

**MOTOR CORTICAL ACTIVITY RELATED TO THE  
COMBINED CONTROL OF FORCE AND MOTION**

by

**Scott Kennedy**

B.S. in Mechanical Engineering

Washington University in St. Louis, 2009

Submitted to the Graduate Faculty of  
the Swanson School of Engineering in partial fulfillment  
of the requirements for the degree of

**Doctor of Philosophy**

University of Pittsburgh

2018

UNIVERSITY OF PITTSBURGH  
SWANSON SCHOOL OF ENGINEERING

This dissertation was presented

by

Scott Kennedy

It was defended on

July 12, 2018

and approved by

Andrew B. Schwartz, Ph.D., Distinguished Professor, Department of Neurobiology

Peter L. Strick, Ph.D., Distinguished Professor, Department of Neurobiology

Patrick J. Loughlin, Ph.D., Professor and Associate Chair, Department of Bioengineering

Neville Hogan, Ph.D., Sun Jae Professor, Department of Mechanical Engineering,

Massachusetts Institute of Technology

Dissertation Director: Andrew B. Schwartz, Ph.D., Distinguished Professor, Department of

Neurobiology

Copyright © by Scott Kennedy  
2018

# MOTOR CORTICAL ACTIVITY RELATED TO THE COMBINED CONTROL OF FORCE AND MOTION

Scott Kennedy, PhD

University of Pittsburgh, 2018

Using tools, writing, and eating are all important behaviors that involve manipulating objects. Successful manipulation requires the control of both the force exerted on the object and its resultant motion. Both have been associated with neural activity in the motor cortex and we are interested in the extent to which neural firing rates in this brain region are related to their combined control. The mechanical relation between force and motion is impedance and we hypothesized that motor cortical activity encodes an impedance signal that reflects the force and motion demands of behavior. We examined this possibility with a paradigm in which subjects manipulated a handle that moved along a track. The handle was locked in place until the subject exerted enough force to cross a specific threshold; it was then released and moved along the track. We hypothesized that this ballistic-release task would encourage subjects to modify their arm impedance in anticipation of the upcoming movement.

We modeled the behavior as a physical dynamical system and found that one component of model impedance, stiffness, varied in a way that matched the behavioral demands of the task and that stiffness could be dissociated from changes in force and displacement. We recorded activity from a population of motor cortical neurons and found that the temporal and time-averaged neural responses encoded information about motion and force. We also could decode model impedance parameters that we then used to approximate the time-varying force exerted on the handle. The force exerted on the handle and the model stiffness depended on muscle activity and we found components of muscle activity related to both force and model stiffness. Additional components of motor cortical activity were also related

to both force and stiffness, suggesting a possible parceling of muscle-related representations in motor cortical activity. In addition to extending current models of neural activity to include manipulations, this study may be helpful in understanding how information encoded in motor cortical activity might be transformed into muscle activity during object interaction.

## TABLE OF CONTENTS

<b>PREFACE</b> . . . . .	xiv
<b>1.0 INTRODUCTION</b> . . . . .	1
1.1 Mechanics . . . . .	2
1.2 Biomechanics . . . . .	2
1.3 Neuroscience . . . . .	3
1.4 Control strategies . . . . .	4
1.4.1 Impedance control . . . . .	5
1.5 Mechanics, Biomechanics, Neuroscience, Control . . . . .	11
1.6 Contributions . . . . .	11
<b>2.0 DETAILED METHODS</b> . . . . .	14
2.1 Behavioral paradigm and Experimental design . . . . .	14
2.1.1 Force thresholds . . . . .	16
2.1.2 Target zones . . . . .	17
2.1.3 Task conditions . . . . .	18
2.2 Signal measurement . . . . .	21
2.2.1 Real-time position signal . . . . .	21
2.2.2 Force signals . . . . .	21
2.2.3 Motion tracking . . . . .	21
2.2.4 EMG signals . . . . .	22
2.2.5 Neural signals . . . . .	23
2.3 Data preprocessing . . . . .	24
2.3.1 EMG . . . . .	24

2.3.2	Motion tracking	24
2.3.3	Joint angles	24
2.3.4	Neural firing rates	25
2.4	Behavior alignment	25
2.4.1	Force ramp	25
2.4.2	Movement	26
2.5	Physical dynamical system	27
2.5.1	Model selection	28
2.6	EMG electrodes	33
2.6.1	Fabrication	33
2.6.2	Surgical procedures	34
2.7	Neural electrodes	36
2.7.1	Array design	36
2.7.2	Surgical procedures	38
<b>3.0</b>	<b>IMPEDANCE AS A CONTROL FACTOR DURING OBJECT MA-</b>	
	<b>NIPULATION</b>	<b>39</b>
3.1	Abstract	39
3.2	Introduction	40
3.3	Methods	43
3.3.1	Subjects	43
3.3.2	Behavioral paradigm and Experimental design	43
3.3.3	Data collection	44
3.3.4	Physical dynamical model	44
3.3.5	EMG analysis	46
3.3.5.1	Regressing force against EMG	46
3.3.5.2	Separating potent EMG and null EMG	46
3.3.6	Testing the effect of arrest position on the regression of stiffness on force threshold	47
3.4	Results	48

3.4.1	Force and motion varied with the four force thresholds and four target zones . . . . .	48
3.4.2	Arm model impedance varied with force and motion . . . . .	50
3.4.3	Arm posture before movement did not vary across task conditions . . . . .	56
3.4.4	EMG patterns co-varied with force threshold and with target zone . . . . .	56
3.4.5	EMG decomposed into potent and null dimensions . . . . .	57
3.5	Discussion . . . . .	62
3.5.1	Arm impedance estimation . . . . .	62
3.5.2	Muscle-related stiffness during the ballistic movement . . . . .	64
3.5.3	Impedance control . . . . .	65
3.5.4	Conclusion . . . . .	66
<b>4.0</b>	<b>NEURAL ENCODING MODELS FOR COMBINED FORCE-MOTION CONTROL</b> . . . . .	<b>75</b>
4.1	Abstract . . . . .	75
4.2	Introduction . . . . .	75
4.3	Methods . . . . .	79
4.3.1	Subjects . . . . .	79
4.3.2	Behavioral paradigm and Experimental design . . . . .	79
4.3.3	Data collection . . . . .	79
4.3.4	Physical dynamical system . . . . .	81
4.3.5	Model stiffness fit to force threshold and hold position . . . . .	81
4.3.6	De-mixed principal components analysis . . . . .	82
4.3.7	Encoding model for force threshold and target zone . . . . .	83
4.3.8	Decoding the physical model parameters . . . . .	84
4.4	Results . . . . .	86
4.4.1	Force and motion control varied according to force threshold and target zone . . . . .	86
4.4.2	Physical dynamical model fit transient and steady-state behavior . . . . .	87
4.4.2.1	Zero-force position matched hold position . . . . .	87
4.4.2.2	Model impedance depended on both force and motion . . . . .	90

4.4.3	Neural firing rates showed general trends that reflected task parameters	92
4.4.3.1	Time-varying neural responses	92
4.4.3.2	Time-averaged neural responses	96
4.4.4	Neural encoding and decoding models for physical model parameters and force	100
4.5	Discussion	105
4.5.1	Ballistic-release behavior	105
4.5.2	Physical dynamical system	106
4.5.3	Information encoded in neural activity	106
4.5.4	Conclusion	108
<b>5.0</b>	<b>MUSCLE ACTIVITY AND MOTOR CORTICAL ACTIVITY RELATED TO FORCE AND STIFFNESS DURING OBJECT MANIPULATION</b>	<b>111</b>
5.1	Abstract	111
5.2	Introduction	111
5.3	Methods	113
5.3.1	Subjects	113
5.3.2	Behavioral paradigm and Experimental design	114
5.3.3	Data collection	114
5.3.4	Physical dynamical system	117
5.3.5	Ridge regression	117
5.3.5.1	Regressing force on muscle and neural activity	118
5.3.6	Output-potent and output-null dimensions in linear regression	119
5.3.6.1	Force-potent and stiffness-potent muscle activity	119
5.3.6.2	Force-potent and stiffness-potent neural activity	121
5.3.7	Reduced-rank regression	122
5.4	Results	123
5.4.1	Force and position varied with task conditions	124
5.4.2	Model stiffness depended on both force threshold and target zone	124
5.4.3	Muscle activity was modulated with force threshold and target zone	126

5.4.4	Force-potent and stiffness-potent muscle activity . . . . .	128
5.4.5	Force-potent and stiffness-potent motor cortical activity . . . . .	134
5.4.6	Communication components between motor cortex and muscles . . . . .	137
5.5	Discussion . . . . .	139
5.5.1	EMG modulation related to task conditions . . . . .	139
5.5.2	Force-potent and stiffness-potent components of muscle and M1 activity	139
5.5.3	Communication components between motor cortex and muscles . . . . .	140
5.5.4	Conclusion . . . . .	141
<b>6.0</b>	<b>DISCUSSION</b> . . . . .	<b>143</b>
6.1	Directions of force, position, and stiffness . . . . .	144
6.2	Time-varying estimations of the zero-force position and impedance components	144
6.3	Arm, joint, and muscle impedance . . . . .	145
6.4	Conclusion . . . . .	146
	<b>APPENDIX. FORMATTED DATA</b> . . . . .	<b>147</b>
	<b>BIBLIOGRAPHY</b> . . . . .	<b>153</b>

## LIST OF TABLES

1	Force thresholds . . . . .	16
2	Maximum voluntary force . . . . .	17
3	Target zones . . . . .	18
4	Individual muscles from which EMG was recorded . . . . .	34
5	Success rate for each task condition . . . . .	52
6	Individual muscles and their abbreviations from which epimysial EMG was recorded in monkey S. . . . .	116

## LIST OF FIGURES

1.1	Feedback control . . . . .	6
1.2	Impedance control . . . . .	10
2.1	Behavioral paradigm . . . . .	20
2.2	Physical dynamical system . . . . .	30
2.3	Physical dynamical model selection . . . . .	30
2.4	Human force predictions . . . . .	31
2.5	Monkey force predictions . . . . .	32
2.6	Epimysial EMG electrodes . . . . .	35
2.7	EMG patch electrodes and leads . . . . .	36
2.8	Array placement . . . . .	37
3.1	Human methods . . . . .	45
3.2	Human behavioral results . . . . .	51
3.3	Human equilibrium position . . . . .	53
3.4	Human impedance . . . . .	54
3.5	Human stiffness F-test . . . . .	55
3.6	Human EMG . . . . .	58
3.7	Stiffness effect on null EMG . . . . .	60
3.8	Stiffness effect on EMG histogram . . . . .	61
3.9	All human behavior . . . . .	67
3.10	All human equilibrium position . . . . .	68
3.11	Human trial stiffness . . . . .	69
3.12	All human impedance . . . . .	70

3.13 Human EMG approximates force . . . . .	71
3.14 Time-averaged EMG . . . . .	72
3.15 Force in all 3 dimensions . . . . .	73
3.16 Correlation between off-axis force and on-axis stiffness . . . . .	74
4.1 Monkey methods . . . . .	80
4.2 Monkey behavioral results . . . . .	88
4.3 Monkey zero-force position . . . . .	89
4.4 Monkey impedance . . . . .	91
4.5 Neural histogram of days and trials . . . . .	93
4.6 Example firing rate time profiles . . . . .	94
4.7 dPCA for task parameters . . . . .	97
4.8 Example mean firing rates . . . . .	98
4.9 Threshold-target coefficients . . . . .	99
4.10 Decoded physical parameters . . . . .	101
4.11 Estimated force from decoded physical model parameters . . . . .	103
4.12 Time-varying force estimated from the decoded physical parameters . . . . .	104
4.13 Monkey I array placement controls . . . . .	109
4.14 Threshold-target coefficients and laterality . . . . .	110
5.1 Monkey methods . . . . .	115
5.2 Monkey force and position . . . . .	125
5.3 Monkey EP and Stiffness . . . . .	127
5.4 Monkey S EMG . . . . .	129
5.5 Monkey I EMG . . . . .	130
5.6 EMG fit to force . . . . .	132
5.7 EMG fit to stiffness . . . . .	133
5.8 FR fit to force . . . . .	135
5.9 FR fit to stiffness . . . . .	136
5.10 FR-EMG communication components . . . . .	138
5.11 Monkey S Extended EMG . . . . .	142

## PREFACE

The end always seems far away, until it's not, and it comes racing toward you. I must start by thanking my advisor, Andy, for the many thoughtful years of development dedicated to this project and the somewhat rushed finish. I'm looking forward to hearing about the project as it continues to improve and to sharing in the new discoveries that I'm sure will emerge. Andy's tireless dedication and relentless pursuit of "good" science has been a constant motivation and inspiration for me as I developed my own sense of the scientific method.

Over the years, all the members of the Schwartz lab have been a tremendous help. I leaned heavily on the early guidance from older graduate students (Sam, Michael, Andrew, Sagi). Meel and Angus taught me the electrical engineering and coding practices that would serve me well throughout my time building devices and writing programs. And now, the post-docs (Jordan and Hongwei) and younger graduate students (Steve, Josue, Rex, and Ivana) are always willing to help me problem solve and offer insightful suggesting for data analysis and interpretation. Of course, the lab couldn't function without Emrah to keep together the many programs that run the experiments and Kathy to make sure we have all the equipment we need.

I would particularly like to thank Steve Chase, Aaron Batista, Byron Yu and all the members of the Scabby journal club. Your patience and willingness to help me understand advanced computational methods has shaped the way I think about data and challenged my ability to understand the interaction of experimental and computational methods.

Finally, I need to thank my family. My daughters, Isla and Charlotte, have been a constant source of the brain's awe-inspiring ability to learn and perform complex actions with seaming ease. My wife, Lisa, has been the foundation of our family these final four years, and this wouldn't have been possible without her support. Words can't express how much she deserves to be a part of this accomplishment.

## 1.0 INTRODUCTION

Object manipulation is a fundamental human behavior that underlies many aspects of communication, tool use, and activities of daily living. These tasks often involve using the arm and hand to move an object. Tool use and handwriting have played especially important roles in the specialization of the arm as a manipulator, propelling the development of human culture and knowledge. Simple machines arose from the manipulation of simple objects, making it possible for early humans to break free from the limitations of the physical body. Handwriting supports art, math, and medicine, and the handwritten signature is still a defining aspect of a person's identity. Furthermore, daily activities such as brushing teeth, getting dressed, and eating with utensils all depend on the ability to manipulate objects. Infants, in the first months of life, start exploring how their actions can change the environment around them: dropping cups, smearing food, banging toys, etc. Later in life, motor deficits caused by stroke and other neurological disorders make it difficult to accomplish the most mundane tasks and navigate a typical day's activities.

The importance of object manipulation in human behavior is paralleled by the complex anatomy that supports it. Dexterous mechanics in the arm are essential for exerting forces to ensure precision, accuracy, and stability (Hogan, 1985a). Flexible biomechanics in the musculoskeletal system ensure successful manipulation under a wide variety of environmental conditions (Latash, 2018). Vast neural and, more specifically, cortical areas devoted to the arm hint at the information bandwidth needed to control the system (Shadmehr and Krakauer, 2008; Kalaska, 2009; Schwartz, 2016). Although there has been considerable progress in understanding how the individual aspects of musculoskeletal biomechanics and

cortical neural activity relate to object manipulation, a thorough investigation of a control framework that can unify these pieces is missing and would be an important goal to better understand fundamental human behavior.

## 1.1 MECHANICS

Moving an object to perform a task requires that force acts upon the object. Isaac Newton described the Laws of Motion as ([Newton, 1962](#)):

1. an object at rest will stay at rest and an object in motion will stay in motion unless forces act upon it;
2. change in an object's motion is proportional to the force acting on the object;
3. for every action there is an equal and opposite reaction; or, the mutual actions of two objects upon each are always equal and directed in opposite directions.

The first law is important in this study because it describes how the manipulative motion can be achieved: namely, by exerting force on the object. The third law is also important because it emphasizes that the force exerted on the object will act on the hand in the opposite direction, i.e. the physical interaction between the hand and object is fundamentally bi-directional.

## 1.2 BIOMECHANICS

The musculoskeletal system supports the mechanics of manipulation. Muscles govern the forces we use to affect the world around us, with muscle fibers contracting as actin and myosin proteins create cross bridges between the fibers and creep past each other. The movement causes the ends of the muscles to move closer together, pulling on bones that rotate about joints. The bones provide the rigid structures that effectively transfer forces from muscles to an object.

A large number of muscles and bones are involved in most manipulations. The force exerted on the object is governed by rotational forces at the joints called joint torques, which in turn are governed by muscle forces. Many different combinations of muscle forces can produce the same joint torques, and many different combinations of joint torques can exert the same force on the object. The result is immense flexibility in the pattern of muscle forces and joint torques that could perform a manipulation task (Bernstein, 1967; Latash, 2018).

The combinations of muscle forces and joint torques used to manipulate various objects have been described by biomechanists using engineering principles (Todorov, 2000; Scott et al., 2015) and have contributed to advanced prosthetics and robots (Dollar and Herr, 2008; Clites et al., 2018). However, there remain significant challenges for describing manipulation in the real world, i.e. a world where change and uncertainty are prevalent. In fact, by the time most children start school, they are more capable of sophisticated real-world manipulation than advanced robots. To close the gap between engineered robot manipulations and dexterous manipulations of young children, we need to gain a better understanding of how the nervous system controls the musculoskeletal system.

### 1.3 NEUROSCIENCE

Muscle contractions are activated by alpha motor neurons in the spinal cord, each of which is embedded within a local spinal circuit. Spinal circuits receive sensory feedback from the periphery, as well as supra-spinal structures, with some remote structures projecting directly to alpha motor neurons (e.g. cortico-motoneuronal projections) (Rathelot and Strick, 2006). The motor cortex is a cortical region that sends direct and indirect output to alpha motor neurons (Phillips and Porter, 1977).

Electrical excitation of the motor cortex elicits muscle contractions, leading to the historical establishment of the motor cortex as a primary source of information about motor behavior (Fritsch and Hitzig, 1870; Ferrier, 1886). However, the flexibility of the musculoskeletal system, the complex anatomical connections contributing to muscle activation,

and poorly understood control strategies make it difficult to describe the role of the motor cortex in controlling behavior. Motor cortical activity has been shown to encode: muscle activity (Morrow and Miller, 2003), the displacement and torque of single joints (Evarts, 1968; Thach, 1978), the displacement and torque of multiple joints (Reina et al., 2001; Kurtzer et al., 2006), the displacement and force of the hand (Georgopoulos et al., 1986; Kalaska et al., 1989), and many other motor signals. However, the majority of these studies were conducted with the subject performing a task that isolated either force or motion. Behavioral information within motor cortical activity could support a wide variety of control strategies and, because object manipulation often involves the combined control of both force and motion, it is unclear to what extent motor cortical activity encodes different control strategies in the context of object manipulation.

## 1.4 CONTROL STRATEGIES

The effective control of object manipulation involves appropriate motor commands being sent by the central nervous system to the musculoskeletal system, and appropriate interactive behavior between the musculoskeletal system and the object. These two aspects of control can be described in terms of a “signal processing” approach in the case of motor commands and an “energy processing” approach in the case of interaction (Hogan, 2014). Importantly, these two approaches are not mutually exclusive and often are complementary during behavior.

Conventional feed-back control is an example of a signal processing approach, where the system is composed of three processing components: a controller, a plant, and a sensor (Figure 1.1). This conventional feed-back control loop has been successful in describing many aspects of human behavior (Scott, 2004; Diedrichsen et al., 2010; Scott et al., 2015; Schwartz, 2016). In this case, the intended arm movement might be the input signal  $\hat{x}$  with the actual arm movement as the output signal  $x$ . The actual arm movement is compared to the intended arm movement and a discrepancy would be sensed by the sensory systems (sensor) and sent as a feed-back signal to a sensory integrator which generates an error

signal  $e$ . The motor system (controller) would then process the error signal to send motor commands (control signal) to the musculoskeletal system (plant), which corrects the error by generating actual arm movement as the output.

However, the difficulty lies in the delays of the signal transmission. Because of conduction delays in the nervous system and excitation/contraction delays in the muscles, the feedback sensory signal is delayed and degraded and the motor commands generate movement at a delay. The delays could be compensated, for example, by an internal model (Shadmehr and Mussa-Ivaldi, 1994; Lackner and Dizio, 1994) that predicts the future sensory signals from current motor commands, but the internal model depends on learning the dynamics of the musculoskeletal system and the environment.

A compliment to the signal processing approach is the energy processing approach, which can be considered as a feed-forward strategy and is particularly important for object manipulation because it emphasizes the physical interaction between the components of the system (Hogan, 2014). Therefore, when the arm and hand interact with an object during manipulation, the control strategy would be to regulate the combined dynamics of the musculoskeletal system and the object. Because this approach emphasizes the physical interaction, it is subject to Newton's third law which states that for every action between physical components there is an equal and opposite reaction. The interaction can be described in terms of the energy exchanged between the arm and the object, with the rate of energy exchanged being power. In simpler terms, the interaction involves force and motion (velocity), where mechanical power is force multiplied by velocity. The end result can be viewed as the arm exerting force and the object moving with some velocity. The mechanical relation between force and velocity is impedance.

#### 1.4.1 Impedance control

Impedance control emphasizes the physical interaction between components of a system, making it suitable for object manipulation where the arm (one component) exerts force to move an object (a second component). A critical aspect of this approach is the bidirectionality of the interaction, where the object exerts an equal and opposite force on the

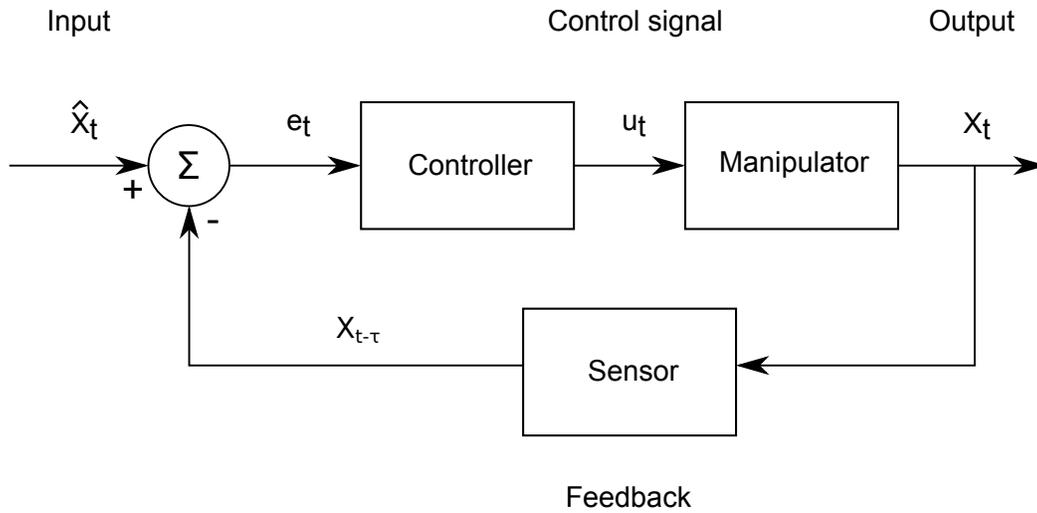


Figure 1.1: **Feedback control.** The user supplies an input  $\hat{x}$  at time  $t$  that is compared with at some delay  $\tau$  with the feedback from the sensor  $x_{t-\tau}$ . Discrepancies between the input and output signals result in an error signal  $e$  that is sent to the controller, which determines a control signal  $u$  that will reduce the error and sends the signal to the plant. The plant executes the control signal and the output  $x$  is fed back by a sensor and becomes  $x_{t-\tau}$ , completing the loop.

hand and the hand/object move in coupled motion. The energy flow during this interaction occurs with no delay, making it especially relevant when it is difficult for an internal model to predict the arm/object dynamics and the interaction is susceptible to errors caused by signal delays. This type of difficulty is present when first contacting an object and when manipulating an unfamiliar or unstable object. The physical interaction between the arm and the object leads to the primary postulate of impedance control, which states:

It is impossible to devise a controller which will cause a physical system to present an apparent behavior to its environment which is distinguishable from that of a purely physical system (Hogan, 1985c).

An intuitive understanding of impedance control can be gained by describing a physical system with multiple individual components, which collectively are referred to as impedance components (Figure 1.2). Although a strict definition of mechanical impedance would describe the mapping of velocity to force, it might also be important to consider position and acceleration variables. Therefore, we refer to impedance as the general mapping between motion, as input, and force, as output. The impedance component associated with position is similar to the stiffness  $K$  of a spring, where  $x_0$  is one end of the spring (reference position),  $x$  is the other end (attached to the object), and  $F$  is the force exerted by the spring on the object.

$$F = K(x_0 - x)$$

When interaction with the object causes the spring to stretch away from its reference position, the spring's stiffness maps the displacement to force. Similarly, the impedance component associated with velocity is like a damper that exerts force to reduce velocity.

$$F = -D\dot{x}$$

Finally, the impedance component associated with acceleration is like a mass that exerts force to reduce acceleration.

$$F = -A\ddot{x}$$

The net effect of these constitutive relations are described in equation 1.1, where  $F$  is the force exerted on the object;  $K, D, A$  are coefficients related to stiffness, damping, and inertia, collectively referred to as impedance;  $x_0$  is the reference position; and  $x, \dot{x}, \ddot{x}$  are the actual

position, velocity, and acceleration. The reference position  $x_0$  is called the zero-force position because zero force is exerted on the object when it is at rest in this position. In other words, the force  $F$  exerted on the object will drive it toward the position  $x_0$  in a way that depends on the impedance coefficients  $K, D, A$ . In essence, the controller is operating as a physical system composed of a spring, damper, and mass whose motion is perturbed when it interacts with the object.

$$F = K(x_0 - x) - D\dot{x} - M\ddot{x} \quad (1.1)$$

Under impedance control, the interaction between the arm and the object is of primary interest. When the object causes the arm to deviate from the zero-force position, the arm exerts force on the object. When the object’s position matches the zero-force position, zero force is exerted on the object. If the zero-force position is stationary, the arm exerts force to maintain the attached object at a steady position according to the impedance coefficients and the motion error as described in equation 1.1. If the zero-force position flows along a trajectory, then the arm exerts force to direct the object along that trajectory.

Critically, combined force-motion control is achieved by changing the impedance coefficients and the zero-force position. With extremely high impedance, motion is controlled accurately. Small deviations from the specified zero-force position map to large corrective forces, reducing error. Low impedance corresponds to force control. Large deviations from the specified zero-force position map to small forces, leading to a mismatch between the actual and controlled position. In this case, the zero-force position is merely an abstract value used to control force. Importantly, impedance is continuously valued and directional, making it possible to modulate the force/motion emphasis in different directions as the behavioral context changes.

In this study, we focused on impedance control to regulate the physical interaction between the arm and the object. We made this decision because we are interested in short time scales where the signal delays in conventional feed-back control might produce errors, such as initially contacting an object. Additionally, impedance control more easily accommodates a variety of object and environment dynamics by conceptually separating the object from the arm (in conventional feed-back control, the object is part of the plant/arm). In practice, it’s likely that some form of feed-back control, in combination with an internal model,

would be used to send motor commands to set the arm's impedance. However, because we are interested in fast time scales and variable object dynamics, we decided that impedance control would be the relevant control strategy for object manipulation in this study.

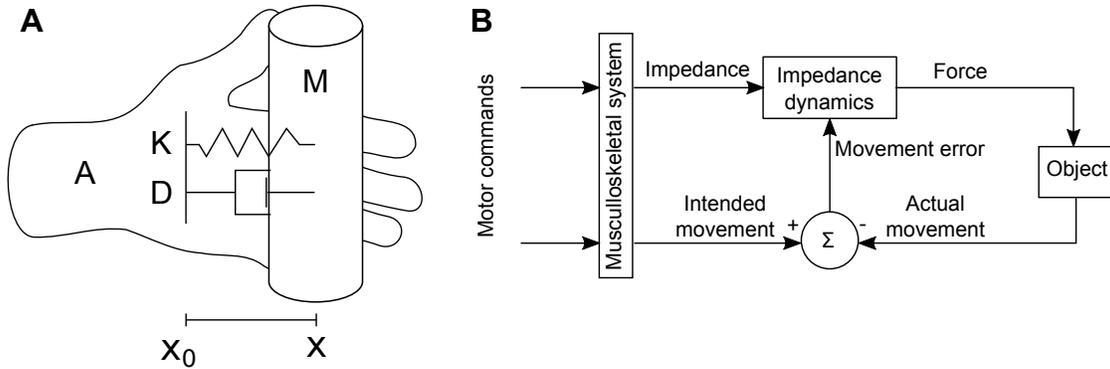


Figure 1.2: **Impedance control.** Impedance control emphasizes the physical interaction between the arm and object. **(A)** The arm is modeled as a physical system consisting of a spring with stiffness  $K$ , a damper with damping  $D$ , and a mass with inertia  $A$ . Collectively, these components describe the impedance of the arm. The arm exerts force on the object when the object's position  $x$  causes the physical system to deviate from its reference position,  $x_0$ . **(B)** This is a control diagram summarizing the physical model. The impedance of the arm and the intended movement (reference position) are set by the motor commands. The impedance dynamics describe the physical interaction with energy exchanged via force and motion at no delay.

## 1.5 MECHANICS, BIOMECHANICS, NEUROSCIENCE, CONTROL

Humans manipulate objects in an unpredictable world, often relying on the spring-like properties of muscles to quickly react to unpredictable interactions and collisions (Franklin and Wolpert, 2011). This type of control can be described as a physical system composed of a spring, mass, and damper. One particular advantage of this description is the combined control of force and motion, critical for many cases of tool use and activities of daily living.

Impedance control is executed by the musculoskeletal system, whose rigid bones and force-producing muscles exert force on the object (Hogan, 1984b). The muscles behave like the spring and damper in the physical system. The bones and body segments behave like the mass. Using different combinations of muscle activations and assuming different arm configurations, the arm can behave like a wide range of physical systems, depending on the goals of the task at hand.

The nervous system provides the motor commands that dictate the impedance of the arm (Latash, 2008). Co-activating antagonist muscles changes the effective stiffness and damping of the arm. Changing the configuration of the arm also affects stiffness by changing muscle length (overlap of the myofibers: shorter length, more overlap, higher stiffness). These factors depend on the activation of the muscle by the alpha motor neurons in the spinal cord. The alpha motor neurons are part of spinal networks influenced by the motor cortex, which encodes high-level information about task goals, such as object motion. In this project we probe the premise that the motor cortex contributes to this impedance framework during object interactions.

## 1.6 CONTRIBUTIONS

The long-term goal of the project is to develop models of neural activity during object manipulation that will expand our understanding of the control strategies employed by the motor system. A step toward that goal, and the objective of this thesis, is to describe the role of the motor cortex in combined force-motion control. Our central hypothesis is that motor

cortical activity encodes an impedance signal and a motion signal that vary according to task conditions. The hypothesis was based on two sets of relations. First, that M1 activity is closely related to motion and the musculoskeletal system. Second, that the musculoskeletal system governs arm impedance. The results presented in this thesis contribute to an understanding of how the motor cortex participates in combined force-motion control, extending the flexibility of neural encoding models to more closely resemble the control observed during healthy behavior. These models can be used in brain-machine interfaces to investigate cortical principles of learning and performance during many activities of daily living.

The objective of this thesis was achieved by addressing three *specific aims*:

1. **Establish conditions under which force and impedance were modulated independently.**
2. **Develop neural encoding models that describe information about impedance and motion.**
3. **Determine the extent to which separate components of muscle activity and neural activity are related to force and stiffness.**

We adopted a behavioral paradigm to study impedance as a primary factor in the combined control of force and motion during object interaction. In addition, by exploiting this paradigm, we found signs of a strategy used for the combined control of force and motion. These results will help form a neural framework for understanding object manipulation. Such a framework is expected to *advance the field* by making it possible to study neural principles involved in regulating interaction dynamics, as detailed in the following chapters.

- Chapter 2 describes the details of the experimental design, equipment fabrication, and surgical procedures.
- Chapter 3 presents the validation of the experimental paradigm and the development of a novel method for relating muscle activity to arm impedance.
- Chapter 4 develops neural encoding models that capture impedance and motion information and predict time-varying behavior using minimal parameters.
- Chapter 5 explores the possibility that the relation between motor cortical activity and muscle activity depends on the functional role of the muscle activity.

- Chapter 6 discusses the context of the results, limitations of the current study, and opportunities for improvement in future work.

## 2.0 DETAILED METHODS

This chapter provides detailed information about the overall thesis experiment and methods, augmenting the individual methods sections of the following chapters. An initial experiment was conducted on five human subjects, four men and one woman between the ages of 20 and 40. All subjects were right handed and had no known neurological deficits. The protocol was approved by the University of Pittsburgh’s Institutional Review Board. We used the human experiment to address any issues in the experimental equipment and design before conducting the experiment on two rhesus monkeys. Monkey I weighed 14 kg and monkey S weighed 12 kg. The protocol was approved by the University of Pittsburgh’s Institutional Animal Care and Use Committee. Most experimental parameters remained the same across all subjects, but some were modified and are noted in the text.

### 2.1 BEHAVIORAL PARADIGM AND EXPERIMENTAL DESIGN

The objective of this thesis was to describe the role of the motor cortex in combined force-motion control. Our central hypothesis was that motor cortical activity encodes an impedance signal and a motion signal that vary according to task conditions. To test this hypothesis, we needed an experimental paradigm to dissociate force, motion, and impedance. We adopted a pre-loaded, ballistic-release task where subjects were required to overcome four force thresholds and move their arm to one of four different target zones ([Viviani and Terzuolo, 1973](#); [Elliott et al., 1999](#)).

Subjects were seated with their torsos restrained to minimize extraneous movement during the task. A linear manipulandum was mounted in front of the subject at shoulder height (Figure 2.1A/D). The manipulandum consisted of a handle mounted on a sled that could be moved along a straight track. The device was oriented in the frontal plane with the starting end of the track aligned to the subject's left shoulder. An electromagnet (Rectangular, 12V DC, 8W, McMaster-Carr) was activated to lock the sled in place with a microcontroller (Mega 2560, Arduino) using custom software (dragonfly-msg.org). A start button was near the subject's right shoulder (see Figure 2.1A/D)

Each subject was instructed to (1) use their right hand to press the start button and then reach to grasp the handle, (2) pull on the handle with enough force to unlock it, and (3) position the handle within a specified target zone. Real-time feedback about the handle's location and the target zone was displayed on a monitor in front of the subject. However, the subject did not receive any direct feedback about the pulling force exerted on the handle or about the force necessary to unlock the handle.

To be successful, the subject needed to pull with enough force to unlock the handle and hold it in the specified target zone for 300 ms (Figure 2.1B/E). Exiting the target zone before 300 ms had elapsed caused the trial to fail. An auditory cue indicated success or failure. The subject then returned the handle to the lock position and initiated the next trial by pressing the start button.

For human subjects, successful and unsuccessful trials were both included in the analyses. For monkey subjects, only successful trials were included. Subjects were given 6 seconds to unlock the handle. The trial was aborted if the subject failed to move the handle in this period and the trial was re-initiated by pressing the start button. These incidents were not counted as a failure nor included in the analyses.

In the task, subjects were required to arrest their movement in one of four target zones and to overcome one of four force thresholds. In theory, an efficient strategy would be similar to attaching a spring to the locked handle and stretching the other end of the spring (the zero-force position) to the target zone; when the handle released, it would come to rest at the zero-force position (Feldman, 1966, 1986; Polit and Bizzi, 1979). For the same target zone/zero-force position, the force exerted on the handle could be increased by increasing the

spring’s stiffness (arm impedance). Thus, the zero-force position and arm impedance could be preset for a given target zone and force threshold, eliminating the need to rely on corrective interventions during the movement. Instead of emphasizing corrective interventions, this simple paradigm focuses on the anticipatory control that is likely to take place during object manipulation.

### 2.1.1 Force thresholds

Four force thresholds were chosen for each subject (Table 1). To cross the threshold, only force along the movement direction was considered. Force in the vertical and anterior/posterior directions was left unconstrained. Because of hardware limitations, there was an unavoidable random 10-30 ms delay between threshold detection and magnet release. This delay was not measured by the recording system, leading to the decision to align time on movement onset (see Section 2.4).

Table 1: Force thresholds

Subject		Force threshold			
		Lowest			Highest
Human	(% MVF)	5	20	35	50
Monkey I	(N)	10	14	17	21
Monkey S	(N)	-5	10	20	30

The four force thresholds for each human subject were chosen as a percentage of the subject’s maximum voluntary force (MVF) (Table 2). At the beginning of the experimental session, each subject was asked to grasp the handle and pull as hard as possible while the handle was locked in place. We measured the maximum force along the direction of motion.

Force thresholds for the monkeys were determined during training. The highest threshold was set at the highest level for which the monkey could still reliably succeed at all task conditions. The remaining force thresholds were spaced to provide approximately even sampling.

Table 2: Maximum voluntary force

Subject	Maximum Voluntary Force (N)
Human subject 1	160
Human subject 2	100
Human subject 3	160
Human subject 4	200
Human subject 5	200

The negative force threshold for monkey S ensured that the handle was unlocked before the monkey grasped the handle, serving as a control condition where the handle immediately started to move and no force threshold needed to be crossed.

### 2.1.2 Target zones

Four target zones (Table 3) were displayed on a monitor at eye-level in front of the subject (Figure 2.1A/D). The target was shown as soon as the start button was pressed.

For the human subjects, all four target zones began at the same start position, near the lock position of the handle, and the end of the four targets were spaced along the track. The target zones were chosen to have different widths because we wanted to encourage the subjects to modulate arm impedance and arm impedance has been shown to increase with positional accuracy (Gribble et al., 2003; Selen et al., 2006).

The target zones for the monkey subjects were originally the same as the human subjects. However, whereas the humans tended to stop the handle in the middle of the target zone, the monkeys tended to stop the handle in the same position, regardless of the target zone. Therefore, four target zones were selected with the same width but different center positions.

Table 3: Target zones

Subject	Target Zones (cm)			
	Closest			Farthest
Human	2.50 - 5.00	2.50 - 7.50	2.50 - 12.50	2.50 - 22.50
Monkey I	2.50 - 10.00	6.25 - 13.75	10.00 - 17.50	13.75 - 21.25
Monkey S	2.50 - 8.75	6.88 - 12.50	10.63 - 16.88	15.00 - 21.25

### 2.1.3 Task conditions

A single task condition was composed of a force threshold and a target zone. There were four force thresholds and four target zones, resulting in 16 task conditions (Figure 2.1C/F). The task conditions were presented in blocks of repeated trials. The order in which the task conditions were presented remained the same for each subject.

For the human subjects, each task condition was composed of 20 repeated trials. The task conditions began with the lowest threshold and farthest target. This block was followed by another with the same threshold, but with the second-farthest target. Two more blocks were completed with targets that moved progressively closer to the handle’s lock position. The next lowest threshold was then presented for the farthest target and the pattern continued. The last task condition was the highest threshold and the closest target.

For the monkey subjects, each task conditions was composed of 30 and 20 repeated trials (monkey S and I). The task conditions also began with the lowest threshold and farthest target. However, this block was followed by another with the same *target*, but with the second-lowest *threshold*. The last task condition was again the highest threshold and the closest target. Grouping the task conditions by target zone made it easier for the monkeys to adjust to the transition between new task conditions because the transition between different targets was more difficult than the transition between different thresholds.

All subjects were able to rest whenever necessary to prevent fatigue. In addition, a 30-second rest period was given between task conditions. A longer 60-second rest was given between every fourth task condition.

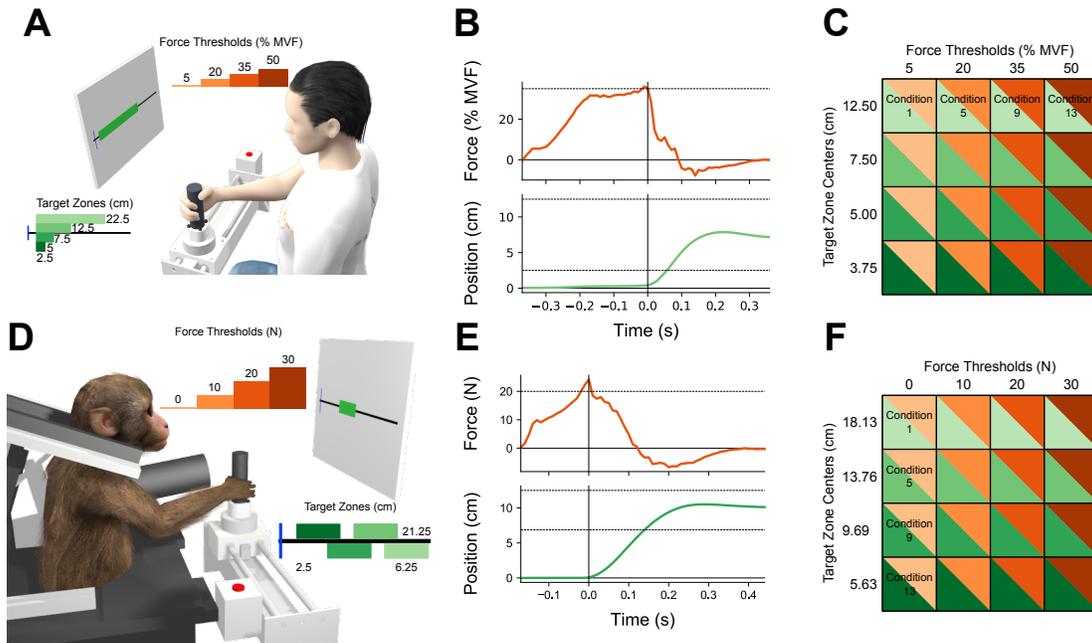


Figure 2.1: **Behavioral paradigm.** **(A-C) Human subjects.** **(A)** The subject was seated in front of a handle with the torso restrained. A monitor displayed the position of the handle along a track (vertical blue line) and a target zone (green rectangle). To initiate a trial, the subject first pressed the start button (to the subject’s right) and then grasped the handle. The subject pulled on the handle while it was locked in place until the force threshold was crossed. The handle was then unlocked to move freely along the track and the subject had to stop and hold the handle within the target zone for 300 ms. **(B)** Time-series of force and position were measured for each trial (a single representative trial is displayed here). At time 0, the handle was unlocked and free to move along the track. The dashed line on the force plot was the force threshold and the pair of dashed lines on the position plot were the near and far boundaries of the target zone. **(C)** A single task condition was composed of a force threshold and a target zone. **(D-F) Monkey subjects.** The structure of each panel is the same as the human subjects.

