

# A Dynamic Primitives Hypothesis: A Descriptive Model of Human Physical Interaction

by

James Hermus

B.S., Biomed. Eng., University of Wisconsin - Madison (2016)  
S.M., Massachusetts Institute of Technology (2018)

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

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2023

© Massachusetts Institute of Technology 2023. All rights reserved.

Author .....  
Department of Mechanical Engineering  
September 11, 2022

Certified by .....  
Neville Hogan  
Sun Jae Professor of Mechanical Engineering  
and Professor of Brain and Cognitive Sciences  
Thesis Supervisor

Accepted by .....  
Nicolas G. Hadjiconstantinou  
Chairman, Department Committee on Graduate Theses



# A Dynamic Primitives Hypothesis: A Descriptive Model of Human Physical Interaction

by

James Hermus

Submitted to the Department of Mechanical Engineering  
on September 11, 2022, in partial fulfillment of the  
requirements for the degree of  
Doctor of Philosophy

## Abstract

Physical interaction is a key aspect of activities of daily living. These tasks require simultaneous regulation of both force and motion. For example, even a task as simple as opening a door presents a challenge to the development of prosthetics, exoskeletons, and human-robot interaction. Motor neuroscience has reported systematic patterns of motion during free reaching and force during static posture. However, similar results do not extend to physical interaction. A descriptive model is required.

The paradox of human performance: Despite large feedback delays, and many degrees of freedom, humans are incredibly dexterous and excel at physical interaction with complex objects. To accomplish such performance, we hypothesize that motor behavior, with and without physical interaction, is constructed using a limited set of primitive dynamic behaviors, including oscillations, submovements, and mechanical impedance. We not only propose that these ‘building blocks’ exist but that their connectivity is important – a Norton equivalent network.

This thesis is composed of four components that systematically investigated this hypothesis. (1) Through the study of crank turning we presented evidence for dynamic primitives. (2) To test the hypothesis, a method to estimate impedance during crank turning was developed. (3) When kinematic redundancy was substantial, a dynamic primitive-based control resolved redundancy without compromising performance. (4) This hypothesis led to the development of an experiment which falsified a common assumption, that humans can directly regulate force during motion.

While it is fundamentally hard to prove hypotheses in human motor control, the hypothesis of dynamic primitives can descriptively account for systematic patterns in constrained motion. Furthermore, the value of this hypothesis was demonstrated in robotics by simplifying the management of kinematic redundancy and force regulation.

Thesis Supervisor: Neville Hogan  
Title: Sun Jae Professor of Mechanical Engineering  
and Professor of Brain and Cognitive Sciences





## Acknowledgments

First, I would like to thank my advisor, Professor Neville Hogan. His knowledge of controls and dynamics is unparalleled. It is an honor to work with someone with such vast knowledge of human motor control research. Our meetings are never rushed, he takes time to discuss the fundamental ideas underlying our work, even if it means being late for a different appointment. After meeting with Neville, I always have a clear understanding of how to move forward with my work. I have benefited immensely from his constructive feedback of my writing. His extensive comments, undoubtedly, show that he has thoroughly read every word of my draft. It is a true privilege to discuss research and ideas with Neville; his enthusiasm and passion for the field is contagious. His guidance continues to help me develop as a critical thinker, writer, and scientific storyteller.

Throughout graduate school, I have slowly become aware of how academia systematically places graduate students in vulnerable positions – where things can, and do, go wrong. Some of my friends have experienced these difficult situations, resulting from advisors who prioritized titles, publications, flashy results, or personal recognition. I cannot express how fortunate I have been to know Neville truly values both me as a person and the scientific quality of our work.

I sincerely thank my committee members, Professor Harry Asada and Professor Alberto Rodriguez, who provided their perspectives and expertise.

I would like to thank Dr. Meghan Huber, who was a post-doctoral researcher in our group and is now a Professor at UMass Amherst. Meghan has spent countless hours discussing statistics, recommending research articles, deliberating results, and theoretically considering experimental designs. Meghan has become a close friend; she has listened to my confusion and frustration along the way, provided a positive outlook, and offered advice whenever it was needed.

I have had the great fortune to work and collaborate with a tremendous group of lab mates over the last 6 years. I would like to thank David Verdi, Jongwoo Lee, Johannes Lachner, and David Mercado for their support; they made hanging out

in lab, well after midnight, a true joy! I would also like to thank Michael West, Rika Sugimoto Dimitrova, Federico Tessari, Moses Nah, Stephan Stansfield, Kaymie Shiozawa, Kuangen Zhang, Nic Arons, Jessie Thorup, Julie Ochoa, and many UROPs. Be it white boarding ideas or attending conferences, I truly enjoyed the times we shared.

In reflecting on this accomplishment, I must recognize three people at the University of Wisconsin-Madison who gave me the opportunity to begin this journey: Professor Charles Mistretta, Professor Tim Szczykutowicz, and Professor Darryl Thelen. They sparked my interest in graduate study and provided tremendous undergraduate research opportunities.

Additionally, I want to thank the fantastic MechE, MakerWorks, and MITOC (MIT Outing Club) communities; their friendship and support, outside the lab, has kept me going!

Finally, I want to thank my parents, Amy and Russ Hermus, for their endless love and encouragement. Without the support provided by my mother, Amy, throughout my many challenges in the educational system, as a student with dyslexia, this incredible opportunity would not have been possible.

This research was performed in the Eric P. and Evelyn E. Newman Laboratory for Biomechanics and Human Rehabilitation at the Massachusetts Institute of Technology. It was supported by the National Science Foundation National Robotics Initiative Grant No. 1637824, and National Institute of Health Grant No. R01-HD087089, and a Harrington Fellowship.

## Members of the Committee

Professor Neville Hogan – Thesis Supervisor  
Sun Jae Professor of Mechanical Engineering and  
Professor of Brain and Cognitive Sciences  
Massachusetts Institute of Technology

Professor H. Harry Asada  
Ford Professor of Engineering  
Massachusetts Institute of Technology

Professor Alberto Rodriguez  
Associate Professor of Mechanical Engineering  
Massachusetts Institute of Technology



# Contents

<b>1</b>	<b>Introduction</b>	<b>21</b>
1.0.1	Personal Aside . . . . .	21
1.0.2	Motivation . . . . .	22
1.1	The Paradox of Human Performance . . . . .	24
1.2	Literature Review . . . . .	25
1.2.1	Coordinate System for Movement Planning . . . . .	26
1.2.2	Smoothness . . . . .	28
1.2.3	Movement Intermittency . . . . .	32
1.2.4	Oscillations . . . . .	34
1.2.5	Impedance . . . . .	35
1.2.6	Constrained Motion Studies . . . . .	37
1.3	Dynamic Primitives . . . . .	39
1.4	Summary of Thesis . . . . .	40
<b>2</b>	<b>Separating Neural Influences from Peripheral Mechanics: The Speed-Curvature Relation in Mechanically-Constrained Actions</b>	<b>43</b>
2.1	Introduction . . . . .	43
2.2	Methods . . . . .	46
2.2.1	Participants . . . . .	46
2.2.2	Experimental Apparatus and Procedure . . . . .	47
2.2.3	Data Processing and Analysis . . . . .	49
2.2.4	Trials at Preferred Speed . . . . .	50
2.2.5	Trials at Instructed Speeds . . . . .	50

2.2.6	Simulation of Passive Inertial Mechanics . . . . .	51
2.2.7	Zero-Force Trajectories . . . . .	52
2.2.8	Speed - Curvature Relation . . . . .	54
2.2.9	Sensitivity to Impedance Assumptions . . . . .	55
2.3	Results . . . . .	55
2.3.1	Speed and Force at Preferred Speed . . . . .	55
2.3.2	Speed and Force at Instructed Speeds . . . . .	58
2.3.3	Zero-Force Trajectory . . . . .	60
2.3.4	Sensitivity Analyses of the Zero-force Trajectory and its Effect on the Speed-Curvature Relation . . . . .	64
2.4	Discussion . . . . .	65
2.4.1	Patterns of Speed and Force . . . . .	66
2.4.2	The Zero-Force Trajectory . . . . .	67
2.4.3	Speed Curvature Relation in the Zero-Force Trajectory . . . . .	69
2.5	Conclusion . . . . .	71
<b>3</b>	<b>Dynamic Primitives in Constrained Action: Systematic Changes in the Zero-Force Trajectory</b>	<b>73</b>
3.1	Introduction . . . . .	73
3.1.1	Motivation for the Experiments . . . . .	74
3.2	Methods . . . . .	75
3.2.1	Dependent Measures . . . . .	75
3.2.2	Statistical Analysis . . . . .	77
3.2.3	Sensitivity to Impedance Assumptions . . . . .	78
3.3	Results . . . . .	78
3.3.1	Variation of Force and Motion . . . . .	78
3.3.2	Zero-Force Trajectory Orientation . . . . .	79
3.3.3	Orientation – Sensitivity Analysis . . . . .	83
3.3.4	Hand Speed Coefficient of Variation . . . . .	84
3.4	Discussion . . . . .	85

3.4.1	Zero-force Trajectory Orientation . . . . .	87
3.4.2	Variability of Hand Speed . . . . .	88
3.4.3	Limitations of this Study . . . . .	89
3.4.4	Implications . . . . .	90
3.5	Conclusions . . . . .	91
<b>4</b>	<b>Estimating Impedance during Constrained Motion – The Configuration-Dependent Ensemble</b>	<b>93</b>
4.1	Introduction . . . . .	93
4.1.1	Brownian Noise . . . . .	96
4.1.2	Evidence of Brownian Noise during Crank-Turning . . . . .	98
4.1.3	System Identification in the Presence of Brownian Noise . . . . .	100
4.1.4	Configuration-Dependent Ensemble Method . . . . .	102
4.2	Methods . . . . .	106
4.2.1	Participants . . . . .	106
4.2.2	Experimental Apparatus and Procedure . . . . .	106
4.2.3	Robot Control . . . . .	108
4.2.4	Robot-only . . . . .	108
4.2.5	Simulations . . . . .	109
4.2.6	Data Processing . . . . .	112
4.3	Results . . . . .	113
4.3.1	Robot . . . . .	114
4.3.2	Robot and Human . . . . .	116
4.3.3	Simulations . . . . .	119
4.4	Discussion . . . . .	122
4.4.1	Configuration-Dependent Ensemble Method . . . . .	123
4.4.2	InMotion Hysteresis . . . . .	124
4.4.3	Subtracting Robot Dynamics . . . . .	125
4.4.4	Frequency Separation . . . . .	125
4.5	Conclusions . . . . .	127

<b>5</b>	<b>Exploiting Redundancy to Facilitate Physical Interaction</b>	<b>129</b>
5.1	Introduction . . . . .	129
5.1.1	Managing Redundancy . . . . .	130
5.1.2	Factors that Influence Controller Design . . . . .	133
5.1.3	Summary . . . . .	136
5.2	Methods . . . . .	136
5.2.1	Experimental Setup . . . . .	137
5.2.2	Impedance Controller . . . . .	138
5.2.3	Test Conditions . . . . .	142
5.2.4	Data Analysis . . . . .	145
5.2.5	Statistical Analysis . . . . .	147
5.3	Results . . . . .	148
5.3.1	Standard Condition, Unconstrained . . . . .	148
5.3.2	Standard Condition, Constrained . . . . .	150
5.3.3	Fast Motion . . . . .	151
5.3.4	Low Joint-Space Stiffness . . . . .	152
5.3.5	Statistical Comparisons . . . . .	153
5.4	Discussion . . . . .	156
5.4.1	Limitations . . . . .	158
5.4.2	Analogy to Polynomial Kernel Methods? . . . . .	160
5.4.3	Understanding Human Motor Control . . . . .	161
5.4.4	Applications to Robotics . . . . .	161
5.4.5	Future Work . . . . .	162
5.5	Conclusions . . . . .	162
<b>6</b>	<b>Dynamic Primitives Limit Human Force Regulation during Motion</b>	<b>165</b>
6.1	Introduction . . . . .	165
6.2	Methods . . . . .	169
6.2.1	Participants . . . . .	169
6.2.2	Experimental Procedures . . . . .	169



6.2.3	Experimental Design . . . . .	171
6.2.4	Dependent Measures and Data Processing . . . . .	171
6.2.5	Data Analysis and Statistics . . . . .	172
6.3	Results . . . . .	173
6.3.1	Performance Improvements . . . . .	173
6.3.2	Existence of Motion-Dependent Force Errors . . . . .	175
6.4	Discussion . . . . .	179
6.4.1	Force Error . . . . .	179
6.4.2	Dynamic Primitives . . . . .	180
6.4.3	Limitations . . . . .	181
6.4.4	Implications . . . . .	182
6.5	Conclusions . . . . .	183
<b>7</b>	<b>Conclusions</b>	<b>185</b>
7.1	Areas of Potential Future Work . . . . .	187
7.2	Speculation . . . . .	189
7.3	Closing Remarks . . . . .	191
<b>A</b>	<b>Modeling a Two-Link Manipulator Coupled to a Crank</b>	<b>193</b>
<b>B</b>	<b>Ensemble Statistics – Assumptions and 1D Simulations</b>	<b>199</b>
B.1	Review of Assumptions . . . . .	199
B.1.1	Probability Density Functions . . . . .	199
B.1.2	Ensemble Statistics . . . . .	201
B.1.3	Stationary Statistics . . . . .	202
B.2	Input-Output Relations . . . . .	203
B.2.1	Stationary Method for Estimating IRF . . . . .	204
B.2.2	Ensemble Method for Estimating IRF . . . . .	205
B.2.3	Example – Time-Varying $x_0(t)$ and Constant $Z$ . . . . .	207
B.2.4	Example – Time-Varying $x_0(t)$ and Time-Varying $Z(t)$ . . . . .	210

<b>C Evidence of Brownian Noise – Static Posture</b>	<b>213</b>
C.1 Introduction . . . . .	213
C.2 Methods . . . . .	213
C.3 Results . . . . .	214
C.4 Discussion and Conclusions . . . . .	215
<b>D InMotion Robot–Inertial Model</b>	<b>217</b>
<b>E InMotion Robot– Static Feed-forward Force Production</b>	<b>219</b>
E.1 Introduction . . . . .	219
E.2 Methods . . . . .	219
E.3 Results . . . . .	220
E.4 Discussion and Conclusions . . . . .	222
<b>F InMotion Robot–Hysteresis</b>	<b>223</b>
F.1 Introduction . . . . .	223
F.2 Methods . . . . .	223
F.3 Results . . . . .	224
F.4 Discussion and Conclusions . . . . .	226
<b>G Quantitative Analysis of the Quasi-Static and Dynamic Speeds</b>	<b>227</b>
<b>H Simulations of Redundancy</b>	<b>229</b>
<b>I Computation of Rotational Displacements</b>	<b>233</b>

# List of Figures

1-1	Zero-force trajectory block diagram . . . . .	24
1-2	Minimum jerk trajectory . . . . .	29
1-3	Speed curvature relation . . . . .	30
1-4	Power law depends on hand path . . . . .	32
1-5	Location of the intermittency source . . . . .	34
1-6	Joint reversals . . . . .	39
2-1	Human crank turning – Experimental setup . . . . .	47
2-2	Human crank turning – Experimental design . . . . .	48
2-3	Comfortable speed – Representative subject speed and force . . . . .	56
2-4	Comfortable speed – Dependent measures . . . . .	57
2-5	Instructed speed – Subject average speed and force . . . . .	58
2-6	Instructed speed – Dependent measures . . . . .	59
2-7	Zero-force trajectory estimates . . . . .	61
2-8	Speed curvature representative subject binned data and $\Delta$ histograms	63
2-9	95% CI for the $\Delta$ parameter across subjects and conditions . . . . .	64
3-1	Ellipse eccentricity and major axis angle vs. log-ratio coordinates . . . . .	76
3-2	Intuition for log-ratio coordinates . . . . .	76
3-3	Systematic changes in velocity and force with respect to crank position	79
3-4	Zero-force trajectory estimates – representative subject . . . . .	80
3-5	Zero-force trajectory estimates – average over all subjects . . . . .	81
3-6	Distribution of ellipse parameters in log-ratio coordinates . . . . .	82
3-7	Log-ratio estimates when impedance parameters varied . . . . .	84

3-8	Coefficient of variation of hand speed . . . . .	85
3-9	Out of phase sinusoid . . . . .	88
4-1	Crank angle . . . . .	96
4-2	Brownian noise . . . . .	98
4-3	Crank angle error . . . . .	99
4-4	Crank – Spectra . . . . .	100
4-5	Structure in position not time . . . . .	101
4-6	Computing gamma . . . . .	103
4-7	Ensemble – Time vs. Configuration . . . . .	105
4-8	Human system ID – Experimental setup . . . . .	107
4-9	Simulation diagram . . . . .	111
4-10	Raw measures – Robot . . . . .	115
4-11	Estimates - Robot . . . . .	116
4-12	Raw measures – Robot and human . . . . .	118
4-13	Estimates – Robot and human . . . . .	119
4-14	Raw measures – Simulations . . . . .	120
4-15	Estimates – Simulations . . . . .	121
4-16	Estimates – Simulations vs. Human . . . . .	122
5-1	Robot crank turning – Experimental setup . . . . .	138
5-2	Robot nominal Posture . . . . .	141
5-3	Standard condition, unconstrained . . . . .	149
5-4	Standard condition, constrained . . . . .	151
5-5	Fast condition . . . . .	152
5-6	Low stiffness condition . . . . .	153
5-7	Difference in RMSE between superposition and nullspace projection .	154
5-8	Influence of weighting matrix . . . . .	155
6-1	Direct vs. indirect force control . . . . .	167
6-2	Human force control – Experimental setup . . . . .	168

6-3	Mean RMS of force across participants and conditions . . . . .	174
6-4	Mean RMS of force error from the last 9 trials . . . . .	175
6-5	Raw force data in each block for a representative subject . . . . .	177
6-6	Lags of maximum auto-correlation . . . . .	178
A-1	Model of two-link manipulator coupled to crank . . . . .	194
B-1	Ensemble . . . . .	201
B-2	Example constant Z – Measurements . . . . .	209
B-3	Example constant Z – IRF . . . . .	210
B-4	Example constant Z – Estimates . . . . .	210
B-5	Example varying Z – Measurements . . . . .	211
B-6	Example varying Z – IRF . . . . .	212
B-7	Example varying Z – Estimates . . . . .	212
C-1	Static experiment – Position . . . . .	214
C-2	Static experiment – Power spectra . . . . .	215
D-1	InMotion inertial model . . . . .	217
E-1	InMotion clamped – Experimental setup . . . . .	220
E-2	InMotion clamped – Magnitude . . . . .	221
E-3	InMotion clamped – Coherence . . . . .	222
F-1	Experimental setup – Spring test . . . . .	224
F-2	Hysteresis . . . . .	225
H-1	Simulated planar 3 DOF manipulator . . . . .	231



# List of Tables

5.1	Optimization criteria for nullspace projectors . . . . .	136
5.2	Control parameters . . . . .	142
5.3	Test conditions and notations . . . . .	145
A.1	Human limb inertia properties . . . . .	197
D.1	InMotion Inertial Properties . . . . .	218
H.1	Robot crank turning – Controller parameters . . . . .	230





# Chapter 1

## Introduction

“The harmony of natural law reveals an intelligence of such superiority that, compared with it, all the systematic thinking and acting of human beings is an utterly insignificant reflection.”

–Albert Einstein

While human physical interaction may not fall into the category of explicitly formulated physical laws to which Einstein was referring, there is something inherently beautiful about the human ability to physically interact with the world.

### 1.0.1 Personal Aside

Unique life experiences personally motivate my research. Spending time in a special-education room while growing up, I observed different types of adaptive devices. These devices, designed by engineers and scientists my friends and I had never met, liberated us. Since early middle school, I have been working to become one of those engineers. I have strived to obtain a unique perspective in the field of dynamic systems and control. I hope that my research at MIT, this thesis, is part of the platform from which to step forward in this dream. In my next steps, I hope to continue learning but also to solidify the application of the ideas presented herein.

## 1.0.2 Motivation

An increasing amount of attention has been directed to the development, both in academia and industry, of hardware designed to interact with humans (e.g. prosthetics, exoskeletons, rehabilitation robotics). This field has the potential to change the daily lives of millions of people. Each year in the United States 185 thousand people experience amputations, resulting in an annual cost of \$8-12 billion [Ziegler-Graham et al., 2008]. Furthermore, with an aging population robotic stroke rehabilitation has great potential. Each year in the United States 795 thousand people experience a stroke, resulting in an annual cost of \$53 billion [Tsao et al., 2022]. The use of exoskeletons for rehabilitation and augmentation has yet to see widespread consumer use. However, this area is expected to grow rapidly in the coming years. Lastly, the more general field of robotics is growing rapidly. Currently, more than 400 thousand robots are used in industry, with an annual cost of \$27 billion [Grau et al., 2021].

Science fiction worlds where robots seamlessly integrate with humans have captured many people’s imagination (e.g. Marvel’s Iron Man). It may appear that the primary obstacle preventing these possibilities is the development of lighter, stronger, more back-drivable systems, and better power sources. However, even if this futuristic hardware existed, control algorithms which work with the human motor system are still in their infancy. This is in part a reflection of the limited understanding of human motor control – especially during contact. The work presented herein is a step forward in this direction.

Humans excel at physical interaction with objects, even when those objects introduce complex dynamics and kinematic constraints. Furthermore, human dexterity exceeds that of most modern robots, despite the fact that the human neuro-mechanical system is considerably slower than its robotic counterparts [Kandel et al., 2013, Loram et al., 2005, Shepard and Metzler, 1971]. In contrast, motor neuroscience has primarily focused on the examination of elementary behavior under strict experimental control. This is in part due to the approach used in many studies: search for consistent patterns observed across a variety of conditions. Some studies investigated conditions

with free reaching (substantial motion, negligible force). In free reaching, systematic patterns have been widely reported; examples include Fitts' law [Fitts, 1954] and the speed-curvature relation [Lacquaniti et al., 1993]. Conversely, other studies have investigated force exertion during static posture (substantial force, negligible motion). In this case, an analog of the speed-curvature relation has been reported [Massey et al., 1992]. However, in everyday life physical interaction frequently involves both substantial motion and substantial force. In physical interaction, when both motion and force vary, robust patterns have not been reported. This indicates that human physical interaction is not well understood.

When a behavior involves substantial contact, force and motion are no longer independent variables. Force exerted on an object depends not only on neural activity, but also on the object's motion. One way to describe the dynamics of interaction is with the mechanical impedance operator  $Z\{\cdot\}$  [Hogan, 1985b].

Mathematically, impedance is the relation between displacement and the force it evokes, a dynamic generalization of stiffness. The force time-function  $F(t)$  can be computed from the displacement time-function  $x(t)$ ,  $F(t) = Z\{\Delta x(t)\}$ . Displacement is defined as  $\Delta x(t) = x_0(t) - x(t)$ , where  $x(t)$  is position, and  $x_0(t)$  is a 'zero-force trajectory'. In principle, knowledge of mechanical impedance, in combination with simultaneous measurement of force and motion during object manipulation, would allow us to 'subtract off' or 'peel back' peripheral biomechanics to uncover a summary of one aspect of the underlying neural influences expressed in terms of motion,

$$x_0 = Z^{-1}\{f\} + x. \tag{1.1}$$

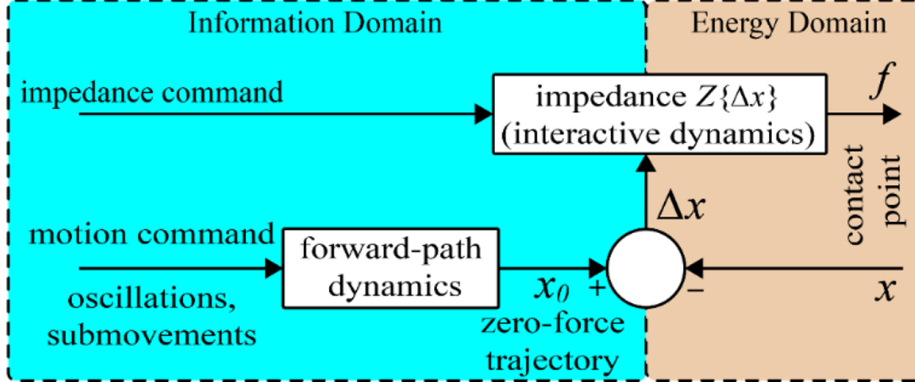


Figure 1-1: Block diagram depicts the relationship between the impedance  $Z$ , force,  $f$ , position,  $x$ , central nervous system (CNS) and the zero-force trajectory,  $x_0$ . Figure replicated from [Hogan, 2014].

In practice, mechanical impedance is nonlinear and time-varying. Measuring it during action can be achieved but is challenging [Kearney and Hunter, 1990]. In this thesis, we will explore these ideas through three approaches: (1) assuming the impedance, (2) measuring the impedance, and (3) choosing the impedance in robotic experiments.

A kinematic constraint provides an intermediate stage between unconstrained (free) motion and interaction with complex dynamics. Moreover, curved constraints are ubiquitous in activities of daily living (ADLs) (e.g. turning a steering wheel or opening a door). In fact, interacting with doors was shown to be the most common ADL [Petrich et al., 2022]. This thesis investigated humans and robots physically interacting with a circular constraint (i.e. turning a crank).

## 1.1 The Paradox of Human Performance

Human neural pathways and muscles are orders of magnitude slower than computer processors and robotic actuators, yet humans outperform robots in tasks requiring physical interaction [Hogan, 2017, Kandel et al., 2013, Churchland and Sejnowski, 2017]. Neural action potentials propagate at speeds which rarely exceed 100 m/s, orders of magnitude slower than electronic communication. Not only is the neural

transmission slow, human muscles contract at frequencies less than 10 Hz, while their robotic counterparts easily achieve movements well beyond 100 Hz.

Furthermore, the control of high-degree-of-freedom robots is frequently viewed as a challenge, especially using popular optimization-based methods, which scale poorly with system dimension. This is the infamous ‘curse of dimensionality’ [Bellman, 1966]. Meanwhile, biological systems regularly articulate appendages vastly more complex than state-of-the-art robots. Furthermore, ‘excess’ anatomical degrees of freedom are commonplace in biological systems; the human arm has between 9 and 10 degrees of freedom, the human hand has more than 20 degrees of freedom, and the entire body has more than 200 degrees of freedom.

In the face of these delays and kinematic complexity, even an activity as simple as opening a door with one hand, holding a cup of coffee in the other hand, and maintaining a stable upright posture would seem impossible. This is the paradox of human performance.

The long delays indicate that humans rely heavily on a predictive or feedforward control strategy, based on some internal representation or model of the system to be controlled. However, it seems unlikely that humans develop a detailed ‘engineering style’ dynamic model for each object they interact with. Without perfect prediction, neuromuscular delays are a fundamental constraint which may yield an evolutionary preference for stereotyped, smooth, and predictable behavior.

## 1.2 Literature Review

Investigations of human motor control have been ongoing for more than a century [Woodworth, 1899]. The following will present a review of the extensive literature which forms the core pillars that this thesis rests upon.

To learn fundamental concepts about human motor control, simple but competent models have been proposed. These models have then been iteratively refined as new observations have been made. When studying the human circuitry and controller, it is rarely possible or ethical to ‘open the box’ – enabling direct measurement of signals

from neurons<sup>1</sup>. In addition, models of human biomechanics are incomplete. Currently, it is not feasible to measure or compute exact quantities such as muscle force<sup>2</sup>, inertial properties, muscle moment arms, and neural activation signals. Thus, researchers rely on the outputs of the motor system which can be observed or estimated, in order to make inferences about the processes that generate them. However, the modeling is notoriously difficult. Human muscles are highly nonlinear [Zahalak, 1990] and the neural control system is complex and largely uncomprehended [Kandel et al., 2013]. Researchers search for ways to ‘peel back’ biomechanics to ‘reveal’ information about the underlying controller. One way this can be done is by looking for outputs of the motor system which are robust to the activities or action [Hogan, 2017]. In the last 50 years, several key ‘robust’ observations have been reported; these ‘robust’ observations have established a basis for our knowledge of human motor control.

### 1.2.1 Coordinate System for Movement Planning

The human body has hundreds of degrees of freedom. If we focus on the reaching movements of the hand, the upper limb has 9-10 degrees of freedom. However, only 6 degrees of freedom are required to specify the position and orientation of the hand – thus, the system is redundant. Yet, humans seem to solve this problem effortlessly. To gain insight into how humans solve the redundancy problem, researchers have investigated movement trajectory formation.

The human limb is comprised of largely rotational joints actuated by muscles that contract to produce joint torque. To move the hand, one could plan movements in endpoint, joint, or muscle length coordinates (while many other choices exist). Several studies have been cleverly designed to tease out clues to understand the coordinate frame humans use to plan movements.

In a study by Morasso [Morasso, 1981] subjects’ hands were confined to move in a

---

<sup>1</sup>In experiments where neural measurements are possible, only a fraction of the total population can be recorded.

<sup>2</sup>Recent work by Martin et al. applied force pulses to the Achilles tendon and measured the speed of the induced wave [Martin et al., 2018]. A measure of wave speed and tendon cross sectional area made it possible to estimate the force acting through the Achilles tendon during walking and running. However, this method can not be applied to all tendons.

plane. Subjects' shoulders and elbows were fixed – only rotation about the elbow and shoulder was allowed. Subjects were instructed to perform point-to-point reaching movements. The hand trajectory, and the joint angular trajectories were computed. The joint trajectories changed markedly for different movements. The hand end-point paths were relatively straight, with smooth bell-shaped velocity profiles. The observation of simple and consistent straight paths in hand end-point space indicated that subjects formulated motor commands in terms of hand coordinates.

In a study by Lackner and Dizio [Lackner and Dizio, 1994] subjects were instructed to perform forward reaching to a goal. During the task the subjects could not see their hand. When unperturbed, subjects made straight reaching movements. However, when the entire room was rotated, producing Coriolis forces on the limb, their hand paths were distorted. After multiple repetitions, subjects again began making straight hand paths. This indicates that humans plan movements in endpoint space.

In another study, subjects grasped a manipulandum which induced a curl field at the hand [Shadmehr and Mussa-Ivaldi, 1994]. When the curl field was turned off, subjects made straight reaching movements. However, when the curl field was turned on, their hand paths were distorted. With practice, despite the presence of the curl field, subjects began to produce straight hand movements in Cartesian coordinates. The recovery of straight movements in the presence of a disturbing force field is referred to as motor adaptation. Interestingly, when the curl field was subsequently removed, subjects produced the mirror image of the effect observed when the curl field was first experienced (before motor adaptation had occurred). This supports the idea that humans plan movements in end-point coordinates. Notably, the motor adaptation was observed in other sections of the work space that had not been explored.<sup>3</sup> This observation suggests that motor adaptation may result from a learned internal model of the limb dynamics, not simply a relationship between visited states and experienced forces.

In a fourth study, subjects received visual feedback about their hand position. The subjects' hands were occluded, and visual feedback of their hand position was

---

<sup>3</sup>However, the effects of motor adaptation were not generalized to all regions of the work space.

provided on a digital display. In certain trials, a nonlinear transformation was applied to the visual feedback, such that when the subjects made straight lines with their hands, curved lines were observed in the visually perceived space – the digital display. With practice, while the nonlinear transformation was applied, subjects made curved trajectories in real hand space in order to produce straight trajectories in visually perceived space – the digital display [Flanagan and Rao, 1995]. This demonstrated that subjects not only plan movements in endpoint coordinates, they plan movements in visually-perceived endpoint coordinates.

These unconstrained motion studies of multi-joint reaching clearly document the importance of hand kinematics. When force fields or visual distortions are applied, humans spontaneously adjust muscle force to restore straight hand paths in visually perceived space.

### 1.2.2 Smoothness

In point-to-point unconstrained motions, approximately straight-line hand paths with smooth bell-shaped velocity profiles are observed. A model that accurately explains the qualitative and quantitative features of human reaching movements was proposed by Hogan [Hogan, 1982, Hogan, 1984b] and later confirmed experimentally [Nelson, 1983, Flash and Hogan, 1985]. When performing unconstrained movements, the hand trajectory was described extremely well by minimizing the squared jerk – the derivative of the acceleration (plots of example movements can be seen in Figure 1-2). To calculate the mean squared jerk trajectory for a planar 2D movement, the cost function,  $Cost$ , must be minimized over a time interval from 0 to the final time,  $t_f$ . This can be written as

$$Cost = \frac{1}{2} \int_0^{t_f} \left( \left( \frac{d^3x}{dt^3} \right)^2 + \left( \frac{d^3y}{dt^3} \right)^2 \right) dt \quad (1.2)$$

where  $x(t)$  and  $y(t)$  are functions of time  $t$  and Cartesian coordinates of hand space. This observation was robust to movement amplitude, duration, initial/final position translation, and rotation of the line joining the two points. This model required that the movements were relaxed and well within the limits of the human motor



system. Consistent with the previous observation that humans plan movements in hand coordinates, this model only predicted the correct hand trajectory when the minimum jerk cost function was formulated in endpoint coordinates, not in joint coordinates. This model is competent to explain the observations, but does not imply that the brain contains the machinery to perform optimization, nor does it imply the brain uses optimization to plan movement trajectories.




Figure 1-2: The (a)position, (b)velocity, and (c)acceleration profiles for a 6-inch minimum-jerk motion of 700 ms duration between two equilibrium positions. Horizontal axes represent time in seconds; vertical axes represent (a) degrees, (b) degrees per second, and (c) degrees per second squared. Figure from Hogan [Hogan, 1984b].

In a study of unconstrained two joint upper limb reaching, Abend et al. [Abend et al., 1982] reported that human hand trajectories exhibit a coincidence between curvature peaks and velocity valleys in drawing movements – speed decreases when the curvature increases as seen in Figure 1-3.




Figure 1-3: The subject was instructed to move using a curved path to reach the target. (left) The hand path is illustrated by plotting the hand location every 10 ms. Arrows indicate direction of movement along the trajectory. (right) The hand speed and curvature profile are plotted. The curvature profile is the shallower of the two. Arrowheads indicate points of local curvature maxima along the trajectory; these curvature maxima are also denoted by vertical lines over the curvature profiles. Figure from Abend et al. [Abend et al., 1982].

When studied in explicitly elliptical trajectories, the detected velocity-curvature relationship was described by a two-thirds power law [Lacquaniti et al., 1993], such that the covariation between instantaneous angular velocity,  $A(t)$  and the curvature,  $C(t)$  was expressed as

$$A(t) = KC(t)^{2/3} \quad (1.3)$$

where  $K$  is a constant. The law can also be written in terms of a one-third power law relating the instantaneous tangential velocity,  $V(t)$  and the radius of curvature,  $R(t) = 1/C(t)$ , expressed as

$$V(t) = KR(t)^{1/3} \quad (1.4)$$

The power law seemed to hold for drawing ellipses and several other geometric figures (such as a lemniscate); however, the reason for this relationship was not clear. Viviani and Terzuolo [Viviani and Terzuolo, 1982] suggested the relationship to be a result of central computational constraints, which occur during the translation of movement trajectories into the appropriate motor parameters. However, others wondered if this relationship could simply be a result of biomechanics. It intuitively makes sense that the limb slows down, to some extent, when accelerating around a corner because

the arm is an inertial body actuated by muscles – finite force actuators with limited stiffness.

Two cleverly-designed studies strongly support that the power law is not a result of biomechanics<sup>4</sup>. Massey et al. [Massey et al., 1992] performed an experiment where subjects grasped a 3D isometric handle and exerted forces continuously to draw circles, ellipses, and lemniscates with and without visual feedback. Interestingly, the two-thirds power law was observed even when the hand did not move, indicating that inertial biomechanics is not the cause of the curvature-velocity relationship.

Research has shown that movement planning requires the cooperative interaction of large neuronal populations [Georgopoulos et al., 1986]. Schwartz [Schwartz, 1994] used a population vector method to transform neuronal activity to the spatial domain. This information was then used to visualize motor cortical representation of the hand trajectories made by monkeys as they drew spirals. In this cortical representation, a power law relating speed and curvature was observed. The findings of Massey et al. and Schwartz provide strong evidence that the power law relating velocity and curvature results from neural commands – not biomechanics.

Later research indicated that minimum jerk trajectory formation predicts the two-thirds power law relationship [Viviani and Flash, 1995]. Moreover, a study by Huh et al. [Huh and Sejnowski, 2015] minimized jerk during trajectory formation when a larger variety of paths were prescribed. They predicted, and then subsequently observed, different power laws, or alternatively a combination of power laws, depending on the prescribed hand path (see Figure 1-4). This curvature/velocity relationship appears to be a limitation of the human control system, and is hypothesized to be imposed by the controller used by humans to plan movements.

---

<sup>4</sup>Work by Hicheur et al. demonstrated the presence of the speed-curvature relation in human walking trajectories [Hicheur et al., 2005].

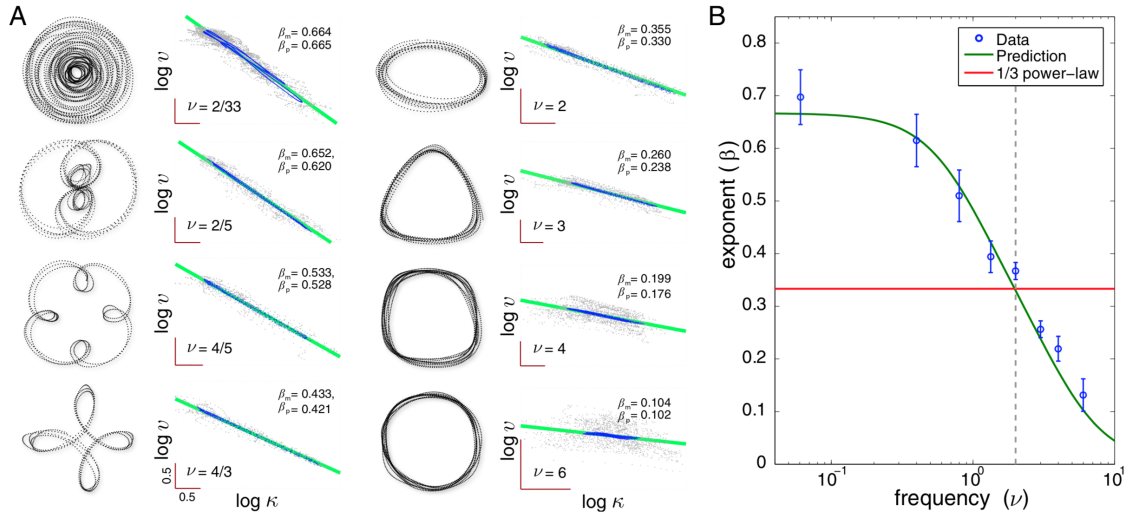


Figure 1-4: Measurements of humans drawing pure frequency curves compared with parameter-free model predictions. (A) Typical movement examples from a single subject shown to the left. Black dots: movement trajectory data. Log speed vs. log curvature shown to the right. Gray dots: raw curvature-speed data. Blue dots: bandpass-filtered curvature-speed data. Green line: minimum-jerk model prediction. Predicted  $\beta_p$  and measured  $\beta_m$  are given in the inset for each  $\nu$ . (Scale bar: 0.5.) (B) Power law exponents ( $\beta$ ) as a function of angular frequency ( $\nu$ ). Blue circles: mean exponent values from  $N = 8$  subjects measured from the slopes of  $\log(v) - \log(\kappa)$  plots. Error bars: SD. Red line: 1/3 power law. Green line: minimum-jerk model prediction. Figure from Huh and Sejnowski [Huh and Sejnowski, 2015].

Maurice et al. [Maurice et al., 2018b] extended this idea to investigate human interaction with a robot handle. They asked subjects to minimize the force exerted on a robot manipulandum. The manipulandum moved along an ellipse with different speed profiles. Interestingly, when the robot moved with a trajectory consistent with the speed-curvature relation, subjects were able to produce the smallest forces.

### 1.2.3 Movement Intermittency

Movement intermittency, the inability to move smoothly and continuously, has been observed in slow human movements [Hogan et al., 1999]. In addition, studies with stroke patients during rehabilitation have documented ‘fragmented’ movements composed of highly-stereotyped sub-movements [Krebs, 1997, Rohrer et al., 2004]. This is vastly different from normal human movements, which are usually quite smooth and

follow a minimum jerk trajectory [Flash and Hogan, 1985]. These findings are not new. For more than 100 years researchers have documented movement intermittency during cyclical movements [Crossman and Goodeve, 1983, Woodworth, 1899], eye saccades [Collewijn et al., 1988], and slow movements [Vallbo and Wessberg, 1993]. At slow speeds, these observations are unexpected. As speed decreases, all effects which make motion control difficult – feedback delays, inertial dynamics, and muscle noise – also decline. Therefore, this variability cannot result from human biomechanics. Thus, this unexpected observation appears to be a result of the neural controller.

Doeringer and Hogan [Doeringer and Hogan, 1998a] designed a single joint experiment in order to determine where the source of this intermittency is introduced into the system. They saw three possible explanations. First, the intermittency was in the feedback pathway (Figure 1-5a). Second, the intermittency could not be bypassed and was part of the forward pathway (Figure 1-5b). Lastly, the intermittency was in the forward pathway but could be bypassed (Figure 1-5c). They occluded the hand in order to test if the intermittency resulted from the visual feedback pathway. They also provided explicit visual feedback about hand velocity on a display. They then asked the question, could subjects eliminate intermittency if they were made explicitly aware of their errors via visual feedback. They found that in single-joint movements, subjects were unable to eliminate the intermittency, and that the intermittency did not result from the presence of visual feedback. This indicated that the intermittency was not in the visual feedback pathway, and could not be bypassed, as represented by Figure 1-5b.

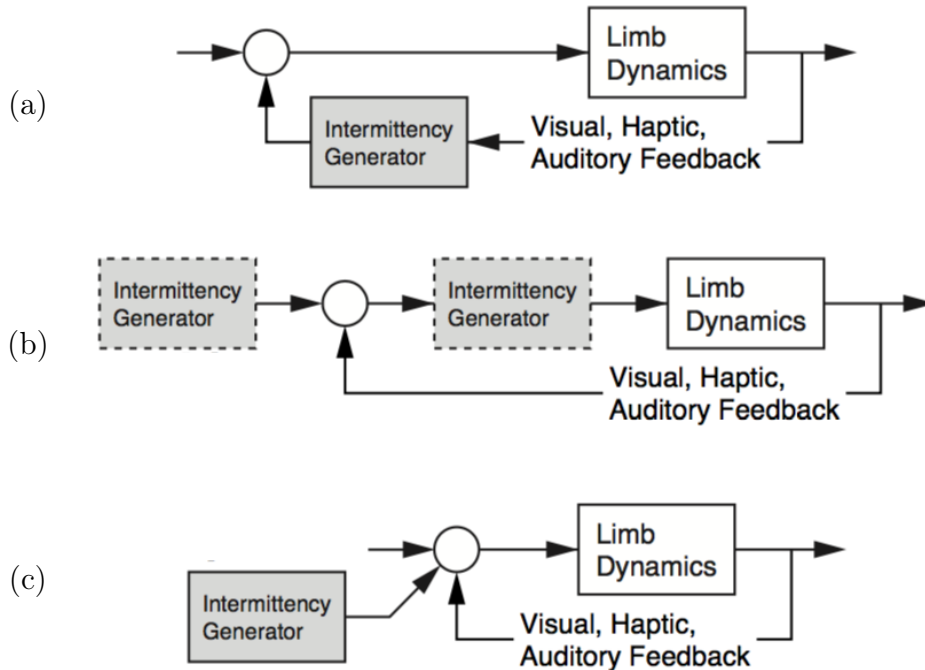


Figure 1-5: Three possible locations for an ‘intermittency generator’. In the top diagram, (a) the generator is in the feedback path, (b) the generator is in the forward path, (c) the generator is in the forward path but can be bypassed. Figure from Doeringer and Hogan [Doeringer, 1999].

This intermittency is believed to result from the composition of a motion plan made up of discrete submovements [Rohrer and Hogan, 2003]. In a study of point-to-point reaching movements by Park et al. [Park et al., 2017], subjects could not execute slow smooth rhythmic movement; instead they ‘defaulted’ to a different approach, composing motion using overlapping submovements.

### 1.2.4 Oscillations

Humans prioritize smoothness and predictability [Bazzi et al., 2018, Maurice et al., 2018a, Sternad, 2017, Zhang et al., 2018]. Arguably, the most predictable action is periodic; in theory, strictly periodic actions are infinitely predictable. Rhythmic movement is old phylogenetically, and available evidence indicates that oscillatory behavior can be evoked and sustained with minimal intervention from the central nervous system [Brown and Sherrington, 1911, Brown, 1914, Grillner and Wallen,

1985].

A rhythmic movement could be composed of multiple discrete movements. A study by Schaal et al. [Schaal et al., 2004] investigated the difference between rhythmic and discrete tasks using functional magnetic resonance imaging (fMRI). They showed that in addition to areas activated in rhythmic movement, discrete movement involved several higher cortical areas. These results provide neuroscientific evidence that rhythmic arm movement cannot be part of a more general discrete movement system.

Lastly, there is a limit to rhythmic actions: When the period is longer than about 2 to 5 seconds, an action can no longer be perceived nor executed as periodic [Fraisse, 1984, James, 1890].

In summary, studies of unconstrained motions at preferred speeds have found that humans move in generally straight-line paths, planned in visually-perceived hand coordinates. Humans move with bell-shaped velocity profiles which are well described by minimizing the jerk of the hand. As predicted by minimum jerk, a relationship between speed and curvature is observed in the hand trajectory. Neuroscientific evidence suggests there is a difference between rhythmic and discrete actions. Curiously, human movements unexpectedly become jerky at slow speeds.

### 1.2.5 Impedance

An early study by Cannon et al. [Cannon and Zahalak, 1982] identified the frequency response of the elbow joint by applying small amplitude oscillations using position disturbances. Until this point, a majority of muscle research investigated the dynamics of individual muscles. This was one of the first works to show that in practice, despite the complexity of neurophysiology, the apparent behavior of an entire joint may be well described by a linear second-order system.

In 1985, Hogan proposed Impedance control [Hogan, 1985a, Hogan, 1985b, Hogan, 1985c]. Manipulation fundamentally requires the mechanical coupling of the manipulator and the object being manipulated. Hogan proposed a theoretic framework to control the dynamic interaction between the manipulator and the environment.

Hogan rigorously extended this idea to the neural, muscular, and skeletal systems to control multiple degree of freedom posture and movement [Hogan, 1985d]. The work of Mussa-Ivaldi et al. [Mussa-Ivaldi et al., 1985] experimentally observed some of the first evidence consistent with these theoretical predictions. They reported measurements of stiffness during multiple-degree-of-freedom static posture. An early review by Kearney and Hunter [Kearney and Hunter, 1990] summarizes many of the key works, methods, and concerns within the estimation of limb dynamics.

A study by Bennett et al. [Bennett et al., 1992] produced one of the first time-varying estimates of mechanical impedance. They applied small pseudo-random force disturbances to the elbow during voluntary movement. The elbow impedance was well described by a quasi-linear second-order model.

Not long after, Lacquaniti et al. [Lacquaniti et al., 1993] estimated the dynamics of the elbow and wrist during ball-catching. This was one of the first studies to estimate time-varying multi-joint impedance. They found evidence that impedance changed before expected contact.

In addition to regulating end-point impedance by altering muscle activity, Hogan theoretically demonstrated that apparent stiffness is also a function of limb configuration. End-point stiffness results from a combination of both muscle stiffness and kinematic stiffness. One of the first works which required kinematic stiffness to be taken into account experimentally was that of Perreault et al. [Perreault et al., 2001]. They experimentally investigated how two factors, force level and limb configuration, influenced impedance during static posture.

In a planar point-to-point reaching task, Burdet et al. [Burdet et al., 2001] added a destabilizing force field. They observed that subjects changed their muscle activity in a way that increased limb stiffness to stabilize the unstable task. A related experiment was developed to look at kinematic stiffness. A destabilizing control was applied to a manipulandum. The task was designed to allow for kinematic redundancy. In this case, subjects systematically changed their limb configuration to increase kinematic stiffness in the destabilized direction [Trumbower et al., 2009, Krutky et al., 2013]. These were some of the first experiments to explicitly show that humans regulate



impedance by changing muscle activity and skeletal configuration to accomplish stable interaction – consistent with the predictions of Hogan [Hogan, 1985d].

Numerous other works have estimated limb impedance in different combinations of single-joint/multi-joint and static/motion conditions [Tsuji et al., 1995, Tsuji, 1997, Franklin et al., 2003, Franklin et al., 2007, van de Ruit et al., 2020].

Many of the developments in the algorithms for estimating impedance have been reported in a slightly different area of research, the study of the ankle joint. Some of the researchers that have made key contributions include Kearney et al. [Kearney and Hunter, 1990, MacNeil et al., 1992, Lortie and Kearney, 2001, Ludvig et al., 2011, Guarín and Kearney, 2017], Rouse et al. [Rouse et al., 2013, Rouse et al., 2014, Ficanha et al., 2015, Nalam and Lee, 2017], and Lee et al. [Lee et al., 2014a, Lee et al., 2014b, Lee and Hogan, 2015]. These methods are discussed in depth later in this thesis.

Interestingly, unconstrained upper-limb research has remained somewhat removed from estimates of impedance. One exception is a study by Flash [Flash, 1987]. In this work she demonstrated that the slight irregularities displayed in a center-out reaching task could be explained in remarkable detail. She implemented a joint-space impedance controller by assuming a constant joint-space stiffness based on static postural estimates [Mussa-Ivaldi et al., 1985]. In addition, she assumed that the zero-force trajectory  $x_0$  was described by a point-to-point minimum jerk trajectory. Prior to the content presented in this thesis, no other work has estimated impedance and/or the zero-force trajectory during a task which involves both motion and force.

## 1.2.6 Constrained Motion Studies

There are few studies that have investigated tasks involving curved constraints. Russell and Hogan [Russell and Hogan, 1989] performed studies on vertical crank turning, and documented non-zero radial forces exerted by subjects.

With the wrist joint immobilized and shoulder constrained, the task of turning a crank has one degree of freedom. Yet, the human arm has more than six muscles which may be used to execute the crank turning task. This results in muscle redundancy. Several works have investigated optimality criteria as a solution to muscle redundancy.

Svinin et al. [Svinin et al., 2001] developed a dynamic model of the crank rotation task. They wanted to know if humans modulate the rotational stiffness of the crank and arm system to introduce an additional constraint. Ohta et al. [Ohta et al., 2004] used the same model to investigate optimality criteria, and to explain how humans deal with muscle redundancy. Two cost functions explained observed velocity, force, and muscle activation during crank turning. One cost function combined hand contact force and joint torque with a weighting coefficient. The other cost function combined hand contact force and muscle force with a weighting coefficient.

McIntyre and colleagues performed experiments where subjects interacted with a concave hemispherical surface constraint. This work suggested that subjects used a feedforward controller [McIntyre et al., 1995], preferred to move along a geodesic path [Berman et al., 2014], and that limb impedance was changed between unconstrained and constrained motion [Damm and McIntyre, 2008].

In a study by Koeppen et al. [Koeppen et al., 2017], subjects interacted with a circular constraint imposed by a robotic manipulandum. They predicted that normal forces would be observed if subjects minimized muscular effort to perform the task – consistent with previous observations. However, at certain crank positions muscles transition from lengthening to shortening or vice versa – a joint reversal (see Figure 1-6). When a joint reversal occurs, the muscle cannot contribute to the task. Therefore, Koeppen et al. predicted that the muscle should be inactive at the reversal points. As subjects turned at an extremely slow speed, 13.3 seconds per revolution, the system was considered quasi-static. Thus, they predicted that there should not be a difference between clock-wise (CW) and counter clock-wise (CCW) turning directions. Contrary to these two predictions, they observed that muscles were not inactivated at reversal points, and that there was a significant effect of turning direction. Thus, they concluded that humans do not minimize muscular effort to perform the crank-turning task.




Figure 1-6: Colored radial lines depict reversal positions for the 3 muscle types. Dark and light shading denotes two distinct reversals for each muscle group. At these positions along the circular constraint, the respective muscle cannot generate tangential force. If subjects are minimizing muscle effort to perform this task, Koeppen et al. predicted that the respective muscles would be deactivated at these positions along the constraint. Figure from Koeppen [Koeppen et al., 2017].

### 1.3 Dynamic Primitives

Mounting evidence indicates that human motor control is modular, composed of primitive actions [Giszter, 2015, Kargo and Giszter, 2008, Nah, 2020, Park et al., 2017, Schaal et al., 2004, Sternad et al., 2000, Sternad and Schaal, 1999]. We have proposed that to achieve highly dynamic and dexterous performance despite neuro-mechanical limitations, human behavior is composed of dynamic primitives [Hogan, 2017, Hogan and Sternad, 2007, Hogan and Sternad, 2012, Hogan and Sternad, 2013, Schaal et al., 2000]. These are envisioned as dynamic attractors (for example, limit-cycle oscillations) that emerge from nonlinear interactions between neural and mechanical parts of the system and, once evoked, require minimal continuous intervention from higher levels of the central nervous system. In this theory, dynamic primitives are ‘building

blocks’ of complex actions. The parameters of these ‘building blocks’ are encoded; this facilitates human learning, performance, and retention of complex skills. The primitives are simultaneously and sequentially combined to produce force and motion. In practice, this may be described by defining a zero-force trajectory composed of sub-movements and/or oscillations interacting with impedances in a Norton equivalent network [Hogan, 2014].

Movement primitives have also been defined at the functional level, as an action to be performed on a specific object such as reaching, grasping, or pushing [Chavandafle and Rodriguez, 2015, Doshi et al., 2020, Giszter, 2015, Hogan et al., 2020]. The discussion of dynamic primitives in this work is distinct from these functional task definitions of primitive, in the sense that this is a default control that always emerges and is not bypassed – a fundamental building block. In the same way that language is composed of multiple levels of structure such that letters compose words, and words compose sentences, dynamic primitives may be thought of as the ‘letters’, while movement primitives based on a functional task could be ‘words’. The important point is that the existence of dynamic primitives does not preclude further organization or simplification at higher levels.

## 1.4 Summary of Thesis

Chapter 2 presents a study of humans turning a horizontal planar crank. To address this challenge, we estimated the zero-force trajectory ( $x_0$ ) using measures of force, measures of motion, and assuming a model of limb impedance. Here we demonstrate the value of this method: remarkably, a speed-curvature relation emerged similar to that seen in unconstrained movements, even though the hand moved at nearly constant speed along a constant-curvature path.

In Chapter 3 the crank turning experiments revealed that when the hand was constrained to move in a circle, non-zero forces against the constraint were measured. In this Chapter, further analysis from this experiment exposed two artifacts that could not result from mechanics alone, but may be attributed to neural control via primi-

tive dynamic actions. Estimated zero-force trajectories were approximately elliptical, and their orientation significantly (and substantially) differed with turning direction. This observation is consistent with control using underlying oscillations to generate an elliptical zero-force trajectory. However, for periods longer than 2 to 5 seconds, motion can no longer be perceived or executed as periodic. Instead, it decomposes into a sequence of submovements, and that results in increased variability, which we observed at the slowest speeds. These quantifiable performance limitations support the hypothesis that humans simplify this constrained-motion task by exploiting at least three primitive dynamic actions: oscillations, submovements, and mechanical impedance.

In Chapter 4, we set out to measure mechanical impedance during the task of crank-turning. Initially, time-based ensemble methods were employed. However, a key assumption—the stationarity of the noise processes—was not justified. A ‘work-around’ was developed: (1) high pass filter to remove the influence of the underlying time-varying zero-force trajectory; (2) assume impedance is a function of configuration and identify a configuration-dependent ensemble. These methods identified the impedance parameters (mass, damping, and stiffness) in the normal direction with an error of less than 5% in simulation and 20% in robot-only experiments. However, in the tangential direction, a fundamental limitation was discovered in practice: impedance may be so low during movement that the natural frequency of impedance at any one configuration overlaps the frequency content of the non-stationary noise process present in the zero-force trajectory. This was not anticipated. The lack of frequency separation poses a fundamental challenge to system identification of mechanical impedance – both in simulation and on hardware.

In the crank-turning experiments the subject’s wrist was braced, the elbow was supported by a sling, and the shoulder was strapped to a chair. This confined the task to two degrees of freedom. While this is advantageous for modeling, it drastically simplifies human behavior. In Chapter 5 [Hermus et al., 2022], a robotic experiment was designed to probe how kinematic complexity may be managed. We assumed the impedance and programmed a Kuka LBR iiwa to interact with an InMotion

planar robot while it simulated a virtual circular constraint – as in the human subject experiments. Interestingly we found that when the nullspace dimension increased a method based on impedance superposition was comparable to nullspace projection (a review of impedance superposition was presented in [Verdi, 2019]). This suggest that high dimensionality may be a ‘blessing’ not a ‘curse’.

In robotics, hybrid control allows simultaneous, independent control of both motion and force and it is often assumed that humans can also modulate force independent of motion. In Chapter 6, we experimentally tested that assumption. Participants were asked to apply a constant 5 N force on a robot manipulandum that moved along an elliptical path. After initial improvement, force errors quickly plateaued, despite practice and visual feedback. Within-trial analyses revealed that force errors varied with position on the ellipse, rejecting the hypothesis that humans have independent control of force and motion. The findings are consistent with a feed-forward motion command composed of two primitive oscillations acting through mechanical impedance to evoke force.

