the endcap with respect to the horizontal exceeds the friction angle (~ 0.5 radians). We solved for $\hat{\alpha}$ and $\hat{\beta}$ so that the average of the undesired equilibria (that represents friction angle; solid red curve in Figure 2.8 on Page 46) for the range of experimentally observed F_s was equal to 0.5 and the least value of this representative friction angle was 0.4 (80% of the friction angle). We found $\hat{\alpha} = 0.006855$ and $\hat{\beta} = 0.2766$. Note that the bifurcation diagram shown in Figure 2.8 and the equilibria given by Equation (2.7) do not depend on the exact values of α , β and γ , but only on the ratios, α/γ and β/γ . Therefore, we will postpone the determination of the numerical value of γ , hence α and β to Chapter 3.

⁴The stability of any equilibrium point of the dynamical system described by Equation (2.5) can be found by easily by looking at the direction of flow of trajectories in Figure 2.6.




Figure 2.8: This bifurcation diagram shows the loci of both stable (green solid curves) and unstable (red dashed curves) equilibria for the thumb+spring+nervous system’s closed loop dynamics without noise or time-delays as the spring is compressed. This figure is a succinct description of the underlying deterministic (no time-delays / noise) dynamics of the subcritical pitchfork bifurcation’s normal form equation. Say, $F_s = 2.0$ N, there is a region (shown in cyan, bounded by the red dashed curves) around $\theta = 0$ (the endcap orientation we want to stabilize), in which the endcap is attracted toward $\theta = 0$. If the endcap strays too far outside this region, then it will be rapidly attracted toward the points far out (solid red curves, close to 0.5 rad), which is representative of a spring slip. In other words, the cyan region in the center of the figure bounded by red dashed curves on either side is the “domain of attraction” for the closed-loop system at $F_s = 2.0$ N.

2.4 Results

The results presented below are a comparison of experimental data (see Figure 2.3 on Page 29 for the experimental setup) with the one-dimensional mathematical model shown in Equation (2.5) on Page 44 and Figure 2.8 on Page 46. It is also important to remember that “performance” is quantified by F_s and a lower value of F_s necessarily means a less challenging task that is more stable in the sense of being more tolerant to external perturbations (from the thumbpad itself or otherwise).

2.4.1 Dimensional collapse at the edge of instability

Three types of analyses were performed to test our first hypothesis that the spatiotemporal dynamics of the endcap rotation in our experimental data are well represented by a one-dimensional nonlinear dynamic model (Equation (2.5) on Page 44 and Figure 2.8 on Page 46). For this analysis, data from all trials on Day 2 (with/without vision/thumbpad sensation) were pooled together, since if our claim that the low-order model can capture the dominant dynamics of manipulation at the edge of instability is true, it must be able to do so independently of the sensory channels available.

Over 90% of endcap rotation was along a single axis

We found that dynamics at the edge of instability were nearly one-dimensional, thus experimentally validating the use of a low - order normal form. We used principal components analysis (PCA) on the Cartesian coordinates of the unit normal vector to the spring’s endcap to determine if the rotation of the endcap could be meaningfully described by just one variable (endcap rotation from horizontal, i.e., the elevation angle). Keep in mind that even if a single variable suffices, PCA cannot estimate the order of the dynamics of this variable, e.g., whether a first- or second-order differential equation is a better fit. PCA of the spring’s endcap rotation revealed that the system’s dynamics was indeed restricted primarily along one dimension, partially justifying⁵ our choice of the one-dimensional subcritical pitchfork bifurcation normal form. The first principal component alone explained 94.5% of the variance in motion of the unit normal vector (99.9% confidence interval = [91.2%,96.6%]). Physically, this means that rotation about a single fixed axis was sufficient to describe almost 95% of the rotation of the endcap when at the edge of instability. A weaker, but physically meaningful way to test the one-dimensional nature of the endcap rotation is to examine the “eccentricity” of the motion of the unit normal vector. We defined “eccentricity” as the ratio of the range of motion of the unit normal vector along the first vs. second principal components. Mean eccentricity was 2.93 (mean > 1, $p < 1e - 6$, Figure 2.9). Hence, not only was the

⁵It is only a partial justification since we are yet to see the results of analyzing how the range of endcap rotation scaled with F_s .

variance predominantly along one direction, but the complete range of motion was also predominantly along one direction, albeit to a smaller amount since a single outlier can bias the eccentricity drastically.




Figure 2.9: Eccentricity is the ratio of range of endcap rotation projected along the first vs. the second principal component. Mean eccentricity ($= 2.93$) was statistically significantly greater than 1 ($p < 1e-6$) as shown by its 99% confidence interval. This provides preliminary support for our modeling approach since the endcap dynamics seem to be restricted primary along one dimension.

Endcap rotation along first principal component agrees with model predictions

The endcap dynamics did not show any sustained periodic oscillations and we captured all the relevant spatiotemporal dynamics of the endcap that depend on the compressive spring force F_s by restricting our model to the first principal component. The data-range for endcap rotation along the first principal component agreed with our subcritical pitchfork bifurcation normal form. For a given value of F_s , the central region around $\theta = 0$ bounded by the red dashed curves in Figure 2.8 on Page 46 (e.g., the blue region at $F_s = 2.0$ N in Figure 2.8) is the domain of attraction of $\theta = 0$, i.e., it is the normal form's prediction for the data-range for endcap rotation that will not destabilize the system at that value of F_s . The

relationship is readily found using Equation (2.7) on Page 45 to be,

$$\theta_{\text{range}} = 2\sqrt{\frac{\hat{\beta}}{2} - \sqrt{\left(\frac{\hat{\beta}}{2}\right)^2 + \hat{\alpha}(F_s - K)}} \quad (2.8)$$

where, $\hat{\alpha}$, $\hat{\beta}$ are parameters that we obtained using the frictional constraints of the problem and K represents the maximum load that subjects could reach.

To test the validity of our normal form model, we independently estimated $\hat{\alpha}$, $\hat{\beta}$ and K from experimental data using a non-negative nonlinear least squares regression. The experimentally obtained dependence of the range of endcap rotation on F_s along the first principal component (solid blue regression curve in Figure 2.10, $\hat{\alpha}_{\text{estimate}} = 0.0017$, $\hat{\beta}_{\text{estimate}} = 0.11$, $K_{\text{estimate}} = 3.45$, $R^2 = 0.32$) closely resembled our normal form model's prediction (dashed red curve in Figure 2.10, $\hat{\alpha}_{\text{estimate}} = 0.0069$, $\hat{\beta}_{\text{estimate}} = 0.28$, $K_{\text{estimate}} = 3.3$).




Figure 2.10: The dynamics (range of endcap rotation) along the first principal components is nearly identical to model prediction for a subcritical pitchfork bifurcation (see Figure 2.8 on Page 46). The low R^2 means that the system is very noisy, but the highly significant slope supports the reliability of the regression fit.

Endcap rotation along second principal component showed no association with F_s

Our model is further validated by the result that a linear regression of the range of endcap rotation along the second principal component showed no association

with F_s (Figure 2.11; slope = 0, $p = 0.96$, $R^2 \approx 0$). That is, variance in range of rotation of the endcap along the second principal component cannot be explained by variance of the compressive spring force F_s .




Figure 2.11: The lack of association between F_s (slope indistinguishable from 0) and the dynamics (range of endcap rotation) along the second principal component further supports our one-dimensional model.

2.4.2 Strength independent limit of sensorimotor performance

The variability in the sensorimotor ability to sustain maximal compression of the spring across subjects was much smaller than variability in hand strength. Moreover, the limit of sensorimotor ability was independent of hand strength. We measured maximal isometric key and opposition pinch strength in the subjects using a pinch meter. We found that the coefficient of variation ($\text{COV} = \text{mean}/\text{standard deviation}$) of pinch strength (mean = 99.41 N, $\text{COV} = 13.8\%$) was almost three times greater than that of for normal sensibility on day 2 (mean = 2.99 N, $\text{COV} = 5.2\%$). Note that F_s is $< 5\%$ of pinch strength. Based on a non-significant linear multiple regression analysis of F_s for all treatment conditions on day 2 vs. static pinch strength (key and opposition; $p = 0.969$, and $p = 0.338$ respectively, multiple regression $R^2 = 0.66$), we conclude that thumb strength did not correlate with F_s .

2.4.3 Lack of measurable safety-margin

Subjects did not maintain any consistent “safety-margin” while reaching for their maximal sustainable F_s , indicating that they were truly at their sensorimotor limit when sustaining maximal compression of the spring. When we compared the change in performance between trials when the spring slipped and when it did not for all four sensory occlusion conditions we found no significant difference in F_s between such trials ($p > 0.530$, Table 2.1). We also compared the difference (change) in maximum compressive spring force ($F_{\max}$) between successful compressions and unsuccessful (slipped) compressions to account for the fact that the endcap might have slipped because of a rapid increase in compressive force just before slipping. Once again, we found no significant difference in $F_{\max}$ between slipped and non-slipped trials ($p > 0.197$, Table 2.1).

Table 2.1: Statistical post-hoc planned comparisons of F_s (sustained load) and $F_{\max}$ (maximal load) for trials when the spring slipped vs. when the spring did not slip. The above values are all for Day 2. These results indicate that subjects did not maintain a “safety margin” during the successful (i.e., no-slip) performance of this dynamic manipulation task. The ANOVA across conditions, was significant with $p < 0.0001$. Note that $\Delta F_{s \text{ or } \max} = F_{s \text{ or } \max}^{\text{slip}} - F_{s \text{ or } \max}^{\text{no-slip}}$ in the table.

Subject Condition	ΔF_s in N	p-value	$\Delta F_{\max}$ in N	p-value
Normal	0.01	0.912	0.03	0.708
No vision	0.15	0.530	-0.07	0.688
Nerve-blocked	0.05	0.667	0.12	0.197
No vision, Nerve-blocked	-0.01	0.900	0.03	0.700

2.4.4 No measurable effect of training

There was no significant learning effect after rehearsing the task numerous times over two days. All sensory channels were intact throughout the training phase as well as during performance measurements made before and after training. On day 1, we tested subjects before and after 100 trials of training. On day 2, we did not give the subjects any additional training before measuring their performance. We found that training on day 1 did not cause any significant increase in F_s ($\Delta F_s^{\text{Day1}} = F_s^{\text{post-train}} - F_s^{\text{pre-train}} = 0.08 \text{ N}$, $p = 0.248$). The performance before any additional training on day 2 was once again not statistically different from either before or after training on day 1. At the start of day 2, F_s was on average, in between performance on day 1 (greater by 0.05 N and lower by 0.03 N than the start and

end of training on day 1, respectively). See Figure 2.12 for a summary of these results.




Figure 2.12: Box-plot showing the effect of training. On Day 2 we measured the initial performance before any additional exposure to the task, labeled as “Initial” in the figure above. The horizontal mid-line shows the sample median, the notches its robust 95% confidence interval and the box itself shows the 25th and 75th percentiles. The whiskers show the minimum and maximum values to indicate any outliers on this plot. None of the differences in F_s calculated using planned comparisons from repeated measures ANOVA is significant at a significance level of 0.025.

2.5 Discussion

Our novel use of bifurcation theory of nonlinear dynamical systems allowed the creation of a mechanics-based low-order model whose spatiotemporal dynamics was indistinguishable from experimental measurements at the edge of instability for the complex nonlinear behavior of dynamic manipulation. We also found that the sensorimotor limit of performance at submaximal forces was independent of pinch strength. Finally, our data supported our hypothesis that during dynamic manipulation there is no safety-margin.

We found that the spatiotemporal dynamics of the endcap rotation lay primarily along one spatial dimension at the edge of instability, indicating that the dominant dynamics at the limit of performance had undergone a dimensional collapse as predicted by bifurcation theory. This enabled the use of a low-order normal form to model this high-dimensional system of dynamic manipulation. Two other important results strongly reinforce the validity of this novel modeling approach. (i) The dependence of the range of motion along the first principal component on compressive force was indistinguishable from that of the normal form and, (ii)

the endcap dynamics along the second principal component of the eccentric spatial dynamics had no association with compressive force, indicating that the one-dimensional model was indeed sufficient to capture the dominant dynamics at the edge of instability. It is worth noting the ability of our nonlinear low-order model to reproduce the spatiotemporal dynamics of the task at the limit of performance of a system as complex as dynamic manipulation. Importantly, all parameters were obtained from basic frictional constraints and experimentally observed maximum compressive load.

We found no relationship between dynamical performance (i.e., F_s) and pinch strength. The fact that our experimental paradigm focuses on very low forces ($F_s < 5\%$ of pinch strength) allows us to conclude that we are investigating the limit of hand sensorimotor integration, independently of muscle strength. Importantly, we found that the limit of performance in our task was much more consistent (lower COV) than pinch strength.

The independence from pinch strength, the submaximal nature of the limit of performance (in terms of force production capability), and the consistent limit of performance across people, are all insightful results. For example, in everyday life we sometimes find that even a weak (low muscular strength) individual can be much more “dexterous” than a strong person — a 10 year old could twirl a pen much better than a professional rock climber. On the flip side, with sufficient training, both the 10 year old and the rock climber may become equally adept at twirling a pen using the fingers. This heuristic and lay-person notion of “dexterity” seems to possess (at least qualitatively), the same features as our choice of a definition for dexterity, namely, the ability to regulate fingertip force and motion in order to control an instability, independently of muscle strength.

Subjects did not have a measurable safety-margin when compressing the spring, contrary to the current dominant thinking for static manipulation (Cole and Abbs, 1988; Johansson and Birznieks, 2004; Eliasson et al., 1995; Cole et al., 1998; Johansson et al., 1992a; Augurelle et al., 2003; Monzee et al., 2003). In other words, when instructed to dynamically compress the spring to the best of their ability, subjects did not simply stop short of the force at which the spring had slipped previously. Instead, they compressed the spring to a force that was indistinguishable from the load at which the spring had slipped before. For example, when vision was absent, the mean value of $F_{\max}$ (maximum spring force) was in fact higher than when the spring slipped. This provides strong evidence that subjects were able to detect an imminent slip before actually reaching their maximal instability, possibly using methods similar to those previously reported by other authors in either abstract dynamical systems or physical / chemical systems (Alvarado et al., 1994; Chen et al., 2000; Kim and Abed, 2000; Moreau et al., 2003; Rico-Martinez et al., 2003; Wiesenfeld, 1985; Wiesenfeld and McCarley, 1990; Wiesenfeld and Moss, 1995)⁶. This ability to detect an imminent slip using feedback is vital to ev-

⁶Several of these methods rely on the (often counter-intuitive) response of nonlinear dynamical systems at the edge of instability when subjected to noisy exci-

everyday dynamical sensorimotor function. When walking on a slippery surface (say ice), it is important to know when one is going to slip to prevent a fall, without actually falling down. Similarly, while handling objects and performing dynamic or fine manipulation, it is important to know when the object is prone to fall, without actually dropping it. A more detailed, but speculative discussion of the ramifications of this experimental finding is presented in Chapter 5.

Finally, it is intriguing that training had no effect on performance in healthy adults. However, investigating why we could not measure any effect of learning warrants a separate study by itself and beyond the scope of this thesis. Some possible explanations for the apparent lack of any effect of training are: (i) our task was mechanically so similar to everyday object manipulation that subjects were already fully trained, (ii) training was insufficient to detect improvement in performance, (iii) change in performance was in some aspect of the sensorimotor control strategy that our metric could not detect, or (iv) the changes with training were so slight that trial-to-trial variability in task performance might have washed out any small systematic changes in performance with training. This is by no means an exhaustive list of possible explanations, but can provide a starting point for future studies on sensorimotor development, training and rehabilitation using this dynamic manipulation task.

2.6 Conclusions

The main messages to take away from the results presented in this chapter are the following:

1. The average spatiotemporal dynamics of the task of dynamic manipulation at the edge of instability closely resemble that of a subcritical pitchfork bifurcation. Given the facts that our choice of subcritical pitchfork bifurcation model was motivated by the mechanics of the task and parameters found solely using frictional properties/experimentally measured best-performance, we have strong reason to believe that our model truly reflects the dominant dynamics of manipulation at the edge of instability.
2. The lack of a safety-margin under all sensory occlusion conditions provide evidence that our subjects could detect an impending instability, and more importantly this detection did not rely only on thumbpad sensation or vision. This suggests that sensory channels other than digital sensors (i.e. belonging to the thumb) or vision — which we will refer to henceforth as non-digital sensors — provided useful information about the dynamics of the buckling spring to the nervous system.

tation. Specifically, these studies find characteristic ‘signatures’ for different types of incipient bifurcations by examining the power-spectral density of the specific outputs of the dynamical system when subjected to continuous noisy inputs.

Chapter 3

Task optimal multisensory integration during dynamic manipulation

3.1 Introduction

Why is it that we usually handle objects effortlessly without looking at them, but rely heavily on vision if our fingers are numb? This contextual use of vision is a phenomenon of potentially great clinical and neurological importance¹ that is observed everyday and also in several psychophysical experiments (Cole and Abbs, 1988; Monzee et al., 2003; Augurelle et al., 2003). However, no one has provided any mathematical explanation beyond heuristic word-arguments to explain this phenomenon. This phenomenon begs a mathematical explanation in the framework of the theories of optimal state estimation. Here, we provide a model-based explanation for this apparently selective reliance on vision depending on the availability (or quality) of digital cutaneous sensors.

In this chapter, we present novel results to reveal how task-optimal multisensory integration emerging from the combined effect of sensorimotor noise and time-delays explains this important neurophysiological phenomenon. We start with the same experimental task and mathematical model developed in Chapter 2 and incorporate time-delays and noise into our normal form equation. We saw in the previous chapter how our model’s spatiotemporal dynamics were indistinguishable from experimental measurement. We also reviewed in Chapter 1 on Page 8 that current understanding of multisensory integration in the nervous system is restricted to examining only the effects of noise and cannot address the effect of time-delays. Using the nonlinear dynamic manipulation task at the edge of instability that we developed in Chapter 2, we will present evidence in this chapter that provides support for our hypothesis that hitherto unexplored effects of time-delays on multisensory integration and previously known effects of noise together explain why people rely on vision during dexterous manipulation only when the fingerpads become numb. We will end this chapter with a brief discussion about how this model-based explanation sheds light on why we drop objects more often when we are old or cold!

¹This common observation gains gravity with the realization that several peripheral and central neuropathies such as carpal tunnel syndrome or multiple sclerosis adversely affect the quality of cutaneous receptors. One other common “condition” which is technically not a disease, but can adversely affect lifestyle, is aging. Aging is thought to affect the quality of cutaneous sensors and sensorimotor processing (Cole et al., 1998, and references therein).

3.1.1 Background – recapitulation of Bayesian inference

In the framework of optimal feedback control and for the broad area of sensorimotor control in general, gaining insight into factors that influence multisensory integration during dynamic sensorimotor tasks is very important (Ernst and Bulthoff, 2004; Knill and Pouget, 2004, and references therein). Two main elements that affect the quality of sensory information and hence influence multisensory integration are sensorimotor noise and time-delays. Both noise and time-delays are pervasive in the nervous system (Beuter, 2003; Harris and Wolpert, 1998; Kelso, 1995; van Beers et al., 2002; Wolpert et al., 1995), and it is known that both affect sensorimotor control (Beuter, 2003; Cabrera and Milton, 2002; Collins and Deluca, 1994). Previous studies on multisensory integration (Ernst, 2004; Ernst and Banks, 2002a,b; Ernst and Bulthoff, 2004; Körding and Wolpert, 2004; Knill and Pouget, 2004; Kuo, 1995, 2005; Säfström and Edin, 2004; Sober and Sabes, 2003, 2005; van Beers et al., 2002; van der Kooij et al., 1999; Wolpert et al., 1995) have not revealed or quantified the expected effect of time-delays (Clark and Yuille, 1990), but rather focused on the effect of noise. We previously saw that these studies are restricted to examining the effects of noise either because they used static task goals (see Section 1.1.3) or used optimal estimators that could robustly compensate for time-delays using internal forward models (see Section 1.1.3). Hence, there is a need for studies that can examine the combined effects of noise and time-delays on multisensory integration during dynamic sensorimotor behavior.

3.1.2 Why bother about time-delays?

As we saw in Section 1.1.3, optimal feedback controllers explicitly or implicitly represent the forward dynamics of the system being controlled and hence can compensate for time-delays in sensory feedback. Given the remarkable success of optimal feedback control as a succinct theory to explain human sensorimotor control, a natural question arises, “why bother about time-delays if an optimal feedback controller can compensate for time-delays”? Section 1.1.3 contains a more detailed answer to this question. However, for sake of completeness of the preamble to the work presented in this chapter, we repeat some of the points mentioned earlier.