## 2.2 SIGNAL MEASUREMENT

During the task, we measured: (1) the real-time position of the handle; (2) the force exerted on the handle; (3) the 3D position of optical markers placed on the handle and the subject's hand, lower arm, upper arm, and torso; (4) the electromyographic (EMG) activity of muscle groups or individual muscles; and, for the monkey subjects, (5) neural activity in the motor cortex.

### 2.2.1 Real-time position signal

The handle's real-time position along the track was represented on the monitor as a blue bar. The handle's position was measured by a microcontroller (Mega 2560, Arduino) that sampled a linear potentiometer (SoftPot 300.00mm, Spectra Symbol) at 100 Hz. The voltage output of the linear potentiometer changed as a wiper, attached to the sled, moved along the potentiometer's surface. The mapping from voltage to position was approximately 6.25 mV/cm.

### 2.2.2 Force signals

Force was measured using a six degree-of-freedom force transducer (Delta FT, ATI Industrial Automation, Inc.) and sampled at 100 Hz. The transducer was mounted between the handle and the sled. In the following text, we refer to "force" as the linear force component exerted by the hand on the handle in the direction co-linear with the track.

### 2.2.3 Motion tracking

Motion was measured using a passive, infrared motion tracking system (Nexus 1.8.5, Vicon, Inc.) that sampled at 100 Hz. The system consisted of twelve cameras and eighteen markers that were either 5 or 10 mm in diameter. One marker was placed on either the handle or the sled. The remaining markers were placed on the subject.

For the human subjects, four markers were on the subject's torso at the right acromion, left acromion, jugular notch, and xyphoid process. A set of three markers was placed as a rigid triangle on the lateral side of the upper arm, another set on the lateral side of the lower arm, and the final set on the dorsal side of the hand.

For both monkey subjects, a set of three markers was placed as a rigid triangle on the torso (approximately at the xyphoid process), another set attached to the lateral side of the upper arm, another set attached to the lateral side of the lower arm, and the final set was placed on the dorsal side of the hand.

Although the handle's real-time position was displayed to the subject using the signal from the linear potentiometer, this signal was noisy. For the human subjects and monkey I, the marker location on the handle was used for analysis. However, there was some bending of the handle (on the order of mm) for high forces, and the marker was moved to the sled for monkey S.

#### **2.2.4 EMG signals**

Bipolar EMG activity was recorded in two different ways. Surface EMG was recorded from the human subjects and from monkey I. Epimysial EMG was recorded from monkey S. The surface EMG system (AMT-8; Bortec Biomedical) included a DIO card (DAQ-2208, ADLink Technology Inc.) and MATLAB (Mathworks). The raw surface signal was bandpass filtered at 10-1000 Hz and sampled at 2000 Hz. The epimysial EMG system consisted of a wireless transceiver and neural signal processor (Mira Link and Grapevine, Ripple). The raw epimysial signal was low pass filtered at 500 Hz and sampled at 2000 Hz.

For the human subjects, eight bipolar surface EMG signals were recorded using sixteen electrodes (Pediatric electrodes, Vermed Inc.). Two electrodes were placed over the following muscle groups: wrist flexors, wrist extensors, elbow flexors, elbow extensors, anterior deltoid, posterior deltoid, pectoralis major, and rotator cuff. The electrodes were placed using anatomical landmarks and verified by displaying the signal on an oscilloscope and asking the subject to activate the different muscle groups. A global reference electrode was placed on the back of the right hand.

For monkey I, six bipolar surface EMG signals were recorded using twelve electrodes (Pediatric electrodes, Vermed Inc.). Two electrodes were placed over the following muscle groups: wrist flexors, wrist extensors, elbow flexors, elbow extensors, anterior deltoid, and posterior deltoid. A global reference electrode was placed on the distal, posterior area of the right rib cage.

For monkey S, thirty-two raw epimysial EMG signals were recorded using separate electrodes (see Section 2.6 for details). A pair of electrodes were placed over individual muscles in the lower arm, upper arm, and torso. Bipolar epimysial EMG signals were obtained by subtracting the two channels on each muscle off-line. Separate ground and global references were positioned subcutaneously on the distal, posterior area of the right rib cage.

### 2.2.5 Neural signals

For the monkey subjects, neural activity was recorded from two micro-electrode arrays placed in the arm/hand area of the motor cortex. The neural recording system included a headstage, amplifier, and neural signal processor that high-pass filtered the raw signal at 200 Hz and sampled at 40 kHz (Cerestage, DigiAmp, OptiPlex; Plexon Inc.).

Both monkeys had a 96-channel array placed on the pre-central gyrus (Utah array, Blackrock Microsystems). In addition, monkey I had a 64-channel and monkey S had a 128-channel array placed anterior to the rostral bank of the central sulcus (MatrixHD array, NeuroNexus).

Prior to recording, channel thresholds were initially set at -4 times the standard deviation of a 10-second recording of the signal and some channels were adjusted to help discriminate waveforms. Afterward, all thresholds remained set throughout recording. For each channel, we recorded the time of each threshold crossing and an 800 micro-second waveform, beginning 200 micro-seconds before the threshold crossing.

## 2.3 DATA PREPROCESSING

### 2.3.1 EMG

For each muscle, the bipolar EMG signals were high-pass filtered at 100 Hz. They were then mean-centered and scaled by their standard deviation across time (z-scored). Finally, each signal was squared, low-pass filtered at 30 Hz, square-root transformed, and multiplied by 2 (to recover the signal power). The resulting signal envelopes were down-sampled to 100 Hz.

### 2.3.2 Motion tracking

The motion tracking markers were manually labeled offline. For the human subjects and monkey S, gaps in each marker’s trajectory were filled using spline or source-fitting tools available in the system software (Nexus 1.8.5 and Nexus 2.6.1; Vicon, Inc.). For monkey I, gaps in each marker’s trajectory were filled using a custom Kalman filter.

The reference frame of the track was defined according to the principal components of the handle’s 3-D position in the motion tracking reference frame. The first component was the direction along the track and the remaining components were ordered so that the second dimension pointed away from the subject and the third dimension pointed upwards. The origin of the reference frame was the handle/sled’s start position.

Time derivatives were calculated from position using successive application of a Savitzky-Golay filter with a window length of 7 time bins, a polynomial order of 3, and a derivative order of 1 (`scipy.signal.savgol_filter`).

### 2.3.3 Joint angles

Joint angles were calculated using musculoskeletal modeling software (OpenSim) (Delp et al., 2007) and a generic musculoskeletal model of a human torso (Holzbaaur et al., 2005) and a monkey torso. First, the appropriate generic model was scaled for each subject using markers placed on bony landmarks. The marker positions were recorded during a single static

recording period. Markers were placed on the following bony landmarks: right acromion, left acromion, jugular notch, and xyphoid process; medial and lateral condyles of the humerus; radial and ulnar styloid processes.

Next, joint angles were calculated using OpenSim’s inverse kinematics algorithm and the marker positions recorded during the behavioral task. In brief, the algorithm found the joint angles that minimized the error between the measured marker positions and a set of virtual marker positions placed on the model. Joint centers were calculated using the joint angles and the scaled model.

### **2.3.4 Neural firing rates**

Offline, the threshold crossings of triggered voltage waveforms on each channel were sorted into neural units. Cross-channel artifacts were removed by identifying threshold crossings that occurred at the same time on 60% of the channels. Waveforms on each channel were aligned to the minimum voltage and algorithmically sorted (Valley seek sort, Offline Sorter; Plexon Inc, Dallas, TX). Principal component analysis (PCA) was performed on the temporal dimensions of the waveforms and sorting was manually adjusted using the distribution of the waveforms in PCA space.

Spike times were converted to fractional-interval firing rates with 10 millisecond time bins ([Schwartz, 1992](#)). The same neural unit was identified across recording sessions ([Fraser and Schwartz, 2012](#)).

## **2.4 BEHAVIOR ALIGNMENT**

### **2.4.1 Force ramp**

Each trial included a time period when the subjects exerted force isometrically on the handle, which we refer to as the force ramp. The force ramp began when the force exerted on the handle rose above 1 N for the last time and ended at movement onset. Movement onset began when velocity rose above 10% of the trial’s maximum velocity for the first time.

For the human subjects, the duration of the force ramp varied primarily as a function of the force threshold, with lower thresholds having a shorter duration than longer thresholds. To compare task conditions with the same force threshold and different target zones, we averaged the force ramp duration across trials with the same force threshold and scaled the duration of each trial’s force ramp to match the average using a piecewise cubic hermite interpolating polynomial. The task condition’s trial-averaged data during the force ramp were then used in the analyses.

For the monkey subjects, the force ramp was consistent across all task conditions, although considerably shorter in duration than the human subjects. Because of the consistency, the force ramp duration for each trial was matched to the average force ramp duration across all trials using a piecewise cubic hermite interpolating polynomial. Trials with a prolonged force ramp that started before the trial-averaged force rose above 1 N were excluded

#### **2.4.2 Movement**

Behavior following movement onset was not scaled in time because of the importance of the position time derivatives in the impedance analysis. The impedance hypothesis predicts that, in this task, the zero-force position is where the handle position will be stopped and held, making the hold position an important factor in our stiffness analyses. For the monkey subjects, movement ended when velocity remained below 2.5 cm/s for 100 ms. The hold position was the handle’s position at the end of movement.

Although the human subjects were instructed to hold the handle within the target zone for 300 ms, it was common for them to return the handle to the start position without holding for unsuccessful trials. In addition, the wide human target zones made it possible for the human subjects to enter the target zone and begin to return the handle to the start position while still remaining in the target zone for 300 ms (a successful trial). Because both successful and unsuccessful trials were included in the human analyses, these unanticipated behaviors made it difficult to consistently evaluate the estimated zero-force position across all task conditions.

To address this difficulty, we examined the position trajectory when the human subjects did hold the handle at a final position and found that the position when velocity first changed sign, which we call the arrest position, was a close approximation to the hold position. When the handle asymptotically approached the hold position, the arrest position was nearly identical to the hold position. When the handle oscillated slightly toward the end of the movement, the arrest position was slightly farther from the start position than the hold position. Therefore, the arrest position was used for the human subjects to evaluate the zero-force position and to determine the effect of movement constraints on stiffness.

## 2.5 PHYSICAL DYNAMICAL SYSTEM

We modeled the behavior of the arm as an equivalent physical dynamical system with impedance components that exert force proportional to position and its time derivatives. Figure 2.2 is a graphical depiction of a specific form of the model with three impedance components. A spring with stiffness  $K$  exerts force proportional to position, a damper with damping  $D$  exerts force proportional to the first derivative of position, and a mass with inertia  $A$  exerts force proportional to the second derivative of position.

The model makes it possible to describe specific characteristics of how the arm interacts with the handle. The arm exerts force on the handle when the zero-force position  $x_0$  is different than the handle position  $x$ . Equation 2.1 describes the force exerted by the arm on the handle, where  $F$  is force;  $K, D, A, J$  are impedance coefficients related to stiffness, damping, inertia, and a third-order term;  $x, \dot{x}, \ddot{x}, \overset{\cdot\cdot\cdot}{x}$  are the handle's position and its first 3 time derivatives; and  $x_0$  is the zero-force position, assumed to be constant. Because we were unable to measure the impedance of the arm directly, the impedance coefficients are related to, but not necessarily accurate, representations of the stiffness, damping, and inertia of the arm. The third-order impedance component (not shown in the graphical model) was included to help explain variability in force due to time derivatives of position higher than second order.

$$F(t) = K(x_0 - x(t)) - D\dot{x}(t) - A\ddot{x}(t) - J\overset{\cdot\cdot\cdot}{x}(t) \quad (2.1)$$

The motion of the handle depends on the difference between the force exerted by the arm and the force  $F_{ext}$  exerted by the magnet. When the magnet locks the handle in place,  $x = \dot{x} = \ddot{x} = \dddot{x} = 0$  and  $F(t)$  depends only on  $x_0$  and  $K$ . After the magnet unlocks the handle, the exerted force depends on the motion, with zero force exerted when the handle is held at the zero-force position.

The free parameters of the model are  $x_0, K, D, A, J$  and can be found by fitting the force and motion data after movement onset to equation 2.1. However, it's possible that the data could be better explained using a first-, second-, or higher-order model. Equation 2.1 can be reduced with fewer terms or expanded with additional terms corresponding to additional time-derivatives of position. The best model can be selected by comparing the cross-validated error of the different models.

### 2.5.1 Model selection

Eight different models were tested for the humans and nine different models for the monkeys according to equation 2.1 with order ranging from 1 to 10. Each model was trained and tested on 200 ms of data for the humans and 500 ms of data for the monkeys, beginning at movement onset and using four-fold cross-validation. The data averaged across trials for one fold was used for testing and the data averaged across the remaining trials was used for training. Across all folds, the model parameters were fit to the training data and then used with the motion test data to predict force. Each model was evaluated using the root mean squared error between the predicted and actual force across folds, conditions, and time.

Higher-order models can often explain more variance in the data, reducing the error. However, some of this variance is noise. If the model parameters are fit to noise in the training data, then the error in the testing data will be high. The errors in figure 2.3 generally follow this trend, reaching a minimum at model order 8 or 9, indicating that variability explained by higher order models is likely noise. However, decreases in the human error after order 3 are modest, suggesting that a third-order model is sufficient to explain most of the variance while retaining a simple, interpretable form. Although the monkey error does not exhibit the same plateau, the advantages of an eighth-order model are outweighed by the difficulty

of interpretation. A third-order model provides terms related to stiffness, damping, and inertia, in addition to a higher-order term. For the monkeys, the terms related to stiffness, damping, and inertia were qualitatively similar between the third- and ninth-order models. Furthermore, the third-order model provided a good qualitative fit (Figures 2.4 and 2.5, example from one subject and one fold of the testing data).

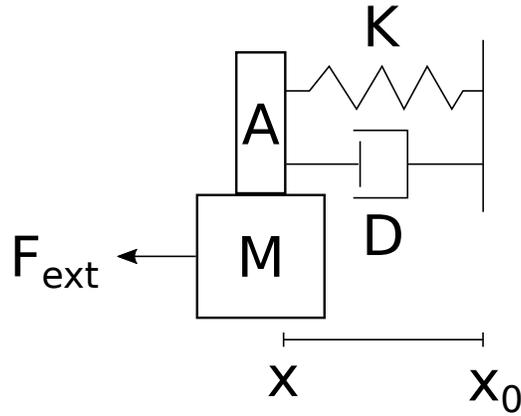


Figure 2.2: **Arm modeled as a physical dynamical system.** The arm is modeled as a physical dynamical system composed of a spring with stiffness  $K$ , damper with damping  $D$ , and mass with inertia  $A$ . Moving the zero-force position  $x_0$  causes the system to exert force on the handle. The motion of the handle depends on the difference between the force exerted by the arm and the force  $F_{ext}$  exerted by the magnet.

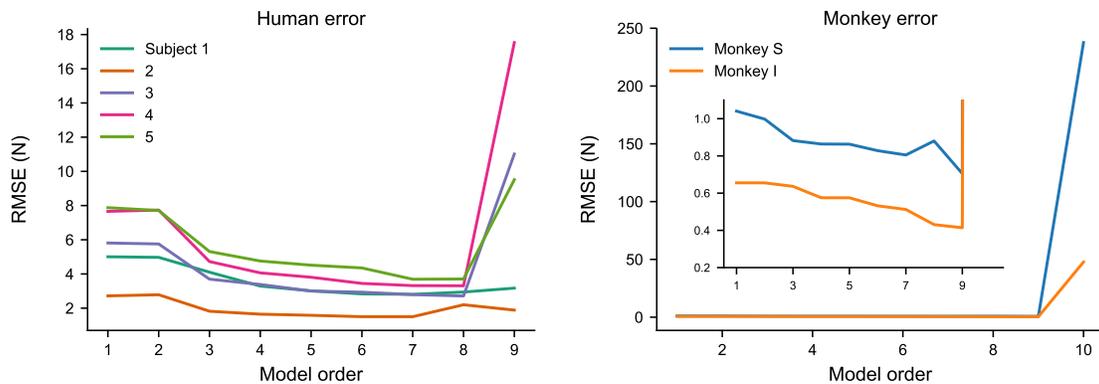


Figure 2.3: **Prediction error for model selection.** Each model was evaluated using four-fold cross validation.

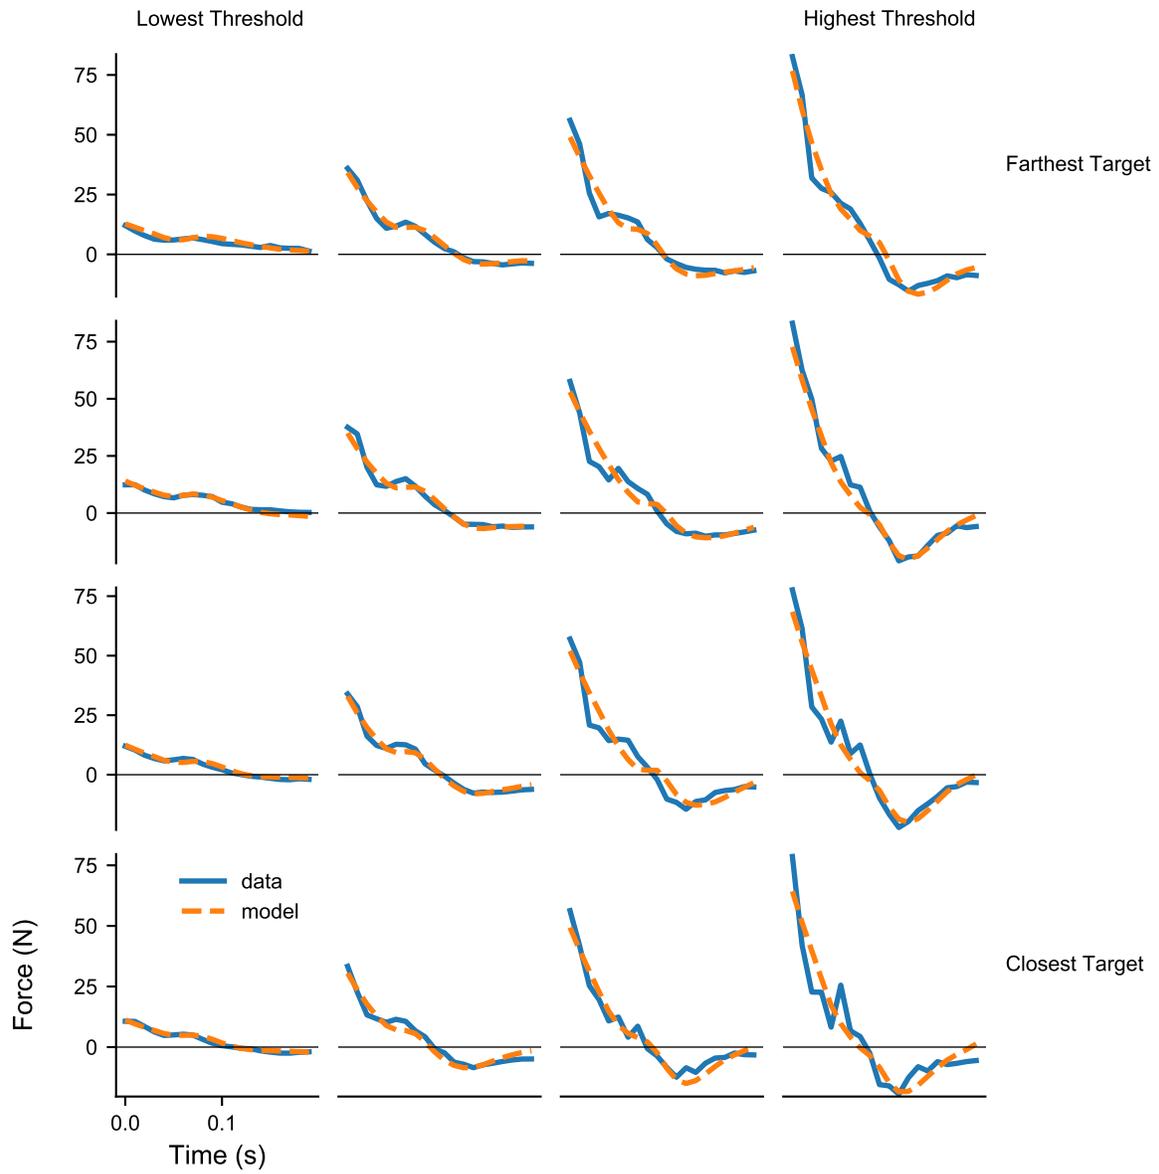


Figure 2.4: **Human model predictions match the data.** Example model predictions and testing data for one human subject and one fold. Predictions were made using 200 ms of data, beginning at movement onset.

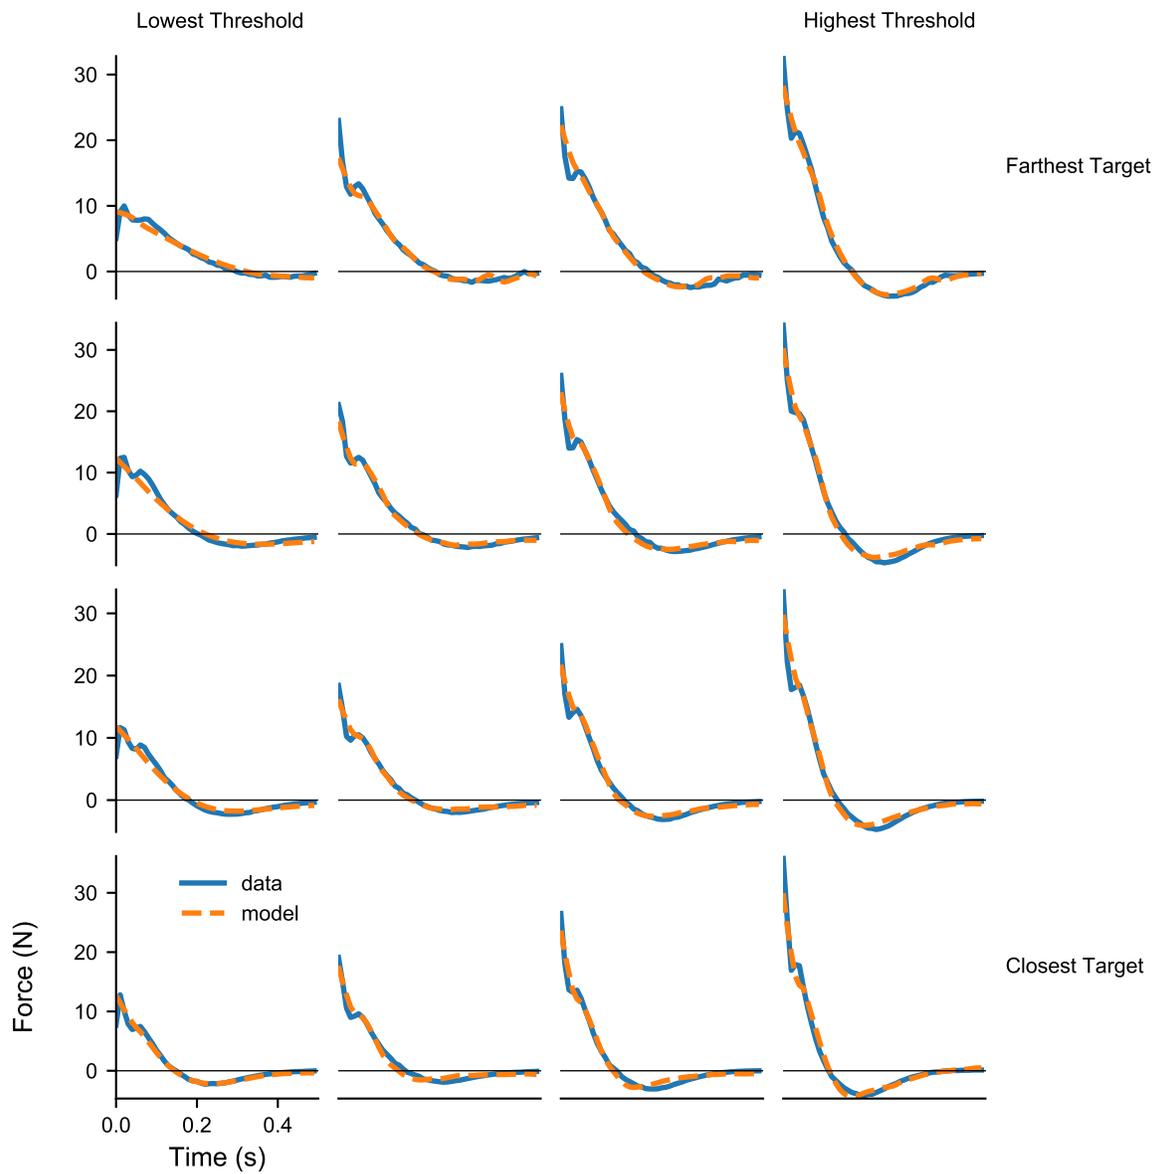


Figure 2.5: **Monkey model predictions match the data.** Example model predictions and testing data for one monkey subject and one fold. Predictions were made using 500 ms of data, beginning at movement onset.

## 2.6 EMG ELECTRODES

Subcutaneous, epimysial electrodes recorded the EMG signal for monkey S (Figure 2.6). These electrodes were custom made according to (Perel, 2012), adapted from (Loeb and Gans, 1986) and (Miller et al., 1993). Individual muscles were selected in the lower arm, upper arm, and torso according to previous literature and the relevance of the muscle’s mechanical action to the task (Table 4).

### 2.6.1 Fabrication

From monkey S, the bipolar EMG from each muscle was recorded using an electrode patch with two leads. The distance from the proposed transceiver location in the back to each muscle was measured and two leads (AS632, Cooner Wire) were cut per muscle, with an additional 10% added to ensure the leads would not limit arm mobility. The ends of the leads were exposed using hot tweezers (Meisei Corp.) and threaded through a silastic patch (SH-21001-007, Bioplexus) (Figure 2.7a). Patch size was either 15 x 15 mm (large) or 10 x 10 mm (small). The exposed wire was 5 mm for the large patches and 3 mm for the small patches. Spacing between the leads was also 5 mm for the large patches and 3 mm for the small patches. Large patches were used for the torso muscles and small patches were used for the remaining muscles. After the exposed leads were threaded through the patch, the ends were folded back and a small amount of silicone adhesive (A-100, Factor II) was placed on the back of the patch and covered with a second, thin silastic layer (SH-20001-002, Bioplexus) to secure and insulate the leads. After the glue was allowed to dry, the leads were twisted together, secured, and allowed to rest until the twist was set. Electrode impedance was tested by placing each patch into a saline bath and connecting the two leads to an impedance meter (1 kHz sine wave, Model IMP-1, BAK Electronics). Typical impedance values were between 1-5 k $\Omega$ .

The electrodes were connected to a transceiver (Mira Link, Ripple) that wirelessly transmitted the signal through the skin. The transceiver leads were exposed using a scalpel and hot tweezers. The stainless steel leads were soldered together using silver solder and liq-

Table 4: Individual muscles from which EMG was recorded

Torso	Upper arm	Lower arm
Pectoralis	Triceps lateral head	Extensor digitorum communis
Latissimus dorsi	Triceps medial head	Extensor carpi unlaris
Anterior deltoid	Biceps short head	Extensor carpi radialis longus
Posterior deltoid	Brachialis	Flexor carpi radialis
Infraspinatus		Palmaris longus
		Flexor digitorum superficialis
		Abductor pollicis longus

uid flux (Stay-brite silver bearing solder kit, Harris Products Group) (Figure 2.7b). The corrosive flux was thoroughly removed (Flux-Off Rosin, Chemtronics). The solder joint was strengthened and insulated with super glue and a curing accelerator (4014 and 7452, Loctite), and then covered with bio-compatible silicone adhesive (A-100, Factor II).

### 2.6.2 Surgical procedures

The electrodes were placed beneath the skin via small incisions 1.5-3 cm in length and oriented proximal-distal. Care was taken to minimize the number of incisions by placing them where multiple muscles could be accessed, but not directly over any muscle where an electrode would be placed. A blunt metal rod, 30 cm long and 1 cm in diameter, was used to tunnel between the incisions and pass the electrodes proximally to distally.

Beginning distally, the muscles were identified anatomically and by electrical stimulation (1 mA/V, 1V pulse, 200  $\mu$ s, 2.5 Hz). The corner of each electrode patch was sutured (3-0 silk) to the epimysium and oriented so that the long-axis of the muscle fiber spanned the two electrodes. The incisions were closed subcutaneously (3-0 polysorb).

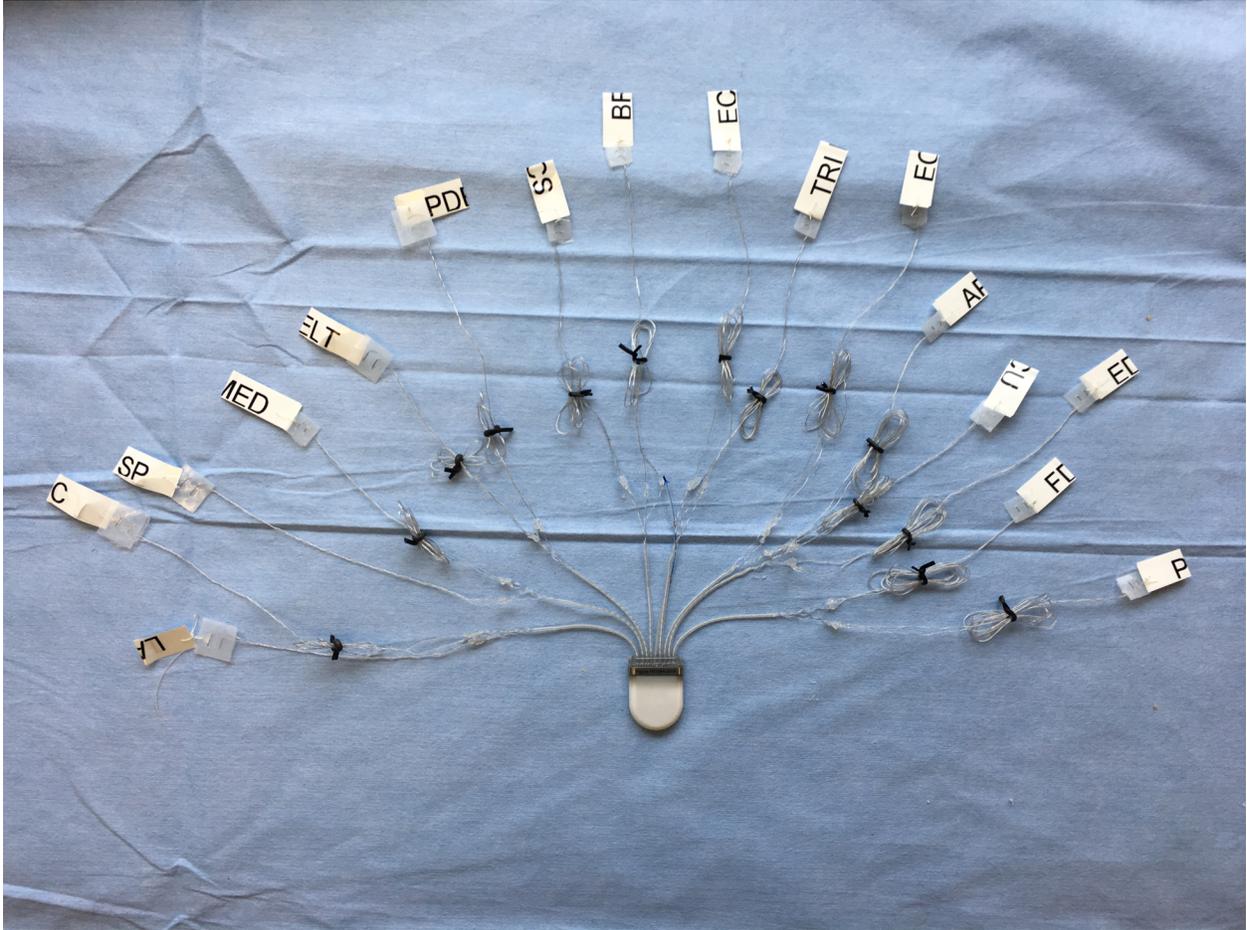
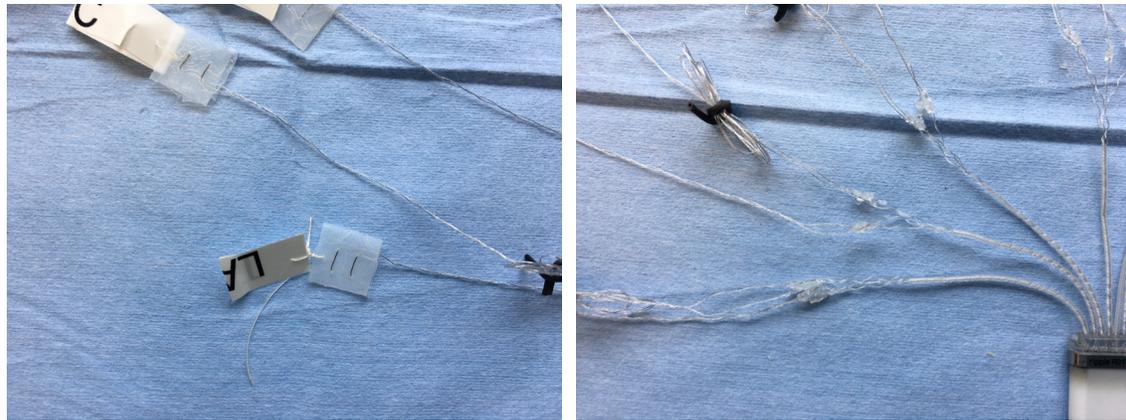


Figure 2.6: **Epimysial EMG electrodes.** Patch electrodes were sutured to the epimysium of each muscle. Electrode leads were spliced to a wireless transmitter in the back that sent the signals through the skin to a receiver.



(a) Patch electrode

(b) Solder joint

Figure 2.7: **EMG leads and patch electrodes.** Patch electrodes were stainless steel wires threaded through a plastic patch. The leads of the electrodes were soldered to the transmitter and covered for biocompatibility.

## 2.7 NEURAL ELECTRODES

Two micro-electrode arrays were placed in the arm/hand area of the motor cortex (left hemisphere) for each monkey (Figure 2.8). A 96-channel array was placed in the gyrus of the arm/hand area of the motor cortex. A second array was placed near the central sulcus of the arm/hand area of the motor cortex and was designed to penetrate along the rostral bank of the central sulcus.

### 2.7.1 Array design

Both monkeys had a 96-channel array (Utah array, Blackrock Microsystems) that was a grid of 10 x 10 shanks, each 1.5 mm long. The array was 4 mm on each side with 400  $\mu\text{m}$  between each shank and a platinum recording site at the tip of each shank.

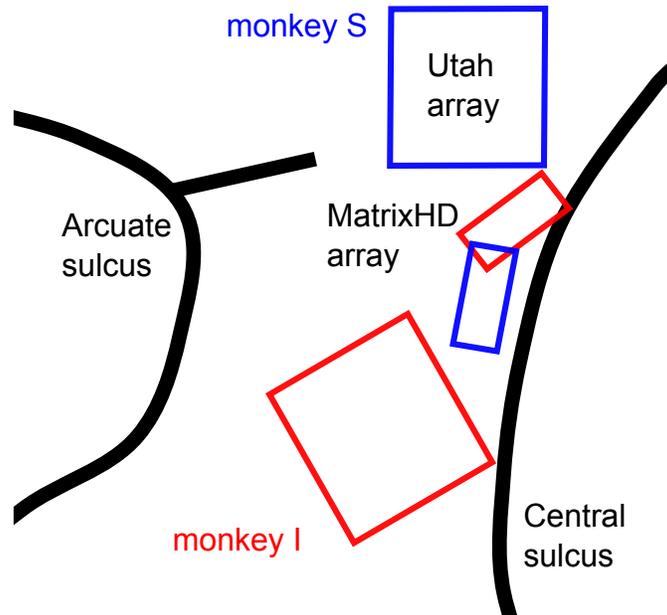


Figure 2.8: Array placement

Monkey I had a 64-channel array (MatrixHD array, NeuroNexus) that had six shanks arranged in two rows of three. The outer two shanks on each row were 8 mm long with 10 recording sites that were  $703 \mu\text{m}^2$  and spaced  $500 \mu\text{m}$  apart. The middle shank on each row was similar to the outer shanks, except for a trio of  $121 \mu\text{m}^2$  recording sites 2.5 mm from the tip of the shank. The three shanks on each row were spaced  $900 \mu\text{m}$  apart and the two rows were spaced  $300 \mu\text{m}$  apart.

Monkey S was implanted with a 128-channel array (MatrixHD array, NeuroNexus) that had 16 shanks arranged in four rows of four. Two rows had shanks that were 2 mm long with 8 recording sites that were spaced  $100 \mu\text{m}$  apart. The two other rows had shanks that were 5 mm long with 8 recording sites that were spaced  $150 \mu\text{m}$  apart. Shanks were spaced  $200 \mu\text{m}$  apart with recording sites of  $703 \mu\text{m}^2$ . All rows were spaced  $300 \mu\text{m}$  apart. The two rows with the 2 mm shanks were placed rostrally to the two rows with the 5 mm shanks. The intention was that this configuration would provide recording sites in the rostral bank of the central sulcus.

### 2.7.2 Surgical procedures

The center of the craniotomy was located 11 mm anterior of the ears and 17 mm lateral of the midline with approximate dimensions 2.5 cm anterior/posterior and 2 cm medial/lateral.

The Utah array was placed on the pre-central gyrus in the arm/hand region of the motor cortex. It was inserted into the cortex using a pneumatic inserter (Blackrock Microsystems).

The MatrixHD array was inserted just anterior to the rostral bank of the central sulcus. The sub-surface shape of the sulcus was inspected using an MRI of each monkey and the insertion was planned so that the recording sites would be located along the rostral bank of the central sulcus. The orientation of the MatrixHD array was set using a motion tracking system (Brainsight Veterinary, Rogue Research Inc.) for monkey I and a stereotax micro-manipulator (Kopf Instruments) for monkey S. Both arrays were inserted using a micro drive with a speed between 0.1 and 0.4 mm/s (IST Motor, NeuroNexus).

### 3.0 IMPEDANCE AS A CONTROL FACTOR DURING OBJECT MANIPULATION

This chapter was submitted for publication to the *Journal of Neurophysiology* and published as a preprint on bioRxiv.

Scott D. Kennedy, Andrew B. Schwartz. Stiffness as a control factor for object manipulation. *bioRxiv* 339101, doi: 10.1101/33910.

Note: Here we refer to the zero-force position as the equilibrium position. They are conceptually the same, but the equilibrium position was used here because we decided later that the zero-force position was a more descriptive name.

#### 3.1 ABSTRACT

We act on the world by producing forces that move objects. During manipulation, force is exerted with the expectation that an object will move in an intended manner. This prediction is a learned coordination between force and displacement. Mechanically, impedance is a way to describe this coordination. As an efficient control strategy, object interaction could be anticipated by setting impedance before the hand moves the object. We examined this possibility with a paradigm in which subjects moved a handle to a specific target position along a track. The handle was locked in place until the subject exerted enough force to cross a specific threshold; then the handle was abruptly released and could move along the track. We hypothesized that this ballistic-release task would encourage subjects to modify their arm

impedance in anticipation of the upcoming movement. If we consider the handle as an object, this paradigm loosely approximates the uncertainty encountered at the end of a reach when contacting a fixed object. We found that one component of arm impedance, stiffness, varied in a way that matched the behavioral demands of the task and we were able to dissociate stiffness from changes in force and displacement. We also found separate components of muscle activity that corresponded to stiffness and to changes in force. Our results show that subjects used a robust and efficient strategy to coordinate force and displacement by modulating muscle activity in a way that was behaviorally relevant in the task.

### 3.2 INTRODUCTION

Manipulating objects is fundamental to human behavior and requires flexible, coordinated control of both force and movement ([Kawato, 1999](#); [Wolpert and Ghahramani, 2000](#); [Flanagan et al., 2006](#); [Franklin and Wolpert, 2011](#)). Identifying a control scheme that takes place during this behavior would be a step toward detecting and understanding the brain signaling underlying the way we interact with objects. Depending on the task conditions, multiple strategies can be used to achieve this coordination. The most challenging behaviors are those in which rapid, precise manipulation takes place. While humans perform these movements with great skill, the control principles underlying this behavior is a topic of interest for both scientists and engineers.

Roboticians use control schemes that utilize rapid feedback to monitor ongoing changes in displacement and force. To manipulate an object, motion of the robotic effector is controlled precisely by exerting the force needed to achieve the movement. In real-world conditions, this can be problematic when unexpected collisions take place or if interaction with the object leads to large changes in force – for instance, when displacement is measured inaccurately. The generation of these large forces makes robots dangerous in the workplace, less than ideal for tasks that rely on rapid and precise object interaction, and renders robots unsuitable for interaction with humans.

In contrast, human manipulative behavior can be complex, precise and rapid- despite noisy sensory information, long feedback delays and muscles with slow, nonlinear dynamics. How this takes place despite these biological constraints is an open question. For slow movements, a feedback strategy could be used to continuously monitor the object’s movement and to adjust the exerted force (Kalaska and Crammond, 1992; Scott, 2004; Scott et al., 2015). If interaction with the object is predictable, then fast movements could be performed using a feedforward strategy to plan the time-varying forces that produce the desired movement (Kawato, 1999). However, there is always some amount of uncertainty when interacting with an object (Shadmehr and Mussa-Ivaldi, 1994; Burdet et al., 2001; Rancourt and Hogan, 2001; Takahashi et al., 2001; Milner and Franklin, 2005).