In Chapter 8, the main findings of the work are reviewed and future work directions are discussed.

As this thesis includes peer-reviewed published papers (e.g. in Chapters 2, 5, and 6), there may be a certain amount of repetition. This repetition was retained as it may help to re-cap the context of the separate studies.

## Chapter 2

# Separating Neural Influences from Peripheral Mechanics: The Speed-Curvature Relation in Mechanically-Constrained Actions

This chapter is an adapted version of [Hermus et al., 2020] published in the Journal of Neurophysiology.

### 2.1 Introduction

Experimental studies of motor neuroscience should start with macroscopic behavior as there are some remarkably robust patterns that can guide the investigation of neural control [Krakauer et al., 2017]. One such pattern robustly observed in curved motions is a power-law relation between hand path curvature and speed: hand speed decreases as curvature increases. This finding has been widely reported in unconstrained movements for more than 30 years, though its implications for neural control have been controversial [Abend et al., 1982, Dayan et al., 2007, Gribble and Ostry, 1996, Hicheur et al., 2005, Huh and Sejnowski, 2015, Massey et al., 1992, Schaal and Sternad, 2001, Schwartz, 1994, Viviani and Flash, 1995, Viviani and Stucchi, 1989, Vi-

viani and Terzuolo, 1982]. Specifically, discussion of the power law debated whether this robust relation is due to biomechanical or neural influences. The present study examined a simple constrained-motion task to test whether this pattern persists even when mechanical conditions prevent its expression. It did, suggesting that the speed-curvature relation is present in the neural organization of action.

While human dexterity vastly exceeds that of modern robots, the human neuro-mechanical system is orders of magnitude slower than its robotic counterparts [Kandel et al., 2013, Loram et al., 2005, Shepard and Metzler, 1971]. Slow neural transmission and muscle response implies that humans have to rely heavily on feed-forward (i.e. predictive) control, especially when physically interacting with objects and environments where bidirectional interaction forces arise. Prior work on the control of dynamically complex objects showed that humans adjust their behavior to prioritize predictability of the object dynamics [Bazzi et al., 2018, Maurice et al., 2018b, Nasserolelami et al., 2014, Sternad, 2017, Zhang et al., 2018]. Here we extend the study of predictability to kinematically constrained actions.

One way to move predictably is to move smoothly; and smoothness, quantified by mean squared jerk, has been demonstrated in numerous studies of unconstrained reaching [Flash and Hogan, 1985, Hogan, 1982, Hogan, 1984b, Nelson, 1983, Sternad et al., 2010]. Smoothness can also account for the relation between hand path curvature and tangential speed [Abend et al., 1982, Viviani and Terzuolo, 1982]. For simple curves (e.g. ellipses) this is the so-called ‘two-thirds power law’ [Schaal and Sternad, 2001, Sternad and Schaal, 1999, Viviani and Flash, 1995]. For more complex curves, smoothness requires different exponents and more complex relations (a ‘spectrum of power laws’) [Huh and Sejnowski, 2015], but in all cases a temporal coincidence between extrema of tangential speed and path curvature is derived from theory and observed in practice. A power-law speed-curvature relation is consistent with a preference for smoothness and predictability.

Much of the previous neuroscience literature has focused on simple unconstrained movements. While these paradigms have rendered manageable data for analysis and modeling, it is difficult to generalize the insights gained to understand tasks that



involve physical interaction, essential in any kind of tool use. While some animals are capable of making and using tools, this ability is vastly more developed in humans [Boesch and Boesch, 1990, Hunt, 1996, Johnson-Frey, 2004, Kenward et al., 2005]. Many tools and daily interactions with the environment introduce kinematic constraints, such as turning a steering wheel, opening a door, and turning a crank. For example, opening a door or turning a steering-wheel confines the hand to a planar circular path with constant curvature; if the speed-curvature relation generalizes to object manipulation, no speed fluctuations should be evoked. However, if that relation is of primarily neural origin, it may still influence behavior even in this situation. This study tested whether a speed-curvature relation – a corollary of predictability – was present in the neural command even in those cases.

A central challenge in motor neuroscience is to tease apart the contributions of neural control and biomechanics, a problem that is even more pronounced during interactive tasks. During physical interaction, a limb is fundamentally subject to at least two sets of inputs for every output: Force exerted on an object depends not only on neural activity but also on the object’s motion. The relation between displacement and the force it evokes may be characterized mathematically by mechanical impedance (a dynamic generalization of stiffness). In principle, knowledge of mechanical impedance combined with simultaneous measurement of force and motion during object manipulation would allow us to ‘subtract off’ or ‘peel back’ peripheral biomechanics to uncover underlying neural influences. In practice, mechanical impedance is nonlinear and time-varying and measuring it during movement, though possible, is challenging [Kearney and Hunter, 1990]. Moreover, measurement inevitably introduces perturbations that may alter behavior.

An alternative (pursued here) is to approximate mechanical impedance using a plausible mathematical model based on measurements made under static postural conditions. Given this model and observations of actual motion and exerted force, we define a construct that we call the ‘zero-force trajectory’ the trajectory that would have been followed if the external forces were zero. The zero-force trajectory summarizes one consequence of neural activity and expresses it as a quantity that may

be compared with actual motion. To evaluate its sensitivity to the uncertainty about mechanical impedance, the results were tested over a wide range of parameter variations.

Motion under a circular constraint, i.e. turning a crank, imposes a constant curvature hand path. If a speed-curvature relation applies to hand motion during object manipulation, turning a crank should not elicit any variation of hand speed. However, the zero-force trajectory is not confined to constant curvature. If a speed-curvature relation, consistent with smoothness and predictability, is a feature of neural control, it should be manifest in the consequence of neural action that is expressed by the zero-force trajectory. The central hypothesis tested in this study was that the zero-force trajectory exhibits a systematic relation between curvature and speed.

To improve the ability to subtract off the influences of inertial mechanics and neuro-muscular dynamics, crank-turning was studied at three speeds covering a wide range: 1) a speed close to the fastest that subjects could manage, such that inertial effects predominate; 2) a speed near subjects' preferred or comfortable speed; 3) an extremely slow speed, such that the action was quasi-static and all dynamic effects were negligible. To encourage movements with constant speed along the constant curvature path, we provided a visual display of instructed speed together with visual feedback of actual hand speed. As inertial behavior is strongly directional and varies with limb position, we examined clockwise and counter-clockwise rotation.

Remarkably we found that, independent of the direction of rotation, when peripheral biomechanics were subtracted, the zero-force trajectory exhibited a coincidence of curvature and velocity extrema resembling that reported for unconstrained motion.

## **2.2 Methods**

### **2.2.1 Participants**

Ten healthy male college-age students were recruited for the study. All participants were right-handed, and none reported any biomechanical injury to their arm nor any

neurological problems. Prior to participating in the study, they were informed about the experimental procedure and signed the informed consent document approved by MIT's Institutional Review Board.

### 2.2.2 Experimental Apparatus and Procedure

The crank used in this experiment is shown in Figure 2-1. The crank arm was mounted on a high precision incremental optical encoder/interpolator set (Gurley Precision Instruments encoder #8335-11250-CBQA, interpolator #HR2-80QA-BRD) with a resolution of 0.0004 degrees per count. A six-axis force transducer (ATI Model 15/50) was attached to the end of the crank, with a handle mounted on it. A spool managed the force transducer cable.




Figure 2-1: Experimental setup. The crank displayed in the inset was used to provide a circular constraint. Vision of the arm and crank was occluded but the subject was provided with visual speed feedback. The wrist was braced, the elbow was supported by a sling, and the shoulders were strapped to a chair.

During the experiment, the subject's arm was occluded from view by a wooden structure, which did not limit the range of motion. The upper arm was suspended

by a canvas sling connected to the ceiling using a steel cable; upper and lower arm were in the plane of the crank. The subject sat in a chair with a rigid back, while the shoulder was constrained by a harness attached to the back of the chair. The subject was positioned such that the crank, with radius 10.29 cm, was well within the workspace of the arm.

Data acquisition was controlled by a computer running the QNX real-time operating system on an Intel Pentium 100 processor. The encoder, sampling at 200 Hz, was connected to a set of counters and to the computer via digital I/O. The ATI force transducer's signal, sampled at 100 Hz, was processed by its embedded controller and input to the computer through the digital I/O. The visual display, also generated by the computer, was on a 17-inch monitor (311 x 238 mm, resolution 1280 x 1024, 76 Hz) which was mounted approximately 75 cm from subjects' eyes. The experiment was divided into two sections: trials at subjects' preferred or 'comfortable' speed and trials at a visually-instructed speed. The design of the experimental conditions is graphically overviewed in Figure 2-2.

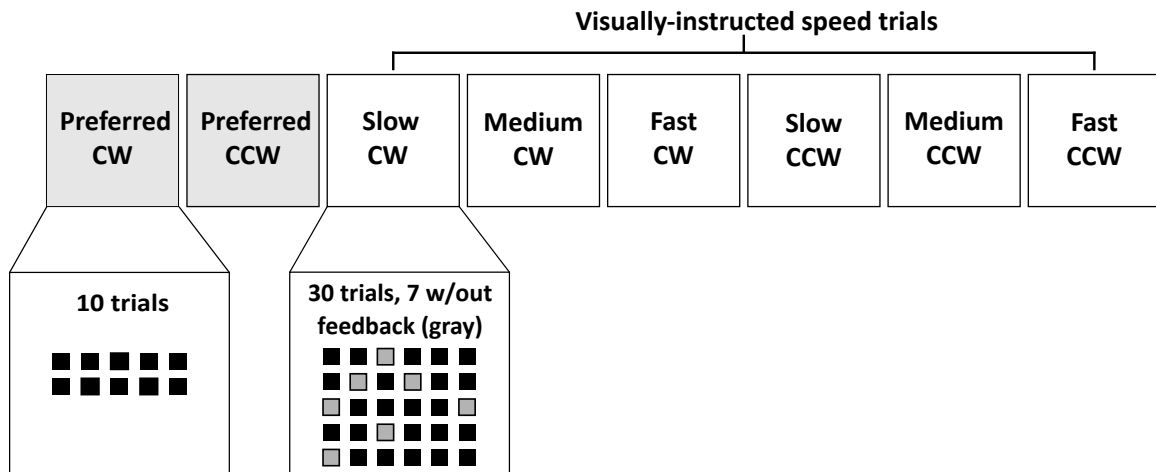


Figure 2-2: Experimental design. Each of the 10 subjects completed 20 trials at their preferred speed, 10 trials in clockwise (CW) and counterclockwise (CCW) directions. In the instructed-speed trials subjects completed 30 trials in each condition. Visual feedback was not provided on the display for 7 of the 30 trials during each of the instructed-speed blocks.

At the start of the experiment, subjects performed 20 trials at their preferred speed, 10 trials in clockwise direction (CW) and 10 in counterclockwise direction

(CCW); each trial lasted 8 seconds. Subjects were not provided any visual feedback during these trials. Thereafter, subjects performed six blocks of 30 trials each with visual specification of one of three target speeds (slow: 0.075, medium: 0.5, and fast: 2.0 revolutions per second), in either clockwise or counter-clockwise directions (Figure 2-2). The order of the speed and direction blocks was randomized across subjects. The three speeds were selected to cover a significant range: 0.075 rev/s was extremely slow (required over 13 seconds per revolution), 0.5 rev/s was close to subjects' preferred speed, and 2.0 rev/s was close to the fastest subjects could turn the crank. Visual feedback on the monitor displayed the target speed, as well as subjects' real-time hand speed. The horizontal axis was time, and the vertical axis was speed. Target speed was displayed as a continuous horizontal line in the middle of the screen. Subjects' speed was estimated using a backward finite difference algorithm. The relationship between crank motion and screen display was re-scaled for every block; the width of the screen corresponded to the time of the trial, which was a function of the desired crank speed. Seven trials in each block were 'blind' catch trials, in which visual feedback of the actual hand speed was removed, while the display of target speed was retained.

In the slow-speed conditions, each trial lasted 45 seconds; in the medium-speed conditions, each trial lasted 16 seconds; in the fast-speed conditions, each trial lasted 4 seconds. This yielded 8 turns of the crank for the fast and medium conditions, but only about 3.4 turns of the crank for the slow condition. The duration of the slow-speed trials was chosen as a compromise between acquiring adequate data and avoiding subject fatigue.

### **2.2.3 Data Processing and Analysis**

Crank speed and normal force were computed with respect to crank angle divided into 360 angle bins. As the different target speeds were sampled at the same rate with different time durations, they contained a different number of samples per bin. Therefore, the data were interpolated to the largest number of samples, which was 9001 samples at the slowest target speed, 0.075 rev/s. The interpolation was performed

using a piece-wise cubic Hermite interpolating polynomial [Fritsch and Carlson, 1980]. After interpolation, the speed or force profiles were binned into 360 sections corresponding to a full rotation of the crank. Then the mean of each bin was taken, resulting in a speed or force profile over 360 crank angles. In all trials, the first 1.5 seconds were discarded to eliminate transients before calculating summary measures. Dependent measures were mean speed, standard deviation (SD) of speed, mean normal force, and SD of normal force.

#### **2.2.4 Trials at Preferred Speed**

To characterize subjects' preferred behavior, we tested whether the mean normal force was significantly different from zero, and whether turning direction influenced the dependent measures. The speed mean and standard deviation (SD) were computed for the initial 10 trials performed at the preferred speed. Then these quantities were averaged to produce a mean speed and SD for each subject. Mean force and SD of force were similarly averaged for each subject. One subject moved much faster than all others (beyond two standard deviations of the other subjects). His data were excluded from subsequent analysis of the preferred speed trials.

Paired-sample t-tests were carried out to test for significant effects of turning direction on mean speed and normal force. Student's t-tests were used to detect whether the mean normal force was significantly different from zero. Significance values of post-hoc t-tests were adjusted using the Sidak-Bonferroni procedure, where the original significance level was defined as  $\alpha$ , the number of t-tests was  $m$ , and the corrected Sidak-Bonferroni significance values were  $\alpha_{SID} = 1 - (1 - \alpha)^{1/m}$ . All statistical analyses were performed using the SPSS statistical software package (SPSS Inc., Chicago IL); the significance level was set to 5%.

#### **2.2.5 Trials at Instructed Speeds**

In the first of the trials in which speed and direction were instructed, subjects often did not achieve the task goal; they occasionally turned in the wrong direction, did

not complete a full cycle, or turned at the wrong speed. Hence, the first trial in each block was discarded. The seven ‘catch’ trials, without visual feedback, were omitted from initial statistical analysis. The speed mean and SD were computed for each of the remaining trials.

To quantify the influence of speed and direction, a linear mixed model was employed; it was then tested using analysis of variance (ANOVA). The linear model which represented the observed dependent measure  $Y_{i,j,k}$  was expressed as

$$Y_{i,j,k} = \mu_T + \alpha_j + \beta_k + \gamma_l + (\alpha\beta)_{j,k} + (\alpha\gamma)_{j,l} + (\beta\gamma)_{k,l} + (\alpha\beta\gamma)_{j,k,l} + E_{i,(j,k,l)} \quad (2.1)$$

where the grand mean is  $\mu_T$ , the fixed effect of speed is  $\alpha_j$ , where  $j$  is an index from 1 to 3, the fixed effect of direction is  $\beta_k$ , where  $k$  is an index from 1 to 2, the random effect of subject is  $\gamma_l$ , where  $l$  is an index from 1 to 10, and the stochastic sampling effect  $E_{i,j,k}$ , where  $i$  is an index from 1 to 22 (representing the trial number excluding the first and the catch trials). As above, all statistical analyses were performed using the SPSS statistical software package (SPSS Inc., Chicago IL); the significance level was set to 5%.

## 2.2.6 Simulation of Passive Inertial Mechanics

Physical interaction with a kinematic constraint changes the nonlinear inertial mechanics with which the neuro-muscular system interacts. Despite the familiarity of this constrained-motion task, a failure to compensate perfectly for these changes might account for some of the observed patterns of force and motion. To provide insight about these highly nonlinear and non-intuitive effects we simulated the fluctuations of speed and force that would be predicted from the configuration-dependent variation of inertial mechanics subject to this kinematic constraint, without any contribution from muscle action. The arm was modeled as a two-link planar serial linkage, with no gravitational or frictional effects. Inertial parameters were estimated based on the cadaver studies of Dempster [Miller and Nelson, 1973, Plagenhoef, 1971]. The shoulder joint was located at the thorax which was assumed to be stationary. This

approach is the same as that used by Ohta et al. [Ohta et al., 2004]. The variation of motion due to closed-chain inertial dynamics was simulated. This was done by setting muscle-generated joint torques and crank damping to zero in the model. The system was initialized at the zero-degree position, with the initial angular velocities of the three target speeds: slow (0.075 rev/s), medium (0.5 rev/s), and fast (2.0 rev/s). Numerical integration was performed using MATLAB’s ode45 algorithm [Shampine and Reichelt, 1997].

### 2.2.7 Zero-Force Trajectories

One way to describe the dynamics of interaction uses a mechanical impedance operator  $Z\{\cdot\}$  (Hogan, 1985a, 1985b). The force time-function  $F(t)$  can be computed from the displacement time-function  $\Delta x(t)$ ,  $F(t) = Z\{\Delta x(t)\}$ . Displacement is defined as  $\Delta x(t) = x_0(t) - x(t)$  where  $x(t)$  is the actual hand position and  $x_0(t)$  is a zero-force trajectory. Though muscle force production is a complex function of many factors, its dominant behavior can be well-described by a function of muscle length and its rate of change [Hill, 1938, Joyce et al., 1969, Rack and Westbury, 1969]. Accordingly, a simplified model of muscle mechanical impedance was used, a linear spring and viscous damping element with common displacement [Hogan, 1984a]. To implement this model on a two-joint arm, joint stiffness was assumed to be a  $2 \times 2$  symmetric matrix, independent of configuration. Joint damping, also a  $2 \times 2$  symmetric matrix, was proportional to joint stiffness. This is similar to the muscle model previously used by Flash (1987), but in this case we used a damping term which was defined relative to the zero-force trajectory.

The joint torque was defined by

$$\boldsymbol{\tau} = \mathbf{K}(\mathbf{q}_0 - \mathbf{q}) + \mathbf{B}(\dot{\mathbf{q}}_0 - \dot{\mathbf{q}}) \quad (2.2)$$



The stiffness in units of N-m/rad was defined as

$$\mathbf{K} = G \begin{bmatrix} K_{11} & K_{12} \\ K_{21} & K_{22} \end{bmatrix} = G \begin{bmatrix} 29.5 & 14.3 \\ 14.3 & 39.3 \end{bmatrix} \quad (2.3)$$

The viscous damping in units of N-m-s/rad was defined as

$$\mathbf{B} = \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix} \quad (2.4)$$

$\mathbf{K}_{11}$  and  $\mathbf{B}_{11}$  were the net shoulder joint stiffness and damping,  $\mathbf{K}_{12}$ ,  $\mathbf{B}_{12}$ ,  $\mathbf{K}_{21}$ , and  $\mathbf{B}_{21}$  were the two-joint parameters, and  $K_{22}$  and  $B_{22}$  described the elbow parameters;  $G$  was a dimensionless scalar. The values for joint stiffness and damping were consistent with those of Flash [Flash, 1987], such that  $B = \beta K$ . The  $\beta$  term had units of time, consistent with a first-order model of muscle impedance [Hill, 1938]. A gain of  $G = 0.5$  was used in the slow and medium cases, and a gain of  $G = 1.5$  was used in the fast case. Damping was derived from stiffness by multiplication by a constant factor,  $\beta$ , which was 0.05 s for the slow and medium cases, and 0.1 s for the fast cases.

Substituting Equation 2.2 into Equation A.2, A.3, and A.4 (from Appendix A), the equation could be manipulated to solve for  $\dot{\mathbf{q}}_0$ :

$$\dot{\mathbf{q}}_0 = \mathbf{B}^{-1} \left[ \mathbf{M} \mathbf{J}^{-1} \left[ \{ \mathbf{J} \mathbf{M}^{-1} \mathbf{J}^T + r^2 \mathbf{I}^{-1} \mathbf{e} \mathbf{e}^T \} \mathbf{F} - \dot{\mathbf{J}} \dot{\mathbf{q}} - r \dot{\theta} (\dot{\theta} \mathbf{n} + b_c \mathbf{I}^{-1} \mathbf{e}) \right] + \mathbf{h} - \mathbf{K}(\mathbf{q}_0 - \mathbf{q}) \right] + \dot{\mathbf{q}} \quad (2.5)$$

Integrating Equation 2.5 allowed for the computation of the zero-force trajectory corresponding to a prescribed position, velocity, acceleration, and force. The velocity and force signals were filtered with a second-order zero-phase-lag Butterworth filter using a cutoff frequency of 10 Hz, except in the slow condition. The tangential force in the slow condition was small in magnitude. At slow speeds, a large number of samples with a magnitude close to the resolution of the sensor were observed. This resulted in artifactual step changes in the force measurements. To eliminate this artifact, the

tangential force in the slow condition was filtered with a cutoff frequency of 0.5 Hz, far faster than the turning frequency of the slow task (0.075 rev/s).

## 2.2.8 Speed - Curvature Relation

In unconstrained motions, the relation between curvature and speed has been well documented in human arm trajectories. This study investigated if this relation was evident in the zero-force trajectory. To compute the speed and curvature of the zero-force trajectory, it was transformed into Cartesian coordinates,  $x_0$  and  $y_0$ . Using the methods of Dohrmann et al. [Dohrmann et al., 1988] the derivatives of the Cartesian position of the zero-force trajectory were computed. Smoothing parameters of 0.03,  $10^{-7}$ , and  $10^{-11}$  were used for the slow, medium, and fast trials, respectively. These derivatives were used to compute tangential speed,  $V_0$ , and curvature,  $\kappa_0$

$$V_0 = \sqrt{\dot{x}_0^2 + \dot{y}_0^2} \quad (2.6)$$

$$\kappa_0 = \frac{\dot{x}_0\ddot{y}_0 - \dot{y}_0\ddot{x}_0}{(\dot{x}_0^2 + \dot{y}_0^2)^{3/2}} \quad (2.7)$$

Positive curvature would be consistent with a positive z rotation, according to the right-hand rule. However, for ease of computation the radially inward direction was defined as positive curvature.

To test the relation between speed and curvature, local minima in tangential speed and local maxima in curvature were identified and their temporal relation was quantified. For each local minimum in tangential speed, the nearest local maximum in curvature was found. The signed distance between the two extrema was normalized by the target speed and defined as the interval  $\Delta$ . This was computed for each subject in each condition, and the data were pooled across trials. The 95% confidence interval for each subject in each condition was computed.

### 2.2.9 Sensitivity to Impedance Assumptions

To the best of our knowledge, limb impedance measurements during physical interaction with a constraint have not been reported. Thus, impedance parameters were based on unconstrained static arm stiffness and single joint damping measurements. Consequently, the impedance parameters, the gain term for the stiffness,  $G$ , and the proportional damping term,  $\beta$ , were varied to understand if the results were sensitive to the impedance values used to compute the zero-force trajectory. When the gain and damping terms changed, the zero-force trajectory was expected to change. However, the main question was whether any speed-curvature relation was affected when the zero-force trajectory changed shape. To this end, the 95% confidence interval for  $\Delta$  was computed when the  $G$  and  $\beta$  term were each varied over a 3:1 range.

## 2.3 Results

### 2.3.1 Speed and Force at Preferred Speed

The preferred speed trials were collected to establish a baseline and quantify human performance without visual feedback. A representative subject's turning speed and normal force are plotted with respect to crank angle in Figure 2-3. His speed fluctuated about 0.5 rev/s and systematic variations with angular position were evident (Figure 2-3, top panel). Further, the subject's force alternated between tension (positive) and compression (negative) depending on the crank angle (Figure 2-3, bottom panel). Some minor differences appear between the two turning directions.




Figure 2-3: Representative data from one subject’s tangential speed (top) and normal force (bottom) with respect to crank position at ‘comfortable’ speed. Blue lines indicate clockwise (CW) trials and red lines indicate counter-clockwise (CCW) trials. Each line represents binned speed or force values for one of the ten trials. The shading indicates one SD from the mean across trials.

Despite the appearance of occasional differences between clockwise and counter-clockwise performance, there was no statistically significant difference between mean speed in the clockwise ( $0.37 \pm 0.11$  rev/s), and counter-clockwise ( $0.40 \pm 0.10$  rev/s) conditions ( $p = 0.539$ ). There was also no significant difference between mean normal force in the clockwise ( $-0.58 \pm 1.65$  N), and counter-clockwise ( $-0.12 \pm 1.43$  N) conditions ( $p = 0.462$ ). Furthermore, the mean normal force was not significantly different from zero in the clockwise condition ( $p = 0.325$ ,  $\alpha_{SID} = 0.0253$ ), or counter-clockwise condition ( $p = 0.811$ ,  $\alpha_{SID} = 0.0253$ ). Figure 2-4 displays these dependent measures as a function of mean speed, illustrating the lack of differences between the two rotation directions.

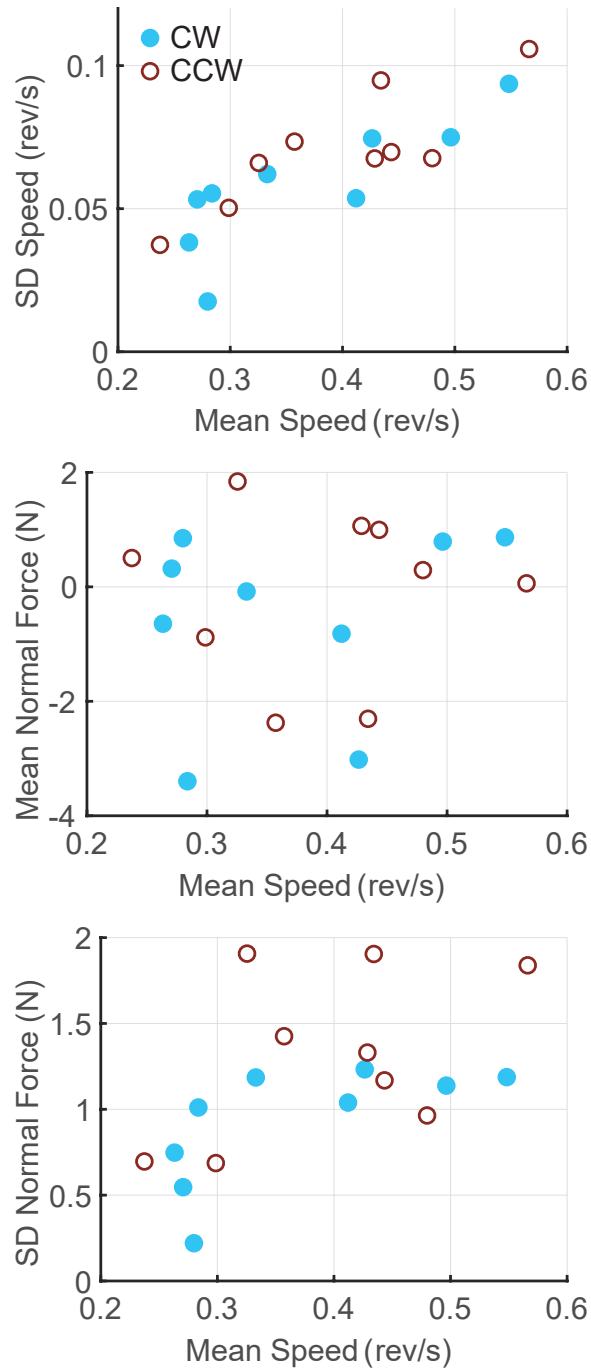


Figure 2-4: Dependent measures with respect to mean speed for preferred speed trials: standard deviations of speed (top), mean normal force (middle), and standard deviations of normal force (bottom). Blue dots: CW; red circles: CCW.

### 2.3.2 Speed and Force at Instructed Speeds

In these trials subjects were instructed to turn at one of three speeds, in both clockwise and counter-clockwise directions, with visual feedback provided. Figure 2-5 shows mean speed and force of all subjects with respect to crank angle for the slow, medium and fast conditions, revealing a noticeable pattern with respect to crank angle. Figure 2-6 summarizes the dependent measures for the three different target speeds and directions.

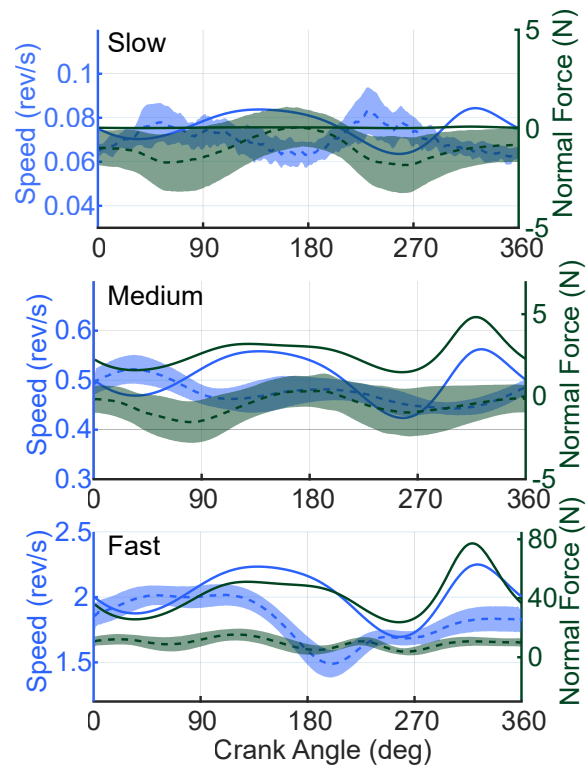


Figure 2-5: Mean and standard deviations of speed (dotted blue) and normal force (dotted green) versus crank angle for all subjects. The shading indicates one standard deviation from the mean across subjects. The solid lines show fluctuations of speed (blue) and normal force (green) due to passive inertial mechanics. Data are from the clockwise direction trials. Note the different scales on the ordinates.

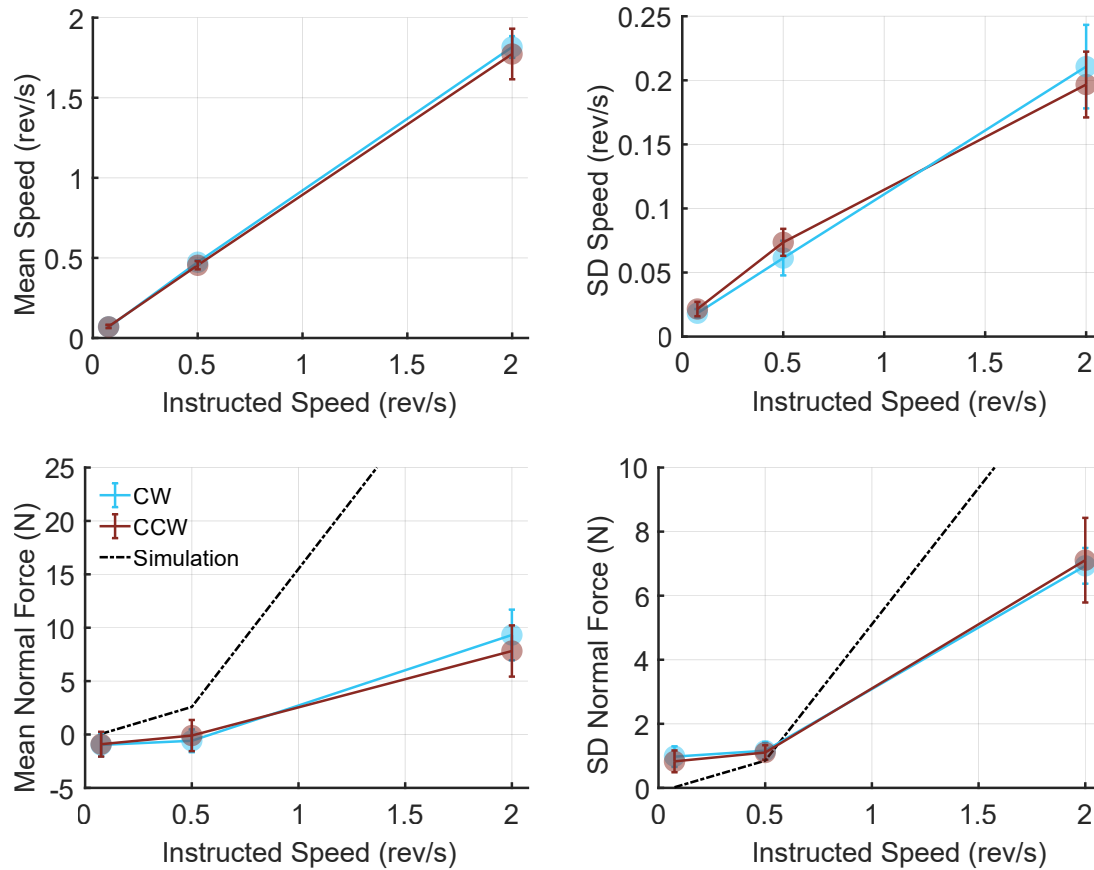


Figure 2-6: Dependent measures, mean speed, SD speed, mean normal force, and SD normal force. Error bars indicate the SD across subjects. Clockwise trials are depicted in blue, counter-clockwise trials are in red; black dashed lines indicate the fluctuations which result from only inertial effects. As inertial dynamics are not a function of direction there is only one line for the simulations (see Appendix II).

The dependent measures were analyzed with respect to the experimental conditions speed and direction and displayed in Figure 2-6. A main effect of target speed on mean speed was observed ( $F_{2,0,18} = 3579.513, p < 0.001$ ), indicating that, as expected, subjects could successfully perform this task. The SD of speed showed a significant interaction between speed and direction ( $F_{2,0,18} = 5.317, p < 0.015$ ) and a main effect of speed ( $F_{2,0,18} = 477.497, p < 0.001$ ). Two paired-sample t-tests compared the SD of speed at different instructed speeds. There was a significant difference between the slow and medium conditions ( $p < 0.001, \alpha_{SID} = 0.0253$ ), and the medium and fast conditions ( $p < 0.001, \alpha_{SID} = 0.0253$ ). Thus, target speed was the predominant factor influencing the mean and SD of speed.

Analysis of the mean normal force (Figure 2-6 lower left) revealed a significant interaction between speed and direction ( $F_{2,0,18} = 3.752$ ,  $p = 0.043$ ), and a main effect of target speed ( $F_{2,0,18} = 258.878$ ,  $p < 0.001$ ). There was no significant effect of direction. Paired-sample t-tests identified a significant difference between the slow and medium conditions ( $p < 0.001$ ,  $\alpha_{SID} = 0.0253$ ), and the medium and fast conditions ( $p < 0.001$ ,  $\alpha_{SID} = 0.0253$ ). The large difference between the medium and fast conditions was clear, while the significant difference between the slow and medium speeds was less pronounced. Lastly, analysis of the SD of normal force revealed only a significant effect of speed ( $F_{2,0,18} = 749.292$ ,  $p < 0.001$ ). Target speed was the predominant factor influencing the mean normal force.

In the instructed speed trials, occasional catch trials occurred in which visual feedback was withheld. Comparing the dependent measures for the catch trial and the immediately preceding trial with visual feedback assessed whether visual feedback affected subject performance. These comparisons resulted in several significant interactions, including three-way interactions. However, the magnitudes of these effects were very small.

### 2.3.3 Zero-Force Trajectory

The bidirectional interaction between the forces exerted on the constraint and the force observed at the hand was described by mechanical impedance. A model of upper-limb mechanical impedance was assumed, and the zero-force trajectory (defined above) was computed. Figure 2-7 shows one example from a representative subject in the three speed conditions in the left column; the right column shows the mean paths from all subjects. It is clear in all subjects that the path of the zero-force trajectory deviated from circular, both closer to and further from the center of the crank circle and was roughly elliptical. The slowest speed revealed the most elliptical path. Further, zero-force speed varied systematically along the elliptical trajectory. Visually, the speed peaks corresponded well to the curvature valleys. This zero-force trajectory enabled a test of the hypothesis that a systematic speed-curvature relation is found in the motor consequences of neural activity.



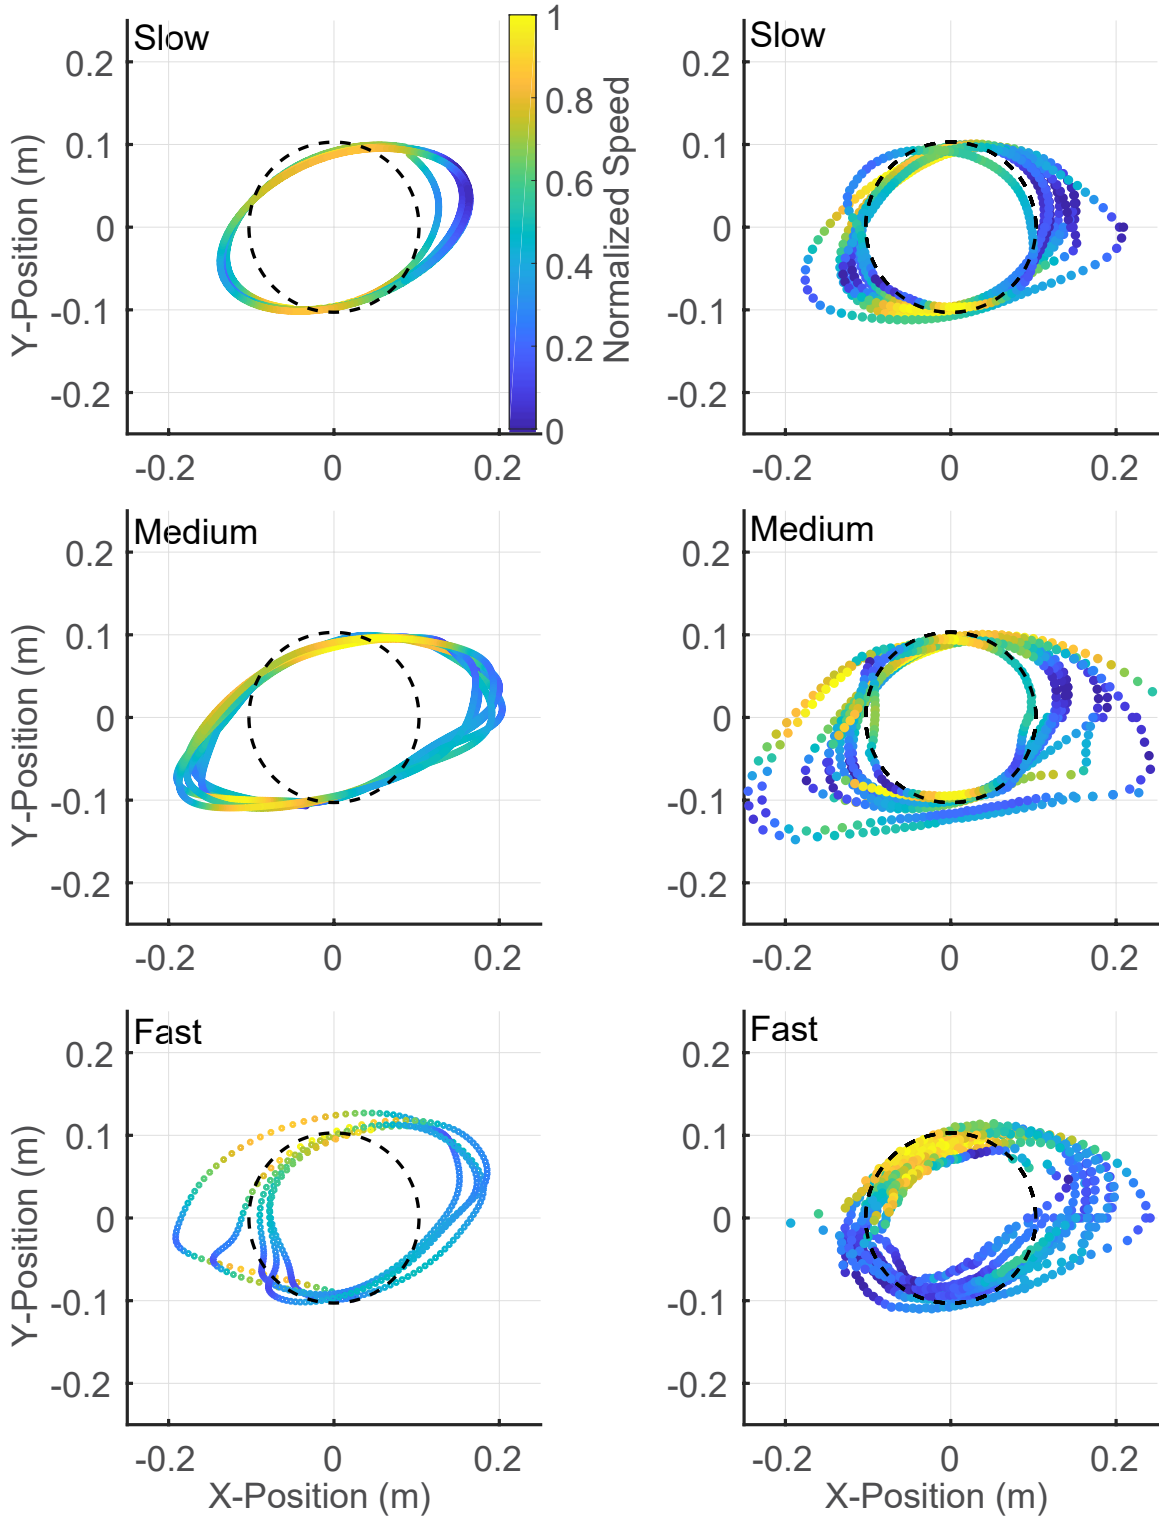


Figure 2-7: Figure 7: Left column: Representative trial from one subject in each of the speed conditions. Right column: Average zero-force trajectories for all 10 subjects. The path defined by the constraint is shown by the black dashed circle. The zero-force trajectories are shown by variable color lines. The plots are from the clockwise direction trials. The color bar indicates speed along the zero-force path normalized by its range.

To test this relation quantitatively, the extrema of speed and curvature were determined and the difference in crank angle between speed and curvature extrema  $\Delta$  was computed. Exact temporal coincidence would lead to zero  $\Delta$  values. Figure 2-8 shows representative profiles of the speed and curvature of the zero-force trajectory for a single trial performed by one subject at each of the three speed conditions in the clockwise direction. The circles at the extrema highlight the temporal coincidence of the corresponding extrema in speed and curvature. The data of the same subject are summarized in the histograms for the three speed conditions. The 95% confidence intervals for the  $\Delta$  parameter of all subjects in the 6 conditions were less than 3% of a revolution, indicating that the curvature peaks corresponded to the speed valleys. This coincidence of extrema recapitulates observations of hand motion in unconstrained movements.

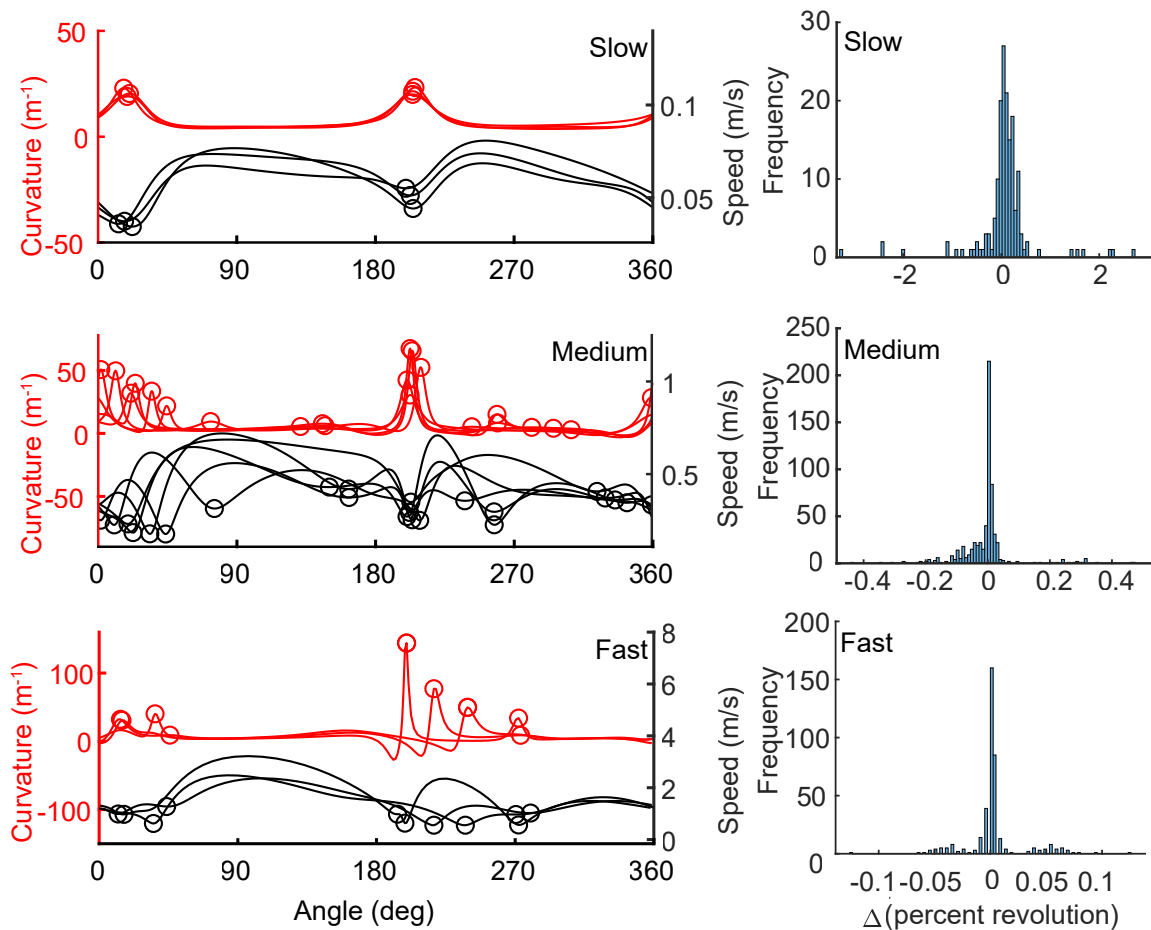


Figure 2-8: Left column: Zero-force trajectory speed (black) and curvature (red) for a single trial performed by one representative subject at each of the three speed conditions in the clockwise direction. The circles highlight the extrema in both profiles. Right column: Histograms of  $\Delta$  for all trials performed by one representative subject at each of the three speed conditions in the clockwise direction. Top row: slow, middle row: medium, bottom row: fast.




Figure 2-9: 95% confidence intervals of the mean of  $\Delta$  for all 10 subjects (S1-S10) in each speed and direction. Confidence intervals are grouped by speed, and the color indicates turning direction [blue: clockwise (CW), and brown: counter-clockwise (CCW)].

### 2.3.4 Sensitivity Analyses of the Zero-force Trajectory and its Effect on the Speed-Curvature Relation

The zero-force trajectory is a construct derived from our experimental observations based on several assumptions combined with parameter values from the published literature. To assess the sensitivity of this construct to the assumptions used to compute it, key parameters of the model were varied over a 3:1 range. A linear time-invariant first-order model of mechanical impedance was assumed, with damping proportional to the assumed stiffness. Values for the stiffness term  $G$  were [0.25, 0.5, 0.75] (slow and medium), and [0.75, 1.50, 2.25] (fast); values of the proportionality constant  $\beta$  were [0.025 s, 0.05 s, 0.075 s] (slow and medium), and [0.05 s, 0.1 s, 0.15 s] (fast). We anticipated that if the impedance was varied, the zero-force trajectory would change, and it did. Nevertheless, the coincidence of the curvature and speed extrema was still observed, even when the zero-force path changed shape. The 95%

confidence interval for  $\Delta$  was always less than 4% of a revolution from zero. Hence, the coincidence between curvature and speed extrema was not sensitive to the particular values of stiffness and damping and the observed results were robust.

## 2.4 Discussion

Long-standing experience in robotics has shown that physical interaction is challenging [Colgate and Hogan, 1989, Colgate and Hogan, 1988, Paul, 1987, Whitney, 1977]. Specifically, contact with a kinematic constraint may compromise stability. The present study examined how humans managed physical interaction by investigating their patterns of motion and force while turning a crank. This action is typical of many activities of daily living that humans perform with ease, with no evidence of the instability that robots typically exhibit. However, it has been proven that properties of interactive dynamics (mechanical impedance) afford a general solution to the stability problem: coupled stability is guaranteed if mechanical impedance mimics that of an energetically passive object [Colgate and Hogan, 1989, Colgate and Hogan, 1988]. Remarkably, human limb impedance appears to have this property [Lee et al., 2014a, Lee et al., 2014b].

Even if stability is assured, if the action is coded in terms of motion control, negotiating a constraint may require detailed knowledge of the constraint and precise musculo-skeletal coordination. Evidence from reaching studies suggests that the control of unconstrained hand motion involves some form of internal model. That model includes kinematics, relating hand position to joint angle or muscle length; and dynamics (of at least the musculo-skeletal system) relating muscle force to hand motion [Kluzik et al., 2008, Lackner and Dizio, 1994, Wolpert et al., 1995, Wolpert et al., 1998]. However, constrained motion introduces a ‘closed chain’ and that changes both the kinematic and dynamic models and makes them much more challenging [Slotine and Asada, 1992].

One proposed robotic approach to constrained motion is ‘hybrid control’: parse the workspace into independent directions and control force when motion is zero,

i.e. normal to a constraint, and control motion when force is zero, i.e. tangent to a constraint [Khatib, 1987, Mason, 1981, Raibert and Craig, 1981]. While consistent with recent work by Chib et al. [Chib et al., 2009], it requires precise knowledge of the constraint in order to identify the normal and the tangent. Such precision in human movement control seems implausible. Instead, humans may exploit the ‘softness’ of neuro-muscular mechanical impedance to reduce the required knowledge and precision. Our experiments were designed to explore the plausibility of such a simplified approach. For that reason, this disarmingly simple task may provide significant insight about how humans manage more sophisticated tools.

### 2.4.1 Patterns of Speed and Force

When subjects turned the crank at their preferred rate, small fluctuations of speed about a mean value were observed. Normal force fluctuated about zero, and both varied systematically with crank angle, consistent with previous work [Ohta et al., 2004, Russell and Hogan, 1989]. The direction of rotation had no significant effect on the means of these measures. When subjects moved at constant speed guided by a visual target, fluctuations of speed and normal force that varied systematically with crank position were again observed. There was a statistically significant but weak interaction between instructed speed and direction for the standard deviations of speed and for mean normal force. However, this interaction merely reflected the sensitivity of the statistical analysis as the magnitude of the directional effect was negligible (Figure 2-6). To test the influence of visual feedback we compared the dependent measures on the catch trials (in which visual feedback was removed) with the immediately preceding trials. While some statistically significant differences were observed, they were negligibly small. Removing visual feedback had no substantial effect.

Might these fluctuations reflect imperfect compensation for musculo-skeletal dynamics? Inertial forces increase with the speed of crank rotation, with a concomitant increase of muscle force, and hence muscle noise. Consistent with imperfect compensation, the standard deviations of speed and normal force both increased with

instructed speed. To test whether the systematic variation of speed and force with crank angle might also be due to imperfect compensation, simulations of the passive inertial dynamics of the arm constrained by the crank, but with no muscle action, were performed. These simulations exhibited systematic variation of both force and speed with crank angle, but with patterns that could not account for our observations. Most notably, at the slowest speed the simulated normal force was essentially zero; that was one reason why such a slow speed (over 13 seconds per revolution) was chosen. The pattern displayed by human subjects was quite different (Figure 2-5).

### 2.4.2 The Zero-Force Trajectory

The zero-force trajectory is a construct based on measured force and motion, combined in a model of peripheral neuro-mechanics. It allows us to ‘peel back’ the peripheral neuro-mechanics to uncover one consequence of the underlying neural commands, and that consequence is expressed in terms of motion. It is similar to, but distinct from the virtual trajectory of the equilibrium-point hypotheses [Bizzi et al., 1982, Bizzi et al., 1984, Feldman, 1966, Feldman, 1986]. That is because the forward-path dynamics between neural input and actual motion is, in general, quite different from the interactive dynamics (mechanical impedance) used to construct the zero-force trajectory. As a result, the zero-force trajectory may differ from the virtual trajectory [Gribble et al., 1998]. More importantly, unlike the equilibrium-point trajectory, we remain agnostic about whether the central nervous system actually encodes the zero-force trajectory. Instead, this quantity is a way to interpret measured force and motion, in combination with a reasonable, albeit simplified, model of peripheral neuro-mechanics.

If the exact neuro-mechanical impedance were known for each subject, this construct would be an exact (i.e. noise-free) measure of the motion consequences of neural commands. However, as we used average measures of neuro-mechanical impedance obtained from different subjects and during static conditions, this construct is at best a noisy and uncertain estimate based on a number of assumptions. It assumed a model of neuro-muscular dynamics that is time-invariant, first-order, and linear.

All of these assumptions are questionable, or even demonstrably incorrect, but they served as workable approximations. The most critical assumptions used to compute the zero-force trajectory were that impedance was linear and time-invariant, with a constant stiffness and damping in joint coordinates. Linearization is valid if the actual (nonlinear) interactive dynamics are differentiable and deviations from the linearization operating point are small. The proximity of the zero-force trajectory to the circular constraint path suggests that the linear approximation was reasonably accurate. The time-invariant parameters were based on measurements made under static postural conditions [Flash and Hogan, 1985, Mussa-Ivaldi et al., 1985]. The slow speed condition was quasi-static – close to postural conditions. The elliptical zero-force path seen in this condition was least sensitive to the assumption of time-invariance. The analysis also assumed that stiffness and damping were symmetric, proportional to each other, connected with a specific topology<sup>1</sup>, and that the same values of stiffness and damping could be applied for all subjects. The assumptions of symmetric stiffness and damping coefficients may be justified as this is sufficient to ensure a stable interaction [Colgate and Hogan, 1988]. Moreover, the study from which we derived the stiffness parameters also reported a symmetric stiffness. The proportionality between stiffness and damping implies first-order interactive dynamics, consistent with experimental observations [Hill, 1938]. However, it also implies that a single time-constant describes the interactive dynamics of all muscles contributing to the overall mechanical impedance.

Given the uncertain and approximate nature of these assumptions, the regularity of the pattern that emerged is striking. Unlike Gomi and Kawato’s [Gomi and Kawato, 1996] findings, our results estimated a zero-force trajectory close to the actual hand path, consistent with the work of Won and Hogan [Won and Hogan, 1995]. In addition, the statistical reliability of our estimates is quite remarkable. It suggests that, from macroscopic behavioral measurements, we can approximate at least some consequences of the underlying neural processes, and that they are informative [Krakauer et al., 2017].

---

<sup>1</sup>In this context, ‘topology’ refers to how the stiffness and damping are connected in the model.



### 2.4.3 Speed Curvature Relation in the Zero-Force Trajectory

In a study of unconstrained two-joint upper-limb reaching Abend et al. [Abend et al., 1982] reported that human hand trajectories exhibited a coincidence between curvature maxima and speed minima-velocity decreased when the curvature increased. When studied in explicitly elliptical trajectories, the speed-curvature relation was described by a two-thirds power law [Lacquaniti et al., 1983], such that the relation between instantaneous angular velocity,  $A(t)$  and the curvature,  $C(t)$  was expressed as  $A(t) = K_C C(t)^{2/3}$ , where  $K_C$  is a constant (the so-called velocity gain factor). The law can also be written in terms of a one-third power law relating the instantaneous tangential velocity,  $V(t)$  and the radius of curvature,  $R(t) = 1/C(t)$ , was expressed as  $V(t) = K_R R(t)^{1/3}$ , where  $K_R$  is a constant. For this reason, it is often referred to as the two-thirds or one-third power law, respectively. This power law was reported for drawing ellipses and several other geometric figures (such as lemniscates); however, the reason for this relation was not clear. Viviani and Terzuolo [Viviani and Terzuolo, 1982] originally suggested it to result from central computational constraints which occur during the translation of movement trajectories into the appropriate motor parameters. However, others showed that this relationship might simply emerge as a result of neuro-mechanics [Gribble and Ostry, 1996, Schaal and Sternad, 2001, Sternad and Schaal, 1999]. Using the lambda-equilibrium-point model, Gribble and Ostry [Gribble and Ostry, 1996] generated an input command which did not follow a curvature-velocity relation. Nevertheless, the simulated motion exhibited a curvature-velocity relation. These authors also tested the alpha-equilibrium-point model as simulated by Flash [Flash, 1987] and observed the same result. In other work, Schaal and Sternad [Schaal and Sternad, 2001, Sternad and Schaal, 1999] showed that a power-law relation between hand velocity and curvature relation emerges from simple oscillatory joint motions in a 7-degree-of-freedom arm model. More importantly, this account predicted systematic deviations from a power-law relation for large motions, which was confirmed by experimental observation.