Time-delays become important in the context of sensorimotor control because: (i) sensorimotor control is regarded highly for its remarkable performance at the limit of performance (at least compared to robotic controllers) and disease is regarded as devastating because of the dramatic loss of ability to reach previously achievable limits of performance, and (ii) given the unavoidable and ubiquitous nature of sensorimotor noise, any time-delay compensation techniques are rendered ineffective at the edge of instability. Although optimal controllers are capable of compensating for time-delays, this ability is not robust to noise when the closed-loop system (controller + body + world) is close to instability (Stein, 2003), which is typically the case when the sensorimotor system is at its limit of performance

(with the exception of static activities like producing maximum force).

3.1.3 Hypotheses

We hypothesized that a task-optimal sensory weighting strategy emerging from the combined effect of noise and time-delays is used by humans during dynamic manipulation of objects. The rationale for this hypothesis rests on theories of Bayesian inference for state estimation in humans, which have been very successful at explaining effects of noise on multisensory integration during static tasks. We tested our hypothesis by measuring the consequence of sensory occlusion (thumbpad sensation/vision) when subjects performed the experimental task of compressing a slender spring that is explained in detail in Chapter 2 on page 31. We then used our model of dynamic manipulation at the edge of instability to test whether numerical optimization to find performance maximizing sensory weights could replicate our experimental results. Additionally, to quantify the hitherto unexplored effects of time-delays on multisensory integration, we compared task-optimal sensory weights (that incorporate the effects of both time-delays and noise) vs. sensory weights that minimize the effect of noise alone.

We proposed and tested this hypothesis to explain the prevalent and neuro-physiologically important observation of the selective use of vision during object manipulation depending on the quality (and presence) of digital cutaneous sensors.

3.2 Computational model

The mathematical model developed in Chapter 2 is a deterministic model with no noise or time-delays. Moreover, it is a model of the thumb + spring + nervous system at the edge of instability, with no clear separation between the constituent elements of this fused dynamic system. However, for testing the consequences of sensory noise and time-delays, three extensions to the model are necessary:

1. **Separation of “controller” (nervous system) from “plant” (thumb + spring).** To achieve this separation, we will make simplifying assumptions as outlined below, while ensuring that the closed-loop dynamics still resemble that of the normal form equation for a subcritical pitchfork bifurcation.
2. **Inclusion of multiple sensory channels.** To incorporate multiple sensory channels, we need some mechanism of combining or ‘integrating’ the signals from multiple sensory channels before using them to apply a control action on the plant.
3. **Inclusion and calculation of sensory noise and time-delays.** Finally, we need to incorporate noise and time-delays in each sensory channel, since that is the central portion of the work presented here. Once again, some simplifying assumptions about the nature of noise and time-delays will be made.

3.2.1 Separating “controller” from “plant”

In a low-order model of the type developed by us, it is not always clear how to separate the controller from the plant. The main reason for that is the distinction between the different approaches to model a system — emergent/dynamical vs. goal-oriented (see Section 1.4 on Page 18). The low-order model of a subcritical pitchfork bifurcation is a close approximation to the dominant dynamics of the entire thumb + spring + nervous system that emerge at the edge of instability. From that point of view, this model does not distinguish between controller and plant. However, theories of sensorimotor control and most descriptions of sensorimotor control in contemporary literature separate the nervous system as a distinct module from the body and the outside world. This leads to sensorimotor control as a goal-oriented behavior, where the nervous system implements strategies for achieving a desired goal through the interaction of the body and the world. Although we are interested in such goal-oriented behavior, we resorted to an emergent dynamical model for the sake of mathematical tractability. This puts us in the quandary of how to separate controller and plant, since there is no unique or obvious/natural way to perform this separation.

First, we present the normal form equation for the subcritical pitchfork bifurcation once again to help read this thesis without frequently turning pages (which incidentally is an incredible achievement of dexterous manipulation!). The deterministic, no time-delay model that we developed in Chapter 2 is,

$$\dot{\theta} = \alpha(F_s - K)\theta + \beta\theta^3 - \gamma\theta^5 \quad (3.1)$$

where, θ represents the rotation angle of the endcap from the desired orientation ($\theta = 0$), $K = 3.3$, $\alpha/\gamma = 0.006855$ and, $\beta/\gamma = 0.2766$. Remember that we are yet to determine the value of γ , since so far, none of the results presented using this model depended on the value of γ .

As a first approximation, we chose to perform this separation by treating K as the representative of the controller. In other words, this equation can be rewritten as,

$$\dot{\theta} = \overbrace{\alpha F_s \theta + \beta \theta^3 - \gamma \theta^5}^{\text{thumb + spring dynamics}} + \overbrace{u(\hat{\theta})}^{\text{feedback}} \quad (3.2)$$

$$u(\hat{\theta}) = \alpha K (\theta_{\text{desired}} - \hat{\theta}) \quad (3.3)$$

$$\theta_{\text{desired}} = 0 \text{ without loss of generality}$$

where, $\hat{\theta}$ is the sensed angle of the spring’s endcap. Note that we recover our original normal form Equation (3.1) if $\hat{\theta} = \theta$. In Equation (3.3), K represents “control strategy” in response to the perceived error in endcap orientation, i.e., to a first approximation, the complicated action of the nervous system through the thumbpad is assumed to resemble constant proportional feedback². If there were

²A basic concept obtained from automatic feedback control theory is the enor-

no time-delays or noise and if perfect (unbiased, i.e., accurate on average) state estimation were possible, then $\hat{\theta} = \theta$. However, in a real sensory system $\hat{\theta}$ is noisy and time-delayed at the least and could also be biased (different from θ even on average). Below we will represent $\hat{\theta}$ as a noisy, time-delayed, unbiased estimate of θ by combining multiple sensory channels.

A note on the assumption of constant proportional feedback is warranted here. A little thought reveals that using low bandwidth actuators (viscoelastic thumbpad-endcap contact that is actuated by low-bandwidth muscles) and a soft finger contact to emulate constant proportional feedback control is no easy feat in terms of either engineering implementation or mathematical analysis. This is because, to mimick proportional control entails moving the effective center of pressure on the endcap sufficiently rapidly and in an appropriate manner. Since we do not permit macroscopic slips in our experiment, this can only be achieved by “rolling” the thumbpad on the endcap. The technical challenges in either implementing or analyzing a rolling digit contact can form the basis of an entire study and so will not be covered here. It suffices to realize that the apparently simple separation of controller and plant that we carried out above is not simple in either its implementation or analysis (Bicchi, 2000; Murray et al., 1994; Yoshikawa, 1990, and several references cited therein).

3.2.2 Multisensory feedback and linearly weighted sensor fusion

Now that we have separated the controller and plant in our low-order normal form, we are ready to incorporate multisensory integration with noise and time-delays. We will use a minimal model of noisy time-delayed sensors by assuming that each sensor is polluted by additive white Gaussian noise and has a constant time-delay. Our model reflects the experiments and includes three sensory channels, namely, thumbpad sensors³ (sensor 1), non-digital sensors⁴ (sensor 2) and vision (sensor 3), as shown in Figure 3.1. The signals from all three sensory channels are

mous advantage of using derivative control action in addition to proportional control, especially when time-delays are inherent in the system (Ogata, 2002; Doyle et al., 1992). However, for a multitude of reasons that are outlined in Section 5.2 on Page 93, we have chosen to restrict our attention to a proportional controller.

³Thumbpad sensors refers to cutaneous mechanoreceptors such as fast- and slow-adapting cutaneous pressure sensors.

⁴Non-digital sensors refers to cutaneous mechanoreceptors present outside of the digit, Golgi tendon organs, muscle spindles, etc.




Figure 3.1: A block diagram showing feedback control using three sensory modalities, namely, (1) thumbpad sensation, (2) non-digital sensors and, (3) vision. The mathematical model is a specific implementation of this block diagram that uses simple proportional feedback control and the entire closed-loop dynamics of this system is modeled as a subcritical pitchfork bifurcation. Time-delays (τ) and noise (ν) are explicitly labeled only for the sensory branches since for the purpose of this study, only the sensory time-delays and noise affect the relative usefulness of various modalities.

combined using a linear weighted sum to obtain $\hat{\theta}$, i.e.,

$$\begin{aligned} \overbrace{\hat{\theta}(t, t - \tau_1, t - \tau_2, t - \tau_3)}^{\text{state estimate with multiple delays}} &= \sum_3^{i=1} \omega_i \overbrace{(\theta(t - \tau_i) + \sigma_i \nu_i(t))}^{\text{noisy, time-delayed, unbiased sensors}} \\ \nu_i &\sim \underbrace{N_i(0, 1)}_{\text{Gaussian white noise processes}} \end{aligned} \quad (3.4)$$

where, ω_i are the sensory weights that will be optimized to maximize F_s in the model. We also assumed that all sensors are unbiased estimators of θ , i.e., in steady state (if θ was constant), the mean value of each sensory signal is equal to the actual value of θ .

3.2.3 Noise and time-delay estimates

I first present a general discussion of various aspects that were considered in estimating the noise and time-delays associated with each sensory channel. Sections 3.2.4 and 3.2.5 (see below) will then provide detailed estimates for the relevant parameters, including citations to existing experimental measurements of sensory variances and time-delays.

To estimate sensory variance (σ_i^2) for each sensation, we resorted to neurophysiological data reported in the literature. However, the data are not reported in terms of ability to distinguish rotation of a planar object (which is what we are interested in). For variance of vision, we took data on visual acuity reported in terms of angle subtended on the retina and converted it first to discrimination ability of spatial displacement for a known viewing distance. We then converted this to an equivalent rotation of the spring's endcap based on the geometry of the endcap. We used a similar procedure for determining variance in tactile discrimination of endcap rotation. Tactile discrimination ability is usually reported in terms of the ability to distinguish closely spaced pressure points on the thumbpad. We asked one consenting subject to “roll” their thumbpad on a flat surface attached to a 6-axis load cell (model 20E12A-I25. JR3, Inc., Woodland, CA) while they were pressing down with a near constant 3N vertical force. By simultaneously recording translation of the center of pressure on the flat surface and the angle of “roll” of the thumbpad, we were able to use reported center of pressure discrimination ability to find the equivalent endcap rotation angle discrimination ability. There are admittedly limitations to this process of estimating variance of thumbpad sensation. However, as we shall see below, this leads to physiologically tenable values for the variance. Finally, since we could not find any data that could be used to calculate the variance of non-digital sensors we used the reported fact that non-digital sensors were much less reliable (more noisy) than digital cutaneous sensors during slip-grip tasks (Macefield and Johansson, 1996; Häger-Ross and Johansson, 1996) to estimate sensory variance for non-digital sensors.

3.2.4 Sensory variance estimates

Human visual acuity is known to be at least 1' arc at a viewing distance of 250 mm (Liang and Westheimer, 1993; Saunders and Knill, 2003, 2004). Subjects were typically at a distance of 100-170 mm in our experiment. Together with an endcap radius of 40 mm, we estimated the standard deviation of visual sensation of endcap rotation to be $\sigma_3 = 0.0009$ rad. We estimated the standard deviation of thumbpad sensation based on reported tactile discrimination ability (Wheat et al., 1995, 2004) to be $\sigma_1 = 0.0007$ rad. (almost half the variance of vision). This agrees with the fact that tactile discrimination is typically considered to be more precise than vision (Cole and Abbs, 1988; Johansson and Birznieks, 2004; Johansson et al., 1992b). To account for the higher variance and hence lesser reliability of non-digital mechanoreceptors (Macefield and Johansson, 1996; Häger-Ross and Johansson, 1996), we used $\sigma_2 = 0.003$ rad (10x the variance of vision).

3.2.5 Sensory time-delay estimates

Neurophysiological data is more readily available for sensor time-delays and nerve-conduction time-delays than for time-delays arising out of neural sensory processing (see Figure 3.1). The time-delays arising from sensor and nerve-conduction delays for the three sensory channels are, $\tau_1 = 65$ ms (thumbpad sensors: Cole and Abbs, 1988; Johansson and Birznieks, 2004; Johansson et al., 1992b; Eliasson et al., 1995), $\tau_2 = 65$ ms (non-digital sensors: Gandevia and McCloskey, 1976; McCloskey, 1978; Kandel et al., 2000), $\tau_3 = 120$ ms (vision: van Beers et al., 2002; Paillard, 1996; Prablanc and Martin, 1992; Day and Brown, 2001). However, non-digital sensors are not only known to be noisy, but Häger-Ross and Johansson (1996) and Macefield and Johansson (1996) suggested that the response to object slip using only non-digital sensors might be as slow as 120ms. They do not state any reason for why this might be the case. However, their experimental results indicate the possibility that non-digital sensors might have greater time-delays than just sensor and nerve-conduction delays. This could be due to extra time-delay from sensory processing of noisy and non-collocated sensors that are distributed across different muscles/tendons and different patches of non-digital skin. We will discuss the potential sources for additional time-delays in more detail at the end of this chapter.

To account for the possibility of extra time-delays in non-digital sensors, we let τ_2 be a free parameter in our model that could vary from 65ms (minimum delay arising from sensor and nerve-conduction delays) to 120ms (the largest proposed delay by past studies). We tuned this parameter so that task-optimal performance of the model agreed best with experimentally measured effects of occluding both thumbpad sensation and vision (see Section 3.3).

3.2.6 Metrics for stability: “Survival times” and “Success rates”

We now define stability for the noisy, time-delayed model so that it agrees with the experimental notion of stability, namely, a compression was successful so long as it did not slip for a finite period of time. Stability in linear and nonlinear deterministic dynamical systems is a well-defined notion, either in the sense of asymptotic or Lyapunov stability (Doyle et al., 1992; Ogata, 2002; Guckenheimer and Holmes, 1983). For example, defining stability in the local sense (near a specific state of the system) is easily done for a hyperbolic fixed point⁵ if the system under consideration is described either by ordinary differential equations (Doyle et al., 1992; Ogata, 2002; Guckenheimer and Holmes, 1983) or by delayed differential equations (i.e., systems with time-delay) (Kolmanovskii and Nosov, 1986; Kolmanovskii and Myshkis, 1999; Engelborghs et al., 2000, 2002). However, when there is some source of noise in the dynamic system, stability is often defined in terms of stationary distributions, i.e., using steady-state distributions of time spent in various parts of the phase space of the dynamical system. For some dynamical systems that are modeled using stochastic differential equations, the stationary distribution can be analytically derived using the Fokker-Planck equations (Soize, 1994). However, for most complex dynamical systems, the true distributions are approximated using statistical histograms that are obtained through large ensemble simulations of the given stochastic dynamic system. For example, one could define stability for a noisy system based on distributions of the time spent by trajectories of a stochastic system in different parts of its phase space (Arnold, 1998). Numerically, this could be calculated by simulating large ensembles of the noisy dynamic system and thus obtaining histograms of time spent in different regions (if these distributions converge to stationary distributions). Peaks in the distribution (i.e., “representative” locations) can then be called as “stable” points in the phase space of the dynamical system.

However, in the context of our system, there is an alternate “natural” definition of stability that arises from the task requirement for subjects in the experiments. We called the experimental behavior as “stable” or “successful” if the subjects could prevent the spring from slipping for a finite time period (7s). This definition of “success” in our task naturally lends itself to “be studied in the context of a survival, or first passage, time problem” (Cabrera and Milton, 2004a), terms that we define below.

⁵The term hyperbolic refers to the requirement that none of the eigenvalues of the linearization near the fixed point of interest lie on the imaginary axis. In other words, the system can be stable or unstable in different directions, but not marginally stable in any direction.

Definitions of success rate and survival time

The first observation is that if a trajectory (time-series of θ) leaves the domain of attraction (region enclosed by dashed red curves in Figure 2.8 on Page 46) and never returns inside it during the 7s period (let us name it T^*), then it almost certainly reached one or the other undesirable stable fixed points (at $\theta \approx 0.5\text{rad}$; solid red curves in Figure 2.8) and thus, the spring “slipped”. So, we can define ‘success-rate’ as the probability that the time (t_{exit}) at which the θ trajectory exits the domain of attraction (to never return again) is greater than T^* (the desired duration of the hold phase, namely 7s).

$$t_{\text{exit}} = \min \{T; \text{s.t. } \theta(t > T) \notin [\theta_0 - \delta\theta, \theta_0 + \delta\theta]\} \quad (3.5)$$

$$p_{\text{success}} = p(t_{\text{exit}} \geq T^*) \quad (3.6)$$

where, $\delta\theta$ defines the domain of attraction and p_{success} is the ‘success-rate’. The time t_{exit} is called the ‘survival time’. Based on experimental data that after training subjects slipped in approximately 20% of the trials, we chose a nominal value of $p^* = 0.8$ for the success-rate to define a “successful compression” in our model. The utilization of p^* in our model will become clear when we define F_s below. However, it is important to note that in our simulations we calculate p_{success} using large ensembles of simulations and calculating the fraction of the ensemble that are ‘stable’ (p_{success}) in the sense that $t_{\text{exit}} \geq T^*$.

Definition of F_s in the model

For given sensory weights, the success rate (p_{success}) depends on the value of F_s . Symbolically, $p_{\text{success}} = p_{\text{success}}(F_s)|_{(\omega_1, \omega_2, \omega_3)}$. Now, we can define F_s for a successful compression in our model as the solution to the ‘root finding’ problem,

$$p_{\text{success}}(F_s)|_{(\omega_1, \omega_2, \omega_3)} = p^* \quad (3.7)$$

Numerically, we implemented this root finding problem using an adapted version of the Newton-Raphson method. Note that by defining F_s in this manner, it is implicitly (through the definition of p_{success}) an expected value, i.e., a metric of average performance and not single-trial performance. We have thus explained how F_s is defined. We will explain how to calculate sensory weights (i.e., $(\omega_1, \omega_2, \omega_3)$) that maximize F_s in Section 3.2.7 below.

Estimating γ

Using the definition of success rate as stated above, we estimated the last remaining parameter in our model – γ – so that when all sensory modalities were intact, the task-optimal performance (sensory weights that maximize F_s) for a success rate of 80% (p^*) was similar to experimental measurement.

Numerical integration of stochastic delay differential equations

Numerical integration of the one-dimensional stochastic delay differential equation (SDDE) was carried out using a simple Euler integration scheme (Küchler and Platen, 2000). As shown by Küchler and Platen (2000), for the case of additive noise, the Euler integration scheme has a strong order of convergence 1.0. The term ‘strong order’ just refers to the fact that if the ‘true’ solution to the SDDE was known for a specific instance of the noisy processes in the system, then, with smaller and smaller time-steps, the numerically integrated solution converges to the true solution. This is different from weakly convergent numerical techniques, where the average of some function of the solution converges to the ‘true’ value, but each individual solution might itself not converge. We will not say more about numerical techniques for integrating SDDEs since the paper by Küchler and Platen (2000) and the references cited by them provide a good reference for numerical integration of SDDEs.

3.2.7 Numerical optimization: sensory weights that maximize F_s

We now outline the numerical optimization procedure used to compute task-optimal sensory weights. As we saw in Section 3.2.2 and Figure 3.1, there are only three sensory weights that need to be found by our optimization routine that maximizes F_s . Given the additional constraint that the sum of the sensory weights is one, the optimization problem reduces to a 2-parameter optimization problem, namely,

$$\max_{\omega_1, \omega_2, \omega_3} F_s \text{ s.t. } \sum_{i=1}^3 \omega_i = 1 \text{ and } p_{\text{success}}(F_s)|_{\omega_1, \omega_2, \omega_3} = 0.8 \quad (3.8)$$

This is amenable to a global parameter search. We discretized the plane defined by $\omega_1 + \omega_2 + \omega_3 = 1$ in the positive octant of the space of sensory weights using a fine grid and numerically calculated F_s at each grid point. Thus, we found task-optimal performance and sensory weights for every sensory occlusion condition.

3.2.8 Sensory weights that minimize the effects of noise alone

To quantify the impact of time-delays on sensory weighting, we performed simulations using noise-minimizing sensory weights in addition to task-optimal sensory weights (that emerge from the combined effect of noise and time-delays; outlined in Section 3.2.7). Any deficit in performance (F_s) and deviation from experimental measurements that arise from using sensory weights that minimize the effects of noise alone can then be attributed to time-delays. The sensory weights that minimize the effect of noise alone are obtained using the same equation as

that of Bayesian inference in static tasks (see Section 1.1.3, Page 10), i.e.,

$$\omega_i = \frac{1/\sigma_i^2}{\sum_j (1/\sigma_j^2)} \quad (3.9)$$

3.2.9 *Best-fit* vs. *Baseline* simulations

As stated earlier in Section 3.2.5, we allowed the time-delay for non-digital sensors to be a free parameter. We varied its value in 5ms increments from 65ms to 120ms and performed the numerical optimization described in Section 3.2.7 at each increment.