Modulating the arm’s mechanical impedance by coordinating force and kinematics has been proposed as a strategy for handling uncertain interaction dynamics between the hand and external forces (Bizzi et al., 1982, 1984; Hogan, 1984c; Flash and Hogan, 1985). Mechanical impedance is the force which opposes changes in movement, i.e., position, velocity, acceleration, etc. (Hogan, 1984a, 1985b). The current length and activation of the musculotendon tissue, specified by ongoing neural activity, has intrinsic stiffness and damping that exerts force instantaneously to impede changes in position (Hill, 1950; Joyce and Rack, 1969; Rack and Westbury, 1974; Nichols and Houk, 1976). Additionally, changing the configuration of the arm changes muscle length and the distribution of mass, affecting stiffness, damping, and inertia (Trumbower et al., 2009). A compliant arm and hand that yields predictably upon object interaction may minimize the need for moment-by-moment updates to a control signal. This type of control is likely a contributing factor to the fast and robust movements characteristic of human manipulation.

Arm impedance has been studied in a variety of experimental conditions. In posture-control paradigms, subjects held a manipulandum at an equilibrium position to resist randomly imposed displacements (Mussa-Ivaldi et al., 1985). When instructed to “resist” or “not resist” the displacements, subjects modulated their arm impedance to generate the required force needed to return to the equilibrium position (Lacquaniti et al., 1982). In addition, subjects could modulate arm impedance when instructed to co-activate different groups of antagonist muscles (Gomi and Osu, 1998; Osu and Gomi, 1999). However, it is

unclear how the results of these studies can be extrapolated to real-world movements in which both posture and the equilibrium position change (Gomi and Kawato, 1997; Darainy et al., 2007).

A variation of the posture-control paradigm required subjects to exert an isometric force. Under this paradigm, it was found that subjects adopted an arm impedance that was proportional to force (McIntyre et al., 1996). Again, different values of arm impedance could be achieved by co-activating different groups of antagonist muscles (Gomi and Osu, 1998), but the range of this modulation was constrained by the level of isometric force (Perreault et al., 2002).

In addition to posture maintenance, arm impedance can also be used to constrain the arm spatially along an equilibrium trajectory as it moves toward a target (Bizzi et al., 1982, 1984; Hogan, 1984c; Flash and Hogan, 1985). A path toward the target would consist of a series of equilibrium positions. An arm following this trajectory would be resistant to perturbations away from these positions. If the equilibria were specified sequentially, the arm would be propelled toward the target by the continuous change in impedance. In this way, the control strategy could be modeled as a spring-mass-damper (stiffness-inertia-damping) that moves along the equilibrium trajectory, exerting force to pull the object along behind it. In subsequent studies testing this hypothesis, subjects were instructed to relax their arms as much as possible and, although it was found that arm impedance varied during movement, the equilibrium trajectories were complex (Gomi and Kawato, 1996, 1997). In later studies, subjects moved their arms through unstable force fields with varying amounts of uncertainty (Takahashi et al., 2001; Perreault et al., 2002; Osu et al., 2003; Milner and Franklin, 2005). Again, the subjects were able to modulate the impedance of their arms to complete the movements successfully. The finding that impedance could be steered to compensate for directionally-specific instability (Burdet et al., 2001; Franklin et al., 2007; Kadiallah et al., 2011) was taken as evidence for an explicit impedance controller with an internal model of environmental instability (Franklin et al., 2007).

In general, internal models are invoked to explain how predictions can be made to account for the inherent delays of the motor system. In particular, setting a value of impedance in anticipation of object interaction could be useful in minimizing the effect of these delays.

Evidence that subjects change their impedance during a task to anticipate object contact was found using a task in which subjects caught a falling ball (Lacquaniti et al., 1993). Pseudorandom force pulses were used to measure arm impedance throughout the task and it was found that subjects changed the direction and magnitude of that impedance shortly before the ball made contact with the hand.

In order to further characterize anticipatory changes in impedance, we designed a ballistic release paradigm that required subjects to pull on a handle with different levels of force and move the handle to various target positions. We found that our arm model’s impedance varied with task conditions in a way that could be separated from its linkage to force. In addition, the model’s impedance and force were related to separate components of muscle activity. Not only do these results support the conclusions of previous studies showing that impedance is controlled explicitly, they suggest that it may be specified before an expected object interaction takes place. The preset impedance allows the forces following the perturbation to direct the arm and hand to an intended target.

### 3.3 METHODS

#### 3.3.1 Subjects

Five subjects – four men and one woman between the ages of 20 and 40 – performed a ballistic-release task with their right arm. All subjects provided written and informed consent and the protocol was approved by the University of Pittsburgh’s Institutional Review Board.

#### 3.3.2 Behavioral paradigm and Experimental design

Each subject was instructed to (1) use their right hand to press a start button and then reach to grasp a handle, (2) pull on the handle with enough force to unlock it, and (3) position the handle within a specified target zone (Figure 3.1). To be successful, the subject needed to exert enough force to unlock the handle, then move and hold the handle in the specified target zone for 300 ms. In theory, an efficient strategy would be similar to attaching a spring

to the locked handle and stretching the other end of the spring (the equilibrium position) to the target zone; when the handle released, it would come to rest at the equilibrium position (Feldman, 1966, 1986; Polit and Bizzi, 1979). For the same target zone and equilibrium position, the force exerted on the handle could be increased by increasing the spring’s stiffness (arm impedance). Thus, the equilibrium position and arm impedance could be preset for a given target zone and force threshold, eliminating the need to rely on corrective interventions during the movement. See section 2.1 for more details.

### 3.3.3 Data collection

During the task, we measured three signals: (1) the force exerted on the handle; (2) the 3D position of optical markers placed on the handle and the subject’s hand, lower arm, upper arm, and torso; and (3) the surface electromyographic (EMG) activity of 8 muscle groups. All data were synchronized in time and analyzed at 100 Hz. Data from individual trials were aligned on movement onset. See sections 2.2, 2.3, and 2.4 for more details.

### 3.3.4 Physical dynamical model

To characterize how arm impedance varied with task conditions, we estimated arm impedance by modeling the behavior as an equivalent physical dynamical system with impedance components that exert force proportional to position and its first 3 time derivatives (equation 2.1):

$$F(t) = K(x_0 - x(t)) - D\dot{x}(t) - A\ddot{x}(t) - J\dddot{x}(t),$$

where  $F$  is the force exerted on the handle,  $x_0$  is the equilibrium position,  $x, \dot{x}, \ddot{x}, \dddot{x}$  is position and its first 3 time derivatives, and  $K, D, A, J$  are impedance coefficients. The model maps a motion signal to the force exerted on the handle. The equilibrium position is the reference position for the model and represents the steady state of the system, which can be interpreted as the hold position in this task. The impedance components represent the transition from the initial state to the steady state and make it possible to exert different forces while keeping the equilibrium position (hold position) constant. See section 2.5 for more details.

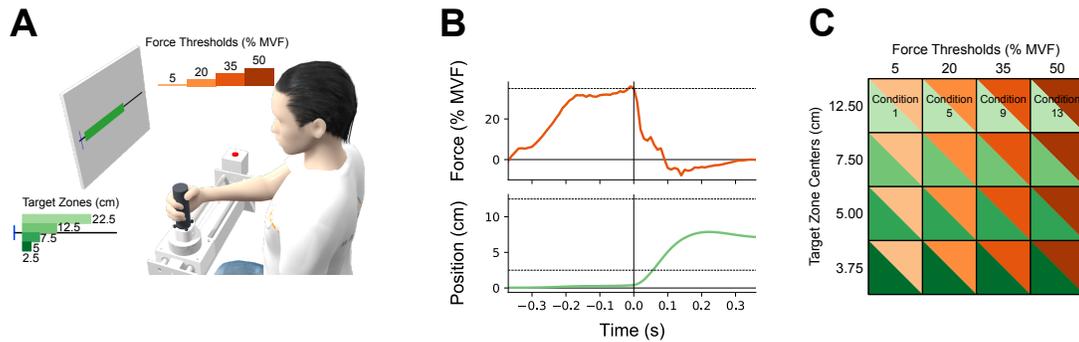


Figure 3.1: **Behavioral paradigm (A)** To initiate a trial, the subject first pressed the start button (to the subject’s right) and then grasped the handle. The subject pulled on the handle while it was locked in place until the force threshold was crossed. The handle was then unlocked to move freely along the track and the subject had to stop and hold the handle within the target zone for 300 ms. **(B)** Time-series of force and position were measured for each trial (a single representative trial is displayed here). **(C)** A single task condition was composed of a force threshold and a target zone.

### 3.3.5 EMG analysis

We were interested in characterizing the EMG pattern that correlated with arm stiffness, which can increase when antagonist muscle groups co-activate. When this is the case, EMG can change without a change in the force exerted on the handle. In our analysis, we distinguished between changes in the EMG pattern that were most correlated with changes in force (potent EMG activity) from those patterns that were less correlated (null EMG activity).

**3.3.5.1 Regressing force against EMG** We regressed force against EMG during the force ramp using trial-averaged force and EMG signals from the 16 task conditions. We first fit the linear model according to equation 3.1

$$F(t) = \beta_0 + W \cdot EMG(t), \quad (3.1)$$

where  $F$  is force exerted along the track,  $EMG$  is muscle activity from 8 muscle groups,  $W$  is the coefficient vector, and  $\beta_0$  is a constant offset. We then used this model to fit force to EMG for each task condition.

**3.3.5.2 Separating potent EMG and null EMG** We isolated the EMG pattern that most correlated with changes in force using singular value decomposition (SVD) (Kaufman et al., 2014) on the  $[1 \times M]$  coefficient vector in equation 3.1,

$$U, \Sigma, V = SVD(W),$$

where  $V$  is an  $[M \times M]$  rotation matrix,  $\Sigma$  is a  $[1 \times M]$  scaling matrix,  $U$  is a  $[1 \times 1]$  matrix, and  $M$  is the number of muscle groups. EMG projected onto the first dimension (row) of  $V$  was most correlated with changes in force. We called this the “potent EMG” activity:

$$EMG_{potent}(t) = V_{potent} \cdot EMG(t),$$

where  $V_{potent}$  is a  $[1 \times M]$  vector that is the first row of  $V$ . If the relation between EMG and force was perfectly linear, then EMG projected onto the remaining 7 dimensions, i.e. rows 2 through  $M$  of  $V$ , would not be correlated with changes in force. We summarized this EMG

activity by performing principal components analysis on the non-potent EMG and found the single dimension that captured the most variance. EMG was projected onto this dimension and the result was considered “null EMG”:

$$EMG_{null}(t) = P_{pc1} \cdot V_{null} \cdot EMG(t),$$

where  $V_{null}$  is a  $[M-1 \times M]$  matrix that is rows 2 through  $M$  of  $V$  and  $P_{pc1}$  is a  $[1 \times M-1]$  eigenvector of  $V_{null} \cdot EMG(t)$ .

### 3.3.6 Testing the effect of arrest position on the regression of stiffness on force threshold

Although it is known that stiffness depends on force (McIntyre et al., 1996), the impedance control hypothesis predicts that stiffness would also depend on position. We tested this hypothesis using the partial F-test, which determines if additional parameters improve the explanatory power of a regression model. It employs two nested models, a full model and a restricted model. The restricted model consists of a subset of the parameters from the full model.

The two models’ residual sum of squares are compared according to

$$F = \frac{RSS_r - RSS_f}{RSS_f} \frac{n - P_f}{P_f - P_r},$$

where  $F$  is the F-statistic,  $RSS_f$  and  $RSS_r$  are the residual sum of squares of the full and restricted model,  $n$  is the number of observations, and  $P_f$  and  $P_r$  are the number of parameters in the full and restricted model. A higher F-statistic indicates more explanatory power in the full model. The null hypothesis assumes a value of 0, indicating that the full model does not have more explanatory power than the restricted model. Statistical significance is tested using the F-distribution with  $(P_f - P_r, n)$  degrees of freedom.

The partial F-test was used to test the hypothesis that adding the arrest position as a parameter would improve the regression of stiffness on force threshold. Equation 3.2 describes the full model and equation 3.3 describes the restricted model,

$$K = \beta_0 + \beta_1 Threshold + \beta_2 AP + \beta_3 (Threshold)(AP) \tag{3.2}$$

$$K = \beta_0 + \beta_1 Threshold \tag{3.3}$$

where, across all task conditions,  $K$  is stiffness,  $Threshold$  is the force threshold,  $AP$  is the arrest position (defined as the position after movement onset when velocity falls below 10% of the maximum), and  $\beta$  are constants.

### 3.4 RESULTS

Five subjects performed a ballistic-release task that required them to pull on a handle with different levels of force and then move the handle to different target zones. Successful trials were achieved by exerting enough force to unlock the handle while still controlling the subsequent movement. Below, we highlight the results from subject 1 and include the remaining subjects in the supplementary material. Summary statistics across all subjects are also presented.

#### 3.4.1 Force and motion varied with the four force thresholds and four target zones

The subject pulled on the handle with four levels of force and moved the handle to four target zones (Figure 3.2). Movement onset began at time 0 and, at this time, force for the same threshold was similar across targets (Figure 3.2A). At or near the end of the movement, 500 ms after movement onset, position was separated across targets for the same threshold (Figure 3.2B).

For each signal and task condition, we found the median and 95% confidence interval of the trial-averaged data. We used a bootstrap to re-calculate the trial average 1,000 times. Each calculation used a random sample of the trials, with replacement.

For this subject, force increased rapidly and began to plateau as it approached the threshold. This pattern was more pronounced for higher thresholds and closer targets. The plateau suggests that the subject could approximate the force that would unlock the handle

but was unable to accurately predict the exact timing of when the handle would be unlocked. After movement onset, the force decreased rapidly, falling below 0 to slow the handle and then approaching 0 as the movement of the handle reached a steady position. This trend was consistent across all subjects (Supplementary Figure 3.9). The MVF for this subject was 160 N.

For this subject, the movements were usually arrested around the center of the target, after some initial overshoot. The overshoot remained in the target zone, except for the closest target and the highest threshold combination, where it often extended beyond the target zone, reflecting the extreme difficulty of this task condition. The overshoot was more pronounced for closer targets and higher thresholds which corresponds to a physical dynamical system having various values of impedance. Before movement onset, the position of the handle was fixed at 0, with some handle bending (on the order of millimeters) for high thresholds. Because all trials were included in the analysis, there was substantial across-subject variability in the positions of the handle at the end of the trial (Supplementary Figure 3.9). For a given subject and task condition, the across-trial variability of the positions was small, suggesting that the subject had selected a consistent movement strategy for that condition.

The timing of the maximum velocity was similar across all task conditions (Figure 3.2C). However, the magnitude of the maximum velocity depended on both the target zone and the force threshold, increasing for targets farther from the lock position and for higher thresholds.

As expected, the magnitude of the maximum acceleration depended on the force threshold (Figure 3.2D). The initial acceleration values are inaccurate due to the numerical differentiation of position. For a pre-loaded release, the initial acceleration should be nearly a step function and begin at time 0. The initial deceleration was similar across target zones with the same force threshold.

Task conditions varied from extremely easy to nearly impossible (Table 1, statistics across subjects). The first four task conditions began with the lowest threshold and farthest target (top left element in Table 1 and proceeding down). Success rates for these initial conditions were lower than expected because of some occasional confusion about the exact requirements of the task. However, after a few trials, the subjects moved quickly and smoothly (Figure

3.2B and Supplementary Figure 3.9), although the 95% confidence interval was slightly wider for the handle position during these first four task conditions (different targets with the lowest threshold) compared with the remaining the task conditions.

### 3.4.2 Arm model impedance varied with force and motion

We modeled the arm as a physical dynamical system consisting of an equilibrium position and four impedance components. For each task condition, we averaged the force exerted on the handle and the handle’s movement across trials. We then used least-squares optimization to find the model parameters that minimized the difference between the actual and predicted force over 200 ms beginning at movement onset (see Methods, equation 2.1).

The equilibrium position (EP) of the physical model represents the position of the handle for which zero force would be exerted. In the context of this task, the EP can be considered the position where the handle would come to rest and could reflect the intended hold position (Figure 3.3 and Supplementary Figure 3.10). The following describes the results from Subject 1. We observed that the distance to the EP increased for arrest positions farther from the handles lock position (Figure 3.3). However, the distance to the EP was consistently shorter than the distance to the arrest position. The difference between the arrest position and the EP is suggestive of an under-damped system. Although individual subjects sometimes arrested their movements at locations that were outside the specified target zones, the match between EP and the arrest position followed the general trend displayed by this subject (see also Supplementary Figure 3.10).

The task was designed so that the same target was specified for different force thresholds. If the EP depended on the target, then the different force thresholds could be crossed by selecting different arm impedances. This would be an efficient control strategy for the ballistic-release task. We found that the model’s stiffness remained consistent across trials for each task condition (Supplementary Figure 3.12) and that stiffness increased with force threshold for a given target zone (Figure 3.4 and Supplementary Figure 3.12), consistent with

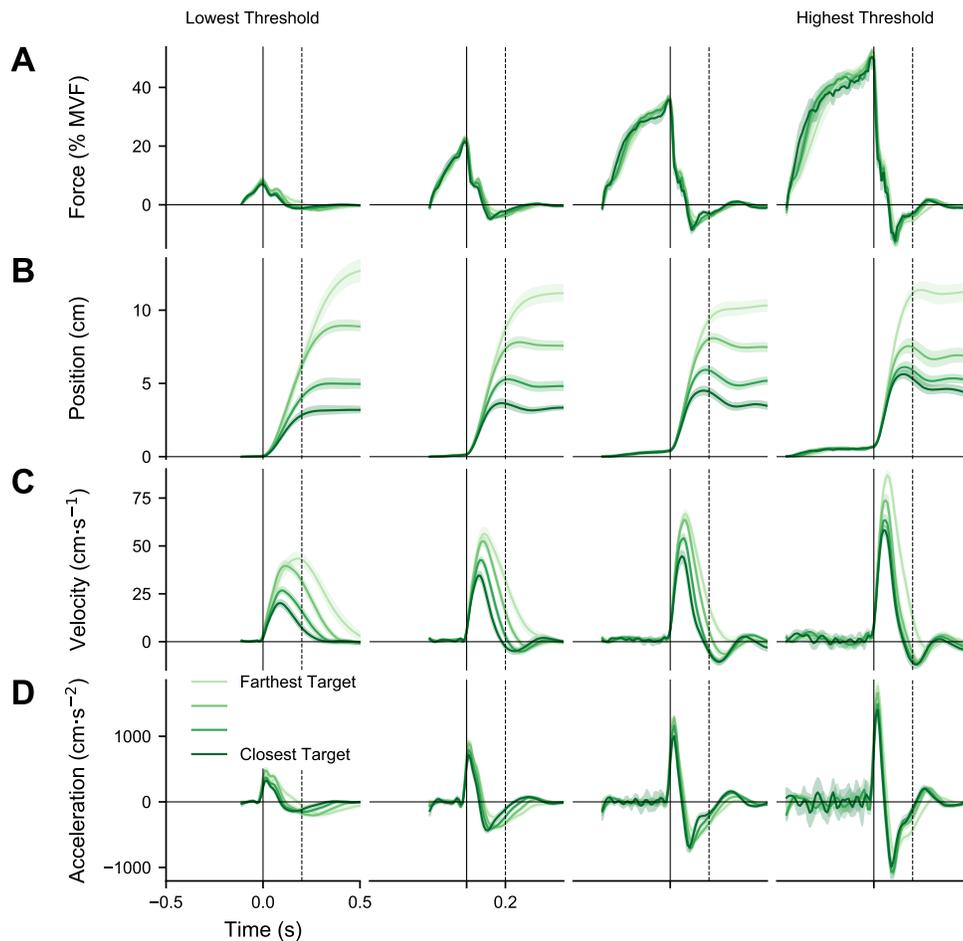


Figure 3.2: **Behavioral results.** The data are trial-averaged from one subject. Shading represents the median and 95% confidence interval. Each plot depicts four target zones with the same force threshold. Movement onset began at time 0 (solid vertical line). **(A)** Force varied across thresholds. For a given threshold, force was similar across targets. **(B)** Position varied across targets. Subjects tended to overshoot the target zone as the threshold increased and the target moved closer. **(C)** Velocity varied with target and threshold. The maximum velocity was related to both target and threshold. **(D)** Acceleration varied with threshold. For a given threshold, initial deceleration was similar across targets. Results for the remaining subjects are shown in Supplementary Figure 3.9.

Table 5: **Success rate for each task condition.** On successful trials, the subject moved the handle into the target zone and remained there for 300 ms. Exiting the target zone before 300 ms caused the trial to fail. The difficulty of the task generally increased as the force threshold increased and as the distance to the target zone decreased. The order of the task conditions started in the upper left element and proceeded down the column. Entries are mean  $\pm$  standard error across subjects.

	Lowest threshold		Highest threshold	
Farthest target	91.4% $\pm$ 2.4	100% $\pm$ 0	100% $\pm$ 0	100% $\pm$ 0
	97% $\pm$ 1.8	99% $\pm$ 0.9	99% $\pm$ 0.9	94% $\pm$ 3.6
	82.4% $\pm$ 7.6	95% $\pm$ 2.0	77.3% $\pm$ 10.8	54% $\pm$ 16.5
Closest target	71% $\pm$ 15.4	68.3% $\pm$ 7.8	39% $\pm$ 14.4	5% $\pm$ 3.5

previous results (McIntyre et al., 1996). Although this relation was approximately linear for a given target zone, across all task conditions there was distinct structure in the effect of target zone on the relation between stiffness and force threshold.

For each subject, we quantified the target-dependent structure by comparing the linear fit of two nested regression models using a partial F-test (Figure 3.5). The full model regressed stiffness on both threshold and arrest position (grey lines in Figure 3.4). The restricted model regressed stiffness on threshold alone (black lines in Figure 3.4). Adding the arrest position to the regression increased the predictive power as evidenced by the F-statistics p-value, which was below 0.05 for every subject: 0.004, 0.002, 0.008, 0.03, 0.0006; F-distribution with (2, 12) degrees of freedom.

Although the model’s first- and second-order impedance coefficients generally increased with force threshold, there was not a consistent relation between these impedance components and target zone (Figure 3.4). When stiffness increases more than damping, the physical dynamical system can become under-damped and tend to oscillate, which may explain the target overshoot and movement oscillations seen in Figure 3.2.

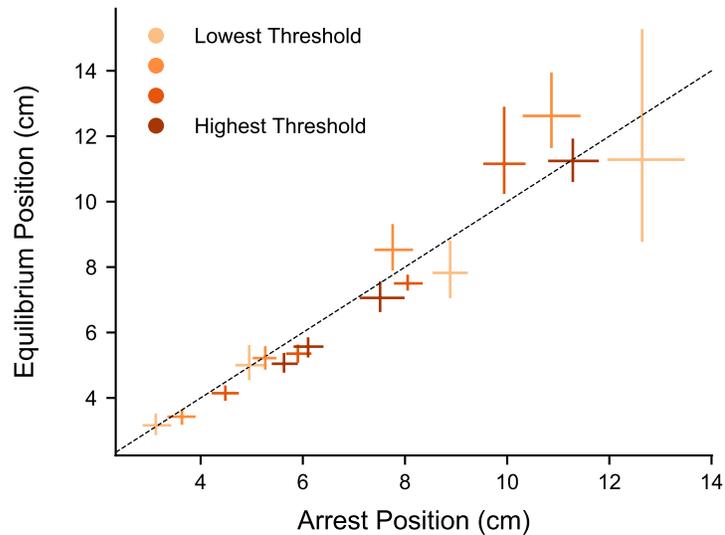


Figure 3.3: **The equilibrium position (EP) co-varied with the arrest position.** The data are from subject 1. The distance to the EP increased as the distance to the arrest position increased. However, the distance to the EP was less than the distance to the arrest position. This pattern was consistent across all subjects (Supplementary Figure 3.10) and is suggestive of an under-damped system. Error bars indicate median and 95% confidence interval. Dashed line is the unity line. The model was fit using 200 ms of data beginning at movement onset.

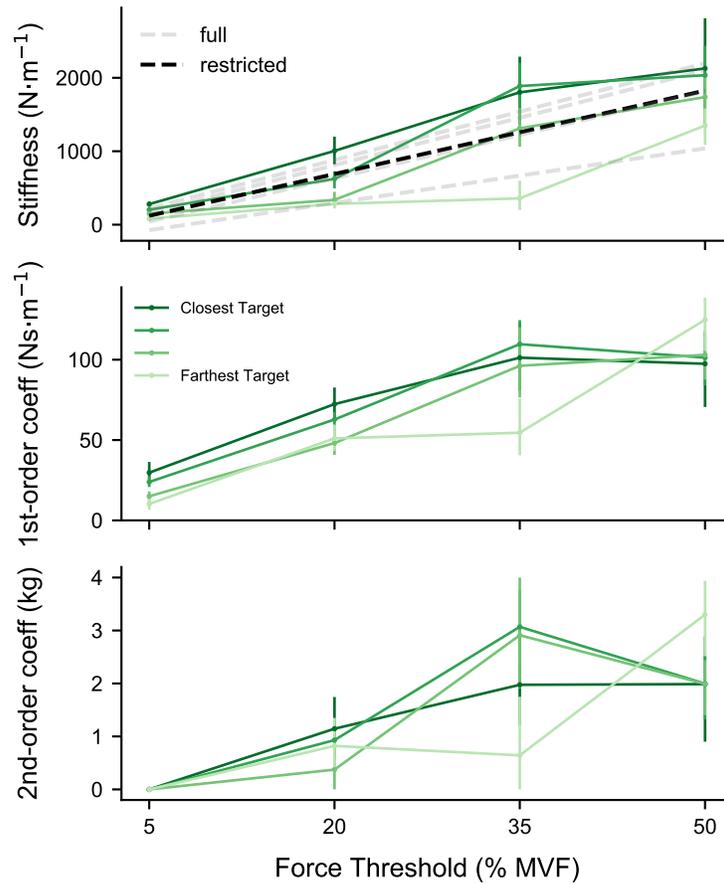


Figure 3.4: **Model stiffness co-varied with both force threshold and arrest position.** The data are from subject 1. Although the model’s stiffness was proportional to threshold for a given target, across all task conditions the linear fit that included arrest position was better able to explain stiffness. The dashed grey lines represent the linear fit of the full regression of stiffness on force threshold and arrest position. The dashed black line represents the linear fit of the restricted regression of stiffness on force threshold alone. The model’s first- and second-order impedance coefficients increased with force threshold. However, there was not a strong relation between these impedance components and the arrest position. Error bars indicate median and 95% confidence interval. The physical dynamical model was fit using 200 ms of data beginning at movement onset.

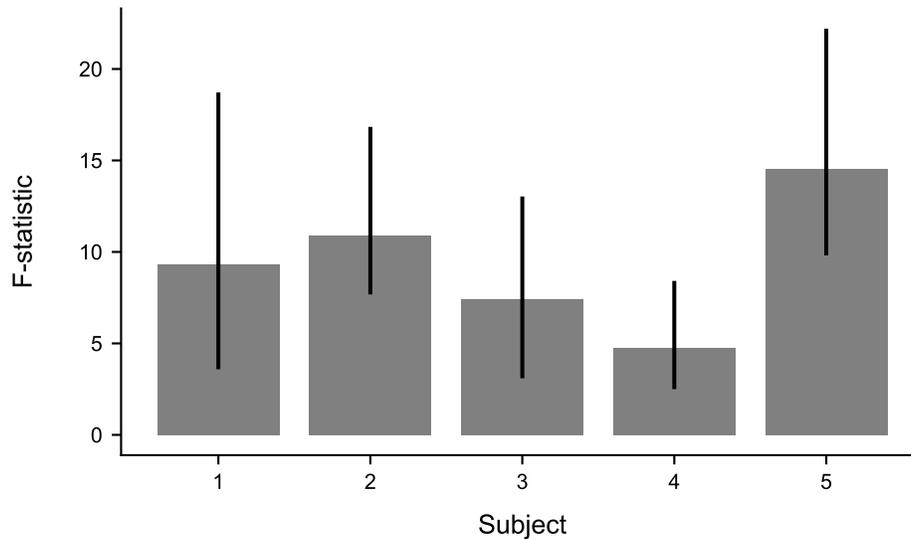


Figure 3.5: **Model stiffness variability is best explained by both force and movement information.** F-statistic from a partial F-test comparing the regression of model stiffness on force threshold with and without additional arrest position information. The additional arrest position information improved the stiffness regression for all subjects, with subject 1 and 5 showing the largest effect. Error bars represent median and 95% confidence interval.

### 3.4.3 Arm posture before movement did not vary across task conditions

Although we suspected that subjects might vary their arm posture for different task conditions (Trumbower et al., 2009), we found the arm’s configuration to be remarkably consistent. For subject 1, the deviation of the joint centers during the 200 ms before movement onset across all task conditions were as follows: torso 0.98 cm [0.21, 1.99], shoulder 1.23 cm [0.33, 2.39], elbow 1.1 cm [0.27, 2.57], wrist 0.75 cm [0.29, 1.49]; median [95% confidence interval]. Similar values were found for the other subjects (not reported).

Figure 3.1A is representative of the arm posture before movement onset. Specifically, the joint center positions for subject 1 during the 200 ms before movement onset across all task conditions were as follows: jugular notch (11.34 cm, -31.62 cm, -7.3 cm), right shoulder (32.72 cm, -33.23 cm, -6.86 cm), right elbow (27.37 cm, -13.13 cm, -20.27 cm), and right wrist (4.78 cm, -0.11 cm, -10.01 cm). The origin was the handle’s lock position, the +x-direction pointed to the subject’s right, +y-direction pointed ahead of the subject, and +z-direction pointed upward. Similar values were found for the other subjects (not reported).

### 3.4.4 EMG patterns co-varied with force threshold and with target zone

We recorded bipolar surface EMG from 8 muscle groups in the arm and shoulder (Figure 3.6) and found that muscle activity gradually ramped up before movement onset at time 0 (solid vertical line in Figure 3.6), particularly for task conditions with the highest force threshold. Additionally, EMG activity of most muscle groups was elevated before and after movement onset, suggesting a level of co-activation of antagonist muscle groups. For each muscle and task condition, we found the median and 95% confidence interval of the trial-averaged data. We used a bootstrap to re-calculate the trial average 1000 times. Each calculation used a random sample of the trials, with replacement.

The elbow flexor, posterior deltoid, and rotator cuff were particularly active before movement onset, i.e., while the subject was increasing the force they exerted on the handle. The activity of these muscle groups, and their biomechanical actions, suggest that they might be acting to exert force to cross the force threshold and accelerate the handle along the track.

The elbow extensor, anterior deltoid, and pectoralis activity showed a distinct burst about 70 ms after the handle started to move (dashed vertical line in Figure 3.6), consistent with a stretch reflex and the functional use of these muscles to decelerate the handle after it started to move.

The wrist flexor and extensor were the muscle groups most consistently modulated across task conditions. One possibility for this modulation is that the co-activation of these antagonist muscle groups would stiffen the wrist joint and improve the force transfer from the arm to the handle. Another possibility is that this activity may reflect how tightly the subject squeezed the handle, and this may also stiffen the linkage. With our current data, we cannot distinguish between these possibilities.

There are oscillations in the EMG before and after the handle was released. These tend to increase with larger force thresholds (e.g., posterior deltoid). These oscillations and those in the force and acceleration traces are evident in Figures 3.2 A and D and appear to be phase-locked to movement onset. The acceleration pulse occurring at movement onset is in phase with the preceding oscillations, suggesting that they may play a role in overcoming the force threshold.

Comparing a muscle group's EMG across the force thresholds shows that muscle activity increases with force. In addition, there is a tendency for greater muscle activity for near targets (especially evident at high force thresholds). For a given threshold, the exerted force is similar for different target zones (Figure 3.2A), despite different EMG patterns, suggesting that the subject was modulating the co-activation of antagonist muscle groups.

### **3.4.5 EMG decomposed into potent and null dimensions**

We regressed force on EMG during the force ramp and found that a linear model could explain much of the variance (RMSE =  $11.68 \text{ N} \pm 1.44$ , mean  $\pm$  sem across subjects, Supplementary Figure 3.13). Because there were more dimensions of EMG (8 muscle groups) than dimensions of force (only one along the movement direction), the mapping from EMG to force was redundant. If the mapping from EMG to force was perfectly linear, there would

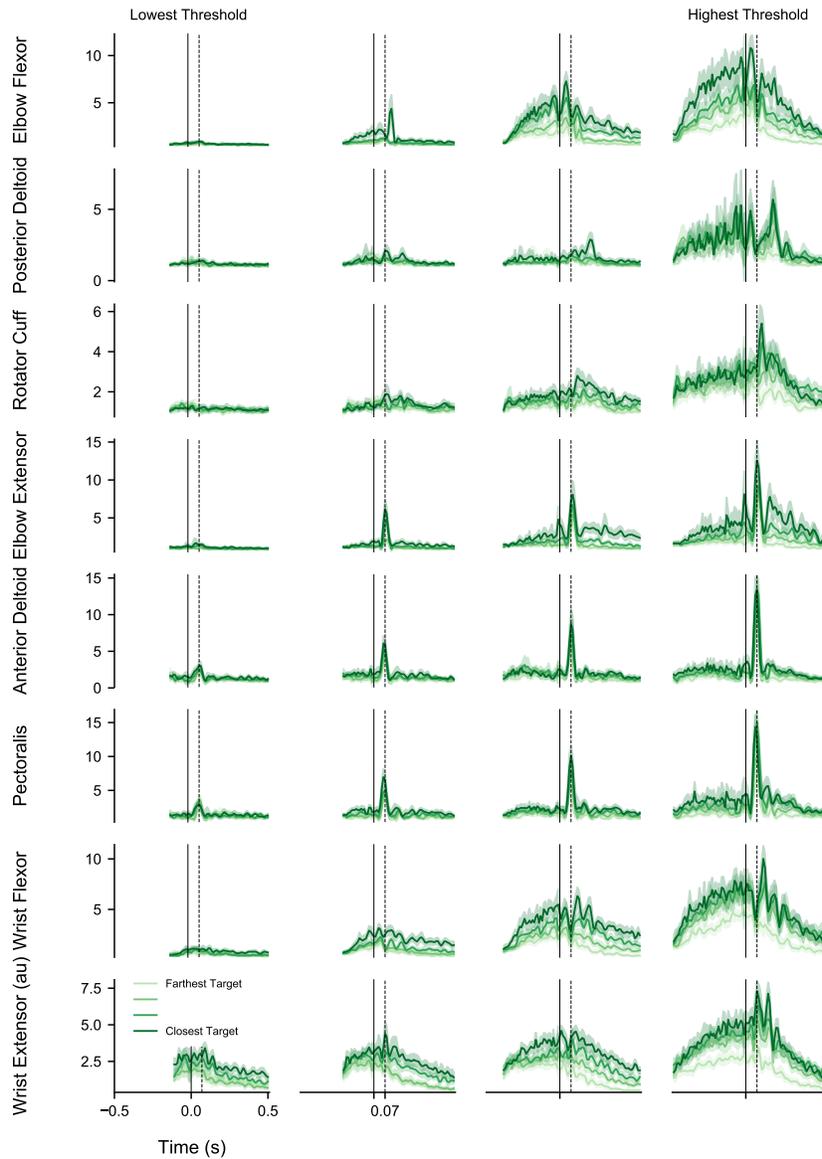


Figure 3.6: **Time-varying EMG across task conditions.** The data are from subject 1. Each plot depicts four target zones with the same force threshold. Movement onset was at time 0 (solid vertical line). Some muscles exhibited a strong burst about 70 ms after movement onset (dashed vertical line). The wrist and elbow muscle groups appeared to co-activate, particularly for task conditions that included the highest force threshold. Shading represents the median and 95% confidence interval.

be one direction (i.e., muscle combination) in the multiple-dimensional EMG space that correlated with changes in force. EMG orthogonal to this direction, i.e. in the remaining dimensions of EMG space, would not correlate with force changes.

Using the regression coefficient matrix (equation 3.1), we found a rotation in the multi-dimensional EMG space that maximized EMG-force correlation in a single dimension. We called EMG projected onto this dimension “potent EMG.” EMG projected onto the remaining 7 dimensions, i.e., dimensions that theoretically did not correlate with changes in force, were summarized using principal component analysis and called “null EMG.”

We found that potent EMG accounted for  $21.28\% \pm 4.35$  of the total variance (mean  $\pm$  sem across subjects). The remaining non-potent variance occupied the dimensions that did not correlate as well with changes in force. Null EMG accounted for  $77.93\% \pm 2.44$  of the non-potent variance and  $61.71\% \pm 4.41$  of the total variance (mean  $\pm$  sem across subjects).

We designed the null-space analysis under the assumption that balanced changes in the EMG of antagonist muscles produce changes in EMG without concomitant changes in net force. Increases in this type of balanced muscle activation should correspond to changes in stiffness. Although we found that force was generally related to the arm model’s stiffness (Figure 3.4), we hypothesized that the model’s stiffness would likely be better related to null EMG than to potent EMG. We tested this by regressing the time-averaged potent and null EMG (Supplementary Figure 3.14) against the model’s stiffness and compared the slopes, re-sampling the data 1000 times, with replacement. The median values, displayed in figure 3.7, show that increases in stiffness corresponded to greater increases in null EMG compared to potent EMG for 4 out of the 5 subjects (EMG offsets were matched to highlight the slope). This was further confirmed by a histogram of the difference between the null EMG slope and the potent EMG slope (Figure 3.8).

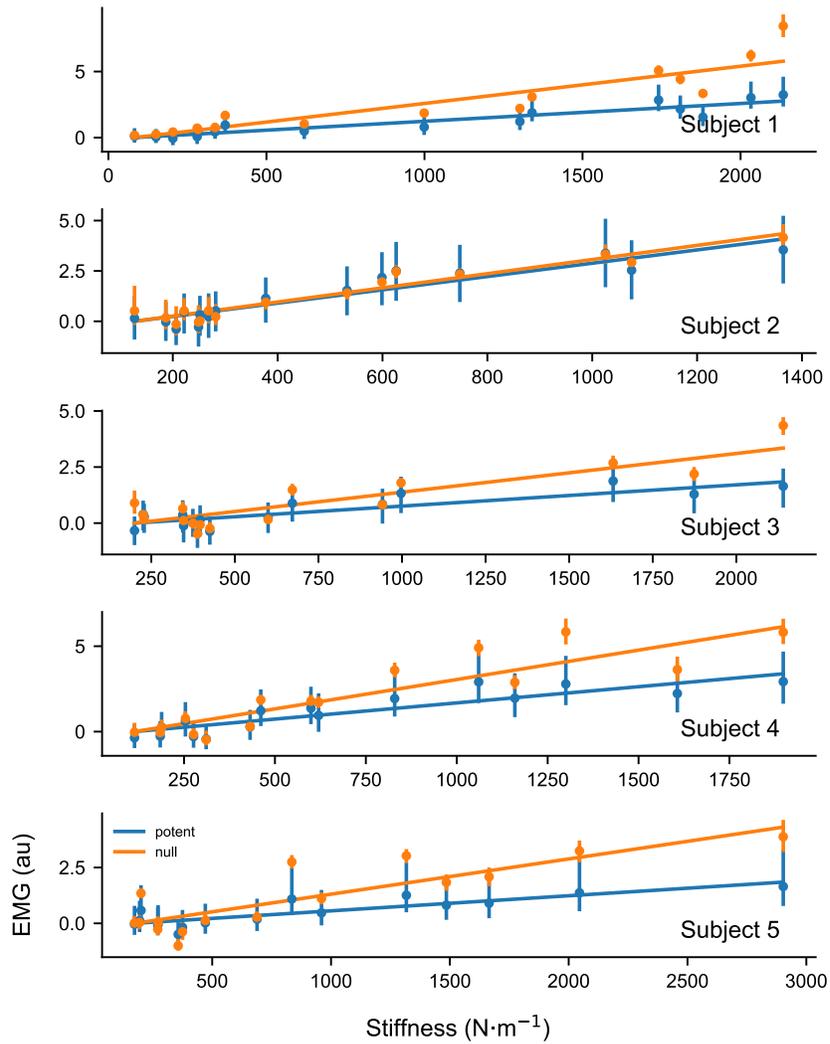


Figure 3.7: **Model stiffness has a greater effect on null EMG compared to potent EMG.** Each plot depicts the time-averaged EMG for all sixteen conditions. Error bars are the median and 95% confidence interval. EMG offsets were matched to highlight the slope. Subject 1 shows the most distinct difference between null and potent EMG, while subject 2 shows no difference. The potent EMG was the direction in EMG space that was most correlated with force. The null EMG was the orthogonal direction that captured the most variance in the non-potent dimensions.

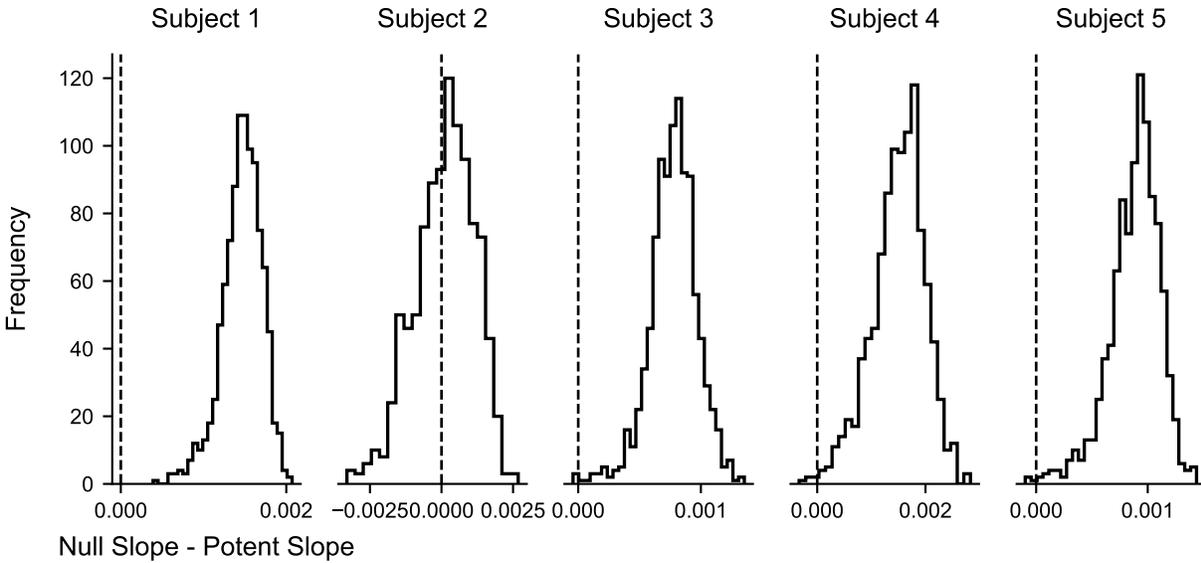


Figure 3.8: **Model stiffness has a larger effect on null EMG compared to potent EMG.** The slope relating null EMG to stiffness was greater than the slope relating potent EMG to stiffness. The data were resampled 1000 times, with replacement. For each resample, we calculated the difference in slopes relating potent and null EMG to stiffness. Vertical dashed lines represent no difference in the slopes, indicating no difference in the effect of stiffness on potent and null EMG. Positive values indicate that the null EMG slope was greater than the potent EMG slope.