Other observations across multiple tasks and extremities, as well as in the percep-

tion of actions, provided evidence that the power-law relation is of non-mechanical, i.e. neural origin [Dayan et al., 2007, Hicheur et al., 2005, Maurice et al., 2018a, Viviani and Flash, 1995]. For example, Massey et al. [Massey et al., 1992] performed an experiment where subjects grasped a 3D isometric force-sensing handle and exerted forces continuously to draw circles, ellipses, and lemniscates on a screen, with and without visual feedback. The power-law curvature-velocity relation was observed even when the hand did not move, indicating that inertial dynamics alone could not be its cause. Nevertheless, a later Gribble and Ostry study [Gribble and Ostry, 1996] showed that biomechanics may account for a curvature-speed relation even during isometric conditions. Other research has shown that movement production involves the cooperative interaction of large neuronal populations [Georgopoulos et al., 1986]. Schwartz [Schwartz, 1994] used a population vector method to interpret one aspect of neural activity as representing spatial motions. This information was then used to visualize motor cortical representations of wrist trajectories which were remarkably similar to the actual wrist motions made by monkeys as they drew spirals. Surprisingly, a power law relating velocity and curvature was observed in this cortically-derived representation of motion. Neuro-mechanics alone cannot account for these observations.

A possible resolution of this controversy is that the neural controller and the neuro-mechanical periphery co-evolved to embody similar dynamic structures and constraints. Long communication delays are a core challenge of human motor control. Coping with these delays requires a controller that favors predictability, even over energetic cost in some cases [Bazzi et al., 2018, Koeppen et al., 2017, Maurice et al., 2018b, Sternad and Hasson, 2016]. Smoothness quantified by minimizing mean squared jerk provides a measure of predictability and maximizing smoothness has been shown to account for the coordination of simple reaching movements, including an account for the two-thirds power law [Flash and Hogan, 1985, Schaal and Sternad, 2001, Sternad and Schaal, 1999, Viviani and Flash, 1995]. More recent work by Huh showed that maximizing smoothness along a curved path yields a spectrum of power-law speed-curvature relations [Huh and Sejnowski, 2015].

Periodic actions are (in principle) infinitely predictable and have been proposed as one class of dynamic primitives used to construct motor behavior [Hogan, 2017, Hogan and Sternad, 2012, Hogan and Sternad, 2013, Ronsse et al., 2009, de Rugy and Sternad, 2003, Schaal and Sternad, 1998, Sternad, 2008, Sternad et al., 2000]. An elliptical path may be generated by combining two sinusoids of the same frequency at a non-zero relative phase, though possibly of different amplitudes. The resulting trajectory (i.e. time-course of speed along the elliptical path) necessarily exhibits a power-law relation between velocity and curvature. Generating an elliptical zero-force trajectory via out-of-phase sinusoids is consistent with our observations. The same model (a zero-force trajectory composed of simple dynamic primitive oscillations) may account for the observations of Massey et al. [Massey et al., 1992] and possibly those of Schwartz [Schwartz, 1994]. Perhaps more important, it is also consistent with prioritizing predictability over other aspects of performance, including the variability of actual hand speed and exerted normal force.

## 2.5 Conclusion

Unconstrained curved movements exhibit a systematic relation between speed and curvature. This has been attributed to either neural or mechanical causes, without resolution of the debate. To obtain a new perspective on this open question, this study examined movements against a simple kinematic constraint. By design, the experiment confined the hand to a circular path with constant curvature; speed was controlled by instruction with visual feedback. If control is focused on hand trajectories, there should be no reason for a systematic variation of speed as curvature was constant. Nevertheless, systematic fluctuations of speed and normal force were observed at all speeds, unaffected by the presence or absence of visual feedback. When the influence of peripheral neuro-mechanics was subtracted based on a biomechanical model to identify a zero-force trajectory, the widely-observed relation between speed and curvature extrema re-emerged: Extrema of curvature and tangential speed were systematically time-aligned. This relation persisted even when the stiffness and

damping parameters of the model were varied over a wide range. These findings provide evidence that the speed-curvature relation is due, at least in part, to neural processes, possibly reflecting their co-evolution with peripheral dynamic structures. These results are also a first demonstration that the zero-force trajectory may successfully dis-entangle neural and biomechanical factors underlying motor actions.

## Chapter 3

# Dynamic Primitives in Constrained Action: Systematic Changes in the Zero-Force Trajectory

This is an adapted version of a paper to be submitted to the Journal of Neurophysiology. Furthermore, a conference paper presented at BioRob 2020 [Hermus et al., 2020] covers a portion of the content presented herein.

### 3.1 Introduction

Previously, a method which estimated the zero-force trajectory was proposed. It was applied to the task of turning a planar crank. A model of limb impedance was assumed, and the zero-force trajectory was computed. The zero-force trajectory showed evidence of the speed-curvature relation reported in unconstrained motion [Hermus et al., 2020] suggesting that our approach reveals information about neural control. However, structure beyond that of the speed-curvature relation was also evident in those data. The work reported here presents further investigation of the zero-force trajectory in the task of crank turning.

### 3.1.1 Motivation for the Experiments

If dynamic primitives control physical interaction, quantifiable limitations may be evident in human performance, from which dynamic primitives may be inferred. Prior work showed that humans adjust their behavior to prioritize predictability [Bazzi et al., 2018, Maurice et al., 2018b, Nasserolelami et al., 2014, Nayeem et al., 2021, Sternad, 2017]. Arguably, the most predictable action is periodic; in theory, strictly periodic actions are infinitely predictable. Negotiating a circular constraint at constant speed requires periodic hand motion in each degree of freedom. For these reasons, we anticipated that crank-turning might preferentially be executed as a combination of oscillatory actions.

Two non-collinear sinusoids of the same period but different amplitude and phase describe an ellipse. This system can produce a subset of the so-called Lissajous plots. Our previous study on circular crank turning revealed zero-force trajectories with a roughly elliptical shape [Hermus et al., 2020]. Constant-speed circular hand motion requires sinusoidal motion in orthogonal directions with a phase offset of  $\pm 90^\circ$  depending on the direction of motion. However, the motions evoked by neural oscillations would lag those neural oscillations to an extent determined by the dynamic behavior of the neuro-mechanical periphery. This motivated **Hypothesis 1**: The zero-force trajectory will describe a roughly-elliptical path with different orientations for CW and CCW rotations.

Despite the predictability of the actions required to turn a crank, imperfect execution may be anticipated due to sensory noise, motor noise, and/or inadequate prediction of inertial dynamics and neuro-muscular response. The consequences of these imperfections should decline (precipitously) with decreasing speed: all inertial forces decline with the square of speed; velocity-dependent muscle dynamics decline in proportion to speed; motor noise declines in proportion to speed; and slower speed allows ample time for feedback corrections. However, rhythmic actions with periods longer than 2 to 5 seconds are neither perceived nor executed as periodic (Fraisse 1984; James 1890). Previous research has shown that slower movements, even if pe-

riodic, ‘break down’ into a sequence of stereotyped submovements [Park et al., 2017], another class of dynamic primitives. Consequently, despite the dynamic advantages of moving slowly (reduced influence of dynamic complexities, more time for feedback corrections), we anticipated that performance would be compromised during slow movements. This motivated **Hypothesis 2**: Variability of hand speed will increase in the slowest movements.

## 3.2 Methods

The participants, experimental apparatus, procedure, and computation of the zero-force trajectory were the same as described in Chapter 2 [Hermus et al., 2020].

### 3.2.1 Dependent Measures

#### Orientation of Zero-force Trajectory

In our prior work, the zero-force trajectory could be approximated by an ellipse. To quantify the orientation of an ellipse, the choice of coordinates required care. Two potential candidates are shown in Figure 3-1. Orientation and eccentricity were intuitively understandable, but they were not always numerically well-behaved. The major axis angle was not defined for an eccentricity of unity and, for eccentricities near unity, the major axis angle was expected to have a nearly flat distribution. To address this concern we used log-ratio coordinates. Though less intuitive (see Figure 3-2), they have numerical properties better suited to statistical analysis [Fasse et al., 2000, Fasse and Hogan, 1996].




Figure 3-1: Two possible sets of coordinates: (left) Ellipse eccentricity and major axis angle; (right) log-ratio coordinates. The major axis is poorly defined when eccentricity approaches unity. Log-ratio coordinates alleviate this problem.




Figure 3-2: Graphical illustration of how the log-ratio coordinates change as the ellipse varied.

For each subject in each speed and direction condition the zero-force trajectory was binned into 200 angular position bins. The average radial position for each bin was computed. This polar representation was used to estimate the lengths  $l_x, l_y, l'_x, l'_y$  at zero degrees and 45 degrees, respectively. For each subject and each condition



$\ln(r_0)$  and  $\ln(r_{45})$  were estimated. To decrease the dimensionality of the statistical analysis, principal component analysis (PCA) was performed to identify coordinates which explained the most variance (see Figure 3-6). In order to test Hypothesis 1, this principal component,  $\ln(r)_{PC}$ , served as a dependent measure for statistical analysis.

### Variability of Hand Speed

To quantify the variability of hand speed the crank velocity data from all trials within a speed and direction condition were binned into 200 position bins. In each bin the coefficient of variation (CV) was computed (standard deviation divided by the mean). Then the average of this CV across bins served as the second dependent measure for statistical analysis.

### 3.2.2 Statistical Analysis

To quantify the influence of speed and direction, a linear mixed model was employed; it was then tested using analysis of variance (ANOVA). The linear model which represented the observed dependent measure  $Y_{i,j,k}$  was expressed as

$$Y_{i,j,k} = \mu_T + \alpha_j + \beta_k + \gamma_l + (\alpha\beta)_{j,k} + (\alpha\gamma)_{j,l} + (\beta\gamma)_{k,l} + (\alpha\beta\gamma)_{j,k,l} + E_{i,(j,k,l)} \quad (3.1)$$

where the grand mean is  $\mu_T$ , the fixed effect of speed is  $\alpha_j$ , where  $j$  is an index from 1 to 3, the fixed effect of direction is  $\beta_k$ , where  $k$  is an index from 1 to 2, the random effect of subject is  $\gamma_l$ , where  $l$  is an index from 1 to 10. For the first dependent measure,  $\ln(r)_{PC}$ , the stochastic sampling effect is  $E_{i,(j,k,l)}$ , where  $i$  is an index from 1 to 22 (representing the number of trials excluding the first and the 7 catch trials). For the second dependent measure, the CV of speed, the index  $i$  was 1. This was because computing the CV over bins ignored individual trials. The significance level was set to 5% for all statistical tests. Paired-sample t-tests were carried out to interpret the results of the ANOVA.

### 3.2.3 Sensitivity to Impedance Assumptions

To the best of our knowledge, two-joint limb impedance measurements during physical interaction with a constraint have not been reported. Thus, impedance parameters were based on unconstrained static arm stiffness and single joint damping measurements. Consequently, the impedance parameters, the gain term for the stiffness,  $G$ , and the proportional damping term,  $\beta$ , were varied to test whether the results were sensitive to the impedance values used to compute the zero-force trajectory. When the gain and damping terms changed, the zero-force trajectory was expected to change. However, the main question was whether the dependent measures were robust to the changes in the stiffness and damping. To this end, the  $\ln(r)_{PC}$  was computed when the  $G$  and  $\beta$  terms were each varied over a 3:1 range.

## 3.3 Results

### 3.3.1 Variation of Force and Motion

Despite instructions and the availability of visual feedback, speed fluctuated, displaying a systematic pattern with respect to crank angle. At the same time, non-zero force normal to the constraint was observed, again displaying a systematic pattern with respect to crank angle. To explore the possible cause of these patterns, we combined these observations by computing zero-force trajectories (see Figure 3-3).

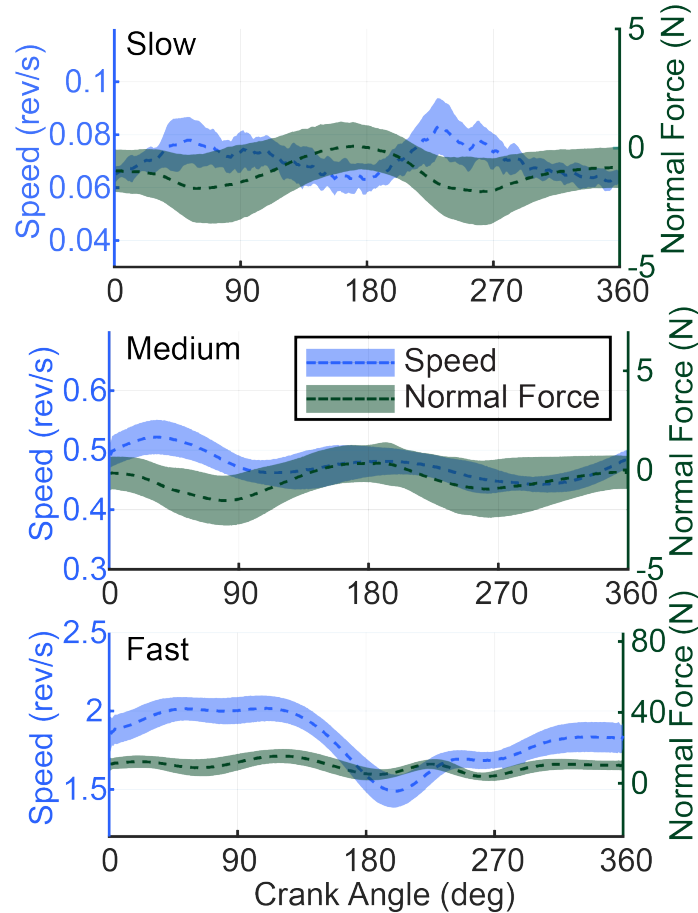


Figure 3-3: Mean and SDs of speed (blue dashed line) and normal force (green dashed line) versus crank angle for all subjects. The shading indicates 1 SD from the mean. Data are from the clockwise direction trials. Systematic fluctuations in speed and force were observed with respect to crank position. Note the different scales on the ordinates.

### 3.3.2 Zero-Force Trajectory Orientation

A zero-force trajectory from one representative subject in each direction and speed condition is presented in Figure 3-4. To provide a sense of the data across all subjects, the average ZFT for each subject binned by position, in each direction and speed condition, is presented in Figure 3-5. From these two figures it is clear that (1) the path shapes were approximately elliptical; (2) there was a consistent pattern of speed fluctuations with position along the path; and (3) the shapes displayed a clear difference of orientation between the two directions.




Figure 3-4: One representative trial from one subject in each of the slow, medium, and fast speed conditions. Left: slow speed; middle: medium speed; right: fast speed. Top: clockwise direction trials; bottom: counter-clockwise direction trials. The path defined by the constraint is shown by the black dashed circle. The zero-force trajectories are shown by lines with varying color that indicates speed along the zero-force path (normalized by its range). Importantly, the zero-force trajectory is roughly elliptical and its orientation differs with direction.




Figure 3-5: The average ZFT binned over angular position for each of the 10 subjects in the, slow, medium, and fast speed conditions. Left: slow speed; middle: medium speed; right: fast speed. Top: clockwise direction trials; bottom: counter-clockwise direction trials. The path defined by the constraint is shown by the black dashed circle. The zero-force trajectories are shown by lines with varying color that indicates speed along the zero-force path (normalized by its range).

The ellipse parameters in log-ratio coordinates are presented in Figure 3-6. Despite the variability of the data, the difference between CW and CCW directions is visually evident. Statistically, the mean  $\ln(r)_{PC}$  in the CW conditions were  $0.098 \pm 0.12$  (slow),  $0.23 \pm 0.15$  (medium) and  $0.58 \pm 0.07$  (fast). The mean  $\ln(r)_{PC}$  in the CCW conditions were  $-0.21 \pm 0.16$  (slow),  $-0.27 \pm 0.19$  (medium) and  $-0.45 \pm 0.09$  (fast).

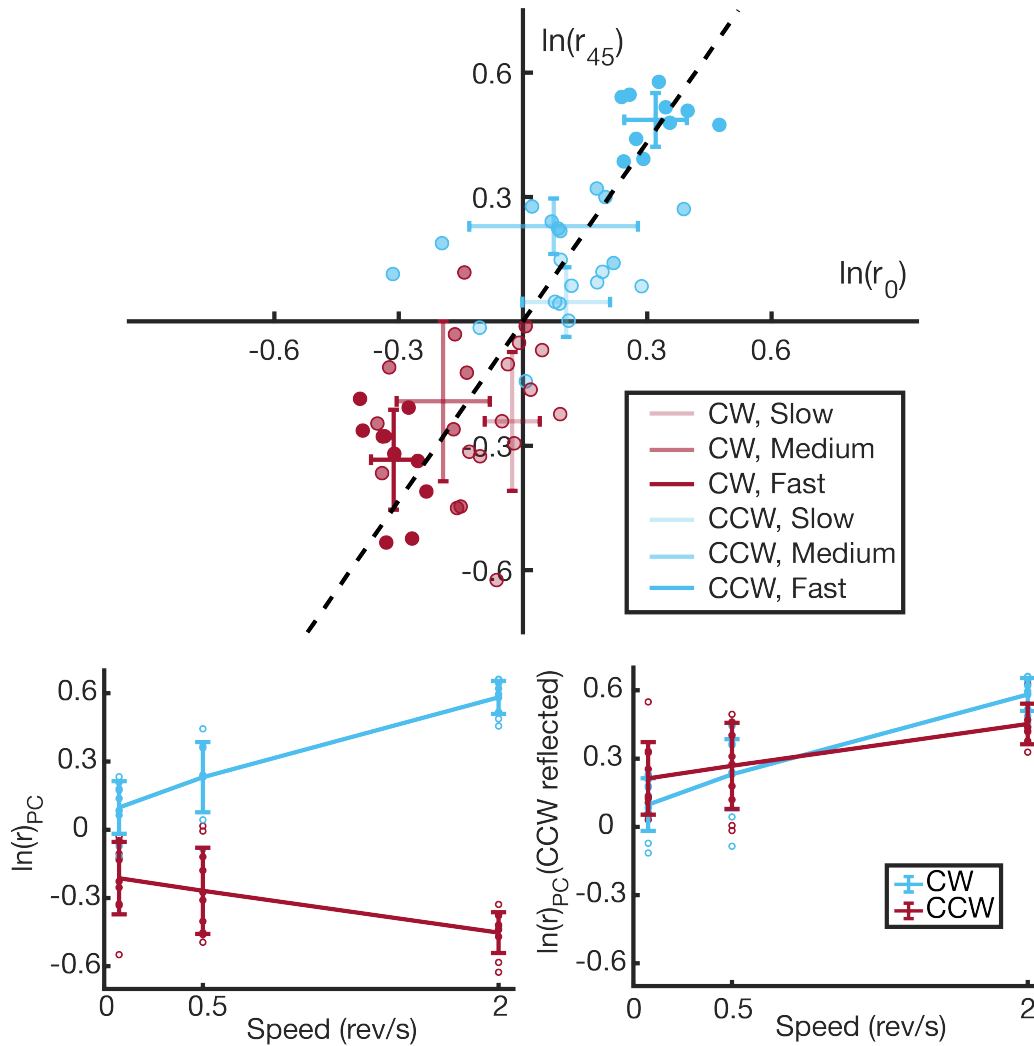


Figure 3-6: Distribution of ellipse parameters in log-ratio space. The black dashed line represents the direction of the first principal component. The CW direction is represented by blue, and the CCW direction is represented by red. Lighter shades indicate slower speeds and darker shades indicate faster speeds. The crosses represent plus or minus one standard deviation from the mean across subjects. (bottom left) Distribution of dependent measure  $\ln(r)_{PC}$ . Note the clear difference between the CW and CCW directions. (bottom right) Distribution of the dependent measure  $\ln(r)_{PC}$  with CCW data reflected.

A significant interaction between speed and direction ( $F_{2,0,18,0} = 106.388$ ,  $P \ll 0.001$ ) together with a main effect of speed ( $F_{2,0,9,0} = 6.062$ ,  $p = 0.006$ ), and direction ( $F_{1,0,9,0} = 212.879$ ,  $P \ll 0.001$ ), were detected. For each speed post-hoc pairwise t-tests were run between the CW and CCW directions. A significant difference between directions was observed in all three cases: slow ( $P \ll 0.001$ ), medium ( $P \ll 0.001$ ),

and fast ( $P \ll 0.001$ ).

To clarify the effect of direction, the CCW data were reflected along the  $\ln(r)_{PC}$ . The result is displayed in the lower right panel of Figure 3-6. For comparison, the original (un-reflected) data are displayed in the lower left panel. After reflection, a significant main effect of speed ( $F_{2,0,9,0} = 106.388$ ,  $P \ll 0.001$ ), and a significant interaction between speed and direction ( $F_{2,0,18,0} = 6.0622$ ,  $P = 0.006$ ) were still detected. However, this action, reflecting the CCW about the  $\ln(r_0)$  and  $\ln(r_{45})$  axes, removed the main effect of direction.

### 3.3.3 Orientation – Sensitivity Analysis

The zero-force trajectory is a construct derived from our experimental observations based on several assumptions combined with parameter values from the published literature. To assess the sensitivity of this construct to the assumptions used to compute it, key parameters of the model were varied over a 3:1 range. A linear time-invariant first-order model of mechanical impedance was assumed, with damping proportional to the assumed stiffness. Values for the gain term  $G$  were [0.25, 0.50, 0.75] (slow and medium), and [0.75, 1.50, 2.25] (fast); values of the proportionality constant  $\beta$  were [0.025 s, 0.05 s, 0.075 s] (slow and medium), and [0.05 s, 0.1 s, 0.15 s] (fast). We anticipated that, if the impedance was varied, the zero-force trajectory would change, and it did. Nevertheless, the differences in ellipse orientation between conditions remained when the impedance parameters were varied (see Figure 3-7). Hence, the zero-force trajectory orientation was not sensitive to the particular values of stiffness and damping; the observed results were robust.




Figure 3-7: Distribution of dependent measure  $\ln(r)_{PC}$  when the impedance parameters were varied over a 3:1 range. Error bars indicate the standard deviation (SD) between subjects. Lighter shades indicate lower stiffness and damping and darker shades indicate higher stiffness and damping.

### 3.3.4 Hand Speed Coefficient of Variation

The CV of hand speed is presented in Figure 3-8. There was a substantial increase in hand speed CV at slow speeds. Statistically, the mean CV in the CW conditions were  $0.26 \pm 0.06$  (slow),  $0.13 \pm 0.02$  (medium) and  $0.11 \pm 0.03$  (fast). The mean CV in the CCW conditions were  $0.28 \pm 0.06$  (slow),  $0.16 \pm 0.02$  (medium) and  $0.11 \pm 0.02$  (fast).

A significant main effect of speed ( $F_{2,0,9,0} = 116.076$ ,  $P \ll 0.001$ ), direction ( $F_{1,0,9,0} = 5.755$ ,  $P = 0.014$ ), and an interaction between speed and direction ( $F_{2,0,18,0} = 3.489$ ,  $P = 0.038$ ) were detected. Note that, consistent with visual inspection of Figure 7, the effect of direction and the interaction, though significant, were weaker than the influence of speed. For each direction, post-hoc pairwise t-tests were run between adjacent speeds. A significant difference was detected between the slow and medium speed for CW ( $P \ll 0.001$ ) and CCW ( $P \ll 0.001$ ) motion. A significant difference was also detected between medium and fast speeds CW ( $P = 0.003$ ) and CCW



( $P \ll 0.001$ ).

A significant interaction between speed and direction ( $F_{2.0,18.0} = 3.489$ ,  $P = 0.038$ ) together with a main effect of speed ( $F_{2.0,9.0} = 116.076$ ,  $P \ll 0.001$ ), and direction ( $F_{1.0,9.0} = 5.755$ ,  $P = 0.014$ ), were detected. Note that, consistent with visual inspection of Figure 3-2, the effect of direction and the interaction, though significant, were weaker than the influence of speed. For each direction, post-hoc pairwise t-tests were run between adjacent speeds. A significant difference was detected between the slow and medium speed for both CW ( $P \ll 0.001$ ) and CCW ( $P \ll 0.001$ ) directions. A significant difference was also detected between medium and fast speeds in the CW ( $P = 0.003$ ) and CCW directions ( $P \ll 0.001$ ).




Figure 3-8: Coefficient of variation (CV) of hand speed. Error bars indicate the standard deviations between subjects. Blue lines indicate the clockwise trials (CW), red lines indicate the counter-clockwise trials (CCW).

### 3.4 Discussion

Kinematically constrained motion presents an intermediate step between the widely-studied unconstrained motions, such as reaching and pointing, and the sparsely-studied physical interaction with dynamically complex objects. Several works have

investigated crank turning from a variety of perspectives including early descriptive studies (Russell and Hogan 1989), optimal control studies with muscle level modeling [Davoudabadi Farahani et al., 2016, Ohta et al., 1998, Ohta et al., 2004, Svinin et al., 2001, Zheng and Wang, 2017], studies that reported minimizing muscular effort [Koeppen et al., 2017], studies that summarized muscle behavior as an impedance [Hermus et al., 2020, Hermus et al., 2020], as well as studies in robotics [Kazanzides et al., 1989, Williamson, 2003].

To understand the interaction with a kinematic constraint our approach was to estimate an underlying zero-force trajectory. The zero-force trajectory is a construct based on measured force and motion, combined in a model of peripheral neuro-mechanics. It allows to ‘peel back’ the peripheral neuro-mechanics to uncover one consequence of the underlying neural commands; that consequence is expressed in terms of motion. The zero-force trajectory is similar to, but distinct from, the virtual trajectory of the equilibrium-point hypothesis [Bizzi et al., 1982, Feldman, 1966, Feldman, 1986]. The virtual trajectory was postulated to be encoded in neural commands descending from the higher central nervous system to the periphery. The forward-path dynamics between neural input and actual motion is, in general, quite different from the interactive dynamics (mechanical impedance) used to construct the zero-force trajectory. Because of neural transmission delays, excitation-contraction coupling and other dynamic effects in the forward path, the zero-force trajectory may differ substantially from the virtual trajectory [Gribble et al., 1998]. The zero-force trajectory is an objective observation defined by measurements of hand force and motion combined with an estimate of interactive dynamics. While it is a consequence of neural activity, it remains unknown whether this quantity is represented in higher-level activity of the central nervous system. An unconstrained trajectory is also a consequence of neural activity, but an observed trajectory is not sufficient to determine how neural activity controls or encodes this action. Nevertheless, observations of actual motion have proven to be informative. The zero-force trajectory is a way to interpret measured force and motion, in combination with a reasonable, albeit simplified, model of peripheral neuro-mechanics. As with actual unconstrained trajectories, observations

of the zero-force trajectory are informative.

### 3.4.1 Zero-force Trajectory Orientation

This study revealed a statistically significant difference in orientation of the zero-force trajectory between the CW and the CCW conditions – consistent with Hypothesis 1. We anticipated that turning a crank might be executed using oscillatory primitives. Constant-speed circular hand motion requires sinusoidal motion in orthogonal directions with a phase offset of  $\pm 90^\circ$  (depending on direction, CW vs. CCW). However, the oscillatory zero-force trajectory required to produce this motion would lead hand motion by an extent determined by the slow dynamic response of the neuro-mechanical periphery interacting with the inertia of the skeleton. That lead in time would manifest as a lead in phase that differed in opposite turning directions, and that would result in different performance in CW and CCW crank turning, just as we observed.

Consider a zero-force trajectory with two orthogonal components,  $x_0$  and  $y_0$ , constructed from two out-of-phase sinusoids with the same frequency,  $\Omega$ , same magnitude, and a phase difference,  $\phi$ .

$$\begin{cases} x_0 = \sin(\Omega t) \\ y_0 = \sin(\Omega t + \phi) \end{cases} \quad (3.2)$$

A perfect circle can be drawn in the CW or CCW direction with a phase difference of  $\pm 90^\circ$  (see Figure 3-9). However, slow peripheral neuro-muscular dynamics interacting with skeletal inertia would require the zero-force trajectory to move ahead of the actual trajectory, contributing an additional phase difference. Moreover, because of the anisotropy of skeletal inertia and neuro-muscular impedance, the magnitudes of the two sinusoidal components of the zero-force trajectory would differ. This would result in zero-force trajectories with an elliptical shape that is oriented differently for CW and CCW motion—just as we observed. The medium and fast speed behavior strongly supports control using oscillatory primitives.

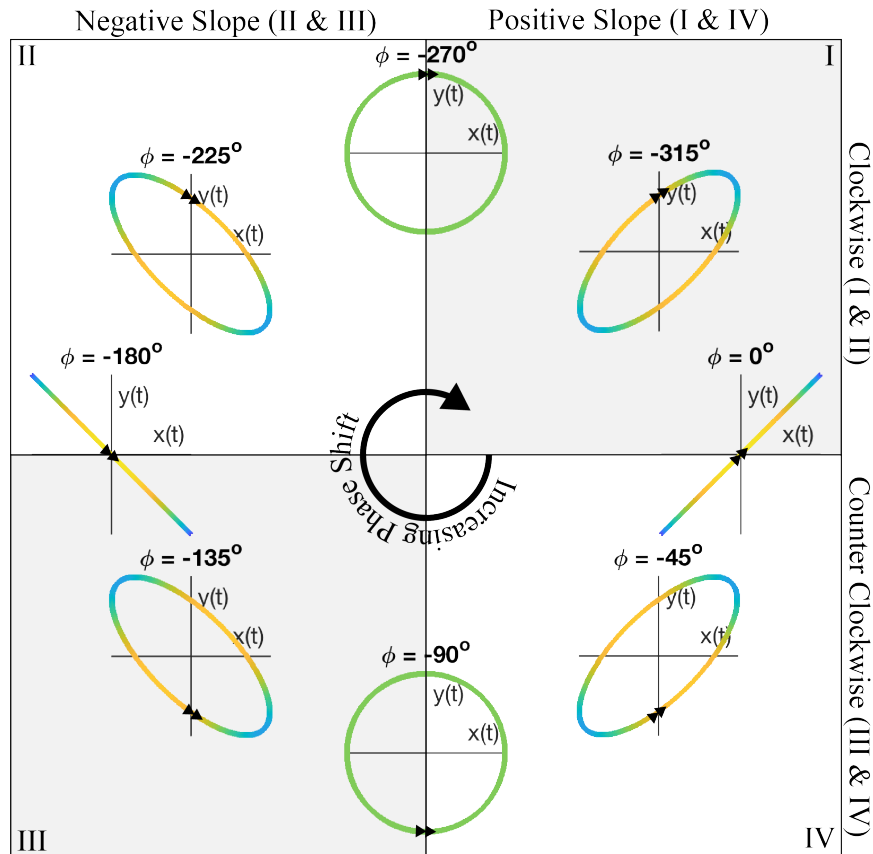


Figure 3-9: Trajectories resulting from out of phase sinusoids. To draw a perfect circle  $\pm 90^\circ$  of phase is required depending on direction CW or CCW. However, the limb dynamics add phase, thus zero-force trajectories in the III and II quadrants are observed for the CW and CCW directions respectively. \*This figure is replicated from [Al-Khazali and Askari, 2012].

### 3.4.2 Variability of Hand Speed

In our experiment, the coefficient of variation of hand speed was largely similar between the fast and medium speeds (both with periods no longer than 2 sec), but increased dramatically at the slowest speed (period of 13.3 sec) – consistent with Hypothesis 2 and providing evidence of primitive submovements.

For more than 100 years researchers have documented movement intermittency, the inability to move smoothly and continuously, during slow cyclical movements [Crossman and Goodeve, 1983, Doeringer and Hogan, 1998b]; saccadic eye movements [Collewijn et al., 1988]; and slow discrete movements [Hogan et al., 1999, Park

et al., 2017, Vallbo and Wessberg, 1993, van der Wel et al., 2009, Woodworth, 1899]. In addition, studies with stroke patients during rehabilitation have documented ‘fragmented’ movements composed of highly stereotyped submovements [Krebs et al., 1999, Rohrer et al., 2004]. This is vastly different from normal human reaching movements, which are usually quite smooth and follow a minimum-jerk trajectory (Flash and Hogan 1985).

An oscillation could, in theory, be a composite of opposite-direction submovements. However, physiological evidence indicates that this is not the case [Brown and Sherrington, 1911, Brown, 1914, Grillner and Wallen, 1985, Schaal et al., 2004]. A well-documented observation in human psychology and motor control is that when the period of a rhythmic action is longer than 2 to 5 seconds it can no longer be perceived nor executed as periodic [Fraisse, 1984, James, 1890]. Slower movements, even if periodic, ‘break down’ into a sequence of stereotyped submovements (Park et al. 2017). These submovements may have limitations, e.g. amplitude, duration, and a minimum ‘refractory period’ – time between initiation of adjacent submovements. Production of constant speed motion, with intermittent control, would require submovements to overlap [Loram et al., 2014, Markkula et al., 2018]. As movement slows, the refractory period and minimum submovement duration make constant motion impossible and individual submovement peaks appear. This causes greater variability at slow speeds – a quantifiable limitation of human behavior that may account for our observations.

### **3.4.3 Limitations of this Study**

The zero-force trajectory is an estimate based on a large number of assumptions. It assumes a model of neuro-muscular dynamics that is (1) time-invariant; (2) first-order; and (3) linear. All of these assumptions are demonstrably incorrect, but they serve as a first workable approximation. The analysis also assumes that stiffness and damping are (1) connected with a specific topology; (2) symmetric; (3) proportional and (4) that the same values of stiffness and damping may be applied to all subjects. Given the highly approximate nature of these assumptions, the regularity of the pattern

that emerged is striking. Unlike [Gomi and Kawato, 1996] showed for a reaching movement, the path of the zero-force trajectory that emerged in this study is ‘close’ to the actual hand path. In particular, the effect of turning direction was observed even when the parameters scaling the magnitude  $G$  and relative size  $\beta$  of stiffness and damping were each varied over a 3:1 range.

While our observations are consistent with control based on a composition of underlying oscillations and submovements, we cannot rule out alternative explanations. Moreover, while we refer to these as dynamic primitives, we provide no direct evidence of their dynamic nature. However, a robust oscillation in the presence of sensory and motor noise implies that it is a stable attractor of a nonlinear dynamic system (a limit cycle); and the stereotypy reported for submovement speed profiles suggests that they are transient responses of dynamic systems, also nonlinear as submovements have finite duration. A zero-force trajectory composed of oscillations and/or submovements provides a parsimonious, even elegant, account of our observations—substantially nonzero normal force even though it contributes no mechanical work; and patterns of force and velocity error that are clear with respect to configuration but not time.

### 3.4.4 Implications

While our results support a role for dynamic primitives in the production of kinematically constrained motions, they also present a puzzle. Slow speed behavior cannot be executed by oscillatory dynamic primitives, yet it exhibited a feature easily explained by oscillatory primitives at medium and fast speeds—direction dependence. The slowest speed, 13.33 seconds per revolution was quasi-static by any reasonable definition of the term [Hermus et al., 2020]. In that case, all inertial effects were negligible, as were any delays in the neuro-muscular dynamics. What might account for the observed direction dependence at the slowest speed? One possibility is that the CNS develops some form of geometric representation of the behavior and uses it to generate a nominal slow-speed zero-force trajectory that serves as a target for feedback correction. Combined with an error threshold below which no action is taken, this may result in intermittent control, manifest as a sequence of submove-

ments and resulting in an elevated coefficient of variation of hand speed [Markkula et al., 2018, Park et al., 2017]. The validity of these speculations is a topic for future study.

### 3.5 Conclusions

This study examined kinematically constrained motion as an intermediate step to bridge the gap between (widely-studied) unconstrained motions and (sparsely-studied) physical interaction with dynamically complex objects. We investigated the detailed patterns of motion and force that human subjects exhibited when performing a simple constrained-motion task, turning a circular crank. Turning the crank in both CW and CCW directions and at slow, medium, and fast speeds exposed several ‘artifacts’ that could not result from mechanics alone. We speculate that they may arise from underlying dynamic primitives.

The zero-force trajectory displayed clear differences in orientation when turning in different directions. This is consistent with control using oscillatory primitives to generate an elliptical zero-force trajectory. The difference in orientation with direction was observed even at the slowest speed. However, when the oscillation period became longer than 2 to 5 seconds, motion can no longer be perceived or executed as periodic. Instead, it “breaks down” into a sequence of submovements, and that results in increased variability, which we observed at the slowest speeds. This reinforces previous evidence [Park et al., 2017] of a transition from smoothly rhythmic to intermittent control as actions slow.

Using dynamic primitives may allow humans to ‘work around’ the shortcomings of their slow muscles and neural communication to perform complex physical interaction tasks. However, the advantages of this approach imply concomitant disadvantages. An elliptical zero-force trajectory that does not coincide with the circular constraint results in non-zero normal force applied to the crank [Hermus et al., 2020]. This was observed even at the slowest speed when dynamic effects were negligible and there was ample opportunity for feedback correction. This reinforces recent reports that human

subjects interacting with a robot cannot suppress a dependence of interaction force on motion, even with visual feedback and opportunity to practice and learn [Maurice et al., 2018a, West et al., 2022]. These limitations are important to understand the interaction between humans and the devices we develop, from complex robotic systems to simple hand-held tools.



# Chapter 4

## Estimating Impedance during Constrained Motion – The Configuration-Dependent Ensemble

In this Chapter, we set out to measure mechanical impedance during the task of crank-turning. Initially, time-based ensemble methods were employed. However, a key assumption—the stationarity of the noise processes—was not satisfied. A ‘work-around’ was developed: (1) high-pass filter to remove the influence of the underlying time-varying ZFT; (2) identify a configuration-dependent ensemble. The method was tested in three ways: (1) simulation; (2) using a robot but simulating a model of human behavior; (3) using a single human subject. Results demonstrated an opportunity for further improvement but also uncovered a fundamental limitation that may preclude usable estimation in some circumstances.

### 4.1 Introduction

In previous Chapters 2 and 3, a linear first-order time-invariant model of limb impedance was assumed based on static postural measurements [Mussa-Ivaldi et al., 1985]. In practice, however, impedance varies as a function of many factors including: muscle activation [Cannon and Zahalak, 1982], movement [Bennett et al., 1992], activity

preparation [Lacquaniti et al., 1993], force exertion level [Perreault et al., 2001], task stabilization [Burdet et al., 2001], limb configuration [Trumbower et al., 2009], and walking gait state [Lee et al., 2016, Lee and Hogan, 2015]. Direct measurement of impedance may provide insight which could clarify or falsify portions of the proposed dynamic primitive hypothesis. This led to the following research question: is it possible to measure impedance directly in the crank-turning task?

An experimental paradigm was developed where a human subject turned a virtual crank simulated by a robot. The goal of this experiment was to perturb human behavior, measure the response, and estimate the interactive dynamics of the limb. When estimating limb dynamics, perturbation-based system identification methods fall into two categories: (1) Estimate impedance directly – apply position displacements and measure the resulting forces; (2) Estimate admittance directly – apply force perturbations and measure the resulting position. After estimating admittance directly, invert admittance to estimate impedance. There is a vast literature which has applied both methods to estimate impedance in a variety of conditions – a majority of which has studied a static case or a single joint. A review of the literature prior to 1990 was presented by Kearney and Hunter [Kearney and Hunter, 1990]. While perturbation-free approaches exist [Hodgson and Hogan, 2000, Lee et al., 2021] perturbation-free methods require assumptions about the forward-path or interactive dynamics which are difficult to justify in a crank-turning paradigm.

One of the primary reasons for studying human limb impedance is to better understand the awesome feats of stable manipulation humans achieve every day. Activities of daily living frequently involve significant physical interaction – simultaneous motion and force production. Considering that over 30 years has passed since Kearney and Hunter’s review, surprisingly little research has been performed which estimates limb impedance during physical interaction. In order to investigate one such task, we used a robot to simulate a compliant constraint. An impedance controller rendered the constraint while force perturbations were super imposed – consistent with case 2.

Originally, an ensemble method [MacNeil et al., 1992, Lortie and Kearney, 2001, Perreault et al., 1999] was applied to estimate the time-varying impedance of the limb.

This approach worked well when tested in simulation. However, it failed when applied to human crank-turning data. It appeared that important aspects of the human data were missing from the simulations used to generate the testing data. Time-varying ensemble methods assume that a time-dependent impulse response function relates the inputs to the outputs. However, in the previous experiments, presented in Chapter 2 and 3, humans did not turn the crank with a constant period, even when instructed to move at a constant speed and provided with velocity feedback (see Figure 4-1 and 4-5).

Instead, substantial position drift was observed, even when velocity feedback was provided. Experimentally, variations in crank position after one period (crank revolution) could be as large as plus or minus  $1/3^{rd}$  of a crank revolution (See Figure 4-1). Variation in crank position on this scale is problematic. Geometric changes in limb configuration would result in considerable error in impedance estimation simply due to the change in inertia.

This drift does not appear to be an artifact, but is a feature of the human motor controller. As a ‘sanity check’ a separate experiment was performed to determine if drift was present in a static posture task. The drift was observed during static posture as well, even when visual position feedback was provided (See Appendix C).

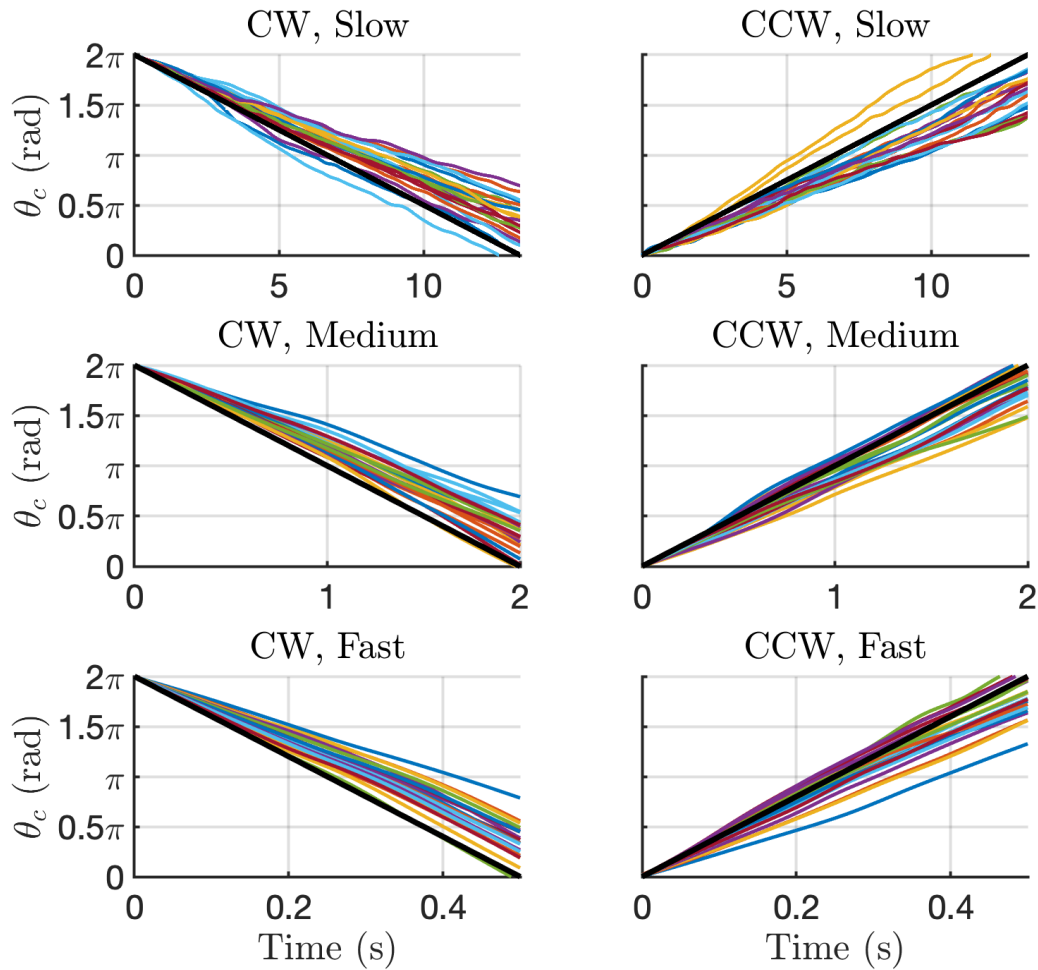


Figure 4-1: A representative subject’s crank position with respect to time. The black line represents the target speed. The varying color lines represent different trials. The left column corresponds to clockwise rotation. The right column corresponds to counter-clockwise rotation. The slow, medium, and fast speeds correspond to the top to bottom rows respectively. Note that the time axes were chosen to cover one revolution. Thus, for each row the duration displayed is different. A substantial drift is observed relative to the zero crank position after only one revolution.

#### 4.1.1 Brownian Noise

This drift was found to be well-described by Brownian noise. Brownian motion is the random motion of particles suspended in a medium (e.g. a liquid or gas) [Brown, 1828]. This motion is named after the botanist Robert Brown, who first described the phenomenon in 1827, while looking through a microscope at pollen of the plant *Clarkia*

pulchella immersed in water. In 1905, Einstein [Einstein, 1905] studied Brownian motion and showed that the mean squared displacement of a one dimensional random walk was related to the time interval  $\Delta t$  by the expression

$$\Delta x^2 = 2D\Delta t \tag{4.1}$$

where  $D$  is a diffusion coefficient. In this case the diffusion coefficient is an average measure of the stochastic activity of the random walker. This served as the first evidence that atoms and molecules exist. This type of behavior is also often referred to as a Wiener process, in honor of Norbert Wiener who later rigorously derived the Einstein observation. A historical review of this research was presented by Nelson [Nelson, 1967].

Given a random white noise process  $X_w(t)$ , such that its spectral density function  $S_{X_w}(f) = k_b$  where  $k_b$  is a constant for all frequencies  $f$ , Brownian noise can be obtained by integration of the white noise process

$$X_b(t) = \int_0^\infty X_w(\tau)d\tau. \tag{4.2}$$

This is demonstrated in Figure 4-2. Brownian noise has a power spectral density characterized by,

$$S_{X_b}(f) = \frac{k_b}{f^2} \tag{4.3}$$

For this reason, Brownian noise has a power spectrum with a slope of exactly -20 dB/dec. Without knowledge of the process value at some initial time  $x(t_o)$ , the non-stationary behavior of a Brownian process makes estimating the first central moment difficult. The study of processes with poorly-defined first central moments (or even all moments) has been the subject of a substantial amount of work applying statistical mechanics to study neurophysiological systems [Tuckwell, 1989, Holden, 2013]. Unfortunately, to the author's knowledge, many of these methods are developed for gradient-based first order systems. The extension of these ideas to second-order or higher dynamics requires care.




Figure 4-2: (top) Four instances of a white noise process, and (bottom) four instances of a Brownian noise process are presented. Each  $r^{th}$  replication of the Brownian noise process was produced by integrating the corresponding replication of the white noise process.

### 4.1.2 Evidence of Brownian Noise during Crank-Turning

In order to quantify the presence of Brownian noise during the task of crank-turning the experimental data presented in Chapter 2 and 3 were further analyzed.

#### Methods

The crank angle error was defined as the  $\theta_{error} = \theta_c - \theta_{target}$ , where  $\theta_{target}$  is the displacement that would have been observed if subjects moved with the instructed constant speed (slow 0.075 rev/sec, medium 0.5 rev/sec, and fast 2 rev/sec). The power spectrum of the error in crank angle was estimated using 21 hamming windows. These windows corresponded to the 21 trials of the experiment. The length of the windows was chosen to be as long as possible, given a starting crank angle of 0 rad (the 3 o'clock position). In practice, this allowed for a slow (16.8 s), medium (9.3 s), and fast (1.9 s) window length.

## Results

The presence of temporal drift was visually evident in the human crank angle and angle error (see Figure 4-1 and 4-3). The power spectrum of crank angle error displayed a -20 dB/dec slope over the entire frequency range for which the spectra could be estimated. This observation was robust to both speed and direction.

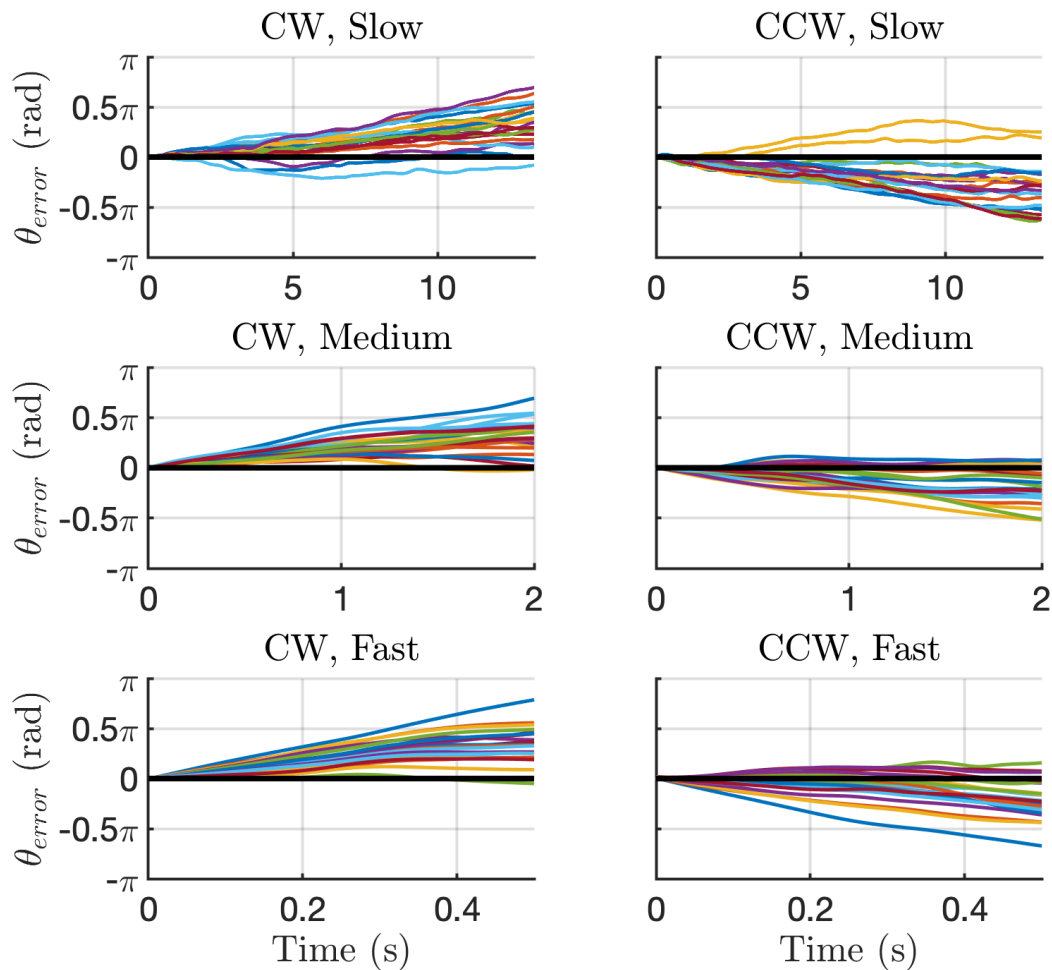


Figure 4-3: A representative subject's position error with respect to time ( $\theta_{error} = \theta_c - \theta_{target}$ ). The differently colored lines represent different trials. The black line represents the zero error case. The left column corresponds to clockwise rotation and the right column corresponds to counter clockwise rotation. The slow, medium, and fast speeds correspond to the top to bottom rows respectively.

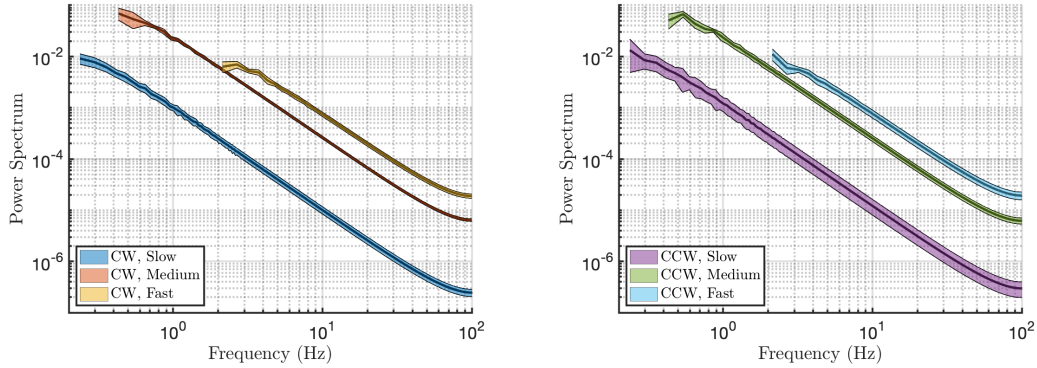


Figure 4-4: Power spectral estimates of the crank error for the (left) CW and (right) CCW turning directions. The mean is the average over subjects, and the error bars represent plus or minus one standard deviation from the mean. Note the clear -2 mag/dec slope in all 6 spectral estimates.

### 4.1.3 System Identification in the Presence of Brownian Noise

Processes which drift are non-stationary. Rather than an inclusive categorization, non-stationarity is defined as the lack of stationarity. Analogous to nonlinearity, there are many underlying structures in data which can result in a violation of stationarity. For this reason, there is no single well-defined method to deal with this class of problems. In practice, two common approaches to deal with violations of stationarity which will be relevant to this work are: leveraging multiple replications (assuming an ensemble) or temporally averaging within a replication to estimate the underlying moments (local stationarity).

In the prior analysis, the key to dealing with the time varying nature of  $x_0(t)$  and  $Z(t)$  was to enforce a specific assumption that the noise influencing the estimate at a particular time,  $t$ , did not drift. However, the Brownian motion present in the crank-turning data resulted in drift. In the experimental human crank data, a specific time index may occur at substantially different configurations from one cycle to the next (if the time index is reset at the start of each crank cycle). Interestingly, the raw velocity and force data exhibited systematic fluctuations with respect to crank position – suggesting a solution (see Figure 4-5).






Figure 4-5: Plot of a representative subject’s data from the crank-turning studies presented in Chapter 2 and 3. Blue lines represent multiple crank cycles in the clockwise direction. Red lines represent multiple cycles in the counter-clockwise directions. Velocity and normal force are presented with respect to time and configuration variable  $\theta$  (crank angle). Qualitatively there is little to no pattern in the velocity or normal force with respect to time. However, clear systematic patterns are present with respect to the configuration variable.

The Brownian motion present in the crank-turning data results in drift. In the experimental human crank data, a specific time index from one cycle to the next may occur at substantially different configurations.

We observed that structured variations were more evident with respect to configuration than time. In the following method we propose to differentiate the non-stationary position measurement. If a process is Brownian, the derivative of the process is white noise with zero mean and constant variance – stationary by definition. (see Figure 4-2). However, a task with constant speed motion in the desired degree of freedom cannot have a zero mean velocity. This was the case in the crank-turning experiment. Before differentiation, high-pass filtering the position measurement allowed for the steady-state velocity to be removed.

#### 4.1.4 Configuration-Dependent Ensemble Method

The standard ensemble method for time-varying system identification was altered in two ways: (1) the position measurement was high-pass filtered, and (2) a configuration-dependent ensemble index was used. The high-pass filter removed the influence of the underlying time-varying ZFT while the configuration-dependent ensemble avoided artifacts resulting from temporal drift. These alterations required additional assumptions:

- The information required to estimate impedance is present in the frequency range above that of the high-pass filter.
- The apparently random fluctuations in the zero-force trajectory can be described as a Brownian noise process.
- The time-variation of impedance arises exclusively from the configuration-dependence of impedance.
- The human zero-force trajectory is smooth and differentiable.

In theory, given these assumptions, the influence of a zero-force trajectory corrupted by Brownian noise can be removed, leaving only the influence of the configuration-dependent mechanical impedance. Further discussion of the fundamental assumptions of this work, presentation of prior ensemble methods, and simulations using ensemble methods are presented in Appendix B.

Assume  $R$  replications of a configuration-dependent process such that each replication has a variable number of time samples  $N_r$ . In this case,  $l$  is the index of discrete uniformly-spaced time steps ranging from 0 to  $N_r$ . In the following definitions, the subscript  $r$  notation indicates that a variable is a function of replication. Assume it is possible to measure or separately estimate the configuration variable  $\theta_r(l)$ , an input  $u_r(l)$ , and noisy measurement of the output  $z_r(l)$  while the noiseless output measurement  $x_r(l)$  is not directly measurable. Given  $N$  evenly-spaced configuration bins  $\theta_{bin}(i)$  between 0 to  $2\pi$  the  $i^{th}$  configuration bin of an ensemble can be defined

as

$$\gamma_r^{(i)} = \underset{l \in [1, N_r]}{\operatorname{argmin}} |\theta_r(l) - \theta_{bin}(i)|. \quad (4.4)$$

Switching the relation around yields  $\theta_{bin}(i) = \theta_r(\gamma_r^{(i)})$ . Provided these indices an  $N$ -dimensional configuration-dependent subset is defined for each replication. In this case,  $\gamma_r^{(i)}$  is the index of the  $z_r(\gamma_r^{(i)})$  which provides the value of the noisy ensemble measurement closest to the  $i^{th}$  crank position bin. This concept is graphically displayed in Figure 4-6.