The task-optimal performance with 100ms total time-delay for non-digital sensors (i.e., 35 ms extra delay beyond sensing and conduction delays) best reproduced the experimentally measured performance of combined loss of thumbpad sensation and vision. Henceforth, we will refer to this simulation with 100ms total time-delay in non-digital sensors as the *Best-fit* simulation.

We will refer to the simulation with no extra time-delay, i.e., total time-delay of 65ms for sensing and conduction delays of non-digital sensors as the *Baseline* simulation.

3.3 Experimental methods

The experimental setup, protocol and statistical analyses are already outlined in Section 2.2 on Page 31. Here, we will provide only the details of how we occluded thumbpad sensation and vision. Recall that the experimental protocol involved testing subjects over two days. The first day was the training phase that consisted of over 100 trials in all. Hence, all experimental data we report in this chapter are from the second day (post-training) unless otherwise stated.

3.3.1 Sensory occlusion

On the second day, we measured subjects' performance with normal thumbpad sensibility, both with and without vision. We blocked vision by asking subjects to wear goggles with opaque black tape pasted over them. Even in blinded trials, subjects could use vision until they placed their thumbpad on the spring's endcap. Then, an experienced hand surgeon⁶ administered 5cc of 1% Lidocaine solution on the ulnar and radial sides of the base of the thumb (just below the metacarpophalangeal – MCP-joint of the thumb) to obtain a digital nerve-block without affecting any musculature (and associated sensors) that actuate the thumb. Cutaneous sensation proximal to the thumb MCP joint was also unaffected. We regarded the

⁶We acknowledge here the critical support provided by Dr. Stephanie Roach, M.D.

nerve-block as complete when vision occluded subjects could not detect a 10 g load randomly applied using a pointed tip anywhere on their thumbpad or moved across their thumbpad. However, subjects were still able to detect large loads applied using the pointed tip. Subjects were also able to detect the small pointed projection at the center of the spring’s endcap by pressing their thumbpad forcefully against the endcap and used that to place their thumbpad suitably. We then measured performance after the occlusion of digital cutaneous sensation, both with and without vision.

Note that we did not measure maximal pinch strength after administering the nerve-block to prevent accidental damage to joints or soft tissue in the thumb from excessive force production since subjects could not feel pain once the nerve-block was effective.

3.4 Results

We found that with only one free-parameter (time-delay for non-digital sensors) and neurophysiologically tenable values for the remaining time-delays and noise variances, task-optimal behavior of our *Best-fit* simulation closely resembled the experimentally observed effects of sensory occlusion namely:

1. Occlusion of vision had no impact on performance if thumbpad sensation was intact.
2. Occlusion of thumbpad sensation caused a large drop in performance.
3. Loss of vision caused a further drop in performance when thumbpad sensation was absent.

Task-optimal sensory weights replicated experimental results of sensory occlusion much better than noise minimizing sensory weights alone. Moreover, unlike the relatively robust task-optimal sensory weights, noise minimizing sensory weights yielded performance that was extremely sensitive to the exact parameter value chosen for non-digital time-delays. This is of importance because we faithfully incorporated experimentally reported uncertainty in the time-delay for non-digital sensors (Häger-Ross and Johansson, 1996; Macefield and Johansson, 1996). Since our model, like all others, is only a cartoon of reality, i.e., an approximation to reality, contrasting the high parameter sensitivity of noise minimizing sensory weights to the robustness (to parameter variability) of task-optimal sensory weights lends further credibility to our conclusion that the selective use of vision is best explained by the combined effect of noise and time-delays.

3.4.1 Experimental results of sensory occlusion show selective use of vision

Loss of thumbpad sensation was severely detrimental to experimentally measured performance. When vision was present, loss of thumbpad sensation caused a

significant drop in F_s ($F_s^{\text{normal}} - F_s^{\text{numb}} = 0.18\text{N}$, $p = 0.021$; Figure 3.3.1). Loss of thumbpad sensation caused an even larger drop in F_s when vision was also absent ($F_s^{\text{normal}} - F_s^{\text{numb}} = 0.57\text{N}$, $p < 0.0001$; Figure 3.2).




Figure 3.2: This box plot summarizes the experimental results of sensory occlusion. The horizontal bars inside the box plots is the median, the notches are its robust 95% confidence interval, the boxes are bounded by the 75th percentile and 25th percentile, the numeric values next to each box is the sample mean, and the whiskers represent the entire range of the data. The differences in F_s shown in magenta are significant at a level of 0.025.

Visual occlusion did not affect performance when thumbpad sensation was present ($F_s^{\text{vision}} - F_s^{\text{blind}} = -0.07\text{N}$, $p = 0.423$; Figure 3.2). However, visual occlusion caused a large drop in performance when thumbpad sensation was absent ($F_s^{\text{vision}} - F_s^{\text{blind}} = 0.31\text{N}$, $p < 0.0001$; Figure 3.2) suggesting that non-digital sensors were not very effective in compensating for the loss of vision.

3.4.2 Combined effects of noise and time-delays drive the selective use of vision

We present the results of our numerical simulations in two parts. In the first part, we present the performance (F_s) arising from task-optimal sensory weights for both *Baseline* and *Best-fit* simulations. In the second part, we present the performance arising from noise minimizing sensory weights for both *Baseline* and *Best-fit* simulations. This will help examine the effect of additional “computational” time-delay for non-digital sensors and the effect of time-delays in addition to noise on performance. The results from both sections are summarized in Figure 3.3. The

sensory weights used to obtain the results shown in Figure 3.3 are given in Table 3.4.2. Keep in mind that the *Best-fit* simulation had an additional time-delay of 35ms for non-digital sensors over and above known sensor and nerve-conduction delays of 65ms.

Table 3.1: Task-optimal and noise minimizing sensory weights for all simulation results. The columns of sensory weights in this table correspond to the columns in Figure 3.3.

<i>Baseline simulation</i>							
Vision	Blind		Seeing		Seeing		Blind
Thumbpad	Intact		Intact		Numb		Numb
Strategy[§]	TO	NM	TO	NM	TO	NM	TO = NM
Digital (ω_1)	0.96	0.95	0.96	0.65	–	–	–
Non-digital (ω_2)	0.04	0.05	0.04	0.03	0.35	0.09	1
Visual (ω_3)	–	–	0	0.32	0.65	0.91	–

<i>Best-fit simulation</i>							
Digital (ω_1)	0.99	0.95	0.95	0.65	–	–	–
Non-digital (ω_2)	0.01	0.05	0	0.03	0.27	0.09	1
Visual (ω_3)	–	–	0.05	0.32	0.73	0.91	–

[§] TO = Task-Optimal; NM = Noise minimizing

Task-optimal sensory weights

The results relevant to this part are shown in blue markers (‘O’ and ‘+’ data markers) in Figure 3.3.

Much like the experimental results, both the *Best-fit* and *Baseline* simulations showed little or no effect of loss of vision if thumbpad sensation was present (Figure 3.3: ΔF_s between columns 1 and 2 for blue ‘O’ and ‘+’ markers). Interestingly, when thumbpad sensation was present, there was no difference in F_s between the




Figure 3.3: This figure shows the effect of time-delays in addition to noise on task-optimal multisensory integration. Only simulations with no additional “computational” time-delay for non-digital sensors (Baseline simulation, ‘+’) and with 35ms of extra “computational” time-delay for non-digital sensors that best replicated experimental data (Best-fit simulation, ‘O’) are shown above. The task-optimal strategy yields better performance (larger F_s) than the noise minimizing strategy demonstrating the effect of time-delays on multisensory integration. Baseline simulation (crosses) fails to agree with experimental results in the absence of thumbpad sensation (Nerve-blocked), since an additional loss of vision (cf. 3rd vs. 4th columns) has either a very small effect on performance (blue cross), or an unrealistic increase in performance (magenta cross). The Best-fit simulations (circles) that includes computational time delay for non-digital mechanoreceptors agree with experimental data (again, cf. 3rd vs. 4th columns). Note that when both thumbpad sensation and vision are absent there is no multisensory integration required since only the non-digital mechanoreceptors modality remains, hence the cross and circle with two colors. The task-optimal sensory weights used in the simulations are listed in Table 3.4.2.

Best-fit and *Baseline* simulations. Moreover, numerical simulations yielded F_s values very similar to experimentally measured performance.

When tactile sensation was blocked in our simulations (i.e., $\omega_1 = 0$), both the *Best-fit* and *Baseline* simulations showed a large drop in F_s (Figure 3.3: ΔF_s between columns 2 and 3 for blue ‘O’ and ‘+’ markers), much like in our experimental results. Importantly, task-optimal F_s for the *Best-fit* simulation closely resembled experimental results (*Best-fit* simulation: $F_s = 2.80\text{N}$, Experimental result: mean $F_s = 2.81\text{N}$).

Finally, both the *Best-fit* and *Baseline* simulations showed a further drop in performance if vision was lost in addition to thumbpad sensation. However, the decrease in F_s for the *Baseline* simulation was only 0.08N (from 2.97N to 2.89N) in contrast to a decrease of 0.27N (from 2.80N to 2.53N) for the *Best-fit* simulation. The additional decrease in F_s for the *Best-fit* simulation is due to the inclusion of 35ms additional time-delay for non-digital sensors. Importantly, experimentally measured drop in F_s was very close to that of the *Best-fit* simulation (Experimental: mean $\Delta F_s = 0.31\text{N}$, *Best-fit*: $\Delta F_s = 0.27\text{N}$).

In summary, task-optimal sensory weights for the *Best-fit* simulation closely reproduced experimental results. Although the *Baseline* simulation was not as successful as the *Best-fit* simulation in reproducing experimental data, it showed the same qualitative trends, namely the selective dependence on vision based on the availability of thumbpad sensation.

Noise minimizing sensory weights

We now present the results of simulations using noise minimizing sensory weights to examine whether the selective use of vision shown by task-optimal sensory weights is driven by noise, time-delays or the combined effect of both time-delays and noise. Simulation results using noise minimizing sensory weights are summarized in Figure 3.3 as pink ‘O’ and pink ‘+’ data markers for *Best-fit* and *Baseline* simulations, respectively and the corresponding sensory weights are given in Table 3.4.2.

With intact thumbpad sensation, noise minimizing sensory weights yielded performance that was indistinguishable from task-optimal performance although the underlying sensory weights in these simulations were different (Table 3.4.2: columns 1 and 2). The dominance of thumbpad tactile sensation is once again reinforced by the observation that noise minimizing sensory weights also found that the lack of vision did not affect F_s when thumbpad sensation was intact (Figure 3.3: columns 1 vs. 2, pink ‘O’ and ‘+’ markers).

Once thumbpad sensation is lost, the similarities in the results between task-optimal and noise minimizing sensory weights are lost. For the *Baseline* simulation, noise minimizing sensory weights yield F_s comparable to experimental results (Experimental: mean $F_s = 2.81\text{N}$, *Baseline* simulation: $F_s = 2.84\text{N}$). Similarly, F_s after loss of thumbpad sensation was comparable to experimental results for the *Best-fit* simulation also, when using noise minimizing weights (*Best-fit* simulation:

$F_s = 2.74\text{N}$).

However the loss of vision over and above thumbpad sensation yields the unrealistic effect of an increase in F_s from 2.84N to 2.89N for the *Baseline* simulation using noise minimizing sensory weights. This is because the noise-minimizing strategy (by definition) ignores the impact of time-delays on performance. As a consequence, when vision is available, it over-weights vision with no regard for the larger visual time-delay relative to non-digital sensors. However, for the *Best-fit* simulation using noise minimizing sensory weights, occluding vision in addition to thumbpad sensation did cause a decrease in F_s (from 2.74N to 2.53N), albeit not in as good agreement with experimental results like task-optimal sensory weights. In the *Best-fit* simulation, the time-delay for non-digital sensors is comparable to vision (although still smaller), thus when the noise-minimizing strategy disregards the effect of time-delays, we find unrealistic outcomes like observed in the *Baseline* simulation. Importantly, the effect on performance with loss of vision in addition to thumbpad sensation for noise minimizing sensory weights showed a severe dependence on the time-delay for non-digital sensors, going from an unrealistic increase in F_s to a smaller than experimental decrease in F_s .

In summary, noise minimizing sensory weights do not agree as well with experimental results as task-optimal sensory weights. Moreover, the selective dependence on vision based on availability of thumbpad sensation may or may not exist for noise minimizing sensory weights depending on the exact value of the time-delay for non-digital sensors. This lack of robustness for noise minimizing weights, but not for task-optimal weights leads us to conclude that the selective use of vision is best explained by the combined effects of both noise and time-delays.

Fitness landscape representation of simulation results

The results of the global optimization are shown as contour plots for both the *Baseline* (Figure 3.4) and *Best-fit* (Figure 3.5) simulations as an alternate representation of the data shown in Figure 3.3 and Table 3.4.2. The triangular planar surface in the contour plots are the set of feasible sensory weights, i.e., sensory weights that satisfy both the constraints $\omega_1 + \omega_2 + \omega_3 = 1$ and $\omega_i > 0$ for $i = 1, 2, 3$. The color coding depicts F_s according to the definition given in Equation (3.7) on Page 64 at each point on the plane.

3.5 Discussion

Experimental results of sensory occlusion exhibited the expected and commonly observed selective use of vision that we originally sought to explain. Numerical simulations closely resembled experimental measurements, and provided evidence that the selective use of vision is best explained by task-optimal performance that considers the combined effect of both noise and time-delays. Thus, we provide a model-based insight for multisensory integration during dynamic manipulation.




Figure 3.4: Results of the global optimization using the *Baseline* simulation. The edges corresponding to the no vision and no thumbpad sensation conditions are marked in the figure. Note how tactile sensation dominates the landscape when it is available (dark red region). Keep in mind that the vertices of the triangular planar surface of feasible sensory weights are the case when one sensory channel is used exclusively.




Figure 3.5: Results of the global optimization using the *Best-fit* simulation. Note how the peak in F_s for the no thumbpad sensation condition shifts slightly closer to the “vision corner” for the *Best-fit* simulation compared the *Baseline* simulation (Figure 3.4).

Importantly, the novel results we have presented in this chapter bring together previously disparate results from dexterous manipulation in humans and the optimal control hypothesis of sensorimotor control.

Our model revealed that a task-optimal sensory weighting strategy indeed explains the selective reliance on vision depending on the quality of digital cutaneous receptors. As we saw in the previous chapter, our novel use of bifurcation theory of nonlinear dynamical systems allowed the creation of a mechanics-based low-order model whose spatiotemporal dynamics was indistinguishable from experimental measurements at the edge of instability for the complex nonlinear behavior of dynamic manipulation. Numerical optimization of this model closely replicated our experimental results of sensory occlusion. It revealed that task-optimal sensory weights emerging from the combined effect of time-delays and noise better explained experimental data than noise minimizing sensory weights. Our results also lend support to the applicability of optimal state estimation to understand human sensorimotor control. Finally, the *Best-fit* simulation required an additional time-delay for non-digital mechanoreceptors over and above known sensory and conduction latencies. This naturally leads to the physiological interpretation that computational time-delays arising from sensory processing of noisy, non-collocated sensors such as non-digital mechanoreceptors could be significant to neuromuscular performance.

We found that sensory noise in non-digital sensors alone could not account for the large drop in performance if vision was lost in addition to thumbpad sensation, but an additional time-delay of at least 35ms together with large noise could account for this large drop. Since this additional time-delay is over and above known sensory and nerve-conduction latencies for non-digital sensors (Gandevia and McCloskey, 1976; McCloskey, 1978), it can only be due to a “computational” time-delay (Figure 3.1). Thus, we propose that sensory processing of non-collocated and noisy sensors like non-digital mechanoreceptors in our experiment could result in physiologically tenable (Johansson and Birznieks, 2004; Cole et al., 1998; Häger-Ross and Johansson, 1996; Macefield and Johansson, 1996) computational time-delays.

With the available data, we can only speculate about the source of this additional time-delay. It could likely arise from computations needed to transform signals from distributed sensors into an aggregate state estimate. The delay might also be a consequence of the excessive noisiness of non-digital sensors. For instance, Johansson and Birznieks (2004) reported that processing of noisy tactile signals from fingerpads to correctly discriminate the external sensory cue using a probabilistic inference procedure caused a small but measurable time-delay. It is possible that some variant of this might be involved in continuous-time state-estimation using noisy non-digital sensory signals. Finally, it is also possible that sensory processing of noncollocated non-digital sensors might involve areas of the nervous system other than the spinal cord, namely areas of the cerebral cortex, cerebellum, etc. In that case there will be significant increase in transmission delays and thus lead to the extra time-delay. Nevertheless, it remains that the data

from this study cannot conclusively state that an extra time-delay exists or point to its source.

Independent of the exact value or existence of this extra time-delay we saw that task-optimal sensory weights always lead to a selective dependence on vision. Thus, any uncertainty related to the time-delay for non-digital sensors does not affect the primary result of the work presented in this chapter. See Chapter 5 for a more detailed discussion of whether this extra-time delay could just be an artifact of our simple model and how we can test that experimentally.

It is worth noting the ability of our nonlinear low-order model to capture both the effects of sensory occlusion (results presented in this chapter) and the spatiotemporal dynamics (results from Chapter 2) of the task at the limit of performance of a system as complex as dynamic manipulation. Importantly, our model had only one free parameter (time-delay for non-digital sensors) and all other parameters were obtained from basic frictional constraints and previously reported neurophysiological data for sensory time-delays and noise variances. In fact, the one free parameter was itself chosen based on the physiologically tenable and previously suggested (Cole and Abbs, 1988; Cole et al., 1998; Häger-Ross and Johansson, 1996; Macefield and Johansson, 1996) rationale that sensory processing from noisy, non-located sensors might result in onerous computational time-delays.

3.6 Conclusions

Our results lead us to conclude that during dynamic sensorimotor behavior at the edge of instability, small changes in the underlying system produce measurable and mathematically tractable changes in performance without compromising the complexity of the task or oversimplifying the analysis. Specifically, time-delays do have the predicted effect on multisensory integration in addition to the previously known effects of noise. We also propose that computational time-delays in sensory processing have measurable consequences to task performance.

Another relevant observation about our modeling approach and sensory weighting is the fact that the task-optimal sensory weights that balance the effects of both time-delays and noise do not emerge from some form of multiobjective optimization. In other words, we do not try to manually define cost functions that weigh the relative impact of time-delays and noise to task performance. The relative impact of time-delays and noise are captured by a scalar performance metric (F_s) thus enabling us to solve an unambiguously defined optimization problem.

3.6.1 Do we better understand the effects of growing old or cold?

Finally, we end this chapter with a brief discussion about the ramification of the results we found to the question of “why we drop objects more often when we grow old or cold”.

Let us first address the question of growing cold. If you stepped out during the winter or washed your hands in cold water and immediately tried to do some fine manipulation task, you will most probably fail. Coming back home on a cold winter day and fumbling with the keys and even dropping them is a familiar sight for people that live sufficiently far from the equator! When your hands grow cold, cutaneous sensors on your fingerpads don't respond to stimuli, similar to our experimental setup where we numbed the thumbpad by administering a local anesthetic. It is not immediately apparent if cooling down cutaneous sensors "pollutes" them by injecting significant sensor latencies, making them more erratic (noisy), or simply desensitizing them like a local anesthetic. However, as we saw in this chapter, either or both of these mechanisms could render these sensors less effective. The outcome would then be a greater reliance on vision and more importantly, higher sensory variance. A larger number of redundant sensors helps reduce the variance of the net estimate of the measured sensory cue (so long as the noise processes polluting the sensors are independent). Hence, loss of one sensor increases the net "noise-level" in the sensory system. In the context of our model, the increase in noise-level will decrease the 'success-rate' for a specific level of F_s . In other words, greater noise in the system will lead to a greater probability for slips. No doubt this explanation is speculative and heuristic, but it provides some useful insight into how sensor "quality" might influence dynamic manipulation demands of everyday life. In addition to the effect on sensor efficacy, a low temperature will undoubtedly affect several other electrochemical and mechanical properties of biological tissues such as muscles, tendons, ligaments, etc.

With that discussion on the effect of cold hands, the consequence of aging is apparent. Aging degrades sensors in the body and also causes overall "neural" slowing in the CNS (Cole et al., 1998, and references therein). Hence, replacing 'cold' by 'old' in the previous paragraph provides insight into why aging could impair our dexterous manipulation ability and make us drop objects more often.

Chapter 4

Clinical implications

4.1 Introduction

This chapter deals with a clinical application of the techniques developed so far in this thesis to reveal possible reasons behind the effectiveness of an upcoming treatment for osteoarthritis. There are some medical terms that will be frequently used in this chapter. To avoid repeated use of long phrases, I will resort to abbreviations. An effort has been made to minimize area-specific terminology and abbreviations to ease the burden on the reader. However, some terminology and abbreviations are unavoidable. I present a short list of abbreviations that will be used in this chapter together with a brief explanation of each term and relevant citations from the literature.