## 3.5 DISCUSSION

Successfully interacting with an object relies on the ability to coordinate the movement of, and the force exerted on, an object. Modulating arm impedance could be an efficient strategy for this type of control. We have designed a ballistic-release paradigm that encourages this strategy while making it possible to dissociate the relations between arm impedance, force, and motion. We modeled the arm as a physical dynamical system and, of the model's impedance components studied, stiffness was the only one that varied consistently across both force threshold and target zone. Other studies have shown that stiffness co-varies with force (McIntyre et al., 1996; Gomi and Osu, 1998; Perreault et al., 2002). However, in our task, the linear force-stiffness relation had structure that was target-dependent. Model stiffness was related to the target specified before the movement began and was larger for targets that required shorter movements. Our method for relating EMG to stiffness found that separate components of muscle activity varied with force and with stiffness. In designing this study, our aim was to emphasize the use of impedance control as a way of completing a task. These results, from the ballistic-release paradigm, expand upon studies of arm impedance that investigated the isolated effect of force (McIntyre et al., 1996; Gomi and Osu, 1998; Perreault et al., 2002) or motion (Gomi and Kawato, 1997; Burdet et al., 2001; Darainy et al., 2007; Franklin et al., 2007; Piovesan et al., 2013) and extend the concept that impedance can be specified predictively for object interaction (Lacquaniti et al., 1993; Damm and McIntyre, 2008).

### 3.5.1 Arm impedance estimation

We found model stiffness values that were higher than those reported previously from movements with slower arm speeds and smaller forces (Gomi and Kawato, 1997; Darainy et al., 2007). Higher forces and arm speeds are both known to increase stiffness (Latash and Gottlieb, 1991; McIntyre et al., 1996). In a study with similar arm speeds, values of stiffness are similar to the values we report (Piovesan et al., 2013).

Typically, arm impedance is estimated by perturbing the arm in a way that is external to the task, either by applying a force pulse (Gomi and Kawato, 1996, 1997; Piovesan et al., 2013) or a position displacement (Mussa-Ivaldi et al., 1985; Burdet et al., 2001; Darainy et al., 2007; Franklin et al., 2007). In these studies, it is assumed that subjects do not intervene during the perturbation. Arm impedance can be estimated by measuring the force exerted on the object when the arm’s movement is perturbed away from its equilibrium position. However, this method relies on accurately estimating equilibrium positions. In posture maintenance paradigms, position, velocity, and acceleration are zero at this position (Mussa-Ivaldi et al., 1985; Gomi and Osu, 1998; Perreault et al., 2002). It is more difficult to estimate impedance during movement (Gomi and Kawato, 1997; Darainy et al., 2007). Our task was designed so that the force perturbation, unlocking the handle, was a specific component of the task. Instead of using movement perturbations interspersed throughout the task, we modeled the arm as a physical dynamical system (Viviani and Terzuolo, 1973; Gomi and Kawato, 1997; Burdet et al., 2001; Darainy et al., 2007) based on the force exerted on the handle and the initial displacement of the hand as it was released. Impedance, in our model, was calculated with the assumption that the equilibrium position (the target) was constant, and that at this position, velocity and acceleration were both zero (Polit and Bizzi, 1979). This assumption was also used in a similar ballistic-release paradigm (Viviani and Terzuolo, 1973).

For a given force threshold, the force exerted toward the subject was correlated with model stiffness (Supplementary Figure 3.15). The directional nature of stiffness and force makes it highly unlikely that the off-axis force was causal to the modulation of on-axis stiffness. Instead, a simple correlation can be described (Supplementary Figure 3.16). Two imaginary muscles could exert force in opposite directions along the track, but in the same direction toward the subject. If the projection of the two muscle forces along the track canceled out, then the on-axis force would be zero and the force exerted toward the subject would be non-zero. Proportional modulation of the muscle forces would increase both the stiffness along the track and the force exerted toward the subject. For this reason, the

definition of potent and null EMG also considered only force along the track. Considering force in all three dimensions would have prevented certain behavioral strategies and was not part of the instructions to the subjects.

### 3.5.2 Muscle-related stiffness during the ballistic movement

Reaching toward a target is composed of an initial rapid stereotypic movement followed by a homing phase composed of multiple small submovements (Woodworth, 1899; Meyer et al., 1988). The initial phase is considered too rapid for concurrent corrections (Elliott et al., 2010) and the arm effectively behaves as a spring-mass-damper system (Viviani and Terzuolo, 1973; Hogan, 1985a). Upon release, our subjects moved in a way that was similar to this initial reaching phase in that the task constraints encouraged a behavior that was ballistic. This is consistent with a control strategy characterized by preset stiffness and damping.

The first 200 ms following movement onset took place in the absence of corrective interventions to the movement. Since those movements depended on threshold and target, it is likely that the combination of muscle activations were pre-adjusted to reflect the behavioral conditions of the task, as muscle activation contributes to both the force exerted on the handle and to arm impedance (Hogan, 1985a). During this time window, it is likely that feedback-related changes in EMG take place, as evidenced by the EMG response we saw at 70 ms (Figure 3.6). These responses may be composed of spinal and cortical reflexes and could contribute to both stiffness and damping at a time delay. These responses are known to change with task requirements (Kurtzer et al., 2009; Dimitriou et al., 2013; Pruszynski et al., 2014) and could be preset by the nervous system to control the arm during the movement. However, further analysis will be needed to detail the differences between force responses that occur at a time delay because of sensory feed-back and force responses that occur at no time delay because of arm impedance.

### 3.5.3 Impedance control

Results from a similar paradigm suggest that the equilibrium position is set before the force begins to increase (Elliott et al., 1999). In that study, the handle was unlocked in random catch trials before or after the force threshold was crossed. Since force increased continuously before the handle was released, a pure force-dependent strategy would result in shorter displacements for earlier releases. However, the handle’s displacement was not related to the time of release in the catch trials, showing that a simple predictive force strategy was probably not used in the task. Furthermore, force at high magnitudes is susceptible to signal-dependent noise (Harris and Wolpert, 1998; Franklin et al., 2004), making it difficult, as in our task, to predict when force would cross the threshold and unlock the handle, again arguing against the idea of a scheme relying only on predictive force control.

Although our model results suggest that arm impedance is set before the movement takes place, in theory, subjects could perform the ballistic-release paradigm without changing impedance. If a subject could predict when the handle would be unlocked, it would not be necessary to change impedance for different targets. Instead, a set of muscles could be activated to generate the force needed to unlock the handle, followed by the activation of a different set that would generate the precise force needed to decelerate the handle to stop in the target zone. In contrast to co-activation, this type of control would be energetically efficient (Franklin et al., 2004) and could be implemented using implicit and/or explicit knowledge of the physical plant’s mechanics (i.e., mapping activation to force) along with the object’s properties (i.e. the force needed to unlock the handle). The subject’s internal model would encompass this knowledge and could be used to precisely control the transition from isometric force to movement control. Although this would be energetically efficient, precise timing would be required. The additional details needed for this scheme would increase information loading in the system and could result in slower movement.

Our results suggest that subjects chose to control both force and movement together via impedance control, an idea consistent with other studies which found impedance control to be preferred when the relation between force and movement is uncertain (Thoroughman and Shadmehr, 1999; Takahashi et al., 2001; Franklin et al., 2003). These concepts also fit under

the umbrella of the equilibrium point hypothesis (Feldman, 1966, 1986) and its extensions (Bizzi et al., 1984; Hogan, 1985a; Flash, 1987). In this framework, the force exerted on the handle in our task is controlled only indirectly. It depends on the desired movement, the actual movement imposed by the object on the arm, and the preset impedance.

### 3.5.4 Conclusion

The ballistic-release paradigm used in this task encouraged subjects to adopt a strategy in which they simultaneously activated their muscles, creating a virtual spring that arrested a fast arm movement in a specified target. The physical model of the arm suggested that subjects adjusted impedance to achieve the desired displacement for each combination of force threshold and target zone. By modeling the arm’s impedance in the short interval following release, we found that model stiffness changed in a way that anticipated the displacement needed to reach the target. We assessed the relation between model impedance and muscle activation, finding EMG patterns that were less correlated with changes in force. This “null” component was, instead, correlated with stiffness, suggesting that subjects used their muscles to modulate arm impedance without changing force. The ability to separate changes of force, position, and stiffness, and the manner in which these variables are associated with different components of muscle activation, suggests that anticipatory changes in impedance may be a cardinal feature of manipulative control. Because this paradigm demonstrates this aspect of control, it will be useful for studying the relation between cortical neural activity and impedance (Humphrey and Reed, 1983). Future enhancement of this paradigm to include multiple directions (Darainy et al., 2007) and time-varying estimates of impedance (Lacquaniti et al., 1993; Piovesan et al., 2013) will make it possible to generalize the control of object interaction to a wider range of behaviors.

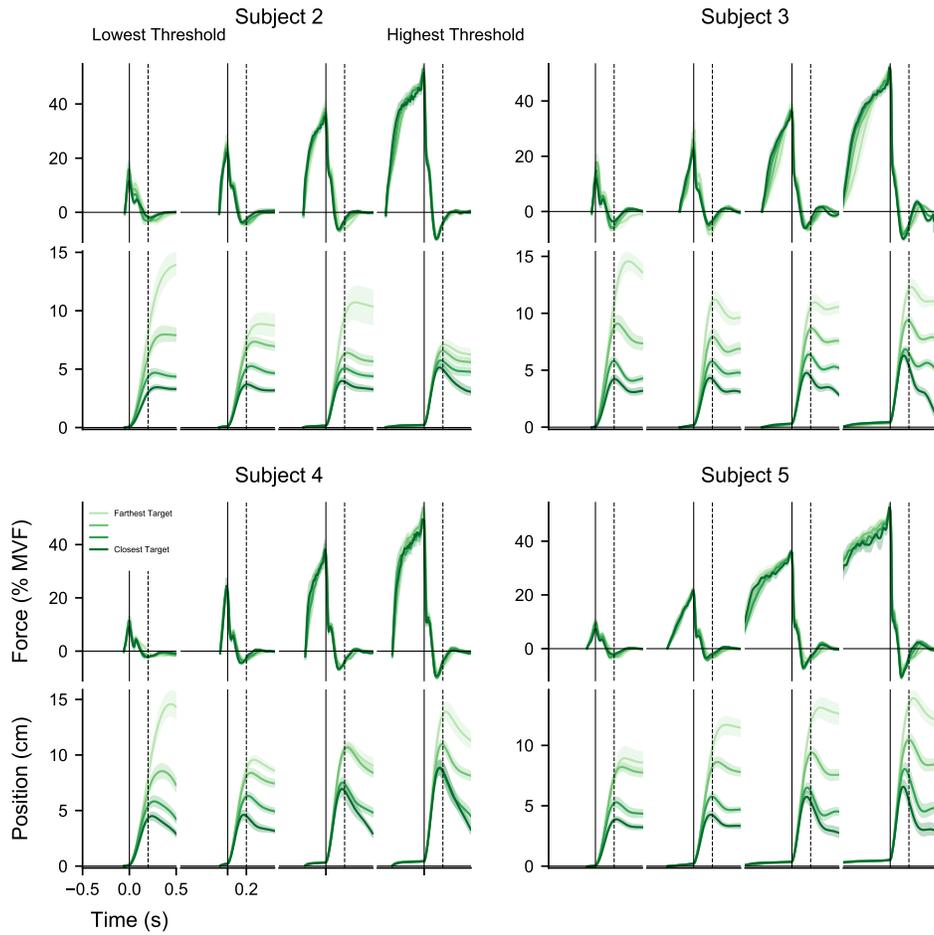


Figure 3.9: **Behavioral data for all subjects.** Movement varied considerably across subjects. However, across-trial variability was much less. For some conditions, the subjects returned the handle to the start position without holding the handle in place (see Subject 4, highest threshold). Additionally, Subject 2 adopted the strategy for the highest threshold of arresting the handle in the same position, regardless of the target. The data are trial-averaged and shading represents the median and 95% confidence interval. Each plot depicts four target zones with the same force threshold. Movement onset began at time 0 (solid vertical line). The maximum voluntary force (MVF) for Subject 2-5 was 100, 160, 200, and 200 N.

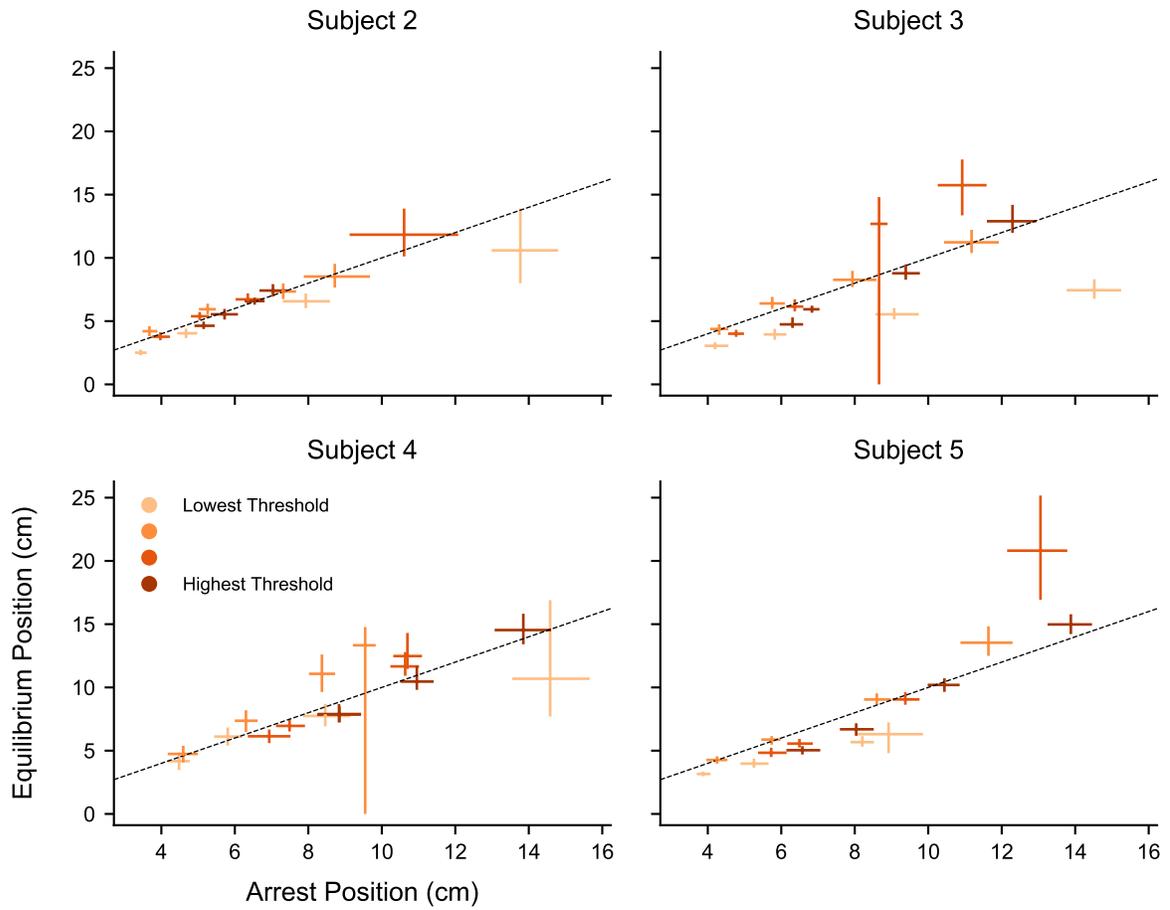


Figure 3.10: **Equilibrium position for all subjects.** The distance to the equilibrium position (EP) was similar to the distance to the arrest position, suggesting that the physical dynamical model could describe important aspects of the behavior. The distance to the EP was generally a little less than the distance to the arrest position, especially for Subject 5, which could be explained by an underdamped physical system and result in the position overshoot observed in Supplementary Figure 1. Markers are the median and 95% confidence interval.

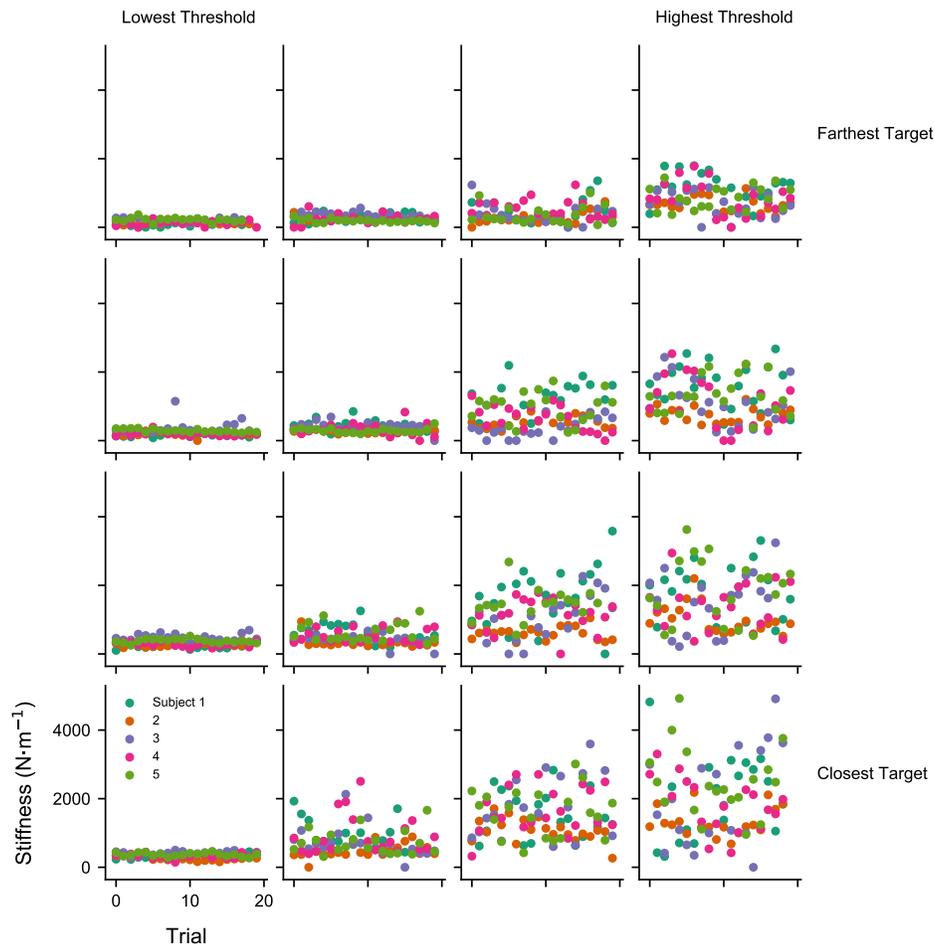


Figure 3.11: **Stiffness across trials for all subjects and task conditions.** Subjects did not appear to modulate their stiffness consistently across trials for a given condition. However, variability within a given condition did increase with stiffness, suggesting signal-dependent noise that could be linked to similar signal-dependent noise in muscle force.

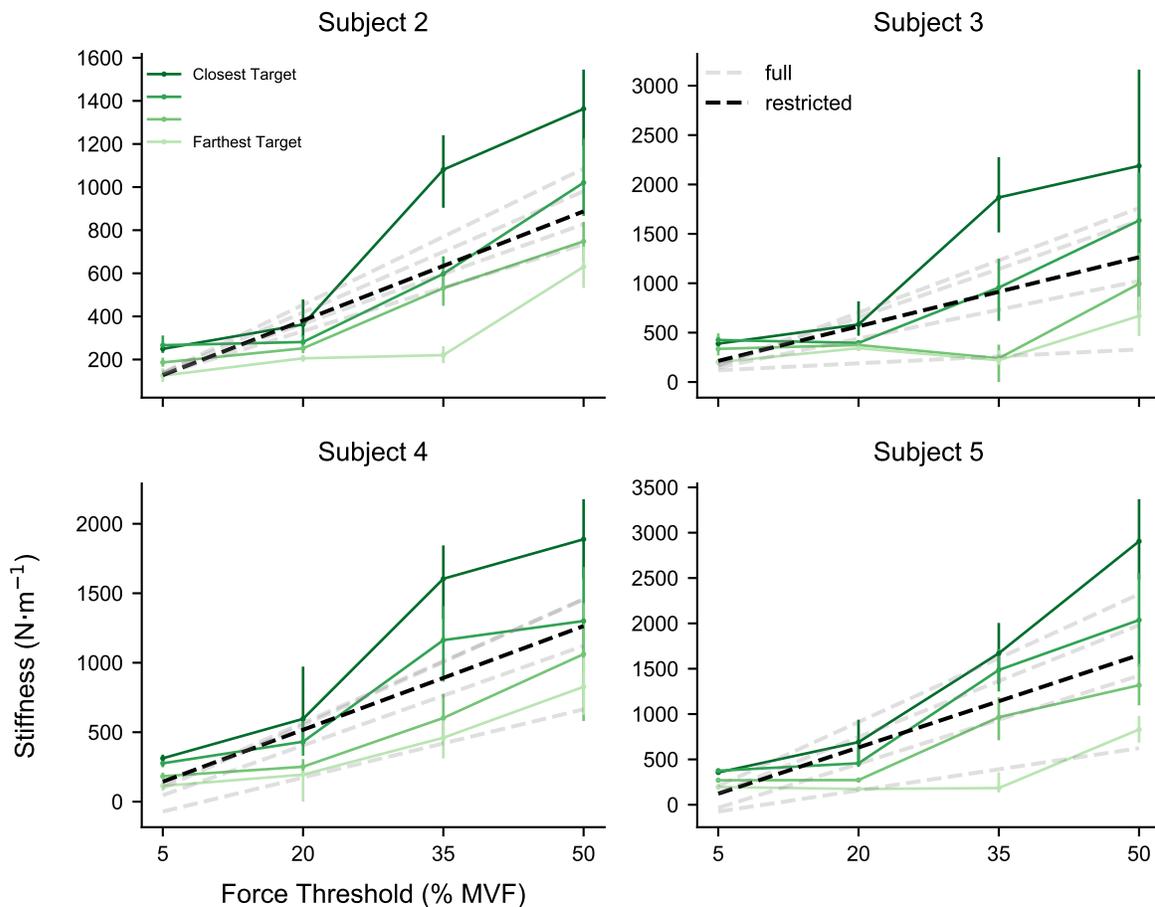


Figure 3.12: **Stiffness depended on both force and movement.** Stiffness exhibited a strong linear relation with force threshold, particularly for Subjects 3 and 5. However, the slope and offset of the linear relation depended on the target. Error bars indicate median and 95% confidence interval. The dashed grey lines represent the linear fit of the full regression of stiffness on force threshold and arrest position. The dashed black line represents the linear fit of the restricted regression of stiffness on force threshold alone. The physical dynamical model was fit using 200 ms of data beginning at movement onset.

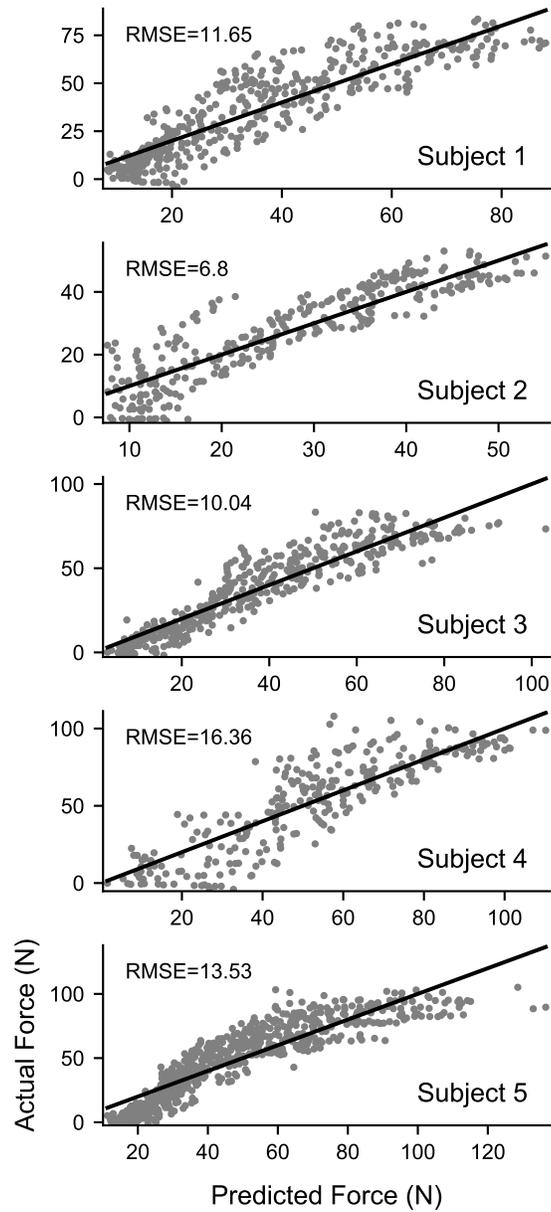


Figure 3.13: **EMG approximates force.** Time-varying force during the force ramp was regressed on the EMG from all 8 muscles. The linear model was able to consistently capture variability across all subjects. The data displayed are the median predicted and actual values from the bootstrap. The dark black line is the unity line representing a perfect fit. Although the small deviations from the unity line could be due to the non-linear relation between EMG and force, a linear model is a good approximation.

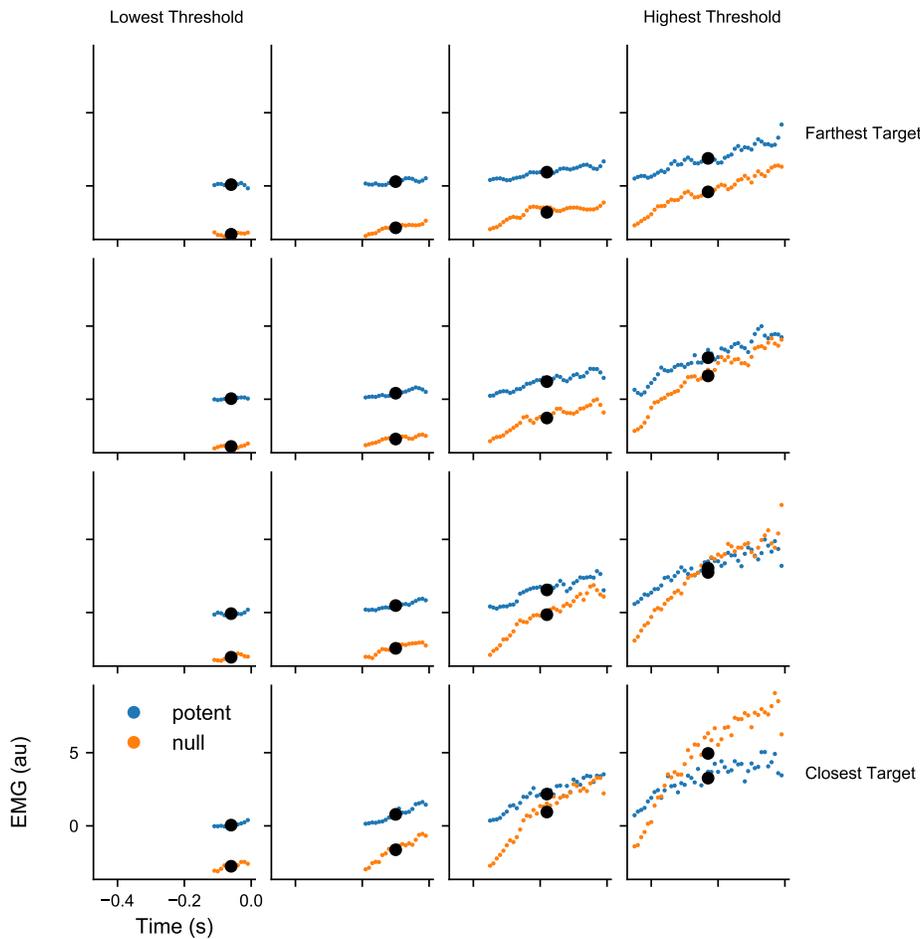


Figure 3.14: **Example of time-averaged EMG for subject 1.** There was a single value of stiffness for each task condition. To determine the effect of stiffness on potent and null EMG, we averaged EMG across time before movement onset. The result was a single potent and null EMG value for each task condition that was first averaged across trials and then averaged across time. The confidence intervals were calculated by resampling 1000 times from the 20 trials per condition. The colored data are the time series EMG for a single sample and the large black dots are the averages.

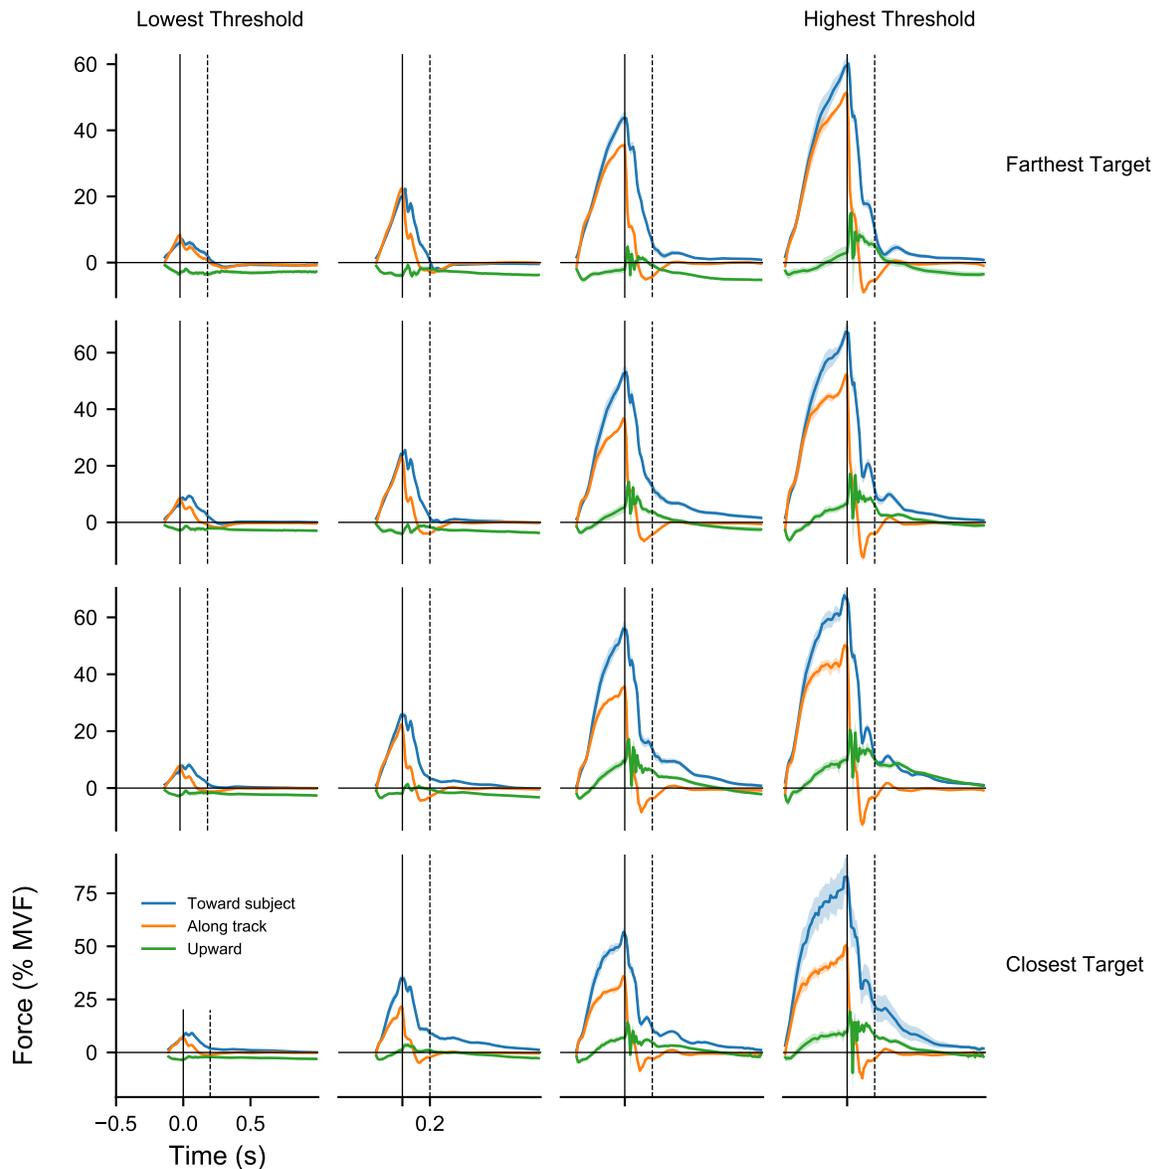


Figure 3.15: **Force in all 3 dimensions.** Although force exerted along the track did not vary across conditions with the same force threshold, force exerted toward the subject increased as the target moved closer to the start position. The correlation between the force exerted toward the subject and stiffness was likely due to the biomechanics of the arm, where modulation of off-axis muscle force enabled modulation of on-axis co-activation.

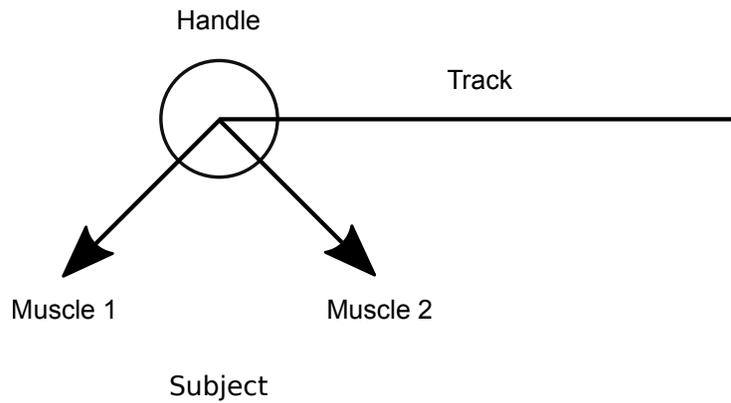


Figure 3.16: **Off-axis forces could correlate with on-axis stiffness.** Two imaginary muscles exert force on the handle according to the labeled vectors. If the magnitude of the muscle forces are equal, then the sum of the two muscle forces would result in zero force along the track and a non-zero force toward the subject. In addition, increasing the muscle force would increase both the force toward the subject and the stiffness along the track, without changing the force along the track.

## 4.0 NEURAL ENCODING MODELS FOR COMBINED FORCE-MOTION CONTROL

### 4.1 ABSTRACT

Combined force-motion control is important for manipulating objects flexibly in various contexts. Neural control of the musculoskeletal system is particularly effective at performing this type of behavior. Cortical activity in the motor cortex has been found to encode force and motion signals, but the extent to which combined force-motion control is encoded is unclear. We leveraged a ballistic-release task that systematically varied the control of force and motion while recording from a population of motor cortical neurons. We found that the temporal and time-averaged neural responses encoded information about both motion and force. By assuming the structure of the physical system used to produce the movement, we could decode impedance parameters and approximate the time-varying force exerted on an object. This suggests that impedance is a component of neural signaling that makes the motor control of object manipulation efficient.

### 4.2 INTRODUCTION

Manipulating objects to perform activities of daily living necessitates the combined control of force and motion ([Rancourt and Hogan, 2001](#); [Burdet et al., 2001](#)). Combined control is particularly important as a unified framework that achieves both motion in the free space and interaction force upon contact ([Hogan, 1985c](#)). The established representation of force ([Ashe, 1997](#)) and motion ([Schwartz, 2016](#)) in motor cortical (M1) activity makes this brain region

a likely source of information related to object manipulation. Although the representation of individual force or motion signals in M1 has been extensively studied, the representation of the combined control of force *and* motion has yet to be thoroughly investigated.

When M1 is stimulated by electrical impulses, muscles contract and limbs move ([Fritsch and Hitzig, 1870](#); [Ferrier, 1886](#)). The connection between M1 and motor behavior was first established anatomically via electrical excitation ([Leyton and Sherrington, 1917](#); [Phillips and Porter, 1977](#); [Asanuma and Rosen, 1972](#)) and later via viral tracing ([Rathelot and Strick, 2006, 2009](#)). Subsequent studies further established the functional connection by recording M1 activity from awake behaving monkeys ([Evarts, 1968](#); [Thach, 1978](#); [Cheney and Fetz, 1980](#); [Georgopoulos et al., 1982, 1986](#); [Kalaska et al., 1989](#); [Georgopoulos et al., 1992](#); [Schwartz, 1994](#); [Moran and Schwartz, 1999](#); [Kakei et al., 1999](#); [Sergio et al., 2005](#); [Griffin et al., 2015](#)).

Monkeys were trained to move their hand to various spatial targets using their whole arm. M1 activity was modulated with the direction of hand motion ([Georgopoulos et al., 1982](#)) and the relation was broadly tuned so that a unique direction could only be specified by incorporating the directional information from a population of neurons ([Georgopoulos et al., 1986](#)). Combining the directional information across time traced a continuous motion trajectory that predicted the actual motion of the hand ([Schwartz, 1994](#)). The speed with which the hand traveled along this trajectory was also found to be represented in the M1 population activity ([Moran and Schwartz, 1999](#)). Neural encoding models developed from these results have made it possible to predict the intended trajectory of hand motion in free space during reaching and drawing.

In other studies, monkeys were trained to exert force on a manipulandum in various contexts. In one context, external loads were applied to aid or hinder the manipulandum's motion, dissociating force and motion. M1 activity was found to reflect the change in force more than the motion ([Evarts, 1968](#); [Kalaska et al., 1989](#)). In another context, hand motion was eliminated altogether and monkeys were trained to exert isometric force on a stationary handle ([Georgopoulos et al., 1992](#)). The force was mapped to the motion of a cursor and M1 activity was again found to be related to the change in force. Finally, [Sergio et al. \(2005\)](#)

found that M1 activity reflected the change in force direction when moving inertial loads. The consistency of these results across various contexts indicates that force information is encoded in M1 activity.

Combined, these studies highlight the representation of both force and motion in M1 activity. Furthermore, these representations are similar, maintaining a strong directional component; M1 activity for a given neuron is often maximal in the same direction for both force and motion (Kalaska et al., 1989). This has led to a possible interpretation that M1 activity is a combination of the two representations, with a potential gain modulation in different contexts (Kalaska et al., 1989; Georgopoulos et al., 1992). We interpret this dual representation as potential for M1 to encode information about the combined control of force and motion. However, this hypothesis proves difficult to test because previous task designs emphasized either force or motion control. To explore the extent to which combined force-motion control is encoded in M1 activity, task conditions would need to systematically vary the parameters of force and motion.

We designed an experimental paradigm, based on impedance control, to emphasize the interaction between force and motion (Hogan, 1984c, 2014). Conceptually, impedance control is similar to stretching a spring, where one end of the spring is called the zero-force position and the other end of the spring is attached to an object. Stretching the spring (displacing the zero-force position away from the object) causes the spring to exert force on the object according to both the displacement and the stiffness. In this task, motor commands could be sent to the musculoskeletal system to set a zero-force position and an impedance, which is the generalization of stiffness to the time derivatives of position. The force exerted by the arm on the object would drive the actual position toward the zero-force position, just like when one end of the spring pulls the object toward the other. For motion in free space, displacement from the zero-force motion should be minimal and the controlled behavior resembles motion control. However, interactions with an object typically displace the position from the zero-force position, causing the manipulator to exert force on the object.

Despite the advantages of impedance control, relatively little is known about how it could be encoded in cortical activity. One barrier is the difficulty in estimating the zero-force position. Indirect estimates of the zero-force position have been found using two approaches.

First, the interaction force and arm impedance have been used to estimate the zero-force position for both simulated (Katayama and Kawato, 1993) and real arm movements (Gomi and Kawato, 1996). Second, because arm impedance generates interaction force in response to displacements from the zero-force position, the force exerted on the object is zero when the object position equals the zero-force position. To test the second approach, subjects were instructed to grasp a manipulandum and reach to different spatial targets (Hodgson and Hogan, 2000). The manipulandum’s trajectory was iteratively adjusted to minimize the interaction force. Importantly, the subjects were instructed to maintain the same motor commands for the original reach and not to adjust to the manipulandum’s iterative trajectories. The final position trajectory was then estimated as the zero-force position trajectory associated with the original, unperturbed reach. Although both approaches provide compelling estimates of the zero-force position, the author of the second approach suggested that it may be possible to record the motor commands directly from M1 (Hodgson and Hogan, 2000).

Another barrier to determining how cortical activity is related to impedance control is the difficulty in measuring arm impedance. Typically, arm impedance is measured by perturbing the arm in a way that is external to the task, either by applying a force pulse (Gomi and Kawato, 1996, 1997; Piovesan et al., 2013) or a position displacement (Mussa-Ivaldi et al., 1985; Burdet et al., 2001; Darainy et al., 2007; Franklin et al., 2007). Again, this approach depends on subjects not intervening during the perturbation.

Our goal was to employ a task that did not involve an external perturbation, but instead leveraged a ballistic-release paradigm that encouraged behavior that could be modeled as a physical dynamical system (see Chapter 3). We estimated arm impedance based on the force exerted on the object and the initial motion of the object as it was released. We hypothesized that M1 activity could encode a zero-force position and impedance signal that could be used in impedance control. We found that the temporal and time-averaged neural responses encoded information about both motion and force. Additionally, we were able to decode impedance parameters that made it possible to estimate time-varying force using few parameters.