Figure 4-6: The variable  $\gamma_r^{(i)}$  is the time index of the  $r^{th}$  replication closest to a the  $i^{th}$  crank position bin ( $\theta_{bin}(i)$ ). The red, blue, and yellow sample records represent different exemplary replications of crank angle sampled in time. The horizontal black dotted line represents the crank position  $\theta_{bin}(i)$ . The vertical dotted lines are the time positions for each of the three replications which are closest to the  $i^{th}$  crank position bin.

$$z_r(\gamma_r^{(i)}) = \Delta t \sum_{j=0}^M \hat{h}(\gamma_r^{(i)}, j) u_r(\gamma_r^{(i)} - j) \quad (4.5)$$

where  $\hat{h}(\gamma_r^{(i)}, j)$  is an impulse response function (IRF) estimate with finite lag length  $M + 1$ , where  $j$  is a lag index. This assumes  $\hat{h}(\gamma_r^{(i)}, j) = 0$  for  $j > M$ , which is true for causal systems. The Wiener-Hopf equation can be obtained by multiplying both

sides of Equation 4.5 with  $u_r(\gamma_r^{(i)} - j)$  and summing over all realizations

$$\frac{1}{R} \sum_{r=1}^R z_r(\gamma_r^{(i)}) u_r(\gamma_r^{(i)} - k) = \Delta t \sum_{j=0}^M \hat{h}(\gamma_r^{(i)}, j) \frac{1}{R} \sum_{r=1}^R u_r(\gamma_r^{(i)} - j) u_r(\gamma_r^{(i)} - k). \quad (4.6)$$

Given the input auto-correlation

$$\hat{\phi}_{uu}(i, \tau) = \frac{1}{R} \sum_{r=1}^R u_r(\gamma_r^{(i)}) u_r(\gamma_r^{(i)} + \tau) \quad (4.7)$$

and the input-output cross-correlation

$$\hat{\phi}_{zu}(i, \tau) = \frac{1}{R} \sum_{r=1}^R z_r(\gamma_r^{(i)}) u_r(\gamma_r^{(i)} + \tau), \quad (4.8)$$

substituting Equation 4.7 and Equation 4.8 into Equation 4.6 yields,

$$\hat{\phi}_{zu}(i, -k) = \Delta t \sum_{j=0}^M \hat{h}(\gamma_r^{(i)}, j) \cdot \hat{\phi}_{uu}(i - k, k - j) \quad (4.9)$$

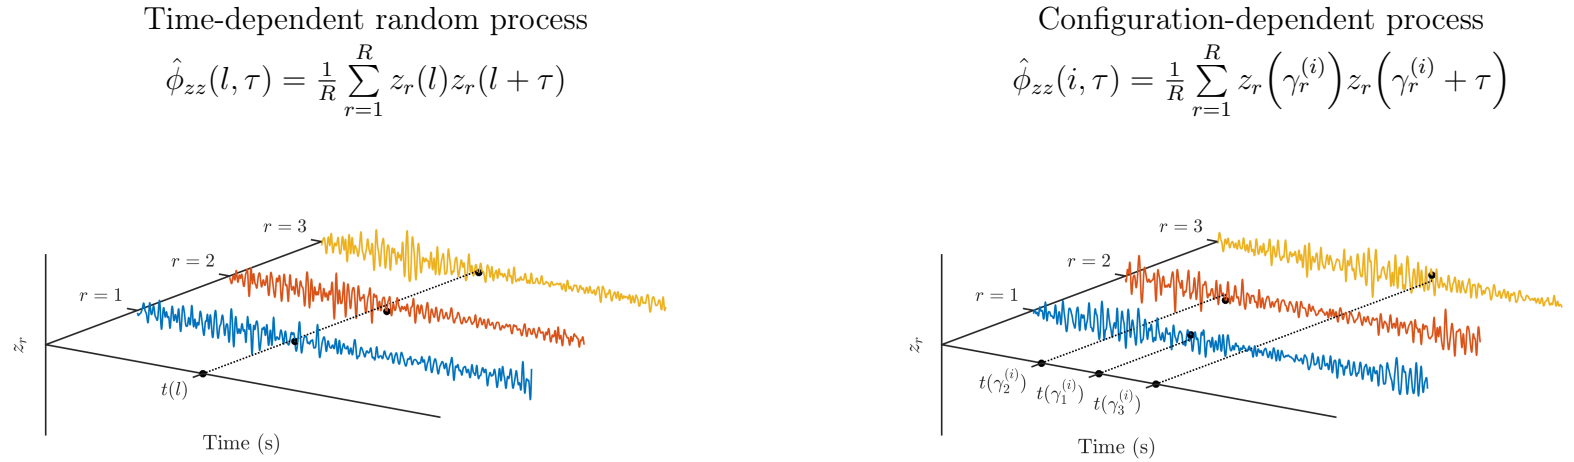


Figure 4-7: (left) An ensemble of a time-dependent process. (right) An ensemble of a configuration-dependent process. The red, blue, and yellow sample records represent different exemplary replications sampled in time. Note, in both cases the variance changes as a function of time. In the configuration-dependent case, the variance change is not temporally aligned across ensembles. Leveraging knowledge of the configuration,  $\gamma_r^{(i)}$ , allows for the cross-correlation to be computed. Note,  $l \in \mathbb{Z}$  is the time  $t$  index that corresponds to uniform steps in time, where  $l \in [1, N_l]$  while  $i \in \mathbb{Z}$  is the configuration  $\theta_c$  index that corresponds to uniform steps in configuration, where  $i \in [1, N]$ .

A matrix expression can be constructed by changing the lag index  $j$  and  $k$  from 0 to  $M$ , where  $\hat{\Phi}_{uu}(i) \in \mathbb{R}^{L \times L}$  and  $\hat{\Phi}_{zu}(i) \in \mathbb{R}^{L \times 1}$ . In this case the IRF estimate,  $\hat{\mathbf{h}}(i)$ , can be computed using a Moore-Penrose pseudo-inverse matrix operation.

$$\hat{\Phi}_{zu}(i) = \Delta t \hat{\Phi}_{uu}(i) \hat{\mathbf{h}}(i) \quad (4.10)$$

$$\hat{\mathbf{h}}(i) = \Delta t \hat{\Phi}_{uu}(i)^{-1} \hat{\Phi}_{zu}(i) \quad (4.11)$$

## 4.2 Methods

### 4.2.1 Participants

A single-subject pilot experiment was performed to test this configuration-dependent ensemble method. One healthy male subject (age 30) was recruited for the study. The participant was right-handed, and did not report any biomechanical injury to their arm nor any neurological problems. Prior to participating in the study, he was informed about the experimental procedure and signed an informed consent document approved by MIT’s Institutional Review Board.

### 4.2.2 Experimental Apparatus and Procedure

This experimental design is similar to that of Chapter 2 and 3. All experiments were performed with an InMotion2 Shoulder-Elbow robot (Interactive Motion Technologies Inc.). The InMotion is a highly back-drivable torque controlled  $x - y$  planar robot, which was designed for stroke rehabilitation and human motion research. The custom control system for the InMotion robot was implemented on a CompactRIO 9034 controller, with low-level functionality implemented at 2 kHz on a CompactRIO FPGA, and high-level functionality implemented at 350 Hz on a CompactRIO real-time processor. InMotion joint positions were measured by a 16-bit/rev encoder and interaction forces were measured at the end-effector of the InMotion robot using an ATI Gamma force/torque transducer [Hogan et al., 1995, Thorup, 2018].

During the experiment, the subject's arm was occluded from view by a wooden structure, which did not limit the range of motion. The upper arm was suspended by a canvas sling connected to the ceiling using a steel cable; the upper and lower arm moved in a horizontal  $x - y$  plane. Anatomically this corresponded to a transverse plane. The subject sat in a chair with a rigid back, while the shoulder was constrained by a harness attached to the back of the chair. The subject was positioned such that the virtual crank, with radius 10 cm, was well within the work-space of the arm. The visual display, also generated by the computer, was on a 38 cm monitor mounted approximately 75 cm from the subjects' eyes. The experimental setup is displayed in Figure 4-8.



Figure 4-8: Experimental setup. The InMotion2 simulated a virtual circular constraint. Vision of the arm and robot were occluded. The subject was provided with visual speed feedback. The wrist was braced, the elbow was supported by a sling, and the shoulders were strapped to a chair.

The experiment comprised three 15-minute constant-target-speed turning trials. Each trial began by initializing the robot, zeroing the force transducer, and instructing the subject to grasp the handle. Perturbations imposed by the robot began 10 seconds after the robot was initialized. The subject was instructed to begin turning in the counter-clockwise direction after the force perturbations started. The subject was instructed to match the line representing their tangential velocity to the constant-target-speed line in the visual display. Achieving this goal would correspond to moving with a constant period of 13.33 seconds per revolution. In practice, the subject completed approximately 65-70 revolutions per trial. In all conditions the robot control was the same, simulating a virtual constraint and applying small force perturbations in the normal and tangential directions. The subject was given a three-minute break between each of the three trials.

### **4.2.3 Robot Control**

The control of the robot was separated into the normal and tangential direction of the virtual constraint. In the normal direction, the stiffness was 1500 N/m and the damping was 40 Ns/m while in the tangential direction the stiffness was zero and the damping was 5 Ns/m. A random square-wave force perturbation was applied in each of the normal (-1.5 N and 1.5 N) and tangential (-0.5 N and 0.5 N) directions. The power spectrum of the force perturbation was flat up to 20 Hz (See Figure 4-10 upper left). Preliminary experiments were performed to determine the amplitude of the perturbations. These experiments are presented in Appendix E.

### **4.2.4 Robot-only**

The robot-only experiment was conducted in order to estimate the impedance of the robot without the human interacting with it. Without the human exerting forces on the handle, an additional control term was required. A virtual human stiffness of 500 N/m was added in the tangential direction. Stiffness was not added in the normal direction as the dynamics of the virtual crank confined the system in this degree of



freedom. The virtual stiffness was connected to a nominal position  $\mathbf{x}_0$  which moved around the constraint with constant speed. This speed achieved the task goal exactly, moving the manipulandum with a period of 13.33 seconds.

### 4.2.5 Simulations

A simulation experiment was also performed. A model of a human was based on the previous crank-turning experiments presented in Chapters 2 and 3. This ideal data set allowed potential problems due to hardware to be separated from the system identification implementation. This also served as a way to test whether the system ID methods worked properly. The human was modeled as a two-link manipulator interacting with a circular constraint. The human inertial parameters were the same as in the previous work (see Appendix A). The system was simulated in human joint-space coordinates by numerically integrating the expression to solve for  $\dot{\mathbf{q}}$  and  $\ddot{\mathbf{q}}$ .

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} = \boldsymbol{\tau} + \mathbf{J}^T(\mathbf{F}_{pret} + \mathbf{F}_c) - \mathbf{C}(\mathbf{q}, \dot{\mathbf{q}}) \quad (4.12)$$

The term  $\mathbf{M}(\mathbf{q})$  was the mass matrix and  $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})$  was the Coriolis and centrifugal terms for the entire linkage system including the human and InMotion robot. The inertial parameters of the InMotion robot are presented in Appendix D. The forward kinematic map from joint-space to hand space was denoted  $L$  and the inverse kinematic map from hand space to joint-space was denoted  $L^{-1}$ . The torque resulting from the constant joint impedance was defined as

$$\boldsymbol{\tau} = \mathbf{K}_q(L^{-1}(\mathbf{x}_0) - \mathbf{q}) + \mathbf{B}_q(\mathbf{J}^{-1}\dot{\mathbf{x}}_0 - \dot{\mathbf{q}}). \quad (4.13)$$

The joint-space stiffness was constant and the damping was proportional to the stiffness ( $\mathbf{B}_q = 0.05\mathbf{K}_q$ ). The stiffness in units of N-m/rad was defined as

$$\mathbf{K} = G \begin{bmatrix} K_{11} & K_{12} \\ K_{21} & K_{22} \end{bmatrix} = G \begin{bmatrix} 29.5 & 14.3 \\ 14.3 & 39.3 \end{bmatrix} \quad (4.14)$$

where  $G$  was equal to 0.5. This choice of impedance is consistent with Chapters 2 and 3. The zero-force trajectory,  $\mathbf{q}_0$ , was defined in hand space and transformed to joint-space coordinates. In hand space it was assumed to be composed of two out-of-phase sinusoids with the same amplitude  $A$  of 0.1 m, frequency  $\Omega$  equal to the turning frequency  $2\pi(0.075)$  rad/s, and with a phase offset.

$$\mathbf{x}_0 = \begin{bmatrix} A \sin(\Omega(t + w_b) + \phi_1) \\ A \sin(\Omega(t + w_b) + \phi_2) \end{bmatrix} \quad (4.15)$$

The phase in the  $x$  direction  $\phi_1$  was  $-1.49\pi$ , and the phase in the  $y$  direction  $\phi_2$  was  $2\pi$ . This zero-force trajectory produced qualitatively similar velocity and force fluctuations to those observed previously in Chapter 2. Phase drift was produced through the addition of a Brownian noise term  $w_b$ . The Brownian noise was generated by integrating normally-distributed pseudo-random noise scaled by  $2.5 \times 10^{-2}$  and low-pass filtering with a 4<sup>th</sup> order Butterworth filter with a cutoff frequency of 2 Hz. These parameters were selected such that the crank angle drift and the power spectrum were comparable to that of the human subject.

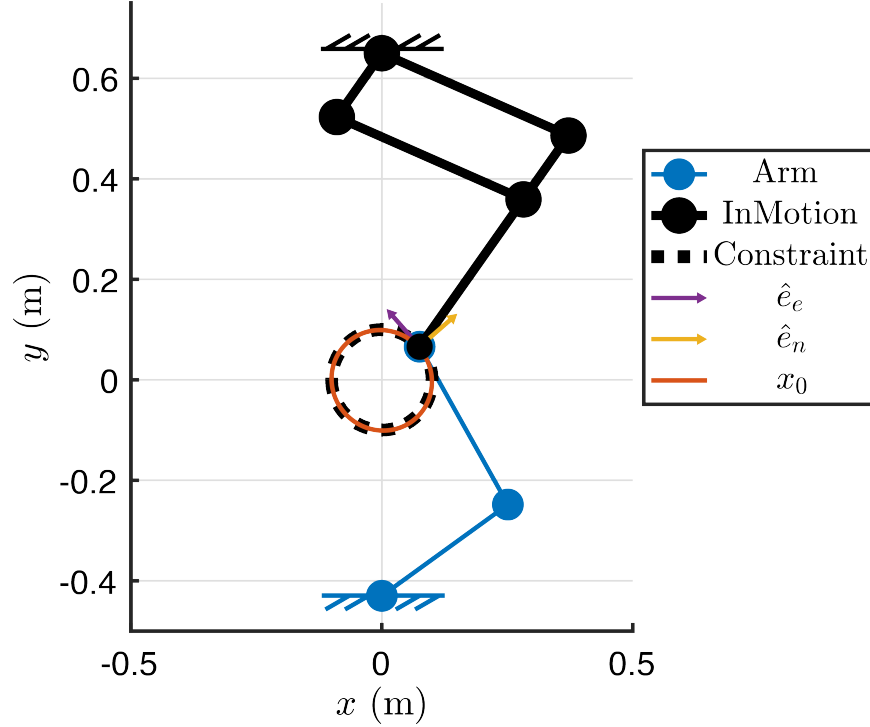


Figure 4-9: Diagram of the system model used for simulation

A pseudo-random perturbation with a magnitude of 1N or -1N was denoted  $F_{pert}$ . The force exerted by the circular constraint was denoted  $F_c$ . The perturbations were applied in the normal  $\hat{e}_n$  and tangential  $\hat{e}_e$  directions. Consistent with the robot experiment, the normal stiffness of the constraint was 1500 N/m and the damping was 40 Ns/m while a 5 Ns/m damping was applied in the tangential direction.

The simulations were performed at 200 Hz. Seven crank-turning cycles (slow speed 13.33 seconds per revolution) were simulated. The first cycle was discarded. This was repeated to generate 300 crank cycles of data from which to perform system identification.

The rotation matrix as a function of crank angle  $\theta$  transformed the Cartesian coordinates with unit vectors  $\hat{e}_x$  and  $\hat{e}_y$  to the normal and tangential coordinates with unit vectors  $\hat{e}_n$  and  $\hat{e}_e$ . This notation is displayed graphically in Figure 4-9. This transformation was denoted

$$\mathbf{R}(\theta) = [\hat{e}_n, \hat{e}_e] = \begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix}. \quad (4.16)$$

## 4.2.6 Data Processing

The data were re-sampled uniformly in time with a sampling frequency of 200 Hz for  $R$  replications. A replication corresponded to a single crank cycle. In practice, the number of time samples  $N_r$  in a replication varied. The configuration variable was defined as crank angle  $\theta_r(l)$ . The system input was the commanded force perturbation in the normal  $u_r^n$  or tangential  $u_r^e$  directions. The measured position was denoted  $z_r^n$  in the normal and  $z_r^e$  in the tangential direction. The position measurement was high-pass filtered with a cut-off frequency of 0.1 Hz. The binned configuration variable  $\theta_{bin}$  was defined with  $N = 300$  bins evenly spaced from 0 to  $2\pi$  rad. The perturbation applied at the handle caused small high-frequency deviations in crank angle. In the computation of  $\gamma_r^{(i)}$  the crank angle position was low-pass filtered with a 4<sup>th</sup> order Butterworth filter, cut-off frequency of 0.1 Hz.

The configuration-based ensemble method was applied in the normal and tangential directions. Clearly the real systems have coupling. However, in practice the coupling terms have a low stiffness. Without inertial compensation, this results in extremely noisy estimates of the coupling terms. This noise is so substantial that the matrix inversion as a MIMO problem resulted in poor estimates even on the diagonal terms. Thus, in the following analysis, the behavior in the normal and tangential directions were processed separately. In the normal direction the drift was small. For this reason the additional step of differentiating the input and output position to deal with the Brownian noise was unnecessary and only increased noise. In the normal direction the IRF was estimated directly from equation 4.11. The presence of drift in the tangential direction was substantial. For this reason the input and output were differentiated. Differentiation amplifies noise, and for this reason the tangential estimates were low-pass filtered before differentiation with a 4<sup>th</sup> order Butterworth filter with a cut-off frequency of 30 Hz. After this filtering the equation used to solve for  $\hat{\mathbf{h}}_e(i)$  was

$$\hat{\mathbf{h}}^e(i) = \Delta t \left( \hat{\Phi}_{\dot{u}^e \dot{u}^e}(i) \right)^{-1} \hat{\Phi}_{z^e \dot{u}^e}(i) \quad (4.17)$$

The estimated IRFs were filtered in both arguments – lag and configuration. In

the configuration direction, a moving-average filter with a window width of 15 bins was applied. In the lag direction a 4<sup>th</sup> order Butterworth filter with a frequency of 15 Hz was applied to the normal estimates while a 4<sup>th</sup> order Butterworth filter with a frequency of 7 Hz was applied to the tangential estimates. A different cut-off frequency was used for the normal and tangential directions because the underlying dynamics to be identified had a different natural frequency.

For each time step  $i$ , the best fit second-order impulse response function ( $h_{model}(i)$ ) was fit by minimizing the mean squared error between  $\hat{h}(i)$  and the model using a bound-constrained optimization. This model is described by parameters ( $\hat{m}(i)$ ,  $\hat{b}(i)$ ,  $\hat{k}(i)$ ). To quantify the model fit the % $VAF_{IRF}$  between  $\hat{h}(i)$  and  $h_{model}(i)$  was assessed.

$$\%VAF_{IRF}(r) = 100 \times \left( 1 - \frac{\text{var}(\hat{h}(i) - h_{model}(i))}{\text{var}(\hat{h}(i))} \right) \quad (4.18)$$

To quantify the error in the parameter estimates, the average percent error between the expected and estimated values were computed for the robot-only and the simulation experiments.

The kinematic stiffness term was neglected during this analysis. The kinematic stiffness term results from  $\frac{\partial J}{\partial q} F$ . The term  $\frac{\partial J}{\partial q}$  was less than 1 in the works-space occupied by the virtual crank. In addition, the forces observed in simulation and in the experiments were less than 5N. Thus, the magnitude of the kinematic stiffness term in this experiment was smaller than 5 N/m – smaller than the resolution of our measurements.

### 4.3 Results

The results from two experiments and one simulation study are presented. First, the robot-only experiment with the InMotion robot; then, a human interacting with the InMotion robot; finally, simulations of a robot coupled to a human.

### 4.3.1 Robot

In order to estimate the impedance of the robot independent of the human subject, trials were collected where the robot turned with a ‘virtual’ human stiffness in the tangential direction. The results are presented in Figure 4-10 and 4-11.

It is clear from the plot of crank position vs. time (Figure 4-10 upper right) that the InMotion turned with a constant period without evidence of Brownian noise. Consistent with this observation, the power spectra of the input and output signals have a slope of zero at low frequencies.

The estimates of mass and damping (Figure 4-11) were compared to the expected inertial dynamics of the robot, the commanded damping of 40 Ns/m in the normal direction, and the commanded damping of 5 Ns/m in the tangential direction. Deviation from the dashed black lines is artifact, presumably due to the system ID method or hardware limitations. These estimates resulted in an average mass error of 8.0% in the normal and 6.9% in the tangential direction, a damping error of 11.4% in the normal and 17.0% in the tangential direction, and a stiffness error of 17.6% in the normal and 7.0% in the tangential direction. The estimated stiffness was consistently larger than the programmed 1500 N/m stiffness in the normal direction.

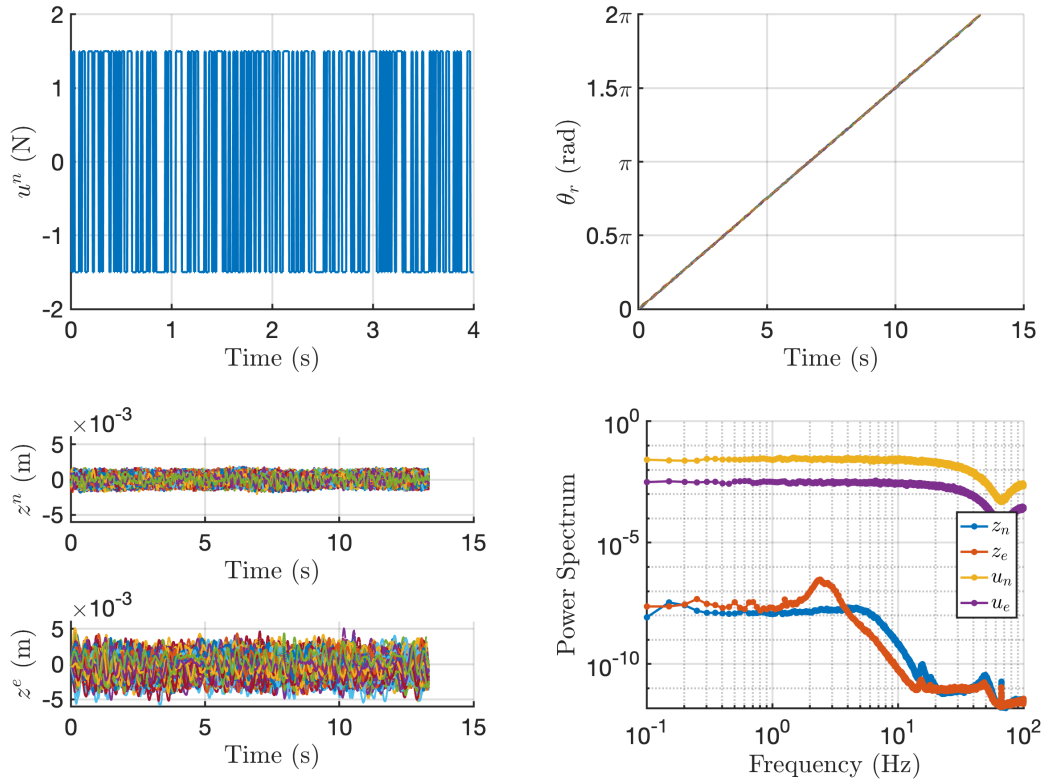


Figure 4-10: (top left) Force perturbations from four seconds of a single replication in the normal direction. (top right) The measured configuration  $\theta_r$  with respect to time. Note that there is no drift in the robot control. Each  $r^{th}$  realization was a different color line. Note it appears that there is only one line because the realizations lie on top of one another. (lower left) Normal  $n$  and tangential  $e$  displacements after high-pass filtering. (lower right) Power spectra for the force perturbations and measured displacements in both the normal and tangential directions.

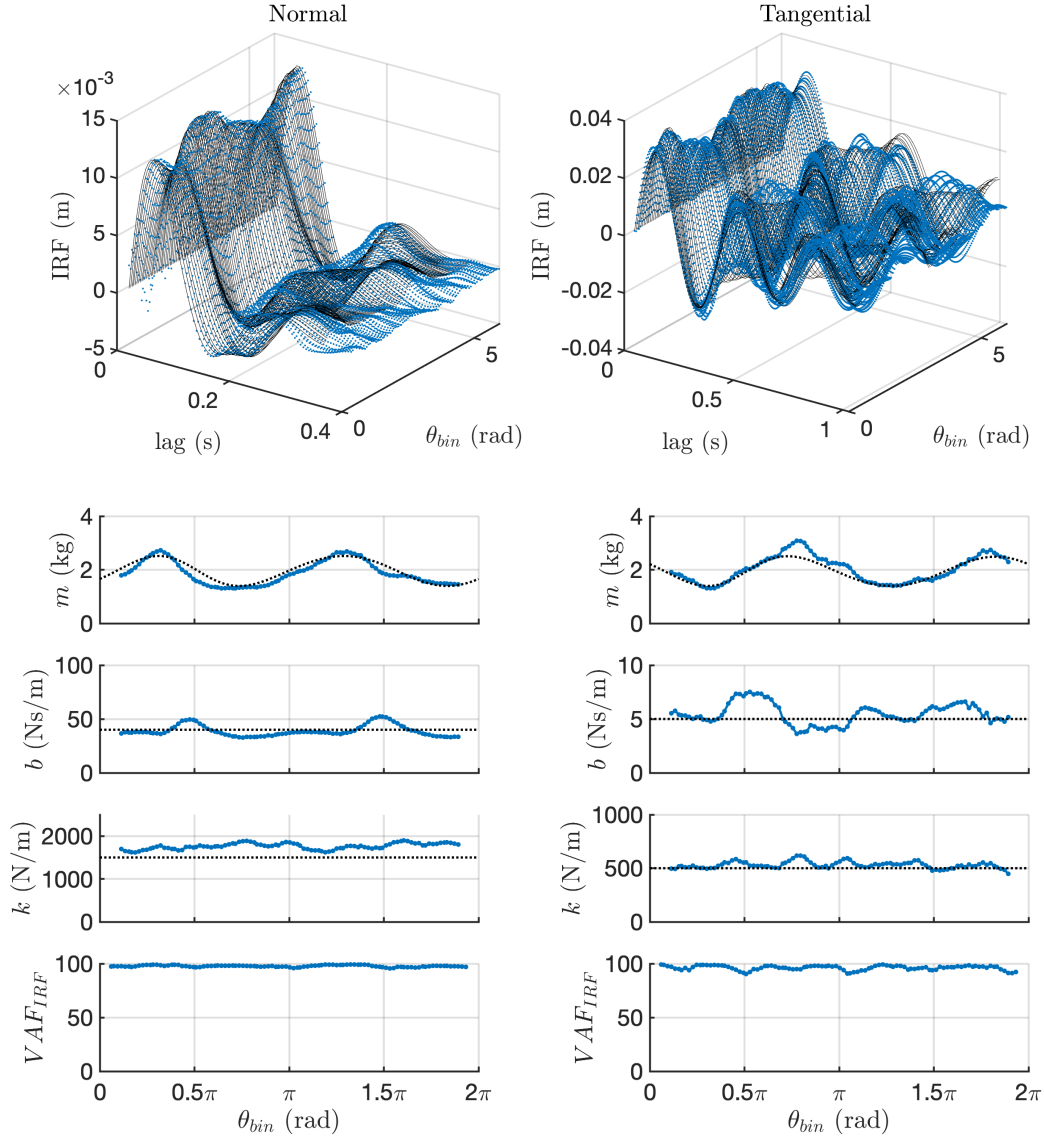


Figure 4-11: (top row) The blue lines represent the IRF estimates and the black lines represent the best-fit second-order model. (lower rows) Blue lines represent the parameter estimates from the best-fit model. The dashed black lines represent the expected inertia, damping, and stiffness. The left column contains the results from the normal direction. The right column contains the results from the tangential direction.

### 4.3.2 Robot and Human

The results from the subject interacting with the virtual crank are presented in Figure 4-12 and 4-13. In this case, the drift present in the human subject data is clearly



evident in the crank angle vs. time plot (upper left of Figure 4-12) in addition to the -20 dB/dec slope present in the power spectrum of the tangential direction (lower right of Figure 4-12).

In Figure 4-12 the red lines represent estimates from experimental data. In contrast, the black dashed lines provide insight as to what would be expected from the model assumed in Chapter 2 and 3. The approximate inertia of the combined human and robot system was computed and plotted for reference (black dashed line of Figure 4-13). In addition, a reference damping and stiffness were plotted. These lines represent the stiffness and damping which result from the superposition of (1) the programmed robot impedance and (2) a model of human impedance. The model of human impedance was assumed to result from a constant joint-space impedance which was the same as that used in the simulations (presented in the next section).

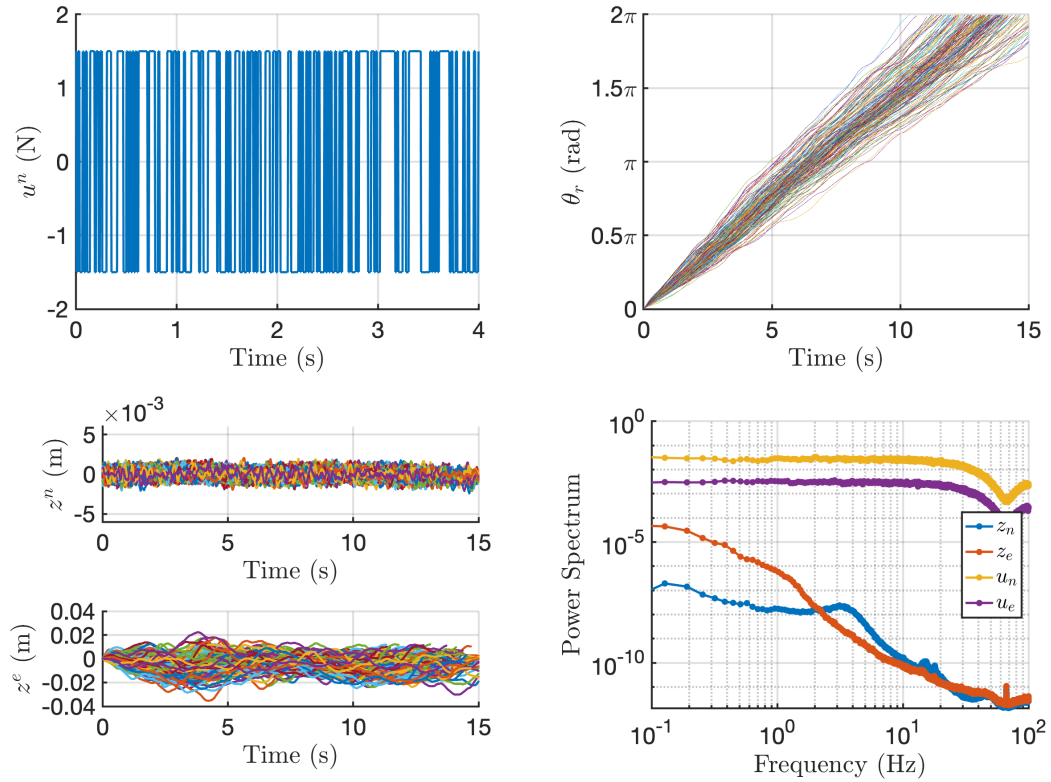


Figure 4-12: Raw measures from the experiment where the subject interacted with the robot. (top left) Force perturbations from four seconds of a single replication in the normal direction. (top right) The measured configuration  $\theta_r$  with respect to time. Note the substantial drift. Each  $r^{th}$  realization was a different color line. (lower left) Normal  $n$  and tangential  $e$  displacements after high-pass filtering. (lower right) Power spectrum for the force perturbations and measured displacements in both the normal and tangential directions.

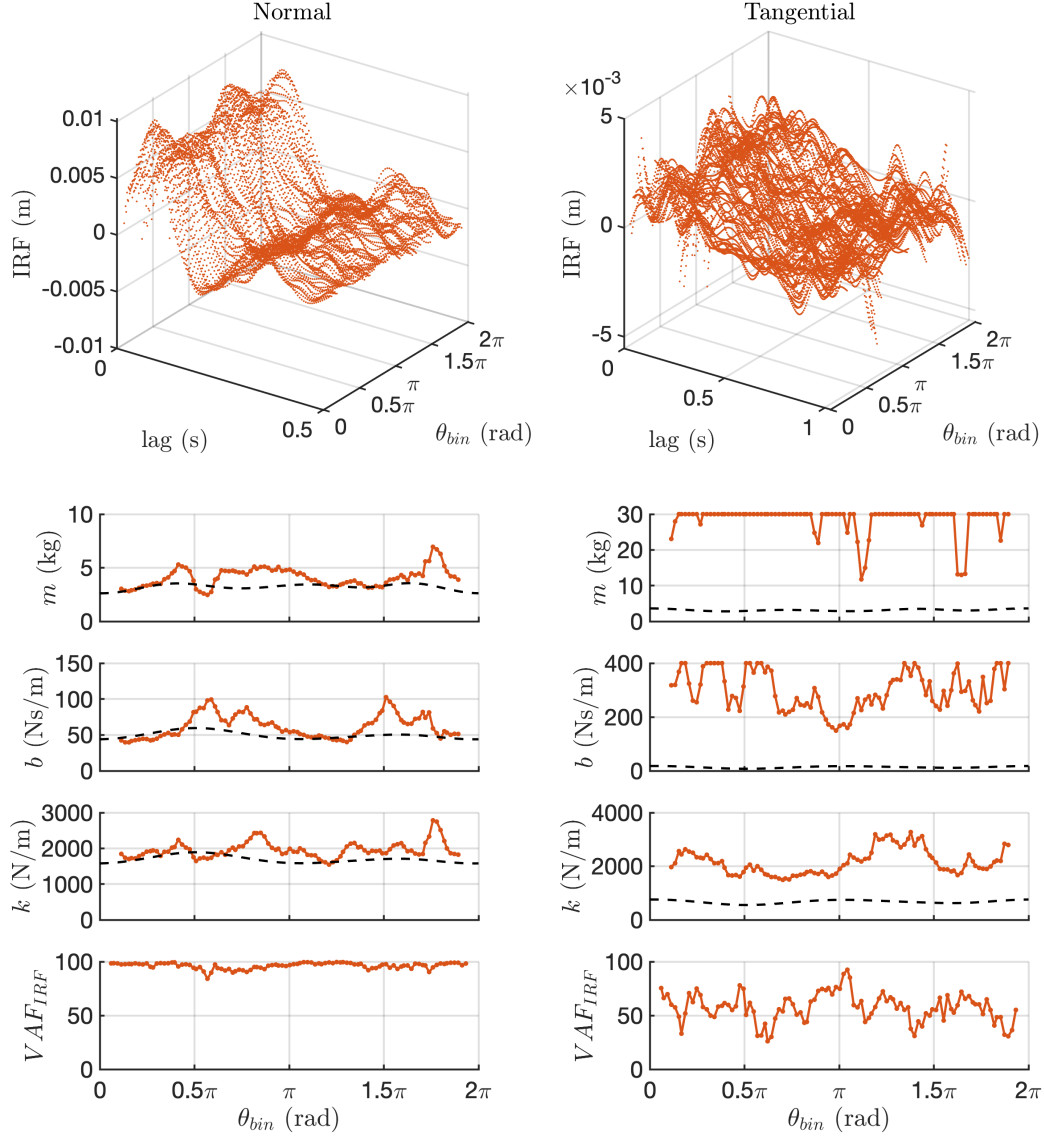


Figure 4-13: (top row) The red lines represent the IRF estimates. (lower rows) Red lines represent the parameter estimates from the best-fit model. The dashed black lines represent the inertia, damping, and stiffness used in the simulations. The left column contains the results from the normal direction. The right column contains the results from the tangential direction.

### 4.3.3 Simulations

Simulations were developed to validate the algorithm independent of the robotic hardware. Constant joint-space stiffness consistent with the previous work was assumed and a zero-force trajectory with Brownian noise was generated. A substantial amount

of drift, produced by the addition of Brownian noise, was present in the simulations. This drift is visible in the crank angle vs. time plot (upper right of Figure 4-14) and the power spectrum of the tangential direction (lower right of Figure 4-14)

In Figure 4-15 the IRF estimates in the normal direction are well described by a second-order model and the parameters closely match those of the expected inertia, stiffness, and damping, while the tangential direction presented more variability. These estimates resulted in an average: mass error of 2.2% in the normal and 28.3% in the tangential direction, a damping error of 3.1% in the normal and 39.5% in the tangential direction, and a stiffness error of 2.4% in the normal and 39.1% in the tangential direction.

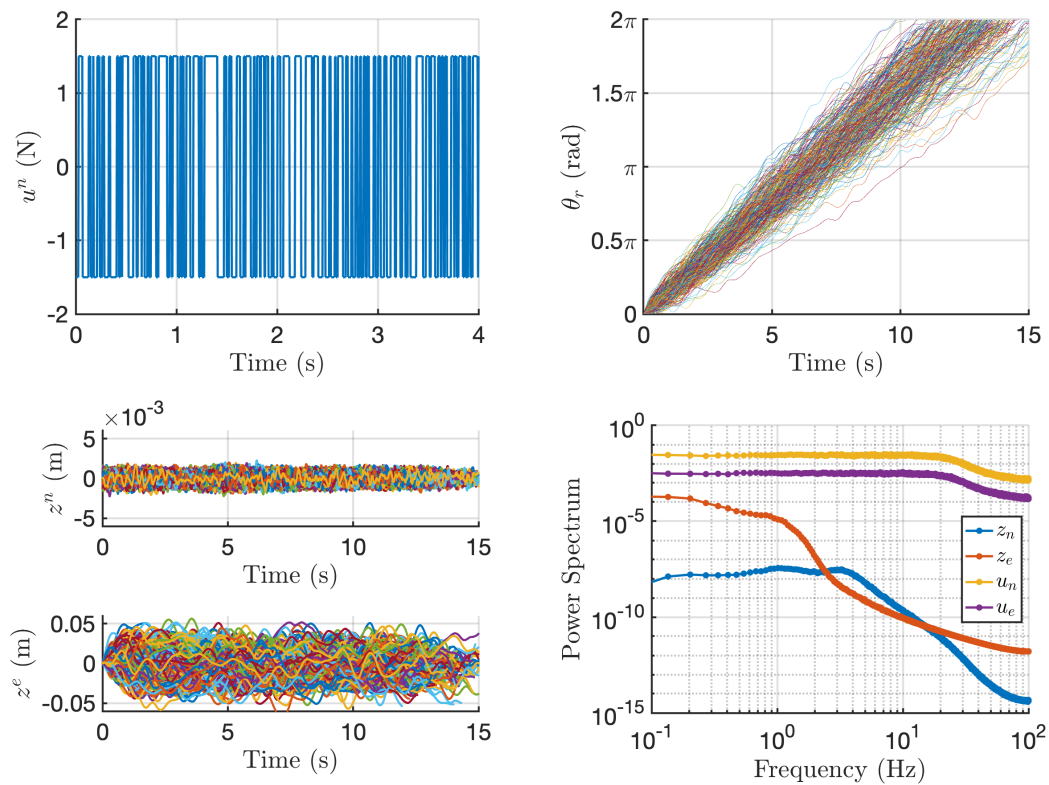


Figure 4-14: (top left) Force perturbations from four seconds of a single replication in the normal direction. (top right) The measured configuration  $\theta_r$  with respect to time. The drift was chosen to be comparable to the human experiments. Each  $r^{th}$  realization is a different color line. (lower left) Normal  $n$  and tangential  $e$  displacements after high-pass filtering. (lower right) Power spectra for the force perturbations and measured displacements in both the normal and tangential directions.

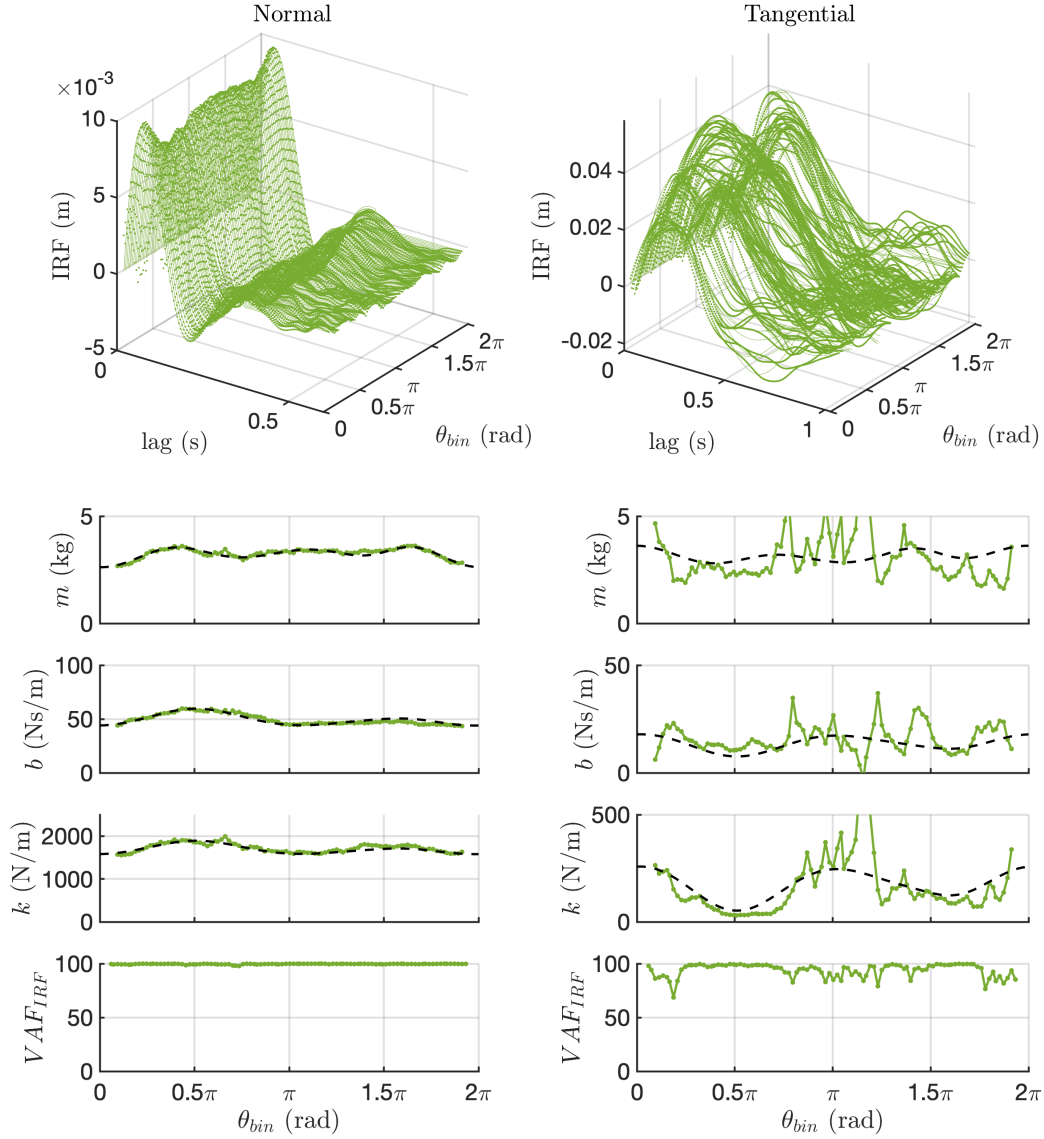


Figure 4-15: (top row) The green lines represent the IRF estimates and the black lines represent the best-fit second-order model. (lower rows) Green lines represent the parameter estimates from the best-fit model. The dashed black lines represent the expected inertia, damping, and stiffness. The left column contains the results from the normal direction and the right column contains the results from the tangential direction.

In Figure 4-16, the overlaid IRFs from the simulated case and the combined human and robot case demonstrate the amplitude difference. Especially in the tangential direction, these errors are substantial and suggest a hardware problem.

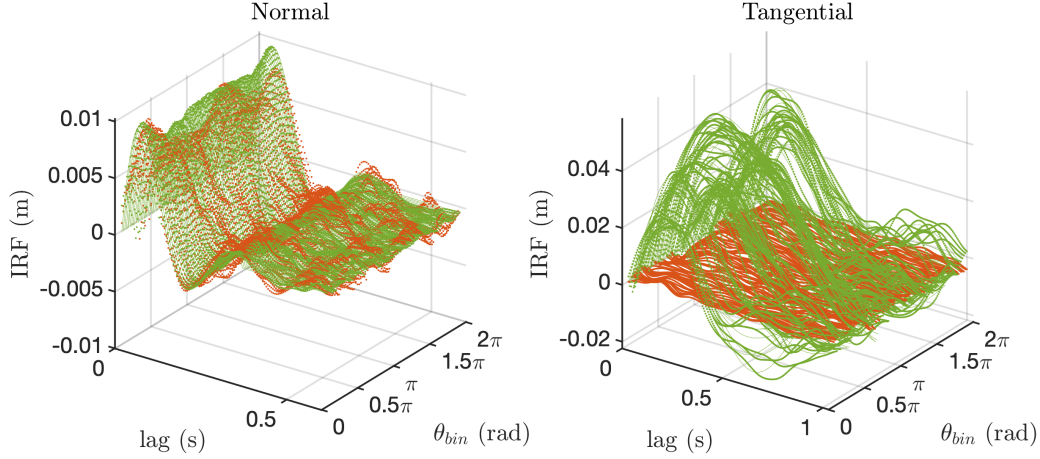


Figure 4-16: Overlaid IRF estimates for the simulations displayed in green and the combined human and robot displayed in red. (left) Results from the normal direction and the (right) the tangential direction.

## 4.4 Discussion

This experiment tested the configuration-dependent ensemble method in simulation and on hardware. This led to several insights. (1) The configuration-based ensemble method, was able to estimate impedance in the normal direction with a error of less than 5% in simulation and 20% in the robot-only experiment. (2) A fundamental limitation was discovered in practice: impedance may be so low that the natural frequency of interactive dynamics at any one configuration may overlap the frequency content of the non-stationary noise process present in the zero-force trajectory.

Variable-impedance actuators are fundamentally nonlinear. A general nonlinear actuator model from equivalent circuit theory, the Norton-type network in the impedance operational form, allows for the separation of forward-path dynamics ( $x_0$ ) from interactive dynamics ( $Z\{\cdot\}$ ) [Hogan, 2014]. Likewise, the Norton-type network, is competent to describe the human arm in tasks which involve interaction. It assumes the system has a well-defined zero and the zero-dynamics of the impedance operator are stable.

Furthermore, there are two especially notable properties of a Norton-type network in the impedance operational form: (1) The network is invariant to any change of the

reference frame which does not change the identity of the ‘interactive’ and ‘forward-path’ dynamics (2) If a task can be developed where zero interaction force is observed, the motion source is unambiguously identifiable. This suggests a solution alternative to estimating impedance directly. Hodgson and Hogan [Hodgson and Hogan, 2000] developed a method which achieved this goal in the study of human point-to-point reaching. If the subject maintains the same  $x_0$  across conditions, a similar idea maybe applied to the study of human crank-turning. This is an area of future work.

In the work presented herein we presented an approach with the goal of removing the influence of the forward-path dynamics to allow the estimation of the interactive dynamics. If successful, given an estimate of the interactive dynamics the forward-path dynamics could then be computed – consistent with the method presented in Chapter 2 and 3.

#### 4.4.1 Configuration-Dependent Ensemble Method

The drift in crank angle during a single cycle was frequently as large as  $1/3^{rd}$  of a crank revolution. This could be replicated in the simulation by adding Brownian noise to the forward-path, zero-force trajectory. The inertial dynamics of the system are configuration-dependent. Drift on the order of  $1/3^{rd}$  of a revolution would result in artifactual time-dependent IRF estimates.

There were several reasons to suspect that modeling limb interactive dynamics as a configuration-dependent quantity would be informative. Static postural estimates of the upper limb [Mussa-Ivaldi et al., 1985] can be well-described by constant joint stiffness. Constant joint stiffness results in a hand stiffness which changes with configuration. Furthermore, Perreault et al. [Perreault et al., 2001] demonstrated that, given a particular configuration and force level, the subjects’ hand impedance was well-described by a constant endpoint impedance. In previous Chapters 2 and 3, observations of unperturbed human crank-turning exhibited systematic patterns in velocity and force with respect to crank angle – not with respect to time (see Figure 4-5). This motivated the development of the configuration-dependent ensemble method.

The configuration-based ensemble method resulted in an unbiased estimate of the IRF given several assumptions. It assumed that human limb stiffness changes only as a function of crank angle; that the behavior is smooth such that differentiation is possible; that the zero-force trajectory can be described by a Brownian noise process; and that frequency separation between the interactive dynamics and the forward-path dynamics is possible.

#### 4.4.2 InMotion Hysteresis

The increased errors in the impedance parameter estimates of the robot-only case, in contrast to the simulation case, suggested that there may be an artifact in the InMotion position measurement or force production. To address this concern a follow-up experiment was performed. This experiment is presented in its entirety in Appendix F. It revealed a force-production hysteresis of  $\sim 0.5$  N. The hysteresis was present throughout the force range encompassed by the experiment. At the present time it is unclear whether this hysteresis results solely from friction, the current control amplifier, or another cause.

The virtual crank-turning experiment with force perturbations requires a robot that can produce a feedforward force with a resolution far less than 0.5N. Given that the perturbations used in the experiment were between  $\pm 0.5$ -1.5N the observed hysteresis could result in errors in force production on the order of 10-50%.

While the InMotion is not a new system [Hogan et al., 1995, Thorup, 2018], it is direct-drive with low inertia. To our knowledge there is not another system of comparable or better feedforward force rendering. The WAM (Barret Technology, Newton, MA) is an example of another state-of-the-art mechanically back-drivable system due to its clever cable-driven design. However, it has substantial friction and inertial dynamics in the force and frequency range required for the virtual crank experiment. On the other side of the spectrum, the Kuka LBR iiwa [Schreiber et al., 2010] is a state-of-the-art system in terms of electro-mechanical compensation for inertial dynamics. However, the holonomic gear reduction mechanically amplifies actuator inertia. Even with compensation, it is not possible to render highly back-drivable behavior in the



range of parameters used for the virtual crank-turning experiment (e.g. stiffness of 0 N/m, a damping of 5 Ns/m, and an apparent inertia of 2-4 kg in the tangential direction). At the present time, it is unclear whether 2D feedforward force production, independent of motion, with a resolution in the 0.1 N range is achievable with an existing robotic system.

### 4.4.3 Subtracting Robot Dynamics

Ideally, the ‘robot-only’ impedance estimates would be subtracted from the ‘robot and human’ estimates to yield the ‘human-only’ impedance. The robot stiffness was programmed to be 1500 N/m in the normal direction. A human stiffness in the range of 20-500 N/m was anticipated. In the ‘robot and human’ case the human stiffness comprised 1% - 25% of the stiffness. In this case, the observed variation of the robot-only estimates of 20% are problematic. They could cause errors equal to or greater than the anticipated human stiffness. In practice, when the ‘robot-only’ values were subtracted from the ‘robot and human’ estimates, the ‘human-only’ estimate was negative for a substantial number of configuration bins. This was not anticipated. We suspect this error resulted from the hysteresis in the InMotion controller.

### 4.4.4 Frequency Separation

In the study of the motor cortex, the firing rate of neurons is modulated to control action (e.g. muscle activation). The key for this control strategy to work is that, a substantial separation exists between the frequency content of the neural firing rate and the frequency at which the firing rate is modulated. In this case, there exists a frequency range over which the observable behavior is locally stationary, and that frequency range does not include the modulation frequencies associated with muscle activation [Kandel et al., 2013].

In the virtual crank-turning experiments, the subject moved at a rate of approximately one crank revolution per 13.33 seconds. This was extremely slow, slow enough to be considered quasi-static. A priori, it was not obvious that a similar frequency

separation, comparable to that of neural firing rate, would not exist between the interactive dynamics  $Z\{\cdot\}$  and forward-path dynamics  $x_0$ .

The frequency content of Brownian noise increases as frequency decreases. Given the presence of Brownian noise in the forward-path dynamics  $x_0$  the configuration-based ensemble method requires a high-pass filter pre-processing step. The filter cut-off frequency must be below the resonant frequency of the interactive dynamics. The interactive dynamics of the arm have a low resonance frequency ( $\sim 0.5$ -2 Hz in the work-space of the crank). The observed Brownian noise has substantial signal power in this same frequency range. This makes disentangling forward and interactive dynamics a challenge and may account for the problems with the configuration-based ensemble method, when applied in the tangential direction. This challenge is consistent with that observed in the power spectrum plotted in the lower right of Figure 4-12. In this figure the tangential displacement (red) approaches the amplitude of the tangential perturbation (purple) at low frequencies.

One solution to this problem is to add an additional stabilizing impedance. This added impedance could eliminate drift and increase the natural frequency. A higher natural frequency would make frequency separation easier. This was the case in the normal direction. In part this may be the reason that other studies have added a stabilizing stiffness during impedance estimation [Lee and Hogan, 2015, Palazzolo et al., 2007].

This suggest that looking at short-duration discrete tasks might avoid the drift problem. The ball-catching work of Lacquaniti et al. [Lacquaniti et al., 1983], the single-joint motion work of Bennet et al. [Bennett et al., 1992], and the point-to-point reaching work of Burdett et al. [Burdet et al., 2001] investigated discrete tasks shorter than 1.5 seconds. While continuous crank-turning makes it possible to collect the large number of cycles required for the ensemble methods, in our experiment one crank cycle took 13.33 seconds. With this duration the drift is substantial. Furthermore, even if a cycle took less than 1.5 seconds, the crank turning was not discrete. Continuous crank turning does not have a clear start and end of each cycle.

## 4.5 Conclusions

In this chapter, a configuration-based ensemble method was presented and tested on hardware and in simulation. The configuration-based ensemble method was able to estimate impedance in the normal direction with an error of less than 5% in simulation and 20% in a ‘robot-only’ experiment. However, hysteresis was observed in the robot control. For this reason conclusions about human impedance could not be drawn. If this error in control could be alleviated, the estimates in the robot-only case might be improved. With an improvement in the robot-only estimate it may allow the robot-only dynamics to be subtracted from that of the combined ‘human and robot’. This process would yield an estimate of the human-only impedance during a continuous physical interaction task – one with substantial force and motion.

Furthermore, a key observation was the lack of frequency separation between impedance and the zero-force trajectory. This lack of separation appears to have clear structure, consistent with Brownian noise in the zero-force trajectory. The lack of frequency separation poses a fundamental challenge to the identification of the impedance in the tangential direction.



# Chapter 5

## Exploiting Redundancy to Facilitate Physical Interaction

This chapter is an adapted version of [Hermus et al., 2022] published in IEEE Transactions on Robotics. This work was done in collaboration with Johannes Lackner and David Verdi.

### 5.1 Introduction

The human arm was one motivation for many of the recently developed seven-degree-of-freedom robotic manipulators (e.g. Kuka LBR iiwa, ABB Yumi - IRB 14000, Franka Emika, and Rethink Robotics Sawyer). They provide an additional kinematic degree of freedom during the performance of any end-effector task. The work reported in this paper investigated approaches to manage this redundancy, not only during free motion but also in tasks which involve forceful physical interaction. Remarkably, we found that with a sufficient excess of robot degrees of freedom over task degrees of freedom, a superposition of simple impedances performed as well as more complex null-space projection methods.

### 5.1.1 Managing Redundancy

One way to approach the control of a robot with many degrees of freedom is to express the desired robot behavior in the space of its end-effector actions<sup>1</sup>. This representation is bounded by a maximum number of independent variables ( $m \leq 6$ ). If the differential map  $\mathbf{J}(\mathbf{q}) \in \mathbb{R}^{m \times n}$  (Jacobian) from configuration variables  $\mathbf{q} \in \mathbb{R}^n$  to end-effector variables is known and the desired end-effector behavior can be expressed as a force  $\mathbf{f} \in \mathbb{R}^m$ , a unique map to joint torques  $\boldsymbol{\tau} \in \mathbb{R}^n$  will always exist:

$$\boldsymbol{\tau} = \mathbf{J}(\mathbf{q})^T \mathbf{f}. \quad (5.1)$$

This is a beneficial feature of torque-controlled robots, since eq. (5.1) also holds for kinematically redundant robots with  $n > m$ .