OA *Osteoarthritis*: is a degenerative¹ joint disease (Altman et al., 1986) that is almost inevitable in the elderly population (it is estimated that 80% of the population in the U.S. over 65 years will have OA; Green, 2001).

CMC *Carpometacarpal*: joint is the name for a joint that lies at the base of the thumb (Santos and Valero-Cuevas, 2006; Valero-Cuevas et al., 2003a) and is often the victim of OA.

HA *Hyaluronic acid* or *Hyaluronan*: is a carbohydrate polymer (Meyer, 1950b,a; Meyer and Fellig, 1950) ubiquitous in the fluid medium (called ‘matrix’) that is prevalent around cells in most tissues of our body. It provides lubrication for articulating joints (Fraser et al., 1997) among the myriad physiological functions it is purported to serve.

DASH *Disabilities of arm, shoulder and hand*: is a 30-question questionnaire used to evaluate upper-extremity function in clinics (Hudak et al., 1996; Amadio, 1997; Navsarikar et al., 1999). The questionnaire asks about the patient’s symptoms as well as ability to perform certain activities using the upper extremity, which has to be answered based on the patient’s condition over the past week. Each response is assigned a score from 1 to 5 – mildest to worst symptom or intact function to severe dysfunction, as appropriate to the question asked. The average score from all 30 questions is converted to a 0-100 scale, where a higher score means a greater disability.

VAS *Visual analog scale (for pain)*: is a tool used to assess the level of pain a person feels (Carlsson, 1983; Ohnhaus and Adler, 1975; Price et al., 1983). It is simply a 100mm (10cm) long line segment, where the lower/left end represents “no pain” and the upper/right end represents “the worst imaginable

¹A degenerative disease will progressively worsen over time, unlike an infectious disease can improve over time.

pain”. In a clinic, a patient would mark off the location on this line that they think best represents the pain they feel. The clinician who administers this test for pain would record the distance in mm (cm) from the no-pain end to score the level of pain as a number between 0 and 100 (10).

We will examine the effect of a specific treatment (intra-articular HA injection) on patients diagnosed with OA of the CMC joint of the thumb. Specifically, we test our hypothesis that the self-reported (using the DASH questionnaire) improvement in hand function after treatment is primarily because of pain relief and not because of any improvement in innate sensorimotor capability. Importantly, this study takes us one crucial step toward having a clinical tool for mechanically meaningful quantitative assessment of hand function.

4.1.1 Why does hand function improve in CMC OA patients after treatment with intra-articular HA injection?

Osteoarthritis is a very prevalent disease in the elderly. It is a degenerative disease that can affect any joint in the body in theory, but most commonly affects joints of the hand, feet, hips, knees and spine. A characteristic of OA is the wear and thinning of the cartilage that cover and form the articulating surfaces of a joint. As the cartilage wears over time (e.g., as a person ages), the bone surfaces become less protected, make direct contact and this results in “joint pain”, which is often excruciating given the rich innervation of the periosteum. Hence, not surprisingly, the most common symptom of OA is chronic pain. Unfortunately, there is no evidence of a definitive cure for OA so far (Brandt et al., 2000).

4.1.2 Intra-articular HA – a potential treatment for CMC OA

There is much evidence for beneficial effects of intra-articular HA injection to patients suffering from knee OA (Brandt et al., 2000; Fraser et al., 1997; Listrat et al., 1997; Peyron, 1993; Wobig et al., 1998) although the efficacy and mechanism are still hotly debated. There are also some preliminary reports that intra-articular HA injection might improve manipulation ability in patients suffering from CMC OA of the thumb (Fuchs et al., 2006; Venkadesan et al., 2005; Wei et al., 2002). It has been suggested by several studies that injecting HA into the intra-articular space (the joint space between articulating surfaces of the bones) might have beneficial effects (Altman and Moskowitz, 1998; Brandt et al., 2000; Fuchs et al., 2006; Listrat et al., 1997; Peyron, 1993; Stahl et al., 2005; Wei et al., 2002; Wobig et al., 1998) by achieving two important goals — alleviating pain and improving “function”. Note that most studies of effectiveness of OA treatment using HA have been restricted to the knee, with very little evidence for effects of treatment on

CMC OA patients. There are of course several other established and customary treatments for OA that range from surgical intervention at one extreme to intra-articular steroid injections, over-the-counter nonsteroidal anti-inflammatory drugs (Green, 2001) such as Advil™ or even simply splinting. However, HA is considered to be more than just a symptomatic treatment due to its ubiquitous natural presence in healthy joints and empirically observed decline in concentration in OA joints (Fraser et al., 1997). It suffices to say that it has been often reported that HA is effective in pain reduction in the short term at the least (Brandt et al., 2000) and some studies claim that HA might cause beneficial changes in articular surface properties (Listrat et al., 1997) and might provide long term pain reduction (Altman and Moskowitz, 1998; Brandt et al., 2000; Listrat et al., 1997; Wei et al., 2002; Wobig et al., 1998). Critics of HA indicate that local administration of HA may have mostly a placebo effect because HA remains in the joint space for a few days at most (Brandt et al., 2000).

These past results on using HA to treat knee OA are the motivation for us to try and explain if and why DASH scores (self-reported improvement in physical function) improve in CMC OA patients treated with intra-articular HA injections².

4.1.3 Existing outcome measures do not quantify impairment and recovery of dynamic manipulation

There have been some studies (Fuchs et al., 2006; Stahl et al., 2005), including the precursor to the work presented in this chapter (Venkadesan et al., 2005) that have reported improvement in hand “function” after intra-articular HA injection. With the exception of Venkadesan et al. (2005) the other studies relied solely on indirect “dexterity” evaluation tests such as the DASH (Hudak et al., 1996; Amadio, 1997; Navsarikar et al., 1999), maximal pinch strength (Fess, 1995; Mathiowetz et al., 1985; Totten and Flinn-Wagner, 1992), Purdue Pegboard Test, Jebsen-Taylor Hand Function Test (Fess, 1995; Totten and Flinn-Wagner, 1992), clinical evaluation of the fingers (e.g., range of motion) (Jones, 1989), etc., which are limited in their ability to quantify thumb function as we shall see below.

Everyday manipulation tasks are dynamic and typically involve low forces, thus the validity of maximal pinch strength as an indication of manipulation ability is doubtful. In fact, we have reported earlier that maximal pinch strength cannot even distinguish between people with CMC OA and those with healthy hands (Valero-Cuevas et al., 2003b).

Tests of hand-eye or reach-to-grasp coordination such as the Purdue Pegboard Test or the Jebsen Taylor Test are not specific to finger function and involve

²The work reported in this Chapter was carried out in collaboration with researchers and clinicians at the Hospital for Special Surgery, New York, NY. The researchers at New York that were instrumental to this study are, Dr. Lisa Mandl, Dr. Sherry Backus, Dr. Stephen Lyman, Dr. Robert Hotchkiss, Alison Swigart, Dr. Margaret Peterson and, Lily Ariola.

the entire upper extremity, thus being susceptible to compensatory and adaptive strategies that can mask hand impairment.

Clinical evaluation of the fingers such as range of motion or subjective questionnaires that rely on verbal reports of hand function such as the DASH do not quantify the sensorimotor ability of object manipulation using the fingers.

4.2 Hypotheses

As part of the first RCT on the use of HA for CMC OA, we proposed two working hypotheses³:

1. We hypothesized that HA injection would lead to self-reported improvement in hand function and pain relief, but not improve dynamic manipulation ability at low forces.
2. We also hypothesized that the self-reported improvement in hand function is primarily due to pain relief; and patients with better current dynamic manipulation ability at low forces will benefit more from pain relief. In other words, dynamic manipulation ability is a predictor of treatment effectiveness of HA injection for CMC OA patients, provided there is pain relief.

The rationale for our hypotheses rests on past studies which reported positive results of HA injection using subjective questionnaires. From Section 4.1.1, recall that chronic pain is the clinical hallmark of CMC OA and past studies have repeatedly found consistent short and long term pain relief after HA treatment (VAS score) and consistent improvement in function according to self-reported function (DASH score). However, these studies report inconsistent improvement in objective metrics of manipulation ability (using methods such as pinch meters, range of motion of the thumb, etc.) and we saw in Section 4.1.3 that these past studies did not use any objective, finger-specific metrics of dynamic manipulation ability. Finally, since chronic and severe pain could restrict CMC OA patients from exploiting their innate sensorimotor control ability to the fullest, and HA seems to provide short and long term pain relief, we hypothesized that: (i) after HA treatment, DASH scores (self-reported function) would improve primarily due to pain relief and not because the innate dexterous manipulation ability improved, and (ii) since pain could mask the innate sensorimotor control ability, patients with better dynamic manipulation ability at low forces should function better after treatment, i.e., report higher DASH scores after achieving some pain relief due to HA injection.

³The work to test these hypotheses was carried out at the Hospital for Special Surgery, New York, NY, under the supervision of Dr. Lisa Mandl and was supported by, grants from NIH R21-HD048566, NSF 0237258 and the Whitaker Foundation to Dr. Francisco Valero-Cuevas and grants from NIH K23-AR50607 and Wyeth Pharmaceuticals to Dr. Lisa Mandl.

Notice that we repeatedly stress the need for “low forces” when referring to dynamic manipulation ability. This is because we are interested in quantifying the innate sensorimotor control ability of the hand for tasks of daily living, independent of pain. For patients with degraded joint surfaces, higher forces that result in higher joint contact forces will likely lead to greater pain.

An alternate explanation for self-reported improvement in DASH scores after treatment with HA is as follows. It is possible that degradation of CMC joint articular surfaces due to OA could cause not only pain, but loss of sensorimotor control ability because some (unknown) aspect of the joint mechanics could have altered due to the articular surface degradation. Hence, it is possible that if HA improves articular surface properties (Listrat et al., 1997), it could result in improvement of self-reported hand function because the innate sensorimotor control ability improved and not just because of pain relief. However, given the tentative nature of the evidence for any mechanism for lasting change in articular surface properties due to HA injection (Brandt et al., 2000), we expected that the innate sensorimotor control ability is not affected by HA injection treatment and the self-reported improvement in hand function is primarily due to pain relief. Moreover, existing research on biomechanics of the thumb do not shed any light on the relationship between articular surface properties and sensorimotor function.

We tested our hypotheses using the method of quantifying dynamic manipulation that we have presented in the previous chapters of this thesis. The results shown below have been presented at a conference (Venkadesan et al., 2005).

4.3 Methods

In our study, 32 consenting patients with CMC OA received intra-articular injection of Hylan G-F 20 for 3 consecutive weeks after approval by the Institutional Review Board at the Hospital for Special Surgery, New York, NY, U.S.A., where the study was conducted under the supervision of Dr. Lisa Mandl. Functional evaluation at baseline (prior to 1st injection), and 26 weeks later included VAS for pain; DASH Questionnaire; key and opposition pinch strength; and the mean sustained maximal load during the task of compressing a slender spring, i.e., F_s – a metric of dynamic manipulation ability at low forces. The dynamic manipulation task was nearly identical to that presented in Chapters 2 and 3, using the same spring and setup used by us previously. It consisted of compressing a slender spring prone to buckling using the thumbpad to maximize vertical compressive load, and holding it for 5 seconds (see Figure 2.3 on Page 29). We chose to set a holding time of 5s instead of 7s that we used previously since the thumbs of the patients we tested could become painful and we were trying to minimize patient discomfort, while garnering useful information at the same time. Given the time constraints in the clinic and because our past studies showed no significant effect of training, we did not impose beyond 10 compressions of the spring on any patient.

4.4 Statistical Analyses

At the time of analysis of the data, only 19 patients had completed the 26 week followup and one patient did not compress the spring slowly and long enough to provide reliable data. The average age of the patients was 64 years (range 46-79), with 22 females and 10 males. We always verified the necessary assumption for parametric statistics that the residual values should be independently, identically and normally distributed.

4.4.1 Testing Hypothesis 1: Change in metrics of hand function after treatment

We first tested to see if any of the outcome measures, namely — DASH, VAS, key pinch strength, opposition pinch strength or, F_s changed from baseline (before first injection) to final-followup (after 26 weeks). The comparisons were performed as five paired two-sided t-tests.

4.4.2 Testing Hypothesis 2: Multiple linear regression

To test whether the variance in observed improvement of DASH scores after treatment with HA injection could be explained by improvement in pain and innate sensorimotor control ability, we performed a multiple linear regression analysis. Specifically, we posed the linear model for this statistical test as follows:

$$\Delta\text{DASH} = \alpha_1(\Delta\text{VAS}) + \alpha_2(F_s^{\text{baseline}}) + \epsilon \quad (4.1)$$

where, ΔDASH is positive if the DASH score decreased after treatment, since a lower DASH score means better upper extremity function; ΔVAS was defined to be positive if there was pain relief since a lower VAS score means lesser pain; ϵ is the residual error term, which satisfied assumptions for parametric statistics, namely, normality, independence and identical distribution. Note that α_1 and α_2 — the slopes of the regression equation — need to be significantly different from 0 in order for us to accept the hypothesis that change in pain and baseline dynamic manipulation ability are associated with the self-reported change in function. If both are not only different from 0, but greater than 0, then we can conclude that our data provide support for Hypothesis 2.

It is important to note that we do not include an intercept in our model and force the regression plane defined by Equation (4.1) to pass through origin. The reason for this becomes apparent by noting that $F_s^{\text{baseline}} = 0$ means that the person's ability was at manipulation was so poor that they could not compress the spring even as well as a small (say, $< 1\text{N}$) dead weight, and no change in pain just lends further credibility to our assumption that the regression plane passes through the origin. In other words, we assumed that (no pain-relief) + (completely dysfunctional thumb) = (no effect of treatment). We verified the

validity of this assumption by performing a separate multiple linear regression analysis that allowed for a non-zero intercept. We found that our data could not provide support for a non-zero intercept and the intercept obtained from the regression was not statistically different from 0 ($p=0.743$).

4.4.3 Stepwise linear regression

The test for Hypothesis 2 outlined above cannot rule out the possibility that the variance in ΔDASH can be better explained by including other outcome measures such as pinch strength, etc. in addition those stated above. To further strengthen our test of Hypothesis 2, we performed a stepwise regression analysis. Stepwise regression analysis is a data-mining procedure to explore and compare alternate choices for linear regression models to determine how best to explain a data set. To understand this further, we outline below, the details of the stepwise linear regression analysis for our data set.

A stepwise regression analysis is similar to a standard linear regression analysis in all regards but for how the independent or predictor variables (regressors) are treated. A standard regression analysis uses a fixed number of predetermined independent regressors based on the researcher's insight about the problem at hand. A stepwise regression analysis, however, performs additional statistical testing to find the least number of independent variables that possess meaningful explanatory power. Specifically, a stepwise regression procedure starts with a master-set of independent variables as potential candidates and eliminates all but those variables that possess significant explanatory power, i.e., independent variables with non-zero partial R^2 and a slope significantly different from 0 (p -value less than pre-set threshold). So, a stepwise procedure finds the minimal subset of the master-set that we started with so that every independent variable that belongs to the subset explains some non-zero variance in ΔDASH with statistical significance greater (i.e., p -value smaller) than a preset threshold (typically, $\simeq 0.20$). Remember that both DASH and VAS scores numerically decrease with improvement. Hence, they were defined as (baseline – final) so that an improvement yields a positive number for ΔDASH or ΔVAS .

For our study, the dependent variable was ΔDASH that could depend on several factors. We list below in Table 4.1 a list of the seven factors that we measured and considered as potential candidates for the master-set of our stepwise regression analysis.

Table 4.1: This table gives a list of potential factors that could explain any improvement in DASH scores after treatment with intra-articular HA injection.

Predictor	Explanation for inclusion
ΔVAS	Intra-articular HA injection is presumed to provide pain relief. Moreover, the DASH questionnaire has several questions that explicitly ask the patient to rate the severity of pain that they feel in different parts of the upper extremity.
ΔF_s	If sensorimotor control ability improved with treatment then it is very likely that patients will report improvement in upper extremity function using their DASH questionnaire.
$\Delta(\text{KEY or OPP})$	Using the same reasoning as ΔF_s , if pinch strength improves, it is possible that subjects might report better “function”, since they feel stronger!
F_s^{baseline}	Pain impairs manipulation ability. Since HA is presumed effective in pain relief, we expected that patients with innately better sensorimotor ability (F_s at baseline) will report greater improvement in function, since they are less limited once pain relief allows them to utilize their innate ability.
$(\text{KEY or OPP})^{\text{baseline}}$	Similar to F_s^{baseline} , having higher pinch strength at baseline could lead to better self-reported improvement with treatment if the musculature is not yet affected by pain-driven disuse or atrophy.

In summary, our statistical model using the master-set was,

$$\begin{aligned}
\Delta\text{DASH} = & \alpha_1 (\Delta\text{VAS}) \\
& + \alpha_2 (\Delta F_s) \\
& + \alpha_3 (\Delta\text{KEY}) \\
& + \alpha_4 (\Delta\text{OPP}) \\
& + \alpha_5 (F_s^{\text{baseline}}) \\
& + \alpha_6 (\text{KEY}^{\text{baseline}}) \\
& + \alpha_7 (\text{OPP}^{\text{baseline}}) \\
& + \epsilon
\end{aligned} \tag{4.2}$$

where, $\Delta(\text{DASH}, \text{VAS}) = (\text{DASH}, \text{VAS})^{\text{baseline}} - (\text{DASH}, \text{VAS})^{\text{final}}$, and the other quantities (pinch strengths – KEY or OPP, F_s) were defined as (final – baseline). Once again, ϵ , the model residue needs to satisfy the assumption $\epsilon \sim \text{iid}N(0, \sigma^2)$, i.e., it should be independently, identically, normally distributed.

4.4.4 Pairwise correlation analysis

A necessary and crucial assumption for stepwise regression analysis is the mutual independence of predictor variables (the right-hand side of Equation (4.2)). We tested this assumption by calculating pairwise Spearman correlation coefficients for all pairs of the independent variables listed above. We found that $\text{KEY}^{\text{baseline}}$ showed significant correlation with ΔVAS , ΔKEY and $\text{OPP}^{\text{baseline}}$. ΔVAS is critical to test our hypothesis and we can substitute $\text{KEY}^{\text{baseline}}$ with baseline opposition pinch strength ($\text{OPP}^{\text{baseline}}$). In fact, we found that $\text{KEY}^{\text{baseline}}$ strongly positively correlated with $\text{OPP}^{\text{baseline}}$ (Table 4.2; $\rho = 0.705$, $p=0.001$). Hence, we dropped $\text{KEY}^{\text{baseline}}$ from our stepwise regression analysis since it becomes redundant. We also found that ΔF_s significantly correlated with F_s^{baseline} . However, to leave open the possibility that genuine improvement in dynamic manipulation ability might have led to the improved DASH score, we performed two separate stepwise regression analyses — with and without the inclusion of ΔF_s . The results of the pairwise correlation analysis are presented in Table 4.2.

4.5 Results and Discussion

We found that only DASH and VAS improved significantly (statistically and clinically) after treatment with intra-articular HA injection. Importantly, the improvement in DASH was explained only by improvement in VAS and baseline F_s . In other words, patients who got the most pain relief reported the greatest improvement in hand function and patients with innately better sensorimotor control ability had greater benefits from pain relief.

Table 4.2: Results of pairwise Spearman correlation for all the predictors from the master-set in the stepwise regression model. For each row, the number on the top is the Spearman correlation coefficient (ρ) and the number in the bottom is the p - value for this correlation coefficient. If the p - value was < 0.05 , we regarded that correlation as significant. For want of space, the ‘baseline’ superscript for all variables has been replaced by ‘0’, for example F_s^0 in the table below refers to F_s^{baseline} .

	ΔVAS	ΔF_s	ΔKEY	ΔOPP	F_s^0	KEY^0	OPP^0
ΔVAS							
ΔF_s	-0.334 (0.175)						
ΔKEY	0.342 (0.165)	0.001 (0.997)					
ΔOPP	0.033 (0.897)	0.425 (0.079)	0.364 (0.137)				
F_s^0	0.272 (0.275)	-0.688 (0.002)	-0.206 (0.412)	-0.237 (0.345)			
KEY^0	-0.492 (0.038)	0.026 (0.919)	-0.492 (0.039)	0.024 (0.925)	0.041 (0.870)		
OPP^0	-0.169 (0.503)	0.200 (0.427)	-0.274 (0.272)	-0.342 (0.165)	0.126 (0.618)	0.705 (0.001)	

4.5.1 DASH and VAS are the only metrics that showed improvement after treatment

Similar to what was found in past studies on HA injection for patients with CMC OA, we found that patients had good long-term pain relief (VAS) and good self-reported improvement in hand function (DASH). In support of our first hypothesis, we found that treatment did not improve either pinch strength (both key and opposition) or sensorimotor control ability (F_s). See Table 4.3 for a summary of the results of our t-tests.