## 4.3 METHODS

### 4.3.1 Subjects

Two rhesus monkeys were trained to perform a ballistic-release task with their right arm. All procedures were approved by the University of Pittsburgh’s Institutional Animal Care and Use Committee.

### 4.3.2 Behavioral paradigm and Experimental design

Each monkey was trained to (1) use its right hand to press a start button and then reach to grasp a handle, (2) pull on the handle with enough force to unlock it, and (3) position the handle within a specified target zone (Figure 4.1). To be successful, the monkey needed to pull with enough force to unlock the handle and hold the handle in the specified target zone for 300 ms. An effective control strategy would be similar to modeling the arm as a third-order physical system, where the arm would move to the correct target zone if the zero-force position was within the target zone (Feldman, 1966, 1986; Polit and Bizzi, 1979). For the same target zone and zero-force position, the changes in force on the handle would correspond to changes in impedance. Thus, a zero-force position and arm impedance could be preset for a given task condition (target zone and force threshold), obviating the need for corrective interventions during the movement. See section 2.1 for more details.

### 4.3.3 Data collection

During the task, we measured: (1) the force of the monkey pulling on the handle; (2) the 3D position of optical markers placed on the handle and the monkey’s hand, lower arm, upper arm, and torso; and (3) neural activity in the motor cortex. All data were synchronized in time and analyzed at 100 Hz. Data from individual trials were aligned on movement onset. See sections 2.2, 2.3, and 2.4 for more details.

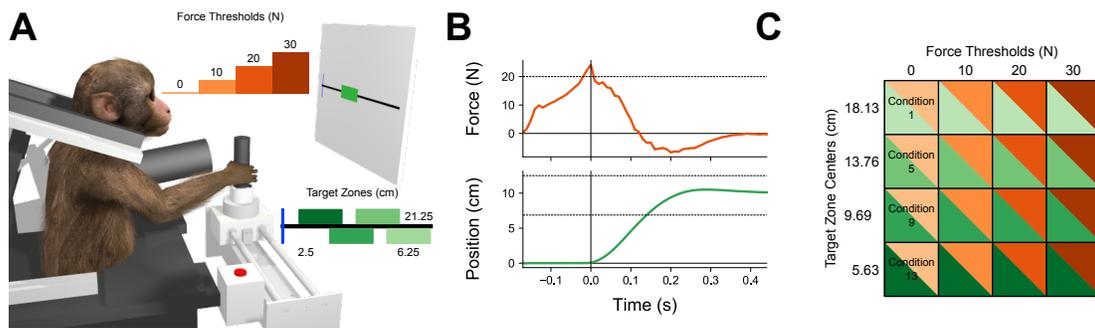


Figure 4.1: **Behavioral paradigm (A)** To initiate a trial, the monkey first pressed the start button (to the monkey’s right) and then grasped the handle. The monkey pulled on the handle while it was locked in place until the force threshold was crossed. The handle was then unlocked to move freely along the track and the monkey had to stop and hold the handle within the target zone for 300 ms. **(B)** Time-series of force and position were measured for each trial (a single representative trial is displayed here). **(C)** A single task condition was composed of a force threshold and a target zone.

#### 4.3.4 Physical dynamical system

To estimate impedance, we modeled the behavior of the arm as an equivalent physical dynamical system with a zero-force position and 4 impedance components that exert force proportional to position and its first 3 time derivatives (equation 2.1):

$$F(t) = K(x_0 - x(t)) - D\dot{x}(t) - A\ddot{x}(t) - J\dddot{x}(t),$$

where  $F$  is the force exerted on the handle ;  $x, \dot{x}, \ddot{x}, \dddot{x}$  is position and its first 3 time derivatives;  $x_0$  is the zero-force position, and  $K, D, A, J$  are impedance coefficients. Although this model describes the arm's behavior, it does not directly probe arm impedance, meaning that the first- and second-order impedance components are related to, but not measurements of, damping and inertia. The model was fit to the trial-averaged data for each task condition during the first 500 ms after movement onset.

The model maps a motion signal to the force exerted on the object. The zero-force position is the reference position for the model and governs where the handle would eventually come to rest, which can be interpreted in this task as the hold position. The impedance coefficients govern the transition period and make it possible to exert different forces while keeping the zero-force position (hold position) constant. When provided with the actual motion as input, the physical model can provide an estimated force as output, according to equation 2.1. See section 2.5 for more details.

#### 4.3.5 Model stiffness fit to force threshold and hold position

We have previously demonstrated in humans that variability in the model's stiffness can be explained by force and motion task parameters (Section 3.4.2). We replicate this finding in monkeys by fitting model stiffness to force threshold and hold position according to equation 4.1

$$K = \beta_0 + \beta_1 Threshold + \beta_2 HP + \beta_3 (Threshold)(HP), \quad (4.1)$$

where  $K$  is a  $[1 \times C]$  matrix of model stiffness values ( $C$  is the number of task conditions),  $Threshold$  is a  $[1 \times C]$  matrix of force thresholds,  $HP$  is a  $[1 \times C]$  matrix of hold positions (defined as the position corresponding to the time when velocity remained below 2.5 cm/s for 100 ms), and  $\beta_{0,1,2,3}$  are regression constants.

#### 4.3.6 De-mixed principal components analysis

The firing rates of individual motor cortical neurons in a population often vary together, and the different modes of shared variability can be interpreted as different common drivers influencing each neuron in the population. The effect of the common drivers on the shared variability can be measured by performing principal components analysis (PCA). In this task, force threshold and/or target zone might be common drivers and the shared variability explained by each task parameter would then be a measure of the extent to which the task parameters are encoded in the population. However, the common drivers revealed by PCA typically represent a mixture of task parameters, making it difficult to determine how much of the shared variability can be attributed to force threshold or target zone.

De-mixed PCA (dPCA) is an approach to measuring the shared variability that emphasizes the contribution of different parameters (Kobak et al., 2016). We performed dPCA by smoothing the firing rates to remove high-frequency noise using a Gaussian filter with a 30 ms standard deviation, averaging the firing rates across successful trials, and then z-scoring the firing rates of each neural unit across time and task conditions. We separated different parameters by grouping the result according to the 4 force thresholds and 4 target zones, resulting in a data matrix that was  $[R \times F \times P \times T]$ , where  $R$  represents the number of neural units,  $F$  represents the 4 force thresholds,  $P$  represents the 4 target zones, and  $T$  represents the number of 10 ms time bins. The middle two parameters,  $F$  and  $P$ , represent the 16 task conditions.

The effect of either force threshold or target zone on the shared variability can be isolated by averaging across different parameters. The population response that is independent of force threshold and target zone, i.e. only driven by time, is isolated by averaging across both the force threshold and target zone parameters. The population response driven by force

threshold is isolated by removing the time response and then averaging across the target zone parameter. Likewise, the population response driven by target zone is isolated by removing the time response and then averaging across the force threshold parameter. Finally, the population driven by the interaction of force threshold, target zone, and time is isolated by removing the time, threshold, and target responses and then averaging across time.

The dPCA algorithm uses a special form of reduced-rank regression to reconstruct the population response driven by the time, force threshold, target zone, and interaction parameters while at the same time maximizing the shared variability. The algorithm returns scores that correspond to the shared variability driven by each parameter and the percent of the total variance explained by each parameter. The scores summed across parameters approximate the scores corresponding to the first principal component in traditional PCA. Likewise, the percent variance explained summed across parameters approximates the percent variance explained by the first principal component in traditional PCA. The algorithm was performed using code from <http://github.com/machenslab/dPCA>. We report dPCA results for 1 component separated into the parameters of time, force threshold, target zone, and interaction. The algorithm default returns time-independent threshold/target terms and threshold/target-time interaction terms. Because we expected all neural components to change with time, the time-independent and threshold/target-time interaction terms were combined according to the algorithm’s documentation.

#### **4.3.7 Encoding model for force threshold and target zone**

The task conditions were constant across time and we were interested in how the responses of each neural unit varied across task conditions. Therefore, to aid the comparison, we averaged each unit’s raw firing rate over a time window that began when the monkey grasped the handle (30 or 90 ms before movement onset) and ended 500 ms after movement onset (monkey S and I). Averaging across trials and then across time resulted in an average rate for each neural unit per task condition that maintained large time-scale features that would influence the average.

The task parameter encoding model was defined as

$$FR = \beta_0 + \beta_1 Threshold + \beta_2 Target, \quad (4.2)$$

where  $FR$  is a  $[R \times C]$  matrix of the average firing rate of a single neural unit per task condition ( $R$  is the number of neural units),  $Threshold$  is a  $[1 \times C]$  matrix of force thresholds,  $Target$  is a  $[1 \times C]$  matrix of target zones, and  $\beta_{1,2}$  are each  $[R \times 1]$  matrices of regression coefficients, and  $\beta_0$  is a  $[R \times 1]$  of constant offsets. The model was fit for each neural unit across the 16 task conditions. Each unit's firing rate was first averaged across trials and then across time. To compare the coefficients across neural units, the average firing rates, the force thresholds, and the target zones were z-scored across the 16 task conditions.

Coefficients were deemed different from zero using a bootstrap distribution. The trials per unit per task condition were re-sampled, with replacement, 1000 times and the target and threshold coefficients were calculated across task conditions for each re-sample. A tuned neural unit had at least one of the target or threshold coefficients whose distribution did not include zero ( $p < 0.05$ ). To help distinguish between the influence of each task parameter, coefficients whose distribution included zero were set to zero ( $p < 0.05$ ).

#### 4.3.8 Decoding the physical model parameters

The physical model maps the motion of the handle to the force exerted on the handle. Therefore, the time-varying force exerted on the handle can be estimated by decoding the physical parameters and then using them to map the actual motion to force.

The physical parameters were decoded using an optimal linear estimator (OLE) and evaluated using leave-one-condition-out cross-validation. Because the variance of the physical parameters differed dramatically (see figure 4.4), they were first z-scored across task conditions.

Because the physical parameters were constant across time per task condition, each neural unit's firing rate per task condition was first averaged across trials and then across time to yield an average rate per neural unit per task condition. The decoder was trained by fitting a planar encoding model

$$FR = \mu + \beta_1 x_0 + \beta_1 K + \beta_2 D + \beta_3 A + \beta_4 J \quad (4.3)$$

using least-squares regression, where  $FR$  is a  $[R \times C]$  matrix of average firing rates;  $x_0$  is a  $[1 \times C]$  matrix of the z-scored zero-force positions;  $K, D, A, J$  are each  $[1 \times C]$  matrices of the z-scored impedance coefficients corresponding to position and its first three time derivatives;  $\beta$  are each  $[R \times 1]$  matrices of regression coefficients; and  $\mu$  is a  $[R \times 1]$  matrix of constant offsets.

To invert the encoding model, the regression coefficients  $\beta$  were fit to the firing rates using weighted least-squares regression. We used weighted least-squares regression because it out-performs ordinary least-squares regression when the variance of the errors is non-constant. It follows the same procedure as ordinary least-squares regression, but with an added weight to the errors to balance the difference in variance. The weights  $W$  were an  $[R \times R]$  diagonal matrix of the inverse variance of the residuals from the encoding model.

The decoder consisted of a  $[5 \times R]$  coefficient matrix  $\theta$  that mapped the  $[R \times C]$  matrix of trial- and time-averaged firing rates  $FR$  to a  $[5 \times C]$  matrix of z-scored physical parameters  $Z$  according to

$$Z = \theta(FR - \mu), \quad (4.4)$$

where  $\theta = (\beta^T W \beta)^{-1} \beta^T W$  and  $\mu, \beta$  are from the encoding model in equation 4.3. The decoder was evaluated using the root-mean-squared prediction error (RMSE) across the 5 z-scored physical parameters and 16 task conditions using leave-one-condition-out cross validation.

The time-varying force was estimated by first transforming the decoded z-scored physical parameters back to the original units. The actual motion during a 500 ms time window, beginning at movement onset, was then mapped through equation 2.1 using these parameters, estimating the force exerted on the handle. The estimated force was evaluated using the RMSE across the 16 task conditions and fifty 10 ms time bins.

## 4.4 RESULTS

We trained two monkeys to control force and motion during a ballistic-release task. Data were recorded across 11 days for each monkey, with at most 2 days between recording sessions. Monkey S completed 4030 successful trials and monkey I completed 2347 successful trials. Trial data were aligned on movement onset and averaged across successful trials for each task condition. Bootstrapped confidence intervals were calculated by re-sampling the data 1000 times, with replacement, from the successful trials per condition.

### 4.4.1 Force and motion control varied according to force threshold and target zone

To complete the task, the monkeys pulled on the handle with enough force to cross the force threshold and then moved and held the handle within the target zone (Figure 4.2). The task conditions dissociated force and position by evoking similar position trajectories for each of the force thresholds and similar force trajectories for each of the target zones. Movement onset began at time 0. Signals were very consistent across trials for a given task condition.

The monkeys pulled on the handle quickly, crossing the force threshold in 30 ms and 90 ms on average (monkey S and I). Monkey S tended to initially exert force in the direction opposite motion when it grasped the handle, quickly reversing the force direction to cross the threshold. Both monkeys increased the exerted force rapidly to the threshold with no plateau or apparent anticipation of threshold crossing.

The monkeys moved to the hold position smoothly, overshooting the hold position by a small distance but still staying within the target zone, typical of an arm impedance that is slightly under damped. Task conditions with near targets or high thresholds were associated with larger absolute overshoot, suggesting that arm impedance was modulated across task conditions.

Greater maximum velocities and longer positive velocity durations were associated with farther target zones. However, the time of maximum velocity was similar across targets, with a slightly longer time till maximum velocity for farther targets. Maximum velocity for

monkey S was also positively correlated with force threshold. This might have been a result of the higher force thresholds for monkey S and was also seen in the human subjects, who had relatively high force thresholds (see Figure 3.2).

As expected, maximum acceleration was positively correlated with force threshold. Although the maximum was also slightly related to target zone, the initial decrease in acceleration was similar across targets. Shortly after movement onset the acceleration values were negative, mirroring the force values and indicating that the handle was being slowed.

#### 4.4.2 Physical dynamical model fit transient and steady-state behavior

The arm was modeled as a physical dynamical system according to equation 2.1 and fit to the behavior using 500 ms of data, beginning at movement onset. The physical model consisted of 5 free parameters: a zero-force position and 4 impedance coefficients corresponding to position and its first 3 time derivatives.

**4.4.2.1 Zero-force position matched hold position** In the context of this task, the zero-force position can be interpreted as an estimate of the hold position. We found that the zero-force position matched the hold position across task conditions (RMSE = 0.37 cm [0.34, 0.41] and 0.97 cm [0.93, 1.01], monkey S and monkey I), indicating that the physical model estimated the end of the movement well (Figure 4.3).

The variance in hold position for a given target was slightly less for monkey S compared with monkey I (Figure 4.2). This consistency was also reflected in the zero-force position. Monkey S moved more rapidly and, for the given 500 ms time window, reached the hold position sooner than monkey I. This provided better data for the model fit, which could explain why the zero-force position more accurately matched the hold position for monkey S compared with monkey I.

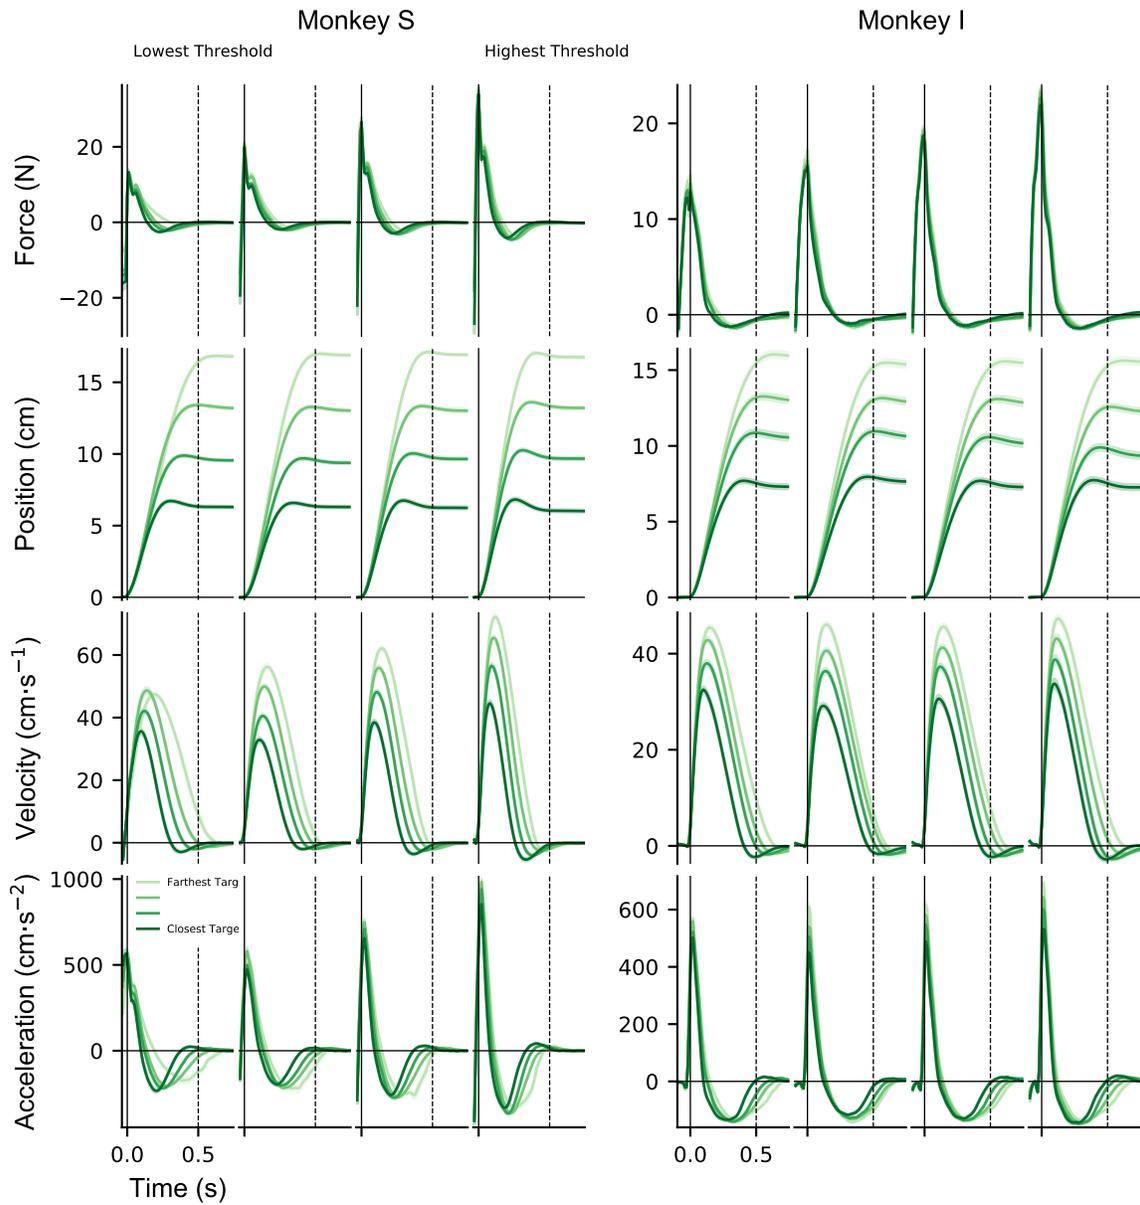


Figure 4.2: **Force and Kinematics.** Both monkeys crossed the force threshold and moved smoothly to each target zone. The force and position profiles varied with force threshold and target zone, respectively, while the velocity and acceleration changes were related to the combination of force threshold and target zone. Movement onset is designated as time 0 and data were averaged across successful trials. Shading represents the median and 95% confidence interval.

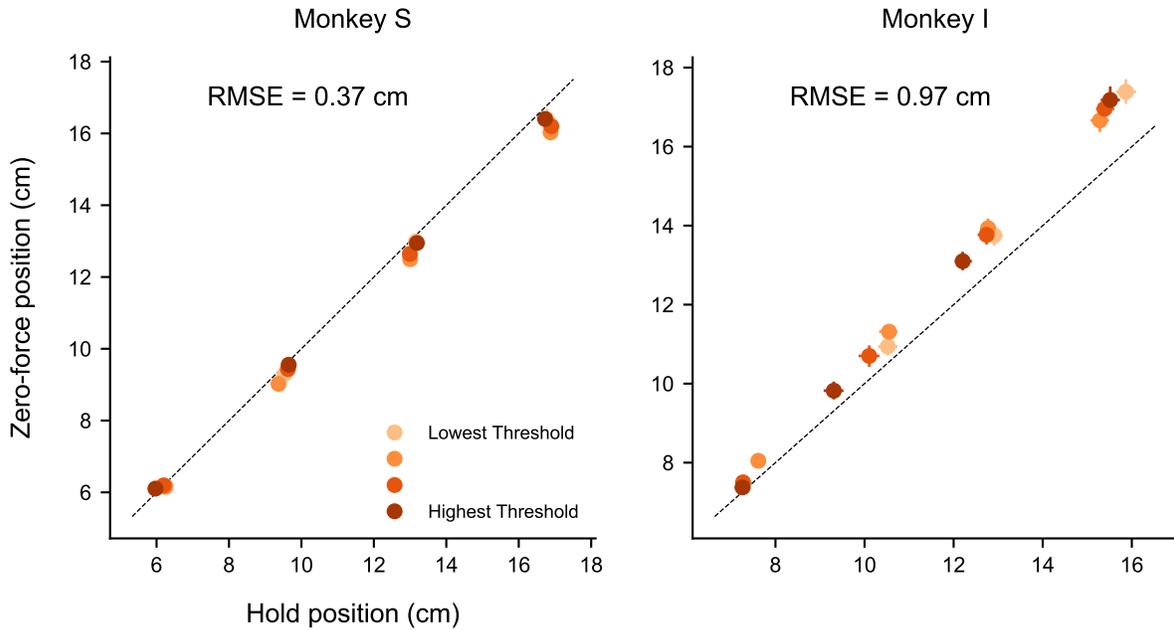


Figure 4.3: **Zero-force position matched the hold position.** The match between the zero-force position and the hold position indicates that the physical model fit the end of the movement well. Although the match was slightly better for monkey S compared with monkey I, the model for both monkeys was a good match for all force thresholds and target zones. Error bars represent the median and 95% confidence interval for both axes.

**4.4.2.2 Model impedance depended on both force and motion** Whereas the zero-force position describes the steady-state behavior of the physical model, the impedance coefficients describe the transition from the initial state to the steady state. For a given target zone, the zero-force position was expected to be similar across force thresholds and we expected impedance to increase as the release force increased.

We found that, for a given target zone, the model's stiffness was approximately linearly related to force threshold (Figure 4.4). However, the slope and offset of the linear relation depended on the motion, showing that, although stiffness depended on force, it also was modulated by movement distance. A regression model (equation 4.1) that fit stiffness to force threshold, hold position (offset), and an interaction term (slope), was able to explain much of the variance ( $R^2 = 0.96, 0.90$  monkey S and I), providing good evidence that the model's stiffness was related to both force and position in this task.

Although the model's first- and second-order impedance coefficients increased with the force threshold, they did not exhibit a consistent pattern across target zones. Because the task requirements emphasized handle position, we expected stiffness to be the primary task-related component of impedance. Although the monkeys were restrained in a chair, they were able to rotate their body slightly. Qualitatively, monkey S tended to rotate its body toward the handle's lock position for more difficult task conditions, while monkey I maintained a more consistent body orientation across task conditions. This could have accounted for the model's unexpectedly large second-order impedance values for monkey S, although additional quantitative analyses are needed to confirm this. The model's second-order impedance coefficient values of 0 for both monkeys at the lowest force threshold could be a result of the combination of low threshold (minimal pre-loading) and inaccuracies in the model fit. The model's second-order impedance coefficient dominates the initial motion, when the load is suddenly released and acceleration is high, and the lower acceleration values for the lowest threshold may have not fit the model well.

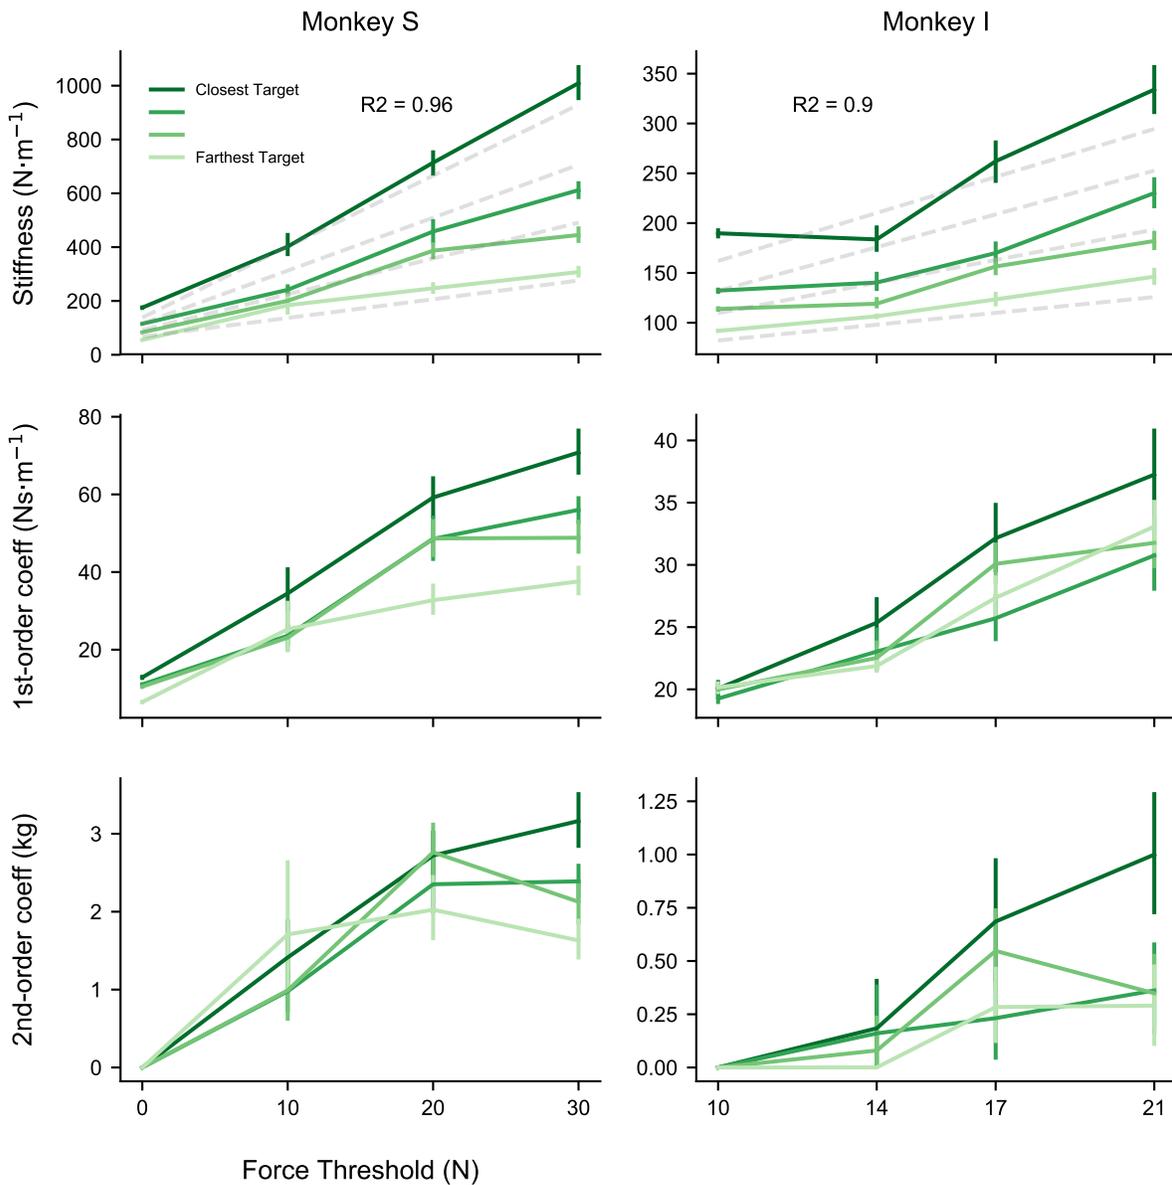


Figure 4.4: **Model stiffness depended on both force threshold and hold position.** The model's impedance values were generally proportional to force threshold. However, stiffness was also related to motion. A regression model that fit stiffness to force threshold and hold position explained much of the variance across task conditions (dashed grey lines). Error bars represent the median and 95% confidence interval.

### 4.4.3 Neural firing rates showed general trends that reflected task parameters

Two micro-electrode arrays were chronically implanted in the arm/hand region of motor cortex (M1). From these arrays, we recorded the time and triggered voltage waveforms of 96-192 channels of neural activity. Waveforms were sorted offline into 213 neural units for monkey S and 127 neural units for monkey I. Although effort was taken to record from both the pre-central gyrus and the anterior bank of the central sulcus, relatively few neural units were identified near the sulcus (14 and 9, monkey S and I). However, no differences in responses were identified and therefore we will not distinguish between units recorded from the gyrus and the bank of the sulcus.

For monkey I, activity on the lateral array during a passive exam was related to the face/lips (Figure 4.13a). Therefore, we performed additional analyses on the different channels and neural units from this array for monkey I to confirm that the results were consistent with those found from monkey S (Figures 4.13b and 4.14).

Each unit was tracked across recording sessions and spike times were first converted to fractional interval firing rates in 10 ms time bins and then averaged across successful trials per condition (Figure 4.5). Most units were tracked for 1-3 sessions, but a handful were tracked for the full 11 sessions.

**4.4.3.1 Time-varying neural responses** Firing rates of individual units varied across time and task conditions (Figure 4.6). Rates were smoothed using a Gaussian filter with 30 ms standard deviation. Temporal profiles could be bell-shaped, inverted bell-shaped, decreasing, or increasing. Although the temporal profiles of some individual firing rates were related to task parameters (motion-top row, force-middle row), most firing rates were complex and not readily interpretable.

We measured the extent to which a given task parameter influenced the population using de-mixed principal components analysis (dPCA). The shared variability of the population that was driven by different parameters is shown as a time-varying score in figure 4.7. Each panel consists of 16 traces corresponding to the 16 task conditions.

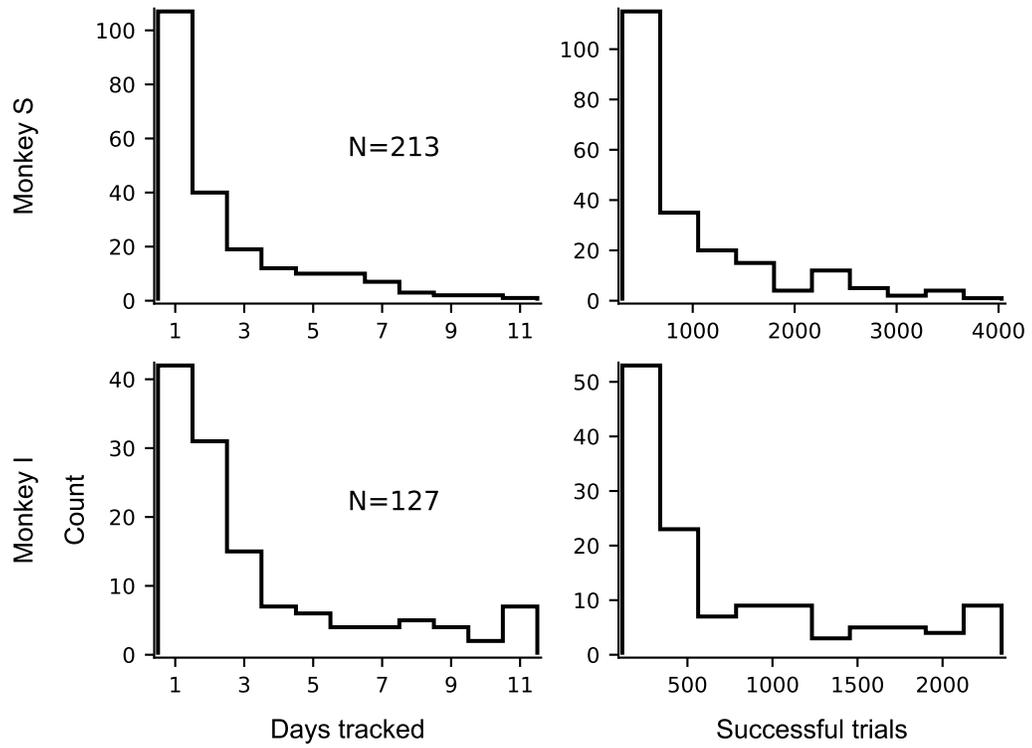


Figure 4.5: Neural histogram of days tracked and successful trials per neural unit. Neural units were tracked across 11 recording sessions.

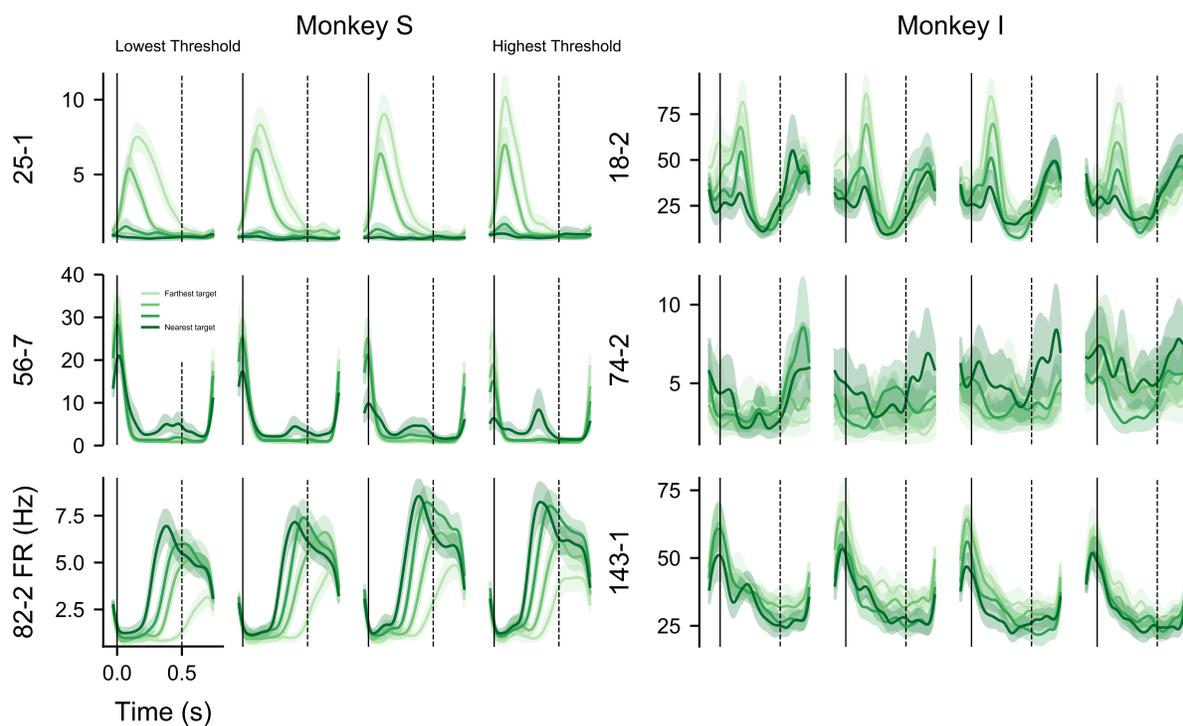


Figure 4.6: **Example firing rate time profiles.** Fractional interval firing rates were calculated in 10 ms time bins and smoothed using a 30 ms Gaussian filter. Although the temporal profiles could be complex, there was a general relation to motion (top row), force (middle row), or a combination of force and motion (bottom row). Shading represents the median and 95% confidence interval. Each response is labelled according to the recording channel and the sorted unit on that channel (channel-unit).

The time score (Figure 4.7, first row) represents shared variability that was independent of task condition, i.e. driven only by time, and explained 36.1% [35.4, 36.8] and 33.2% [31.5, 34.9] of the total variance (monkey S and I). Unsurprisingly, this had the largest relative influence on the population and likely reflected the temporal focus of the behavior around movement onset (time 0).

The threshold score (Figure 4.7, second row) represents shared variability that was driven by force threshold and explained 2.6% [2.3, 3.0] and 2.3% [1.5, 3.6] of the total variance (monkey S and I). The scores for task conditions with the same force threshold varied together tightly, but there was some mixture of the scores for the lowest two force thresholds for monkey S and toward the end of movement for monkey I. The relatively low representation of the threshold score in the population makes it difficult to determine if this mixture reflected a particular signature of neural activity or was merely an artifact of noise.

The target score (Figure 4.7, third row) explained 15.9% [15.3, 16.6] and 6.8% [5.9, 8.0] of the total variance (monkey S and I). The target zone's influence was much stronger than the force threshold's, particularly for monkey S. Other than the closest two targets for monkey I, all target zones were well separated across time. The relative stability of the scores across time, particularly for monkey S, could be representative of the zero-force position, which was assumed constant across time.

The interaction of target, threshold, and time was represented by the interaction score (Figure 4.7, fourth row) and explained 1.5% [1.3, 1.8] and 5.0% [3.3, 6.9] of the total variance (monkey S and I). There was quite a bit of differences between the two monkeys, with monkey S having stable scores across time but which explained little variance, while monkey I had scores that explained more variance but were dynamic across time. In either case, the scores are not particularly interpretable for force threshold, target zone, or a combination of the two.

The summation of scores across parameters, represented by the total score (Figure 4.7, bottom row), approximates the scores corresponding to the first principal component for traditional principal components analysis and explained 56.1% [55.4, 57.0] and 47.5% [45.5, 49.4] (monkey S and I). The time profile of these scores have the same general shape as

the time score, reflecting the strong contribution of that parameter. Additionally, the general organization of scores by target zone for monkey S indicates the particularly strong contribution of the target score.

**4.4.3.2 Time-averaged neural responses** Four broad categories of neural responses were found by comparing the coefficients of equation 4.2. The trial-averaged firing rates were averaged across time and regressed against the task parameters. A target response (1), shown in the top row in figure 4.8, is one whose variability can be explained by target zone (colors), but not by force threshold (x-axis). A force response (2), shown in the middle row, is one whose variability can be explained by force threshold, but not by target zone. A stiffness response (3), shown in the last row, is one whose variability can be explained by a combination of force threshold and target zone. For monkey S, the firing rates in the last row increased as the distance to the target zone decreased and as the force threshold increased, a pattern consistent with a positive correlation to stiffness (compare with figure 4.4). In contrast, the response in the last row from monkey I shows a pattern consistent with a negative correlation to stiffness. Both positive and negative correlations were considered stiffness responses. An other response (4), is one whose response was tuned but did not fall into the other three categories.

The target and stiffness responses were most represented in the tuned units for both monkeys. Monkey S had 176 tuned units of the 213 recorded and monkey I had 101 tuned units of the 127 recorded. A tuned unit had an encoding model with at least one of the target or threshold coefficients in equation 4.2 different from zero ( $p < 0.05$ ). The tuned responses were further categorized according to the coefficients by setting coefficients with  $p > 0.05$  to 0. Figure 4.9 shows the coefficient values for each unit, highlighting the target response (green), threshold response (orange), and stiffness response (purple) from figure 4.8. There were 73 and 49 target responses, and 27 and 13 threshold responses (monkey S and I). Perhaps most striking, of the 76 and 39 responses tuned to both target and threshold, 59 and 29 were stiffness responses (monkey S and I). These results indicate a salient motion and stiffness signal encoded in the population and suggest appropriate signaling for impedance control.

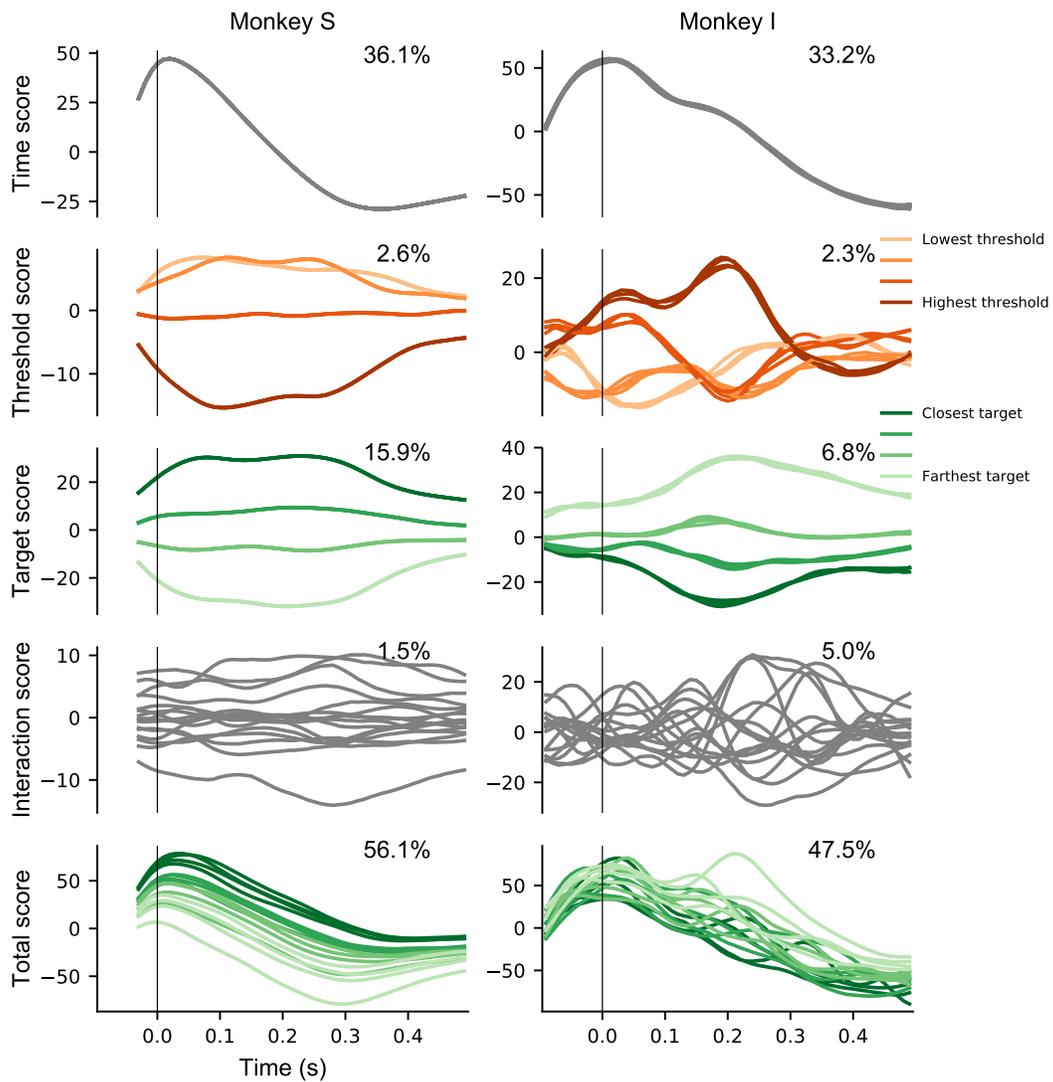


Figure 4.7: **De-mixed PCA revealed shared variability driven by force threshold and target zone.** De-mixed PCA found the shared variability of the population that was driven by different parameters. Each panel contains 16 traces corresponding to 16 task conditions. The percent of the total variability explained by the first four rows is reported. The bottom row is the summation across parameters and approximates the scores from the first traditional principal component; its percent variance explained is the sum of the first four rows.