Finding the end-effector forces  $\mathbf{f}$  that are balanced by a given set of joint torques  $\boldsymbol{\tau}_f \in \mathbb{R}^n$  represents an optimization problem that may be solved by a generalized inverse of  $\mathbf{J}(\mathbf{q})^T$ :

$$\mathbf{f} = (\mathbf{J}(\mathbf{q})^T)^\# \boldsymbol{\tau}_f, \quad (5.2)$$

If  $n > m$ , a nullspace exists in  $(\mathbf{J}(\mathbf{q})^T)^\# \in \mathbb{R}^{m \times n}$ . This means that the end-effector forces can be balanced with infinitely many different joint torque solutions. The kernel of the optimization is the weighting matrix  $\mathbf{W} \in \mathbb{R}^{n \times n}$ :

$$\mathbf{J}(\mathbf{q})^\# = \mathbf{W}^{-1} \mathbf{J}(\mathbf{q})^T (\mathbf{J}(\mathbf{q}) \mathbf{W}^{-1} \mathbf{J}(\mathbf{q})^T)^{-1}. \quad (5.3)$$

Accordingly,  $\mathbf{J}(\mathbf{q})^\#$  yields a joint torque that minimizes the quadratic cost

$$g(\dot{\mathbf{q}}) = \frac{1}{2} \dot{\mathbf{q}}^T \mathbf{W} \dot{\mathbf{q}}. \quad (5.4)$$

The projector in the nullspace of  $(\mathbf{J}(\mathbf{q})^T)^\#$  can be expressed by

$$\mathbf{N}_\tau = \mathbf{I} - \mathbf{J}(\mathbf{q})^T (\mathbf{J}(\mathbf{q})^T)^\#, \quad (5.5)$$

---

<sup>1</sup>Also called task space or work space.

where  $\mathbf{N}_\tau \in \mathbb{R}^{n \times n}$ . All torques  $\boldsymbol{\tau}_{\text{any}} \in \mathbb{R}^n$  that are projected into this nullspace do not interfere with end-effector forces of higher priority and can be used for additional tasks [Antonelli et al., 2009, Dietrich et al., 2012, Dietrich et al., 2013, Siciliano and Slotine, 1991], e.g. to avoid obstacles or joint limits:

$$\boldsymbol{\tau} = \mathbf{J}(\mathbf{q})^T \mathbf{f} + \mathbf{N}_\tau \boldsymbol{\tau}_{\text{any}}. \quad (5.6)$$

Note that  $\boldsymbol{\tau}_{\text{any}}$  can incorporate further projections in the nullspace of lower-priority tasks. In this way arbitrarily many task levels can be produced with either the successive [Dietrich et al., 2012] or augmented [Siciliano and Slotine, 1991, Sentis and Khatib, 2005] methods. The lowest priority level is often chosen to be a joint damper to avoid oscillations due to nullspace motions [Khatib, 1987]. The feasibility of these task levels depends on the dimension of the nullspace, i.e. a one-dimensional nullspace only allows projection of a one-dimensional task. Hence, theoretically, a robot with  $n \gg m$  degrees of freedom is capable of accomplishing multiple tasks, without disturbing the main task.

The literature on redundancy resolution is predominantly concerned with nullspace projection approaches [Liegeois, 1977, Oda et al., 1995, Oriolo, 1994, Albu-Schaffer and Hirzinger, 2002, Albu-Schaffer et al., 2003, Ott, 2008, Ott et al., 2010, Peters et al., 2008, Sentis and Khatib, 2005]. A general overview of nullspace projections was presented in [Dietrich et al., 2015]. Implementations of hierarchical nullspace-projection-based control have been applied to tasks which involve contact [Siciliano and Slotine, 1991, Dietrich et al., 2012, Dietrich, 2016, Platt et al., 2011, Sadeghian et al., 2012, Sadeghian et al., 2013, Sadeghian et al., 2014, Ficuciello et al., 2015, Lin et al., 2015] and systems with multiple contact points [Park and Khatib, 2008, Dietrich et al., 2013, Henze et al., 2017]. Problems with instabilities have been discussed in [Klein and Huang, 1983] and [Hollerbach and Ki Suh, 1987]. This led to the development of conservative nullspace projection methods [Baillieul, 1985, Hu et al., 1995]. Stable nullspace projection methods have also been developed for mobile robot platforms [Antonelli et al., 2005, Antonelli et al., 2006, Antonelli et al., 2008, Antonelli

et al., 2009] and to cope with velocity actuator saturation [Arrichiello et al., 2010]. Energy tank methods have been applied to render nullspace projection methods passive [Dietrich et al., 2016, Dietrich et al., 2017].

## Mechanical Impedance Superposition

While many approaches have been developed to ensure stability when nullspace projections are employed, neither the nullspace projector nor the stabilizing corrections would be required if each controller were formulated as an energetically passive impedance. A simple—even naive—solution to control the desired dynamic robot behavior can be achieved by assigning a set of impedances, which can be visualized as a spring-damper system [Hogan, 1984a, Hogan, 1985a, Hogan, 1985b, Hogan, 1985c]. These impedances can be applied in end-effector space and in joint-space. Even if these impedances are non-linear, they can be superimposed:

$$\boldsymbol{\tau} = \sum_{i=1}^k \mathbf{J}(\mathbf{q})_i^T Z\{\mathbf{x}\}_i + \sum_{j=1}^l Z\{\mathbf{q}\}_j, \quad (5.7)$$

with  $k$  end-effector impedances  $Z\{\mathbf{x}\}_i : \mathbb{R}^m \rightarrow \mathbb{R}^m$  and  $l$  joint impedances  $Z\{\mathbf{q}\}_j : \mathbb{R}^n \rightarrow \mathbb{R}^n$ . If each component impedance is passive, their sum is energetically passive, and since no inverse kinematics are needed, this approach works at kinematic singularities.

If  $n > m$ , the end-effector impedance does not control the nullspace of  $\mathbf{J}(\mathbf{q})$ . To achieve predictable joint motion, a full-rank set of joint-space impedances can be assigned. These impedances push the robot toward a desired configuration  $\mathbf{q}_0$ . However, the joint-space impedance may conflict with the end-effector impedance—which is usually the task of interest—except in the rare cases when the end-effector position corresponds to that configuration,  $\mathbf{q}_0$ . This may be the reason why nullspace projection approaches have rarely [Verdi, 2019] been compared with impedance superposition: in theory nullspace projection approaches should eliminate the end-effector task error, while the simple superposition of impedance controllers may result in task conflict. However, as we show below, due to imperfections in a robot’s kinematic and



dynamic models, in practice nullspace projection may also cause task-space disruption.

One novel aspect of the work reported here is that it assessed task-space errors due to implementing both approaches on hardware. For nullspace projections, we assigned different weighting matrices. The choice of nullspace weighting matrix has been discussed in several contexts: to generate favorable kinematic behavior [Yoshikawa, 1985], to prioritize different motions [Whitney, 1969], to perform motion control with joint constraints [Flacco et al., 2012, Flacco et al., 2015], and to ensure dynamic consistency [Khatib, 1987]. In this work, we compare nullspace projection methods to the superposition of mechanical impedance in both unconstrained motion and during physical interaction with a constraint.

### 5.1.2 Factors that Influence Controller Design

To find an appropriate control approach for a given robot task, many factors should be considered. Some factors are determined by the robotic system and the environment with which it interacts. Another factor is the available information about the robot model, i.e. kinematic and/or dynamic data. Lastly, the desired task may or may not be achievable by the robot. We took several of these factors into account when comparing nullspace projection methods with impedance superposition.

#### Environmental Factors

A robot is influenced by its own controller and the dynamics of the environment which acts on it. Since it is impossible to have a perfect model (or in many cases even a competent model) of the interacting environment [Stramigioli, 2001], most algorithms solely concentrate on the robot’s controller. With an impedance controller, a desired interactive dynamic behavior can be implemented (though perhaps imperfectly). To specify how well a robot performs its tasks both in and out of contact, quantitative measures are required. In unconstrained motion, we assessed the difference between the desired and actual position and orientation. In tasks involving

continuous physical interaction, the deviations between desired and actual forces were used as a quantitative measure of controller performance.

## Nullspace Dimension

Most industrial robots are serial kinematic chains with six degrees of freedom. Since it is desirable to describe the robot task in end-effector space, these robots have the benefit that the mapping between end-effector space and joint-space is bijective. However, in practice many tasks require fewer degrees of freedom. For example, consider a robot with a welding gun: rotation of the gun about its long axis has no influence on task performance. Thus even a six-degree-of-freedom robot is redundant with respect to some tasks. Of course, serial kinematic chain robots with seven degrees of freedom or more always exhibit a nullspace, but the dimension of the nullspace depends on the end-effector task. Nullspace projection methods can take advantage of redundancy by assigning additional tasks in the nullspace of the main task. In the work reported here, we investigated whether there could be advantages to decreasing the task dimensions and thereby increasing the nullspace dimensions.

## Weighting Matrix

Nullspace projection methods require a weighting matrix. This weighting matrix defines the cost function minimized in the optimization, as seen in eq. (5.4). Even though any positive definite matrix can be used, without a meaningful choice, physical insight may be lost [Lachner et al., 2020]. A list of some possible options can be seen in table 5.1. Two common choices are  $\mathbf{W} = \mathbf{I}$  and  $\mathbf{W} = \mathbf{M}(\mathbf{q})$ . The former yields the least-norm solution [Penrose, 1955] and the latter minimizes the kinetic energy [Khatib, 1987, Bruyninckx and Khatib, 2000] produced by nullspace motion. The dynamic consistency provided by the mass matrix is superior [Dietrich et al., 2015, Chang and Khatib, 1995] especially when inertial dynamics are significant. In theory, this is the only nullspace projector that does not produce accelerations that interfere with the main task. Moreover, it is the only projector that does not inject energy during nullspace motion and should therefore have superior stability properties

[Bruyninckx and Khatib, 2000]. However, in practice, without a perfect model of the mass matrix, other choices may be better [Dietrich et al., 2015, Albu-Schaffer et al., 2003, Nakanishi et al., 2008, Peters et al., 2008]. Often when working with low-cost robots, the mass matrix is not well known. In addition, poorly-modeled joint friction and motor rotational inertia, amplified through a gear transmission, may dominate the dynamic response of the robot [Hosford, 2016]. Yet another reasonable choice is  $\mathbf{W} = \mathbf{B}_q$ . By using the joint-space damping matrix  $\mathbf{B}_q \in \mathbb{R}^{n \times n}$  the nullspace motion with least energy dissipation is produced. The choice  $\mathbf{W} = \mathbf{K}_q$  produces the solution that minimizes potential energy at equilibrium. We acknowledge that there are many other weighting matrix choices not considered here, including [Yoshikawa, 1985, Whitney, 1969, Flacco et al., 2012, Flacco et al., 2015]. For notational convenience we also define  $\mathbf{W}^{-1} = \mathbf{0} \in \mathbb{R}^{n \times n}$ , a matrix with only zero entries. With this, the nullspace projection matrix is equal to the identity matrix, corresponding to a superposition of all task levels.

## Inertial Dynamics

A robot’s performance is affected by its inertial dynamics. If the robot moves slowly enough, inertial dynamics can be neglected and the task can be considered quasi-static. At fast speeds, however, inertial dynamics, damping, and stiffness must all be considered. The choice of task execution speed (slow/quasi-static vs. fast/dynamic) was expected to have a substantial impact on task performance for certain choices of nullspace projection weighting matrices.

## Relative Impedance Magnitudes

Using impedance superposition, a large joint-space impedance will result in a substantial conflict with any end-effector task, while a smaller joint-space impedance will evoke a lesser conflict. This prompted the question: if a small joint-space impedance, sufficient to ‘manage the redundancy’<sup>2</sup> is superimposed, how large will the task dis-

---

<sup>2</sup>This was defined operationally as an impedance as small as possible, but still capable of restoring the robot to a position near its nominal configuration after a large null-space position disturbance.

ruption be? When comparing mechanical impedance superposition and nullspace projection methods, the magnitude of the joint-space stiffness was taken into account.

Table 5.1: Optimization criteria for nullspace projectors.

<b>Weighting matrix</b>	<b>Cost</b>	<b>Description</b>
$\mathbf{W}^{-1} = \mathbf{0}$	–	Impedance Superposition
$\mathbf{W} = \mathbf{I}$	$\frac{1}{2}\dot{\mathbf{q}}^T\dot{\mathbf{q}}$	Least Velocity Norm
$\mathbf{W} = \mathbf{M}(\mathbf{q})$	$\frac{1}{2}\dot{\mathbf{q}}^T\mathbf{M}(\mathbf{q})\dot{\mathbf{q}}$	Least Kinetic Energy
$\mathbf{W} = \mathbf{B}_q$	$\frac{1}{2}\dot{\mathbf{q}}^T\mathbf{B}_q\dot{\mathbf{q}}$	Least Energy Dissipation
$\mathbf{W} = \mathbf{K}_q$	$\frac{1}{2}\Delta\mathbf{q}^T\mathbf{K}_q\Delta\mathbf{q}$	Least Potential Energy

### 5.1.3 Summary

The principal aim of this study was to quantify and compare the performance of mechanical impedance superposition and nullspace projection methods to manage redundancy on real hardware in practice. Quantitative assessment was performed during both unconstrained and constrained motion. A secondary aim was to understand and quantify how the dimension of the nullspace—the wealth of degrees of freedom—influenced performance. Our results show that for a nullspace of sufficient dimension, the task conflict from simple impedance superposition was comparable to that of all nullspace projection methods.

## 5.2 Methods

The goal of these experiments was to examine the behavior of a redundant robot placed under an end-effector impedance controller (Task 1), along with a joint-space impedance controller (Task 2). Task 2 was either superimposed directly ( $\mathbf{W}^{-1} = \mathbf{0}$ ) or projected into the nullspace of Task 1 using each of the four weighting matrices listed in Table 5.1. The experiments investigated both unconstrained motion and forceful physical contact with a circular constraint, specified in Task 1. For all five

weighting matrices, performance was quantified for different nullspace dimensions, joint stiffnesses, and task speeds.

### 5.2.1 Experimental Setup

All experiments were conducted with a seven degree of freedom torque-controlled KUKA LBR iiwa R800 (LBR) [Schreiber et al., 2010]. To facilitate measurement, in each experiment the LBR was coupled to a customized InMotion2 Shoulder-Elbow robot (Interactive Motion Technologies Inc.) via a U-joint and bearing connection. The InMotion is a highly back-drivable light-weight torque controlled  $x - y$  planar robot, which was designed for stroke rehabilitation and human motion research. The custom control system for the InMotion robot was implemented on a CompactRIO 9034 controller, with low-level functionality implemented at 2 kHz on a CompactRIO FPGA, and high-level functionality implemented at 1 kHz on a CompactRIO real-time processor. InMotion joint positions were measured by a 16-bit/rev encoder and interaction forces were measured at the end-effector of the InMotion robot using an ATI Gamma force/torque transducer [Hogan et al., 1995, Thorup, 2018]. The experimentally coupled robots are shown in Figure 5-1.

The U-joint (Neapco Components, Pottstown, PA) and bearing connection enabled a  $\pm 45^\circ$  rotational range of motion about the  $x$  and  $y$  axes. Rotation around the  $z$  axis was facilitated by an ultra-low-friction dry-running sleeve bearing, which also enabled translation along the  $z$  axis. Thus, if the center of the U-joint is viewed as the kinematic coupling point, the two robots are constrained relative to each other translationally in the  $x - y$  plane, but not constrained translationally along the  $z$  axis. Furthermore, the U-joint decoupled rotations about all three axes.

Once the two robots were coupled together with the U-joint and bearing system, the total amount of free-play or backlash in the coupling connection was quantified. The brakes were applied on the LBR, nominally fixing it rigidly in space. The handle of the InMotion robot was then lightly perturbed by hand in several directions, and the resulting handle displacements were measured using the InMotion encoders. It was found that the InMotion handle could undergo a displacement of  $\pm 1.5$  mm in the

$x - y$  plane without applying appreciable forces to the LBR.




Figure 5-1: Experimental setup. (Left) Planar depiction of the InMotion Robot and reference frame. The dotted circle denotes the zero-force trajectory of the InMotion robot and the dashed circle denotes the zero-force trajectory of the LBR. (Right) The LBR was coupled with a U-joint and sleeve bearing to the handle of the InMotion robot. The inset in the upper left illustrates the rotational degrees of freedom allowed by the U-joint and sleeve bearing.

## 5.2.2 Impedance Controller

The end-effector and joint-space impedance controllers were implemented on the LBR using the KUKA Fast Research Interface (FRI), via an external PC, with torque commands computed at 200 Hz. The FRI friction and gravity compensation was active throughout all of the experiments. The analytical Jacobian matrix  $\mathbf{J}(\mathbf{q}) \in \mathbb{R}^{6 \times n}$  of the robot was denoted by:

$$\mathbf{J}(\mathbf{q}) = \begin{bmatrix} \mathbf{J}(\mathbf{q})_x \\ \mathbf{J}(\mathbf{q})_\theta \end{bmatrix}. \quad (5.8)$$

Here,  $\mathbf{J}(\mathbf{q})_x \in \mathbb{R}^{3 \times n}$  maps the joint velocities  $\dot{\mathbf{q}} \in \mathbb{R}^n$  to translational end-effector velocities and  $\mathbf{J}(\mathbf{q})_\theta \in \mathbb{R}^{3 \times n}$  maps  $\dot{\mathbf{q}}$  to rotational end-effector velocities. In order to define the controller, three reference frames were defined: a fixed base frame denoted  $\Sigma_b$  (this is displayed in Figure 5-1 as  $x$ ,  $y$ , and  $z$ ), a moving frame fixed to the center of

the U-joint (which was taken to be the robot’s end-effector), denoted  $\Sigma_e$ , and a frame moving with the LBR robot’s zero-force trajectory, denoted  $\Sigma_0$ . Both,  $\mathbf{J}(\mathbf{q})_x$  and  $\mathbf{J}(\mathbf{q})_\theta$  were expressed with respect to the end-effector frame  $\Sigma_e$ . For the end-effector translational impedance controller, the desired control torque  $\boldsymbol{\tau}_x \in \mathbb{R}^n$  was computed by:

$$\boldsymbol{\tau}_x = \mathbf{J}(\mathbf{q})_x^T \left( \mathbf{K}_x(\mathbf{x}_0 - \mathbf{x}) - \mathbf{B}_x \dot{\mathbf{x}} \right). \quad (5.9)$$

$\boldsymbol{\tau}_x \in \mathbb{R}^n$  described a translational spring-damper system with linear stiffness  $\mathbf{K}_x \in \mathbb{R}^{3 \times 3}$  and linear damping  $\mathbf{B}_x \in \mathbb{R}^{3 \times 3}$ . Both  $\mathbf{K}_x$  and  $\mathbf{B}_x$  were chosen to be diagonal matrices. The virtual spring was attached between the  $\Sigma_e$  and  $\Sigma_0$  frame. The position of the end-effector  $\mathbf{x} \in \mathbb{R}^3$  and the zero-force position  $\mathbf{x}_0 \in \mathbb{R}^3$  were represented in the base frame coordinates  $\Sigma_b$ . The zero force trajectory,  $\mathbf{x}_0$ , moved with constant speed around a circular path with a radius of 0.1 m. For the end-effector rotational impedance controller, the desired control torque  $\boldsymbol{\tau}_\theta \in \mathbb{R}^n$  was computed with:

$$\boldsymbol{\tau}_\theta = \mathbf{J}(\mathbf{q})_\theta^T \left( \mathbf{K}_\theta \hat{\mathbf{u}}_0 \theta_0 - \mathbf{B}_\theta \dot{\theta} \right). \quad (5.10)$$

The rotational torque  $\boldsymbol{\tau}_\theta$  aligned the axes of frame  $\Sigma_e$  and moving frame  $\Sigma_0$ . The rotation between  $\Sigma_e$  and  $\Sigma_0$  was expressed by the rotation matrix  ${}^0\mathbf{R}_e \in SO(3)$ . To calculate the rotational torque  $\boldsymbol{\tau}_\theta \in \mathbb{R}^n$ ,  ${}^0\mathbf{R}_e$  was converted to axis-angle representation, with unit axis  $\hat{\mathbf{u}}_0 \in \mathbb{R}^3$  and angle  $\theta_0 \in \mathbb{R}$  [Luh et al., 1980, Natale, 2003]. Thus, a virtual rotational spring with rotational stiffness  $\mathbf{K}_\theta \in \mathbb{R}^{3 \times 3}$  was attached around  $\hat{\mathbf{u}}_0$  to rotate about  $\theta_0$ . This implementation was further discussed in detail in Appendix I. The rotational velocity  $\dot{\theta} \in \mathbb{R}^3$  was damped with dissipating element  $\mathbf{B}_\theta \in \mathbb{R}^{3 \times 3}$ . Note that all vectors and matrices of Equation 5.9 and 5.10 were expressed in  $\Sigma_b$ . The stiffness in the  $z$  direction was chosen to ensure the robot maintained contact with the sleeve bearing in all trials. Finally, the translational and rotational end-effector torques were combined,

$$\boldsymbol{\tau}_e = \boldsymbol{\tau}_x + \boldsymbol{\tau}_\theta. \quad (5.11)$$

A diagonal  $\mathbf{K}_\theta$  and  $\mathbf{B}_\theta$  were chosen to approximate a constant damping ratio for

each rotational direction in the end-effector impedance controller. The damping ratio along each of the three rotational directions was roughly approximated as:

$$\zeta_i = \frac{b_i}{2m_i\sqrt{k_i/m_i}}. \quad (5.12)$$

Here,  $b_i$  and  $k_i$  represent the  $i^{\text{th}}$  diagonal elements of  $\mathbf{B}_\theta$  and  $\mathbf{K}_\theta$  respectively, and  $m_i$  is one of the corresponding three diagonal elements of the rotational end-effector mass matrix  $\mathbf{\Lambda}_\theta \in \mathbb{R}^{3 \times 3}$  [Khatib, 1995], given by:

$$\mathbf{\Lambda}_\theta = (\mathbf{J}_\theta(\mathbf{q})(\mathbf{M}(\mathbf{q}))^{-1}\mathbf{J}_\theta(\mathbf{q})^T)^{-1}, \quad (5.13)$$

where  $\mathbf{M}(\mathbf{q}) \in \mathbb{R}^{n \times n}$  is the manipulator mass matrix. Given our choice for the values of  $k_i$ , the values of  $b_i$  were chosen to yield  $\zeta_i \approx 0.4$ . This was a reasonable balance between an undamped ( $\zeta = 0$ ) and critically damped ( $\zeta = 1$ ) behaviour.

For the joint-space impedance controller, the commanded torque  $\boldsymbol{\tau}_q \in \mathbb{R}^n$  was expressed by

$$\boldsymbol{\tau}_q = \mathbf{K}_q(\mathbf{q}_0 - \mathbf{q}) - \mathbf{B}_q\dot{\mathbf{q}}, \quad (5.14)$$

with joint-space stiffness  $\mathbf{K}_q \in \mathbb{R}^{n \times n}$  and joint-space damping  $\mathbf{B}_q \in \mathbb{R}^{n \times n}$ . The nominal joint position  $\mathbf{q}_0 \in \mathbb{R}^n$  was constant throughout the trial and corresponded to a robot end-effector position at the origin of  $\Sigma_b$  with a  $15^\circ$  rotation about both the  $x$  and  $y$  axes. This configuration was chosen so that the joint-space impedance controller always conflicted with the end-effector impedance controller. The nominal joint-space pose was  $\mathbf{q}_0 = [-56.16, -47.4, 87.7, 83.2, -42.1, -71.9, 28.8]^T$  (degrees) and is shown in Figure 5-2.  $\mathbf{B}_q$  was chosen as a function of  $\mathbf{K}_q$  to yield an approximate damping ratio of 0.32 to 0.42 along each joint. This was done in a manner similar to (5.12), but with  $k_i$ ,  $b_i$ , and  $m_i$  being each of the seven diagonal entries of  $\mathbf{K}_q$ ,  $\mathbf{B}_q$ , and  $\mathbf{M}(\mathbf{q})$  respectively.



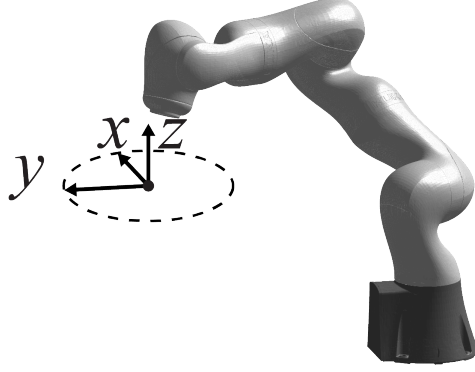


Figure 5-2: The nominal joint-space position ( $\mathbf{q}_0$ ) used in the experiments.

Finally, to vary the nullspace dimension, we composed the sub-controllers in two ways: using the complete 6D rotational and translational end-effector torque  $\tau_e$ , yielding a 1D nullspace; or using only the 3D translational end-effector torque  $\tau_x$ , yielding a 4D nullspace. The nullspace projectors were modified as follows in the 1D and 4D nullspace cases:

$$\tau_{1D} = \underbrace{\tau_e}_{\text{Task 1}} + \underbrace{\mathbf{N}_{1D}\tau_q}_{\text{Task 2}} \quad (5.15a)$$

and

$$\tau_{4D} = \underbrace{\tau_x}_{\text{Task 1}} + \underbrace{\mathbf{N}_{4D}\tau_q}_{\text{Task 2}} \quad (5.15b)$$

Here, the nullspace projector  $\mathbf{N}_{1D} \in \mathbb{R}^{n \times n}$  is defined using the complete Jacobian matrix by

$$\mathbf{N}_{1D} = \mathbf{I} - \mathbf{J}(q)^T (\mathbf{J}(q)^\#)^T \quad (5.16a)$$

to project the torques of the joint-space impedance controller (Task 2) into the nullspace of the six dimensional end-effector impedance controller (Task 1). The nullspace projector  $\mathbf{N}_{4D} \in \mathbb{R}^{n \times n}$  is defined using only the translational component of the Jacobian matrix by

$$\mathbf{N}_{4D} = \mathbf{I} - \mathbf{J}(q)_x^T (\mathbf{J}(q)_x^\#)^T. \quad (5.16b)$$

Likewise, it projects the torques of the joint-space impedance controller (Task 2) into the nullspace of the three dimensional end-effector impedance controller (Task 1).

Table 5.2: Controller Parameters. All non-diagonal stiffness and damping terms were zero.

Variable	Values	Units
$\mathbf{K}_x$	diag([1800, 1800, 2000])	N/m
$\mathbf{B}_x$	diag([43.3, 31.68, 37.19])	N-s/m
$\mathbf{K}_\theta$	diag([260, 260, 100])	N/rad
$\mathbf{B}_\theta$	diag([3.0, 3.4, 1.6])	N-s/rad
$\mathbf{K}_q$	diag([10, 10, 10, 10, 5, 5, 1])	N-m/rad
$\mathbf{B}_q$	diag([2.5, 3.6, 2.1, 2.1, 0.3, 0.2, 0.1])	N-m-s/rad
$\mathbf{K}_{q,low,1D}$	$0.1\mathbf{K}_q$	N-m/rad
$\mathbf{B}_{q,low,1D}$	$\sqrt{0.1}\mathbf{B}_q$	N-m/rad
$\mathbf{K}_{q,low,4D}$	$0.3\mathbf{K}_q$	N-m/rad
$\mathbf{B}_{q,low,4D}$	$\sqrt{0.3}\mathbf{B}_q$	N-m/rad

### 5.2.3 Test Conditions

The performance of impedance superposition ( $\mathbf{W}^{-1} = \mathbf{0}$ ) and a conventional choice of nullspace projector ( $\mathbf{W} = \mathbf{I}$ ) was evaluated under several test conditions. First, each condition was tested when the circular constraint was off and when it was on,

denoted ‘unconstrained’ and ‘constrained’, respectively. Second, in each case the nullspace dimension was either one or four. Third, the influence of the LBR controller parameters was investigated. Fourth, to compare nullspace projection methods, the effect of weighting matrix was evaluated. Each of these conditions is described in detail in this section; a summary is presented in Table 5.3.

### **Circular Constraint**

In all conditions, the LBR was coupled to the InMotion with a U-joint and bearing system described above. In the unconstrained trials the InMotion robot was turned off; it remained passive and highly back-drivable, but still coupled to the LBR. In the constrained trials, the InMotion robot enforced a circular constraint (radius = 0.08 m). This constraint was enforced by an impedance controller, with a normal stiffness of 2,500 N/m, and a normal damping of 40 Nm/s. The accuracy of rendering the constraint was verified by kinematic and end-point force/torque measurements obtained from the InMotion robot. The differing radii of the LBR’s zero-force position (0.1 m) and the constrained circular path enforced by the InMotion robot ensured that the observed behavior occurred under conditions of significant forceful contact. In all experiments, the InMotion encoders and force transducer were used to record the planar interaction-point positions and interaction forces presented herein.

### **Nullspace Dimension**

For all experiments, Task 2, the nullspace-projected joint-space impedance controller, was always active. However, the number of dimensions in Task 1 was varied. In some experiments, only the end-effector translational impedance controller was applied (Equation 5.16b). This meant that the end-effector task was three-dimensional, resulting in a four-dimensional nullspace. For other experiments, both the translational and rotational controllers of Equation 5.16a were applied. This resulted in a six-dimensional end-effector task, leaving a one-dimensional nullspace.

## Weighting Matrix

To quantify performance differences due to the choice of weighting matrix  $\mathbf{W}$ , five different weighting matrices were tested:  $\mathbf{0}$ ,  $\mathbf{I}$ ,  $\mathbf{M}(\mathbf{q})$ ,  $\mathbf{B}_q$ , and  $\mathbf{K}_q$ . A weighting matrix choice of  $\mathbf{W}^{-1} = \mathbf{0}$  which results in  $\mathbf{N} = \mathbf{I}$ ; this is the case of simple superimposition of the end-effector and joint-space tasks, without a nullspace projector. We initially compared that case to a representative weighting matrix,  $\mathbf{W} = \mathbf{I}$ , which corresponds to the well-known Moore-Penrose pseudo-inverse [Penrose, 1955]. Subsequently, performance with the other weighting matrices was quantified.

## Controller

To understand the influence of robot inertial dynamics and the joint-space impedance controller used to manage redundancy, three conditions were examined: moderate joint-space stiffness and slow speed (13 sec/rev), which is hereafter referred to as a ‘standard’ condition; moderate joint-space stiffness and fast speed (4 sec/rev); and low joint-space stiffness and slow speed. A period of 13 sec/rev was extremely slow; at this speed, all dynamic effects were negligible, the motion was quasi-static, and behavior was dominated by the controller stiffness. If joint-space stiffness is sufficiently large relative to task-space stiffness, managing redundancy using joint-space impedance will substantially interfere with the end-effector task. In this work the ‘moderate’ joint-space stiffness was chosen to be large enough to cause substantial deviation from the desired end-effector task (more than 10 cm when unconstrained and more than 10 N when constrained) when the impedances were superimposed (i.e. the  $\mathbf{0}$  projector case). The fast speed, a period of 4 sec/rev, was chosen to elicit significant dynamic effects due to the robot’s inertial dynamics. Quantitative analysis of the quasi-static vs. dynamic speeds is presented in Appendix G.

To test whether a nullspace projection is even required in the first place, a low joint-space stiffness condition was included, executed at the slow speed (13 sec/rev). This stiffness was chosen to be the smallest value that would ‘manage the redundancy’ of the robot within a single cycle. If a large joint-space position disturbance was

applied to the LBR arm, it would return to a set of joint angles near the nominal joint configuration within one crank-turning cycle. This low joint-space stiffness and damping differed when the nullspace dimension varied; the exact values used are reported in Table 5.2.

Table 5.3: Test conditions and notations in this paper.

Condition	Parameter	Notation
LBR		
Weighting Matrix $\mathbf{W}$	See methods	$\mathbf{0}, \mathbf{I}, \mathbf{M}(\mathbf{q}), \mathbf{B}_q, \mathbf{K}_q$
Nullspace Dimension	$n - m = 1$	1D
	$n - m = 4$	4D
Controller	$\mathbf{K}_q, 13 \text{ s/rev}$	Standard
	$\mathbf{K}_q, 4 \text{ s/rev}$	Fast Speed
	$\mathbf{K}_{q,low}, 13 \text{ s/rev}$	Low Stiffness
InMotion		
Circular constraint	–	Unconstrained
	radius 0.08 $m$	Constrained

#### 5.2.4 Data Analysis

In each trial the LBR completed two separate motions of three revolutions each. The abrupt engagement of the robot controller at the start of each motion induced transient behavior in the robot end-effector motion as task 1 was not critically damped. To eliminate possible transients from the data analysis, the first revolution in each of these trials was discarded.

In the standard condition, since the robot moved quasi-statically, dominated by stiffness, the robot was expected to closely follow the target position of the end-effector impedance controller when a nullspace projector was used. To quantify errors in position, the displacement normal to the closest point on the circle was computed,

$$\Delta x = 0.1 - r_n \tag{5.17}$$

where 0.1 m was the desired distance from the origin and  $r_n = \sqrt{x^2 + y^2}$  was distance from the origin to the actual robot position. This was the dependent measure for the unconstrained case.

In the constrained condition, the radius of the zero-force LBR path was  $x_0 = 0.1$  m, while the radius of the InMotion’s virtual constraint was 0.08 m. Thus the LBR robot was expected to move in a perfect circle at a radius determined by equilibrium between the stiffnesses of the LBR and InMotion robots. This required a constant normal force of 20.93 N to be exerted on the InMotion handle. Positive normal force denoted an outward-directed force exerted on the InMotion handle, away from the circular constraint center; conversely, negative normal force was directed inwards, toward the constraint center. Thus the dependent measure in the constrained condition was the deviation of the measured normal force from the expected normal force:

$$\Delta f = 20.93 - f_n. \tag{5.18}$$

For each of the dependent measures, the root-mean-squared-error (RMSE) was computed for four crank cycles. While this metric was expected to approach zero only in the quasi-static case, it was also applied to the fast trials, even though the dynamic effects were expected to be significant and result in a non-zero mean force error, e.g. due to centrifugal acceleration. Nevertheless, it remained a suitable metric with which to compare the behaviors of different projector choices and nullspace dimensions.

### 5.2.5 Statistical Analysis

The difference of RMSE between the superposition of mechanical impedance ( $\mathbf{W}^{-1} = \mathbf{0}$ ) and a standard nullspace projector ( $\mathbf{W} = \mathbf{I}$ ) was computed. This difference was denoted RMSE  $\mathbf{0} - \mathbf{I}$ . When unconstrained, the RMSE difference was computed in terms of position; when constrained, the RMSE difference was computed in terms of normal force. For both the unconstrained and constrained conditions, a two-way analysis of variance (ANOVA) was performed. The ANOVA assessed the effect of nullspace dimension (1D or 4D) and controller condition (moderate stiffness, slow; moderate stiffness, fast; or low stiffness, slow) on the RMSE difference. For both the unconstrained and constrained results, three post-hoc two-sample t-tests were performed to determine if a significant difference existed between the 1D and 4D conditions. Two post-hoc one-sample t-tests were performed to determine if impedance superposition was significantly different from nullspace projection in the constrained 4D nullspace standard and fast conditions.

In this work, statistical analysis aimed to assess differences between multiple conditions, each of which possessed more than one level. To avoid Type I errors (false positive) associated with performing multiple t-tests, ANOVA was first employed to determine if statistically significant differences between the means existed. The test statistic used by ANOVA is described by the F-distribution. The results present the F-statistic, its associated degrees of freedom, and the probability of a non-significant effect. The analysis used in this work is described in several standard texts, e.g. [Keppel and Wickens, 2004].

The differences of RMSE between the identity weighting matrix ( $\mathbf{W} = \mathbf{I}$ ) and all other weighting matrices ( $\mathbf{M}(\mathbf{q})$ ,  $\mathbf{B}$ , and  $\mathbf{K}$ ) were computed. These differences were denoted RMSE  $\mathbf{I} - \mathbf{W}$ . For each nullspace dimension and constraint condition, a two-way ANOVA was performed to assess the effect of nullspace weighting matrix ( $\mathbf{M}(\mathbf{q})$ ,  $\mathbf{B}$ , or  $\mathbf{K}$ ) and controller condition (moderate stiffness, slow; moderate stiffness, fast; or low stiffness, slow). Post-hoc one-sample t-tests were run to identify significant differences between in RMSE  $\mathbf{I} - \mathbf{W}$ .

## 5.3 Results

In this section, the unconstrained results in the standard condition (slow, moderate stiffness) are presented first. Next, constrained results in the standard condition are shown. Third, the results of fast, moderate stiffness trials are reported. Fourth, results with low joint-space stiffness and slow speed are presented. Finally, a comparison of different nullspace weighting matrices is reported.

### 5.3.1 Standard Condition, Unconstrained

In the unconstrained condition, the InMotion robot did not enforce a circular constraint. However, the InMotion robot remained passively coupled to the LBR with the U-joint and bearing system. With no constraint, we hypothesized that when the robot was operating with a nonzero nullspace projection weighting matrix, there would be no disruption of Task 1 from Task 2, regardless of whether the nullspace had one dimension or four. The trajectory was expected to be close to or even overlay the circular robot path  $\mathbf{x}_0$  (bold dashed line in Figure 5-3) and with  $\mathbf{W} = \mathbf{I}$ , this was indeed observed.



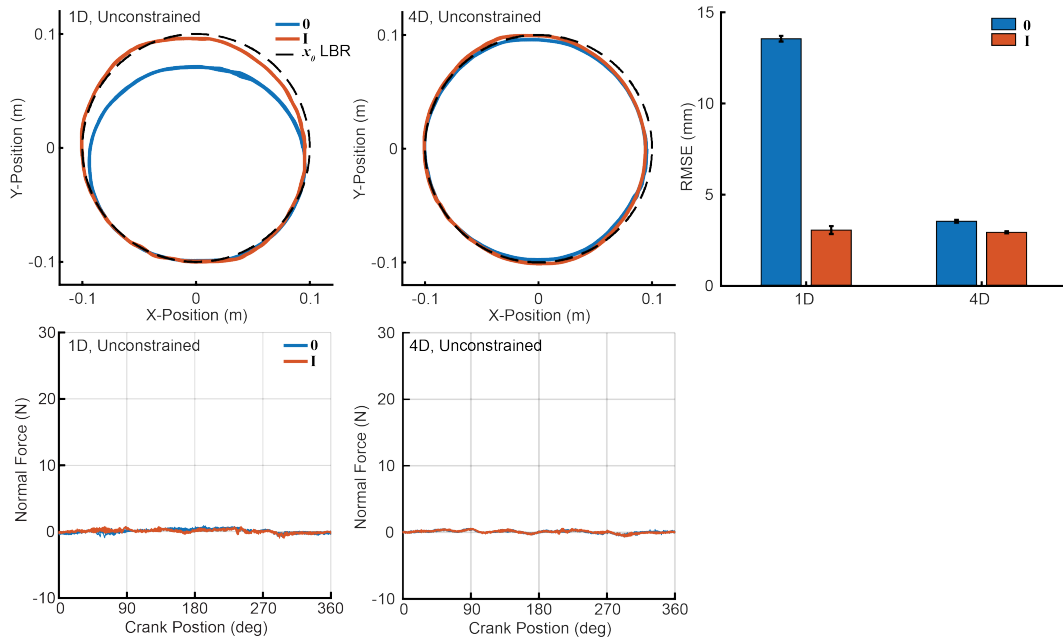


Figure 5-3: **Standard condition, unconstrained**: 1D nullspace (left column) and 4D nullspace (middle column). Trajectory plots (top left and top middle) and the zero-force path of the LBR (bold dashed line). Normal force plots (bottom left and bottom middle). Position RMSE (top right). The superposition of mechanical impedance is denoted by  $\mathbf{W}^{-1} = \mathbf{0}$  (blue) and the Moore–Penrose inverse denoted  $\mathbf{W} = \mathbf{I}$  (red). The superposition of mechanical impedances was substantially less disruptive with a 4D nullspace.

In the  $\mathbf{W}^{-1} = \mathbf{0}$  case, superimposing impedances was expected to result in conflict between Tasks 1 and 2, leading to significant tracking errors in Task 1. This was observed in the 1D nullspace condition (see Figure 5-3 top left). The  $\mathbf{W}^{-1} = \mathbf{0}$  case (blue line) substantially deviated from the circular trajectory of the LBR end-effector task, as indicated by the bold dashed line.

While it was expected that Task 2 would visibly conflict with Task 1 in the absence of a valid nullspace projector, one unexpected result was that when the dimension of the nullspace was increased from one to four (by removing the rotational impedance controller from Task 1), the task conflict was substantially reduced. This can be seen in Figure 5-3 (top middle).

To assess quantitative differences in the unconstrained trials, the RMSE of the robot position was computed. This provided a measure of the deviation from expected behavior as seen in Figure 5-3 (top right) and quantified the qualitative obser-

vations. Remarkably, for both the 1D and 4D nullspace conditions, there appeared to be no appreciable difference between the various non-zero projectors. In the 1D nullspace case, the zero projector (impedance superposition) introduced substantial task conflict, while in the 4D nullspace case, the RMSE it evoked was not appreciably higher than any of the other projectors. These differences are compared statistically in section 5.3.5.

### 5.3.2 Standard Condition, Constrained

In the constrained condition, both force and motion must be considered. In this experiment, the virtual constraint radius enforced by the InMotion was 0.08 *m* and the diameter of the zero-force robot path was 0.1 *m*. Thus, the LBR was expected to move along a constant radius circle between the InMotion (dotted) and LBR (dashed) lines in Figure 5-4 (top). In the standard trials, the motion was quasi-static, well within the stiffness-dominated regime, meaning that all dynamic effects were negligible (Appendix G). The displacement of the handle from its zero-force path was determined by the relative stiffness of the two robots and a constant normal force should have been exerted.

As expected, with a 1D nullspace the superposition of joint-space stiffness substantially disrupted the LBR task-space position and normal force (the solid blue line in the top two panels of Figure 5-4). The disruption of the task was sufficient for the robot occasionally to exert inward normal forces on the virtual constraint. Inward (compressive) normal forces exerted on a constraint surface are inherently destabilizing [Rancourt and Hogan, 2001]. This demonstrates that superposition of joint-space and task-space mechanical impedance may, in some cases, lead to static instability and potential safety concerns.

As with the unconstrained case, the superposition of mechanical impedance ( $\mathbf{W}^{-1} = \mathbf{0}$ ) performed substantially better with a 4D nullspace as seen in Figure 5-4. These differences are compared statistically in section 5.3.5.

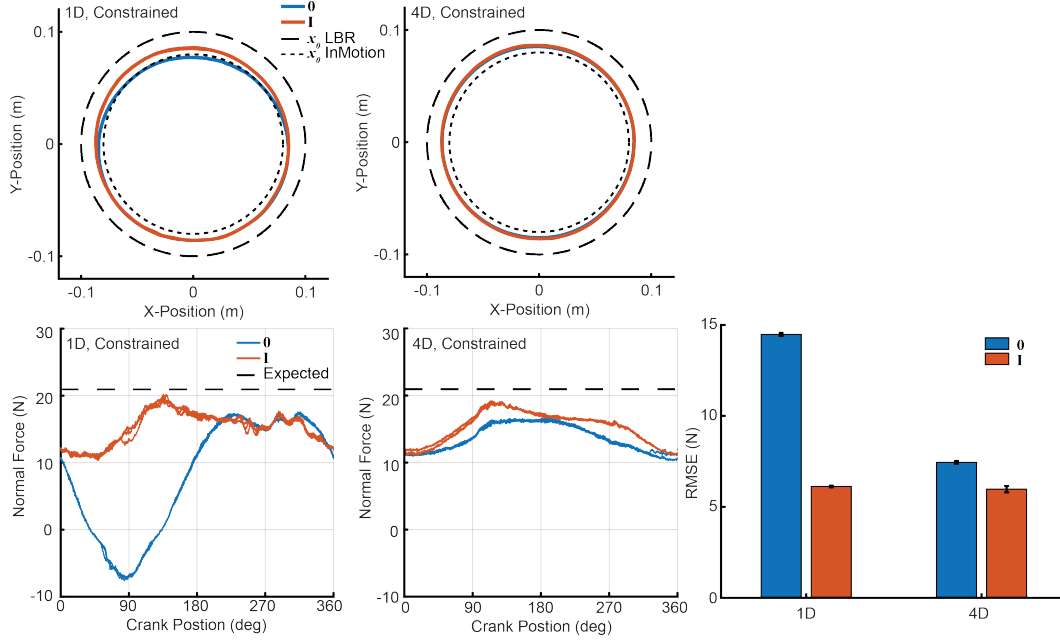


Figure 5-4: **Standard condition, constrained**: 1D nullspace (left column) and 4D nullspace (middle column). Trajectory plots (top left and top middle), the zero-force path of the LBR (bold dashed line), and the zero-force path of the InMotion (dotted line). Normal force plots (bottom left and bottom middle) and the expected normal force (bold dashed line). Normal force RMSE (bottom right). The superposition of mechanical impedance is denoted by  $\mathbf{W}^{-1} = \mathbf{0}$  (blue) and the Moore–Penrose inverse denoted  $\mathbf{W} = \mathbf{I}$  (red). With a 4D nullspace, the superposition of mechanical impedances was substantially less disruptive.

### 5.3.3 Fast Motion

Fast motions were tested to elicit behavior in which inertial dynamics were substantial. With a 1D nullspace, both impedance superposition  $\mathbf{W}^{-1} = \mathbf{0}$  and nullspace projection  $\mathbf{W} = \mathbf{I}$  showed visible deviations from nominal motion when unconstrained (Figure 5-5, top left) and from nominal force when constrained (Figure 5-5, bottom left). With a 4D nullspace, these deviations were substantially reduced, both motions when unconstrained (Figure 5-5, top middle) and forces when constrained (Figure 5-5, bottom middle). A comparison of RMSE for position is shown in Figure 5-5, top right and force in Figure 5-5, bottom right. While impedance superposition was clearly inferior with a 1D nullspace, that disadvantage was nearly eliminated by the 4D nullspace. These differences are compared statistically in section 5.3.5.

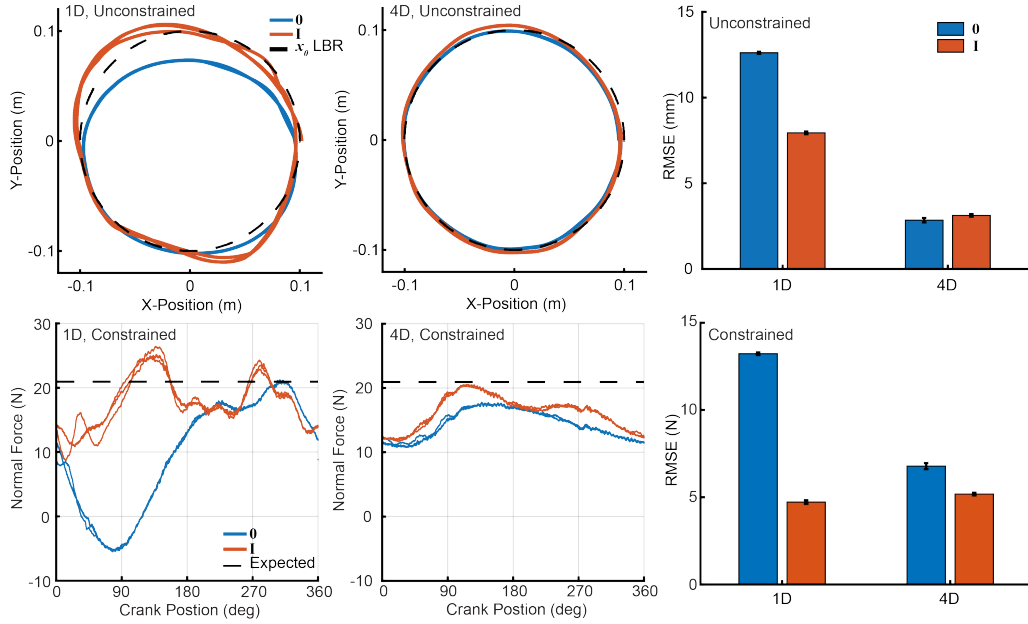


Figure 5-5: **Fast motion:** Comparison of unconstrained (top row) and constrained (bottom row) trials with impedance superposition ( $\mathbf{W}^{-1} = \mathbf{0}$ ) and nullspace projection  $\mathbf{W} = \mathbf{I}$ . Trajectories are presented for the unconstrained trials and normal forces for the constrained trials. Performance with a 1D nullspace is shown in the left column and with a 4D nullspace in the middle column. The right column compares the RMSE of position (unconstrained) and normal force (constrained). With a 4D nullspace, comparable position errors and smaller force errors were achieved with impedance composition.

### 5.3.4 Low Joint-Space Stiffness

For a redundant robot, one way of achieving a predictable motion is to assign joint-space impedances that affect the nullspace motion of the robot in a repeatable manner. As can be seen in section 5.3.1 and 5.3.2, these impedances may cause a conflict between Task 1 and Task 2. Of course, the smaller the joint-space impedance, the smaller the task conflict. However, as the joint-space stiffness is reduced, external perturbations or small errors in friction compensation may cause the redundant degrees of freedom to deviate from the nominal configuration. This may lead to unpredictable or undesirable behavior; for example, joint-space drift may cause the robot to reach joint limits. In this experiment we aimed to determine if an acceptable compromise between these two competing factors could be achieved. We tested whether

joint-space stiffness could be made small enough to reduce task-space disruption to acceptable levels, yet large enough to ensure desirable behavior. In this experiment, the joint-space stiffness of Task 2 was reduced to the point where it was still sufficient to restore the nominal joint-space configuration within one cycle of motion. The exact parameters used are presented in Table 5.2.

The results of using this lower joint-space stiffness can be seen in Figure 5-6. In the unconstrained case with a 1D nullspace, impedance superposition still resulted in greater RMSE position errors than nullspace projection, though, as expected, to a much lesser degree. With a 4D nullspace, any difference became negligible. These differences are compared statistically in section 5.3.5.

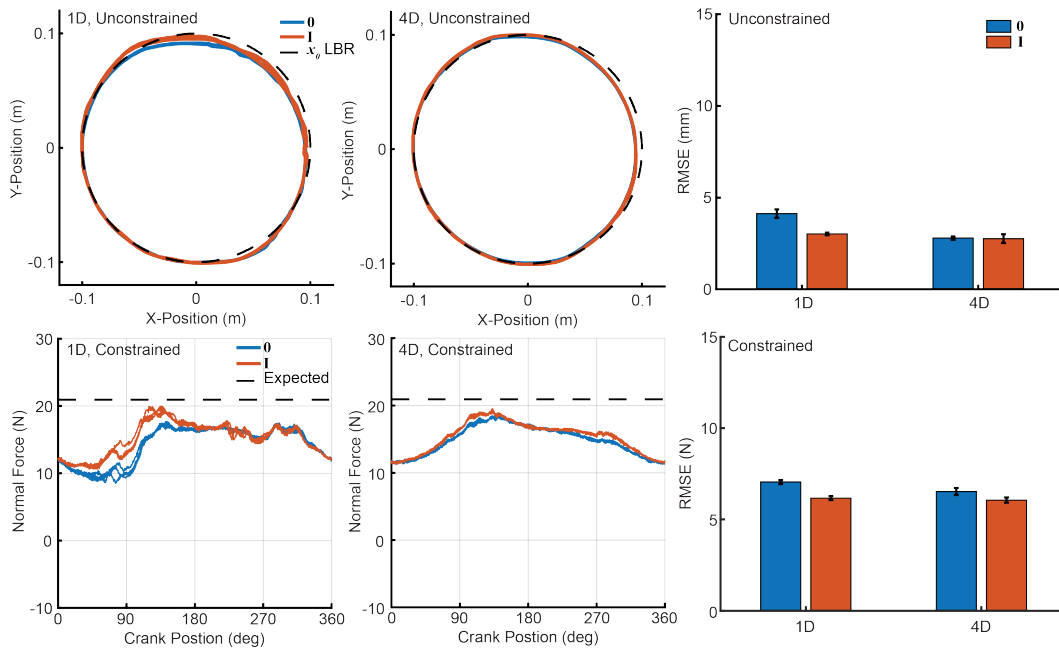


Figure 5-6: **Low stiffness condition:** Performance with low stiffness at slow speed, with a 1D nullspace (left column) and a 4D nullspace (middle column). Trajectories are presented in the unconstrained case (top left and top middle) and normal force in the constrained (bottom left and bottom middle). The right column compares RMSE for position when unconstrained (top row) and force when constrained (bottom row).

### 5.3.5 Statistical Comparisons

The RMSE  $0 - I$  for the position in the unconstrained case showed a significant main effect of nullspace dimension ( $F_{1,18} = 8148.94$ ,  $P \ll 0.001$ ), a main effect of con-

troller ( $F_{2,18} = 2494.55$ ,  $P \ll 0.001$ ), and a significant interaction between nullspace dimension and controller ( $F_{2,18} = 1886.36$ ,  $P \ll 0.001$ ). Post-hoc two-sample t-tests identified significant differences between 1D and 4D nullspace dimension at the standard, fast, and low stiffness levels of the controller conditions. Figure 5-7 (left) shows that the interaction was clearly due to a greater sensitivity to controller with a 1D nullspace. Thus, the effect of increasing nullspace dimension was significant and substantial.

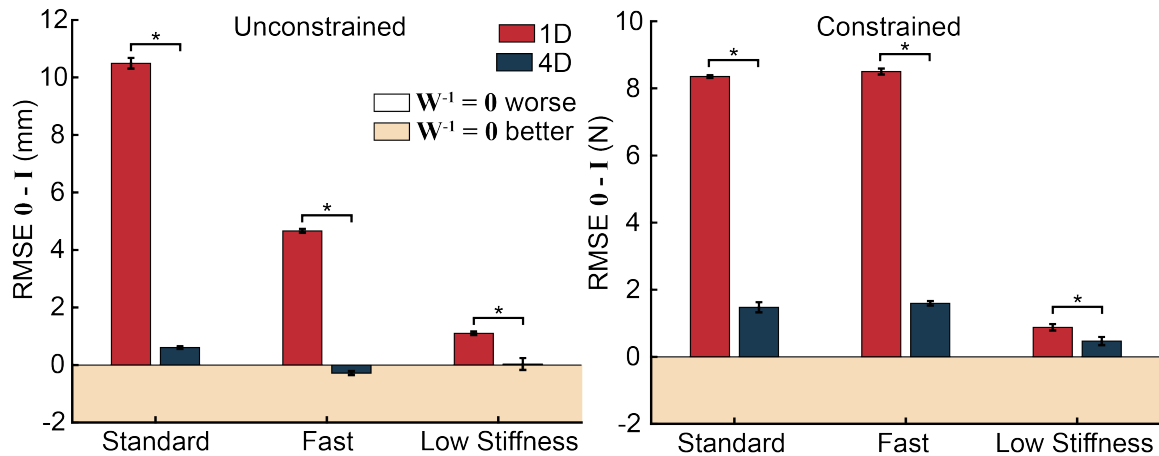


Figure 5-7: Difference in root-mean-squared errors between impedance superposition and nullspace projection for different controller parameters and nullspace dimensions. \* denotes statistical significance with  $P < 0.05$ . Left panel: RMSE  $\mathbf{0} - \mathbf{I}$  of position for unconstrained motions; Right panel: RMSE  $\mathbf{0} - \mathbf{I}$  of normal force for constrained conditions. With a 4D nullspace, both position and force errors were reduced and occasionally impedance superposition was superior to nullspace projection.

The RMSE  $\mathbf{0} - \mathbf{I}$  for normal force in the constrained case showed a significant main effect of nullspace dimension ( $F_{1,18} = 9960.34$ ,  $P \ll 0.001$ ), a main effect of controller ( $F_{2,18} = 3672.30$ ,  $P \ll 0.001$ ), and an interaction between nullspace dimension and controller ( $F_{2,18} = 2078.45$ ,  $P \ll 0.001$ ). Post-hoc two-sample t-tests identified significant differences between 1D and 4D nullspace dimensions at the standard and fast levels of controller. Clearly, the negligible difference at the low stiffness level of the controller was the cause of the interaction as seen in Figure 5-7 (right). Thus, the effect of increasing nullspace dimension on the RMSE  $\mathbf{0} - \mathbf{I}$  normal force was significant and substantial when the joint-space stiffness was not negligible.

## Different Weighting Matrices

We found no evident differences between the various non-zero null-space weighting matrices. The trajectories corresponding to the different nullspace projections all very nearly lay on top of one another. With a 4D nullspace, there appeared to be small, systematic differences in the normal forces but they were minimal.

This was somewhat surprising. Theoretically, using the  $\mathbf{M}(\mathbf{q})$  weighting matrix yields a dynamically-consistent nullspace projector and should therefore show superior behavior for fast robot motions [Dietrich et al., 2015, Chang and Khatib, 1995]. However, without an accurate inertial model of the robot, other projector choices may lead to superior performance in practice [Dietrich et al., 2015, Albu-Schaffer et al., 2003, Nakanishi et al., 2008, Peters et al., 2008]. In the experiments reported here, a negligible difference between projection methods was observed (See Figure 5-8). Of these small differences one notable observation was that using the mass matrix,  $\mathbf{W} = \mathbf{M}(\mathbf{q})$ , was not superior – even in the fast case. Indeed, in many cases performance using the mass matrix was slightly worse than with the other choices.




Figure 5-8: RMSE  $\mathbf{I} - \mathbf{W}$  of position for the unconstrained (top) and RMSE  $\mathbf{I} - \mathbf{W}$  of normal force for the constrained conditions (bottom). Note that the identity weighting matrix was comparable to if not better than the other weighting matrix choices in nearly every condition.

Four ANOVAs were performed to assess the influence of weighting matrix and controller. With a 1D nullspace, the RMSE  $\mathbf{I} - \mathbf{W}$  for position in the unconstrained

case showed no significant effects; however the main effect of controller was nearly significant ( $F_{2,27} = 3.31$ ,  $P = 0.0519$ ).

For constrained motion, the RMSE  $\mathbf{I} - \mathbf{W}$  for normal force with a 1D nullspace showed no significant effect of weighting matrix or controller.