Table 4.3: DASH score and VAS for pain are the only metrics that showed a significant change 26 weeks after the first HA injection. The data are reported as mean $\pm$ standard deviation in the table. Statistically significant differences are bold and italicized.

	Baseline (N=32)	Week 26 (N=19)	p-value Baseline-Week 26
DASH score	29.1 $\pm$ 16.5	14.9 $\pm$ 13.5	<i>0.011</i>
VAS for pain	6.2 $\pm$ 1.8	4.2 $\pm$ 2.3	<i>0.022</i>
Key pinch strength (N)	13.6 $\pm$ 5.6	14.1 $\pm$ 5.3	0.379
Opp pinch strength (N)	10.6 $\pm$ 5.3	11.6 $\pm$ 5.5	0.204
F_s (N)	2.66 $\pm$ 0.26	2.65 $\pm$ 0.27	0.754

4.5.2 Patients with better initial sensorimotor control ability benefited more from pain-relief

The multiple linear regression found that change in VAS and F_s^{baseline} had significant explanatory power for change in DASH. Specifically, the linear model that we found by the multiple linear regression was:

$$\Delta\text{DASH} = \underbrace{3.11N^{-1}}_{p=0.025} F_s^{\text{baseline}} + \underbrace{4.45}_{p=0.0009} \Delta\text{VAS} + \epsilon; R^2 = 0.81 \quad (4.3)$$

The regression revealed that change in VAS and baseline F_s could explain more than 80% of the variance in change in DASH. Moreover, the change in DASH had a statistically significant positive regression slope with both baseline F_s and change in VAS. This provides support for our hypothesis that the primary source of improvement in function is from pain relief and those with innately better sensorimotor control ability showed greater improvement.

4.5.3 Only pain relief and baseline F_s could explain self-reported improvement

Both the stepwise regressions — with or without including ΔF_s — found only ΔVAS and F_s^{baseline} could significantly explain variance in $\Delta DASH$. An automatic data-mining approach such as a stepwise regression analysis is frequently criticized for proposing unrealistic models, i.e., they sometimes select physically unrealistic or unreasonable predictors. In light of this nature of stepwise regressions, our finding that the only predictor variables automatically selected by this data-mining approach agree with Hypothesis 2 lends strong support to the validity of our conclusion that, pain relief and innate dynamic manipulation ability suffice to explain self-reported improvement.

4.5.4 Static pinch strength is not informative of any aspect of functional improvement

In agreement with our past study on OA (Valero-Cuevas et al., 2003b) we found that neither key nor opposition pinch strength changed with treatment. Moreover, neither metric of strength possessed any predictive power of improvement in hand function (change in DASH). This indicates that unlike sensorimotor control ability (F_s), patients with higher strength do not necessarily show greater improvement with treatment.

4.6 Conclusions

The results of this preliminary study lend support to our hypothesis that self-reported improvement in hand function is primarily due to pain relief. Importantly, patients with innately better sensorimotor control ability showed greater improvement, while greater pinch strength did not mean greater improvement. This leads to three implications of clinical importance:

1. Dynamic manipulation ability at the edge of instability is predictive of treatment outcome using HA. This provides a useful clinical tool for assessing treatment outcome.
2. The force required to perform the dynamic manipulation task is very low ($< 3N$), thus probably making it less influenced by pain due to joint contact forces and stimulation of the periosteum. This provides a useful window into the innate dexterity of patients without any confounding effects of pain.
3. Pinch strength should likely be removed from the customary battery of tests administered to the OA hand. For humane reasons, testing maximal pinch strength and inducing pain is unnecessary given that pinch strength has little clinical predictive value.

These data are only preliminary and, because we did not collect any data about the motion of the spring during the compression, we have no way to verify if patients were indeed at the edge of instability. This is of some concern since with painful thumbs, motivation for maximizing task performance could be compromised. Regardless, the spring compression was more informative than any other currently preferred objective outcome measure. Importantly, despite such potential confounding factors, we found supporting evidence for our hypotheses. In fact, refining our experimental methods to use motion capture and provide feedback to the clinical tester about proximity to the edge of instability (based on established dimensional collapse at the edge of instability — Chapter 2) can only help further strengthen the results and conclusions of the study presented in this chapter.

Specifically, our results suggest that HA provides self-reported improvement in function primarily due to pain relief. This provides better-defined directions for future studies to examine the physiological mechanisms behind the apparent benefits of intra-articular HA injections. Finally, this study is the first crucial step to develop a quantitative assessment of hand function and importantly, a means to quantify changes in hand function as part of any randomized controlled trial.

Chapter 5

Future Work and Conclusions

This chapter forms the concluding part of this thesis. It is organized into two parts based on the level of specificity of the proposed future work. The first part will examine the main results of this thesis in a critical manner. Specifically, it will examine the scope, generalizability and limitations of the main conclusions of the work presented in this thesis. I will then present some short-term goals based on the preceding critical evaluation. These short-term goals aim to provide refinements to broaden the scope of this work and also to assess the generalizability of some of our results. The second part of this chapter consists of speculative proposals for long term goals based on some of the results of this thesis. Specifically, it presents an approximate road-map for future exploration of new directions in optimal feedback control of complex high-dimensional systems, be they machines or biological organisms.

5.1 Refinement of the normal form model

The low-dimensional model of dynamic manipulation at the edge of instability that we have developed is very effective and useful in exploring problems related to sensorimotor control and even capable of generating experimentally testable hypotheses (for example, see Sections 5.3 and 5.4). However there are several possible refinements to the model that can make it more representative of the underlying physics of the real system it tries to approximate. Below is a list of some such refinements.

5.1.1 Inclusion of “imperfections”

The normal form equation we used for modeling the transition to instability in the thumb + spring + nervous system is generic only in the class of systems that have an inherent symmetry. However, we already saw in Section 2.3.2 on Page 44 that because the rotational buckling is preceded by the lateral buckling, an ‘imperfection’ is introduced into the endcap rotational bifurcation. We however chose not to include an additional free parameter into our model to keep the model as simple as possible. Moreover, the spring is not perfectly symmetric either because of both a non-zero helix angle and manufacturing imperfections. Importantly, with near certainty, the subjects placed their thumb with a finite offset from the geometric center of the spring. There are several other symmetry-breaking¹ aspects

¹The usage of the phrase “symmetry-breaking” in this context should not be confused with the symmetry-breaking feature of certain bifurcations. For example, the pitchfork bifurcation is sometimes referred to as a “symmetry-breaking” bifurcation. However, we use this term here purely in the physical sense of parity that is introduced in the thumb + spring + nervous system as a consequence of

that could be listed. Let us resort to the nomenclature used in the literature of solid mechanics and call these symmetry breaking features as *imperfections*. These imperfections can then be represented in the normal form equation by a lumped-parameter approximation, that we will refer to as the *imperfection parameter*. For the pitchfork bifurcation, the normal form equation that includes an imperfection is sometimes referred to a ‘versal unfolding’ of the pitchfork bifurcation normal form (Guckenheimer and Holmes, 1983). The imperfect subcritical pitchfork bifurcation normal form is,

$$\dot{\theta} = \alpha(F_s - K)\theta + \beta\theta^3 - \gamma\theta^5 + \epsilon \quad (5.1)$$

where the only new parameter ϵ is a lumped representation of all sources of imperfection. We have shown an example of a bifurcation diagram for a small value of the imperfection parameter in Figure 5.1. Note that the normal form equation given in Equation (5.1) is an example of a 2-parameter normal form that exhibits both codimension - 1 and codimension - 2 bifurcations. Both F_s and ϵ control the nature of bifurcations that occur in this system. The normal form equation that we used in the work presented in this chapter (e.g., Equation (3.1) on Page 58) then becomes a special case of Equation (5.1) with $\epsilon = 0$. In fact, when the constraint of symmetry is removed, the bifurcation that happens at $F_s = K$ and $\epsilon = 0$ is a codimension-2 bifurcation, that can be achieved only by fine-tuning two parameters (F_s and ϵ) simultaneously.

With this discussion about the unfolding of the subcritical pitchfork bifurcation normal form, it becomes apparent that the unfolding is a better representation of the experiment. However, we avoided using the more general form of the equation in an attempt to keep the model simple by minimizing the number of free parameters.

Nevertheless, there is an important feature of the experiment that could be revealed by the inclusion of this imperfection parameter. That is, an explanation for the trial-to-trial variability that we observed in F_s . One way to address this variability is by making the imperfection parameter a random variable. Namely, we can hypothesize that the effective imperfection in the system is primarily a consequence of placement of the thumb on the endcap. Much like the studies on finger-pointing (e.g., Körding and Wolpert, 2004), we could quantify the distribution of errors in effective center of pressure that subjects commit when asked to press down at a specified point on a flat surface. We can then test our hypothesis that the imperfection parameter and in turn, trial-to-trial variability is primarily driven by errors in thumb placement by making the imperfection parameter a random variable that is drawn from the experimentally measured distribution of thumb center-of-pressure. This way, we can address trial-to-trial variability without introducing any additional free parameters in our model. Going a little further, it is readily seen that the increase in variability of F_s after loss of thumbpad sensation has a potential explanation in terms of multisensory integration — loss of redundant sensors naturally increase noisiness of state estimates!

manufacturing and loading offsets.




Figure 5.1: The bifurcation diagram of a subcritical pitchfork bifurcation with the inclusion of a small imperfection. It is seen that the figure looks somewhat like the no-imperfection bifurcation diagram, but there is no longer a well-defined pitchfork bifurcation at $F_s = K$. The pitchfork bifurcation that was present for $\epsilon = 0$ now appears “blurred”. It is also seen that the desired stable configuration ($\theta = 0$) now loses stability a little sooner than if there was no imperfection.

To address the imperfection introduced due to the precedence of a lateral buckling is a harder proposition. To estimate the effective imperfection due to the lateral buckling requires modeling both the lateral and rotational buckling using more sophisticated models than a 1D differential equation. A possible approach to model this could be using discrete approximations of a buckling spring such as an inverted triple-pendulum with viscoelastic joints. Other uses of a more sophisticated model are outlined in Section 5.5 on Page 101. However, the development, calibration and analyses of more sophisticated models is beyond the intended scope of this thesis and better suited for future studies to explore.

5.1.2 Trial-specific parameter estimation

Since the validity of a low-dimensional model on average has already been established in the work presented in this thesis, we can now turn our attention to expanding the utility of this modeling approach as a sophisticated tool for quantifying dynamic manipulation ability in able and impaired populations. Instead of trying to predict average behavior using the low-order model, we can use it as a tool to quantify individual behavior. Specifically, instead of calculating the parameters K , α , β , γ and ϵ (if we include imperfections) like we did in Chapters 2 and 3, we can use time-series analysis techniques to estimate these parameters for individual trials, for individual subjects, or for various clinical conditions.

It is not immediately apparent how best to pose or carry out this parameter estimation problem. However, future work could explore the possibility of trying to use techniques such as extended Kalman filters, maximum likelihood estimators, quasi maximum likelihood estimators, or even genetic algorithms (Alcock and Buggage, 2004; Bishwal, 2003; Hurn and Lindsay, 1999; Hurn et al., 2003; McDonald and Sandal, 1999; Nielsen et al., 1999; Singer, 2002; Timmer, 2000).

Developing robust parameter estimation techniques and appropriate experimental methods (appropriate for parameter estimation) can then create a very useful way to quantify dynamic manipulation. Although a scalar metric like F_s is an attractive option, a multidimensional metric might be more sensitive to and provide more robust tools for quantifying dynamic manipulation ability. For example, although we did not find any change in F_s after training, it might be the case that training caused a systematic change in some facet of dynamic manipulation ability that is not sensitively measured by F_s . Such multidimensional metrics if developed could also find a useful application in clinical diagnosis of hand impairments.

5.2 Proportional-derivative feedback controllers

In Chapter 3 we explicitly defined a “controller” in order for us to explore the effects of sensory noise and time-delays on multisensory integration. As an initial attempt, and to keep our model simple, we chose a constant proportional feedback controller. We also discussed in Section 3.2.1 how the physical/robotic implemen-

tation of a constant proportional feedback control action can be a non-trivial effort. However, the use of a simple proportional controller with no derivative control action can be construed as an over-simplification of the controller, especially in the presence of time-delays in the system. This is because a proportional-derivative (PD) feedback controller can provide some degree of time-delay compensation, i.e., anticipatory-like control because of the derivative action (Doyle et al., 1992; Ogata, 2002). Clearly, a PD controller cannot predictively control the system since it is still reliant on sensory feedback, but the derivative control action can resemble predictive control, at least on short time-scales. Despite this apparent advantage and validity of using a PD controller, there are several reasons why we resorted to a purely proportional controller. Here, we outline some of those reasons:

1. As we discussed in some detail in Section 3.2.1 on Page 58, there is unfortunately no ‘natural’ way to separate the controller from the plant in our normal form model. This separation is probably best carried out in a detailed model of the experimental task (namely, spring compression using the thumbpad), which is not within the scope of this thesis. However, we used the language of ‘controller’ and ‘plant’ solely to provide a physical interpretation for the entry of noise and time-delays in our model. In fact, the proportional controller can be regarded as a 2-term Taylor series truncation of a more generalized controller, while the PD controller is a 3-term truncation. In summary, the important message is that the work presented in this thesis does not aim to identify or dwell on the structure of the controller used by the nervous system, but aims only to identify the effect of sensory noise and time-delays on multisensory integration.
2. The second reason for not using a PD controller is because of additional unknown / unmeasurable / non-estimable free parameters that are introduced in the model. The spirit of this thesis is to develop models with the least possible free parameters (in fact, just one free parameter) without sacrificing explanatory power of the model at the same time. Extra free parameters not only make the model more complex, but can restrict the generality of our results, thus casting doubt on the validity of the model itself. The burden of unknown or non-estimable parameters associated with a PD controller will become self-evident in the simulation results that we present below in Section 5.2.1.
3. It is tempting to include a PD controller and to estimate any unknown parameters as part of the proposed future work. However, we go even further in Section 5.5 on Page 101 and outline the steps needed to develop an optimal feedback controller for the task of spring compression. In light of the discussion on optimal feedback control for sensorimotor integration in Section 1.1 starting on Page 1, the necessity for further pursuing a PD controller becomes superfluous. Importantly, an optimal feedback controller is a superset of a PD controller.

Notwithstanding the three reasons cited above for using only a proportional controller and not a PD controller, we performed a preliminary investigation of the effect of using a PD controller. The details of this preliminary investigation and its results are presented below. This preliminary investigation revealed the following salient aspects of a PD controller:

1. Our original finding (Chapter 3) that a task-optimal multisensory integration strategy automatically leads to a selective use of vision depending on the availability of thumbpad sensors remained unchanged.
2. The model parameters that best agreed with experimental data led to an extra time-delay for non-digital sensors over and above neurophysiologically known sensor activation and nerve-conduction delays. The exact value of this additional time-delay was 40ms, not very different from our previous finding of a 35ms time-delay using a proportional feedback controller.
3. There were at least two additional free parameters that were introduced, despite several simplifying assumptions made in order to minimize extra free parameters. This led to a total of three free parameters in our model. However, the results did not show any dramatic qualitative change for a wide range of values of these free parameters. It should be noted that the exact values for these free parameters were chosen by hand-tuning them.

5.2.1 Proportional-derivative control model

Introducing a PD controller without the addition of too many new free parameters required a few simplifying assumptions as listed below:

1. We assumed that each of the three sensory channels — thumbpad sensors, nondigital mechanoreceptors, and vision — could sense ‘angular velocity’ ($\dot{\theta}$) in addition to the angle (θ) of the endcap.
2. To prevent the addition of three additional sensory weights for combining the three velocity signals, we assumed that the weights used for combining the velocity signals from the three sensory channels are the same weights used to combine the angular sensory signals from them. To put it simply, we assumed that $\omega_i^{\text{angular velocity}} = \omega_i^{\text{angle}}$. This helped us restrict the optimization problem for finding task-optimal sensory weights to a 2D parameter-space and thus enabling a global parameter search.
3. Since, it is not possible to estimate noise variances associated with the velocity sensory signals and these parameter values were not found in existing literature, we opted to assume that the velocity signals were noise-free. However, the presence of noise in the angle sensation causes the true dynamics of the system (and not just the sensory signals) to still remain noisy. Nevertheless, it is important to recognize that by assuming noise-free velocity sensors,

we bias the results to show a greater beneficial effect of including the derivative control action. If we had allowed for noise in the velocity sensations as well, then the relative beneficial impact of the derivative control action can only become lesser, since derivative control amplifies noise. Thus, it is safe to conclude that we do not bias the results of the PD control simulations to be close to the proportional control simulations by simply disregarding noise in velocity sensation.

The PD control model that we developed using the above assumptions is shown in Equation (5.2) below (compare this to the proportional controller in Equation (3.3) on Page 58):

$$\begin{aligned}
 \dot{\theta} &= \overbrace{\alpha F_s \theta + \beta \theta^3 - \gamma \theta^5}^{\text{thumb + spring dynamics}} + \overbrace{u(\hat{\theta}, \hat{\dot{\theta}})}^{\text{PD control action}} \\
 u(\hat{\theta}, \hat{\dot{\theta}}) &= \alpha K(\theta_{\text{desired}} - \hat{\theta}) - K_D \hat{\dot{\theta}} \\
 \theta_{\text{desired}} &= 0 \text{ without loss of generality} \\
 \overbrace{\hat{\theta}(t, t - \tau_1, t - \tau_2, t - \tau_3)}^{\text{angle estimate}} &= \sum_{i=1}^3 \omega_i \overbrace{(\theta(t - \tau_i) + \sigma_i \nu_i(t))}^{\text{noisy, time-delayed, unbiased sensors}} \\
 \nu_i &\sim \overbrace{N_i(0, 1)}^{\text{Gaussian white noise processes}} \\
 \overbrace{\hat{\dot{\theta}}(t, t - \tau_1, t - \tau_2, t - \tau_3)}^{\text{angular velocity estimate}} &= \sum_{i=1}^3 \omega_i \overbrace{\dot{\theta}(t - \tau_i)}^{\text{non-noisy sensors}}
 \end{aligned} \tag{5.2}$$

where the only new parameter introduced into the system is the derivative control gain, namely, K_D . However, caution is called for when introducing derivative control in a first-order dynamical system. The need for caution is readily seen if we examine the equations above in the absence of noise and time-delays. Then, the only effect of the derivative control action would be to scale the parameters α , β and γ by a factor of $1/(1 + K_D)$. This observation was used as a guiding principle in tuning the parameters of this model so that the results of sensory occlusion resembled experimental measurements.

5.2.2 Parameter values for proportional-derivative control model

We used the same values for the parameters $\hat{\alpha} = \alpha/\gamma$, $\hat{\beta} = \beta/\gamma$, K , τ_i , and σ_i ($i = 1, 2, 3$) as used previously in Chapter 3: $\hat{\alpha} = 0.006855$, $\hat{\alpha} = 0.2766$, $K = 3.3$, $\tau_1 = 65\text{ms}$, τ_2 was allowed to be between 65ms and 120ms (the only free parameter with a proportional controller), $\tau_3 = 120\text{ms}$, $\sigma_1 = 0.0007$, $\sigma_2 = 0.0031$ and $\sigma_3 = 0.0009$.

However, K_D and γ are now additional free parameters. It is clear why K_D is a free parameter – we determined the proportional control gain (K) based solely on the maximum observed compressive load in our experiments, but this is not possible for the derivative control gain. Hence, we made this a tunable, free parameter.

We saw above that if we disregarded time-delays and noise, the introduction of K_D causes the previous value of $\gamma = 385$ to be scaled by a factor of $1/(1 + K_D)$. The parameter γ is representative of the time-scale of the unstable dynamics in the system. Thus, reducing γ in effect renders the system more robust to noise and time-delays, which would in turn lead to model performance (F_s) that is unrealistically better than experimental data. To counter this effect of a derivative control action, we tuned γ to be sufficiently high so that the results of our simulations were comparable to experimental data.

By hand-tuning K_D and γ to best replicate experimental data, we found the following values for them: $K_D = 0.6$ and $\gamma = 565$. With this elaborate preamble, we are now ready to examine the results of using numerical optimization on this model to find task-optimal sensory weights and the resultant performance, i.e., F_s .

5.2.3 Results of *Best-fit* simulations

The time-delay for non-digital mechanoreceptors (τ_2) that yielded results in best agreement with experimental data, was 105ms, i.e., an additional 40ms time-delay over and above well-known sensor-activation and nerve-conduction delays for non-digital mechanoreceptors. Table 5.1 summarizes results of the *Best-fit* simulation using a PD controller. It is readily seen that vision is selectively useful only if thumbpad sensation was absent.