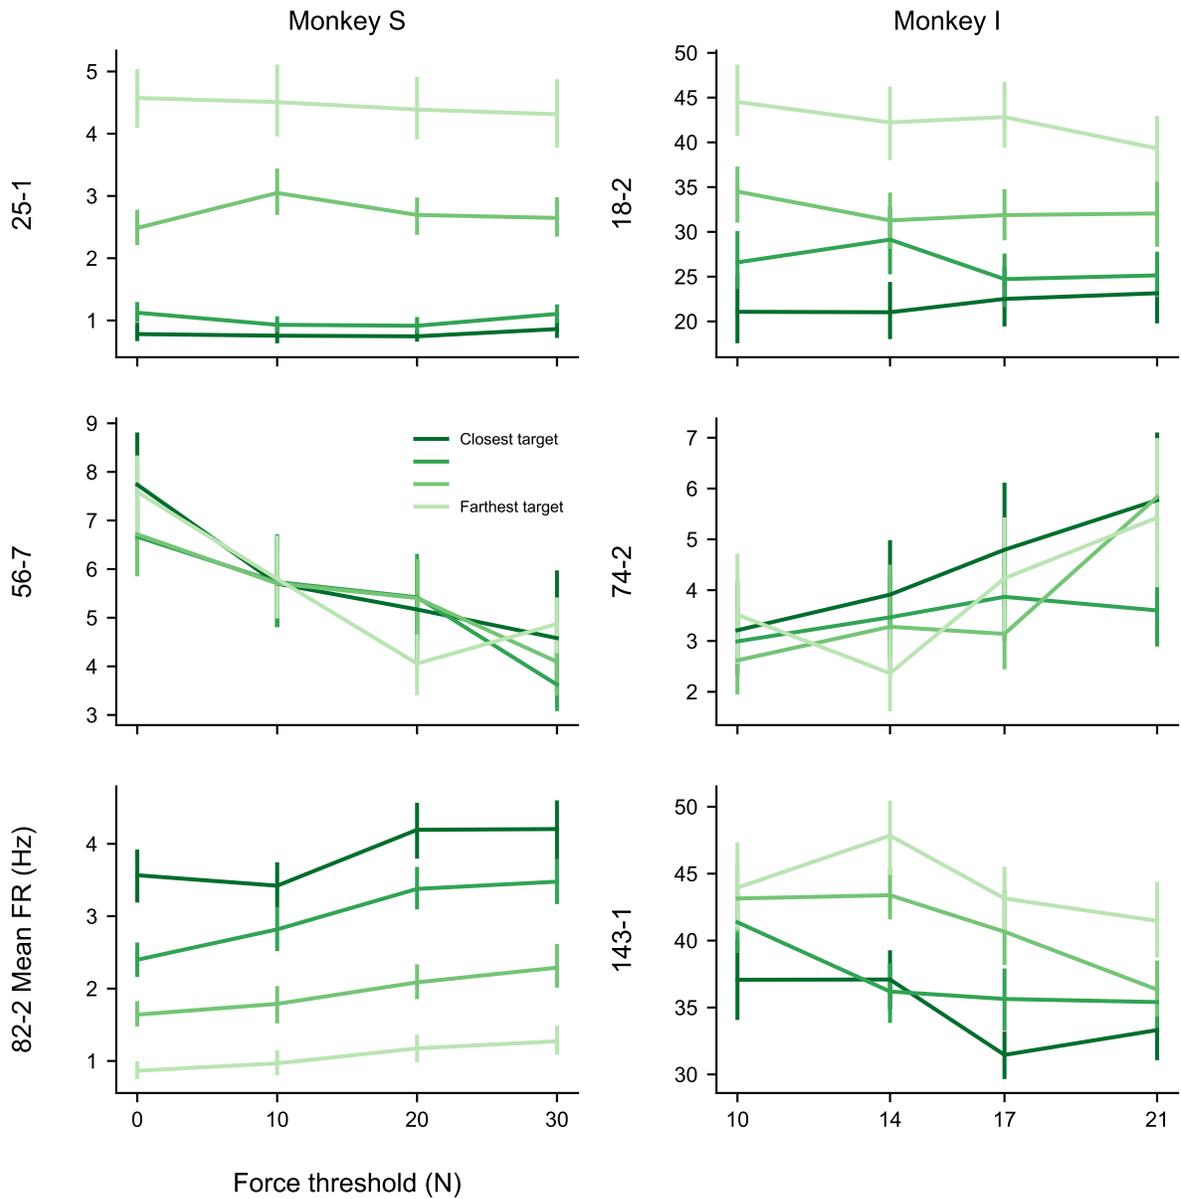


Figure 4.8: **Example mean firing rates.** The time-averaged firing rates could be explained by target zone (top row), force threshold (middle row), or a combination of the target zone and force threshold (bottom row). The neural units are the same as those in figure 4.6. Error bar represent the median and 95% confidence interval.

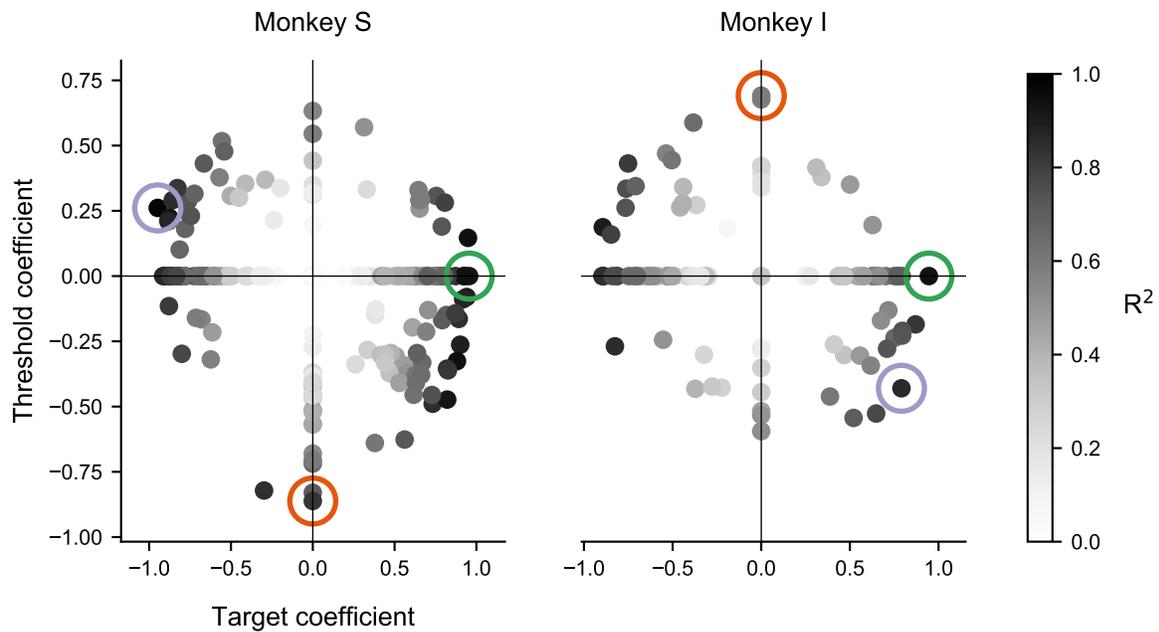


Figure 4.9: **Target and stiffness responses were most prevalent in the tuned units.** Each dot represents one neural unit and corresponds to the target and threshold coefficients from the regression of time-averaged firing rates on target zone and force threshold. The orange, green, and purple circles highlight the threshold, target, and stiffness responses depicted in figures 4.6 and 4.8. Of the tuned units, the target (x-axis) and stiffness (quadrants II and IV) responses were most prevalent.

#### 4.4.4 Neural encoding and decoding models for physical model parameters and force

In a task where the primary success criterion is the position of the object, it might not be surprising that a motion signal is strongly encoded in the neural population. However, it is also true that the judge of many object manipulation tasks is the motion of the object. Despite the importance of motion, mechanics dictate that an object can only move if an external force is exerted on it. A mapping from the motion signal to the exerted force is described under impedance control by the physical model in equation 2.1. This motivated us to use an optimal linear estimator (OLE) to decode the physical model parameters. Using these, we then calculated the force estimated by the physical model.

We estimated the physical model parameters for one condition using a decoder trained on the remaining conditions (leave-one-condition-out cross-validation). The zero-force position and model impedance coefficients were calculated for each task condition and z-scored across conditions. Data from the 15 training conditions were used to find parameters for each unit's encoding model (equation 4.3). The decoder was found by inverting the encoding model according to equation 4.4 and tested by providing the firing rates from the test condition as input and returning the decoded physical model parameters as outputs.

The result of this process was 5 decoded physical model parameters for each of the 16 conditions (Figure 4.10). The decoding results were slightly better for monkey S than for monkey I (RMSE = 0.68 and 0.71), but both monkeys had values that were clustered around the unity line. The physical parameters with lower values were generally over-estimated and the physical parameters with higher values were generally under-estimated. This could have been because the encoding model fit 5 parameters from just 16 observations. More observations would be possible with time-varying estimates of the physical parameters and would likely have resulted in a more confident encoding model and improved the decoding accuracy and precision.

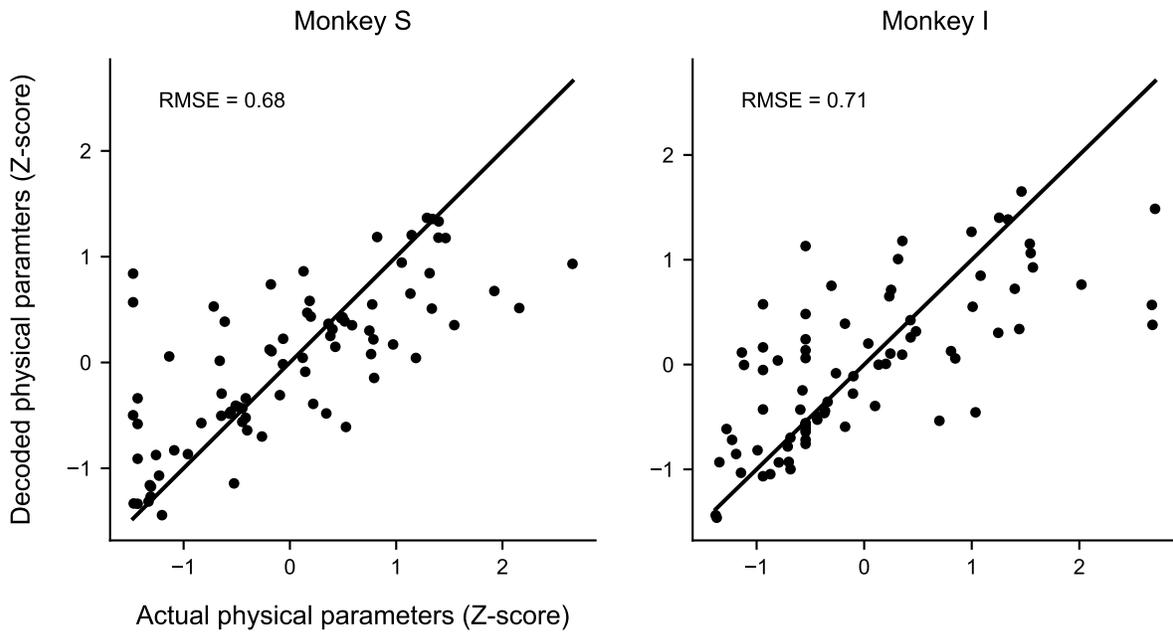


Figure 4.10: **Physical model parameters were decoded from motor cortical firing rates.** The physical parameters were z-scored and used to train an OLE decoder. The cross-validated decoded parameters approximated the actual values across all 5 parameters and 16 task conditions. The decode for monkey S was slightly more accurate than monkey I, perhaps because of the additional neural units recorded from monkey S.

The decoded z-scored physical model parameters were transformed back to their original units and used, with the actual motion, to estimate force after movement onset according to equation 2.1. The estimated force matched the actual force for both monkeys (RMSE = 3.35 N and 1.56 N, monkey S and I) across time, task conditions, and for both low and high values (Figure 4.11).

Furthermore, the match was consistent across task conditions, capturing the rapid decrease in force following movement onset, the negative force that slowed the handle, and the plateau that held the handle in the target zone (Figure 4.12). The dramatic exception was the lowest threshold and farthest target for monkey S (top left panel). This was the one instance when the decoded physical model parameters were unrealistic (negative stiffness). This problem can be addressed in future work.

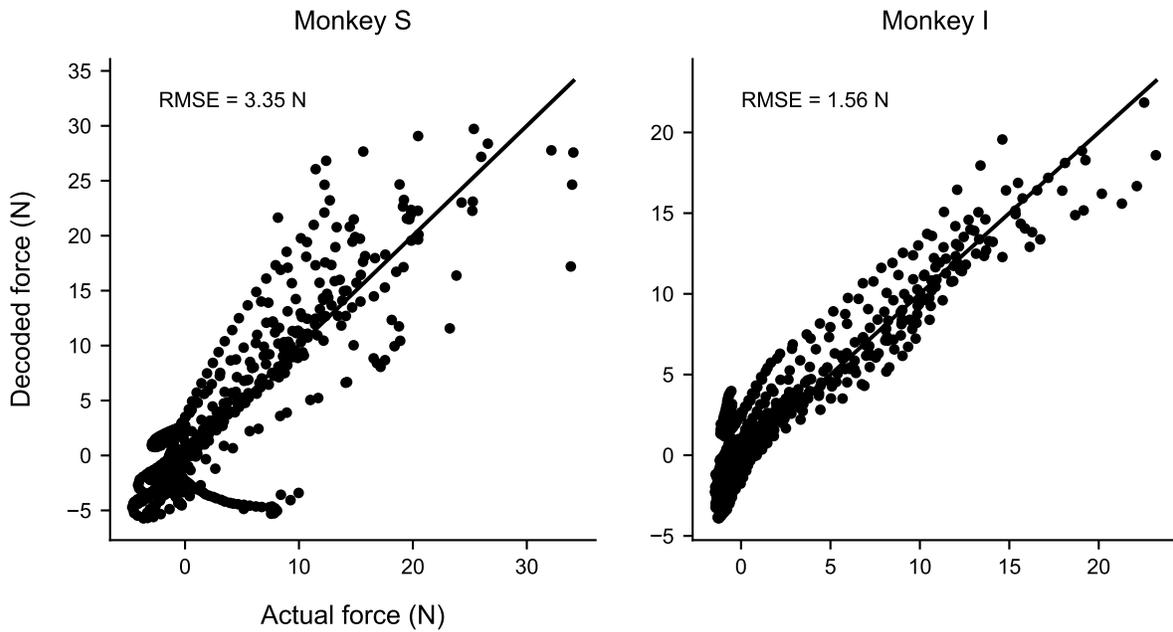


Figure 4.11: **Estimated force from decoded physical model parameters.** The decoded physical model parameters were used to estimate force during the 500 ms after movement onset according to the physical model. The decoded force matched the actual force consistently for both low and high magnitudes.

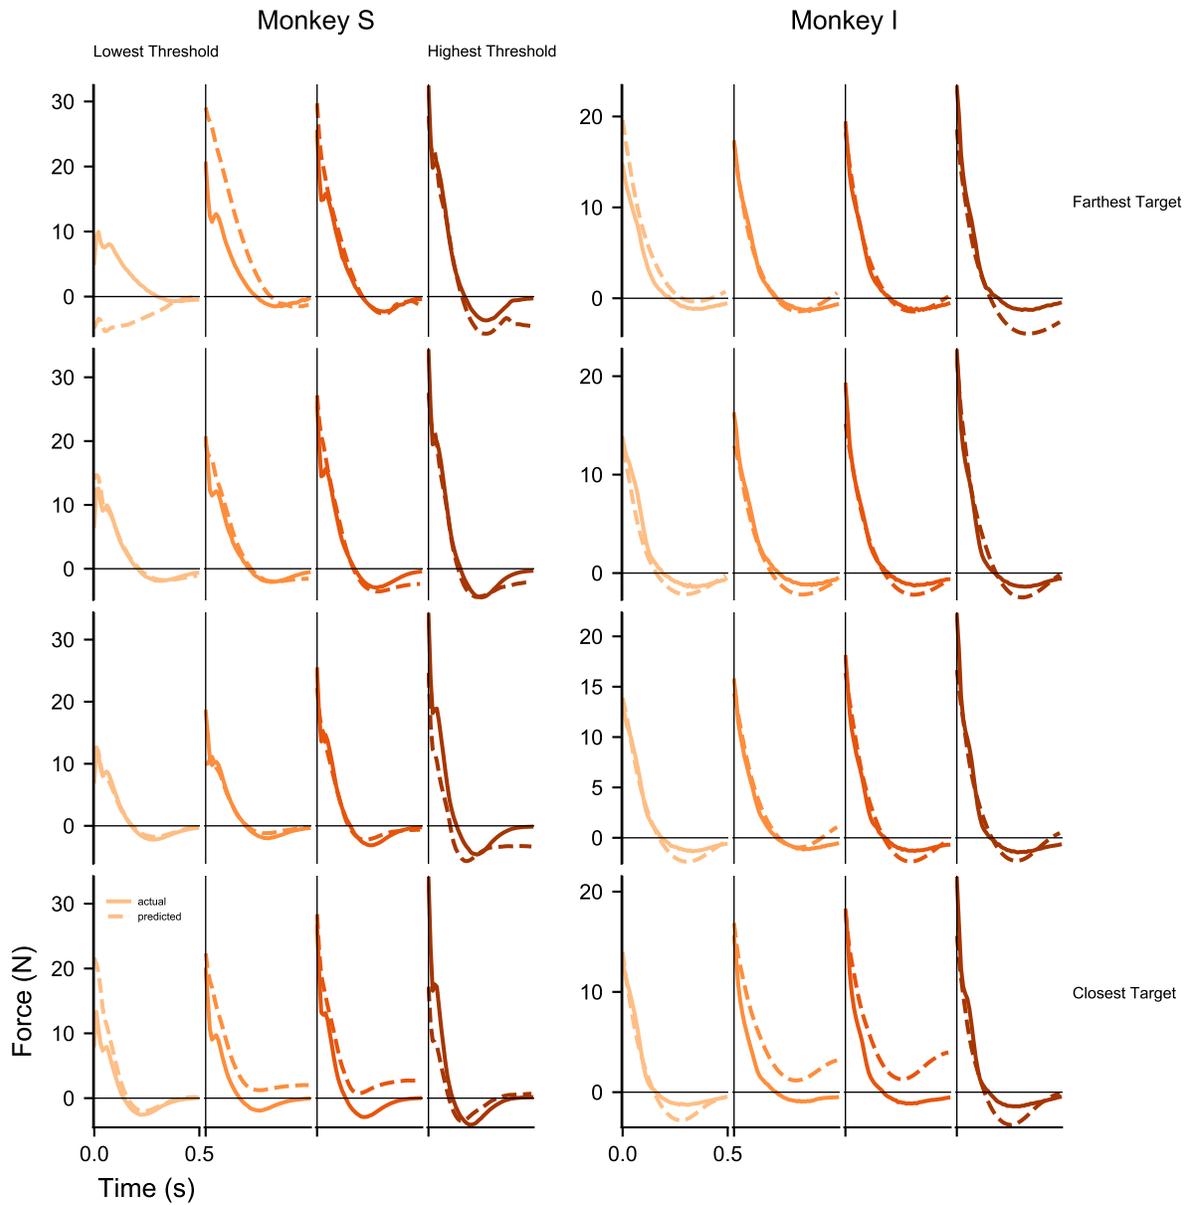


Figure 4.12: **The accuracy of the time-varying force estimated from the decoded physical parameters was consistent across task conditions.** Each panel is a single task condition. The physical parameters for a given task condition were decoded from a decoder that was trained on the other 15 task conditions. The estimated force trajectory reflected the same initial decrease, subsequent negative values, and final plateau as the actual force trajectory. Time 0 was movement onset.

## 4.5 DISCUSSION

We leveraged a behavioral paradigm that emphasized the combined control of force and motion to reveal a motion and impedance signal encoded in motor cortical activity. The behavior was modeled as a physical dynamical system and the model’s impedance and zero-force position were decoded and used within the impedance control framework to estimate a time-varying force that closely approximated the actual force exerted on the object. This approach synthesized existing neural encoding models of either force (Evarts, 1968; Kalaska et al., 1989; Georgopoulos et al., 1992; Kalaska et al., 1989; Sergio et al., 2005) or motion (Georgopoulos et al., 1982, 1986; Schwartz, 1994; Moran and Schwartz, 1999) with a successful force-motion control framework (Hogan, 1985a, 2014) to estimate the force exerted on an object by mapping motor cortical activity to the parameters of a physical dynamical system.

### 4.5.1 Ballistic-release behavior

The reported results are for a ballistic-release task that involved little to no intervention during the movement (Viviani and Terzuolo, 1973; Polit and Bizzi, 1979; Elliott et al., 1999). In impedance control, this assumes that the model’s zero-force position and impedance remained constant throughout the movement. The physical interaction between the arm and handle was described so that the motion of the handle determined the time-varying force exerted by the arm on the handle, similar to stretching a spring. These conditions encourage the musculoskeletal system to be pre-set to a state that will accomplish the task when the handle is unlocked. Although manipulations that are composed of a ballistic phase followed by a homing phase (Meyer et al., 1988) might demonstrate different patterns of neural activity, the general framework explored in this study would still be realizable (Hogan, 1985a).

Both monkeys performed the task very quickly, crossing the force threshold in less than 100 ms. This was much quicker than was observed in the human study (Chapter 3). The short duration of the force ramp could indicate that the monkey was planning its strategy

even before the handle was grasped. In this case, it would be difficult to dissociate neural activity related to reaching for the handle from neural activity related to pre-setting the musculoskeletal system. Future experiments could improve the task design by requiring the monkey to maintain force above the threshold for a random period of time. This would provide a clear time window within which to search for force-related cortical signals and could encourage a slower, more controlled behavior during the force ramp.

#### 4.5.2 Physical dynamical system

Modeling the arm as an equivalent physical dynamical system that changes according to the demands of the behavior assumes that the nervous system is capable of setting the effective impedance of the arm (Hogan, 1984c). Arm stiffness is linked to force (Perreault et al., 2002), but additional stiffness can be modulated by co-activating antagonist muscles in the arm and shoulder (Hogan, 1985a; Lacquaniti et al., 1982; Gomi and Osu, 1998). Mechanically, damping and inertia are related to stiffness through the damping ratio, which describes the transient response of the system (under-, over-, and critically damped). Correspondingly, some modulation of damping and inertia was expected to vary with large changes in stiffness. Damping is affected by muscle activity (Joyce and Rack, 1969; Rack and Westbury, 1974; Nichols and Houk, 1976) and inertia could be modulated by changing the configuration of the arm (Mussa-Ivaldi et al., 1985; Trumbower et al., 2009). However, the additional modulation of stiffness with target zones appeared to be independent of damping and inertia. These concepts generally fall within the equilibrium point hypothesis (Feldman, 1966, 1986) and its extensions (Bizzi et al., 1984; Hogan, 1985a; Flash, 1987).

#### 4.5.3 Information encoded in neural activity

We found that a motion and stiffness signal were more strongly represented in motor cortical activity than a force signal. The relatively weak force signal could be a result of the short duration of the force ramp or a consequence of the few neural units recorded from the rostral bank of the central sulcus (Kalaska et al., 1989; Rathelot and Strick, 2009). However, we did record force levels orders of magnitude higher than previous studies (Evarts, 1968; Kalaska

et al., 1989; Georgopoulos et al., 1992; Sergio and Kalaska, 2003; Sergio et al., 2005), and it may have been expected that such high force levels would have resulted in a large force signal. Ultimately, the short duration of the force ramp prevents any strong conclusions about the force signal encoded in the neural activity.

The shared variability in the population driven by force threshold and target zone were revealed by demixed PCA (dPCA). The algorithm is capable of identifying the transient variability driven by a transient stimulus, and the sustained variability driven by a sustained stimulus (Kobak et al., 2016). The shared variability driven by the target zone in our task was separated across targets at the beginning of the force ramp and sustained that separation throughout the movement. Although not definitive, this could reflect a sustained motion signal driving the motor cortical population, indicative of the movement’s ballistic nature and consistent with impedance control. It should also be noted that dPCA was performed on trial-averaged neural responses that were not recorded simultaneously. While a valid application of the algorithm, a more nuanced description of the shared variability could be achieved with trial-by-trial responses from a simultaneously recorded population.

The neural units were categorized into target, threshold, and stiffness responses according to the coefficients of the task parameter encoding model. We found that the target and stiffness responses were most prevalent in the recorded population for both monkeys. However, the category of an individual neural unit should be interpreted with care. The temporal responses were complex and it’s possible that a more nuanced temporal analysis would have yielded slightly different conclusions about individual units. Instead, our goal was to convey the general prevalence of population information that could be related to each category.

The zero-force position and impedance coefficients of the physical dynamical model were decoded from the population activity. The impedance coefficient for position, stiffness, is generally related to the co-activation of antagonist muscles (Hogan, 1984b; Gomi and Osu, 1998). Stiffening the wrist by co-activating flexor and extensor muscles was found to correlate with the activity of a sub-population of motor cortical neurons (Humphrey and Reed, 1983); the activity of this population decreased when the same muscles were reciprocally activated. A similar functional tuning of some motor cortical neurons for fixating

muscle activity was observed by [Griffin et al. \(2015\)](#). Our ability to decode the impedance coefficients is likely related to these results, but more studies are needed to determine how the different impedance coefficients might be represented within individual neurons in the motor cortex.

We should emphasize that our results describe information contained within motor cortical activity and do not claim a specific function of the motor cortex. The motor cortex is one of several neural structures influencing the spinal network that ultimately controls muscle activity and behavior ([Phillips and Porter, 1977](#); [Lemon, 2008](#)). Future work would be needed to narrow down the potential influence of motor cortical activity on muscle activity and arm configuration that mechanistically modulate the behavior of the arm and hand.

#### **4.5.4 Conclusion**

The combined control of force and motion was systematically explored using a ballistic-release paradigm and neural activity from the motor cortex was recorded, an area of the brain canonically related to both force and motion. The combination of rich behavior and neural recordings enabled us to synthesize traditional neural encoding models with an established control theory framework. We found a strong motion and stiffness signal and could decode physical system parameters that efficiently approximated the force exerted on an object. Extensions of this work will be useful for more detailed models of neural activity during object manipulation in a variety of contexts. It will also add to the versatility of brain-machine interfaces, potentially providing people who are paralyzed the ability to manipulate objects and perform many activities of daily living.

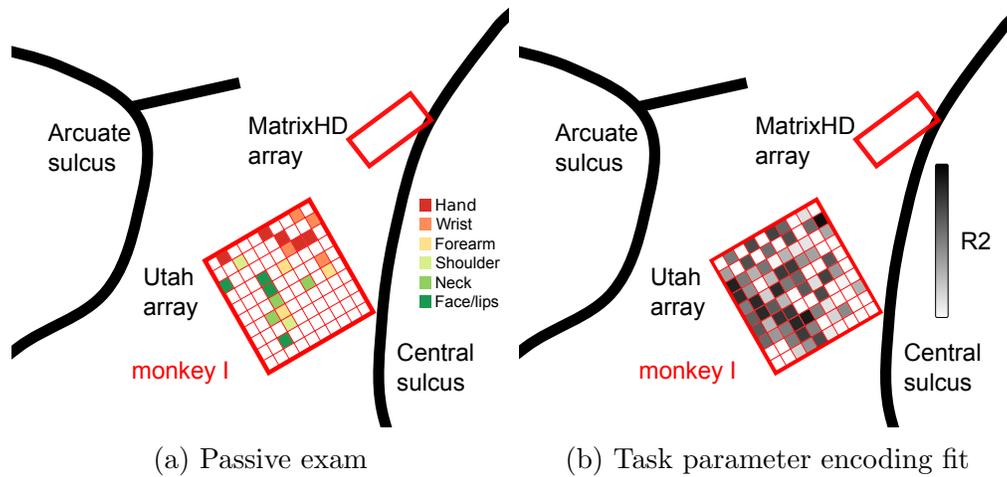


Figure 4.13: The lateral array for monkey I was sub-divided into individual recording channels. **(A)** A passive exam revealed that some channels exhibited activity related to the face and lips. However, the medial portion of the array contained channels that were related to the shoulder, arm, and hand. **(B)** For each channel, the shading represents the highest  $R^2$  on that channel from the regression for equation 4.2. The medial portion of the array, corresponding to the arm and hand area, had regression fits that were consistent with monkey S (see figure 4.9)

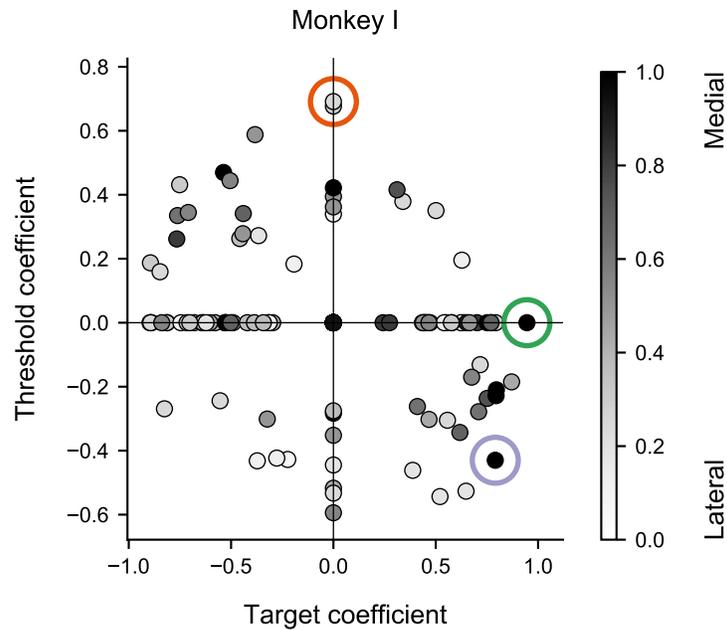


Figure 4.14: **The target and threshold coefficients did not depend on the laterality of the neural units recorded from monkey I.** Each dot represents one neural unit from monkey I and corresponds to the target and threshold coefficients from the task parameter encoding model regression in equation 4.2. The color corresponds to the laterality of the recording channel, with 1 representing the medial-most corner of the array and 0 representing the lateral-most corner of the array (see Figure 4.13).

## 5.0 MUSCLE ACTIVITY AND MOTOR CORTICAL ACTIVITY RELATED TO FORCE AND STIFFNESS DURING OBJECT MANIPULATION

### 5.1 ABSTRACT

To simultaneously control force and motion during object manipulation, muscle activity in the arm exerts force to move the object and changes stiffness to prevent unwanted motion. We explored the extent to which separate components of muscle and motor cortical activity were related to force and stiffness during a ballistic-release task that dissociated these two parameters. We found components of muscle activity and motor cortical activity related to force; different components of muscle activity and motor cortical activity were related to stiffness. A second set of analyses regressed muscle activity against motor cortical activity and revealed two strong components that accounted for most of the correlation between the two. These results support the possibility of a functional mapping between the force and stiffness components of M1 and muscle activity and related experiments could help shed light on the context-dependent nature of connectivity between motor cortex and muscles.

### 5.2 INTRODUCTION

A key aspect of manipulating an object is the control of both force and motion ([Rancourt and Hogan, 2001](#); [Burdet et al., 2001](#)). Executing combined force-motion control often involves activating muscles in the arm to simultaneously exert force on the object while setting stiffness to restrict the resulting motion ([Hogan, 1984b, 1985a,c](#)). Motor cortical (M1) activity may contain signals related to both components of muscle activity and this

signaling likely changes with the functional role of the muscles (Griffin et al., 2015). However, the extent to which these components might affect the relation between M1 activity and that of muscles during object manipulation is an open question.

Muscles can be described in terms of the force they exert on an object, where antagonist muscles exert force in opposite directions. The temporal activation pattern of antagonist muscles govern both the force exerted on the object and the stiffness of the arm. The net difference between the force produced by antagonist muscles determines the force exerted on the object: reciprocal activation exerts reciprocal force (Hoffman and Strick, 1990). In contrast, the balanced force produced by antagonist muscles determines arm stiffness: co-activation stiffens the arm (Hogan, 1984a). Experiments show that subjects can readily adopt strategies in which either type of muscle component is used. While exerting isometric force on a handle, subjects changed the activity of multiple muscles in their arm to match an instructed activity pattern across muscles (Gomi and Osu, 1998; Osu and Gomi, 1999) or to match an instructed stiffness (Perreault et al., 2002).

Because of the complex geometry of the musculoskeletal system, each muscle can be involved in multiple functional roles. The same muscle might be activated to accelerate an object in one direction, decelerate an object moving in the opposite direction, compensate for unwanted forces induced by other active muscles, and/or to stiffen the arm for stabilization. Neuronal projections from M1 contain patterns of activity that may be related to the different functional roles of recipient muscles (Griffin et al., 2015), suggesting that M1 signaling may govern the partitioning of muscle activity for these roles.

In addition to the studies of M1 activity related to the force exerted on objects, at least one study has looked at the ways M1 firing rates are related to arm stiffness. During a center-out reaching task, an artificial neural network was trained to decode muscle activity from M1 (Heliot et al., 2010). The muscle activity was then applied to a musculoskeletal model that estimated stiffness and kinematics (Kim et al., 2007). The estimations matched the actual kinematics well and provided a reasonable qualitative fit for stiffness. The framework was discussed in terms of combined force-motion control in different dynamic environments, but follow-up studies were not conducted and a direct relation between M1 and stiffness was not explored.

In perhaps the most direct study of functional components of muscle and M1 activity related to force and stiffness, monkeys were trained to control the position of the wrist against perturbing sinusoidal forces of different frequencies (Humphrey and Reed, 1983). For low frequencies, the monkeys reciprocally activated flexor and extensor muscles, exerting reciprocal force to resist the perturbations. In contrast, for high frequencies, the monkeys co-activated antagonist muscles, increasing stiffness to resist the perturbations. During both conditions, M1 activity was recorded and the activity of a specific spatial population of M1 neurons was related to the co-activation of antagonist muscles but not reciprocal activation. The specificity of this population was interpreted as a control signal in M1 which could be used to control stiffness. The functional tuning in muscle and M1 activity (Humphrey and Reed, 1983; Griffin et al., 2015) and the ability to decode muscle activity to estimate stiffness (Heliot et al., 2010) represent promising results, but more studies are needed to explore the details and limitations of these results in the context of object manipulation.

The objective of this study was to determine the extent to which separate components of muscle activity and M1 activity are related to exerting force on an object and stiffening the arm during combined force-motion control. A ballistic-release paradigm dissociated force and stiffness, making it possible to fit muscle activity and M1 activity to both parameters. Separate components of muscle activity and M1 activity were related to force and stiffness, demonstrating that signals for both types of signaling were contained in the motor cortical activity.

## 5.3 METHODS

### 5.3.1 Subjects

Two rhesus monkeys were trained to perform a ballistic-release task with their right arm. All procedures were approved by the University of Pittsburgh's Institutional Animal Care and Use Committee.

### 5.3.2 Behavioral paradigm and Experimental design

Each monkey was trained to (1) use its right hand to press a start button and then reach to grasp a handle, (2) pull on the handle with enough force to unlock it, and (3) position the handle within a specified target zone (Figure 5.1). To be successful, the monkey needed to pull with enough force to unlock the handle and hold the handle in the specified target zone for 300 ms.

This behavior is similar to stretching a spring, where one end of the spring is attached to the handle and the other end of the spring is called the zero-force position. If the handle is located at the zero-force position, the spring is relaxed and exerts no force on the handle. Stretching the spring to the target zone (displacing the zero-force position away from the handle) causes the spring to exert force on the handle according to both the displacement and the stiffness. When the handle is unlocked, the force exerted on the handle moves it to stop at the zero-force position, providing the control to move the handle to different target zones. For the same target zone and zero-force position, the force exerted on the handle could be increased by increasing stiffness. Thus, a zero-force position and stiffness could be preset for a given task condition (target zone and force threshold), preventing the need for corrective interventions during the movement. See section 2.1 for more details. Importantly, this behavior would dissociate force and stiffness. Changing the spring's zero-force position would cause different amounts of force to be exerted for a given stiffness value.

### 5.3.3 Data collection

During the task, we measured: (1) the force of the monkey pulling on the handle; (2) the 3D position of optical markers placed on the handle and the monkey's hand, lower arm, upper arm, and torso; (3) the epimysial electromyogram (EMG) of 16 individual muscles (monkey S, names in Table 6) or the surface EMG of 6 muscle groups (monkey I; wrist flexor/extensor, elbow flexor/extensor, anterior/posterior deltoid, pectoralis/rotator cuff); and (4) neural activity in the motor cortex. All data were synchronized in time and analyzed at 100 Hz. Data from individual trials were aligned on movement onset. See sections 2.2, 2.3, and 2.4 for more details.

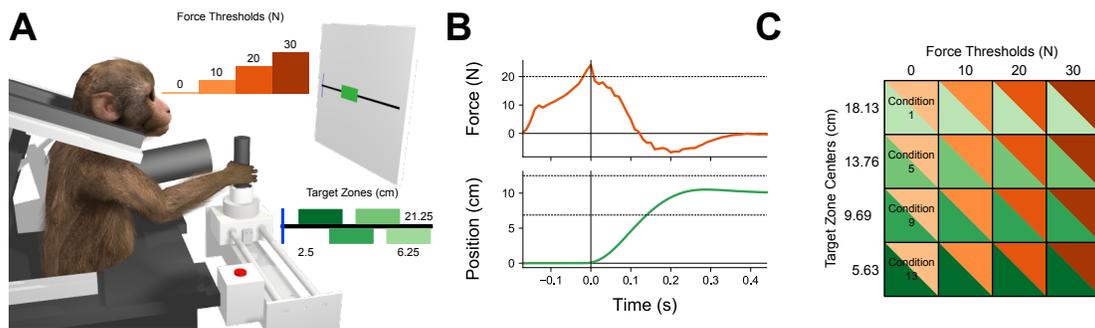


Figure 5.1: **Behavioral paradigm (A)** To initiate a trial, the monkey first pressed the start button (to the monkey’s right) and then grasped the handle. The monkey pulled on the handle while it was locked in place until the force threshold was crossed. The handle was then unlocked to move freely along the track and the monkey had to stop and hold the handle within the target zone for 300 ms. **(B)** Time-series of force and position were measured for each trial (a single representative trial is displayed here). **(C)** A single task condition was composed of a force threshold and a target zone.

Table 6: Individual muscles and their abbreviations from which epimysial EMG was recorded in monkey S.

Muscle name	Abbreviation
Pectoralis	Pec
Latissimus dorsi	Lat
Anterior deltoid	Adelt
Posterior deltoid	Pdelt
Infraspinatus	Infsp
Triceps lateral head	Tril
Triceps medial head	Trim
Biceps short head	Bics
Brachialis	Bra
Extensor digitorum communis	Edc
Extensor carpi unlaris	Ecu
Extensor carpi radialis longus	EcrL
Flexor carpi radialis	Fcr
Palmaris longus	Pl
Flexor digitorum superficialis	Fds
Abductor pollicis longus	Apl

### 5.3.4 Physical dynamical system

To estimate stiffness, we modeled the behavior of the arm as an equivalent physical dynamical system with a zero-force position and 4 impedance components that exert force proportional to position and its first 3 time derivatives (equation 2.1):

$$F(t) = K(x_0 - x(t)) - D\dot{x}(t) - A\ddot{x}(t) - J\ddot{\ddot{x}}(t),$$

where  $F$  is the force exerted on the handle ;  $x, \dot{x}, \ddot{x}, \ddot{\ddot{x}}$  is position and its first 3 time derivatives;  $x_0$  is the zero-force position, and  $K, D, A, J$  are impedance coefficients. Although this model describes the arm's behavior, it does not directly measure arm impedance, and the first- and second-order impedance coefficients are related to, but not measurements of, damping and inertia. The model was fit to the trial-averaged data for each task condition during the first 500 ms after movement onset.

The model maps a motion signal to the force exerted on the object. The zero-force position is the reference position for the model and governs where the handle would eventually come to rest. In this task, the zero-force position would be related to the target zone. Changing the impedance coefficients for a given zero-force position (target zone) would change the force exerted on the handle. See section 2.5 for more details.