With a 4D nullspace, the RMSE  $\mathbf{I} - \mathbf{W}$  for position in the unconstrained case showed a main effect of weighting matrix ( $F_{2,27} = 95.47$ ,  $P \ll 0.001$ ), a main effect of controller ( $F_{2,27} = 98.45$ ,  $P \ll 0.001$ ), and a significant interaction between weighting matrix and controller ( $F_{4,27} = 43.77$ ,  $P \ll 0.001$ ). The RMSE  $\mathbf{I} - \mathbf{W}$  for the normal force showed a main effect of weighting matrix ( $F_{2,27} = 11.36$ ,  $P \ll 0.001$ ), and a main effect of controller ( $F_{2,27} = 9.40$ ,  $P = 0.001$ ). Post hoc t-tests revealed that only the unconstrained fast speed  $\mathbf{B}$  matrix condition was significantly better than the  $\mathbf{I}$  weighing matrix. To our knowledge this was the first time a damping matrix was used as a weighting matrix; it shows promising results.

Despite the statistical significance of some of these comparisons, the magnitude of the differences between projection methods was small (RMSE less than 1.5 mm and less than 1 N respectively in all cases). These values are close to the resolution of the measurement system. Furthermore, the robot variability across trials was extremely low, which increased the sensitivity of the statistical methods. Thus, while significant effects of nullspace weighting matrix were detected, these significant effects were insubstantial and mainly reflect the repeatability of the robot.

## 5.4 Discussion

The literature on redundancy management, particularly in the areas of reaching and manipulation, predominantly considers nullspace projection methods and their application. One direction of research focused on qualitative evaluation of weighting matrix choice in either free-space or contact tasks [Sadeghian et al., 2014, Dietrich et al., 2015, Whitney, 1969, Khatib, 1987, Albu-Schaffer et al., 2003, Baillieul, 1985, Hollerbach and Ki Suh, 1987, Dietrich et al., 2012]. Another direction of research explored large task hierarchies in which a full-rank joint-space redundancy-managing impedance



plays an insignificant role at the bottom of the hierarchy [Antonelli et al., 2009, Siciliano and Slotine, 1991, Dietrich et al., 2012, Dietrich et al., 2013, Dietrich et al., 2015]. In most of the literature, the case of simple impedance superposition is not considered. This is presumably because, in theory, impedance superposition may cause task conflict, while the use of nullspace projectors will not. In practice, however, this may not be the case.

The goal of the present work was to investigate real kinematically redundant robots with all of their non-ideal behavior, including friction, kinematic errors, etc. These robots are increasingly used to manage complex physical interaction. We aimed to quantify the performance in practice of commonly used control methods – specifically nullspace projections – and compare them with impedance superposition. The stability concerns presented by physical interaction motivate understanding any differences between theoretical and actual performance.

The experimental paradigm investigated in this work, turning a crank, embodies a number of key challenges encountered in physical interaction tasks. The first is contact and/or coupled instability; it has been known since the 1970’s that a robot capable of stable unconstrained motion may become unstable on contact with a physical constraint. A well-established solution to this problem is to ensure energetic passivity of the robot’s dynamic interactive behavior [Colgate and Hogan, 1988, Hogan, 1988]. Generally, nullspace projection approaches are not passive, since the projector only acts on the force/torque factor of the product that determines mechanical power and is not power-continuous. For this reason energy tank methods have been applied to ensure the passivity of nullspace projection methods [Dietrich et al., 2016, Dietrich et al., 2017]. Even if the interactive behavior is dynamically passive (by the usual definition) static instability may be induced by the curvature of a kinematic constraint. Forces applied towards the center of curvature (compressive) are statically de-stabilizing while forces applied away from the center of curvature (tensile) are statically stabilizing [Rancourt and Hogan, 2001, Rancourt and Hogan, 2009]. This is an important distinction as the usual definition of passivity and its relevance to coupled stability does not encompass the static (in-)stability that may be induced by exert-

ing forces [Edward Colgate, 1992]. Managing both of these challenges is a minimum requirement for safe and successful physical interaction [Lachner et al., 2021, Hjorth et al., 2020].

This work quantified the performance of existing controllers in practice on real hardware. We report three major findings: (1) Using different nullspace weighting matrices made no substantial difference; (2) In practice, task conflicts were still present even with nullspace projections; (3) Increasing the nullspace dimension dramatically decreased task conflicts resulting from impedance superposition.

It is important to note that, in nearly every condition, the identity weighting matrix achieved comparable or better performance than the more complex weighting matrices. In practice, there may be little need for complex nullspace projections. In this experiment, the simplest approach, identity weighting, worked equally well or better than the other options investigated. Since the identity weighting matrix facilitates computation, this result might be beneficial for controlling robots with many degrees-of-freedom based on nullspace projection methods.

In the work reported here, we implemented only two multi-dimensional tasks, enabling us to manipulate the ‘wealth’ of nullspace degrees of freedom left by the first task. This allowed us to examine the impact of excess degrees of freedom between the primary task and the redundancy-management task. Additionally, we went beyond qualitative comparison, and statistically examined the impact of weighting matrix choice, task speed, impedance magnitude, and degree of task error with simple superposition in both unconstrained motion and contact scenarios. A key result of this study was that increasing the effective nullspace dimension decreased the task conflict when impedance superposition was used.

### 5.4.1 Limitations

There are several potential sources of artifact in this work which include errors in the robot’s kinematic model, dynamic model, friction-compensation model, and the choice of task conflict metric. Using nullspace projection, the task conflict was theoretically expected to be zero. When implementing controllers on real hardware, errors may be

expected due to numerical artifact and/or imperfect low-level torque control. Both of these effects should have been small; errors observed in our experiments were much larger than could be explained by imprecise torque control. Nevertheless, the torque commands may have been influenced by the highly nonlinear effects of friction. In the experiments, the LBR’s friction compensation was active to ameliorate these effects but in practice this compensation was not expected to perform perfectly. However, the same imperfections were present whether impedance superposition or nullspace projection was used. They cannot account for the differences we observed.

It may seem that the choice of impedance parameters could have influenced the results. This was avoided by experimental design. This work investigated a possible conflict between two impedance controllers. The experiment was designed such that the impedance of one controller was held constant and the other was varied. In this experiment, the end-effector impedance was fixed and two levels of joint impedance were investigated. The ‘standard’ condition used a joint stiffness which was deliberately chosen to be large enough to cause a task conflict in the 1D nullspace case. The second ‘low joint stiffness’ condition was chosen by decreasing the stiffness to the lowest value that would ‘resolve the redundancy’ (restore unperturbed motion within one cycle of motion). The main results of the paper were observed in both the standard and low joint stiffness conditions. They cannot be dismissed as due to a fortuitous selection of impedance parameters.

Another possible source of artifact might have been the choice of metric with which to quantify task conflict. For statistical analysis we chose to use the difference in RMSE (of position and force for unconstrained and constrained tasks, respectively) between the two controllers. This avoided any concern related to the absolute magnitude of these RMSE measures.

Another potential source of error was the model assumed by the nullspace projectors. Considering that the task conflict was observed when moving quasi-statically, dynamics could not have been the cause of artifact. All nullspace projections depend on the Jacobian, which requires a model of the robot kinematics. If the kinematic model was incorrect, a difference between the actual robot nullspace and the model

robot nullspace would exist. Our observations suggest that nullspace projection may be more sensitive to errors in the kinematic model than impedance superposition. This would be consistent with previous theoretical work which has shown that the stability and passivity of an impedance controller is remarkably insensitive to errors in the kinematic model of the robot [Hogan, 1988].

Finally, a concern might be raised that our results were peculiar to the mechanics and kinematics of the particular robot, task and configuration that we studied. To address this concern we performed simulations of arguably the simplest hypothetical case that could demonstrate the influence of increasing nullspace dimension (Appendix H). A planar 3 degree-of-freedom linkage performed a 3 degree-of-freedom end-effector task (nullspace dimension 0) and a comparable 2 degree-of-freedom end-effector task (nullspace dimension 1). Figure H-1 clearly demonstrates the substantial influence of nullspace dimension, even with idealized kinematics and zero friction. Our results are unlikely to be an accident of the particular robot, task and configuration that we studied.

### 5.4.2 Analogy to Polynomial Kernel Methods?

When the nullspace dimension was increased, a substantial decrease in task space disruption was observed. One explanation of this result may be that increasing nullspace dimension with respect to a primary task increases the number of poses that the robot can take. This makes the robot more likely to reach a configuration which will result in a smaller task conflict. This approach, which casts a low-dimensional problem into a high-dimensional space, appears loosely analogous to common-practice data-driven methods for classification. It is well known that a low-dimensional problem which is sparsely populated can be non-linearly cast into a higher dimensional space, e.g. using the polynomial kernel method [Farouki, 2012]. This projection increases the likelihood that a problem which was not linearly separable in the low-dimensional space will be linearly separable in the high dimensional space [Cover, 1965]. We suspect that a similar phenomenon may account for our results but testing this speculation requires further investigation beyond the scope of this report.

### 5.4.3 Understanding Human Motor Control

Humans do not simply regulate kinematics. Humans also modulate the interactive dynamics of their limbs [Hogan, 1985a, Hogan, 1985b, Hogan, 1985c]. Human limb impedance varies as a function of many factors including: muscle activation [Cannon and Zahalak, 1982], movement [Bennett et al., 1992], activity preparation [Lacquaniti et al., 1993], force exertion level [Perreault et al., 2001], task stabilization [Burdet et al., 2001], and walking gait state [Lee and Hogan, 2015, Lee et al., 2016]. Despite about three times as many muscles as skeletal degrees of freedom, if the human limb is viewed as an actuator configured to produce an arbitrary time-varying impedance in an arbitrary configuration, it becomes clear that the human limb is profoundly *underactuated*.

However, humans modulate impedance not only with muscle activity, but also kinematics [Hogan, 1985d, Trumbower et al., 2009]. In many tasks the influence of kinematics can be more than an order of magnitude greater than muscle activity or joint torque. Thus, kinematic redundancy increases the range of impedance which the human limb can produce. The kinematic nullspace is an essential aspect of human physical interaction. Even though the work reported here was performed on a robotic platform, it demonstrated one of the many benefits of the high-dimensional skeletal anatomy that humans possess.

This substantial influence of kinematics may be the reason that several example cases, which employ simple models of impedance, have been able to describe observations of human behavior [Flash, 1987, Hermus et al., 2020, Hermus et al., 2020] and achieve human-like performance [Nah, 2020].

### 5.4.4 Applications to Robotics

One notable result was that there may be cases in which there is no need for a nullspace projector. Instead, the simpler approach of mechanical impedance superposition may be applied. This approach may be successful when the task dimension is small relative to the number of joint-space degrees of freedom; or when only small

joint-space stiffness is required; when there are computational limitations; or when dynamic interactive behavior is prioritized over exact position or force accuracy.

However, we do not conclude that impedance superposition is always superior to nullspace projection; in fact our own results showed cases in which it was not. For moderate joint-space stiffness, there was substantial disruption of Task 1 (in end-effector space) by Task 2 (in joint-space) as intended by the experimental design. That disruption was sufficient to exert compressive forces on the constraint, which is inherently destabilizing and potentially unsafe. With sufficient end-effector stiffness, instability can be avoided and passivity preserved [Colgate and Hogan, 1988, Edward Colgate, 1992]. The biological solution to this problem is that muscle stiffness increases in proportion to muscle force (one of the most robust observations about mammalian muscle) [Hoffer and Andreassen, 1981, Maganaris, 2001] but it is unclear whether this is a satisfactory approach for robotic applications. With a 1D nullspace and moderate or greater joint-space stiffness, nullspace projections may be required.

#### **5.4.5 Future Work**

There are several directions of future work which could provide valuable insight to the observations reported here. We highlight two of them: First, in this work only two nullspace dimensions were investigated. Clearly, the nullspace dimension can be modulated in two ways, by either changing the dimension of task one or by changing the dimension of task two. This simple idea motivates a systematic experimentation or simulation to investigate these factors. Second, this experiment was specifically designed to make the task conflict easy to quantify. However, if more complex tasks are to be understood, alternative metrics for quantifying task conflicts may be required [Schettino et al., 2020].

### **5.5 Conclusions**

In robotics, controlling a large number of redundant degrees of freedom has commonly been viewed as a difficult challenge to overcome, especially if control is performed

via optimization-based techniques. A common approach to deal with the control of kinematically redundant robots is the nullspace projection method. A simpler alternative is based on superimposing mechanical impedances, but that approach is vulnerable to task conflict, whereas nullspace projections theoretically avoid this problem. In practice, we observed that both nullspace projections and impedance superposition resulted in measurable task conflict. This surprising observation was minimally influenced by the choice of projection weighting matrix. Remarkably, when the dimension of the nullspace increased, the superposition method showed errors that were comparable to the nullspace projection methods. With no disrespect intended to Richard Bellman, high-dimensional kinematics may be a blessing rather than a curse.





# Chapter 6

## Dynamic Primitives Limit Human Force Regulation during Motion

This chapter is an adapted version of [West et al., 2022] published in IEEE Robotics and Automation Letters. This work was done in collaboration with Michael West, Meghan Huber, and Pauline Maurice.

### 6.1 Introduction

The majority of human neuro-motor control research to date has focused on the control of motion during free unconstrained reaching without physical contact (for review see [Gulletta et al., 2020, Campos and Calado, 2009]). In this case, relating a planned motion to an actual motion is sufficient to describe the control system. In robotics, the mathematics underlying motion control is well understood [Slotine and Asada, 1992]. However, most tasks that humans and robots perform require physical interaction with the external environment; for such interactive tasks, motion control alone is insufficient.

During physical interaction, bidirectional forces between the actor and the environment critically affect the behavior of the coupled system. If humans regulate motion during free reaching, a simple extension of this idea to contact tasks may be

to regulate both force and motion. In robotics, hybrid control allows for simultaneous and independent control of both motion and force in complementary subsets of the workspace [Mason, 1981, Raibert and Craig, 1981]. In human motor control, it is yet unresolved whether humans can control force independent of motion.

Several studies in human motor neuroscience have reported findings in support of such hybrid control. For example, Chib et al. [Chib et al., 2009] found that hybrid motion / force control can describe how humans performed an interaction task in a virtual force field. Casadio et al. [Casadio et al., 2015] presented and experimentally validated a computational model of how the neural system may combine two independent modules that separately control motion and force. Further, neural activity in the motor and parietal cortex of non-human primates indicate that there are separate modules for the control of force and motion [Georgopoulos et al., 1992, Hamel-Pâquet et al., 2006, Sergio and Kalaska, 1998].

On the other hand, it has been shown that the central nervous system (CNS) contains a controller that modulates the coupling of force and motion [Kolesnikov et al., 2011, Piovesan et al., 2019]. Other studies demonstrated that humans modulate the relation between motion and force during upper limb reaching in unstable force fields [Burdet et al., 2001, Milner and Franklin, 2005, Osu et al., 2003, Takahashi et al., 2001]. Additionally, our own previous research showed that exerted force depended on the velocity profile when grasping and following a robot manipulandum. Specifically, participants were asked to trace the motion of a robot manipulandum without exerting force as it moved on an elliptical path with varying velocity profiles [Maurice et al., 2018b]. If force can be controlled independent of motion, the velocity profile should not matter; however, it did.

This study aimed to examine human control of physical interaction that could resolve these seemingly contradictory results. We conducted an experiment in which participants physically interacted with a motion-controlled robot to test whether humans could regulate force independent of motion. We refer to this independent control as ‘direct force control’ (Figure 6-1a). Explicitly, direct force control applies an actual force as a function of only a planned force. This function is an operator that may

be dynamic and nonlinear. If participants can regulate force independent of motion, direct force control can be accepted as a plausible schema for human physical interaction. Conversely, if humans are unable to decouple force from motion, an alternative hypothesis is ‘indirect force control’. With indirect force control, a planned force  $f_p(t)$  may still exist in the forward path, but an impedance term  $Z\{\cdot\}$  is needed to relate the difference between input motion  $x_0(t)$  and actual motion  $x(t)$  to the output force  $f(t)$  (Figure 6-1b). The core feature of indirect force control is that force depends on motion.

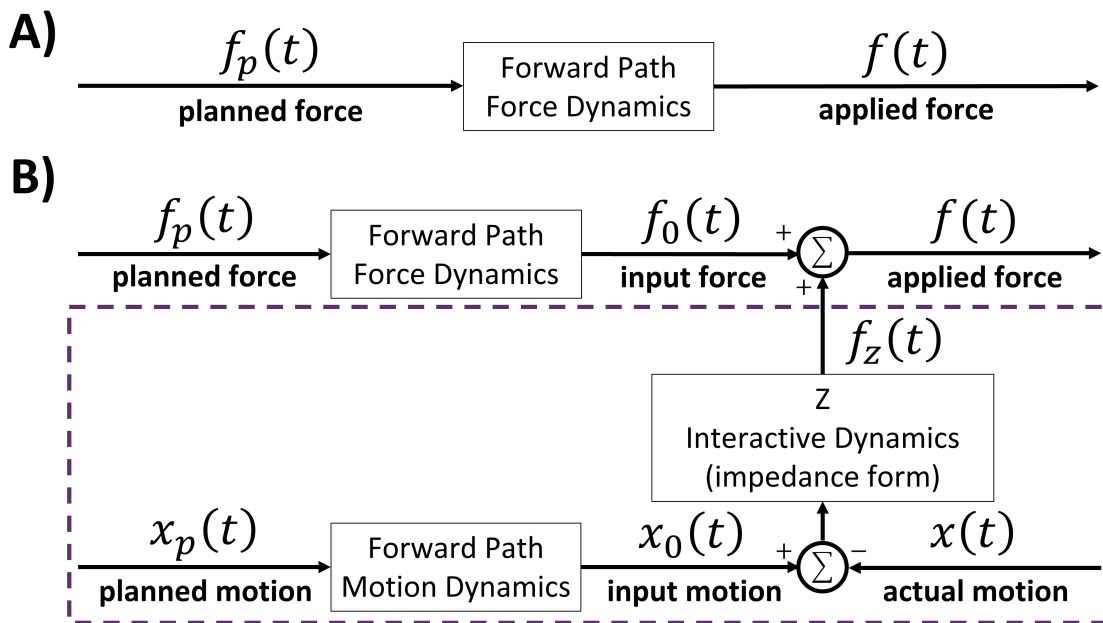


Figure 6-1: (a) **Direct force control.** Applied force is a function of a planned force that is independent of motion. (b) **Indirect force control.** Applied force is a function of a planned force but also depends on motion. Further details are in the Discussion section.

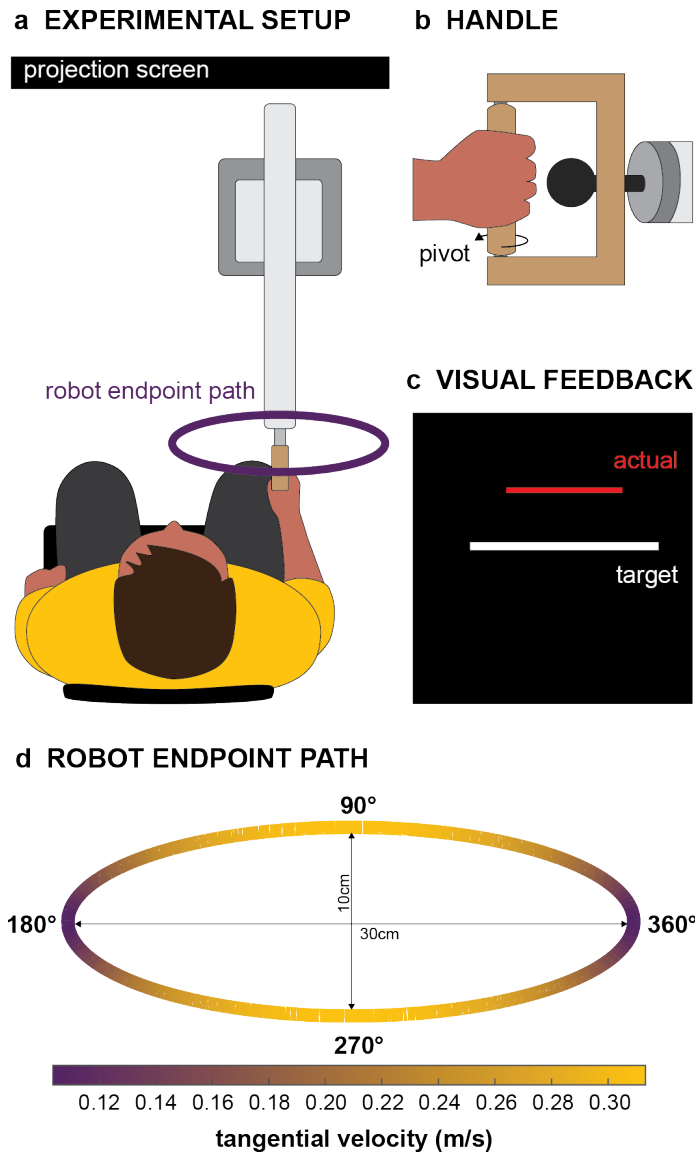


Figure 6-2: **Experimental setup.** (a) Top-down view of the experimental setup. Participants were instructed to hold the handle of a moving robotic manipulandum and apply a constant force in the direction of the robot’s motion. The elliptical path of the robot endpoint is displayed on the figure for clarity. However, participants did not see any visual display of the robot path. (b) Robot handle used to decouple human wrist and robot end-effector orientations. (c) To provide visual feedback, the projection screen displayed a stationary white bar indicating the target tangential force of 5N and a moving red bar indicating the current applied tangential force. Visual feedback was given during Block 1V and Block 2V. Otherwise, the screen was black. (d) Elliptical trajectory of the robot endpoint (i.e., handle) in the horizontal plane. The robot manipulandum moved counterclockwise and followed a velocity profile that was in accordance with the two-thirds power law [Zago et al., 2018]. Tangential velocity is shown by color.

The direct force control hypothesis leads to a testable prediction: Errors in contact force will be independent of motion. Thus in this experiment, participants were instructed to apply a specific constant force on a robot manipulandum in its direction of motion as it moved along an elliptical path. To give participants the best opportunity to complete the task, the robot moved with a velocity profile that matched human movement preferences, i.e., angular velocity scaled with curvature with a power of  $2/3$  [Maurice et al., 2018b, Zago et al., 2018]. Despite visual feedback and some practice, errors in exerted force persisted and were dependent on motion, suggesting (1) a coupling of force and motion, and (2) the existence of an underlying structure in the feedforward motion planning signal. Additional analysis of previous data from [Maurice et al., 2018b] further validated the current results. In sum, this work showed that interactive dynamics are significant and of particular concern in (1) quantification of human performance and in (2) physical human-robot interaction.

## 6.2 Methods

### 6.2.1 Participants

Eleven healthy right-handed individuals (3 females, 8 males; ages from 19 to 35 years old) participated in the experiment for some compensation. All participants signed a consent form which explained the experiments' procedures. The experimental protocol was approved by the Institutional Review Boards of Northeastern University and the Massachusetts Institute of Technology.

### 6.2.2 Experimental Procedures

#### Task and Instructions

Participants were instructed to hold the handle of a moving robotic manipulandum (HapticMaster) [Van der Linde et al., 2002] and apply a constant 5N force in the direction of robot motion (i.e., tangential direction) as it traversed an elliptical trajectory in a horizontal plane (Figure 6-2a).

Participants performed the experiment seated and held the robot through a vertical handle which could pivot around its vertical axis; the pivot decoupled the robot end-effector orientation and the participant's wrist orientation (Figure 6-2b). Participants were positioned such that when holding the robot handle at the  $270^\circ$  position of the ellipse (Figure 6-2d), the right upper arm hung downward slightly away from the torso. This position aligned the forearm with the minor axis of the ellipse. The robot height was adjusted such that the forearm was approximately parallel to the ground. This resulted in an angle between the upper arm and forearm of approximately  $90^\circ$  (Figure 6-2a-b).

### **Visual Display**

Participants sat approximately 2.2m in front of a projection screen (height: 1.8m, width: 2.4m). In conditions with visual feedback, two horizontal bars appeared on the screen (Figure 6-2c). A red horizontal bar moved vertically to indicate the tangential force (averaged over 80ms) applied by the participant onto the robot; the stationary white bar indicated the desired tangential force of 5N. Otherwise, the screen was black.

### **Control of Robot Motion**

The robot handle was commanded to move counterclockwise along an elliptical path (major axis = 30cm, minor axis = 10cm) on a horizontal plane with a period of 3s (Figure 6-2d). The velocity profile of the robot handle followed the so-called  $2/3$  power-law relation [Maurice et al., 2018b, Zago et al., 2018] between path curvature and angular velocity (Figure 6-2d), decreasing in highly curved portions and increasing in less curved portions. The position of the robot handle was controlled with a Cartesian PD controller; a high proportional gain was used such that deviation from the desired trajectory was negligible. The desired position of the robot was updated at 700Hz, and an internal force control loop ran at 2kHz.

### 6.2.3 Experimental Design

To assess the effect of practice and visual feedback on force control, participants performed four experimental blocks; each block consisted of 15 trials. In each trial, the robot continuously traversed the elliptical path four times with a period of 3s per cycle; each trial lasted 12s. Blocks 1V and 2V presented visual feedback as shown in Figure 6-2c. Blocks 1NV and 2NV did not present visual feedback. Participants always performed the four blocks in the following order: 1V, 1NV, 2V, 2NV. In all four blocks, participants were instructed to maintain a constant force of 5N in the tangential direction. At the start of each trial, participants heard three short beeps through a headset, after which the robot began to move. Between blocks, participants were allowed to take a break if needed.

A familiarization block, referred to as Block F, preceded the four experimental blocks. It also consisted of 15 trials of 12s each. There, participants were instructed to maintain a constant level of force in the tangential direction of the robot motion. The exact level of force applied was not specified and there was no visual feedback. After Block F, participants were given 60s to familiarize themselves with the visual feedback. During that time, the robot was in a stationary position and participants could apply forces against the robot. In total, the experiment lasted just over an hour.

### 6.2.4 Dependent Measures and Data Processing

The force that participants applied to the robot handle was measured at  $\sim 560$  Hz with a 3 DoF force sensor mounted at the robot end effector. In each trial, the tangential component of the force applied by the human to the robot was calculated and resampled as a function of robot position along the elliptical path at a resolution of  $1^\circ$ .

Angular position along the elliptical path was defined using the eccentric anomaly<sup>1</sup>,

---

<sup>1</sup>The eccentric anomaly is one of three angles (or "anomalies") identified by Johannes Kepler in his study of celestial mechanics to describe the position of a body that is moving along an elliptical orbit [Kepler, 1609]. The other two angles are the true anomaly and the mean anomaly.

$E$ , such that  $E = \text{atan}(ay/bx)$ , where  $a$  and  $b$  were half the length of the major and minor axes of the ellipse, respectively.  $x$  and  $y$  were the magnitudes of the position vector in the direction of the major and minor axes in Cartesian space, respectively.

For each trial, task performance was summarized by calculating the root-mean-square (RMS) of force error. Force error was defined as the difference between the actual tangential force and the target tangential force of 5N. The tangential force was resampled as a function of robot position along the elliptical path at a resolution of  $1^\circ$ . To avoid the potential influence of transient behavior, the first cycle of each trial was omitted in the calculation of the RMS error.

### 6.2.5 Data Analysis and Statistics

All data were processed and statistical analyses were performed using custom scripts in MATLAB. The significance level for statistical tests was  $\alpha = 0.05$ . Unless stated otherwise, only data from the four experimental blocks (i.e., Blocks 1V, 1NV, 2V, and 2NV) were included in the statistical analyses.

#### Performance Improvements

Prior to testing the effects of practice and feedback in the four blocks, performance improvements were assessed within blocks across trials by calculating linear regressions between RMS force error and trial number. Performance improvement was indicated if the slope was different from zero, i.e., the 95% confidence interval of the slopes did not include zero.

To determine where participants' performance reached steady state, the regression slopes between trial number and the average RMS force error across participants were calculated iteratively for the last 15, 14, 13, trials and so forth until an insignificant slope was found. This occurred when the linear regression was computed over the last 9 trials (i.e. from trial 7 to 15). The lack of a significant slope with the RMS of force error and trial number justified averaging measures over the last 9 trials within a block.



To assess whether visual feedback or practice across the 2 blocks influenced performance, the block means of all participants were calculated over the steady state portion of each block. These block means of RMS force error were submitted to a 2 (block)  $\times$  2 (feedback) repeated-measures ANOVA.

### **Existence of Motion-Dependent Force Errors**

To assess the presence of motion-dependent patterns in the force error, the auto-correlation function of force error (as a function of robot angular position) was calculated for each trial. The lag with the maximum peak in the auto-correlation function (hereafter referred to as maximum lag) and its corresponding auto-correlation coefficient (maximum auto-correlation coefficient) were identified.

Two clusters were identified in the distribution of lags at maximum auto-correlation and their means were determined. Trials where the maximum auto-correlation coefficient was less than 0.1 were omitted from the analysis of position dependency of force error (4 out of 825 trials). From visual inspection, these low maximum auto-correlation coefficient values resulted from isolated uncharacteristic changes in RMS force error during the trial. They also occurred at lag values that were significant outliers.

## **6.3 Results**

### **6.3.1 Performance Improvements**

#### **Change in RMS Force Error Within Blocks**

Inspection of the grouped time series of force error revealed that subjects showed a consistent decline of the force error in the first part of Block 1V (Figure 6-3a). The iteratively computed linear regressions between the average RMS force error across participants and trial for the last 15, 14, 13, and so forth trials identified that the force error values in Block 1V plateaued when calculated over the last 9 trials (i.e., from trial 7 to 15). From trial 7 onwards the regression slopes did not differ from zero. As

this initial drop of error seemed to be a result of familiarization, subsequent analyses only examined the last 9 trials of all four blocks to evaluate the errors reached in each condition.




Figure 6-3: Mean RMS of force error across participants in each trial for (a) all experimental blocks in the main experiment, and (b) Experiment 0N. Yellow dots depict the average value across participants. The shaded region depicts  $\pm 1$  standard deviation across participants. An asterisk (\*) indicates a significant linear relation between the mean RMS force error and trial number calculated across participants for each block.

### Effect of Practice and Visual Feedback on RMS Force Error Across Blocks

To statistically evaluate whether visual feedback and practice had a significant effect on the force error, a 2 (block)  $\times$  2 (feedback) repeated-measures ANOVA was conducted. The force error revealed a significant interaction ( $F_{1,10} = 15.74, p = 2.66e - 03$ ) as the mean RMS error decreased from Block 1V ( $M = 3.63N, SD = 1.12N$ ) to Block 1NV ( $M = 3.30N, SD = 1.12N$ ) and increased from Block 2V ( $M = 3.08N, SD = 1.32N$ ) to Block 2NV ( $M = 3.86N, SD = 1.96N$ ) (Figure 6-4a). However, neither the main effect of block ( $p = 0.99$ ), nor the main effect of feedback ( $p = 0.60$ ) were statistically significant. Recall, all subjects completed the experiment in order: Block 1V, 1NV, 2V, 2NV. Thus, the increase in mean RMS of force in Block 2 was likely the result of cognitive or physical fatigue as the experiment was quite long.




Figure 6-4: Mean RMS of force error from the last 9 trials for (a) all experimental blocks in the main experiment, and (b) Experiment 0N. Yellow dots depict individual participants. Error bars depict  $\pm 2$  standard errors of the mean. The mean RMS of force error was significantly higher in all experimental blocks of the main experiment compared to Experiment 0N (Table I).

### 6.3.2 Existence of Motion-Dependent Force Errors

Given this indifference to feedback and practice, the time series of force error were inspected. As illustrated by the raw force data shown for a representative participant in Figure 6-5, force error was periodic with pronounced peaks at multiples of  $180^\circ$  in all blocks. To quantify this, the  $k$ -means clustering algorithm was used to assign the maximum lag value of each trial to one of two clusters ( $k = 2$ ) and then calculate the centroid (mean) of each cluster. Two clusters were used because the within-cluster sum of squares showed minimal improvement after the number of clusters increased beyond  $k = 2$ . The means of each cluster identified in the maximum lag data of trials in Blocks 1V, 1NV, 2V, and 2NV were  $179.3^\circ$  and  $359.3^\circ$  (Figure 6-6a). The average maximum auto-correlation coefficient was 0.43 ( $SD = 0.13$ ). Trials with maximum lags of  $360^\circ$  indicate that the peaks in force error at half and full cycle were different, while the maximum lag at  $180^\circ$  indicates that the two peaks in force error were similar. Analyses of individual participants revealed that five subjects showed higher

force applied at  $180^\circ$  and six subjects showed higher force applied at  $360^\circ$ . Taken together, these results indicate that force error strongly depended on the phase of the oscillatory robot motion.

Given this pronounced periodicity in the experimental blocks that specified 5N force, we also examined whether this periodicity was present spontaneously. The same autocorrelation analyses were run on the trials of the familiarization block (Block F). Figure 6-6b shows two clusters with mean values of  $179.9^\circ$  and  $359.5^\circ$ . The average maximum auto-correlation coefficient was 0.53 ( $SD = 0.12$ ). As illustrated in Figure 6-6b, these results give strong evidence for a spontaneous coupling of motion and force.

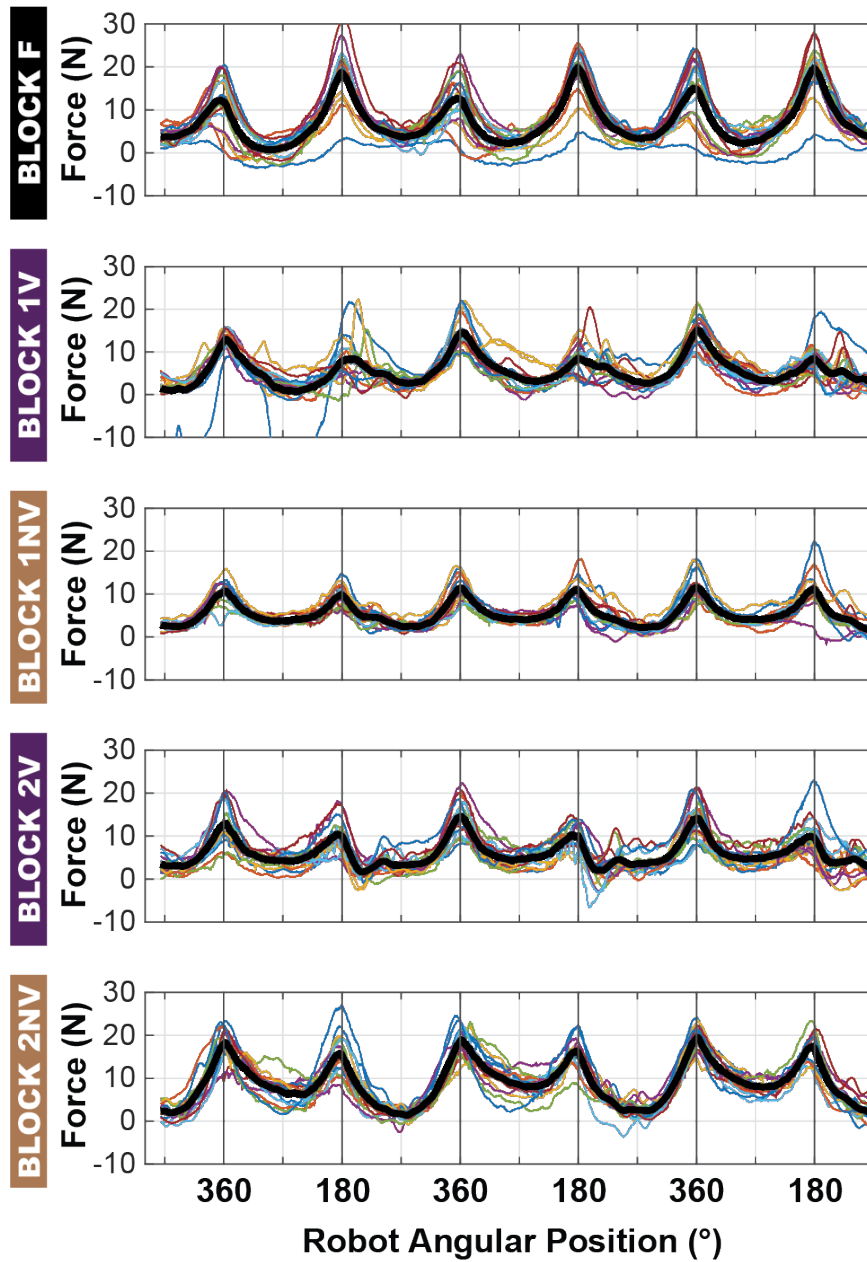


Figure 6-5: **Raw force data in each block for a representative participant.** For each of the 5 blocks in the main experiment, plots of tangential force over robot angular position are shown for the last 3 cycles of every trial. The tangential force was resampled as a function of robot position along the elliptical path at a resolution of  $1^\circ$ . Each trial is depicted with a thin, colored line, and the average across all trials is depicted with a thick, black line.

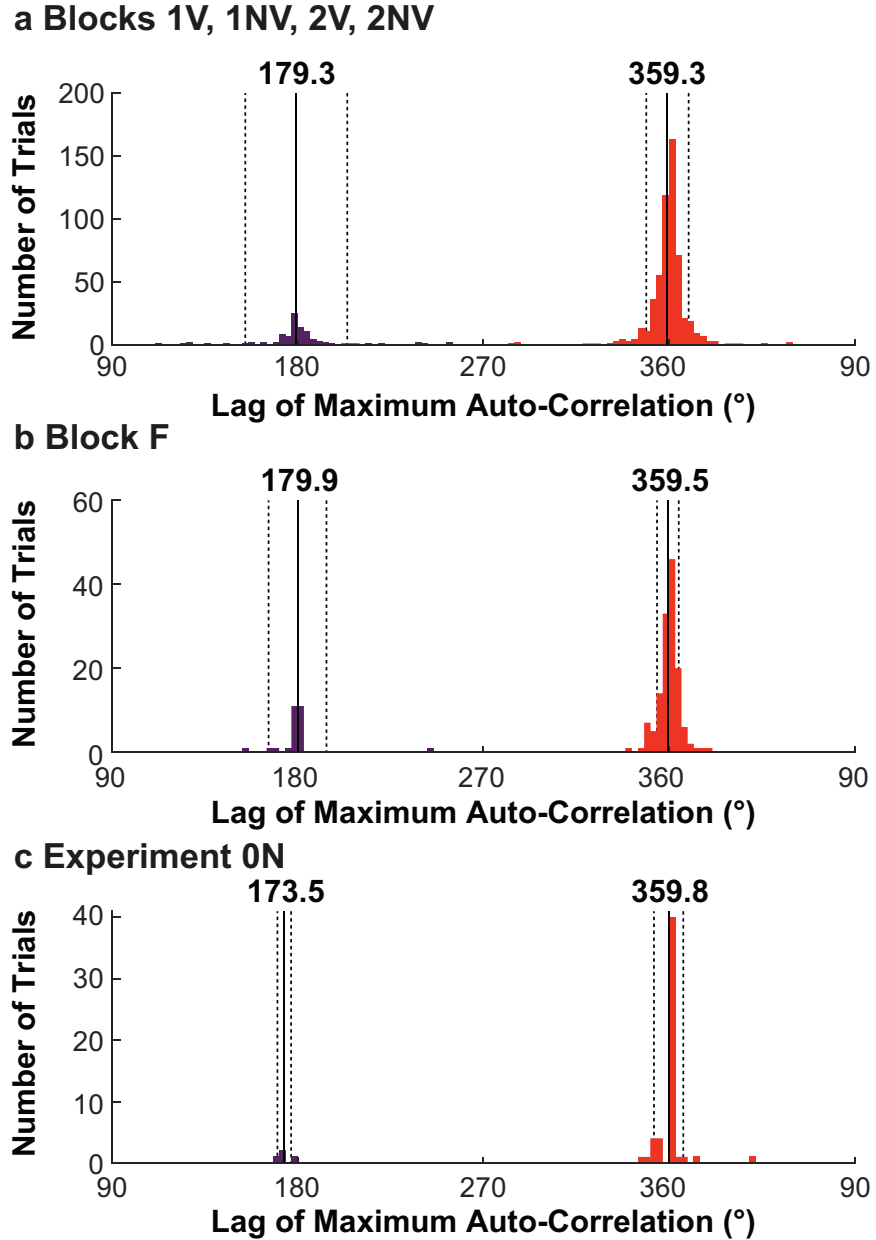


Figure 6-6: Evidence of motion-dependent periodic force errors for all trials in (a) Blocks 1V, 1NV, 2V, and 2NV of the main experiment, (b) Block F of the main experiment (c) Experiment 0N. Histogram of lags of maximum autocorrelation (referred to as maximum lag) values measured in units of robot angular position. In all conditions, two clusters were identified. The solid lines indicate the mean of each cluster, and the dashed lines depict  $\pm 1$  standard deviation of each cluster.

## 6.4 Discussion

This work investigated humans' ability to directly modulate force during motion. Subjects were asked to apply a constant force on a robot manipulandum moving along an elliptical path. The hypothesis of direct force control predicted that errors in contact force would be independent of motion. Here, the force errors observed throughout the entire main experiment depended on motion. Force error showed a periodic pattern consistent with the periodicity of the path; it varied with motion. After initial performance improvements, participants did not reduce force errors with practice, even when visual feedback was provided. Motion-dependent patterns in force error were also observed in Experiment 0N (i.e. Experiment 1B in [Maurice et al., 2018b]), further validating the main results. These findings suggest that force and motion are coupled as schematically shown in Figure 6-1b.

### 6.4.1 Force Error

In the main (5N) experiment subjects were given visual feedback of their tangential force in two of the blocks (Figure 6-2c). In contrast to static tasks, where visual feedback enables subjects to apply a constant force quite accurately [Massey et al., 1992], the elliptic motion of the robot manipulandum in this study significantly compromised the subjects' ability to regulate force. Subjects did not eliminate residual errors, which varied periodically with motion.

Interestingly, the overall magnitude of force errors was significantly lower when the target force was lower (Figure 6-4). There are several plausible explanations why this occurred. One possibility is that greater force applied induced higher noise (i.e., signal dependent noise) [Osu et al., 2004]. Another possibility is that greater force applied induced higher hand impedance [Lipps et al., 2020], which would amplify any errors between the input and actual trajectories. This would provide further support for indirect force control (Figure 6-1b).

Nonetheless, production of actual force  $f(t)$  that equals input force  $f_0(t)$  is possible using the indirect force control strategy of  $f(t) = f_0(t) + Z\{x_0(t) - x(t)\}$  when

$Z\{x_0(t) - x(t)\} = 0$  (Figure 6-1b). This can be achieved in one of two ways: (1) zero interaction dynamics and/or (2) a simultaneous prediction of the input<sup>2</sup> trajectory  $x_0(t)$  that matches the actual trajectory  $x(t)$ . Thus, it is critical to note that if motion-dependent force errors were not observed, it would be impossible to distinguish between the direct and indirect force control strategies. However, the force error we observed was dominated by motion dependency (Figure 6-5). Specifically, the force error was periodic with maximum auto-correlation at lag corresponding to the 180° and 360° ellipse positions (Figure 6-6).

These motion-dependent force errors were also observed in both the familiarization Block F of the main experiment, where subjects were instructed to apply a constant tangential force, (Figure 6-6b) and Experiment 0N (Figure 6-6c), where subjects were instructed to apply zero force. In both, subjects did not receive any visual feedback. Despite some practice with and without visual feedback, the motion dependency of the applied force persisted throughout the main experiment (Figure 6-5). This robust observation suggests an underlying structure in humans' ability to regulate force during motion that limits the performance of this task.

## 6.4.2 Dynamic Primitives

Accurately controlling force would require the central nervous system to acquire an 'internal model' of the task with which to 'compute' predictive forward-path control inputs. The theory of dynamic primitives proposes that motor behavior, with and without physical interaction, is constructed using a limited set of primitive dynamic behaviors that are the 'building blocks' of more complex actions [Hogan and Sternad, 2012, Schaal and Sternad, 2001, Schaal, 2006, Degallier and Ijspeert, 2010]. These 'building blocks' allow for a detailed plan of time-varying neuro-muscular activity to be abstracted to the parameters of a limited set of stereotyped motor patterns. Rhythmic movements can be generated by oscillations, one class of dynamic primitives. The interactive primitive is mechanical impedance. The parameters of these

---

<sup>2</sup>This input trajectory has been referred to as the zero-force trajectory, as it is the motion that would occur in the absence of external forces.



‘building blocks’ may be encoded; this may facilitate human learning, performance, and retention of complex skills.

Dynamic primitives do not preclude arbitrary patterns of force production. A sufficiently accurate internal model might be used to compute both  $f_0(t)$  and the corresponding  $x_0(t)$  (Figure 6-1). However, if the parameters of oscillatory primitives used to plan the motion were limited, the time-course of force production would also be limited. Periodic force errors in our experimental results suggest that the controller appears to be content with ‘good-enough’ performance, which can be obtained using a limited set of ‘primitive’ oscillations and a sufficiently low mechanical impedance. Thus motor behavior constructed by dynamic primitives may result in performance limitations – such as the observed imperfect, periodic force regulation reported here (Figure 6-5).

Other results support this account. A combination of two oscillations (e.g. in two degrees of freedom) generate the two-thirds power law relation between path curvature and angular velocity. Previous studies of crank turning suggest that during physical interaction humans generate an elliptical zero-force trajectory which exhibits a coincidence of speed and curvature extrema [Hermus et al., 2020]. These observations are also consistent with the work of [Maurice et al., 2018b, Schaal and Sternad, 2001, Huh and Sejnowski, 2015, Hermus et al., 2020]. The smallest force errors in [Maurice et al., 2018b] were observed when the velocity profile of the robot followed the two-thirds power-law relation. Moreover, position-dependent errors are evident in the results of other studies on constrained motion [Koeppen et al., 2017, Ohta et al., 2004, Russell and Hogan, 1989]. However, to our knowledge this is the first time that position-dependent force errors have been systematically quantified during a force regulation task with substantial motion.

### 6.4.3 Limitations

In the main experiment, participants experienced the task for approximately one hour (300 cycles). It is possible that participants could learn to better regulate their force with additional practice (e.g., over multiple days). However, investigation of

extensive practice was not the goal of our work. Humans regularly perform a variety of novel forceful interaction tasks with ease and apparently without requiring long-term practice. In fact, task performance slightly worsened at the end of practice in the main experiment, possibly indicating that fatigue set in. Hence, this study aimed to identify the performance that might be expected from intuitive and spontaneous human-robot interaction.

Force errors might also be ascribed to poor perception of the robot’s motion. However, the motion slowly ( $\sim 0.33\text{Hz}$ ) followed a large elliptical path of 66.8cm in circumference. Additionally, if errors in the perception of the robot’s motion led to force errors, we would expect to see differences in error between the blocks that did and did not have visual feedback. Figure 6-4 demonstrates that this was not observed. Motion dependent deviations from the instructed force were persistent throughout the entire main experiment (Figure 6-5 and 6-6a-b).

It is also possible that there may have been too much cognitive demand from mapping the vertical feedback display to the horizontal force. While this argument cannot be directly refuted from the results reported here, it is unlikely to account for our main result. Figure 6-6b shows that subjects force error was motion dependent even before visual feedback had been provided. In short, the position-dependence of force error was consistent throughout the main experiment, despite the presence of visual feedback.

#### **6.4.4 Implications**

Understanding the preferred control strategy employed by humans may guide the design of robot controllers to manage physical interaction. A roboticist may draw upon the proposed ‘building blocks’ to program a simple controller to achieve a complex task [Ijspeert et al., 2002, Stulp et al., 2012, Peters and Schaal, 2006]. For example, a controller based on dynamic primitives has been used successfully (in simulation) to control a 2 DOF arm to manipulate a dynamically complex whip with 50 DOF in a targeting task [Nah, 2020]. Furthermore, the human body has a large number of redundant degrees of freedom. Kinematic redundancy has commonly been viewed

as a difficult challenge to overcome, especially if control is performed via conventional optimization-based techniques. However, redundant degrees of freedom may be controlled by superposition of mechanical impedance primitives. Remarkably, unlike optimization-based methods, as the number of redundant degrees of freedom increased, control based on the superposition of impedance primitives improved; in effect, with greater redundancy control became easier [Hermus et al., 2022].

The account of humans’ motor control strategy proposed here may be especially useful to design controllers for robots intended to interact physically with humans. This paper demonstrated that errors in human force regulation may result from limitations in the way humans compose motor actions (e.g., possibly through dynamic primitives). These limitations should be taken into consideration in all applications involving physical human-robot interaction, including amputation prostheses, assistive exoskeletons, robot-aided rehabilitation, and physical human-robot collaboration.

## 6.5 Conclusions

In this work, we scrutinized a pervasive assumption: force and motion can be controlled independently (an idea referred to here as direct force control). To examine this assumption, subjects were asked to apply a constant force on a robot manipulum that moved along an elliptical path with a speed profile consistent with the preferred pattern of human motion (the two-thirds power law). Results showed that subjects were unable to control force accurately during motion, despite some practice and the presence of visual feedback; errors in force were periodic in response to the periodic motion of the robot. These results point towards an indirect force control formulation (Figure 6-1b), in which commanded motion acts through mechanical impedance to evoke force. Furthermore, the periodic pattern of path-dependent force errors was consistent with commanded motion composed of oscillatory primitives. Taken together, these findings suggest that a relatively simple mathematical model combining dynamic motion primitives with mechanical impedance, as an additional primitive, is competent to describe how humans control contact and physical

interaction. A quantitative model is especially important for designing devices that physically collaborate with humans.

# Chapter 7

## Conclusions

To recap, humans excel at physical interaction with objects, even when those objects introduce complex dynamics and impose kinematic constraints. Furthermore, human dexterity exceeds that of most modern robots, despite the fact that the human neuro-mechanical system is considerably slower than its robotic counterparts. Despite the integral nature of this topic, physical interaction, a key to activities of daily living (ADLs), has remained poorly understood. In part, this thesis demonstrated that, to understand physical interaction the interactive dynamics of the limb must be taken into account.

It has been proposed that, to achieve highly dynamic and dexterous performance despite neuro-mechanical limitations, human behavior is composed of dynamic primitives [Hogan, 2017, Hogan and Sternad, 2007, Hogan and Sternad, 2012]. These are dynamic attractors (for example, limit-cycle oscillations) that emerge from nonlinear interactions between neural and mechanical parts of the system and, once evoked, require minimal continuous intervention from higher levels of the central nervous system. In this theory, dynamic primitives are ‘building blocks’ of complex actions. If dynamic primitives compose action, quantifiable limitations may be evident in human motor control.

While it is fundamentally hard to prove hypotheses in human motor control, this thesis demonstrated that the theory of dynamic primitives can descriptively account for systematic patterns in constrained motion. Furthermore, the value of this hypoth-

esis was demonstrated in robotics by the examples of managing kinematic redundancy and force regulation.

In Chapter 2 [Hermus et al., 2020], a study of planar crank turning was presented. A model of limb impedance was assumed, and the zero-force trajectory was computed. The zero-force trajectory contained evidence consistent with the speed-curvature relation reported in unconstrained motion [Lacquaniti et al., 1983, Huh and Sejnowski, 2015]. This suggests that our novel approach, estimating the zero-force trajectory, reveals information about neural control.

In Chapter 3 [Hermus et al., 2020], further analysis was performed to understand the structure present in the zero-force trajectory. We observed that when subjects turned in different directions (CW vs CCW) the orientation of the zero-force trajectory changed. This is consistent with control based on oscillations. However, in the slow case subjects turned with a period of 13.33 seconds. At this speed the task could not be perceived or executed as an oscillatory action [James, 1890]. In this case, we observed an increased variability in the slow speed condition, suggesting a transition from control using oscillations to control using submovements. These observations are consistent with the hypothesis of dynamic primitives.

While insight can be gained by assuming the mechanical impedance, as in Chapter 2 and 3, direct measurement of mechanical impedance would allow for the hypothesis to be further tested. In Chapter 4, we set out to measure mechanical impedance during the task of crank-turning. Initially, time-based ensemble methods were employed. However, a key assumption – the stationarity of the noise processes – was not justified. A ‘work-around’ was developed: (1) high-pass filter to remove the influence of the underlying time-varying zero-force trajectory; (2) assume impedance is a function of configuration and identify a configuration-dependent ensemble. These methods identified the impedance parameters (mass, damping, and stiffness) in the normal direction with an error of less than 5% in the simulations and 20% in the robot-only case. Measurements of a human subject yielded unacceptable errors, to the extent that the sign of the stiffness was uncertain. However, in the tangential direction, a fundamental limitation was discovered in practice: impedance may be

so low that the natural frequency of impedance at any one configuration overlaps the frequency content of the non-stationary noise process present in the zero-force trajectory. The lack of frequency separation poses a fundamental challenge to system identification of mechanical impedance.

In the crank-turning experiments the subjects' wrist was braced, the elbow was supported by a sling, and the shoulder was strapped to the chair. This confined the task to two degrees of freedom. While this is advantageous for modeling, it drastically simplifies human behavior. In Chapter 5 [Hermus et al., 2022], a robotic experiment was designed to probe how kinematic complexity may be managed. We assumed the impedance and programmed a Kuka LBR iiwa to interact with the InMotion planar robot while it simulated a virtual circular constraint. Interestingly we found that when the nullspace dimension increased, impedance superposition was comparable to nullspace projection. This suggests that dimensionality may be a 'blessing' not a 'curse'.

In Chapter 6 [West et al., 2022], we questioned an assumption so fundamental it is often assumed. Can humans regulate force during motion? Unlike the previous experiments in this thesis, in which subjects controlled motion under a kinematic constraint, this experiment prescribed motion and measured the force humans generated. In the experiments, subjects were surprisingly bad at regulating force during motion (with a target force of 5N subjects frequently exerted force as large as 10-15N). This suggests that when designing a human-robot interaction paradigm, designers should not expect subjects to precisely regulate force during motion. Moreover subjects' force errors correlated strongly with position. This indicates a coupling between force and motion consistent with actions composed of dynamic primitives.

## 7.1 Areas of Potential Future Work

The system ID methods in Chapter 4 were applied in both the normal and tangential directions. These cases represent two extremes. In the tangential direction, there was no stiffness resisting motion. In this direction the stiffness must be generated by the

human. However, in this task, the dynamics of the human limb result in a natural frequency less than 1 Hz. This made it hard to achieve the frequency separation required for this method. However, the normal direction was an approximation of constrained motion. In this case, the addition of the robot dynamics to that of the human dynamics increased the natural frequency of the combined system to 3-5 Hz. This higher stiffness of the interactive dynamics resulted in better estimates as frequency separation was possible. This suggests future work may be successful in the middle ground where subjects interact with simulated objects with a stiffness larger than zero but less than the constrained task of 1500 N/m.

The feedforward force production of the InMotion presents quantifiable direction-dependent hysteresis (see Appendix F). Experimentally, this caused problems with the implementation of the system ID methods in the InMotion human experiments. Two potential directions of future work are to (1) replace hardware components (e.g. the actuator or amplifiers) or (2) develop compensation via control.

Previous work has developed compensators for human experiments on the InMotion. In experiments designed to measure static limb stiffness, Palazzolo et al. [Palazzolo, 2005, Palazzolo et al., 2007] estimated robot dynamics and removed it from the combined human and robot dynamics. This enabled the estimation of the off-diagonal frequency responses. A key aspect in the design of Palazzolo et al.'s experiment was that the robot position was nearly static. The motion case, which involves position dependent friction, is more challenging. Later work by Thorup developed a sliding-mode control to compensate for the inertia of the InMotion during circular constrained motion [Thorup, 2018]. However, this approach was sensitive and quantifying its performance required care. Thus, it was not employed in the experiments presented herein.

To perform experiments with substantial motion, a nonlinear compensator may be required to eliminate the feedforward hysteresis. In an experiment with stochastic perturbations, compensating for hysteresis may be more challenging than inertial compensation. Model-based or learning methods may be able to achieve the required performance. However, the design and validation of such a system is not trivial. The



design of a mechanical system capable of exerting a precise force or impedance on the InMotion during circular constrained motion (or even linear motion) may allow for a control-based compensator to be validated prior to performing human experiments.

In this work the map from feedforward dynamics to measured displacements was used to estimate the dynamics of the combined human and robot system. To investigate the human only system, future work may incorporate the force transducer information. This makes it possible to change the interaction port. In this case, the map from measured interaction force and position displacements could be used for system ID.

An area of future work, related to the investigation of kinematic redundancy, is to increase the nullspace dimension by increasing the joint-space degrees of freedom. In the experiments, the nullspace dimension was increased by decreasing the task-space degrees of freedom. However, the dimension of the nullspace between task space and joint space can also be increased by increasing the joint degrees of freedom. Preliminary simulation-based work showed that the result is robust to this application. However, there may be value in connecting the theoretical understanding of an n-link manipulator to the stiffness of a continuum mechanical beam. This may provide valuable insight in understanding task conflicts in soft robotics.

## 7.2 Speculation

After spending the last 6 years studying human motor control, there is something simple and beautiful about human physical interaction. The more I learn about manipulation research in robotics, the more I am surprised by how the problems are often approached.

While there has been growing success in end-to-end machine learning, when it comes to human-like tool use or manipulation there are still strides to be made. This makes me think that machines could benefit from an understanding of humans.