Table 5.1: This table presents results of the *Best-fit* simulations using a PD controller. The results reveal that a task-optimal strategy explains multisensory integration during dynamic manipulation. The last three columns show the task-optimal sensory weights for each sensory occlusion condition. The *Best-fit* simulations required fine tuning of three free parameters – τ_2 , K_D and γ as explained in Section 5.2.2 on Page 96.

Sensory state	F_s	ω_1	ω_2	ω_3
No vision, No nerve-block	3.26	0.95	0.05	–
With vision, No nerve-block	3.26	0.97	0	0.03
With vision, Nerve-block	2.91	–	0.17	0.83
No Vision, Nerve-block	2.52	–	1	–

It must be emphasized that the close numerical agreement with experimental measurements (See results shown in Figure 3.2 on Page 68) is due to fine-tuning

all three free parameters in this model. This is unlike the proportional control model, where we tuned only one free parameter. It is not readily apparent how these additional free parameters can be independently estimated to verify and validate the generality and impact of the results obtained using a PD controller. Nevertheless, it is worth mentioning that the selective use of vision was not affected by allowing the parameters K_D and γ to vary, although the results no longer numerically match experimental measurements (data not shown). This result serves to further strengthen our previous conclusion in Chapter 3, namely, task-optimal sensory weighting robustly explains multisensory integration strategies used by the nervous system during dynamic manipulation. Needless to say, an exhaustive assessment of the model’s sensitivity to parameter values can and should be carried out.

5.3 “Computational” time-delay — artifact or real?

We now turn our attention from refining the model to the capability of the model to generate testable hypotheses. We saw in Chapter 3 that our model had one free parameter that was tuned to find a *Best-fit* simulation that best reproduced experimental data quantitatively. However, this free parameter was not without physical interpretation. This free parameter represented the uncertainty in past experimental measurements of time-delay associated with non-digital sensors. Specifically, we posited that the value of this free parameter represented time-delays over and above well-known sensor and nerve-conduction time-delays for non-digital sensors. Recall that non-digital sensors is a term we used to describe all sensors that are not located in the thumb, which included, muscle spindles, Golgi tendon organs, and cutaneous receptors in areas of the skin outside the thumb. Specifically, we noted that the non-digital sensors were non-collated with respect to the thumbpad-object interface and might involve extra neural computation compared to collocated sensors like thumbpad tactile sensors or vision. Specifically, we noted that the combination of being non-located and being more noisy than other sensors might cause sensory processing of non-digital sensory information to incur significant “computational” time-delays. In agreement with our admittedly heuristic reasoning, we found that the free parameter reflected an additional 35ms computational time-delay for non-digital sensors. Interestingly, we saw that this extra delay is a neurophysiologically tenable value, falling within the range of time-delays reported by other researchers who examined slip-grip responses after administering a digital nerve-block. Nevertheless, it remains to be seen if this extra time-delay is purely an artifact of our simple model or if it has any physical correlate.

We propose here a potentially simple experiment to test a prediction of our model that relies on the presence of this extra time-delay. We predict that if the time-delay for visual feedback is artificially increased by 20ms, then the present experimental observation of selective use of vision will almost disappear. The

rationale for this hypothesis is based on additional simulations we carried out using the same approach as in Chapter 3. The only difference was, we set the time-delay for vision in the visual-delay simulation to be $\tau_3 = 140\text{ms}$, an additional 20ms time-delay over what was used by us previously. The results of this simulation are summarized in Table 5.2. Hence, we can see that the model predicts that the

Table 5.2: This table presents testable predictions of our model to determine whether “computational” time-delays really exist for non-digital sensors. These are the results of numerical optimization if vision was further delayed artificially by 20ms. The last three columns show the task-optimal sensory weights for each sensory occlusion condition. It is readily seen that after introducing 20ms extra time-delay for vision, the occlusion of vision no longer has any impact on performance.

Sensory state	F_s	ω_1	ω_2	ω_3
No vision, No nerve-block	3.23	0.96	0.04	—
With vision, No nerve-block	3.24	0.95	0.03	0.02
With vision, Nerve-block	2.55	—	0.91	0.09
No Vision, Nerve-block	2.54	—	1	—

absence of vision has almost no impact on performance if we somehow delayed vision externally. This can be achieved externally by providing visual feedback by use of 3D LCD display goggles whose video inputs are delayed signals from miniature cameras that could be attached to the subject. Thus, we can relatively cleanly evaluate whether the “computational” time-delay is real or a model artifact.

5.4 Carpal tunnel syndrome vs. multiple sclerosis

To parallel the testable neurophysiological hypothesis we presented above, the model of multisensory integration can also be used to generate testable hypotheses pertaining to certain neurological diseases. Specifically, we predicted that patients with two debilitating neurological diseases that share very similar symptoms in their early stage, but have completely different pathophysologies — multiple sclerosis vs. carpal tunnel syndrome — will possess measurably different effects of sensory occlusion in our experimental task.

Multiple sclerosis (MS) is a devastating neuro-degenerative disease that causes loss of the myelin sheath around axons of neurons in the brain and spinal cord (Compston and Coles, 2002; McDonald, 1999; Kandel et al., 2000), leading to severe sensorimotor dysfunction, cognitive dysfunction or even death in some cases. However, the principal effect of demyelination even in the early stages of the disease is a slowing down of nerve-conduction, leading to larger time-delays everywhere in

the nervous system. Interestingly, one of the earliest symptom of multiple sclerosis is poor “coordination” that manifests itself as clumsiness in handling objects and occasionally an increased propensity of loss of balance (Compston and Coles, 2002; Lassmann, 1999; Lublin and Reingold, 1996; McDonald and Ron, 1999; Namerow, 1968; Smith and McDonald, 1999).

Carpal tunnel syndrome (CTS) is a prevalent impairment of the hand (Atroshi et al., 1999) caused by compression of the median nerve at the wrist leading to pain, numbness and in its advanced stages, weakness of the hand (Kandel et al., 2000; Atroshi et al., 1999; Katz and Simmons, 2002). In contrast to MS, the effects of CTS are peripheral, localized to the hand, and not as devastating as MS although in advanced stages of CTS, the impact on lifestyle can be severe. CTS unfortunately is a broad term for a variety of clinical symptoms caused by compression of the median nerve. Consequently, there is still no consensus on how best to clinically define CTS (Katz and Simmons, 2002; Rempel et al., 1998). However, what is undeniable is the immediate and early consequence of median nerve compression — it affects tactile sensory signals from all but the 5th digit (commonly known as the “little finger”). Specifically, due to nerve compression, both the noisiness and time-delay of tactile sensory signals are affected (Katz and Simmons, 2002; Rempel et al., 1998).

Based on these known pathophysiologies of the two diseases, we carried out simulations using our model of multisensory integration. For MS, we added 20% extra time-delay for all sensory channels (recall that our model does not even consider motor noise / time-delays) to reflect the overall significant slowing of nerve-conduction, especially in the faster nerves of sensory channels such as tactile / visual or non-digital sensors (Kandel et al., 2000). To simulate CTS, we left untouched all but thumbpad sensors in our model. For thumbpad sensors, we increased the variance by 3 and the time-delay by 20%. Keep in mind that these numbers are by no means accurate or constant across all patients, but a first-attempt at modeling the functional consequences of these diseases. By simulating these two conditions and examining the effect of sensory occlusion, we found astounding differences between MS and CTS. The results, i.e., testable model predictions are summarized in Table 5.3. We find that when all sensory channels are intact, CTS causes loss of performance compared to unimpaired subjects. But, the finding of greatest interest is the qualitative difference between the effect of MS and CTS. The CTS model predicts the same selective use of vision that we observed in unimpaired subjects, but the MS model predicts that this trend will vanish. In other words, much like our visual delay model, the MS model predicts a severe degradation in the usefulness of vision. It is very important to note that in addition to the increased propensity to drop objects, another hallmark of MS is the clinical observation that vision appears to degrade the most during the early stages of the disease (Compston and Coles, 2002; Lublin and Reingold, 1996; McDonald and Ron, 1999).

We end this section by noting that a recent statistical followup study found that “physical function” as quantified by clinical measures of balance control was an

Table 5.3: This table summarizes testable model predictions for the functional consequences of MS vs. CTS. The notation for the performance metric and sensory weights remain unchanged from before.

Sensory	MS				CTS			
State	F_s	ω_1	ω_2	ω_3	F_s	ω_1	ω_2	ω_3
Blind, Intact	3.22	0.97	0.03	–	3.14	0.94	0.06	–
Seeing, Intact	3.22	0.97	0.03	0	3.15	0.74	0.06	0.20
Seeing, Numb	2.08	–	0.99	0.01	2.80	–	0.21	0.79
Blind, Numb	2.08	–	1	–	2.58	–	1	–

excellent predictor of whether or not a person would be a victim of dementia and Alzheimer disease (Wang et al., 2006). Much like our task of dynamic manipulation during spring compression, balance control is a dynamic sensorimotor control task, with a primary goal of avoiding instabilities. Hence, it will be useful to study if performance in our task is a good predictor of degenerative conditions and systemic neurological diseases such as Alzheimer’s disease, Parkinson’s disease or dementia.

5.5 Optimal feedback control

So far, we have not discussed the possibility of developing a detailed model of this experiment — a “first-principles” model according to the definition we used earlier in the thesis. Clearly, current knowledge about either the nervous system or even the mechanics of the thumb are too incomplete to attempt including a detailed model of them. But, we can model the mechanics of the spring fairly well based on known mechanics of spring buckling. There are several excellent treatises that present detailed derivations of dynamical equations of spring buckling, typically using numerical techniques or partial differential equations to describe the dynamics of equivalent columns (Becker et al., 2002; Becker and Cleghorn, 1992; Chassie et al., 1997; Haringx, 1948, 1949a,b,c,d,e; Timoshenko, 1961; Wahl, 1963). Importantly, we can rely on the experimental observation of the mode of spring buckling that is encountered when subjects are asked to maximize compressive spring force to develop a suitable discretization of the spring. Specifically, for the fixed-guided buckling mode that is observed, an inverted triple pendulum with viscoelastic joints will suffice to capture the dynamics of the task. The parameters of this reduced order model of spring buckling can be estimated using experiments that can be performed on the real spring by compressing it using known end-condition and loads on the endcap. The Southwell method of calculating true buckling loads for elastic structures is an example of such experimental techniques

(Southwell, 1932; Haringx, 1949e).

It is important to emphasize the need for caution when using such low order models of elastic (in)stability. For example, consider the lowest order approximation of a buckling column — an inverted simple pendulum with a viscoelastic pivot joint at the base. Under a dead-load, this 1-link approximation undergoes a pitchfork bifurcation much as the case for a buckling column under a dead-load (see Section 2.3.1 on Page 37). However, under a follower load, a real column would undergo a Hopf bifurcation and start ‘fluttering’ as the magnitude of the load increases, but paradoxically, the 1-link approximation would be stable even for arbitrarily large loads so long as it is a perfect follower load. This paradox of the 1-link approximation is called the Pflüger’s paradox (Lobas et al., 2002; Lobas and Lobas, 2004). It was subsequently shown that this paradox is an error that arises from the 1-link simplification and a 2-link model can exhibit both a pitchfork and a Hopf bifurcation.

To summarize, we have reliable ways of experimentally quantifying the overall dynamical behavior of the thumb + spring + nervous system, and have a reliable way of modeling the spring as an inverted triple pendulum with viscoelastic joints. Importantly, we can adopt optimal feedback control as a working hypothesis to represent the output of the nervous system. We can then approximate the action of the thumbpad on the spring’s endcap using a point force whose location and direction with respect to the endcap can be moved in a band-limited way (i.e., with finite bandwidth of actuation). This places us in a unique situation of being able to test whether a specific choice of control strategy can mimic the overall behavior that is measured experimentally. Hence, we can seek to develop a well-posed optimal feedback control problem, say as an iterative linear quadratic Gaussian regulator (Todorov et al., 2005) for the simple first-principles model of dynamic manipulation at the edge of instability.

It is tempting and desirable to extend the above proposed work of optimal feedback control to include more complicated models with thumb-like kinematic models to actuate the spring instead of a point load. However, even with the simplified model described above, the state space is 6-dimensional and pushing the boundaries of technological limitations in using numerical discretization methods for designing optimal feedback controllers. That is because, designing an optimal feedback controller entails the ‘*curse of dimensionality*’ (Bellman, 1961; Basar, 2001, a term coined by Richard E. Bellman), because, for an optimal feedback controller we want to find the optimal path from everywhere in the phase space of the system to the desired (goal) location. For example, to discretize a 3-dimensional phase space with 100 points along each dimension would entail a total of 10^6 grid points. However, for a 6-dimensional phase space, this number balloons into 10^{12} grid points! This is called the (formidable) *curse of dimensionality*. Recently, some alternative methods using hierarchical controllers have been proposed that turn out to be nearly-optimal in numerical simulations. Below, we outline some potential implications of our work to this emerging view of hierarchical nearly-optimal feedback controllers.

5.6 Hierarchical nearly-optimal feedback control

This is the concluding section of this thesis and is the most speculative component of this chapter. It aims to provide approximate research directions for some problems in optimal feedback control that face severe technical challenges because of the curse of dimensionality.

Hierarchical control is a phrase used to describe a specific structure of the controller — namely the presence of multiple levels of controllers. Let us consider an example of a neural control system in a human. Sensorimotor control occurs simultaneously at many levels. ‘Lower-level’ neural circuits in the spinal cord and brainstem receive sensory input from the periphery and generate motor commands before other areas of the brain have had time to react to that input. The ‘higher-level’ areas of the brain then interact with this fused system of the lower-level controllers + periphery + world. This notion of hierarchical control is particularly relevant to neural control systems and there is a large body of evidence that neural control might be organized hierarchically (Collins and Stewart, 1993; Darian-Smith et al., 1999; Li et al., 2004; Loeb et al., 1999; Todorov et al., 2005; d’Avella and Bizzi, 2005; Scott, 2004; Full and Koditschek, 1999; Dickinson et al., 2000) in a manner similar to our simple allegory above.

Recently hierarchical control has been suggested as a means to overcome the curse of dimensionality in optimal feedback control (Li et al., 2004; Todorov et al., 2005). Hierarchical control can help reduce the dimensionality of the control by breaking up the problem of optimal feedback control into two parts according to Todorov et al. (2005). The first part is a low-level controller that receives the state of the plant (body + world, say, x) and generates a task-relevant low-dimensional representation of x that it transmits to a high-level controller. The high-level controller then computes and transmits control action that could be optimal in the low-dimensional representation of x . This can be roughly understood as follows. The high-level controller plans and “optimizes” the control action in terms of relevant task-features alone. The low-level controller is a “dumb” controller that just performs stereotypical transformations from sensory signals to low-dimensional task-relevant features as well as transformations from high-level control commands to detailed motor-unit commands to implement the high-level control action.

This low-dimensional representation of only task-relevant features is motivated by the observation that optimal feedback control minimally intervenes, i.e., an optimal feedback controller responds to disturbances only if it affects task-relevant features. There are a large number of names assigned to this low-dimensional representation — template, task-relevant feature, task-relevant subspace, high-level task goal, task feature, feature, etc. (e.g., Dickinson et al., 2000; Full and Koditschek, 1999; Li et al., 2004; Loeb et al., 1999; Todorov et al., 2005). All of those names refer to the fact there is some low-dimensional aspect of the entire system that the high-level controller is interested in “optimally” controlling. None of these researchers claim that this form of hierarchical control is optimal in any way, but some numerical examples have found that this form of hierarchical control

can be nearly-optimal, if the task-relevant low-dimensional subspace is appropriately chosen. This is the point where the work presented in this thesis has the potential to make a significant contribution in the future. Below, we put forth a proposal for automatically choosing generic² low-level “feature extraction” methods or “anchors” that transform the slew of high-dimensional sensory feedback into a low-dimensional task-relevant representation.

We saw in Section 2.4.3 on Page 51 that the nervous system is adept at detecting incipient instabilities. This motivates us to propose that the nervous system is particularly adept at, and has in-built or learned mechanisms at the level of the spinal cord or the brain-stem (given the short time-scales of the responses) for detecting when the dynamics of body + world is close to a bifurcation. Techniques for bifurcation detection are reasonably well established mathematically (Aguirre and Torres, 2000; Alvarado et al., 1994; Chen et al., 2000; Kim and Abed, 2000, 2001; Moreau et al., 2003; Rico-Martinez et al., 2003) and even in use for detecting voltage collapse and blackouts in power electric systems (Alvarado et al., 1994). These techniques are often based on the sometimes counter-intuitive effects of noise in nonlinear dynamical systems close to a bifurcation (Collins et al., 1995, 1996; Jaramillo and Wiesenfeld, 2000; Jeffries and Wiesenfeld, 1985; Jung and Wiesenfeld, 1997; Moss and Wiesenfeld, 1995; Wiesenfeld, 1985; Wiesenfeld and McCarley, 1990; Wiesenfeld and McNamara, 1986; Wiesenfeld and Moss, 1995). Given the ubiquitous nature of low-dimensional dynamics that ensue near a bifurcation point, a natural choice for a method to extract “task-relevant” low-dimensional features is based on detection of these bifurcations and using possibly stereotypical responses to control or avoid specific bifurcations. This heuristic and speculative proposal of ours gains larger applicability from two other observations.

1. Feedback control and specifically task-optimal feedback control becomes particularly important only when trying to perform a task at the edge of instability. In fact, Stein (2003) argues that the real challenge in engineering control system design lies in controlling inherently unstable plants that operate at their limit of performance. This brings a very large number of performance-critical control tasks into a regime where low-dimensional task-feature extraction using bifurcation detection techniques and nearly-optimal hierarchical feedback control become applicable.
2. It has been observed in experiments of stick balancing that the nervous system chooses to poise itself close to an instability (Cabrera and Milton, 2002, 2004b) and the authors suggest that there are benefits to doing this since a nearly unstable system is sensitive to errors in control, thus enabling better sensory monitoring of the system. They propose other advantages to posing the system at the edge of instability such as on-off intermittency, that enables desirable control actions at time-scales shorter than feedback delays

²We call our proposal *generic* in the sense that it relies on ubiquitous phenomena that are found in most nonlinear dynamical systems.

in the system. In yet another context, Milton et al. (2004) speculate that “cortical neural networks exist on the edge between epileptic seizures on the one hand, and quiescence on the other”. Once again, it is apparent that these observations pertaining to neural control and neural systems enhance the applicability of our proposal.

Hence, we call for designing experiments and mathematical models to test the applicability of nearly-optimal hierarchical feedback controllers that rely on bifurcation detection techniques for extracting low-dimensional features. If such capability is indeed present in biological systems, that might help partially explain the present gap in performance between robotic and biological systems.

5.7 Conclusion

Starting from an apparently simple effort to unambiguously define and quantify “dexterity”, we found several intriguing results that are mathematically, neuro-physiologically and clinically interesting, novel and critical. Going further, we saw how several short term goals that build upon the work presented here can provide rewards both in terms of better scientific appreciation of and insight into complex dynamical systems as well as vital tools for clinical use. Finally, we provided a admittedly speculative, but plausible proposition that can have far reaching effects to not just sensorimotor control, but robotic and general engineering control systems as well. I conclude this thesis reiterating that “living on the edge” does have its benefits!

BIBLIOGRAPHY

- L. A. Aguirre and L. A. B. Torres. Control of nonlinear dynamics: Where do models fit in? *International Journal of Bifurcation and Chaos*, 10(3):667–681, 2000.
- J. Alcock and K. Burrage. A genetic estimation algorithm for parameters of stochastic ordinary differential equations. *Computational Statistics & Data Analysis*, 47(2):255–275, 2004.
- R. M. Alexander. *Optima for animals*. Princeton University Press, Princeton, N.J., rev. edition, 1996.
- R. M. Alexander. Design by numbers. *Nature*, 412(6847):591, 2001.
- R. Altman, E. Asch, D. Bloch, G. Bole, D. Borenstein, K. Brandt, W. Christy, T. D. Cooke, R. Greenwald, M. Hochberg, D. Howell, D. Kaplan, W. Koopman, S. Longley, H. Mankin, D. J. McShane, T. Medsger, R. Meenan, W. Mikkelsen, R. Moskowitz, W. Murphy, B. Rothschild, M. Segal, L. Sokoloff, and F. Wolfe. Development of criteria for the classification and reporting of osteoarthritis - classification of osteoarthritis of the knee. *Arthritis and Rheumatism*, 29(8):1039–1049, 1986.
- R. D. Altman and R. Moskowitz. Intraarticular sodium hyaluronate (hyalgan (r)) in the treatment of patients with osteoarthritis of the knee: A randomized clinical trial. *Journal of Rheumatology*, 25(11):2203–2212, 1998.
- F. Alvarado, I. Dobson, Y. Hu, M. A. Pai, and N. K. Pal. Computation of closest bifurcations in power-systems. *IEEE Transactions on Power Systems*, 9(2):918–928, 1994.
- P. C. Amadio. Outcomes assessment in hand surgery. what’s new? *Clin Plast Surg*, 24(1):191–4, 1997.
- L. Arnold. *Random dynamical systems*. Springer monographs in mathematics. Springer, Berlin ; New York, 1998.
- I. Atroshi, C. Gummesson, R. Johansson, E. Ornstein, J. Ranstam, and I. Rosen. Prevalence of carpal tunnel syndrome in a general population. *Jama-Journal of the American Medical Association*, 282(2):153–158, 1999.
- A. S. Augurelle, A. M. Smith, T. Lejeune, and J. L. Thonnard. Importance of cutaneous feedback in maintaining a secure grip during manipulation of hand-held objects. *Journal of Neurophysiology*, 89(2):665–671, 2003.
- P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality - an explanation of $1/f$ noise. *Physical Review Letters*, 59(4):381–384, 1987.