### 5.3.5 Ridge regression

In this study, multiple linear regression was used to relate components of muscle activity and M1 activity to force and stiffness. However, performing multiple linear regression with many independent variables (number of muscles or neural units) has a tendency to over fit noise in the dependent variable (force or stiffness). This tendency can be reduced by adding a regularization term that penalizes the magnitude of the coefficients. A large penalty shrinks the coefficients and emphasizes the bias. A small penalty permits large coefficients that reduce error but tend to fit noise (variance) in the dependent variable. Ridge regression is a form of regularized regression that balances this bias-variance trade-off.

Linear regression fits the independent variables to the dependent variable by minimizing error according to

$$\sum_{i=1}^n (Y_i - \hat{Y}_i)^2 = \min_{\beta_0^*, W^*} \left( \sum_{i=1}^n (Y_i - (\beta_0^* + W^* \cdot X_i))^2 \right)$$

$$\hat{Y} = \beta_0 + W \cdot X$$

where  $Y$  is the dependent variable,  $\hat{Y}$  is the fitted value,  $X$  is a  $[P \times 1]$  matrix of the independent variables ( $P$  is the number of variables),  $W$  is a  $[1 \times P]$  coefficient matrix,  $\beta_0$  is a constant offset, and  $n$  is the number of observations.

Ridge regression adds an extra penalty for the coefficient magnitudes so that the minimization becomes

$$\sum_{i=1}^n (Y_i - \hat{Y}_i)^2 = \min_{\beta_0^*, W^*} \left( \sum_{i=1}^n (Y_i - (\beta_0^* + W^* \cdot X_i))^2 + \alpha \sum_{j=1}^P W_j^{*2} \right)$$

where  $\alpha$  is the regularization constant.

**5.3.5.1 Regressing force on muscle and neural activity** Ridge regression was used to find the regression of force on muscle activity during the first 200 ms after movement onset according to

$$F(t) = \beta_0 + W_F \cdot EMG(t), \quad (5.1)$$

where  $F$  is a  $[1 \times T]$  matrix of force,  $EMG$  is a  $[M \times T]$  matrix of muscle activity,  $W_F$  is a  $[1 \times M]$  coefficient matrix, and  $\beta_0$  is a constant offset ( $M$  is the number of muscles and  $T$  is the number of 10 ms time bins stacked across task conditions). The regularization constant  $\alpha$  was determined by minimizing the prediction error using leave-one-condition-out cross validation.

The same procedure was used to find the regression of force on firing rates during the first 200 ms after movement onset according to

$$F(t) = \beta_0 + W_F \cdot FR(t), \quad (5.2)$$

where  $FR$  is a  $[R \times T]$  matrix of firing rates,  $W_F$  is a  $[1 \times R]$  coefficient matrix, and  $\beta_0$  is a constant offset ( $R$  is the number of neural units).

### 5.3.6 Output-potent and output-null dimensions in linear regression

Although position constraints can dissociate force and stiffness, the two are often related. We wanted to find separate components of muscle and M1 activity that were related to force and stiffness. Toward that end, we leveraged the output-potent and output-null properties of multiple linear regression. In linear regression with multiple independent  $X$  variables and a single dependent  $Y$  variable, there is a single direction, or component, of  $X$  that best correlates with  $Y$ . This direction in the multi-dimensional space spanned by  $X$  can be considered an output-potent axis and  $X$  variability along this axis, which we call  $X_{output-potent}$ , maps to variability in  $Y$ . In contrast, directions orthogonal to the output-potent axis are considered output-null, and  $X$  variability along these axes, which we call  $X_{output-null}$ , does not map to variability in  $Y$ .

**5.3.6.1 Force-potent and stiffness-potent muscle activity** In Chapter 3 we considered the mapping from muscle activity (EMG) to force exerted on the handle. The output-potent axis described EMG variability that mapped to force variability, which we call force-potent EMG. The output-null axis described EMG variability that did not map to force variability, which we call force-null EMG. Additionally, we found that force-null EMG was related to stiffness.

Here we again find force-potent EMG and force-null EMG. Force-potent EMG is a [1 x T] matrix defined as

$$EMG_{F-potent}(t) = V_{F-potent} \cdot EMG(t)$$

where  $EMG$  is a [M x T] matrix of muscle activity (M is the number of muscles and T is the number of 10 ms time bins stacked across task conditions) and the force-potent axis  $V_{F-potent}$  is a [1 x M] vector from the first row of the [M x M] matrix  $V$ , defined as

$$U, \Sigma, V = \text{SVD}(W_F)$$

where  $W_F$  is the coefficient matrix mapping EMG to force in equation 5.1 and SVD is singular value decomposition. Because the force-potent EMG is simply as scaled version of the fitted values from equation 5.1, we refer to force-potent EMG and the fitted values interchangeably. Force-null EMG is a [M-1 x T] matrix defined as  $EMG_{F-null}(t) = V_{F-null} \cdot EMG(t)$  where the force-null axes  $V_{F-null}$  are a [M-1 x M] matrix from rows 2 through M of  $V$ .

We explicitly find the force-null axis that best correlates with model stiffness and call the [1 x C] matrix of force-null EMG variability along this axis stiffness-potent EMG:

$$EMG_{K-potent} = V_{K-potent} \cdot EMG_{F-null}$$

where  $EMG_{F-null}$  is a [M-1 x T] matrix of force-null EMG, and the stiffness-potent axis  $V_{K-potent}$  is a [1 x M-1] matrix from the first row of the [M-1 x M-1] matrix  $V$ , defined as

$$U, \Sigma, V = \text{SVD}(W_K)$$

where  $W_K$  is the coefficient matrix mapping force-null EMG to stiffness in

$$K = \beta_0 + W_K \cdot EMG_{F-null} \quad (5.3)$$

where  $K$  is a [1 x T] matrix of model stiffness values, with repeated values for each of the 10 ms time bins in a task condition, and  $\beta_0$  is a constant offset. Again, the stiffness-potent EMG is simply as scaled version of the fitted values from equation 5.3, so we refer to stiffness-potent EMG and the fitted values interchangeably.

In summary, we regressed force against EMG. Based on the assumption that variability in muscle activity in the absence of variability in force would be related to the model stiffness in equation 2.1, we performed an analysis that separated combinations of EMG activity that were related to the absence of variability in force (force-null EMG) from the single combination that best corresponded to the variability in force (force-potent EMG). Since only a subset of force-null EMG was likely to be related to stiffness, we used the same procedure to find the single component in the subset best related to variability in stiffness (stiffness-potent EMG).

**5.3.6.2 Force-potent and stiffness-potent neural activity** We apply this same procedure to motor cortical activity, considering the mapping from M1 firing rates (FR) to force exerted on the handle. Force-potent FR was a  $[1 \times T]$  matrix defined as

$$FR_{F-potent}(t) = V_{F-potent} \cdot FR(t)$$

where  $FR$  is a  $[R \times T]$  matrix of motor cortical firing rates ( $R$  is the number of neural units) and the force-potent axis  $V_{F-potent}$  is a  $[1 \times R]$  matrix from the first row of the  $[R \times R]$  matrix  $V$ , defined as

$$U, \Sigma, V = \text{SVD}(W_F)$$

where  $W_F$  is the coefficient matrix mapping FR to F in equation 5.2. Force-potent FR and the fitted values from equation 5.2 are scaled versions of each other and used interchangeably. Force-null FR is a  $[R-1 \times T]$  matrix defined as

$$FR_{F-null}(t) = V_{F-null} \cdot FR(t)$$

where the force-null axes  $V_{F-null}$  are a  $[R-1 \times R]$  matrix from rows 2 through  $R$  of  $V$ .

We find the force-null axis that best correlates with model stiffness and called the  $[1 \times C]$  matrix of force-null FR variability along this axis stiffness-potent FR:

$$FR_{K-potent} = V_{K-potent} \cdot FR_{F-null}$$

where  $FR_{F-null}$  is a  $[R-1 \times T]$  matrix of force-null FR, and the stiffness-potent axis  $V_{K-potent}$  is a  $[1 \times R-1]$  matrix from the first row of  $[R \times R]$  matrix  $V$ , defined as

$$U, \Sigma, V = \text{SVD}(W_K)$$

where  $W_K$  is the coefficient matrix mapping force-null EMG to stiffness in

$$K = \beta_0 + W_K \cdot FR_{F-null} \tag{5.4}$$

where  $K$  is a  $[1 \times T]$  matrix of model stiffness values, with repeated values for each of the 10 ms time bins in a task condition, and  $\beta_0$  is a constant offset. Stiffness-potent FR and the fitted values from equation 5.4 are scaled versions of each other and used interchangeably.

### 5.3.7 Reduced-rank regression

We examined the possibility that there might be a structured mapping of information transmitted from M1 to the muscles of the arm. Specifically, we were interested in whether M1 could signal a pattern of muscle activity related to exerting force or stiffening the arm. Although the previous analysis would reveal partial correlations to force or stiffness present in muscle and M1 activity, those correlations might account for very little of the total correlation between the two signals. An alternative would be to begin with a total correlation between muscle and M1 activity and then determine if that correlation could be pared down to functional components.

The total correlation between M1 activity and muscle activity can be found using linear regression and then systematically restricted to reduce the number of components that map M1 variability to muscle variability. This approach, called reduced-rank linear regression, is similar to canonical correlation analysis but has the added benefit of being able to reduce the number of mapping components and then calculating the cross-validated prediction error for different numbers of components.

Reduced-rank regression is a variation of traditional linear regression that conditions the coefficient matrix to have a certain number of components, or matrix rank. In the context of this task, traditional linear regression uses the full number of correlated components (a full-rank coefficient matrix) to map motor cortical firing rates (FR) to muscle activity (EMG):

$$EMG(t) = \beta_0 + W \cdot FR(t) \quad (5.5)$$

where  $EMG$  is a  $[M \times T]$  matrix of muscle activity,  $FR$  is a  $[R \times T]$  matrix of motor cortical firing rates,  $W$  is a full-rank  $[M \times R]$  coefficient matrix, and  $\beta_0$  is a  $[M \times 1]$  matrix of constant offsets. The rank of  $W$  describes the independent components, or communication components, and is equal to the minimum dimensions in FR and EMG. In this case  $M < R$  and the full-rank version of  $W$  describes  $M$  communication components.

Reduced-rank regression assumes that a small number of communication components can explain a relatively large percentage of the variance. The number of communication components, governed by the rank of  $W$ , can be manipulated by pre-multiplying  $W$  by another coefficient matrix,  $U$ , where  $U$  is a  $[M \times M]$  matrix with rank  $K \leq M$ .

$$EMG(t) = \beta_0 + U \cdot W \cdot FR(t) \quad (5.6)$$

$U$  is defined as

$$\sum_{k=1}^K (u_k^T \cdot u_k)$$

where  $u_k$  is a  $[1 \times M]$  eigenvector, sorted so that  $u_1$  explains the most variance of

$$\Sigma_{mr} \cdot \Sigma_{rr}^{-1} \cdot \Sigma_{rm}$$

where  $\Sigma_{mr}$  is the  $[M \times R]$  cross-covariance matrix between EMG and FR,  $\Sigma_{rr}$  is the  $[R \times R]$  covariance matrix of FR, and  $\Sigma_{rm}$  is the  $[R \times M]$  cross-covariance matrix between FR and EMG. We evaluated the reduced-rank regression using the prediction error of equation 5.6 using leave-one-condition-out cross-validation for each rank  $K = \{1, \dots, M\}$ .

## 5.4 RESULTS

We were interested in muscle activity and motor cortical activity related to the force exerted on an object and the stiffness of the arm during object manipulation. We trained two monkeys to control force and position during a ballistic-release task and recorded the force exerted on a handle, the position of the handle, muscle activity, and motor cortical activity from the same experimental sessions as Chapter 4. To briefly summarize, the data were recorded across 11 days for each monkey: monkey S completed 4030 successful trials and monkey I completed 2347 successful trials. Trial data were aligned on movement onset (time 0) and averaged across successful trials for each task condition. Bootstrapped confidence intervals were calculated by re-sampling the data 1000 times, with replacement, from the successful trials per condition.

#### 5.4.1 Force and position varied with task conditions

The monkeys exerted enough force on the handle to cross the force threshold and then moved and held the handle within the target zone (Figure 5.2). Force and position were dissociated: a given target zone evoked similar position trajectories across force thresholds, a given force threshold evoked similar force trajectories across target zones.

Both monkeys pulled on the handle quickly, with monkey S crossing the force threshold in 30 ms on average and monkey I crossing the force threshold in 90 ms on average. Both monkeys increased the exerted force rapidly to the threshold with no apparent anticipation of threshold crossing, suggesting that the monkeys had pre-planned a behavioral strategy before exerting force on the handle.

The monkeys moved the handle to the hold position smoothly and typically overshoot the hold position by a small distance, staying within the target zone, suggestive of a slightly under-damped system. Task conditions with near targets or high thresholds were associated with larger overshoot, indicating a change in the impedance coefficients across task conditions.

#### 5.4.2 Model stiffness depended on both force threshold and target zone

Because force is often related to stiffness (McIntyre et al., 1996; Perreault et al., 2002), it was important to dissociate these variables before they were related to muscle and motor cortical activity. We modeled the arm as a physical dynamical system (equation 2.1) consisting of 5 free parameters: a zero-force position and 4 impedance coefficients corresponding to position and its first 3 time derivatives. The force was fit to the motion during a 500 ms time window, beginning at movement onset (equation 2.1). The zero-force position was where the handle would come to rest according to the physical model and co-varied with target zone (RMSE = 0.37 cm, 0.97 cm, monkey S and I, Figure 5.3, top row). Stiffness was the model's impedance coefficient corresponding to position and, for a given zero-force position (target zone), model stiffness co-varied with force threshold. Importantly, the same value of stiffness was found for different force thresholds by varying the target zone. For example, for monkey S, a similar

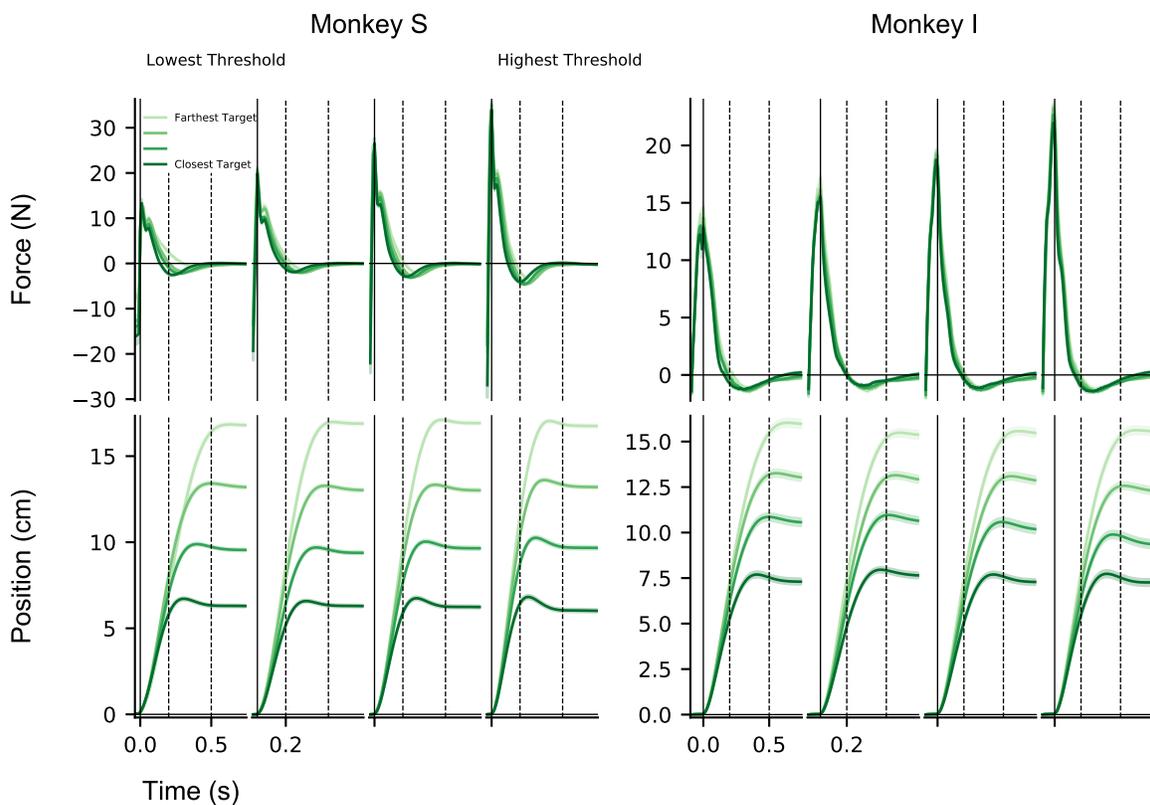


Figure 5.2: **Force and position were dissociated across task conditions.** For a given target zone, position trajectories were similar across force thresholds. For a given force threshold, force trajectories were similar across target zones. Time 0 is movement onset and shading represents the median and 95% confidence interval.

stiffness value was observed for the 10 N force threshold and closest target zone, the 20 N force threshold and second closest target zone, and the 30 N force threshold and second farthest target zone. See chapter 4 for more detailed results.

### 5.4.3 Muscle activity was modulated with force threshold and target zone

Muscle activity was recorded as the epimysial electromyogram (EMG) from 16 muscles in monkey S and as the surface EMG from 6 muscle groups in monkey I (Figures 5.4 and 5.5). Across all muscles, the EMG values started at elevated levels and increased to a maximum value around movement onset (time 0, solid vertical line). The initial elevated values were likely because of the rapid movement from the start button to the handle. After movement onset, the EMG values decreased to sub-initial values, indicating a more relaxed muscle state toward the end of movement than at the beginning of the force ramp.

For monkey S, maximum EMG values across time were related primarily to force threshold and, to a lesser extent, target zone (Figure 5.4, muscle name abbreviations in Table 6). Across most muscles, the EMG values at movement onset were strongly modulated with force threshold, but converged to similar values 200 ms after movement onset, verifying that the recorded muscles were active in this task and likely used to exert force.

EMG patterns clearly changed across different target zones. Some muscles (Pec, Bra, Bics, Tril, Ecu, Fds, Pl), changed before movement onset, supporting the idea that a movement strategy was pre-planned. Shortly after movement onset, the decreasing EMG values of some muscles depended on target zone (Pec, Lat, Adelt, Pdelt, Infsp, Trim, Edc, Ecl, Fcr, Pl, Apl). This activity was most evident 25-50 ms after movement onset (shaded region in figure 5.4), a latency consistent with spinal feedback responses. Additionally, the EMG values for most of these muscles decreased less quickly over time for farther targets (excluding Pec and Pl), suggesting that, even though EMG was decreasing, the muscles were being used to propel the handle toward the target.

Two muscles exhibited an especially strong target relation sustained across the 200 ms after movement onset (Pec and Pl). Further examination showed that this muscle activity returned to a baseline level approximately 300-400 ms after movement onset (Supplementary

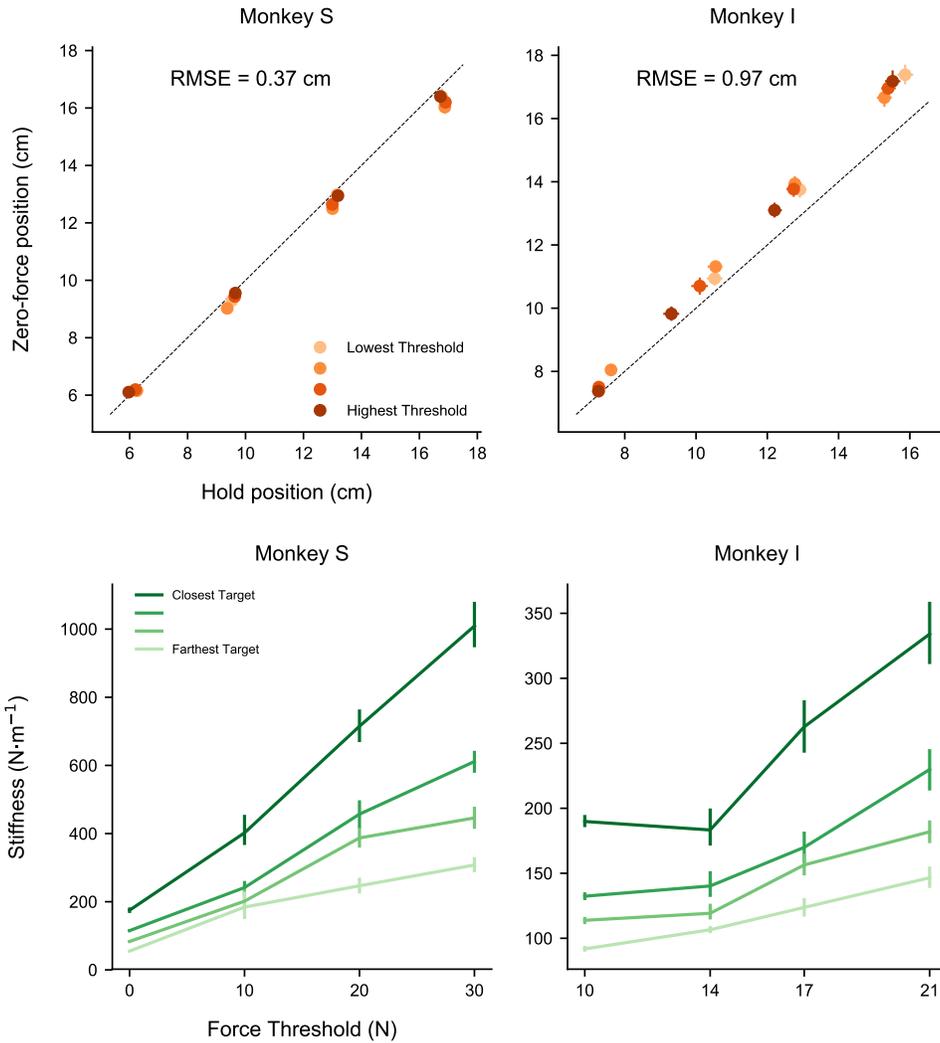


Figure 5.3: **The zero-force position and model stiffness were related to the target zone and force threshold.** (Top row) The zero-force position described where the model would come to rest and matched where the monkey stopped and held the handle (hold position). (Bottom row) Similar values for model stiffness were associated with different combinations of force threshold and target zone, dissociating force and stiffness. Error bars represent the median and 95% confidence interval.

Figure 5.11), near the end of movement (Figure 5.2). The timing suggests that these muscles were involved in stopping the handle in the target zone. This was further supported by the higher EMG values over time for closer target zones, task conditions that required increased force to stop the handle quickly.

For monkey I, the effects of task condition on EMG was similar to monkey S, but less pronounced (Figure 5.5). The main differences were the lack of strong EMG modulation across time and slightly less EMG variability across target zones. Although it's possible that this could have been because of different behavioral strategies employed by the two monkeys, it seems likely that a large contributor was the lower spatial specificity of the surface EMG electrodes for monkey I compared with the epimysial EMG electrodes for monkey S.

The epimysial EMG recorded from monkey S was modulated more across time and was more consistent across trials compared with the surface EMG recorded from monkey I. The distance between the muscle and the surface electrode, as well as the intervening tissue, makes it likely that the signal recorded by the surface EMG represented the combined activity of many muscles and may have contributed to the reduced modulation across time. Additionally, the surface EMG electrodes were replaced each recording session and small deviations across sessions in the electrode placement and the electrical conductivity of the tissue/electrode interface may have contributed to the inconsistencies across trials. Although these considerations limit the interpretation of the present results for monkey I, the strong epimysial EMG recorded from monkey S suggests that future experiments would be better served by epimysial or percutaneous EMG recordings.

#### **5.4.4 Force-potent and stiffness-potent muscle activity**

The modulation of EMG and the dissociation of force and stiffness made it possible to find separate components of EMG related to force and stiffness. We used ridge regression to fit EMG to force during the first 200 ms after movement onset (equation 5.1) and found a good fit across task conditions for both monkeys ( $R^2 = 0.90, 0.82$ , monkey S and I, Figure 5.6). The force-potent EMG captured the rapid increase before movement onset (time 0), the decrease after movement onset, and the modulation of maximum force with force threshold.

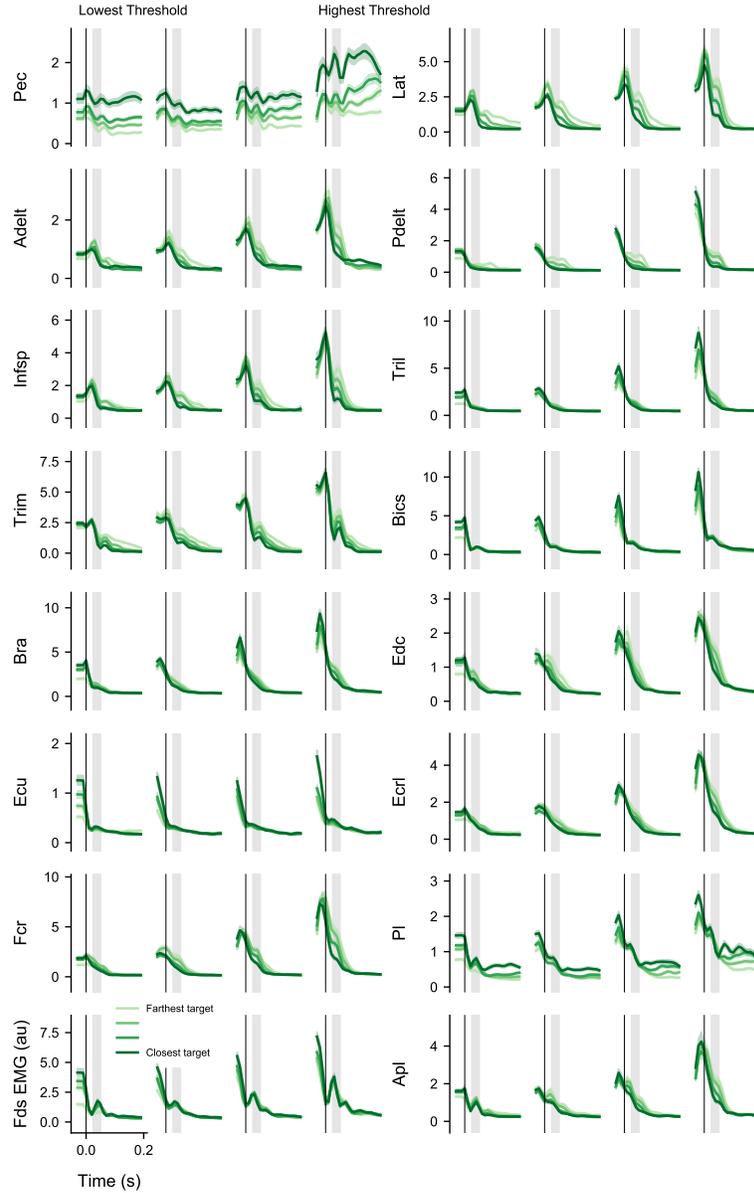


Figure 5.4: **Time-varying epimysial EMG across task conditions.** Each plot depicts four target zones with the same force threshold. The EMG values at movement onset (time 0) were strongly correlated with force threshold. Some muscles exhibited target-dependent responses 25-50 ms after movement onset (grey shaded region). Color shading represents the median and 95% confidence interval.

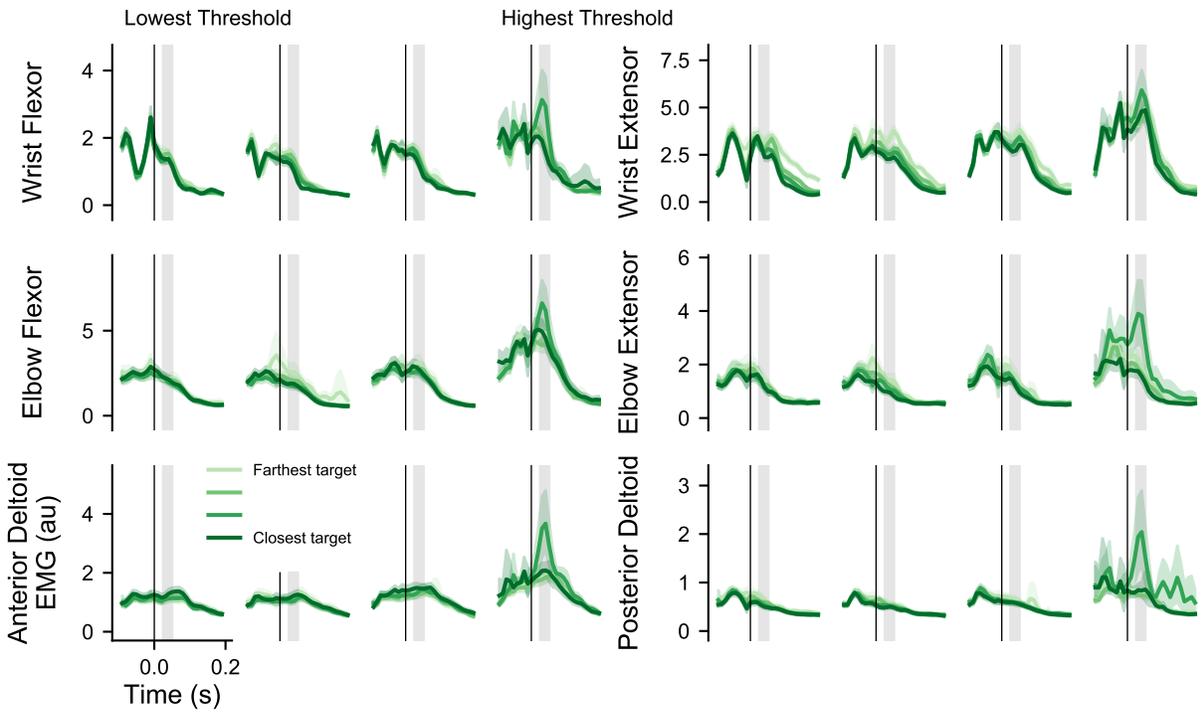


Figure 5.5: **Time-varying surface EMG across task conditions.** Same structure as figure 5.4. Qualitative results were similar to monkey S, although less pronounced, likely because of the different EMG recording techniques.

For monkey I's closest target, the maximum force at movement onset did not match the force-potent EMG as well as for farther targets, particularly for the higher force thresholds. The closest target had high values of stiffness that may have been associated with additional muscle activity from unrecorded muscles (Figures 5.3 and 5.5). The coarse spatial specificity of surface EMG combined with the potential overall increase in muscle activity in the arm could have contributed to a force-potent EMG for these task conditions that was noisier than the other task conditions, supported by the higher signal-to-noise ratio in the epimysial EMG recordings from monkey S and the better match.

The greater number of EMG dimensions (16 muscles and 6 muscle groups, monkey S and I) compared with force dimensions (only one along the track) meant that the mapping from EMG to force was redundant. Mathematically, the redundancy describes a force-potent component of muscle activity whose variability mapped to force variability (equation 5.1) and a force-null component of muscle activity whose variability did not map to force variability (redundant variability). Conceptually, the force-potent component is similar to reciprocal muscle activation and the force-null components are similar to co-activation of antagonist muscles.

Based on previous studies (Chapter 3), we suspected that force-null EMG would be related to stiffness. We regressed the force-null EMG during the first 200 ms after movement onset against model stiffness (equation 5.3) and found a good fit for monkey S ( $R^2 = 0.78$ ). The average of the stiffness-potent EMG across time captured stiffness variance across all task conditions for monkey S (Figure 5.7). It should be noted that this regression assumed a constant model stiffness value across time, and the goodness of fit indicates that this assumption was reasonable. However, the fit between force-null EMG and stiffness for monkey I was not as good ( $R^2 = 0.26$ ). The inferior fit was most likely due to the inferior EMG signal recorded from the surface EMG electrodes. Although not as likely, it may also have been because the assumption of a constant stiffness value across time may have been less appropriate for monkey I.

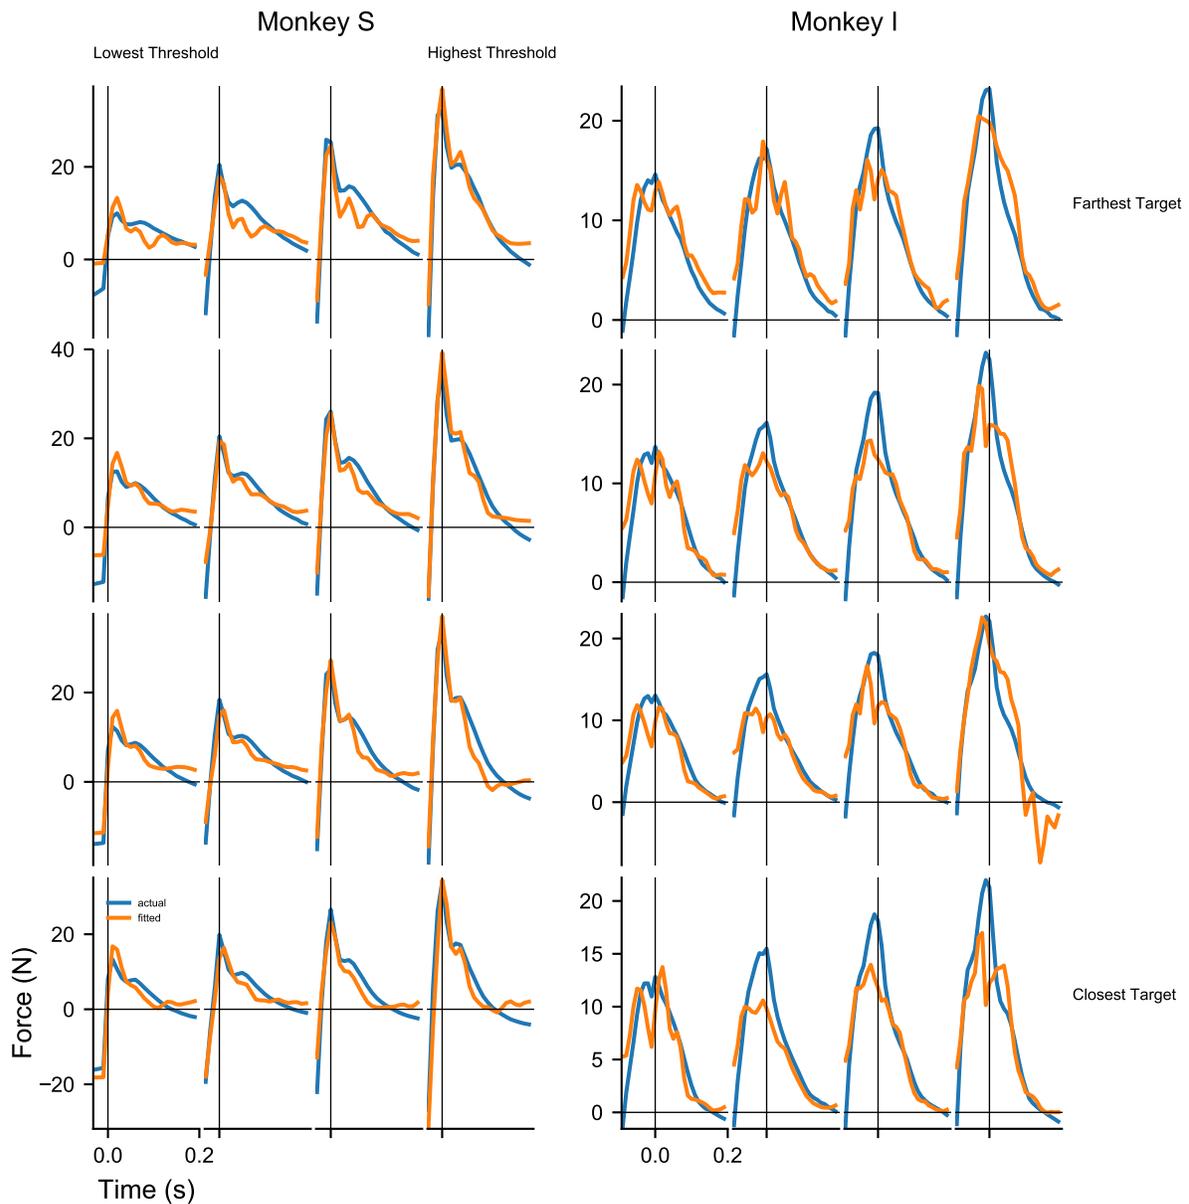


Figure 5.6: **Force-potent EMG fit force across time and task conditions.** Ridge regression was used to fit EMG to force ( $R^2 = 0.90, 0.82$ , monkey S and I). Force-potent EMG captured the temporal profile of the actual force and the modulation of the maximum force with force threshold at movement onset (time 0).

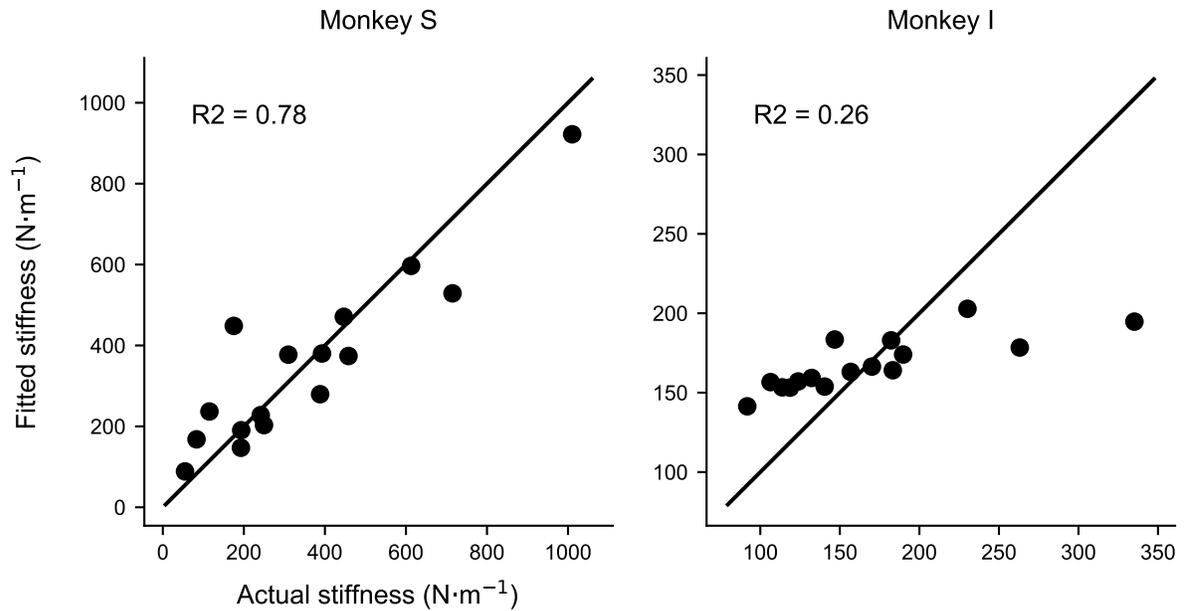


Figure 5.7: **Time-averaged stiffness-potent EMG explained the variance in model stiffness.** Force-null EMG was fit to model stiffness using ridge regression and the average stiffness-potent EMG across time matched stiffness well across all values for monkey S, but did not fit stiffness as well for monkey I. The difference was likely because of the superior quality of the EMG signals for monkey S.

Force-potent EMG and stiffness-potent EMG described separate components of muscle activity correlated with the force exerted on the handle and model stiffness. These components are concurrent and reflect muscle activity modulated to exert force and stiffen the arm, two important factors for the control needed to perform this task.

#### 5.4.5 Force-potent and stiffness-potent motor cortical activity

We were interested in the analogous possibility of separate components of motor cortical (M1) activity that could correspond to force and model stiffness. Two micro-electrode arrays were implanted in the arm/hand region of motor cortex. The data were previously reported in Chapter 4. The fractional interval firing rates were calculated from 213 neural units for monkey S and 127 neural units for monkey I. Each unit was tracked across days and firing rates were averaged across successful trials per condition. The same procedure used to identify force-potent and stiffness-potent EMG was used to identify force-potent and stiffness-potent firing rates during the first 200 ms after movement onset.

M1 firing rates (FR) explained much of the variance in force during the first 200 ms after movement onset (equation 5.2) for both monkeys ( $R^2 = 0.77, 0.90$ , monkey S and I, Figure 5.8). Except for the sharp peak values, the force-potent FR matched the actual force profiles quite well. The force-potent EMG matched the actual force better, and may mean that the FR signaling of force (at least for these samples) is less informative than that for EMG.

The force-null FR was regressed against model stiffness (equation 5.4) and explained much of the stiffness variance ( $R^2 = 0.99, 0.94$ , monkey S and I). The average of the stiffness-potent FR across time was a very close match for the stiffness values (Figure 5.9), likely because of the large number of neural units in the regression. Again, the model stiffness values were assumed constant across time and monkey I's goodness of fit suggests that this assumption was reasonable and the poor fit for stiffness-potent EMG was because of the poor quality of the EMG signal.