If it is not possible to end-to-end learn it all, is there a way to understand some structures so that machine learning can be used to build upon those structures or to

decide which structures work? In essence, could dynamic primitives be a basis for a manipulation language? For example, consider letters: Note that we can spell quite a few words with only a handful of letters. It seems like these letters could compose simple actions such as moving the arm and closing the hand to pick up an object – maybe these simple tasks are a type of word, say a verb. However, the order of words is constrained by the rules of grammar. By analogy, the rules of physics enforce that certain things are not possible or cannot be observed. Moreover, composing a grammatically correct sentence does not imply that the sentence has meaning. You can move a manipulator around all you want following the laws of physics but that does not mean you achieved any type of goal.

Consider a simple perception example, estimating the pose of an object. Even this can be a hard problem in manipulation. Here is an example from Russ Tedrake’s class that got me thinking. It went something like this: consider a box that arrived in the mail sitting on the floor. How many points do you need to find and represent it? Maybe you need one point for the center. Maybe you need three points to specify the orientation. Maybe you need 8 to represent all the corners. I could not help but think, maybe you do not need any because you just woke up and want to go to the kitchen to get coffee which has nothing to do with the box. This type of question seems silly. It seems that, depending on what I want to do, the representation changes.

I think an important thing about the idea of perception is that the way we want to act on representation matters. Thus, the action and perception primitive are somehow related. Perception primitives may be analogous to perception/representation as nouns and actions as verbs. In this case, it is consistent in the way that not all verbs can act on all nouns. When it comes to action and perception, the understanding and goal alter the way the same word is used. If we look at the sentence, ‘I took the dog for a walk.’ The word walk means something on its own but its context is drastically increased when we know it acts upon the word ‘dog’. This is analogous to the previous perception discussion of the box. The number of key points or representation of the box can be entirely different whether I want to avoid it, whether I want to pick it up, or whether I want to open it.

I can not help but wonder if this type of question may be important to consider before trying to find new letters or even words. If we have an expert speller (an engineer) who can correctly spell nearly any word (develop a fancy manipulation task) it would take nearly 20 years of a career to develop the skills required to make these words. Yet if the goal of this expert speller is to do philosophy, reason about complex ideas by composing arguments with sentences and paragraphs—maybe even read multiple books and have an understanding of these arguments—it may seem silly to first spend 20 years teaching all of our philosophers to be expert spellers. Moreover, people who are brilliant at philosophy may not be able to spell well.

Is there a way to tease out a language for manipulation/interaction? I might go as far as to say that we know some of the letters and many advanced research labs have been able to develop some of the words. I wonder, can we find and leverage this language? If so, can individual small sub-pieces (words and sentences) be formulated while following the grammatical structure and writing meaningful sentences? It seems that even kindergarten-level reading may be enough to achieve great things. Especially if this structure allows the system to sound out new words in order to try to figure out how to spell or read words it has never seen before. I am not suggesting modern robots should only be controlled solely by imposing absolute control constrains, e.g. dynamic primitives. However, I do believe that control with dynamic primitives may enable simple stable interaction which can be refined, via optimization and learning. Given the value of a descriptive model of human motor control based on dynamic primitives, areas of future work about which I am especially curious is exoskeleton control as well as task and motion planning [Garrett et al., 2021].

### 7.3 Closing Remarks

As my experience at MIT draws to a close, I am filled with a sense of awe for the truly remarkable design of the control strategies and the hardware that enable humans to accomplish so much more than survival. It is clear that, barring technological advances in methodology or sensing, some of the elusive questions in human motor

control may forever remain a mystery. However, it is also apparent that there are countless critical future areas of inviting investigation. It is my hope that the ideas herein will compose another building block in the foundation of the growing understanding of human motor control.

# Appendix A

## Modeling a Two-Link Manipulator Coupled to a Crank

Knowledge of the subject's inertia is required to develop a model and simulate the system. Inertial parameters were estimated based on the results of the cadaver studies of Dempster [Miller and Nelson, 1973, Plagenhoef, 1971]. The arm and forearm plus hand were denoted by *links* 1 and *link* 2, respectively. Each of the body segments was described by the following parameters: length,  $l$ ; mass,  $m$ ; inertia,  $I$ ; and center of mass,  $c$ . The length  $l_1$  was the distance from the shoulder to the elbow. The length  $l_2$  was the distance from the elbow to the center of the fist. *Link*2 was considered a combination of the forearm and hand. The hand was assumed to be a point mass at the end of the forearm. Limb parameters are reported in Table A.1. The inertia of the crank about the pivot was  $3.716 \times 10^{-3}$  kg/m<sup>2</sup>. Analysis of the mean and median of all subjects' inertial parameters showed that *subject* 1 was representative [Doeringer, 1999].

The model of the arm and crank system was constructed in the same manner as performed by Ohta et al. [Ohta et al., 2004]. Figure A-1 displays the variables and notation used in the development of the model. The system has one degree of freedom; therefore, there is always a kinematic relationship which can be used to transform from Cartesian position,  $\mathbf{x} = [x, y]^T$ , to joint position,  $\mathbf{q} = [q_1, q_2]^T$ , and to

crank position,  $\theta$ , where the center of the crank is defined as  $\mathbf{x}_c = [x_c, y_c]^T$ .

$$\mathbf{x} = \begin{bmatrix} l_1 C_1 + l_2 C_{12} \\ l_1 S_1 + l_2 S_{12} \end{bmatrix} = \begin{bmatrix} r \cos \theta \\ r \sin \theta \end{bmatrix} + \mathbf{x}_c, \quad (\text{A.1})$$

The notations  $S_1, C_1$  denote  $\sin(q_1), \cos(q_1)$  and  $S_{12}, C_{12}$  denote  $\sin(q_1 + q_2), \cos(q_1 + q_2)$ . The radius of the crank is  $r$ , the damping of the crank is  $b_c$ , and the inertia is  $I$ . The upper arm denoted 1, and the forearm denoted 2 are described by length  $l_1, l_2$ , mass  $m_1, m_2$ , inertia about the  $z$  axis  $I_1, I_2$ , and center of mass distance from the joint axis  $c_1, c_2$ . The force on the handle is  $\mathbf{F} = [F_x, F_y]^T$ , with the normal unit vector,  $n$  and tangential unit vector,  $e$ . The joint torque is denoted  $\boldsymbol{\tau} = [\tau_1, \tau_2]^T$ .




Figure A-1: Model of crank rotation task which displays the sign convention and notation used in the computations.

From the sum of moments acting on the crank,

$$I\ddot{\theta} + b_c\dot{\theta} = r\mathbf{e}^T\mathbf{F} \quad (\text{A.2})$$

summation of moments about the shoulder,

$$\mathbf{M}\dot{\mathbf{q}} + \mathbf{h} = \boldsymbol{\tau} - \mathbf{J}^T \mathbf{F} \quad (\text{A.3})$$

and the kinematic relationship that equates the acceleration at the handle to the acceleration at the hand,

$$\ddot{\mathbf{x}} = \mathbf{J}\ddot{\mathbf{q}} + \dot{\mathbf{J}}\dot{\mathbf{q}} = r(\ddot{\theta}\mathbf{e} - \dot{\theta}^2\mathbf{n}) \quad (\text{A.4})$$

a model of the system can be constructed. Parameters comprising these equations include the mass matrix,

$$\mathbf{M}(\mathbf{q}) = \begin{bmatrix} m_1c_1^2 + m_2(l_1^2 + c_2^2 + 2l_1c_2C_2) + I_1 + I_2 & m_2(c_2^2 + l_1c_2C_2) + I_2 \\ m_2(c_2^2 + l_1c_2C_2) + I_2 & m_2c_2^2 + I_2 \end{bmatrix} \quad (\text{A.5})$$

the centrifugal and Coriolis forces,

$$\mathbf{h}(\mathbf{q}, \dot{\mathbf{q}}) = \begin{bmatrix} -m_2l_1c_2S_2(2\dot{q}_1\dot{q}_2 + \dot{q}_2^2) \\ m_2l_1c_2S_2\dot{q}_1^2 \end{bmatrix} \quad (\text{A.6})$$

and the Jacobian relating unconstrained differential arm motions to hand motions.

$$\mathbf{J}(\mathbf{q}) = \begin{bmatrix} -(l_1S_1 + l_2S_{12}) & -l_2S_{12} \\ l_1C_1 + l_2C_{12} & l_2C_{12} \end{bmatrix} \quad (\text{A.7})$$

From Equations 9, 10, and 11 the relationship in Equation 15 can be shown

$$R(\theta)\dot{\theta} + H(\theta, \dot{\theta}) = r\mathbf{e}^T \mathbf{J}^{-T} \boldsymbol{\tau}, \quad (\text{A.8})$$

where the configuration dependent damping is

$$R(\theta) = I + r^2\mathbf{e}^T \mathbf{J}^{-T} \mathbf{M} \mathbf{J}^{-1} \mathbf{e}, \quad (\text{A.9})$$

and the configuration dependent inertia is

$$H(\theta, \dot{\theta}) = b_c \dot{\theta} + r \mathbf{e}^T \mathbf{J}^{-T} \{ \mathbf{h} - \mathbf{M} \mathbf{J}^{-1} (r \dot{\theta}^2 \mathbf{n} + \dot{\mathbf{J}} \dot{\mathbf{q}}) \}, \quad (\text{A.10})$$

From Equations 9, 10, and 11 we can also solve for  $F$ ,

$$\mathbf{F} = \{ \mathbf{J} \mathbf{M}^{-1} \mathbf{J}^T + r^2 I^{-1} \mathbf{e} \mathbf{e}^T \}^{-1} \{ \mathbf{J} \mathbf{M}^{-1} (\boldsymbol{\tau} - h) + \dot{\mathbf{J}} \dot{\mathbf{q}} + r \dot{\theta} (\dot{\theta} \mathbf{n} + b_c I^{-1} \mathbf{e}) \} \quad (\text{A.11})$$



Subject	$l_1$ (cm)	$c_1$ (cm)	$m_1$ (kg)	$I_1$ (kg-m <sup>2</sup> )	$l_2$ (cm)	$c_2$ (cm)	$m_2$ (kg)	$I_2$ (kg-m <sup>2</sup> )
1	36.8	16.1	2.23	0.0313	37.5	23.6	2.25	0.0440
2	36.2	15.8	2.03	0.0276	36.8	23.6	2.10	0.0395
3	32.4	14.1	2.20	0.0239	33.0	20.7	2.23	0.0340
4	36.2	15.8	2.61	0.0354	38.1	23.5	2.55	0.0515
5	34.9	15.2	2.03	0.0257	35.6	22.7	2.10	0.0369
6	29.2	12.7	2.35	0.0208	33.7	20.8	2.35	0.0375
7	37.5	16.3	1.97	0.0287	37.5	24.0	2.05	0.0401
8	34.3	15.0	1.97	0.0240	36.8	23.7	2.05	0.0385
9	36.8	16.1	2.03	0.0286	38.7	24.9	2.10	0.0435
10	34.3	15.0	2.29	0.0279	33.7	21.1	2.30	0.0364

Table A.1: Human limb inertia properties for this experiment.



# Appendix B

## Ensemble Statistics – Assumptions and 1D Simulations

### B.1 Review of Assumptions

A brief didactic review of system identification concepts will be presented in order to develop the context of the proposed method, show that the new method is required, and articulate how it is novel.

#### B.1.1 Probability Density Functions

When a value is not known with precision it can be treated as a random variable. A random variable may take any value of the set  $\{x\}$ . The random variable can be characterized by a cumulative distribution function (CDF)

$$F(x_0) = P(x \leq x_0) = \text{probability that } x \leq x_0. \quad (\text{B.1})$$

If the random variable is continuous, it is often easier to work with the probability density function (PDF) defined as the derivative of  $F$  with respect to  $x$ .

$$p(x, t) = \frac{\partial F}{\partial x} \quad (\text{B.2})$$

The probability that a random variable lies between two values is the integral of its probability density function over the range between those values.

$$P(x_1 < x \leq x_2) = \int_{x_1}^{x_2} p(x)dx \quad (\text{B.3})$$

The probability density function is also known as the likelihood function. A random process is the extension of the idea of a random variable to that of a random function of time. In principle, a random process may be characterized by the likelihood function  $p(x, t)$

$$P(x(t) \leq x_o) = \int_{-\infty}^{x_o} p(x)dx \quad (\text{B.4})$$

In practice, partial descriptors are defined using the expectation operator. The first moment is called the mean

$$\mu(t) = \bar{x}(t) = E\{x(t)\} = \int_{-\infty}^{\infty} x(t)p(x, t)dx. \quad (\text{B.5})$$

For the second and higher moments both the moment and central moment can be defined. The central moment of a distribution measures the deviation from the mean, while the moments of a distribution measure the deviation of the distribution itself. The second moment is the mean square

$$\overline{x^2}(t) = E\{\overline{x^2}(t)\} = \int_{-\infty}^{\infty} x^2(t)p(x, t)dx. \quad (\text{B.6})$$

The second central moment is the variance

$$\sigma^2(t) = \overline{x^2}(t) - \bar{x}^2. \quad (\text{B.7})$$

This concept can be extended to multiple random variables. The joint likelihood function  $p(x_1, t_1, x_2, t_2)$  is defined as,

$$P(x(t_1) \leq x_{1o} \text{ AND } x(t_2) \leq x_{2o}) = \int_{-\infty}^{x_{1o}} \int_{-\infty}^{x_{2o}} p(x_1, t_1, x_2, t_2)dx_1dx_2. \quad (\text{B.8})$$

Higher order joint likelihood functions, relating three or more times, may be defined but are difficult to estimate. An especially important joint moment is the cross-correlation defined as,

$$\phi_{x_1x_2}(t, \tau) = E\{x_1(t)x_2(t + \tau)\} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x_1x_2p(x_1, t, x_2, t + \tau)dx_1dx_2. \quad (\text{B.9})$$

Without exact knowledge of the probability density function, a statistic, an estimate of the moments or central moments of a distribution, can be computed from a finite set of observations. Given a sample, there are multiple ways to compute a statistic. Two general approaches which will be discussed herein are: ensemble methods and stationary methods.

### B.1.2 Ensemble Statistics

An ensemble is an idealization where multiple, usually many, realizations of a system are considered all at once. Each realization represents a possible state of the real system. This approach implies that each realization results from the same underlying likelihood function of the same random process.




Figure B-1: Ensemble of time history records defining a random process. The red, blue, and yellow sample records represent different exemplary replications  $x_r$ .

Given an ensemble, an unbiased estimate of the mean and autocorrelation can be

computed using an approximation of the expectation operator

$$\mu_x(t_1) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N x_k(t_1) \quad (\text{B.10})$$

$$\phi_{xx}(t_1, t_1 + \tau) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=1}^N x_k(t_1)x_k(t_1 + \tau). \quad (\text{B.11})$$

### B.1.3 Stationary Statistics

If the complete set of joint likelihood functions is time-shift invariant, such that  $p(x, t + \tau) = p(x, t)$ ,  $p(x_1, t_1 + \tau, x_2, t_2 + \tau) = p(x_1, t_1, x_2, t_2)$ , and so on, the random process is strictly stationary. In practice, if the first and second moments are time-shift invariant the process is considered wide-sense stationary. In this case an unbiased estimate of the mean and autocorrelation can be computed by

$$\mu_x(t_1) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T x_k(t) dt \quad (\text{B.12})$$

$$\phi_{xx}(\tau, k) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T x_k(t)x_k(t + \tau) dt \quad (\text{B.13})$$

The ergodic hypothesis posits that the expectation over an ensemble is the same as the time-average of a single replication of the ensemble. That is,  $\mu_x(k) = \mu_x$  and  $\phi_{xx}(\tau, k) = \phi_{xx}(\tau)$ . This assumption ties together the two approaches. In practice, many complex physical systems can be studied under the ergodic hypothesis. This is in part why many of the methods used to estimate human limb impedance have been performed in static posture with an additional stabilizing stiffness term.

## B.2 Input-Output Relations

An impulse response function (IRF), is the output of a dynamic system to an impulse input, a function defined by the Dirac delta distribution,

$$\delta(t) = \begin{cases} +\infty, & t = 0 \\ 0, & t \neq 0 \end{cases} \quad (\text{B.14})$$

It is described as a function of time lag from impulse onset as  $h(\tau)$ . Since the impulse function contains all frequencies, the impulse response function defines the entire linear time-invariant system. The goal of most time-domain non-parametric system identification approaches is to identify one or more IRFs.

In steady state, the response of a linear time-invariant system to a sinusoidal input is a sinusoid of the same frequency, which may have different magnitude and phase. The frequency response function,  $H(f)$ , describes this input output relationship. It completely characterizes a linear time-invariant system. There is a one-to-one map, called the Fourier transform, between the impulse response function and the frequency response function. Consider the linear time-invariant system with an impulse response function  $h(\tau)$  and a frequency response function  $H(f)$ . Assume the system subject to input  $u(t)$  from a stationary random process  $u(t)$  and produces an output  $z(t)$ , which will also belong to a stationary random process  $z(t)$ . The output is given by,

$$z(t) = \int_0^{\infty} h(\tau)u(t - \tau)d\tau \quad (\text{B.15})$$

Where  $h(\tau) = 0$  for  $\tau < 0$  for physically realizable systems. The product  $z(t)z(t + \tau)$  is given by

$$z(t)z(t + \tau) = \iint_0^{\infty} h(\alpha)h(\beta)u(t - \beta)u(t + \tau - \alpha)d\alpha d\beta \quad (\text{B.16})$$

Taking the expectation of both sides, substituting the output autocorrelation function

$\phi_{zz}(\tau)$ , and substituting the input autocorrelation function  $\phi_{uu}(\tau)$  yields,

$$\phi_{zz}(\tau) = \iint_0^\infty h(\alpha)h(\beta)\phi_{uu}(\tau + \beta - \alpha)d\alpha d\beta \quad (\text{B.17})$$

Likewise, the product  $u(t)z(t + \tau)$  is given by,

$$u(t)z(t + \tau) = \int_0^\infty h(\alpha)u(t)u(t + \tau - \alpha)d\alpha \quad (\text{B.18})$$

Again, taking the expectation of both sides and substituting the input output cross correlation function  $\phi_{uz}$  yields,

$$\phi_{uz}(\tau) = \int_0^\infty h(\alpha)\phi_{uu}(\tau - \alpha)d\alpha \quad (\text{B.19})$$

Assume the case where  $u(t)$  can be chosen such that it is not correlated with disruptive noise processes  $w(t)$  (e.g. actuator or measurement noise). In this case, when the expectation of Equation B.16 is computed, any additional noise terms would appear as a cross correlation with the input,  $\phi_{uw}(\tau)$ . By definition, we know that these quantities are uncorrelated. Thus, the estimate of the second moment and second joint moment is unbiased. This only leaves the input-output relation  $h(\tau)$  and suggest a possibility of unbiasedly estimating the impulse response function.

### B.2.1 Stationary Method for Estimating IRF

The impulse response of a damped system decays exponentially. For linear damped systems, after  $4\tau_{tc}$  (where  $\tau_{tc}$  is the time constant of the systems slowest mode) the system will remain within less than 2% of the steady state position. For further discussion of system dynamics and time constants, please see [Bendat and Piersol, 2010]. This means that in practice, an impulse response function can be approximated as a finite impulse response function. One way to perform this finite approximation is to estimate the discrete impulse response function. This convolution can be written as a matrix multiplication.



$$\begin{bmatrix} \phi_{uz}(0) \\ \phi_{uz}(1) \\ \vdots \\ \phi_{uz}(N) \end{bmatrix} = \begin{bmatrix} \phi_{uu}(0) & \phi_{uu}(-1) & \cdots & \phi_{uu}(-N) \\ \phi_{uu}(1) & \phi_{uu}(0) & \cdots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{uu}(N) & \vdots & \cdots & \phi_{uu}(0) \end{bmatrix} \begin{bmatrix} h(0) \\ h(1) \\ \vdots \\ h(N) \end{bmatrix} \quad (\text{B.20})$$

Using the Moore-Penrose pseudo inverse [Penrose, 1955] allows for the impulse response function to be unbiasedly estimated [Bendat and Piersol, 2010]. If the process is stationary such that the frequency response can be estimated the analogous operation is defined in the frequency domain. Computing the one-sided spectral density function yields,

$$G_{uz}(f) = H(f)G_{uu}(f) \quad (\text{B.21})$$

The frequency response function can be unbiasedly estimated by

$$H(f) = \frac{G_{uz}(f)}{G_{uu}(f)} \quad (\text{B.22})$$

Note these relations can be defined and applied to the stationary statistics or ensemble statistics presented above. For compactness, the stationary methods will not be presented herein. A discrete implementation of the ensemble-statistics-based method will be presented.

## B.2.2 Ensemble Method for Estimating IRF

For every time step  $i$ , the relationship between noise-free input,  $u_r(i)$ , and the corresponding noisy outputs,  $z_r(i)$ , is given by,

$$z_r(i) = y_r(i) + n_r(i) \quad (\text{B.23})$$

where  $y_r(i)$  is the true output and  $n_r(i)$  is additive white noise for the  $r^{\text{th}}$  realization of the ensemble data. The discrete-time relationship between an input  $u_r(i)$  and the corresponding output  $z_r(i)$  of a linear time-varying system is given by the discrete

convolution equation,

$$z_r(i) = \Delta t \sum_{j=0}^M \hat{h}(i, j) u_r(i - j) + n_r(i) \quad (\text{B.24})$$

where  $\hat{h}(i, j)$  is an impulse response function (IRF) estimate with finite lag length  $L = M + 1$  where  $j$  is a lag index. This assumes  $\hat{h}(i, j) = 0$  for  $j > M$ , which is true for causal systems. In practice, the lag length  $L$  should be long enough that the IRF estimate settles to a near zero value. This suggests a finite impulse response function is appropriate.

The Wiener-Hopf equation [Bendat and Piersol, 2010] can be obtained by multiplying both sides of Equation B.24 with  $u_r(i - j)$  and summing over all realizations ( $R$ ),

$$\frac{1}{R} \sum_{r=1}^R z_r(i) u_r(i - k) = \Delta t \sum_{j=0}^M \hat{h}(i, j) \frac{1}{R} \sum_{r=1}^R u_r(i - j) u_r(i - k) + n_r(i) \quad (\text{B.25})$$

To identify the system dynamics at a specific time  $i$ , the input-output relation can be evaluated using the data across  $R$  realizations and across  $M$  time lags. Equation B.25 can be rewritten with an input-output cross-correlation function estimate ( $\hat{\phi}_{zu}$ ) at time  $i$  and  $i - k$ , and an input auto-correlation function estimate ( $\hat{\phi}_{uu}$ ) at time  $i - k$  and  $i - j$ ,

$$\hat{\phi}_{zu}(i, -k) = \Delta t \sum_{j=0}^M \hat{h}(i, j) \hat{\phi}_{uu}(i - k, k - j) \quad (\text{B.26})$$

A matrix expression can be constructed by changing the lag index  $j$  and  $k$  from 0 to  $M$ , where  $\hat{\Phi}_{uu}(i) \in \mathbb{R}^{L \times L}$  matrix and  $\hat{\Phi}_{zu}(i) \in \mathbb{R}^{L \times 1}$  vector. In this case the IRF estimate,  $\hat{\mathbf{h}}(i)$ , can be computed using a Moore-Penrose pseudo-inverse matrix operation.

$$\hat{\Phi}_{zu}(i) = \Delta t \hat{\Phi}_{uu}(i) \hat{\mathbf{h}}(i) \quad (\text{B.27})$$

$$\hat{\mathbf{h}}(i) = \Delta t \hat{\Phi}_{uu}(i)^{-1} \hat{\Phi}_{zu}(i) \quad (\text{B.28})$$

The IRF estimates which result from this matrix inversion are noisy. In practice, the

IRFs are filtered in both the arguments – lag and time. This approach does require that the system is stable such that the impulse response declines to zero given a computationally tractable value of lag. This is a nonparametric method, which means that the estimation of the IRF requires no further assumption of structure in the data – other than linearity. Visual inspection can provide valuable insight into the system order and types of linear models which may describe the response. If parametric measures are required, once an IRF estimate is produced, system structure can be assumed and a model can be fit to estimate the parameters such as mass, damping, and stiffness.

In practice, a process that is nonergodic, can sometimes be studied under a subset of conditions. This subset is frequently chosen as a compromise between providing insight, while minimizing bias in the statistics (e.g. estimating stiffness in multiple static postures instead of during motion [Mussa-Ivaldi et al., 1985]). In other cases this is not possible. A second approach is to post-process the data, such that the statistics are unbiased. This will be explored in the following examples.

### B.2.3 Example – Time-Varying $x_0(t)$ and Constant $Z$

One way to violate the assumption of stationarity is to try to identify a Norton-type network in the impedance operational form

$$f = Z\{x_o(t), x(t)\} \tag{B.29}$$

where the impedance is constant in time. One of the simplest systems to exhibit this behavior may be,

$$m\ddot{x} = k(x_0 - x) + b(\dot{x}_0 - \dot{x}) + u \tag{B.30}$$

a mass, spring, and damper system with a force acting on it. In this case,  $x_0(t)$ , the end of the spring, is strictly a function of time. The perturbation force  $u(t)$  is a zero-mean random perturbation. This is an example of a multi-input single-output system with non-stationary input  $x_0(t)$  and stationary input  $u(t)$ . One approach to

understand input-output relations is using the Laplace transform.

$$X(s) = \left( \frac{bs + k}{ms^2 + bs + k} \right) X_0(s) + \left( \frac{1}{ms^2 + bs + k} \right) U(s) \quad (\text{B.31})$$

When these inputs pass through the linear system, the output  $x(t)$  will be non-stationary. However, in this case the violation of stationarity has a particular structure. An approach to remove the nonstationary influence of  $x_0(t)$  is to compute the mean across the ensemble

$$\bar{x}_k(i) = \frac{1}{R} \sum_{r=1}^R x_r(i) \quad (\text{B.32})$$

and subtract it (see Figure B-2 middle vs. right). This removes the influence of  $x_0$ , leaving only the input-output relation between the perturbation and measured position  $X(s)/U(s)$ . This estimate has zero mean and constant variance, due to the deterministic nature of the system dynamics. The estimate is stationary in the wide sense. Furthermore, the underlying random process is also a function of time. This result satisfies the ergodic assumption. Either ensemble or stationary methods can be used to produce an unbiased estimate of the dynamics in the form of an impulse response function or frequency response function. Below is an example using the ensemble methods to estimate the IRF.

Simulations were performed at a frequency of 200 Hz, a duration of 2 seconds, a lag length of 0.2 seconds, and 300 realizations. The mass was 1 kg; the damping was 40 Ns/m, and the stiffness was 2500 N/m. In practice,  $x_r$  cannot be measured exactly. For this reason, the noisy measurement of position will be denoted as  $z_r$ . The amplitude of the zero-mean Gaussian additive noise was chosen to result in a signal to noise ratio of 9 dB between  $x_r$  and  $n_r$ . The perturbation  $u_r$  switched between -1 and 1 N randomly, this produced a flat frequency spectrum up to 40 Hz. The ensemble methods for estimating IRF, presented in the previous section, were applied. The impulse response functions were averaged in the temporal direction with 100 ms windows.<sup>1</sup>

---

<sup>1</sup>The  $z_r(i) - \bar{z}(i)$  is stationary in the time coordinate  $i$ . For this reason, averaging over all time, not just a local 100 ms window, could have been easily justified. Alternatively, this same estimate

For each time step  $i$ , the best fit second order impulse response function  $h_{model}(i)$  was fit by minimizing the mean squared error between  $\hat{h}(i)$  and the model using a bound-constrained optimization. This model is described by the parameters for mass  $\hat{m}(i)$ , damping  $\hat{b}(i)$ , and stiffness  $\hat{k}(i)$ . To quantify the model fit the  $\%VAF_{IRF}$  between  $\hat{h}(i)$  and  $h_{model}(i)$  was assessed.

$$\%VAF_{IRF}(r) = 100 \times \left( 1 - \frac{\text{var}\left(\hat{h}(i) - h_{model}(i)\right)}{\text{var}\left(\hat{h}(i)\right)} \right) \quad (\text{B.33})$$

The simulated system behavior is presented in Figure B-2. Note, in the middle and right panels of Figure B-2, that the subtraction of the mean across the ensemble removes the influence of the time-varying  $x_0$ . In Figure B-3, the IRF estimates with respect to time and lag are presented. Lastly, in Figure B-4, the estimates and variance-accounted for are presented and compared.




Figure B-2: (left) Example force perturbation. (middle) The measured position is denoted  $z_r$ . Each  $r^{th}$  realization was represented by a different color line. The mean of  $z_r(i)$  over replications was the thick black line. (right) Measured position with mean subtracted,  $z_r - \bar{z}$ .

---

could have been performed with a stationary statistic as well. However, this analysis was chosen for the sake of consistence with future examples.




Figure B-3: The red dots represent the impulse response function estimates  $\hat{h}$  as a function of time. The blue lines represent the best fit IFR model,  $h_{model}$ . Note, for the sake of visualization, every  $10^{th}$  estimate in the time direction was plotted.

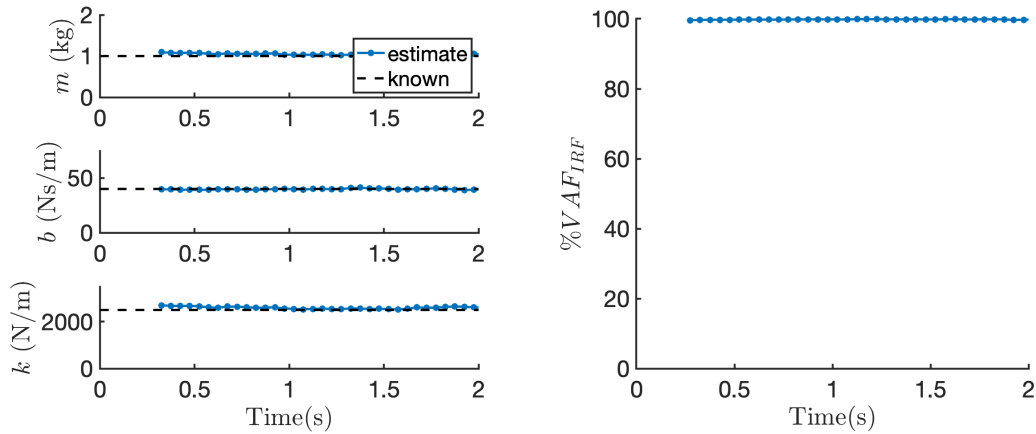


Figure B-4: (Left) Blue dots represent the estimated parameters of  $h_{model}$ . The dashed black lines represent the known values. (Right) The percentage variance accounted for %VAF<sub>IRF</sub>.

### B.2.4 Example – Time-Varying $x_0(t)$ and Time-Varying $Z(t)$

If the system has a time-varying  $x_0(t)$  and a time-varying impedance  $Z(t)$ , the mean across realizations can still be computed and subtracted. This yields an ensemble of signals with zero mean and time-varying second moment. However, the time-varying second moment violates the stationarity assumption. This precludes the use of the stationary methods for computing unbiased statistics. However, this does not

preclude the use of ensemble methods. The ensemble method has been the key to the identification of ankle impedance [Lortie and Kearney, 2001, Ludvig et al., 2011, Lee and Hogan, 2015, Guarín and Kearney, 2017, Guarín and Kearney, 2018].

The same simulation methods from the previous section were applied. However, in this case, the impedance changed as a function of time such that,

$$k(t) = 2500 + 500 \sin(\pi t) \quad (\text{B.34})$$

$$b(t) = 40 + 10 \cos(\pi t) \quad (\text{B.35})$$

with units of N/m and Ns/m respectively.

In this case, it is clear from Figure B-5, that subtracting the mean over replications results in a zero mean signal. However, the variance changes with respect to time. Thus, this signal is not stationary. For this reason, ensemble statistics must be used for the system identification.

In Figure B-6, it is qualitatively evident that the impulse response function varies with respect to time. Despite this lack of stationarity, the estimates presented in Figure B-7 are close to the expected value. Furthermore, the  $\%VAR_{IRF}$  was also high. Consistent with previous works, the ensemble methods work well for identifying the dynamics of a time-varying system in the presence of substantial noise. Unfortunately, this was not directly applicable to upper limb crank turning.




Figure B-5: (Left) Example force perturbation. (Middle) The measured position denoted  $z_r$ . Each  $r^{th}$  realization was a different color line and the mean of  $z_r$  over replications is the thick black line. Clearly, the perturbations do not account for the majority of the variance in the signal. This is because of the time-varying  $x_0$ . (Right) Measured position with mean subtracted,  $z_r - \bar{z}$ .




Figure B-6: The red dots represent the impulse response function estimates  $\hat{h}$  as a function of time. The blue lines represent the best fit IFR model,  $h_{model}$ . Note, for the sake of visualization every  $10^{th}$  estimate in the time direction was plotted.

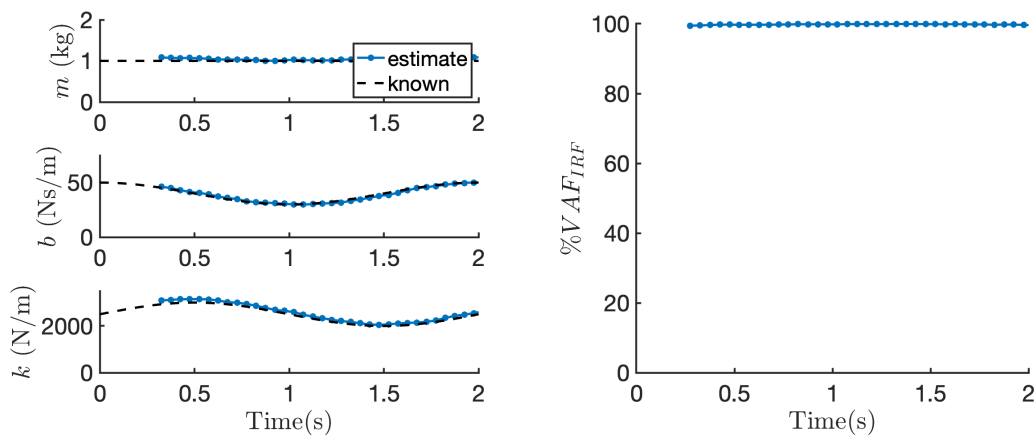


Figure B-7: (Left) Blue dots represent the estimated parameters of  $h_{model}$ . The dashed black lines represent the known values. (Right) The percentage variance accounted for between  $\%VAF_{IRF}$ .



# Appendix C

## Evidence of Brownian Noise – Static Posture

This work reports evidence of Brownian Noise in static human upper limb control. It is not currently published. This work was done in collaboration with Rika Sugimoto Dimitrova and Federico Tessari.

### C.1 Introduction

To ensure the observation of drift present in the mechanical crank-turning experiments was not an artifact, a simple experiment was developed to determine if Brownian noise was present in the control of the upper limb even in an extremely simple case.

### C.2 Methods

One healthy right-handed male subject participated in the study. The subject did not report any biomechanical injury to their arm nor any neurological problems. Prior to participating in the study, he was informed about the experimental procedure and signed the informed consent document approved by MIT's Institutional Review Board. In this experiment, the subject was instructed to hold the handle of the InMotion robot at a constant position for 10 minutes. During these experiments the

elbow was supported by a sling, the wrist was braced, and the hand was occluded from view consistent with the prior crank-turning experiment presented in Chapter 2 and 3. Position was measured by the encoders of the InMotion robot. The actuation torque was set to zero. Experiments were collected with and without visual feedback provided on a display.

The data were re-sampled to a constant sampling frequency of 200 Hz. Then the power spectral density of the position in the  $x$  and  $y$  direction was estimated using Welch’s method [Welch, 1967] (11 windows with 50% overlap).

### C.3 Results

Drift was visually evident in the position measurements with and without visual feedback (see Figure C-1). Consistent with a Brownian noise process, a slope of -20 dB/dec was observed in the power spectrum (see Figure C-2).




Figure C-1: Raw position measurements from InMotion robot when a subject held the handle for 10 minutes (left) visual feedback and (right) no visual feedback. Note the scale is much smaller in the left visual feedback plot than in the right no visual feedback plot.

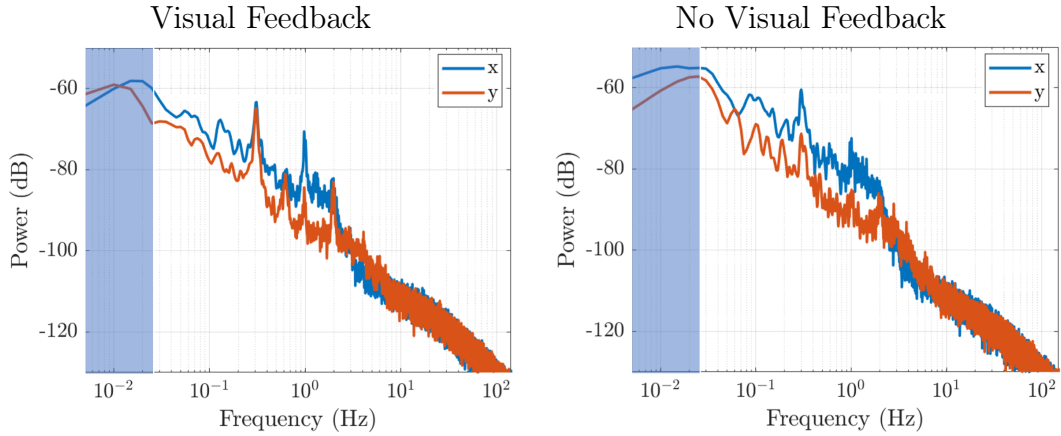


Figure C-2: Power spectral estimates for the (left) visual feedback and (right) no visual feedback conditions are presented. The blue lines denote the  $x$  direction and red lines denote the  $y$  direction. The region shaded in purple represents the region above the 0.01 Hz high pass filter.

## C.4 Discussion and Conclusions

Even in this simple experiment without instructed motion, drift consistent with that of Brownian noise was present. This finding led to key insight in the design of the system identification work (see Chapter 4).

Doeringer and Hogan [Doeringer and Hogan, 1998a] demonstrated that movement intermittency was in the forward path and could not be bypassed. This observation, that visual feedback does not eliminate drift, is consistent with their results.



# Appendix D

## InMotion Robot–Inertial Model

This model was developed by Thorup and is presented in her Master’s thesis [Thorup, 2018]. The InMotion2 robot has four links, shown in Figure D-1. The motors are located at the labeled motor axes and provide torque to links 1 and 4. The angle of the ‘shoulder’ joint measured by the encoder in motor 1 is the angle between link 1 and the motor axis. Link 2 is parallel to link 4 so the angle of the ‘elbow’ joint measured by the encoder in motor 2 is the angle between link 4 and the motor axis.




Figure D-1: Inertial model

The lengths  $l$  and masses  $m$  are presented in Table D.1. The variable  $r$  is the distance between the center of mass of the link and the joint closest to the motor axis. The moments of inertia about the center of mass  $I_c$  and the moments of inertia about the joint closest to the motor axis  $I$  for each component are calculated assuming that

the links are uniform rods and the handle is a point mass. This yields the following joint space mass matrix,

$$\mathbf{M}_\theta(\boldsymbol{\theta}) = \begin{bmatrix} I_1 + m_2 l_1^2 + I_3 & (m_2 l_1 r_2 + m_3 l_4 r_3) \cos(\theta_1 - \theta_2) \\ (m_2 l_1 r_2 + m_3 l_4 r_3) \cos(\theta_1 - \theta_2) & I_2 + m_3 l_4^2 + I_4 \end{bmatrix} \quad (\text{D.1})$$

Property	link 1	link 2			link 3	link 4
		link	handle	total		
length ( <i>m</i> )	0.4064	0.5144	0.5016	N/A	0.4064	0.1555
mass ( <i>kg</i> )	0.756	0.892	1.072	1.964	0.756	0.378
<i>r</i> ( <i>m</i> )	0.2032	0.2572	0.5016	0.3906	0.2032	0.0775
<i>I<sub>c</sub></i> ( <i>kg * m<sup>2</sup></i> )	0.0104	0.0197	0	0.0488	0.0104	0.0007
<i>I</i> ( <i>kg * m<sup>2</sup></i> )	0.0416	0.0787	0.2697	0.3484	0.0416	0.0030

Table D.1: InMotion Link Properties

# Appendix E

## InMotion Robot– Static Feed-forward Force Production

### E.1 Introduction

Prior to performing the virtual crank-turning experiments, tests were performed to determine if the feedforward dynamics of the robot were sufficient or whether a dynamic compensation was required in the control.

### E.2 Methods

In the experiment the robot end-effector was clamped in a vice (see Figure E-1). Two types of feedforward force perturbations were applied. The first case, denoted ‘small’, applied a  $\pm 1.5$  N perturbation in the  $x$  direction and a  $\pm 0.5$  N perturbation in the  $y$  direction. This is consistent with the amplitudes of the perturbations used in Chapter 4 in the normal and tangential direction respectively. In the second case, denoted ‘large’, the amplitude in both the  $x$  and  $y$  direction was  $\pm 5$  N. Consistent with Chapter 4, the amplitude of the perturbation switched randomly between the positive and negative values. The frequency spectrum of the input perturbations was flat up to 20 Hz. For each condition, measurements were recorded for 10 minutes. Spectra were estimated using 330 hamming windows Each window was two seconds

long. Magnitude and partial coherence estimates were computed consistent with the methods of Bendat and Piersol [Bendat and Piersol, 2010].



Figure E-1: Picture of experimental setup used to quantify feedforward force production.

### E.3 Results

The magnitude of the multi-input multi-output relation between commanded force  $u$  and measured force  $f$  was quantified (see Figures E-2 and E-3). From 0.5 - 15 Hz the frequency response of the system was almost flat. Moreover the magnitude was unity for the diagonal terms and zero for the off-diagonal terms. In the frequency range from 0.5-10 Hz the partial coherence was unity on the diagonal terms and close to zero on the off-diagonal terms.






Figure E-2: Magnitude of the relation between commanded force  $u$  and measured force  $f$ . Blue lines represent the small amplitude perturbations used in the experiments. Red lines represented the large amplitude perturbations.




Figure E-3: Coherence of the multi-input multi-output relation between commanded force  $u$  and measured force  $f$ . Blue lines represent the small amplitude perturbations used in the experiments. Red lines represented the large amplitude perturbations.

## E.4 Discussion and Conclusions

Based on the input-output spectral estimates the feedforward force dynamics of the InMotion robot are more than sufficient without compensation. In practice, a perturbation of 5 N was reported to ‘feel’ disruptive to subjects. The partial coherence in the small  $U$  condition was not negligible in the  $\gamma_{xy}$  case. However, for this application, a partial coherence of less than 0.1 is more than acceptable. For this reason, in the experiments, a normal perturbation of 1.5 N and a tangential perturbation of 0.5 N were employed – consistent with the small  $U$  condition.

# Appendix F

## InMotion Robot–Hysteresis

### F.1 Introduction

It was unclear why the inertial estimates of the robot alone were biased when performing the virtual crank-turning experiment. A simple experiment was performed to check the force production of the robot.

### F.2 Methods

The experiment was performed in two subsets. In the first case, denoted ‘robot’, the robot stiffness was set to 80 N/m. The  $\mathbf{x}_0$  position was set to zero in both x and y directions. The damping was also set to zero. The handle of the robot was moved back and forth 3 times in the  $x$  direction. The motion was performed by hand at a period which was slow enough to be considered quasi-static ( $\sim 20$  seconds/cycle).

In the second case, denoted ‘spring’, 3 springs were attached to the InMotion (each with a stiffness of  $\sim 80$  N/m). The spring case served as a control. The three springs were connected in parallel with the first two opposing the third (see Figure F-1). The commanded robot force was set to zero. Again, the handle of the robot was moved back and forth 3 times in the  $x$  direction as done in the previous case. However, in this case the external forces were applied by hand below the force transducer.

The horizontal position,  $x$ , was measured by the robot encoders. The force in

the horizontal direction,  $F_x$ , was measured by the force transducer. In both cases, the dominant behavior should be spring-like, resulting in a constant-slope line when plotted as  $F_x$  vs.  $x$ . Deviation from this behavior would suggest an artifact in either the InMotion position measurement, force production, or force transducer behavior.




Figure F-1: Picture of the three springs connected in parallel with the first to opposing the third.

### F.3 Results

The results are presented in Figure F-2. In the spring case small variability was observed in the force at a given position. This variability was  $\sim 0.1$  N, close to the resolution of the force transducer. This deviation from ideal behavior was considered to be negligible. However, the InMotion displayed evidence of a hysteresis in force production of  $\sim 0.5$  N.




Figure F-2: Plot of force transducer measurement vs. displacement measurement in the  $x$  direction. Blue represents the spring case and red represents the robot case. (Top) Large scale view of total force vs. displacement. Zoomed in view of Figure (bottom left) the spring and (bottom right) the robot. While the axes are translated between the bottom left and bottom right the scales are the same.

## F.4 Discussion and Conclusions

In this experiment the springs which served as a control behaved as expected. This measurement required both the position encoders and the force transducer – indicating that they both worked as expected. The experiments with the InMotion robot clearly demonstrated a substantial hysteresis of  $\sim 0.5$  N for forces  $-8 \text{ N} < x < 8 \text{ N}$ <sup>1</sup>. This range more than encompasses the force range experienced during the crank-turning experiments where subjects exerted forces  $-4 \text{ N} < x < 4 \text{ N}$  [Hermus et al., 2022].

The virtual crank-turning experiment with force perturbations requires force rendering at low amplitude forces. The force perturbations used in the experiments had a magnitude ranging from 0.5 N to 1.5 N. With incorrect force perturbation and/or impedance rendering of the crank dynamics the hysteresis could result in direction-dependent errors on the order of 10-50%. This problem must be addressed before virtual crank turning experiments with force perturbations can be successful.

---

<sup>1</sup>In other experiments, this hysteresis was consistently observed even at loads up to and beyond loads of  $\pm 30$  N.

# Appendix G

## Quantitative Analysis of the Quasi-Static and Dynamic Speeds

Many of the differences between statically consistent and dynamically consistent nullspace projectors will only be elicited if the end-effector task (Task 1) involves significant accelerations with respect to the joint-space task of the robot (Task 2). To gauge the degree to which Task 1 was dynamic with respect to Task 2, the principal natural frequencies of the joint-space task of the robot were quantified.

For any given robot configuration, the local unforced mass-spring behavior of the joint-space task can be approximated as follows:

$$M(\mathbf{q})\ddot{\boldsymbol{\theta}} + \mathbf{K}_q\boldsymbol{\theta} = 0 \quad (\text{G.1})$$

where  $\boldsymbol{\theta} = \mathbf{q} - \Delta\mathbf{q}$ . In order to find the natural frequencies of this system, we can assume solutions of the form:

$$\boldsymbol{\theta} = \mathbf{a}_i \sin(\omega_i t + \phi) \quad (\text{G.2})$$

where  $\mathbf{a}_i$  represents a single mode shape, and  $\omega_i$  represents the corresponding natural

frequency. Substituting this into Eqn. G.1 yields:

$$(-\omega_i^2 \mathbf{M}(\mathbf{q}) + \mathbf{K}_q) \mathbf{a}_i = 0 \quad (\text{G.3})$$

Rearranging this yields:

$$\mathbf{M}(\mathbf{q})^{-1} \mathbf{K}_q \mathbf{a}_i = -\omega_i^2 \mathbf{a}_i \quad (\text{G.4})$$

This has the form of a generalized eigenvalue problem, with  $\mathbf{a}_i$  and  $-\omega_i^2$  being the eigenvectors and eigenvalues of the matrix  $\mathbf{M}(\mathbf{q})^{-1} \mathbf{K}_q$ . For each experiment, the 7 mode shapes and natural frequencies were computed at each time step. For all experiments, the computed natural frequencies remained approximately similar across revolutions. The computed natural frequencies were then averaged across time steps and experiments. The natural frequencies of each mode, listed in ascending order, were  $\omega_n = [0.28, 0.38, 0.81, 1.09, 2.52, 3.65, 7.13]$  (Hz). The end-effector task frequencies were 0.25 Hz for the fast case, and 0.0769 Hz for the slow case, which suggests that in the fast condition, the response was dominated by the two lowest Task 2 natural frequencies. The slow condition was clearly quasi-static with respect to Task 2.



# Appendix H

## Simulations of Redundancy

In this work the benefit of increasing nullspace dimension was clear. However, this might have been a fortuitous accident of our experiment, performed with a particular robot in a particular configuration. In order to determine if the effect of increasing nullspace dimension was generalizable, planar simulations were performed which superimposed an end-effector impedance (task 1) and joint-space impedance (task 2).

In the 0D nullspace condition, the task space consisted of the  $x$ ,  $y$ , and  $\theta$  directions ( $m = 3$ ). In the 1D nullspace condition, the task space consisted of the  $x$  and  $y$  directions ( $m = 2$ ). Both the 0D and 1D conditions are graphically displayed in Figure H-1 (top). In all cases the manipulator had three joints ( $n = 3$ ). The total length of the manipulator was 1 m and the joints were divided into three equal segments. The total mass of all of the links was 1 kg. The links were assumed to be thin rods. The stiffness parameters are reported in Table H.1. Here,  $\mathbf{J}(\mathbf{q})_x \in \mathbb{R}^{2 \times n}$  maps the joint velocities  $\dot{\mathbf{q}} \in \mathbb{R}^n$  to translational end-effector velocities,  $\dot{\mathbf{x}} = [\dot{x}, \dot{y}]^T$  while  $\mathbf{J}(\mathbf{q})_\theta \in \mathbb{R}^{1 \times n}$  maps  $\dot{\mathbf{q}}$  to rotational end-effector velocities,  $\dot{\theta}$ . The torque control laws for the 0D and 1D cases were:

$$\boldsymbol{\tau}_{0D} = \underbrace{\boldsymbol{\tau}_e}_{\text{Task 1}} + \underbrace{\boldsymbol{\tau}_q}_{\text{Task 2}} \quad (\text{H.1a})$$

and

$$\boldsymbol{\tau}_{1D} = \underbrace{\boldsymbol{\tau}_x}_{\text{Task 1}} + \underbrace{\boldsymbol{\tau}_q}_{\text{Task 2}}. \quad (\text{H.1b})$$

The nominal joint configuration  $\mathbf{q}_0$  was constant and is depicted in Figure H-1 (top). The task 1 horizontal position was constant ( $x_0 = 0.75\text{m}$ ), the task 1 orientation (only used in the 0D nullspace condition) was constant ( $\theta = 0^\circ$ ), and the task 1 vertical position was time varying ( $y_0 = \frac{1}{2} \cos(2\pi t)$ ). The simulation was run for one cycle from  $y = 0.5$ , to  $y = -0.5$ , and back to  $y = 0.5$  with a period of 35 seconds which ensured the system was moving quasi-statically.

From the trajectory of the end-effector seen in Figure H-1 (bottom left) it is clear that the 1D nullspace condition resulted in a substantial decrease of task space disruption. This is supported by the  $x$  direction RMSE in Figure H-1 (bottom right).

Table H.1: Appendix H Controller Parameters. All non-diagonal stiffness and damping terms were zero.

Variable	Values	Units
$\mathbf{K}_x$	diag([500,500])	N/m
$\mathbf{B}_x$	$0.1\mathbf{K}_x$	N-s/m
$\mathbf{K}_\theta$	10	N/rad
$\mathbf{B}_\theta$	$0.1\mathbf{K}_\theta$	N-s/rad
$\mathbf{K}_q$	diag([5, 5, 5])	N-m/rad
$\mathbf{B}_q$	$0.1\mathbf{K}_q$	N-m-s/rad

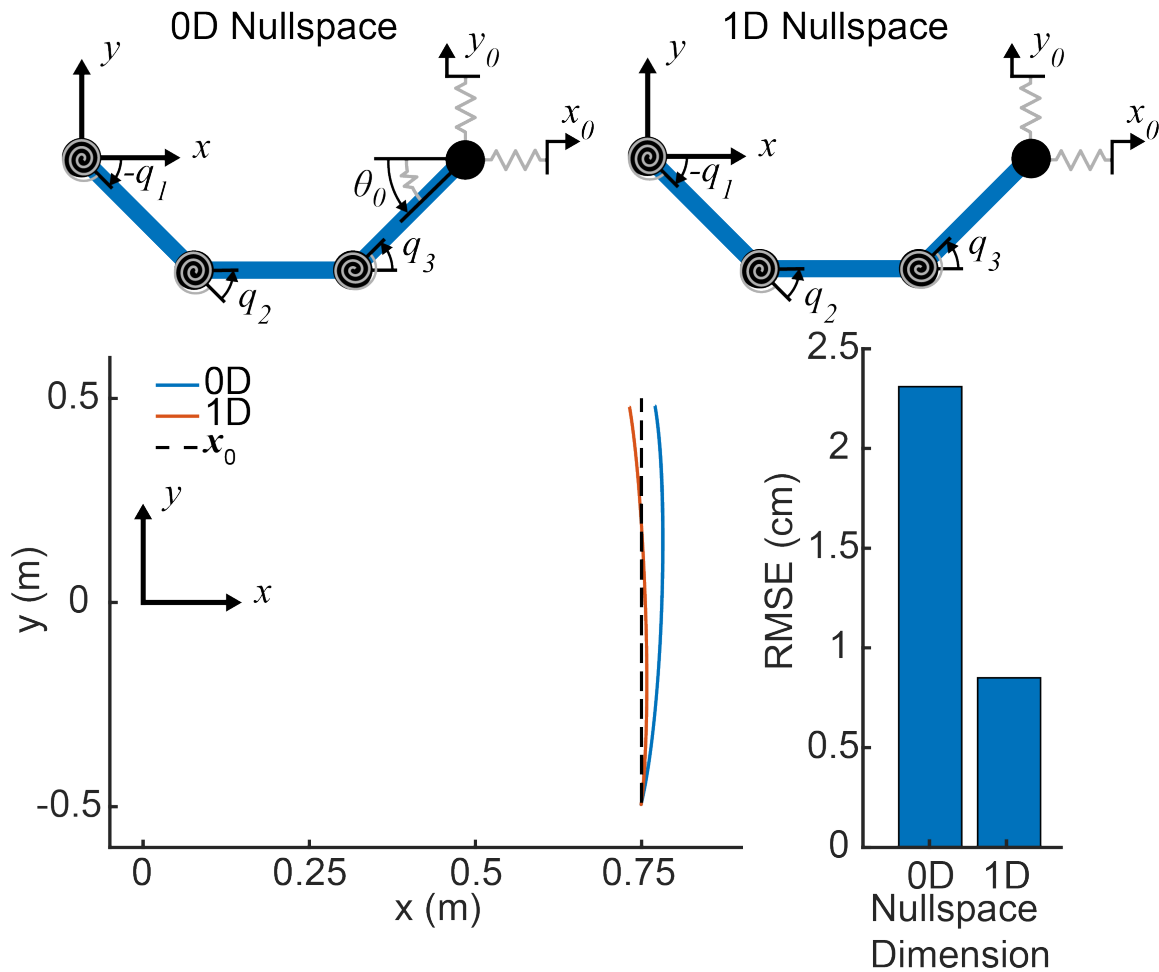


Figure H-1: Graphic depiction of a 3 degree-of-freedom manipulator with a task space impedance controller super imposed with a joint-space impedance controller. A 3D task space yields a 0D nullspace (top left). A 2D task space yields a 1D nullspace (top right). The end-effector path when a cosine motion was tracked in the y direction is shown on the bottom left. The position RMSE in the x direction was much larger with a 0D nullspace than with a 1D nullspace (bottom right).