- P. Bak, C. Tang, and K. Wiesenfeld. Self-organized criticality. *Physical Review A*, 38(1):364–374, 1988.
- T. Basar. *Control theory : twenty-five seminal papers, 1931-1981*. IEEE Press, New York, 2001.
- L. E. Becker and W. L. Cleghorn. On the buckling of helical compression springs. *International Journal of Mechanical Sciences*, 34(4):275–282, 1992.
- L. E. Becker, G. G. Chassie, and W. L. Cleghorn. On the natural frequencies of helical compression springs. *International Journal of Mechanical Sciences*, 44(4):825–841, 2002.
- R. E. Bellman. *Adaptive control processes: a guided tour*. N.J., Princeton University Press, Princeton, 1961.
- N. Berglund and B. Gentz. Pathwise description of dynamic pitchfork bifurcations with additive noise. *Probability Theory and Related Fields*, 122(3):341–388, 2002.
- N. Bernstein. *The co-ordination and regulation of movements*. Pergamon Press, 1967.
- A. Beuter. *Nonlinear dynamics in physiology and medicine*. Interdisciplinary applied mathematics ; v.25. Springer, New York, 2003.
- A. Bicchi. Hands for dexterous manipulation and robust grasping: A difficult road toward simplicity. *IEEE Transactions on Robotics and Automation*, 16(6):652–662, 2000.
- J. P. N. Bishwal. Maximum likelihood estimation in partially observed stochastic differential. system driven by a fractional brownian motion. *Stochastic Analysis and Applications*, 21(5):995–1007, 2003.
- E. Bizzi, N. Accornero, W. Chapple, and N. Hogan. Posture control and trajectory formation during arm movement. *J Neurosci*, 4(11):2738–44, 1984.
- E. Bizzi, N. Hogan, F. A. Mussaivaldi, and S. Giszter. Does the nervous-system use equilibrium-point control to guide single and multiple joint movements. *Behavioral and Brain Sciences*, 15(4):603–613, 1992.
- K. D. Brandt, G. N. Smith, and L. S. Simon. Intraarticular injection of hyaluronan as treatment for knee osteoarthritis - what is the evidence? *Arthritis and Rheumatism*, 43(6):1192–1203, 2000.
- E. Burdet, R. Osu, D. W. Franklin, T. E. Milner, and M. Kawato. The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature*, 414(6862):446–449, 2001.

- E. Burdet, K. P. Tee, I. Mareels, T. E. Milner, C. M. Chew, D. W. Franklin, R. Osu, and M. Kawato. Stability and motor adaptation in human arm movements. *Biological Cybernetics*, 94(1):20–32, 2006.
- J. L. Cabrera and J. Milton. Stick balancing: On-off intermittency and survival times. *Nonlinear Studies*, 11(3):305–317, 2004a.
- J. L. Cabrera and J. G. Milton. On-off intermittency in a human balancing task. *Physical Review Letters*, 89(15):–, 2002.
- J. L. Cabrera and J. G. Milton. Human stick balancing: Tuning levy flights to improve balance control. *Chaos*, 14(3):691–698, 2004b.
- A. M. Carlsson. Assessment of chronic pain .1. aspects of the reliability and validity of the visual analog scale. *Pain*, 16(1):87–101, 1983.
- G. G. Chassie, L. E. Becker, and W. L. Cleghorn. On the buckling of helical springs under combined compression and torsion. *International Journal of Mechanical Sciences*, 39(6):697–704, 1997.
- G. R. Chen, M. L. Moiola, and H. O. Wang. Bifurcation control: Theories, methods, and applications. *International Journal of Bifurcation and Chaos*, 10(3):511–548, 2000.
- V. C. K. Cheung, A. d’Avella, M. C. Tresch, and E. 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.
- J. A. Clack. From fins to fingers. *Science*, 304(5667):57–58, 2004.
- J. J. Clark and A. L. Yuille. *Data fusion for sensory information processing systems*. The Kluwer international series in engineering and computer science ;; SECS 105;. Kluwer Academic Publishers, Boston, 1990.
- K. J. Cole and J. H. Abbs. Grip force adjustments evoked by load force perturbations of a grasped object. *J Neurophysiol*, 60(4):1513–1522, 1988.
- K. J. Cole, D. L. Rotella, and J. G. Harper. Tactile impairments cannot explain the effect of age on a grasp and lift task. *Experimental Brain Research*, 121(3):263–269, 1998.
- J. J. Collins and C. J. Deluca. Random walking during quiet standing. *Physical Review Letters*, 73(5):764–767, 1994.
- J. J. Collins and I. N. Stewart. Coupled nonlinear oscillators and the symmetries of animal gaits. *Journal of Nonlinear Science*, 3(3):349–392, 1993.
- J. J. Collins, C. C. Chow, and T. T. Imhoff. Stochastic resonance without tuning. *Nature*, 376(6537):236–238, 1995.

- J. J. Collins, T. T. Imhoff, and P. Grigg. Noise-enhanced tactile sensation. *Nature*, 383(6603):770–770, 1996.
- A. Compston and A. Coles. Multiple sclerosis. *Lancet*, 359(9313):1221–31, 2002.
- J. G. A. Croll and A. C. Walker. *Elements of structural stability*. Wiley, New York,, 1972.
- I. Darian-Smith, K. Burman, and C. Darian-Smith. Parallel pathways mediating manual dexterity in the macaque. *Exp Brain Res*, 128(1-2):101–8, 1999.
- A. d’Avella and E. Bizzi. Shared and specific muscle synergies in natural motor behaviors. *Proc Natl Acad Sci U S A*, 102(8):3076–81, 2005.
- B. L. Day and P. Brown. Evidence for subcortical involvement in the visual control of human reaching. *Brain*, 124:1832–1840, 2001.
- M. H. Dickinson, C. T. Farley, R. J. Full, M. A. R. Koehl, R. Kram, and S. Lehman. How animals move: An integrative view. *Science*, 288(5463):100–106, 2000.
- J. B. Dingwell, C. D. Mah, and F. A. Mussa-Ivaldi. Manipulating objects with internal degrees of freedom: Evidence for model-based control. *Journal of Neurophysiology*, 88(1):222–235, 2002.
- J. B. Dingwell, C. D. Mah, and F. A. Mussa-Ivaldi. Experimentally confirmed mathematical model for human control of a non-rigid object. *Journal of Neurophysiology*, 91(3):1158–1170, 2004.
- J. C. Doyle, B. A. Francis, and A. Tannenbaum. *Feedback control theory*. Toronto, New York Macmillan Pub. Co., 1992.
- D. J. D. Earn, P. Rohani, B. M. Bolker, and B. T. Grenfell. A simple model for complex dynamical transitions in epidemics. *Science*, 287(5453):667–670, 2000.
- M. S. El Naschie. *Stress, stability, and chaos in structural engineering : an energy approach*. McGraw-Hill, London ; New York, 1990.
- A. C. Eliasson, H. Forssberg, K. Ikuta, I. Apel, G. Westling, and R. Johansson. Development of human precision grip .5. anticipatory and triggered grip actions during sudden loading. *Experimental Brain Research*, 106(3):425–433, 1995.
- K. Engelborghs, T. Luzyanina, and D. Roose. Numerical bifurcation analysis of delay differential equations. *Journal of Computational and Applied Mathematics*, 125(1-2):265–275, 2000.
- K. Engelborghs, T. Luzyanina, and D. Roose. Numerical bifurcation analysis of delay differential equations using dde-biftool. *ACM Transactions on Mathematical Software*, 28(1):1–21, 2002.

- M. O. Ernst. Combining sensory information from vision and touch. *International Journal of Psychophysiology*, 54(1-2):57–58, 2004.
- M. O. Ernst and M. S. Banks. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–433, 2002a.
- M. O. Ernst and M. S. Banks. Using visual and haptic information for discriminating objects. *Perception*, 31:147–147, 2002b.
- M. O. Ernst and H. H. Bulthoff. Merging the senses into a robust percept. *Trends in Cognitive Sciences*, 8(4):162–169, 2004.
- A. G. Feldman. Functional tuning of nervous system with control of movement or maintenance of a steady posture .2. controllable parameters of muscles. *Biophysics-USSR*, 11(3), 1966a.
- A. G. Feldman. Functional tuning of nervous system during control of movement or maintenance of a steady posture .3. mechanographic analysis of execution by man of simplest motor tasks. *Biophysics-USSR*, 11(4):766, 1966b.
- A. G. Feldman and M. L. Latash. Testing hypotheses and the advancement of science: recent attempts to falsify the equilibrium point hypothesis. *Exp Brain Res*, 161(1):91–103, 2005.
- E. Fess. Chapter 12: documentation: essential elements of an upper extremity assessment battery. In J. M. Hunter, E. Mackin, and A. D. Callahan, editors, *Rehabilitation of the hand : surgery and therapy*, pages (xxvii, 1901 p.). Mosby, St. Louis, 4th edition, 1995.
- J. R. Flanagan and S. Lolley. The inertial anisotropy of the arm is accurately predicted during movement planning. *J Neurosci*, 21(4):1361–9, 2001.
- J. R. Flanagan and A. M. Wing. The role of internal models in motion planning and control: Evidence from grip force adjustments during movements of hand-held loads. *Journal of Neuroscience*, 17(4):1519–1528, 1997.
- T. Flash and N. Hogan. The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci*, 5(7):1688–703, 1985.
- H. Forssberg, A. C. Eliasson, H. Kinoshita, R. S. Johansson, and G. Westling. Development of human precision grip. i: Basic coordination of force. *Exp Brain Res*, 85(2):451–7, 1991.
- H. Forssberg, H. Kinoshita, A. C. Eliasson, R. S. Johansson, G. Westling, and A. M. Gordon. Development of human precision grip .2. anticipatory control of isometric forces targeted for objects weight. *Experimental Brain Research*, 90(2):393–398, 1992.

- H. Forssberg, A. C. Eliasson, H. Kinoshita, G. Westling, and R. S. Johansson. Development of human precision grip .4. tactile adaptation of isometric finger forces to the frictional condition. *Exp Brain Res*, 104(2):323–330, 1995.
- J. R. E. Fraser, T. C. Laurent, and U. B. G. Laurent. Hyaluronan: Its nature, distribution, functions and turnover. *Journal of Internal Medicine*, 242(1):27–33, 1997.
- S. Fuchs, R. Monikes, A. Wohlmeiner, and T. Heyse. Intra-articular hyaluronic acid compared with corticoid injections for the treatment of rhizarthrosis. *Osteoarthritis and Cartilage*, 14(1):82–88, 2006.
- R. J. Full and D. E. Koditschek. Templates and anchors: Neuromechanical hypotheses of legged locomotion on land. *Journal of Experimental Biology*, 202(23):3325–3332, 1999.
- S. C. Gandevia and D. I. McCloskey. Joint sense, muscle sense, and their combination as position sense, measured at distal interphalangeal joint of middle finger. *Journal of Physiology-London*, 260(2):387, 1976.
- A. M. Gordon, H. Forssberg, R. S. Johansson, A. C. Eliasson, and G. Westling. Development of human precision grip .3. integration of visual size cues during the programming of isometric forces. *Experimental Brain Research*, 90(2):399–403, 1992.
- V. L. Gracco and J. H. Abbs. Dynamic control of the perioral system during speech: kinematic analyses of autogenic and nonautogenic sensorimotor processes. *J Neurophysiol*, 54(2):418–32, 1985.
- G. A. Green. Understanding nsoids: from aspirin to cox-2. *Clin Cornerstone*, 3(5):50–60, 2001.
- P. L. Gribble and S. H. Scott. Overlap of internal models in motor cortex for mechanical loads during reaching. *Nature*, 417(6892):938–41, 2002.
- J. Guckenheimer and P. Holmes. *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*. Springer, New York, corr. 7th print. edition, 1983.
- C. Häger-Ross and R. S. Johansson. Nondigital afferent input in reactive control of fingertip forces during precision grip. *Exp Brain Res*, 110(1):131–41, 1996.
- H. Haken. Cooperative phenomena in systems far from thermal equilibrium and in nonphysical systems. *Reviews of Modern Physics*, 47(1):67–121, 1975.
- J. A. Haringx. On highly compressible helical springs and rubber rods, and their application for vibration-free mountings .1. *Philips Research Reports*, 3(6):401–449, 1948.

- J. A. Haringx. On highly compressible helical springs and rubber rods, and their application for vibration-free mountings .2. *Philips Research Reports*, 4(1):49–80, 1949a.
- J. A. Haringx. On highly compressible helical springs and rubber rods, and their application for vibration-free mountings .3. *Philips Research Reports*, 4(3):206–220, 1949b.
- J. A. Haringx. On highly compressible helical springs and rubber rods, and their application for vibration-free mountings .4. *Philips Research Reports*, 4(4):261–290, 1949c.
- J. A. Haringx. On highly compressible helical springs and rubber rods, and their application for vibration-free mountings .5. *Philips Research Reports*, 4(5):375–400, 1949d.
- J. A. Haringx. On highly compressible helical springs and rubber rods, and their application for vibration-free mountings .6. *Philips Research Reports*, 4(6):407–448, 1949e.
- C. M. Harris and D. M. Wolpert. Signal-dependent noise determines motor planning. *Nature*, 394(6695):780–4, 1998.
- N. Hogan. Adaptive-control of mechanical impedance by coactivation of antagonist muscles. *IEEE Transactions on Automatic Control*, 29(8):681–690, 1984a.
- N. Hogan. An organizing principle for a class of voluntary movements. *J Neurosci*, 4(11):2745–54, 1984b.
- P. L. Hudak, P. C. Amadio, C. Bombardier, A. Boland, T. Fischer, E. L. Flatow, G. M. Gartsman, D. S. Louis, T. Axelrod, R. Buchbinder, G. Hawker, R. Hotchkiss, J. Katz, T. Bedard, R. Lederman, D. Louis, C. McCormick, S. Odriscoll, D. Richards, R. Richards, and B. Simmons. Development of an upper extremity outcome measure: The dash (disabilities of the arm, shoulder, and head). *American Journal of Industrial Medicine*, 29(6):602–608, 1996.
- A. S. Hurn and K. A. Lindsay. Estimating the parameters of stochastic differential equations. *Mathematics and Computers in Simulation*, 48(4-6):373–384, 1999.
- A. S. Hurn, K. A. Lindsay, and V. L. Martin. On the efficacy of simulated maximum likelihood for estimating the parameters of stochastic differential equations. *Journal of Time Series Analysis*, 24(1):45–63, 2003.
- Y. P. Ivanenko, R. Grasso, M. Zago, M. Molinari, G. Scivoletto, V. Castellano, V. Macellari, and F. Lacquaniti. Temporal components of the motor patterns expressed by the human spinal cord reflect foot kinematics. *J Neurophysiol*, 90(5):3555–65, 2003.

- F. Jaramillo and K. Wiesenfeld. Physiological noise level enhances mechanoelectrical transduction in hair cells. *Chaos Solitons & Fractals*, 11(12):1869–1874, 2000.
- C. Jeffries and K. Wiesenfeld. Observation of noisy precursors of dynamical instabilities. *Physical Review A*, 31(2):1077–1084, 1985.
- R. S. Johansson. Sensory control of dexterous manipulation in humans. In A. M. Wing, P. Haggard, and J. R. Flanagan, editors, *Hand and brain : the neurophysiology and psychology of hand movements*, pages 381–414. Academic Press, San Diego, 1996.
- R. S. Johansson and I. Birznieks. First spikes in ensembles of human tactile afferents code complex spatial fingertip events. *Nature Neuroscience*, 7(2):170–177, 2004.
- R. S. Johansson and K. J. Cole. Sensory-motor coordination during grasping and manipulative actions. *Curr Opin Neurobiol*, 2(6):815–23, 1992.
- R. S. Johansson and K. J. Cole. Grasp stability during manipulative actions. *Can J Physiol Pharmacol*, 72(5):511–24, 1994.
- R. S. Johansson and G. Westling. Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects. *Exp Brain Res*, 56(3):550–64, 1984.
- R. S. Johansson and G. Westling. Significance of cutaneous input for precise hand movements. *Electroencephalogr Clin Neurophysiol Suppl*, 39:53–7, 1987.
- R. S. Johansson, C. Häger-Ross, and L. Backstrom. Somatosensory control of precision grip during unpredictable pulling loads .3. impairments during digital anesthesia. *Experimental Brain Research*, 89(1):204–213, 1992a.
- R. S. Johansson, R. Riso, C. Häger-Ross, and L. Backstrom. Somatosensory control of precision grip during unpredictable pulling loads .1. changes in load force amplitude. *Experimental Brain Research*, 89(1):181–191, 1992b.
- L. A. Jones. The assessment of hand function: a critical review of techniques. *J Hand Surg [Am]*, 14(2 Pt 1):221–8, 1989.
- M. I. Jordan. Computational aspects of motor control and motor learning. In H. Heuer and S. W. Keele, editors, *Handbook of perception and action: Motor skills*, pages xx, 617 p. Academic., London, 1996.
- P. Jung and K. Wiesenfeld. Too quiet to hear a whisper. *Nature*, 385(6614):291–291, 1997.

- E. R. Kandel, J. H. Schwartz, and T. M. Jessell. *Principles of neural science*. McGraw-Hill, Health Professions Division, New York, 4th edition, 2000.
- J. N. Katz and B. P. Simmons. Carpal tunnel syndrome. *New England Journal of Medicine*, 346(23):1807–1812, 2002.
- M. Kawato. Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9(6):718–727, 1999.
- J. A. S. Kelso. *Dynamic patterns : the self-organization of brain and behavior*. MIT Press, Cambridge, Mass., 1995.
- T. Kim and E. H. Abed. Closed-loop monitoring systems for detecting impending instability. *IEEE Transactions on Circuits and Systems I-Fundamental Theory and Applications*, 47(10):1479–1493, 2000.
- T. Kim and E. H. Abed. Stationary bifurcation control of systems with uncontrollable linearization. *International Journal of Control*, 74(5):445–452, 2001.
- D. C. Knill and A. Pouget. The bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12):712–719, 2004.
- V. B. Kolmanovskii and A. D. Myshkis. *Introduction to the theory and applications of functional differential equations*. Mathematics and its applications ;; v. 463; Variation: Mathematics and its applications (Kluwer Academic Publishers) ;; v. 463. Boston, Dordrecht, 1999.
- V. B. Kolmanovskii and V. R. Nosov. *Stability of functional differential equations*. Mathematics in science and engineering ;; v. 180;. Orlando, London, 1986.
- K. P. Körding and D. M. Wolpert. Bayesian integration in sensorimotor learning. *Nature*, 427(6971):244–247, 2004.
- K. P. Körding, S. P. Ku, and D. M. Wolpert. Bayesian integration in force estimation. *Journal of Neurophysiology*, 92(5):3161–3165, 2004.
- U. Kuchler and E. Platen. Strong discrete time approximation of stochastic differential equations with time delay. *Mathematics and Computers in Simulation*, 54(1-3):189–205, 2000.
- A. D. Kuo. An optimal control model for analyzing human postural balance. *IEEE Trans. Biomed. Eng.*, 42(1):87–101, 1995.
- A. D. Kuo. An optimal state estimation model of sensory integration in human postural balance. *J. Neural. Eng.*, 2(3):S235–49, 2005.
- H. Lassmann. The pathology of multiple sclerosis and its evolution. *Phil. Trans. R. Soc. Lond. B*, 354(1390):1635–1640, 1999.

- C. S. Li, C. Padoa-Schioppa, and E. Bizzi. Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron*, 30(2):593–607, 2001.
- W. Li, E. Todorov, and X. Pan. Hierarchical optimal control of redundant biomechanical systems. In *Conference Proceedings - 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC 2004, Sep 1-5 2004*, volume 26 I of *Annual International Conference of the IEEE Engineering in Medicine and Biology - Proceedings*, pages 4618–4621, San Francisco, CA, United States, 2004. Institute of Electrical and Electronics Engineers Inc., Piscataway, NJ 08855-1331, United States.
- J. Liang and G. Westheimer. Method for measuring visual resolution at the retinal level. *J Opt Soc Am A Opt Image Sci Vis*, 10(8):1691–6, 1993.
- V. Listrat, X. Ayral, F. Patarnello, J. P. Bonvarlet, J. Simonnet, B. Amor, and M. Dougados. Arthroscopic evaluation of potential structure modifying activity of hyaluronan (hyalgan(r)) in osteoarthritis of the knee. *Osteoarthritis and Cartilage*, 5(3):153–160, 1997.
- L. Lobas and L. Lobas. Bifurcations, stability, and catastrophes of equilibrium states of a double pendulum subjected to an asymmetric follower force. *Mechanics of Solids*, 39(4):110–18, 2004.
- L. G. Lobas, L. D. Patricio, and I. G. Boruk. Equilibrium of an inverted mathematical double-link pendulum with a follower force. *International Applied Mechanics*, 38(3):372–376, 2002.
- G. E. Loeb, W. S. Levine, and J. He. Understanding sensorimotor feedback through optimal control. *Cold Spring Harb Symp Quant Biol*, 55:791–803, 1990.
- G. E. Loeb, I. E. Brown, and E. J. Cheng. A hierarchical foundation for models of sensorimotor control. *Exp Brain Res*, 126(1):1–18, 1999.
- A. Longtin, J. G. Milton, J. E. Bos, and M. C. Mackey. Noise and critical-behavior of the pupil light reflex at oscillation onset. *Physical Review A*, 41(12):6992–7005, 1990.
- F. D. Lublin and S. C. Reingold. Defining the clinical course of multiple sclerosis: Results of an international survey. *Neurology*, 46(4):907–911, 1996.
- V. G. Macefield and R. S. Johansson. Control of grip force during restraint of an object held between finger and thumb: Responses of muscle and joint afferents from the digits. *Experimental Brain Research*, 108(1):172–184, 1996.
- M. C. Mackey and L. Glass. Oscillation and chaos in physiological control-systems. *Science*, 197(4300):287–288, 1977.

- M. C. Mackey and M. Santillan. Mathematics, biology, and physics: Interactions and interdependence, 2004.
- V. Mathiowetz, N. Kashman, G. Volland, K. Weber, M. Dowe, and S. Rogers. Grip and pinch strength: normative data for adults. *Arch Phys Med Rehabil*, 66(2): 69–74, 1985.
- D. I. McCloskey. Kinesthetic sensibility. *Physiological Reviews*, 58(4):763–820, 1978.
- A. D. McDonald and L. K. Sandal. Estimating the parameters of stochastic differential equations using a criterion function based on the kolmogorov-smirnov statistic. *Journal of Statistical Computation and Simulation*, 64(3):235–250, 1999.
- W. I. McDonald. Introduction to multiple sclerosis: progress and problems. a theme issue published by the royal society. *Phil. Trans. R. Soc. Lond. B*, 354(1390):1613–1614, 1999.
- W. I. McDonald and M. A. Ron. Multiple sclerosis: the disease and its manifestations. *Phil. Trans. R. Soc. Lond. B*, 354(1390):1615–1622, 1999.
- F. Mechsner, D. Kerzel, G. Knoblich, and W. Prinz. Perceptual basis of bimanual coordination. *Nature*, 414(6859):69–73, 2001.
- K. Meyer. The mucopolysaccharides of the interfibrillar substance of the mesenchyme. *Ann N Y Acad Sci*, 52(7):961–3, 1950a.
- K. Meyer. The action of hyaluronidases on hyaluronic acid. *Ann N Y Acad Sci*, 52(7):1021–3, 1950b.
- K. H. Meyer and J. Fellig. [constitution of hyaluronic acid.]. *Experientia*, 6(5):186, 1950.
- J. G. Milton, J. Foss, J. D. Hunter, and J. L. Cabrera. Controlling neurological disease at the edge of instability. In *Quantitative neuroscience: models, algorithms, diagnostics, and therapeutic applications*, pages 117–143. Kluwer Academic Publishers, 2004.
- J. Monzee, Y. Lamarre, and A. M. Smith. The effects of digital anesthesia on force control using a precision grip. *J Neurophysiol*, 89(2):672–83, 2003.
- L. Moreau, E. Sontag, and M. Arcaç. Feedback tuning of bifurcations. *Systems & control letters*, 50(3):229 (12 pages), 2003.
- F. Moss and K. Wiesenfeld. The benefits of background-noise. *Scientific American*, 273(2):66–69, 1995.

- R. M. Murray, Z. Li, and S. Sastry. *A mathematical introduction to robotic manipulation*. CRC Press, Boca Raton, 1994.
- N. S. Namerow. Somatosensory evoked responses in multiple sclerosis patients with varying sensory loss. *Neurology*, 18(12):1197, 1968.
- A. Navsarikar, D. D. Gladman, J. A. Husted, and R. J. Cook. Validity assessment of the disabilities of arm, shoulder, and hand questionnaire (dash) for patients with psoriatic arthritis. *J Rheumatol*, 26(10):2191–4, 1999.
- W. L. Nelson. Physical principles for economies of skilled movements. *Biol Cybern*, 46(2):135–47, 1983.
- J. N. Nielsen, H. Madsen, and P. C. Young. Parameter estimation in stochastic differential equations: An overview. In H. F. Chen, D. Z. Cheng, and J. F. P. O. Zhang, editors, *International Federation of Automatic Control Proceedings of the 14th World congress of IFAC*, pages 289–300, Beijing, 1999. Oxford; Published for the International Federation of Automatic Control by Pergamon; 1999.
- K. Ogata. *Modern control engineering*. Prentice Hall, Upper Saddle River, NJ, 4th edition, 2002.
- E. E. Ohnhaus and R. Adler. Methodological problems in measurement of pain - comparison between verbal rating scale and visual analog scale. *Pain*, 1(4):379–384, 1975.
- A. M. Okamura, N. Smaby, and M. R. Cutkosky. An overview of dexterous manipulation. In *Robotics and automation*, pages 255–262, San Francisco, CA, 2000. IEEE; 2000.
- D. J. Ostry and A. G. Feldman. A critical evaluation of the force control hypothesis in motor control. *Exp Brain Res*, 153(3):275–88, 2003.
- J. Paillard. Fast and slow feedback loops for the visual correction of spatial errors in a pointing task: A reappraisal. *Canadian Journal of Physiology and Pharmacology*, 74(4):401–417, 1996.
- M. G. Pandy, F. E. Zajac, E. Sim, and W. S. Levine. An optimal control model for maximum-height human jumping. *J Biomech*, 23(12):1185–98, 1990.
- J. G. Peyron. Intraarticular hyaluronan injections in the treatment of osteoarthritis - state-of-the-art review. *Journal of Rheumatology*, 20:10–15, 1993.
- C. S. Poon and C. K. Merrill. Decrease of cardiac chaos in congestive heart failure. *Nature*, 389(6650):492–495, 1997.
- C. Prablanc and O. Martin. Automatic-control during hand reaching at undetected 2-dimensional target displacements. *Journal of Neurophysiology*, 67(2):455–469, 1992.

- D. D. Price, P. A. McGrath, A. Rafii, and B. Buckingham. The validation of visual analog scales as ratio scale measures for chronic and experimental pain. *Pain*, 17(1):45–56, 1983.
- A. Prochazka, F. Clarac, G. E. Loeb, J. C. Rothwell, and J. R. Wolpaw. What do reflex and voluntary mean? modern views on an ancient debate. *Exp Brain Res*, 130(4):417–32, 2000.
- D. Rempel, B. Evanoff, P. C. Amadio, M. de Krom, G. Franklin, A. Franzblau, R. Gray, F. Gerr, M. Hagberg, T. Hales, J. N. Katz, and G. Pransky. Consensus criteria for the classification of carpal tunnel syndrome in epidemiologic studies. *American Journal of Public Health*, 88(10):1447–1451, 1998.
- R. Rico-Martinez, K. Krischer, G. Flatgen, J. S. Anderson, and I. G. Kevrekidis. Adaptive detection of instabilities: an experimental feasibility study. *Physica D-Nonlinear Phenomena*, 176(1-2):1–18, 2003.
- S. S. Robertson, J. Guckenheimer, A. M. Masnick, and L. F. Bacher. The dynamics of infant visual foraging. *Developmental Science*, 7(2):194–200, 2004.
- D. Säfström and B. B. Edin. Task requirements influence sensory integration during grasping in humans. *Learn Mem*, 11(3):356–63, 2004.
- V. J. Santos and F. J. Valero-Cuevas. Reported anatomical variability naturally leads to multimodal distributions of denavit-hartenberg parameters for the human thumb. *IEEE Trans Biomed Eng*, 53(2):155–63, 2006.
- J. A. Saunders and D. C. Knill. Humans use continuous visual feedback from the hand to control fast reaching movements. *Experimental Brain Research*, 152(3):341–352, 2003.
- J. A. Saunders and D. C. Knill. Visual feedback control of hand movements. *J. Neurosci.*, 24(13):3223–34, 2004.
- J. P. Scholz and G. Schöner. The uncontrolled manifold concept: identifying control variables for a functional task. *Exp Brain Res*, 126(3):289–306, 1999.
- J. P. Scholz, G. Schöner, and M. L. Latash. Identifying the control structure of multijoint coordination during pistol shooting. *Exp Brain Res*, 135(3):382–404, 2000.
- G. Schöner and J. A. Kelso. Dynamic pattern generation in behavioral and neural systems. *Science*, 239(4847):1513–20, 1988.
- S. H. Scott. Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*, 5(7):534–546, 2004.

- R. Shadmehr and F. A. Mussa-Ivaldi. Adaptive representation of dynamics during learning of a motor task. *J Neurosci*, 14(5 Pt 2):3208–24, 1994.
- J. E. Shigley and C. R. Mischke. *Mechanical engineering design*. McGraw-Hill series in mechanical engineering;. McGraw-Hill, New York, 5th edition, 1989.
- H. Singer. Parameter estimation of nonlinear stochastic differential equations: Simulated maximum likelihood versus extended kalman filter and ito-taylor expansion. *Journal of Computational and Graphical Statistics*, 11(4):972–995, 2002.
- J. B. Smeets and E. Brenner. A new view on grasping. *Motor Control*, 3(3):237–71, 1999.
- J. M. Smith. Optimization theory in evolution. *Annual Review of Ecology and Systematics*, 9:31–56, 1978.
- K. J. Smith and W. I. McDonald. The pathophysiology of multiple sclerosis: the mechanisms underlying the production of symptoms and the natural history of the disease. *Phil. Trans. R. Soc. Lond. B*, 354(1390):1649–1673, 1999.
- S. J. Sober and P. N. Sabes. Multisensory integration during motor planning. *Journal of Neuroscience*, 23(18):6982–6992, 2003.
- S. J. Sober and P. N. Sabes. Flexible strategies for sensory integration during motor planning. *Nat Neurosci*, 8(4):490–7, 2005.
- C. Soize. *The Fokker-Planck equation for stochastic dynamical systems and its explicit steady state solutions*. Series on advances in mathematics for applied sciences ;; v. 17;. Teaneck, N.J., Singapore, 1994.
- R. V. Southwell. On the analysis of experimental observations in problems of elastic stability. *Proc. R. Society*, 135(A):601, 1932.
- R. W. Sperry. Neural basis of the spontaneous optokinetic response produced by visual inversion. *J Comp Physiol Psychol*, 43(6):482–9, 1950.
- M. Srinivasan. *Why walk and run: Energetic costs and energetic optimality in simple mechanics-based models of a bipedal animal*. Ph.d., Cornell University, 2006.
- M. Srinivasan and A. Ruina. Computer optimization of a minimal biped model discovers walking and running. *Nature*, 439(7072):72–75, 2006.
- S. Stahl, I. Karsh-Zafir, N. Ratzon, and N. Rosenberg. Comparison of intraarticular injection of depot corticosteroid and hyaluronic acid for treatment of degenerative trapeziometacarpal joints. *Jcr-Journal of Clinical Rheumatology*, 11(6):299–302, 2005.

- G. Stein. Respect the unstable. *IEEE Control Systems Magazine*, 23(4):12–25, 2003.
- R. B. Stein, E. R. Gossen, and K. E. Jones. Neuronal variability: noise or part of the signal? *Nat Rev Neurosci*, 6(5):389–97, 2005.
- D. Sternad. Debates in dynamics: A dynamical systems perspective on action and perception. *Human Movement Science*, 19(4):407–423, 2000.
- J. Timmer. Parameter estimation in nonlinear stochastic differential equations. *Chaos Solitons & Fractals*, 11(15):2571–2578, 2000.
- S. Timoshenko. *Theory of elastic stability*. Engineering societies monographs. McGraw-Hill, New York,, 2d edition, 1961.
- E. Todorov. Optimality principles in sensorimotor control. *Nature Neuroscience*, 7(9):907–915, 2004.
- E. Todorov and M. I. Jordan. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11):1226–1235, 2002.
- E. Todorov and M. I. Jordan. A minimal intervention principle for coordinated movement. *Advances in neural information processing systems*, 15:27–34, 2003.
- E. Todorov, W. Li, and X. Pan. From task parameters to motor synergies: A hierarchical framework for approximately optimal control of redundant manipulators. *Journal of Robotic Systems*, 22(11):691–710, 2005.
- P. Totten and S. Flinn-Wagner. Functional evaluation of the hand. In B. G. Stanley and S. M. Tribuzi, editors, *Concepts in hand rehabilitation*, volume 9, pages xxiv, 582 p. F.A. Davis Company, Philadelphia, 1992.
- F. J. Valero-Cuevas. Applying principles of robotics to understand the biomechanics, neuromuscular control and clinical rehabilitation of human digits. In *ICRA 2000: IEEE International Conference on Robotics and Automation, Apr 24-Apr 28 2000*, volume 1 of *Proceedings - IEEE International Conference on Robotics and Automation*, pages 270–275, San Francisco, CA, USA, 2000. Institute of Electrical and Electronics Engineers Inc., Piscataway, NJ, USA.
- F. J. Valero-Cuevas. Device for developing and measuring grasping force and grasping dexterity, March 25, 2003 2003.
- F. J. Valero-Cuevas. An integrative approach to the biomechanical function and neuromuscular control of the fingers. *Journal of Biomechanics*, 38(4):673–684, 2005.

- F. J. Valero-Cuevas and V. R. Hentz. Releasing the a3 pulley and leaving flexor superficialis intact increases pinch force following the zancolli lasso procedures to prevent claw deformity in the intrinsic palsied finger. *J Orthop Res*, 20(5): 902–9, 2002.
- F. J. Valero-Cuevas, F. E. Zajac, and C. G. Burgar. Large index-fingertip forces are produced by subject-independent patterns of muscle excitation. *J Biomech*, 31(8):693–703., 1998.
- F. J. Valero-Cuevas, M. E. Johanson, and J. D. Towles. Towards a realistic biomechanical model of the thumb: the choice of kinematic description may be more critical than the solution method or the variability/uncertainty of musculoskeletal parameters. *J Biomech*, 36(7):1019–30, 2003a.
- F. J. Valero-Cuevas, N. Smaby, M. Venkadesan, M. Peterson, and T. Wright. The strength-dexterity test as a measure of dynamic pinch performance. *Journal of Biomechanics*, 36(2):265–270, 2003b.
- R. J. van Beers, A. C. Sittig, and J. J. Gon. Integration of proprioceptive and visual position-information: An experimentally supported model. *J Neurophysiol*, 81(3):1355–64, 1999.
- R. J. van Beers, P. Baraduc, and D. M. Wolpert. Role of uncertainty in sensorimotor control. *Philos Trans R Soc Lond B Biol Sci*, 357(1424):1137–45, 2002.
- R. J. van Beers, P. Haggard, and D. M. Wolpert. The role of execution noise in movement variability. *J Neurophysiol*, 91(2):1050–63, 2004.
- H. van der Kooij, R. Jacobs, B. Koopman, and H. Grootenboer. A multisensory integration model of human stance control. *Biol. Cybern.*, 80(5):299–308, 1999.
- M. Venkadesan, S. Backus, L. A. Mandl, A. Swigart, M. Peterson, S. Lyman, L. Ariola, R. N. Hotchkiss, and F. J. Valero-Cuevas. The strength-dexterity test is a novel and clinically informative measure of treatment outcome in thumb osteoarthritis. *Arthritis and Rheumatism*, 52(9):S516–S516, 2005.
- H. von Helmholtz. *Handbuch der Physiologischen Optik*. Leipzig : Leopold Voss, 1867; 1 v. ; in 8.; DCCC.4.353, 1867.
- H. von Helmholtz, R. Kahl, and ed. *Selected writings of Hermann von Helmholtz*. Conn., Wesleyan University Press, Middletown, [1st] edition, 1971.
- A. M. Wahl. *Mechanical springs*. McGraw-Hill, New York, 2d edition, 1963.
- L. Wang, E. B. Larson, J. D. Bowen, and G. van Belle. Performance-based physical function and future dementia in older people. *Arch Intern Med*, 166(10):1115–1120, 2006.

- N. Wei, S. K. Delauter, and S. J. Beard. Arthroscopic debridement and viscosupplementation - a minimally invasive treatment for symptomatic osteoarthritis involving the base of the thumb. *Jcr-Journal of Clinical Rheumatology*, 8(3):125–129, 2002.
- G. Westling and R. S. Johansson. Factors influencing the force control during precision grip. *Exp Brain Res*, 53(2):277–84, 1984.
- H. E. Wheat, A. W. Goodwin, and A. S. Browning. Tactile resolution: peripheral neural mechanisms underlying the human capacity to determine positions of objects contacting the fingerpad. *J Neurosci*, 15(8):5582–95, 1995.
- H. E. Wheat, L. M. Salo, and A. W. Goodwin. Human ability to scale and discriminate forces typical of those occurring during grasp and manipulation. *J Neurosci*, 24(13):3394–401, 2004.
- K. Wiesenfeld. Noisy precursors of nonlinear instabilities. *Journal of Statistical Physics*, 38(5-6):1071–1097, 1985.
- K. Wiesenfeld and J. S. McCarley. Suppressed fluctuations and incipient instabilities. *Physical Review A*, 42(2):755–760, 1990.
- K. Wiesenfeld and B. McNamara. Small-signal amplification in bifurcating dynamic-systems. *Physical Review A*, 33(1):629–642, 1986.
- K. Wiesenfeld and F. Moss. Stochastic resonance and the benefits of noise - from ice ages to crayfish and squids. *Nature*, 373(6509):33–36, 1995.
- M. Wobig, A. Dickhut, R. Maier, and G. Vetter. Viscosupplementation with hylan g-f 20: A 26-week controlled trial of efficacy and safety in the osteoarthritic knee. *Clinical Therapeutics*, 20(3):410–423, 1998.
- D. M. Wolpert and J. R. Flanagan. Motor prediction. *Curr Biol*, 11(18):R729–32, 2001.
- D. M. Wolpert and Z. Ghahramani. Computational principles of movement neuroscience. *Nat Neurosci*, 3 Suppl:1212–7, 2000.
- D. M. Wolpert, Z. Ghahramani, and M. I. Jordan. An internal model for sensorimotor integration. *Science*, 269(5232):1880–2, 1995.
- D. M. Wolpert, R. C. Miall, and M. Kawato. Internal models in the cerebellum. *Trends in Cognitive Sciences*, 2(9):338–347, 1998.
- D. M. Wolpert, Z. Ghahramani, and J. R. Flanagan. Perspectives and problems in motor learning. *Trends Cogn Sci*, 5(11):487–494, 2001.
- T. Yoshikawa. *Foundations of robotics : analysis and control*. MIT Press, Cambridge, Mass., 1990.

- F. E. Zajac, M. R. Zomlefer, and W. S. Levine. Hindlimb muscular activity, kinetics and kinematics of cats jumping to their maximum achievable heights. *J Exp Biol*, 91:73–86, 1981.
- F. E. Zajac, R. W. Wicke, and W. S. Levine. Dependence of jumping performance on muscle properties when humans use only calf muscles for propulsion. *J Biomech*, 17(7):513–23, 1984.
- V. M. Zatsiorsky, R. W. Gregory, and M. L. Latash. Force and torque production in static multifinger prehension: biomechanics and control. i. biomechanics. *Biol Cybern*, 87(1):50–7, 2002.