# Appendix I

## Computation of Rotational Displacements

Computation of rotational displacements requires care. The computation of the rotational displacement  $\mathbf{u}\Delta\theta \in \mathbb{R}^{3 \times 1}$  are rigorously defined. Given three reference frames: an inertial ‘base’ frame ( $\Sigma_b$ ), a rotating reference frame aligned with the end-effector ( $\Sigma_e$ ) and a rotating reference frame aligned with the zero-force trajectory ( $\Sigma_0$ ) where the rotation matrix describing the end-effector frame in base coordinates is defined as  ${}^b\mathbf{R}_e$  and the zero-force trajectory in base coordinates is defined as  ${}^b\mathbf{R}_0$  the rotational transformation between the two frames can be computed by

$${}^e\mathbf{R}_0 = ({}^b\mathbf{R}_e)^T ({}^b\mathbf{R}_0) \quad (\text{I.1})$$

This rotation matrix was then converted to a quaternion representation. Given a rotation matrix indexed as,  $\mathbf{R} = \{R_{ij}, i = 1, 2, 3, j = 1, 2, 3\}$

$$\eta = \frac{1}{2} \sqrt{R_{11} + R_{22} + R_{33} + 1} \quad (\text{I.2})$$

$$\epsilon = \frac{1}{2} \begin{bmatrix} \text{sgn}(R_{32} - R_{23}) \sqrt{R_{11} - R_{22} - R_{33} + 1} \\ \text{sgn}(R_{13} - R_{31}) \sqrt{R_{22} - R_{33} + R_{11} + 1} \\ \text{sgn}(R_{12} - R_{12}) \sqrt{R_{33} + R_{11} + R_{22} + 1} \end{bmatrix} \quad (\text{I.3})$$

$$\boldsymbol{\epsilon} = \frac{1}{2} \begin{bmatrix} \frac{R_{32}-R_{23}}{4\eta} \\ \frac{R_{13}-R_{31}}{4\eta} \\ \frac{R_{21}-R_{12}}{4\eta} \end{bmatrix} \quad (\text{I.4})$$

The quaternion was normalized then the angle/axis representation was extracted [Natale, 2003]. The angle of rotation was computed

$$\Delta\theta = 2 \cos^{-1}(\eta); \quad (\text{I.5})$$

and the unit vector representing the axis of rotation was computed,

$$\mathbf{u}_e = \frac{1}{\sin(\frac{\Delta\theta}{2})} \boldsymbol{\epsilon} \quad (\text{I.6})$$

$$u_0 = {}^0\mathbf{R}_e \mathbf{u}_e \quad (\text{I.7})$$

# Bibliography

- [Abend et al., 1982] Abend, W., Bizzi, E., and Morasso, P. (1982). Human arm trajectory formation. *Brain : a journal of neurology*, 105(Pt 2):331–348.
- [Al-Khazali and Askari, 2012] Al-Khazali, H. A. and Askari, M. R. (2012). Geometrical and graphical representations analysis of lissajous figures in rotor dynamic system. *IOSR Journal of Engineering*, 2(05):971–978.
- [Albu-Schaffer and Hirzinger, 2002] Albu-Schaffer, A. and Hirzinger, G. (2002). Cartesian impedance control techniques for torque controlled light-weight robots. In *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292)*, volume 1, pages 657–663 vol.1.
- [Albu-Schaffer et al., 2003] Albu-Schaffer, A., Ott, C., Frese, U., and Hirzinger, G. (2003). Cartesian impedance control of redundant robots: recent results with the dlr-light-weight-arms. In *2003 IEEE International Conference on Robotics and Automation (Cat. No.03CH37422)*, volume 3, pages 3704–3709 vol.3.
- [Antonelli et al., 2005] Antonelli, G., Arrichiello, F., and Chiaverini, S. (2005). The null-space-based behavioral control for mobile robots. In *2005 International Symposium on Computational Intelligence in Robotics and Automation*, pages 15–20.
- [Antonelli et al., 2006] Antonelli, G., Arrichiello, F., and Chiaverini, S. (2006). Experiments of formation control with collisions avoidance using the null-space-based behavioral control. In *2006 14th Mediterranean Conference on Control and Automation*, pages 1–6.
- [Antonelli et al., 2008] Antonelli, G., Arrichiello, F., and Chiaverini, S. (2008). Stability analysis for the null-space-based behavioral control for multi-robot systems. In *2008 47th IEEE Conference on Decision and Control*, pages 2463–2468.
- [Antonelli et al., 2009] Antonelli, G., Arrichiello, F., and Chiaverini, S. (2009). Experiments of formation control with multirobot systems using the null-space-based behavioral control. *IEEE Transactions on Control Systems Technology*, 17(5):1173–1182.
- [Arrichiello et al., 2010] Arrichiello, F., Chiaverini, S., Indiveri, G., and Pedone, P. (2010). The null-space-based behavioral control for mobile robots with velocity actuator saturations. *The International Journal of Robotics Research*, 29(10):1317–1337.

- [Baillieul, 1985] Baillieul, J. (1985). Kinematic programming alternatives for redundant manipulators. In *Proceedings. 1985 IEEE International Conference on Robotics and Automation*, volume 2, pages 722–728.
- [Bazzi et al., 2018] Bazzi, S., Ebert, J., Hogan, N., and Sternad, D. (2018). Stability and predictability in dynamically complex physical interactions. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–5.
- [Bellman, 1966] Bellman, R. (1966). Dynamic programming. *Science*, 153(3731):34–37.
- [Bendat and Piersol, 2010] Bendat, J. S. and Piersol, A. G. (2010). *Random Data: Analysis and Measurement Procedures*. John Wiley and Sons, Inc., New Jersey.
- [Bennett et al., 1992] Bennett, D. J., Hollerbach, J. M., Xu, Y., and Hunter, I. W. (1992). Time-varying stiffness of human elbow joint during cyclic voluntary movement. *Experimental Brain Research*, 88(2):433–442.
- [Berman et al., 2014] Berman, S., Liebermann, D. G., and McIntyre, J. (2014). Constrained motion control on a hemispherical surface: path planning. *Journal of Neurophysiology*, 111(5):954–968.
- [Bizzi et al., 1982] Bizzi, E., Accornero, N., Chapple, W., and Hogan, N. (1982). Arm trajectory formation in monkeys. *Experimental Brain Research*, 46(1):139–143.
- [Bizzi et al., 1984] Bizzi, E., Accornero, N., Chapple, W., and Hogan, N. (1984). Posture control and trajectory formation during arm movement. *Journal of Neuroscience*, 4(11):2738–2744.
- [Boesch and Boesch, 1990] Boesch, C. and Boesch, H. (1990). Tool use and tool making in wild chimpanzees. *Folia Primatologica*, 54(1-2):86–99.
- [Brown, 1828] Brown, R. (1828). Xxvii. a brief account of microscopical observations made in the months of june, july and august 1827, on the particles contained in the pollen of plants; and on the general existence of active molecules in organic and inorganic bodies. *The Philosophical Magazine*, 4(21):161–173.
- [Brown, 1914] Brown, T. G. (1914). On the nature of the fundamental activity of the nervous centres; together with an analysis of the conditioning of rhythmic activity in progression, and a theory of the evolution of function in the nervous system. *The Journal of physiology*, 48(1):18–46.
- [Brown and Sherrington, 1911] Brown, T. G. and Sherrington, C. S. (1911). The intrinsic factors in the act of progression in the mammal. *Proceedings of the Royal Society of London. Series B, Containing Papers of a Biological Character*, 84(572):308–319.



- [Bruyninckx and Khatib, 2000] Bruyninckx, H. and Khatib, O. (2000). Gauss’ principle and the dynamics of redundant and constrained manipulators. In *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*, volume 3, pages 2563–2568 vol.3.
- [Burdet et al., 2001] Burdet, E., Osu, R., Franklin, D. W., Milner, T. E., and Kawato, M. (2001). The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature*, 414:446 EP –.
- [Campos and Calado, 2009] Campos, F. M. and Calado, J. M. (2009). Approaches to human arm movement control—A review. *Annual Reviews in Control*, 33(1):69–77.
- [Cannon and Zahalak, 1982] Cannon, S. C. and Zahalak, G. I. (1982). The mechanical behavior of active human skeletal muscle in small oscillations. *Journal of Biomechanics*, 15(2):111 – 121.
- [Casadio et al., 2015] Casadio, M., Pressman, A., and Mussa-Ivaldi, F. A. (2015). Learning to push and learning to move: the adaptive control of contact forces. *Frontiers in Computational Neuroscience*, 9:118.
- [Chang and Khatib, 1995] Chang, K.-S. and Khatib, O. (1995). Manipulator control at kinematic singularities: a dynamically consistent strategy. In *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots*, volume 3, pages 84–88 vol.3.
- [Chavan-Dafle and Rodriguez, 2015] Chavan-Dafle, N. and Rodriguez, A. (2015). Prehensile pushing: In-hand manipulation with push-primitives. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6215–6222.
- [Chib et al., 2009] Chib, V. S., Krutky, M. A., Lynch, K. M., and Mussa-Ivaldi, F. A. (2009). The separate neural control of hand movements and contact forces. *Journal of Neuroscience*, 29(12):3939–3947.
- [Churchland and Sejnowski, 2017] Churchland, P. and Sejnowski, T. (2017). *The Computational Brain*. MIT Press.
- [Colgate and Hogan, 1989] Colgate, E. and Hogan, N. (1989). An analysis of contact instability in terms of passive physical equivalents. In *Proceedings, 1989 International Conference on Robotics and Automation*, pages 404–409 vol.1.
- [Colgate and Hogan, 1988] Colgate, J. E. and Hogan, N. (1988). Robust control of dynamically interacting systems. *International Journal of Control*, 48(1):65–88.
- [Collewijn et al., 1988] Collewijn, H., Erkelens, C. J., and Steinman, R. M. (1988). Binocular co-ordination of human horizontal saccadic eye movements. *The Journal of Physiology*, 404:157–182.

- [Cover, 1965] Cover, T. M. (1965). Geometrical and statistical properties of systems of linear inequalities with applications in pattern recognition. *IEEE Transactions on Electronic Computers*, EC-14(3):326–334.
- [Crossman and Goodeve, 1983] Crossman, E. R. F. W. and Goodeve, P. J. (1983). Feedback control of hand-movement and fitts’ law. *The Quarterly Journal of Experimental Psychology Section A*, 35(2):251–278.
- [Damm and McIntyre, 2008] Damm, L. and McIntyre, J. (2008). Physiological basis of limb-impedance modulation during free and constrained movements. *Journal of Neurophysiology*, 100(5):2577–2588.
- [Davoudabadi Farahani et al., 2016] Davoudabadi Farahani, S., Svinin, M., Andersen, M. S., de Zee, M., and Rasmussen, J. (2016). Prediction of closed-chain human arm dynamics in a crank-rotation task. *Journal of Biomechanics*, 49(13):2684–2693.
- [Dayan et al., 2007] Dayan, E., Casile, A., Levit-Binnun, N., Giese, M. A., Hendler, T., and Flash, T. (2007). Neural representations of kinematic laws of motion: Evidence for action-perception coupling. *Proceedings of the National Academy of Sciences*, 104(51):20582–20587.
- [de Rugy and Sternad, 2003] de Rugy, A. and Sternad, D. (2003). Interaction between discrete and rhythmic movements: reaction time and phase of discrete movement initiation during oscillatory movements. *Brain Research*, 994(2):160 – 174.
- [Degallier and Ijspeert, 2010] Degallier, S. and Ijspeert, A. (2010). Modeling discrete and rhythmic movements through motor primitives: A review.
- [Dietrich, 2016] Dietrich, A. (2016). *Whole-body impedance control of wheeled humanoid robots*. Springer.
- [Dietrich et al., 2012] Dietrich, A., Albu-Schäffer, A., and Hirzinger, G. (2012). On continuous null space projections for torque-based, hierarchical, multi-objective manipulation. In *2012 IEEE International Conference on Robotics and Automation*, pages 2978–2985.
- [Dietrich et al., 2013] Dietrich, A., Ott, C., and Albu-Schäffer, A. (2013). Multi-objective compliance control of redundant manipulators: Hierarchy, control, and stability. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 3043–3050.
- [Dietrich et al., 2015] Dietrich, A., Ott, C., and Albu-Schäffer, A. (2015). An overview of null space projections for redundant, torque-controlled robots. *The International Journal of Robotics Research*, 34(11):1385–1400.
- [Dietrich et al., 2016] Dietrich, A., Ott, C., and Stramigioli, S. (2016). Passivation of projection-based null space compliance control via energy tanks. *IEEE Robotics and Automation Letters*, 1(1):184–191.

- [Dietrich et al., 2017] Dietrich, A., Wu, X., Bussmann, K., Ott, C., Albu-Schäffer, A., and Stramigioli, S. (2017). Passive hierarchical impedance control via energy tanks. *IEEE Robotics and Automation Letters*, 2(2):522–529.
- [Doeringer, 1999] Doeringer, J. A. (1999). *An Investigation into the Discrete Nature of Human Arm Movements*. PhD thesis, Massachusetts Institute of Technology.
- [Doeringer and Hogan, 1998a] Doeringer, J. A. and Hogan, N. (1998a). Intermittency in preplanned elbow movements persists in the absence of visual feedback. *Journal of Neurophysiology*, 80(4):1787–1799.
- [Doeringer and Hogan, 1998b] Doeringer, J. A. and Hogan, N. (1998b). Serial processing in human movement production. *Neural Networks*, 11(7):1345 – 1356.
- [Dohrmann et al., 1988] Dohrmann, C. R., Busby, H. R., and Trujillo, D. M. (1988). Smoothing noisy data using dynamic programming and generalized cross-validation. *Journal of Biomechanical Engineering*, 110(1):37–41.
- [Doshi et al., 2020] Doshi, N., Hogan, F. R., and Rodriguez, A. (2020). Hybrid differential dynamic programming for planar manipulation primitives. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6759–6765.
- [Edward Colgate, 1992] Edward Colgate, J. (1992). Strictly positive real admittances for coupled stability. *Journal of the Franklin Institute*, 329(3):429–444.
- [Einstein, 1905] Einstein, A. (1905). Über die von der molekularkinetischen theorie der wärme geforderte bewegung von in ruhenden flüssigkeiten suspendierten teilchen. *Annalen der physik*, 4.
- [Farouki, 2012] Farouki, R. T. (2012). The bernstein polynomial basis: A centennial retrospective. *Computer Aided Geometric Design*, 29(6):379 – 419.
- [Fasse and Hogan, 1996] Fasse, E. D. and Hogan, N. (1996). Control of physical contact and dynamic interaction. In Giralt, G. and Hirzinger, G., editors, *Robotics Research*, pages 28–38, London. Springer London.
- [Fasse et al., 2000] Fasse, E. D., Hogan, N., Kay, B. A., and Mussa-Ivaldi, F. A. (2000). Haptic interaction with virtual objects. *Biological Cybernetics*, 82(1):69–83.
- [Feldman, 1966] Feldman, A. (1966). Functional tuning of the nervous system during control of movement or maintenance of a steady posture. ii. controllable paramaters of the muscles. *Biophysics*, 11:565–578.
- [Feldman, 1986] Feldman, A. G. (1986). Once more on the equilibrium-point hypothesis ( $\gamma$  model) for motor control. *Journal of Motor Behavior*, 18(1):17–54. PMID: 15136283.

- [Ficanha et al., 2015] Ficanha, E. M., Ribeiro, G., and Rastgaar Aagaah, M. (2015). Instrumented walkway for estimation of the ankle impedance in dorsiflexion-plantarflexion and inversion-eversion during standing and walking. (57267).
- [Ficuciello et al., 2015] Ficuciello, F., Villani, L., and Siciliano, B. (2015). Variable impedance control of redundant manipulators for intuitive human-robot physical interaction. *IEEE Transactions on Robotics*, 31(4):850–863.
- [Fitts, 1954] Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, 47(6):381–391.
- [Flacco et al., 2012] Flacco, F., De Luca, A., and Khatib, O. (2012). Motion control of redundant robots under joint constraints: Saturation in the null space. In *2012 IEEE International Conference on Robotics and Automation*, pages 285–292.
- [Flacco et al., 2015] Flacco, F., De Luca, A., and Khatib, O. (2015). Control of redundant robots under hard joint constraints: Saturation in the null space. *IEEE Transactions on Robotics*, 31(3):637–654.
- [Flanagan and Rao, 1995] Flanagan, J. R. and Rao, A. K. (1995). Trajectory adaptation to a nonlinear visuomotor transformation: evidence of motion planning in visually perceived space. *Journal of Neurophysiology*, 74(5):2174–2178.
- [Flash, 1987] Flash, T. (1987). The control of hand equilibrium trajectories in multi-joint arm movements. *Biological Cybernetics*, 57(4):257–274.
- [Flash and Hogan, 1985] Flash, T. and Hogan, N. (1985). The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of Neuroscience*, 5(7):1688–1703.
- [Fraisse, 1984] Fraisse, P. (1984). Perception and estimation of time. *Annual review of psychology*, 35(1):1–37.
- [Franklin et al., 2007] Franklin, D. W., Liaw, G., Milner, T. E., Osu, R., Burdet, E., and Kawato, M. (2007). Endpoint stiffness of the arm is directionally tuned to instability in the environment. *Journal of Neuroscience*, 27(29):7705–7716.
- [Franklin et al., 2003] Franklin, D. W., Osu, R., Burdet, E., Kawato, M., and Milner, T. E. (2003). Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model. *Journal of Neurophysiology*, 90(5):3270–3282. PMID: 14615432.
- [Fritsch and Carlson, 1980] Fritsch, F. N. and Carlson, R. E. (1980). Monotone piecewise cubic interpolation. *SIAM Journal on Numerical Analysis*, 17(2):238–246.
- [Garrett et al., 2021] Garrett, C. R., Chitnis, R., Holladay, R., Kim, B., Silver, T., Kaelbling, L. P., and Lozano-Pérez, T. (2021). Integrated task and motion planning. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(1):265–293.

- [Georgopoulos et al., 1986] Georgopoulos, A., Schwartz, A., and Kettner, R. (1986). Neuronal population coding of movement direction. *Science*, 233(4771):1416–1419.
- [Georgopoulos et al., 1992] Georgopoulos, A. P., Ashe, J., Smyrnis, N., and Taira, M. (1992). The Motor Cortex and the Coding of Force. *Science*, 256(5064):1692–1695.
- [Giszter, 2015] Giszter, S. F. (2015). Motor primitives — new data and future questions. *Current Opinion in Neurobiology*, 33:156 – 165. Motor circuits and action.
- [Gomi and Kawato, 1996] Gomi, H. and Kawato, M. (1996). Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science*, 272(5258):117–120.
- [Grau et al., 2021] Grau, A., Indri, M., Lo Bello, L., and Sauter, T. (2021). Robots in industry: The past, present, and future of a growing collaboration with humans. *IEEE Industrial Electronics Magazine*, 15(1):50–61.
- [Gribble and Ostry, 1996] Gribble, P. L. and Ostry, D. J. (1996). Origins of the power law relation between movement velocity and curvature: modeling the effects of muscle mechanics and limb dynamics. *Journal of Neurophysiology*, 76(5):2853–2860. PMID: 8930238.
- [Gribble et al., 1998] Gribble, P. L., Ostry, D. J., Sanguinetti, V., and Laboisière, R. (1998). Are complex control signals required for human arm movement? *Journal of Neurophysiology*, 79(3):1409–1424. PMID: 9497421.
- [Grillner and Wallen, 1985] Grillner, S. and Wallen, P. (1985). Central pattern generators for locomotion, with special reference to vertebrates. *Annual review of neuroscience*.
- [Guarín and Kearney, 2017] Guarín, D. L. and Kearney, R. E. (2017). Estimation of time-varying, intrinsic and reflex dynamic joint stiffness during movement. application to the ankle joint. *Frontiers in Computational Neuroscience*, 11:51.
- [Guarin and Kearney, 2018] Guarin, D. L. and Kearney, R. E. (2018). Unbiased estimation of human joint intrinsic mechanical properties during movement. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 26(10):1975–1984.
- [Gulletta et al., 2020] Gulletta, G., Erlhagen, W., and Bicho, E. (2020). Human-like arm motion generation: A review.
- [Hamel-Pâquet et al., 2006] Hamel-Pâquet, C., Sergio, L. E., and Kalaska, J. F. (2006). Parietal Area 5 Activity Does Not Reflect the Differential Time-Course of Motor Output Kinetics During Arm-Reaching and Isometric-Force Tasks. *Journal of Neurophysiology*, 95(6):3353–3370.
- [Henze et al., 2017] Henze, B., Dietrich, A., Roa, M. A., and Ott, C. (2017). Multi-contact balancing of humanoid robots in confined spaces: Utilizing knee contacts. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 697–704.

- [Hermus et al., 2020] Hermus, J., Doeringer, J., Sternad, D., and Hogan, N. (2020). Separating neural influences from peripheral mechanics: the speed-curvature relation in mechanically constrained actions. *Journal of Neurophysiology*, 123(5):1870–1885. PMID: 32159419.
- [Hermus et al., 2022] Hermus, J., Lachner, J., Verdi, D., and Hogan, N. (2022). Exploiting redundancy to facilitate physical interaction. *IEEE Transactions on Robotics*, 38(1):599–615.
- [Hermus et al., 2020] Hermus, J., Sternad, D., and Hogan, N. (2020). Evidence for dynamic primitives in a constrained motion task. In *2020 8th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)*, pages 551–556.
- [Hicheur et al., 2005] Hicheur, H., Vieilledent, S., Richardson, M. J. E., Flash, T., and Berthoz, A. (2005). Velocity and curvature in human locomotion along complex curved paths: a comparison with hand movements. *Experimental Brain Research*, 162(2):145–154.
- [Hill, 1938] Hill, A. V. (1938). The heat of shortening and the dynamic constants of muscle. *Proceedings of the Royal Society of London. Series B - Biological Sciences*, 126(843):136–195.
- [Hjorth et al., 2020] Hjorth, S., Lachner, J., Stramigioli, S., Madsen, O., and Chrysostomou, D. (2020). An energy-based approach for the integration of collaborative redundant robots in restricted work environments. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7152–7158.
- [Hodgson and Hogan, 2000] Hodgson, A. J. and Hogan, N. (2000). A model-independent definition of attractor behavior applicable to interactive tasks. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 30(1):105–118.
- [Hoffer and Andreassen, 1981] Hoffer, J. A. and Andreassen, S. (1981). Regulation of soleus muscle stiffness in preamillary cats: intrinsic and reflex components. *Journal of Neurophysiology*, 45(2):267–285. PMID: 6780665.
- [Hogan et al., 2020] Hogan, F. R., Ballester, J., Dong, S., and Rodriguez, A. (2020). Tactile dexterity: Manipulation primitives with tactile feedback. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8863–8869.
- [Hogan, 1982] Hogan, N. (1982). Control and coordination of voluntary arm movements. In *1982 American Control Conference*, pages 522–528.
- [Hogan, 1984a] Hogan, N. (1984a). Impedance control: An approach to manipulation. *American Control Conference, 1984*, pages 304–313.

- [Hogan, 1984b] Hogan, N. (1984b). An organizing principle for a class of voluntary movements. *Journal of Neuroscience*, 4(11):2745–2754.
- [Hogan, 1985a] Hogan, N. (1985a). Impedance control: An approach to manipulation: Part i—theory. *Journal of Dynamic Systems, Measurement, and Control*, 107(1):1–7.
- [Hogan, 1985b] Hogan, N. (1985b). Impedance control: An approach to manipulation: Part ii—implementation. *Journal of Dynamic Systems, Measurement, and Control*, 107(1):8–16.
- [Hogan, 1985c] Hogan, N. (1985c). Impedance control: An approach to manipulation: Part iii—applications. *Journal of Dynamic Systems, Measurement, and Control*, 107(1):17–24.
- [Hogan, 1985d] Hogan, N. (1985d). The mechanics of multi-joint posture and movement control. *Biological Cybernetics*, 52(5):315–331.
- [Hogan, 1988] Hogan, N. (1988). On the stability of manipulators performing contact tasks. *IEEE Journal on Robotics and Automation*, 4(6):677–686.
- [Hogan, 2014] Hogan, N. (2014). A general actuator model based on nonlinear equivalent networks. *IEEE/ASME Transactions on Mechatronics*, 19(6):1929–1939.
- [Hogan, 2017] Hogan, N. (2017). *Physical Interaction via Dynamic Primitives*, pages 269–299. Springer International Publishing.
- [Hogan et al., 1999] Hogan, N., Doeringer, J. A., and Krebs, H. I. (1999). Arm movement control is both continuous and discrete. *Cognitive Studies*, 6(3):254–273.
- [Hogan et al., 1995] Hogan, N., Krebs, H. I., Sharon, A., and Charnnarong, J. (1995). Interactive robotic therapist. United States patent US 5466213.
- [Hogan and Sternad, 2007] Hogan, N. and Sternad, D. (2007). On rhythmic and discrete movements: reflections, definitions and implications for motor control. *Experimental Brain Research*, 181(1):13–30.
- [Hogan and Sternad, 2012] Hogan, N. and Sternad, D. (2012). Dynamic primitives of motor behavior. *Biol. Cybern.*, 106(11-12):727–739.
- [Hogan and Sternad, 2013] Hogan, N. and Sternad, D. (2013). Dynamic primitives in the control of locomotion. *Frontiers in Computational Neuroscience*, 7:71.
- [Holden, 2013] Holden, A. V. (2013). *Models of the stochastic activity of neurones*, volume 12. Springer Science & Business Media.
- [Hollerbach and Ki Suh, 1987] Hollerbach, J. and Ki Suh (1987). Redundancy resolution of manipulators through torque optimization. *IEEE Journal on Robotics and Automation*, 3(4):308–316.

- [Hosford, 2016] Hosford, L. A. (2016). Development and testing of an impedance controller on an anthropomorphic robot for extreme environment operations. Master’s thesis, Massachusetts Institute of Technology.
- [Hu et al., 1995] Hu, B., Teo, C. L., and Lee, H. P. (1995). Local optimization of weighted joint torques for redundant robotic manipulators. *IEEE Transactions on Robotics and Automation*, 11(3):422–425.
- [Huh and Sejnowski, 2015] Huh, D. and Sejnowski, T. J. (2015). Spectrum of power laws for curved hand movements. *Proceedings of the National Academy of Sciences*, 112(29):E3950–E3958.
- [Hunt, 1996] Hunt, G. R. (1996). Manufacture and use of hook-tools by new caledonian crows. *Nature*, 379(6562):249–251.
- [Ijspeert et al., 2002] Ijspeert, A. J., Nakanishi, J., and Schaal, S. (2002). Movement imitation with nonlinear dynamical systems in humanoid robots. *Proceedings - IEEE International Conference on Robotics and Automation*, 2:1398–1403.
- [James, 1890] James, W. (1890). *The principles of psychology*, volume 2. Henry Holt and Company, NY.
- [Johnson-Frey, 2004] Johnson-Frey, S. H. (2004). The neural bases of complex tool use in humans. *Trends in Cognitive Sciences*, 8(2):71–78.
- [Joyce et al., 1969] Joyce, G. C., Rack, P. M. H., and Westbury, D. R. (1969). The mechanical properties of cat soleus muscle during controlled lengthening and shortening movements. *The Journal of Physiology*, 204(2):461–474.
- [Kandel et al., 2013] Kandel, E., Schwartz, J., Jessell, T., Siegelbaum, S., and Hudspeth, A. (2013). *Principles of Neural Science, Fifth Edition*. Principles of Neural Science. McGraw-Hill Education, fifth edition.
- [Kargo and Giszter, 2008] Kargo, W. J. and Giszter, S. F. (2008). Individual premotor drive pulses, not time-varying synergies, are the units of adjustment for limb trajectories constructed in spinal cord. *Journal of Neuroscience*, 28(10):2409–2425.
- [Kazanzides et al., 1989] Kazanzides, P., Bradley, N. S., and Wolovich, W. A. (1989). Dual-drive force/velocity control: implementation and experimental results. In *Proceedings, 1989 International Conference on Robotics and Automation*, pages 92–97 vol.1.
- [Kearney and Hunter, 1990] Kearney, R. and Hunter, I. (1990). System identification of human joint dynamics. *Critical reviews in biomedical engineering*, 18(1):55–87.
- [Kenward et al., 2005] Kenward, B., Weir, A. A. S., Rutz, C., and Kacelnik, A. (2005). Tool manufacture by naive juvenile crows. *Nature*, 433(7022):121–121.
- [Kepler, 1609] Kepler, J. (1609). *Astronomia Nova*.



- [Keppel and Wickens, 2004] Keppel, G. and Wickens, T. D. (2004). *Design and Analysis A Researcher's Handbook*. Pearson Education, Inc.
- [Khatib, 1987] Khatib, O. (1987). A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal on Robotics and Automation*, 3(1):43–53.
- [Khatib, 1995] Khatib, O. (1995). Inertial properties in robotic manipulation: An object-level framework. *The International Journal of Robotics Research*, 14(1):19–36.
- [Klein and Huang, 1983] Klein, C. A. and Huang, C. (1983). Review of pseudoinverse control for use with kinematically redundant manipulators. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13(2):245–250.
- [Kluzik et al., 2008] Kluzik, J., Diedrichsen, J., Shadmehr, R., and Bastian, A. J. (2008). Reach adaptation: What determines whether we learn an internal model of the tool or adapt the model of our arm? *Journal of Neurophysiology*, 100(3):1455–1464.
- [Koeppen et al., 2017] Koeppen, R., Huber, M., Sternad, D., and Hogan, V. (2017). Controlling physical interactions: Humans do not minimize muscle effort. *ASME 2017 Dynamic Systems and Control Conference*.
- [Kolesnikov et al., 2011] Kolesnikov, M., Piovesan, D., Lynch, K. M., and Mussa-Ivaldi, F. A. (2011). On force regulation strategies in predictable environments. *Annu Int Conf IEEE Eng Med Biol Soc*, 2011:4076–4081.
- [Krakauer et al., 2017] Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., and Poeppel, D. (2017). Neuroscience needs behavior: Correcting a reductionist bias. *Neuron*, 93(3):480 – 490.
- [Krebs, 1997] Krebs, H. (1997). *Robot-aided neurorehabilitation and functional imaging*. PhD thesis, Massachusetts Institute of Technology.
- [Krebs et al., 1999] Krebs, H. I., Aisen, M. L., Volpe, B. T., and Hogan, N. (1999). Quantization of continuous arm movements in humans with brain injury. *Proceedings of the National Academy of Sciences*, 96(8):4645–4649.
- [Krutky et al., 2013] Krutky, M. A., Trumbower, R. D., and Perreault, E. J. (2013). Influence of environmental stability on the regulation of end-point impedance during the maintenance of arm posture. *Journal of Neurophysiology*, 109(4):1045–1054. PMID: 23221409.
- [Lachner et al., 2021] Lachner, J., Allmendinger, F., Hobert, E., Hogan, N., and Stramigioli, S. (2021). Energy budgets for coordinate invariant robot control in physical human–robot interaction. *The International Journal of Robotics Research*, page 02783649211011639.

- [Lachner et al., 2020] Lachner, J., Schettino, V., Allmendinger, F., Fiore, M. D., Ficuciello, F., Siciliano, B., and Stramigioli, S. (2020). The influence of coordinates in robotic manipulability analysis. *Mechanism and Machine Theory*, 146:103722.
- [Lackner and Dizio, 1994] Lackner, J. R. and Dizio, P. (1994). Rapid adaptation to coriolis force perturbations of arm trajectory. *Journal of Neurophysiology*, 72(1):299–313.
- [Lacquaniti et al., 1993] Lacquaniti, F., Corrozzo, M., and Borghese, N. A. (1993). Time-varying mechanical behavior of multijointed arm in man. *Journal of Neurophysiology*, 69(5):1443–1464. PMID: 8509826.
- [Lacquaniti et al., 1983] Lacquaniti, F., Terzuolo, C., and Viviani, P. (1983). The law relating the kinematic and figural aspects of drawing movements. *Acta Psychologica*, 54(1):115 – 130.
- [Lee et al., 2014a] Lee, H., Ho, P., Rastgaar, M., Krebs, H. I., and Hogan, N. (2014a). Multivariable static ankle mechanical impedance with active muscles. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 22(1):44–52.
- [Lee and Hogan, 2015] Lee, H. and Hogan, N. (2015). Time-varying ankle mechanical impedance during human locomotion. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 23(5):755–764.
- [Lee et al., 2014b] Lee, H., Krebs, H. I., and Hogan, N. (2014b). Multivariable dynamic ankle mechanical impedance with active muscles. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 22(5):971–981.
- [Lee et al., 2016] Lee, H., Rouse, E. J., and Krebs, H. I. (2016). Summary of human ankle mechanical impedance during walking. *IEEE Journal of Translational Engineering in Health and Medicine*, 4:1–7.
- [Lee et al., 2021] Lee, J., Zhang, K., and Hogan, N. (2021). Identifying human postural dynamics and control from unperturbed balance. *Journal of NeuroEngineering and Rehabilitation*, 18(1):54.
- [Liegeois, 1977] Liegeois, A. (1977). Automatic supervisory control of the configuration and behavior of multibody mechanisms. *IEEE Transactions on Systems, Man, and Cybernetics*, 7(12):868–871.
- [Lin et al., 2015] Lin, H., Howard, M., and Vijayakumar, S. (2015). Learning null space projections. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2613–2619.
- [Lipps et al., 2020] Lipps, D. B., Baillargeon, E. M., Ludvig, D., and Perreault, E. J. (2020). Quantifying the Multidimensional Impedance of the Shoulder During Volitional Contractions. *Annals of Biomedical Engineering*, 48(9):2354–2369.

- [Loram et al., 2005] Loram, I. D., Maganaris, C. N., and Lakie, M. (2005). Human postural sway results from frequent, ballistic bias impulses by soleus and gastrocnemius. *The Journal of Physiology*, 564(1):295–311.
- [Loram et al., 2014] Loram, I. D., van de Kamp, C., Lakie, M., Gollee, H., and Gawthrop, P. J. (2014). Does the motor system need intermittent control? *Exercise and Sport Sciences Reviews*, 42(3).
- [Lortie and Kearney, 2001] Lortie, M. and Kearney, R. E. (2001). Identification of physiological systems: Estimation of linear timevarying dynamics with non-white inputs and noisy outputs. *Medical and Biological Engineering and Computing*, 39(3):381–390.
- [Ludvig et al., 2011] Ludvig, D., Visser, T. S., Giesbrecht, H., and Kearney, R. E. (2011). Identification of time-varying intrinsic and reflex joint stiffness. *IEEE Transactions on Biomedical Engineering*, 58(6):1715–1723.
- [Luh et al., 1980] Luh, J., Walker, M., and Paul, R. (1980). Resolved-acceleration control of mechanical manipulators. *IEEE Transactions on Automatic Control*, 25(3):468–474.
- [MacNeil et al., 1992] MacNeil, J. B., Kearney, R. E., and Hunter, I. W. (1992). Identification of time-varying biological systems from ensemble data (joint dynamics application). *IEEE Transactions on Biomedical Engineering*, 39(12):1213–1225.
- [Maganaris, 2001] Maganaris, C. N. (2001). Force-length characteristics of in vivo human skeletal muscle. *Acta Physiol Scand*, 172(4):279–285.
- [Markkula et al., 2018] Markkula, G., Boer, E., Romano, R., and Merat, N. (2018). Sustained sensorimotor control as intermittent decisions about prediction errors: computational framework and application to ground vehicle steering. *Biological Cybernetics*, 112(3):181–207.
- [Martin et al., 2018] Martin, J. A., Brandon, S. C. E., Keuler, E. M., Hermus, J. R., Ehlers, A. C., Segalman, D. J., Allen, M. S., and Thelen, D. G. (2018). Gauging force by tapping tendons. *Nature Communications*, 9(1):1592.
- [Mason, 1981] Mason, M. T. (1981). Compliance and force control for computer controlled manipulators. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(6):418–432.
- [Massey et al., 1992] Massey, J. T., Lurito, J. T., Pellizzer, G., and Georgopoulos, A. P. (1992). Three-dimensional drawings in isometric conditions: relation between geometry and kinematics. *Experimental Brain Research*, 88(3):685–690.
- [Maurice et al., 2018a] Maurice, P., Hogan, N., and Sternad, D. (2018a). Predictability, force, and (anti)resonance in complex object control. *Journal of Neurophysiology*, 120(2):765–780. PMID: 29668379.

- [Maurice et al., 2018b] Maurice, P., Huber, M. E., Hogan, N., and Sternad, D. (2018b). Velocity-curvature patterns limit human-robot physical interaction. *IEEE Robotics and Automation Letters*, 3(1):249–256.
- [McIntyre et al., 1995] McIntyre, J., Gurfinkel, E. V., Lipshits, M. I., Droulez, J., and Gurfinkel, V. S. (1995). Measurements of human force control during a constrained arm motion using a force-actuated joystick. *Journal of Neurophysiology*, 73(3):1201–1222.
- [Miller and Nelson, 1973] Miller, D. and Nelson, R. (1973). *Biomechanics of Sport: a research approach*. Lee and Febiger, Philadelphia.
- [Milner and Franklin, 2005] Milner, T. E. and Franklin, D. W. (2005). Impedance control and internal model use during the initial stage of adaptation to novel dynamics in humans. *The Journal of Physiology*, 567(2):651–664.
- [Morasso, 1981] Morasso, P. (1981). Spatial control of arm movements. *Experimental Brain Research*, 42(2):223–227.
- [Mussa-Ivaldi et al., 1985] Mussa-Ivaldi, F., Hogan, N., and Bizzi, E. (1985). Neural, mechanical, and geometric factors subserving arm posture in humans. *Journal of Neuroscience*, 5(10):2732–2743.
- [Nah, 2020] Nah, M. C. (2020). Dynamic primitives facilitate manipulating a whip. Master’s thesis, Massachusetts Institute of Technology.
- [Nakanishi et al., 2008] Nakanishi, J., Cory, R., Mistry, M., Peters, J., and Schaal, S. (2008). Operational space control: A theoretical and empirical comparison. *The International Journal of Robotics Research*, 27(6):737–757.
- [Nalam and Lee, 2017] Nalam, V. and Lee, H. (2017). Design and validation of a multi-axis robotic platform for the characterization of ankle neuromechanics. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 511–516.
- [Nasserouleslami et al., 2014] Nasserouleslami, B., Hasson, C. J., and Sternad, D. (2014). Rhythmic manipulation of objects with complex dynamics: Predictability over chaos. *PLOS Computational Biology*, 10(10):1–19.
- [Natale, 2003] Natale, C. (2003). *Interaction control of robot manipulators: six degrees-of-freedom tasks*, volume 3. Springer Science & Business Media.
- [Nayeem et al., 2021] Nayeem, R., Bazzi, S., Sadeghi, M., Hogan, N., and Sternad, D. (2021). Preparing to move: Setting initial conditions to simplify interactions with complex objects. *PLOS Computational Biology*, 17(12):1–38.
- [Nelson, 1967] Nelson, E. (1967). *Dynamical theories of Brownian motion*. Princeton university press.

- [Nelson, 1983] Nelson, W. L. (1983). Physical principles for economies of skilled movements. *Biological Cybernetics*, 46(2):135–147.
- [Oda et al., 1995] Oda, N., Ohta, H., Murakami, T., and Ohnishi, K. (1995). A robust impedance control strategy for redundant manipulator. In *Proceedings of IECON '95 - 21st Annual Conference on IEEE Industrial Electronics*, volume 2, pages 1254–1259 vol.2.
- [Ohta et al., 1998] Ohta, K., Luo, Z.-W., and Ito, M. (1998). Analysis of human movement under environmental constraints: Adaptability to environment during crank rotation tasks. *Transactions of the Institute of Electronics, Information and Communication Engineers*, 81(6):1392–1401.
- [Ohta et al., 2004] Ohta, K., Svinin, M. M., Luo, Z., Hosoe, S., and Laboissière, R. (2004). Optimal trajectory formation of constrained human arm reaching movements. *Biological Cybernetics*, 91(1):23–36.
- [Oriolo, 1994] Oriolo, G. (1994). Stabilization of self-motions in redundant robots. In *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, pages 704–709 vol.1.
- [Osu et al., 2003] Osu, R., Burdet, E., Franklin, D. W., Milner, T. E., and Kawato, M. (2003). Different Mechanisms Involved in Adaptation to Stable and Unstable Dynamics. <https://doi.org/10.1152/jn.00073.2003>, 90(5):3255–3269.
- [Osu et al., 2004] Osu, R., Kamimura, N., Iwasaki, H., Nakano, E., Harris, C. M., Wada, Y., and Kawato, M. (2004). Optimal impedance control for task achievement in the presence of signal-dependent noise. *Journal of Neurophysiology*, 92(2):1199–1215.
- [Ott, 2008] Ott, C. (2008). *Cartesian impedance control of redundant and flexible-joint robots*. Springer.
- [Ott et al., 2010] Ott, C., Mukherjee, R., and Nakamura, Y. (2010). Unified impedance and admittance control. In *2010 IEEE International Conference on Robotics and Automation*, pages 554–561.
- [Palazzolo, 2005] Palazzolo, J. (2005). *Robotic Technology to Aid and Assess Recovery and Learning in Stroke Patients*. PhD thesis, Massachusetts Institute of Technology.
- [Palazzolo et al., 2007] Palazzolo, J. J., Ferraro, M., Krebs, H. I., Lynch, D., Volpe, B. T., and Hogan, N. (2007). Stochastic estimation of arm mechanical impedance during robotic stroke rehabilitation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 15(1):94–103.
- [Park and Khatib, 2008] Park, J. and Khatib, O. (2008). Robot multiple contact control. *Robotica*, 26(5):667–677.

- [Park et al., 2017] Park, S.-W., Marino, H., Charles, S. K., Sternad, D., and Hogan, N. (2017). Moving slowly is hard for humans: limitations of dynamic primitives. *Journal of Neurophysiology*, 118(1):69–83.
- [Paul, 1987] Paul, R. P. (1987). Problems and research issues associated with the hybrid control of force and displacement. In *IEEE International Conference on Robotics and Automation*, pages 1966–1971.
- [Penrose, 1955] Penrose, R. (1955). A generalized inverse for matrices. *Mathematical Proceedings of the Cambridge Philosophical Society*, 51(3):406–413.
- [Perreault et al., 1999] Perreault, E. J., Kirsch, R. F., and Acosta, A. M. (1999). Multiple-input, multiple-output system identification for characterization of limb stiffness dynamics. *Biological Cybernetics*, 80(5):327–337.
- [Perreault et al., 2001] Perreault, E. J., Kirsch, R. F., and Crago, P. E. (2001). Effects of voluntary force generation on the elastic components of endpoint stiffness. *Experimental Brain Research*, 141(3):312–323.
- [Peters et al., 2008] Peters, J., Mistry, M., Udwardia, F., Nakanishi, J., and Schaal, S. (2008). A unifying framework for robot control with redundant dofs. *Autonomous Robots*, 24(1):1–12.
- [Peters and Schaal, 2006] Peters, J. and Schaal, S. (2006). Policy gradient methods for robotics. *IEEE International Conference on Intelligent Robots and Systems*, pages 2219–2225.
- [Petrich et al., 2022] Petrich, L., Jin, J., Dehghan, M., and Jagersand, M. (2022). A quantitative analysis of activities of daily living: Insights into improving functional independence with assistive robotics. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 6999–7006.
- [Piovesan et al., 2019] Piovesan, D., Kolesnikov, M., Lynch, K., and Mussa-Ivaldi, F. A. (2019). The Concurrent Control of Motion and Contact Force in the Presence of Predictable Disturbances. *Journal of Mechanisms and Robotics*, 11(6). 060903.
- [Plagenhoef, 1971] Plagenhoef, S. (1971). *Patterns of Human Motion: A Cinematographic Analysis*. Prentice-Hall, Englewood Cliffs, NY.
- [Platt et al., 2011] Platt, R., Abdallah, M., and Wampler, C. (2011). Multiple-priority impedance control. In *2011 IEEE International Conference on Robotics and Automation*, pages 6033–6038.
- [Rack and Westbury, 1969] Rack, P. M. H. and Westbury, D. R. (1969). The effects of length and stimulus rate on tension in the isometric cat soleus muscle. *The Journal of Physiology*, 204(2):443–460.

- [Raibert and Craig, 1981] Raibert, M. H. and Craig, J. J. (1981). Hybrid position/force control of manipulators. *Journal of Dynamic Systems, Measurement, and Control*, 102(127):126–133.
- [Rancourt and Hogan, 2001] Rancourt, D. and Hogan, N. (2001). Stability in force-production tasks. *Journal of Motor Behavior*, 33(2):193–204. PMID: 11404214.
- [Rancourt and Hogan, 2009] Rancourt, D. and Hogan, N. (2009). The biomechanics of force production. In *Progress in Motor Control*, pages 645–661. Springer.
- [Rohrer et al., 2004] Rohrer, B., Fasoli, S., Krebs, H. I., Volpe, B., Frontera, W. R., Stein, J., and Hogan, N. (2004). Submovements grow larger, fewer, and more blended during stroke recovery. *Motor Control*, 8(4):472–483.
- [Rohrer and Hogan, 2003] Rohrer, B. and Hogan, N. (2003). Avoiding spurious submovement decompositions: a globally optimal algorithm. *Biological Cybernetics*, 89(3):190–199.
- [Ronsse et al., 2009] Ronsse, R., Sternad, D., and Lefèvre, P. (2009). A computational model for rhythmic and discrete movements in uni- and bimanual coordination. *Neural Computation*, 21(5):1335–1370. PMID: 19018700.
- [Rouse et al., 2014] Rouse, E. J., Hargrove, L. J., Perreault, E. J., and Kuiken, T. A. (2014). Estimation of human ankle impedance during the stance phase of walking. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 22(4):870–878.
- [Rouse et al., 2013] Rouse, E. J., Hargrove, L. J., Perreault, E. J., Peshkin, M. A., and Kuiken, T. A. (2013). Development of a mechatronic platform and validation of methods for estimating ankle stiffness during the stance phase of walking. *Journal of Biomechanical Engineering*, 135(8):081009–081009–8.
- [Russell and Hogan, 1989] Russell, D. and Hogan, N. (1989). Dealing with constraints: a biomechanical approach. In *Images of the Twenty-First Century. Proceedings of the Annual International Engineering in Medicine and Biology Society*, pages 892–893 vol.3.
- [Sadeghian et al., 2012] Sadeghian, H., Keshmiri, M., Villani, L., and Siciliano, B. (2012). Null-space impedance control with disturbance observer. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2795–2800.
- [Sadeghian et al., 2013] Sadeghian, H., Villani, L., Keshmiri, M., and Siciliano, B. (2013). Dynamic multi-priority control in redundant robotic systems. *Robotica*, 31(7):1155–1167.
- [Sadeghian et al., 2014] Sadeghian, H., Villani, L., Keshmiri, M., and Siciliano, B. (2014). Task-space control of robot manipulators with null-space compliance. *IEEE Transactions on Robotics*, 30(2):493–506.

- [Schaal, 2006] Schaal, S. (2006). Dynamic Movement Primitives -A Framework for Motor Control in Humans and Humanoid Robotics. *Adaptive Motion of Animals and Machines*, pages 261–280.
- [Schaal et al., 2000] Schaal, S., Kotosaka, S., and Sternad, D. (2000). Nonlinear dynamical systems as movement primitives. In *IEEE international conference on humanoid robotics*, pages 1–11.
- [Schaal and Sternad, 1998] Schaal, S. and Sternad, D. (1998). Programmable pattern generators. In *3rd International Conference on Computational Intelligence in Neuroscience*, pages 48–51.
- [Schaal and Sternad, 2001] Schaal, S. and Sternad, D. (2001). Origins and violations of the 2/3 power law in rhythmic three-dimensional arm movements. *Experimental Brain Research*, 136(1):60–72.
- [Schaal et al., 2004] Schaal, S., Sternad, D., Osu, R., and Kawato, M. (2004). Rhythmic arm movement is not discrete. *Nature neuroscience*, 7(10):1136–1143.
- [Schettino et al., 2020] Schettino, V., Fiore, M. D., Pecorella, C., Ficuciello, F., Allmendinger, F., Lachner, J., Stramigioli, S., and Siciliano, B. (2020). Geometrical interpretation and detection of multiple task conflicts using a coordinate invariant index. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6613–6618.
- [Schreiber et al., 2010] Schreiber, G., Stemmer, A., and Bischoff, R. (2010). The fast research interface for the kuka lightweight robot. In *IEEE Workshop on Innovative Robot Control Architectures for Demanding (Research) Applications How to Modify and Enhance Commercial Controllers (ICRA 2010)*, pages 15–21. Citeseer.
- [Schwartz, 1994] Schwartz, A. (1994). Direct cortical representation of drawing. *Science*, 265(5171):540–542.
- [Sentis and Khatib, 2005] Sentis, L. and Khatib, O. (2005). Synthesis of whole-body behaviors through hierarchical control of behavioral primitives. *International Journal of Humanoid Robotics*, 02(04):505–518.
- [Sergio and Kalaska, 1998] Sergio, L. E. and Kalaska, J. F. (1998). Changes in the Temporal Pattern of Primary Motor Cortex Activity in a Directional Isometric Force Versus Limb Movement Task. *Journal of Neurophysiology*, 80(3):1577–1583.
- [Shadmehr and Mussa-Ivaldi, 1994] Shadmehr, R. and Mussa-Ivaldi, F. (1994). Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience*, 14(5):3208–3224.
- [Shampine and Reichelt, 1997] Shampine, L. F. and Reichelt, M. W. (1997). The matlab ode suite. *SIAM Journal on Scientific Computing*, 18(1):1–22.



- [Shepard and Metzler, 1971] Shepard, R. N. and Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171(3972):701–703.
- [Siciliano and Slotine, 1991] Siciliano, B. and Slotine, J. (1991). A general framework for managing multiple tasks in highly redundant robotic systems. In *Proceeding of 5th International Conference on Advanced Robotics*, volume 2, pages 1211–1216.
- [Slotine and Asada, 1992] Slotine, J.-J. E. and Asada, H. (1992). *Robot Analysis and Control*. John Wiley & Sons, Inc., New York, NY, USA, 1st edition.
- [Sternad, 2008] Sternad, D. (2008). Towards a unified theory of rhythmic and discrete movements—behavioral, modeling and imaging results. In *Coordination: Neural, behavioral and social dynamics*, pages 105–133. Springer.
- [Sternad, 2017] Sternad, D. (2017). *Human Control of Interactions with Objects – Variability, Stability and Predictability*, pages 301–335. Springer International Publishing, Cham.
- [Sternad et al., 2000] Sternad, D., Dean, W. J., and Schaal, S. (2000). Interaction of rhythmic and discrete pattern generators in single-joint movements. *Human Movement Science*, 19(4):627–664.
- [Sternad and Hasson, 2016] Sternad, D. and Hasson, C. J. (2016). *Predictability and Robustness in the Manipulation of Dynamically Complex Objects*, pages 55–77. Springer International Publishing, Cham.
- [Sternad et al., 2010] Sternad, D., Park, S.-W., Müller, H., and Hogan, N. (2010). Coordinate dependence of variability analysis. *PLOS Computational Biology*, 6(4):1–16.
- [Sternad and Schaal, 1999] Sternad, D. and Schaal, S. (1999). Segmentation of endpoint trajectories does not imply segmented control. *Experimental Brain Research*, 124(1):118–136.
- [Stramigioli, 2001] Stramigioli, S. (2001). *Modeling and IPC control of interactive mechanical systems — A coordinate-free approach*. Lecture Notes in Control and Information Sciences. Springer-Verlag London Limited.
- [Stulp et al., 2012] Stulp, F., Theodorou, E. A., and Schaal, S. (2012). Reinforcement learning with sequences of motion primitives for robust manipulation. *IEEE Transactions on Robotics*, 28(6):1360–1370.
- [Svinin et al., 2001] Svinin, M. M., Ohta, K., Luo, Z. W., and Hosoe, S. (2001). Understanding of human movements in crank rotation. In *Proceedings IEEE RSJ International Conference on Intelligent Robots and Systems*, volume 4, pages 2105–2110.

- [Takahashi et al., 2001] Takahashi, C. D., Scheidt, R. A., and Reinkensmeyer, D. J. (2001). Impedance Control and Internal Model Formation When Reaching in a Randomly Varying Dynamical Environment. <https://doi.org/10.1152/jn.2001.86.2.1047>, 86(2):1047–1051.
- [Thorup, 2018] Thorup, J. (2018). Inertia compensation of a planar robot for human upper limb interaction. Master’s thesis, Massachusetts Institute of Technology.
- [Trumbower et al., 2009] Trumbower, R. D., Krutky, M. A., Yang, B.-S., and Perreault, E. J. (2009). Use of self-selected postures to regulate multi-joint stiffness during unconstrained tasks. *PLOS ONE*, 4(5):1–11.
- [Tsao et al., 2022] Tsao, C. W., Aday, A. W., Almarzooq, Z. I., Alonso, A., Beaton, A. Z., Bittencourt, M. S., Boehme, A. K., Buxton, A. E., Carson, A. P., Commodore-Mensah, Y., Elkind, M. S. V., Evenson, K. R., Eze-Nliam, C., Ferguson, J. F., Generoso, G., Ho, J. E., Kalani, R., Khan, S. S., Kissela, B. M., Knutson, K. L., Levine, D. A., Lewis, T., Liu, J., Loop, M. S., Ma, J., Mussolino, M. E., Navaneethan, S. D., Perak, A. M., Poudel, R., Rezk-Hanna, M., Roth, G. A., Schroeder, E. B., Shah, S. H., Thacker, E. L., VanWagner, L. B., Virani, S. S., Voecks, J. H., Wang, N.-Y., Yaffe, K., and Martin, S. S. (2022). Heart disease and stroke statistics-2022 update: A report from the american heart association. *Circulation*, 145(8):e153–e639.
- [Tsuji, 1997] Tsuji, T. (1997). Human arm impedance in multi-joint movements. In Morasso, P. and Sanguineti, V., editors, *Self-Organization, Computational Maps, and Motor Control*, volume 119 of *Advances in Psychology*, pages 357 – 381. North-Holland.
- [Tsuji et al., 1995] Tsuji, T., Morasso, P. G., Goto, K., and Ito, K. (1995). Human hand impedance characteristics during maintained posture. *Biological Cybernetics*, 72(6):475–485.
- [Tuckwell, 1989] Tuckwell, H. C. (1989). *Stochastic processes in the neurosciences*. SIAM.
- [Vallbo and Wessberg, 1993] Vallbo, A. B. and Wessberg, J. (1993). Organization of motor output in slow finger movements in man. *The Journal of Physiology*, 469:673–691.
- [van de Ruit et al., 2020] van de Ruit, M., Cavallo, G., Lataire, J., van der Helm, F. C. T., Mugge, W., van Wingerden, J., and Schouten, A. C. (2020). Revealing time-varying joint impedance with kernel-based regression and nonparametric decomposition. *IEEE Transactions on Control Systems Technology*, 28(1):224–237.
- [Van der Linde et al., 2002] Van der Linde, R. Q., Lammertse, P., Frederiksen, E., and Ruiters, B. (2002). The hapticmaster, a new high-performance haptic interface. In *Proc. Eurohaptics*, pages 1–5.

- [van der Wel et al., 2009] van der Wel, R. P. R. D., Sternad, D., and Rosenbaum, D. A. (2009). Moving the arm at different rates: Slow movements are avoided. *Journal of Motor Behavior*, 42(1):29–36.
- [Verdi, 2019] Verdi, D. (2019). A compositional approach to robotic impedance control. Master’s thesis, Massachusetts Institute of Technology.
- [Viviani and Flash, 1995] Viviani, P. and Flash, T. (1995). Minimum-jerk, two-thirds power law, and isochrony: Converging approaches to movement planning. *Journal of Experimental Psychology: Human Perception*, 21(1):32–53.
- [Viviani and Stucchi, 1989] Viviani, P. and Stucchi, N. (1989). The effect of movement velocity on form perception: Geometric illusions in dynamic displays. *Perception & Psychophysics*, 46(3):266–274.
- [Viviani and Terzuolo, 1982] Viviani, P. and Terzuolo, C. (1982). Trajectory determines movement dynamics. *Neuroscience*, 7(2):431 – 437.
- [Welch, 1967] Welch, P. (1967). The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics*, 15(2):70–73.
- [West et al., 2022] West, A. M., Hermus, J., Huber, M., Maurice, P., Sternad, D., and Hogan, N. (2022). Dynamic primitives limit human force regulation during motion. *IEEE Robotics and Automation Letters*, pages 1–1.
- [Whitney, 1969] Whitney, D. E. (1969). Resolved motion rate control of manipulators and human prostheses. *IEEE Transactions on Man-Machine Systems*, 10(2):47–53.
- [Whitney, 1977] Whitney, D. E. (1977). Force feedback control of manipulator fine motions. *Journal of Dynamic Systems, Measurement, and Control*, 99(2):91–97.
- [Williamson, 2003] Williamson, M. M. (2003). Oscillators and crank turning: exploiting natural dynamics with a humanoid robot arm. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 361(1811):2207–2223.
- [Wolpert et al., 1995] Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. (1995). Are arm trajectories planned in kinematic or dynamic coordinates? an adaptation study. *Experimental Brain Research*, 103(3):460–470.
- [Wolpert et al., 1998] Wolpert, D. M., Miall, R., and Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences*, 2(9):338 – 347.
- [Won and Hogan, 1995] Won, J. and Hogan, N. (1995). Stability properties of human reaching movements. *Experimental Brain Research*, 107(1):125–136.
- [Woodworth, 1899] Woodworth, R. (1899). The accuracy of voluntary movements. *The Psychological Review*, 3(13):1–114.

- [Yoshikawa, 1985] Yoshikawa, T. (1985). Manipulability and redundancy control of robotic mechanisms. In *Proceedings. 1985 IEEE International Conference on Robotics and Automation*, volume 2, pages 1004–1009.
- [Zago et al., 2018] Zago, M., Matic, A., Flash, T., Gomez-Marin, A., and Lacquaniti, F. (2018). The speed-curvature power law of movements: a reappraisal. *Experimental Brain Research*, 236(1):69–82.
- [Zahalak, 1990] Zahalak, G. (1990). Multiple Muscle Systems. chapter Multiple Muscle Systems. Springer New York, New York, NY.
- [Zhang et al., 2018] Zhang, Z., Guo, D., Huber, M. E., Park, S.-W., and Sternad, D. (2018). Exploiting the geometry of the solution space to reduce sensitivity to neuromotor noise. *PLOS Computational Biology*, 14(2):1–20.
- [Zheng and Wang, 2017] Zheng, Z. and Wang, R. (2017). Arm motion control model based on central pattern generator. *Applied Mathematics and Mechanics*, 38(9):1247–1256.
- [Ziegler-Graham et al., 2008] Ziegler-Graham, K., MacKenzie, E. J., Ephraim, P. L., Trivison, T. G., and Brookmeyer, R. (2008). Estimating the prevalence of limb loss in the united states: 2005 to 2050. *Archives of Physical Medicine and Rehabilitation*, 89(3):422–429.