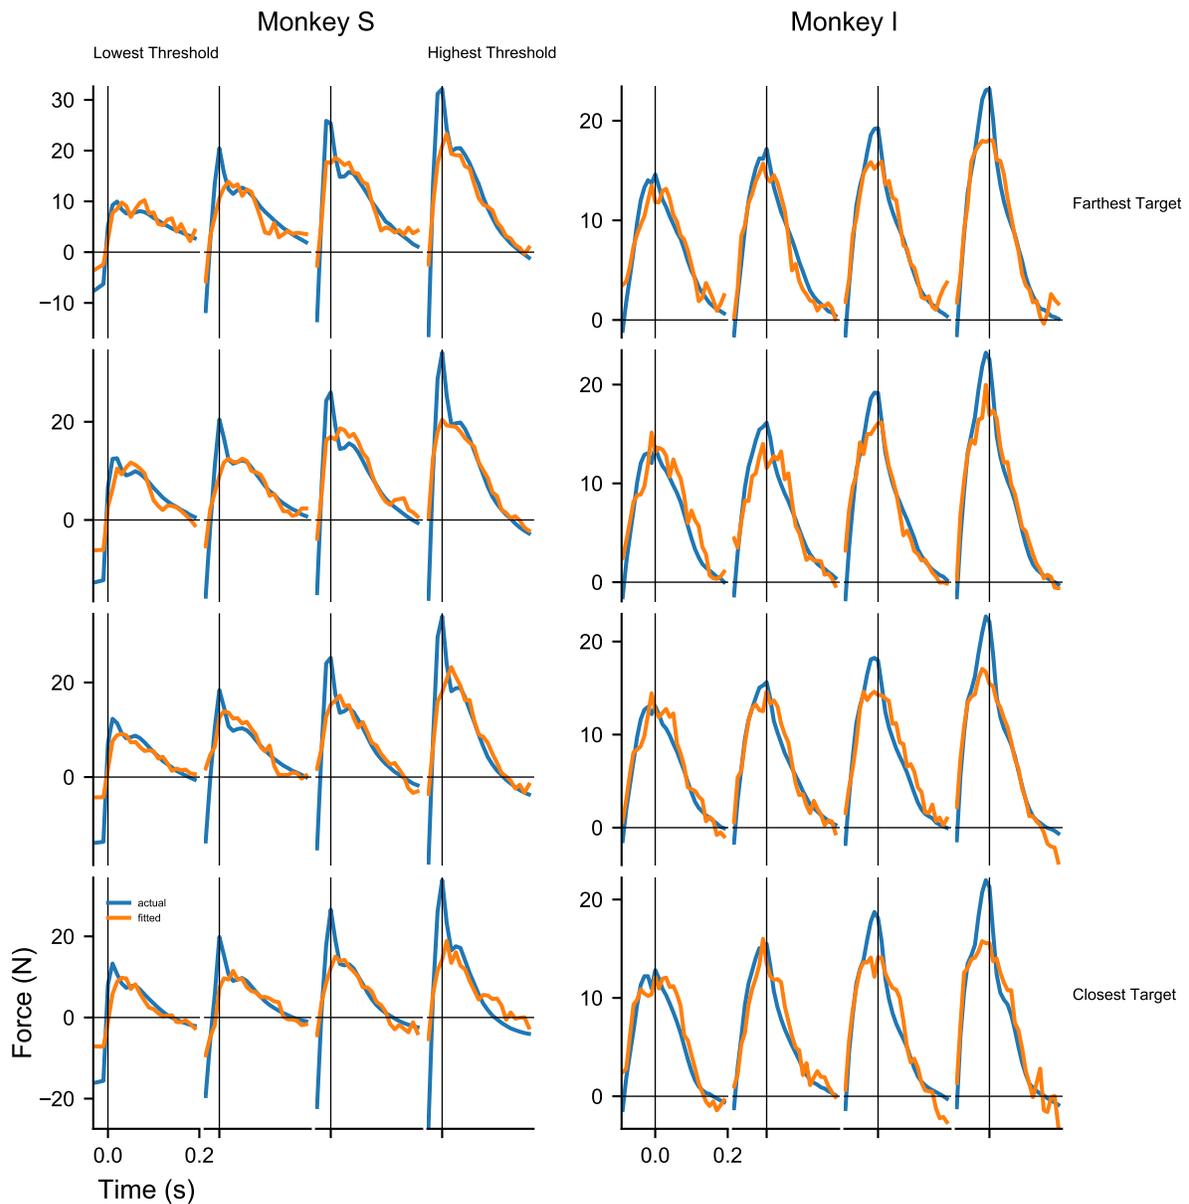


Figure 5.8: **Force-potent FR fit force across task conditions.** Ridge regression was used to fit FR to force ( $R^2 = 0.77, 0.90$ , monkey S and I). Force-potent FR captured the temporal profile of the actual force, but was unable to capture the sharp peak in force at movement onset (time 0).

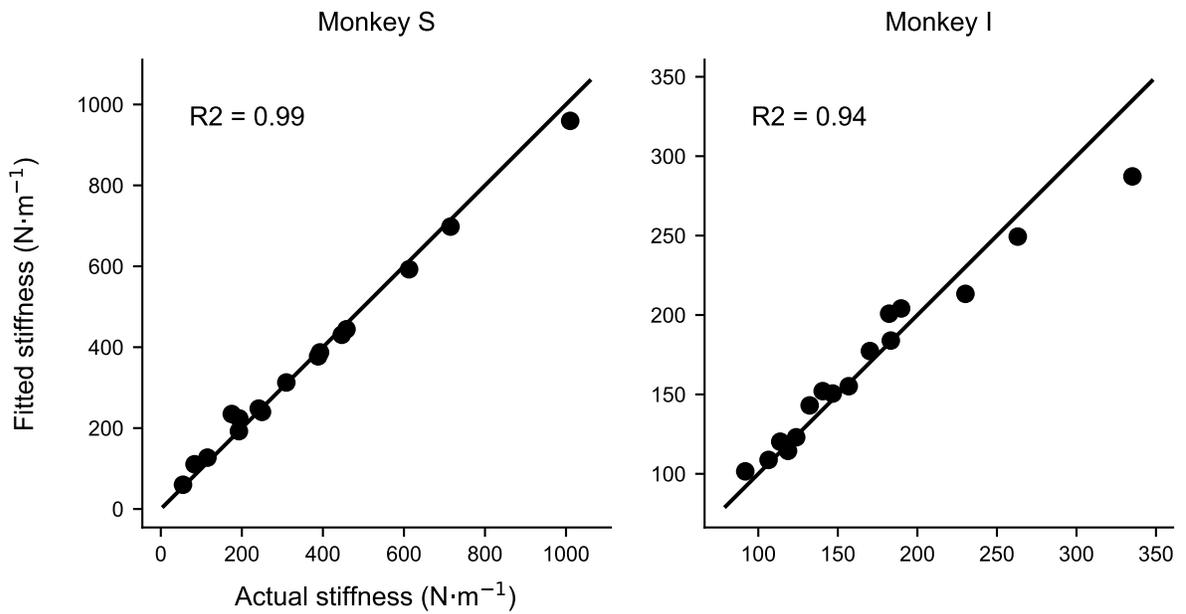


Figure 5.9: **Time-averaged stiffness-potent FR explained the variance in model stiffness.** Force-null FR was fit to model stiffness using ridge regression and the average stiffness-potent FR across time matched stiffness very well across all values for monkey S, likely because of the large number of neural units in the regression.

#### 5.4.6 Communication components between motor cortex and muscles

The possibility of separate components of muscle and motor cortical activity dedicated to exerting force and stiffening the arm suggests distinct communication components between M1 firing rates (FR) and muscle activity (EMG). We define a communication component as correlated components in FR and EMG. To investigate the number of communication components, we first established a functional relationship between EMG and FR during the first 200 ms after movement onset using ridge regression (equation 5.5) and found that FR could explain much of the EMG variance ( $R^2 = 0.93, 0.87$ , monkey S and I). We then used reduced-rank regression to test different numbers of communication components between FR and EMG (equation 5.6) and found evidence for at least two communication components (Figure 5.10). A coefficient matrix of rank two mapped two components of M1 activity to two components of muscle activity and closely approximated the prediction error of a full-rank coefficient matrix (16 or 6 components, monkey S and I) (Figure 5.10). Although we were unable to characterize the correlated components of M1 and muscle activity, this result supports the possibility of functional components of M1 activity mapping to functional components of muscle activity.

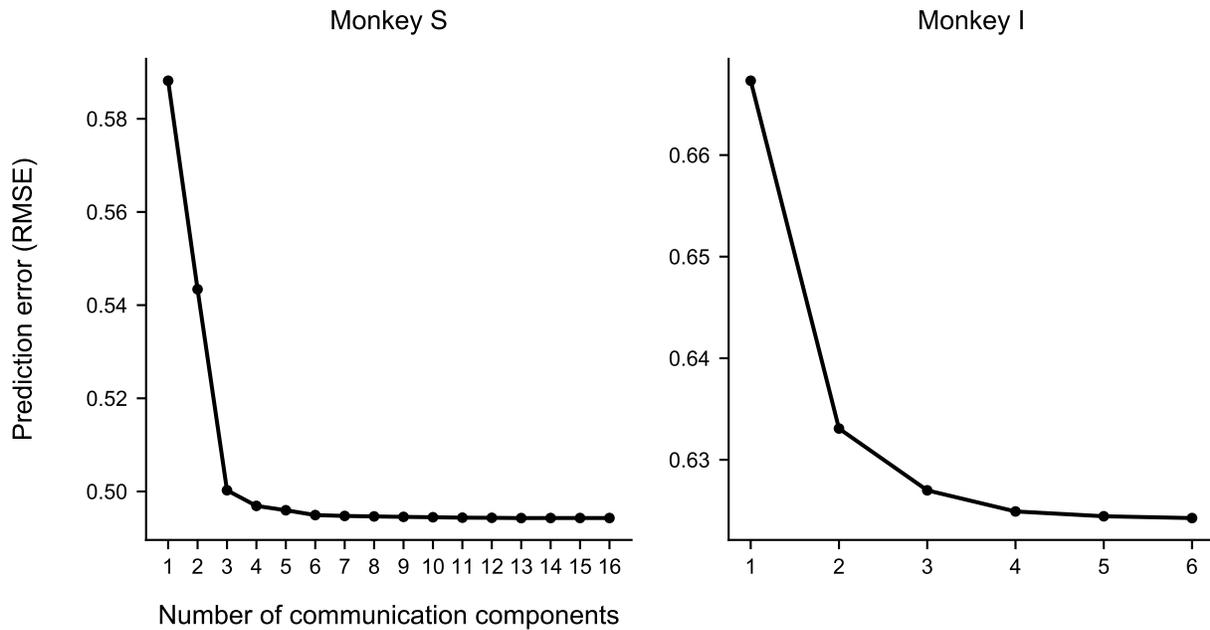


Figure 5.10: **Two communication components approximated the prediction error of the full number of components.** A communication component represents an independent dimension in the coefficient matrix mapping FR to EMG. We calculated the cross-validated prediction error for different numbers of communication components (rank of the coefficient matrix) and found that two communication components approximated the prediction error of the full number of components, equal to the number of recorded muscles (16 and 6, monkey S and I).

## 5.5 DISCUSSION

Muscle activity and motor cortical activity were recorded during a ballistic-release task that dissociated force and stiffness. We found a component of muscle activity correlated with force and a different component of muscle activity correlated with model stiffness. Analogous force and stiffness components were found in the population-based motor cortical activity and these can be considered separate and parallel signals. A second set of analyses examined the correlation between the population of recorded neural units and the simultaneous EMG activity of the sampled muscles. These results showed that there were two strong components that accounted for most of the correlation between M1 activity and muscle activity.

### 5.5.1 EMG modulation related to task conditions

Muscle activity was used to exert enough force to cross the force threshold, accelerate the handle toward the target zone, and then decelerate and stop the handle within the target zone. Additional muscle activity stiffened the arm, likely contributing to the small oscillation in position as the handle neared the hold position, characteristic of an under-damped physical system. The temporal pattern of EMG found in this study matched that of a similar ballistic-release task (Terzuolo et al., 1973) and is consistent with previous findings that EMG can be independently modulated to accelerate and decelerate the hand (Hoffman and Strick, 1993).

### 5.5.2 Force-potent and stiffness-potent components of muscle and M1 activity

Force-potent and stiffness-potent components of muscle and motor cortical activity were found using linear regression models and singular value decomposition. We assumed a linear model because it was a simple model that could still explain much of the variance. However, the relation between force and muscle activity in this task was likely non-linear and there may be aspects of the EMG-force relation that are not captured in this type of model. In addition to the changes in stiffness and force due to muscle activity, both also depend on the arm's skeletal configuration (Milner, 2002; Hu et al., 2012). As a result, the force-potent and stiffness-potent components of EMG should be interpreted as only a partial

accounting of arm stiffness. The same is true for the relation between force and M1 activity. Nonetheless, our results demonstrate that the linearity assumption is reasonable and that skeletal configuration does not obscure the M1 and muscle relation to stiffness.

The physical model used to estimate stiffness assumed a constant stiffness value across time. We have previously discussed that the model’s ability to explain the time-varying force suggests that this assumption is reasonable (Chapter 3). Here, we also made the assumption of constant stiffness values across time for the regression of stiffness against muscle and M1 activity. The goodness of fit in these regressions further supports the constant stiffness assumption in this task and provides an information-efficient strategy for signaling force and motion control in muscle and M1 activity.

### 5.5.3 Communication components between motor cortex and muscles

We found evidence for two components of correlation linking motor cortical activity to muscle activity. There are many connections between motor cortex and muscles that could support correlated components (Phillips and Porter, 1977; Muir and Lemon, 1983; Lemon, 2008). Direct, cortico-motoneuronal connections have been found to depend on the functional role of the muscle (Griffin et al., 2015) and distinct spatial populations have been found to be active depending on the functional role of the muscles (Humphrey and Reed, 1983).

It has been proposed that the motor cortex can be divided into “new” and “old” regions (Rathelot and Strick, 2006, 2009). Old motor cortex is associated with the rostral region of M1 and contains cortico-spinal projections that mainly synapse onto interneurons in the intermediate zone of the spinal cord and can only influence muscle activity indirectly. In contrast, new motor cortex is associated with the caudal region of M1 in the anterior bank of the central sulcus and contains a relatively higher density of cortico-spinal projections that synapse directly onto alpha motor neurons in the ventral horn of the spinal cord (cortico-motoneuronal connections) and can influence muscle activity directly. New M1 and the cortico-motoneuronal connections are a hallmark of mammals capable of advanced manipulative behavior.

The rostral-caudal organization of M1 has also been reported in terms of force (Kalaska et al., 1989). Neurons in the caudal region of M1 were more sensitive to the force exerted during reaching movements than neurons in the rostral region of M1. By interpreting muscles as the actuators that exert force, it is possible to connect the force sensitivity of caudal M1 with the cortico-motoneuronal connections in this same area. Thus, caudal M1, and especially the part of M1 in the anterior bank of the central sulcus, seems to be the most likely location for functional representations of muscle activity in M1. Although our recordings were primarily in the rostral region of M1, future studies could detail the rostral-caudal distribution of force and stiffness components and how they relate to similar components in muscle activity.

#### **5.5.4 Conclusion**

Our findings contribute to the understanding of potential functional components of activity in muscles and the motor cortex. The force-potent and stiffness-potent components of activity in M1 and muscle activity could be mixed to control force and motion during object manipulation. The potential for communication components dedicated to these functional components could represent a parceling of muscle-related representations in M1 activity. These concepts will not only be useful for understanding the context-dependent nature of connectivity between motor cortex and muscles, but also for understanding how information encoded in motor cortical activity might be transformed into muscle activity.

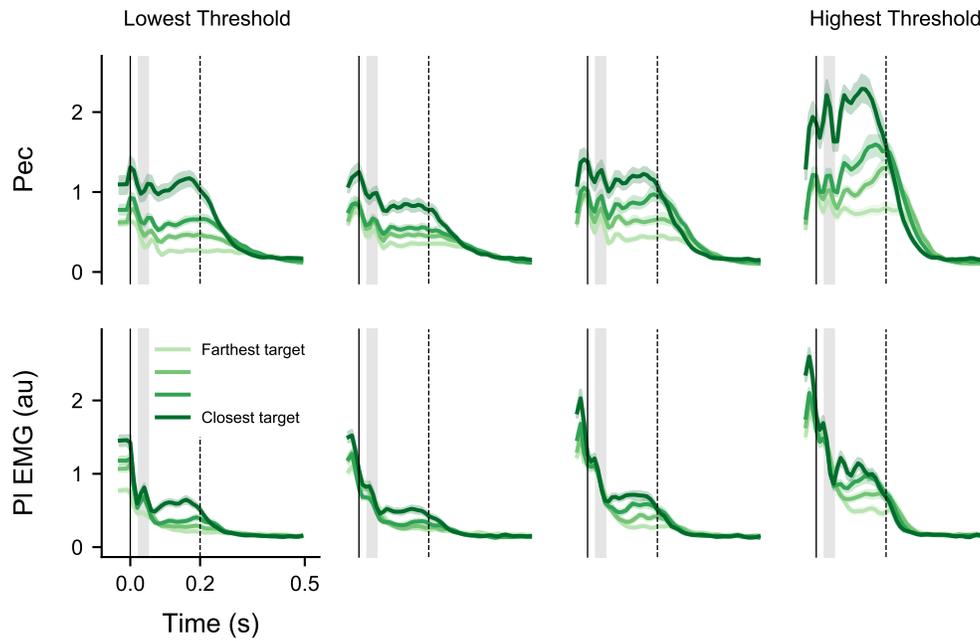


Figure 5.11: **Extending time reveals muscle activity likely related to stopping the handle.** Each plot depicts four target zones with the same force threshold. The EMG values were separated across target zones for approximately 300-400 ms after movement onset (time 0), after which the values converged to a baseline level. The timing corresponded with the end of movement and made it likely that these muscles were used to exert force to stop the handle within the target zone. Shading represents the median and 95% confidence interval.

## 6.0 DISCUSSION

The objective of this thesis was to describe the role of the motor cortex in combined force-motion control. The ability to control the force exerted on, and the resultant motion of, an object is critical to successfully manipulating the object. We hypothesized that modulating arm impedance could be an efficient strategy for this type of control and that information about arm impedance would be encoded in motor cortical activity. To dissociate force, motion, and impedance, we adopted and validated a ballistic-release task and used it to characterize anticipatory changes in impedance. We modeled the behavior of the arm as a physical dynamical system and found that, in this task, the model's stiffness was the only impedance component that varied consistently across both force and position constraints. By relating the force and position constraints to motor cortical activity, we revealed position and stiffness information encoded in motor cortical activity. Impedance and a zero-force position were decoded and used within the impedance control framework to estimate a time-varying force that closely approximated the actual force exerted on the object. We also found separate components of muscle activity and motor cortical activity correlated with force and stiffness.

These results emphasized the anticipatory behavior of a ballistic-release task and explored behavioral, muscular, and neural strategies that could be pre-set to achieve the task goals. The task emphasized the control of both force *and* motion, highlighting the role of arm impedance in object manipulation. This approach unified existing neural encoding models of force (Evarts, 1968; Kalaska et al., 1989; Georgopoulos et al., 1992; Sergio and Kalaska, 2003; Sergio et al., 2005) and motion (Georgopoulos et al., 1982, 1986; Schwartz, 1994; Moran and Schwartz, 1999) with a successful control framework (Hogan, 1985a, 2014).

Finally, it provided a rich behavioral context within which to explore distinct, functional representations of muscle activity in the motor cortex dedicated to controlling force and stiffness.

## 6.1 DIRECTIONS OF FORCE, POSITION, AND STIFFNESS

In this ballistic-release task, stiffness depended on both force and position constraints and could be used as an anticipatory strategy to control the motion of the handle. Importantly, all descriptions of stiffness, force, and motion were in the same direction: parallel with the shoulders. Stiffness is inherently a three-dimensional property (Mussa-Ivaldi et al., 1985) and although stiffness can be modulated in different directions (Kadiallah et al., 2011), the extent to which it can be modulated depends on direction (Perreault et al., 2002). Future enhancement of this paradigm will need to include multiple directions (Darainy et al., 2007) to determine the direction dependency of these results and the effect of different combinations of force and motion constraints in different directions.

## 6.2 TIME-VARYING ESTIMATIONS OF THE ZERO-FORCE POSITION AND IMPEDANCE COMPONENTS

The physical dynamical system used to model behavior assumed constant values across time for the zero-force position and impedance components, and could explain most of the variance in force. Although the impedance components were constant, their contributions to the force exerted on the handle varied in time according to the motion. The stiffness component contributed the most force across time at the moment of movement onset, as movement toward the zero-force position caused the stiffness component to exert decreasing force. Similarly, the second-order impedance component also contributed the most force across time at the moment of movement onset, when the pre-loaded initial conditions caused the acceleration to be at its maximum value. However, the stiffness and second-order impedance

components exerted force in opposite directions at movement onset, a key feature of the dynamic response of the system. The first-order impedance component described changes in force that increased and decreased with velocity.

Although we found that a model with impedance coefficients that were constant across time could describe the behavior during the ballistic-release task, it is possible that the impedance coefficients could vary across time for other types of object manipulation. The voluntary modulation of arm stiffness and damping change on the order of 200 ms ([Lacquaniti et al., 1982](#)) and inertia changes instantaneously as the arm configuration changes. Time-varying estimates of arm impedance ([Lacquaniti et al., 1993](#); [Piovesan et al., 2013](#)) in future experiments would make it possible to generalize the control of object manipulation to a wider range of behaviors and could increase the explanatory power of impedance variables.

### 6.3 ARM, JOINT, AND MUSCLE IMPEDANCE

This thesis has been focused on arm impedance. However, impedance is an additive property, meaning that the joint impedance of multiple joints add together to contribute to arm impedance and the muscle impedance of multiple muscles add together to contribute to joint impedance ([Hogan, 1985b](#)). The arm is a complex system of muscles and joints and the model impedance described in this thesis is similar to the equivalent impedance of the multiple individual components. Ultimately, it is the net effect of all the impedance components that governs the interaction with the object. However, new insights might be gained by considering the way that individual muscles change impedance ([Rack and Westbury, 1974](#); [Kurtzer et al., 2009](#)) or the way that multiple muscles act together to change the impedance of individual joints ([Osu and Gomi, 1999](#); [Perreault et al., 2001](#)).

## 6.4 CONCLUSION

Using a ballistic-release paradigm, this thesis systematically explored the combined control of force and motion. The combination of rich behavior and neural recordings enabled the blending of traditional neural encoding models with an established control theory framework. Our findings contribute to the understanding of motor cortical activity and muscle activity that could be used to control both force and motion. These concepts will not only be useful for understanding the role of the motor cortex in object manipulation, but also for understanding how information encoded in motor cortical activity might be transformed into muscle activity. Extensions of this work will help elucidate more detailed models of neural activity in a variety of contexts, particularly for anticipatory changes relevant for object manipulation. It will also add to the versatility of brain-machine interfaces, providing a scientific foundation to leverage new engineering technology that could help people who are paralyzed manipulate objects and perform many activities of daily living.

## APPENDIX

### FORMATTED DATA

The data from each session were collected as individual data sources (task, force, EMG, motion, neural). Offline, the motion-tracking markers were labeled and the 3D positions were extracted; the neural activity was sorted into neural units. Each source of data was aligned in time using a hardware sync pulse and stored in a Matlab storage file according to the following structured format.

Data

Tags

regular

[1xT logical]: True for regular trials

catch

[1xT logical]: True for catch trials

TaskStateMasks

InterTrial

[1xT logical]: Between trials.

ForceRamp

[1xT logical]: After the start button is pressed and before the force threshold is crossed.

Move

[1xT logical]: After the force threshold is crossed and before the handle enters the target.

Hold

[1xT logical]: After the handle enters the target and before the trial is a success.

Reward

[1xT logical]: When the reward is being given.

Return

[1xT logical]: After the end of the trial and before the next trial. Return the handle to the start position.

EndTrial

[1xT logical]: End of the trial.

Adjusted

Reach

[1xT logical]: After the start button is pressed and before the force remains above 1 N.

ForceRamp

[1xT logical]: After the force rises and stays above 1 N and before the speed rises above 10% of max trial speed.

Move

[1xT logical]: After the speed rises about 10% of max trial speed and before the handle enters the target.

Hold

[1xT logical]: 300 ms beginning when the handle enters the target.

OutcomeMasks

Success

[1xT logical]: True if trial was a success.

ManualSuccess

[1xT logical]: True if the experimenter manually succeeded the trial.

ManualProceed

[1xT logical]: True if the experimenter manually proceeded through a task state.

Failure

[1xT logical]: True if the trial was a failure.

Canceled

[1xT logical]: True if the trial was canceled.

Attempt

[1xT logical]: True if the subject did not attempt the trial.

DidNotStart

[1xT logical]: True if the trial did not proceed through enough task states to start.

Force

data

[6xT double]: Values for the 6 dimensions of force: linear xyz and rotational xyz. +X toward subject, +Y along track, +Z upwards. Optional: after a rotation is applied.

raw

[6xT double]: Values for the 6 dimensions of force before any (if any) rotation was applied.

threshold

upper

[6xT double]: Upper force threshold.

lower

[6xT double]: Lower force threshold.

normLower

[6xT double]: Upper force threshold divided by the maximum voluntary force.

normData

[1xT double]: Values along the track divided by the maximum voluntary force.

Target

pos\_id

[1xT double]: Denso position id. Not used.

Kinematics

MarkerNames

{1xM cell}: Name for each marker.

MarkerIds

MarkerPos

[MxT double]: 3D position (mm) for each marker. XYZ for each marker stacked along 1st dimension. +X along track, +Y away from subject, +Z upward. Origin is the position of the HN DL marker when the handle is locked.

TargetWindows

[2xT double]: Position along the track (mm) of the near and far edges of the target zone.

JointAngle

[JxT double]: Joint angles (deg) as described in the OpenSim model.

JointNames

{1xJ cell}: Name for each joint.

JointCenters

[12xT double]: 3D position (mm) of select joint positions. XYZ for each joint stacked along the 1st dimension.

JointCenterNames

{4 cell}: Name for select joint centers.

Muscle

EMG

[MxT double]: Z-scored and filtered EMG values.

MuscleNames

{1xM cell}: Muscle names

Raw

[MxT double]: Raw bipolar EMG values.

normEMG

[MxT double]: EMG values divided by the 99.9% percentile, excluding intertrial periods.

Spikes

SpikeTS

{Sx1 cell}: Spike times for each neural unit.

FiringRate

[SxT double]: Firing rates for each neural unit.

AverageWaveform

{Sx1 cell}: The average waveform for each neural unit.

Channel

[Sx1 double]: Recording channel number

Unit

[Sx1 double]: Unit label for a given channel

ArrayNames

{1xA cell}: Name of the micro-electrode array

StartStopArrayChannelNums

[2A]: The start and stop channel numbers for each array.

Time

[1xT double]: Time in seconds

TrialNo

[1xT int32]: Trial number

BlockNo

[1xT int32]

ComboNo

[1xT int32]

Dt

double: Time (s) between observations

Version

string: Dragonfly version

Build

string: Dragonfly build date

ConfigName

string: Configuration filename for Dragonfly

SessionNo

int32: Session number

Subject

string: Subject's name

## BIBLIOGRAPHY

- H. Asanuma and I. Rosen. Topographical organization of cortical efferent zones projecting to distal forelimb muscles in the monkey. *Experimental Brain Research*, 14:243–256, 1972.
- James Ashe. Force and the motor cortex. *Behavioural Brain Research*, 86:1–15, 1997.
- N. A. Bernstein. *The co-ordination and regulation of movements*. Pergamon Press Ltd., 1967.
- E Bizzi, N Accornero, W Chapple, and Neville Hogan. Arm trajectory formation in monkeys. *Experimental Brain Research*, 46(1):139–143, apr 1982. doi: 10.1007/BF00238107.
- E. Bizzi, N. Accornero, W. Chapple, and Neville Hogan. Posture control and trajectory formation during arm movement. *Journal of Neuroscience*, 4(11):2738–44, nov 1984.
- Etienne Burdet, Rieko Osu, David W. Franklin, Theodore E. Milner, and Mitsuo Kawato. The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature*, 414(6862):446–449, nov 2001. doi: 10.1038/35106566.
- PD Cheney and EE Fetz. Functional classes of primate corticomotoneuronal cells and their relation to active force. *Journal of Neurophysiology*, 44(4):773–91, oct 1980.
- Tyler R. Clites, Matthew J. Carty, Jessica B. Ullauri, Matthew E. Carney, Luke M. Mooney, Jean François Duval, Shriya S. Srinivasan, and Hugh M. Herr. Proprioception from a neurally controlled lower-extremity prosthesis. *Science Translational Medicine*, 10(443), 2018. doi: 10.1126/scitranslmed.aap8373.
- L Damm and J McIntyre. Physiological basis of limb-impedance modulation during free and constrained movements. *Journal of neurophysiology*, 100(5):2577–2588, 2008. doi: 10.1152/jn.90471.2008.
- Mohammad Darainy, Farzad Towhidkhah, and David J. Ostry. Control of hand impedance under static conditions and during reaching movement. *Journal of Neurophysiology*, 97(4): 2676–85, apr 2007. doi: 10.1152/jn.01081.2006.

- Scott L. Delp, Frank C. Anderson, Allison S. Arnold, Peter Loan, Ayman Habib, Chand T. John, Eran Guendelman, and Darryl T. Thelen. OpenSim: Open-Source Software to Create and Analyze Dynamic Simulations of Movement. *IEEE Transactions on Biomedical Engineering*, 54(11):1940–1950, 2007.
- Jörn Diedrichsen, Reza Shadmehr, and Richard B. Ivry. The coordination of movement: optimal feedback control and beyond. *Trends in Cognitive Sciences*, 14(1):31–39, 2010. doi: 10.1016/j.tics.2009.11.004.
- Michael Dimitriou, Daniel M. Wolpert, and David W. Franklin. The Temporal Evolution of Feedback Gains Rapidly Update to Task Demands. *Journal of Neuroscience*, 33(26):10898–10909, jun 2013. doi: 10.1523/JNEUROSCI.5669-12.2013.
- Aaron M. Dollar and Hugh Herr. Lower Extremity Exoskeletons and Active Orthoses: Challenges and State-of-the-Art. *IEEE Transactions on Robotics*, 24(1):144–158, 2008. doi: 10.1109/TRO.2008.915453.
- Digby Elliott, Matthew Heath, Gordon Binsted, Kathryn L. Ricker, Eric A. Roy, and Romeo Chua. Goal-Directed Aiming: Correcting a Force-Specification Error With the Right and Left Hands. *Journal of Motor Behavior*, 31(4):309–324, dec 1999. doi: 10.1080/00222899909600997.
- Digby Elliott, Steve Hansen, Lawrence E. M. Grierson, James Lyons, Simon J. Bennett, and Spencer J. Hayes. Goal-Directed Aiming: Two Components but Multiple Processes. *Psychological Bulletin*, 136(6):1023–1044, 2010. doi: 10.1037/a0020958.
- Edward V. Evarts. Relation of pyramidal tract activity to force exerted during voluntary movement. *Journal of Neurophysiology*, 31(1):14–27, jan 1968.
- A G Feldman. On the functional tuning of the nervous system in movement control or preservation of stationary pose. II. Adjustable parameters in muscles. *Biofizika*, 11(3):498–508, 1966.
- A G Feldman. Once more on the equilibrium-point hypothesis ( $\lambda$  model) for motor control. *Journal of Motor Behavior*, 18(1):17–54, mar 1986.
- D Ferrier. *The functions of the brain*. Smith, Elder, London, 1886.
- J Randall Flanagan, Miles C Bowman, and Roland S. Johansson. Control strategies in object manipulation tasks. *Current Opinion in Neurobiology*, 16(6):650–659, dec 2006. doi: 10.1016/j.conb.2006.10.005.
- T Flash. The control of hand equilibrium trajectories in multi-joint arm movements. *Biological Cybernetics*, 57:257–274, 1987. doi: 10.1007/BF00338819.
- T Flash and Neville Hogan. The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of Neuroscience*, 5(7):1688–1703, jul 1985.

- David W. Franklin and Daniel M. Wolpert. Computational Mechanisms of Sensorimotor Control. *Neuron*, 72(3):425–442, nov 2011. doi: 10.1016/j.neuron.2011.10.006.
- David W. Franklin, Rieko Osu, Etienne Burdet, Mitsuo Kawato, and Theodore E. Milner. Adaptation to Stable and Unstable Dynamics Achieved By Combined Impedance Control and Inverse Dynamics Model. *Journal of Neurophysiology*, 90(5):3270–3282, nov 2003. doi: 10.1152/jn.01112.2002.
- David W. Franklin, Udell So, Mitsuo Kawato, and Theodore E. Milner. Impedance control balances stability with metabolically costly muscle activation. *Journal of Neurophysiology*, 92(5):3097–105, nov 2004. doi: 10.1152/jn.00364.2004.
- David W. Franklin, Gary Liaw, Theodore E. Milner, Rieko Osu, Etienne Burdet, and Mitsuo Kawato. Endpoint Stiffness of the Arm Is Directionally Tuned to Instability in the Environment. *Journal of Neuroscience*, 27(29):7705–7716, jul 2007. doi: 10.1523/JNEUROSCI.0968-07.2007.
- George Williams Fraser and Andrew B. Schwartz. Recording from the same neurons chronically in motor cortex. *Journal of Neurophysiology*, 107(7):1970–1978, apr 2012. doi: 10.1152/jn.01012.2010.
- G Fritsch and E Hitzig. Ueber die elektrische Erregbarkeit des Grosshirns. *Arch. Anat. Physiol. wiss. Med.*, 37:300–332, 1870.
- Apostolos P. Georgopoulos, John F. Kalaska, R Caminiti, and JT Massey. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *Journal of Neuroscience*, 1982.
- Apostolos P. Georgopoulos, Andrew B. Schwartz, and RE Kettner. Neuronal population coding of movement direction. *Science*, 1986.
- Apostolos P. Georgopoulos, J. Ashe, N. Smyrnis, and M. Taira. The Motor Cortex and the Coding of Force. *Science*, 256(5064):1692–1695, jun 1992. doi: 10.1126/science.256.5064.1692.
- H Gomi and M Kawato. Human arm stiffness and equilibrium-point trajectory during multi-joint movement. *Biological Cybernetics*, 76(3):163–171, apr 1997. doi: 10.1007/s004220050329.
- H Gomi and Rieko Osu. Task-dependent viscoelasticity of human multijoint arm and its spatial characteristics for interaction with environments. *Journal of Neuroscience*, 18(21):8965–8978, nov 1998.
- Hiroaki Gomi and Kawato. Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science*, 272(5258):117–20, apr 1996. doi: 10.1126/science.272.5258.117.

- Paul L. Gribble, Lucy I. Mullin, Nicholas Cothros, and Andrew Mattar. Role of Cocontraction in Arm Movement Accuracy. *Journal of Neurophysiology*, 89(5):2396–2405, may 2003. doi: 10.1152/jn.01020.2002.
- Darcy M. Griffin, Donna S. Hoffman, and Peter L. Strick. Corticomotoneuronal cells are "functionally tuned". *Science*, 350(6261):667–670, nov 2015. doi: 10.1126/science.aaa8035.
- Christopher M. Harris and Daniel M. Wolpert. Signal-dependent noise determines motor planning. *Nature*, 394(6695):780–4, aug 1998. doi: 10.1038/29528.
- Rodolphe Heliot, AL L Orsborn, K Ganguly, and JM M Carmena. System Architecture for Stiffness Control in Brain Machine Interfaces. *IEEE Transactions on Systems, Man and Cybernetics*, 40(4):732–742, 2010.
- A. V. Hill. The series elastic component of muscle. *Proceedings of the Royal Society of London. Series B, Biological sciences*, 137(887):273–280, jul 1950.
- Antony J. Hodgson and Neville Hogan. A model-independent definition of attractor behavior applicable to interactive tasks. *IEEE Transactions on Systems, Man and Cybernetics*, 30(1):105–118, 2000. doi: 10.1109/5326.827459.
- D S Hoffman and Peter L. Strick. Step-Tracking Movements of the Wrist in Humans . II . EMG Analysis. *Journal of Neuroscience*, 10(January):142–152, jan 1990.
- D S Hoffman and Peter L. Strick. Step-tracking movements of the wrist. III. Influence of changes in load on patterns of muscle activity. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 13(12):5212–5227, dec 1993.
- Neville Hogan. Impedance Control: An Approach to Manipulation. In *IEEE American Control Conference*, pages 304–313, 1984a.
- Neville Hogan. Adaptive control of mechanical impedance by coactivation of antagonist muscles. *IEEE Transactions on Automatic Control*, 29(8):681–690, 1984b. doi: 10.1109/TAC.1984.1103644.
- Neville Hogan. An organizing principle for a class of voluntary movements. *Journal of Neuroscience*, 4(11):2745–2754, nov 1984c.
- Neville Hogan. The Mechanics of Multi-Joint Posture and Movement Control. *Biological Cybernetics*, 331:315–331, 1985a.
- Neville Hogan. Impedance Control: An Approach to Manipulation: Part II - Implementation. *Journal of Dynamic Systems, Measurement, and Control*, 107:1–7, 1985b.
- Neville Hogan. Impedance Control: An Approach to Manipulation: Part I - Theory. *Journal of Dynamic Systems, Measurement, and Control*, 107(March), 1985c.

- Neville Hogan. A General Actuator Model Based on Nonlinear Equivalent Networks. *IEEE/ASME Transactions on Mechatronics*, 19(6):1929–1939, dec 2014. doi: 10.1109/TMECH.2013.2294096.
- Katherine R. S. Holzbaur, Wendy M. Murray, and Scott L. Delp. A Model of the Upper Extremity for Simulating Musculoskeletal Surgery and Analyzing Neuromuscular Control. *Annals of Biomedical Engineering*, 33(6):829–840, 2005. doi: 10.1007/s10439-005-3320-7.
- X Hu, WM Murray, and Eric J. Perreault. Biomechanical constraints on the feedforward regulation of endpoint stiffness. *Journal of neurophysiology*, 108(8):2083–91, oct 2012. doi: 10.1152/jn.00330.2012.
- D R Humphrey and D J Reed. Separate cortical systems for control of joint movement and joint stiffness: reciprocal activation and coactivation of antagonist muscles. *Advances in neurology*, 39:347–72, 1983.
- G C Joyce and P M Rack. Isotonic lengthening and shortening movements of cat soleus muscle. *Journal of Physiology*, 204(2):475–491, oct 1969.
- Abdelhamid Kadiallah, Gary Liaw, Mitsuo Kawato, David W. Franklin, and Etienne Burdet. Impedance control is selectively tuned to multiple directions of movement. *Journal of Neurophysiology*, 106(5):2737–2748, nov 2011. doi: 10.1152/jn.00079.2011.
- S Kakei, Donna S. Hoffman, and Peter L. Strick. Muscle and Movement Representations in the Primary Motor Cortex. *Science*, 285(September):2136–2139, 1999.
- John F. Kalaska. *Progress in Motor Control*, volume 629 of *Advances in Experimental Medicine and Biology*. Springer US, Boston, MA, 2009. doi: 10.1007/978-0-387-77064-2.
- John F. Kalaska and DJ Crammond. Cerebral Cortical Mechanisms of Reaching Movements. *Science*, 1992.
- John F. Kalaska, DAD Cohen, ML Hyde, and M Prud’homme. A Comparison of Movement Direction-Related Versus Load Activity in Primate Motor Cortex, Using a Two-Dimensional Reaching Task. *Journal of Neuroscience*, 9(6):2080–2102, 1989.
- M Katayama and M Kawato. Virtual trajectory and stiffness ellipse during multijoint arm movement predicted by neural inverse models. *Biological Cybernetics*, 69(5-6):353–362, 1993. doi: 10.1007/BF00199435.
- Matthew T Kaufman, Mark M Churchland, Stephen I Ryu, and KrishnaV. Shenoy. Cortical activity in the null space: permitting preparation without movement. *Nature Neuroscience*, 17(3):440–8, mar 2014. doi: 10.1038/nn.3643.
- Mitsuo Kawato. Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9(6):718–727, dec 1999. doi: 10.1016/S0959-4388(99)00028-8.

- Hyun K Kim, Jose M Carmena, SJ James Biggs, Timothy L Hanson, Miguel A L Nicolelis, and Mandayam A Srinivasan. The Muscle Activation Method: An Approach to Impedance Control of Brain-Machine Interfaces Through a Musculoskeletal Model of the Arm. *IEEE Transactions on Biomedical Engineering*, 54(8):1520–1529, aug 2007. doi: 10.1109/TBME.2007.900818.
- Dmitry Kobak, Wieland Brendel, Christos Constantinidis, Claudia E. Feierstein, Adam Kepecs, Zachary F. Mainen, Xue-lian Lian Qi, Ranulfo Romo, Naoshige Uchida, and Christian K. Machens. Demixed principal component analysis of neural population data. *eLife*, 5(APRIL2016):1–36, 2016. doi: 10.7554/eLife.10989.
- Isaac L. Kurtzer, Troy M Herter, and Stephen H. Scott. Nonuniform Distribution of Reach-Related and Torque-Related Activity in Upper Arm Muscles and Neurons of Primary Motor Cortex. *Journal of Neurophysiology*, 96(6):3220–3230, dec 2006. doi: 10.1152/jn.00110.2006.
- Isaac L. Kurtzer, J. Andrew Pruszynski, and Stephen H. Scott. Long-Latency Responses During Reaching Account for the Mechanical Interaction Between the Shoulder and Elbow Joints. *Journal of Neurophysiology*, 102(5):3004–3015, nov 2009. doi: 10.1152/jn.00453.2009.
- J R Lackner and P Dizio. Rapid adaptation to Coriolis force perturbations of arm trajectory. *Journal of neurophysiology*, 72(1):299–313, 1994. doi: citeulike-article-id:450102.
- Francesco Lacquaniti, F Licata, and JF Soechting. The Mechanical Behavior of the Human Forearm in Response to Transient Perturbations. *Biological Cybernetics*, pages 291–294, 1982.
- Francesco Lacquaniti, M Carrozzo, and N A Borghese. Time-varying mechanical behavior of multijointed arm in man. *Journal of Neurophysiology*, 69(5):1443–1464, may 1993. doi: 10.1152/jn.1993.69.5.1443.
- Mark L. Latash. Evolution of Motor Control: From Reflexes and Motor Programs to the Equilibrium-Point Hypothesis. *Journal of Human Kinetics*, 19(-1):3–24, 2008. doi: 10.2478/v10078-008-0001-2.
- Mark L. Latash. Abundant Degrees of Freedom Are Not a Problem. *Kinesiology Review*, 7(1):64–72, feb 2018. doi: 10.1123/kr.2017-0058.
- Mark L. Latash and G.L. Gottlieb. Reconstruction of shifting elbow joint compliant characteristics during fast and slow movements. *Neuroscience*, 43(2-3):697–712, jan 1991. doi: 10.1016/0306-4522(91)90328-L.
- Roger N Lemon. Descending pathways in motor control. *Annual review of neuroscience*, 31(Cm):195–218, 2008. doi: 10.1146/annurev.neuro.31.060407.125547.

- A. S. F. Leyton and C. S. Sherrington. OBSERVATIONS ON THE EXCITABLE CORTEX OF THE CHIMPANZEE, ORANG-UTAN, AND GORILLA. *Quarterly Journal of Experimental Physiology*, 11(2):135–222, jul 1917. doi: 10.1113/expphysiol.1917.sp000240.
- Gerald E. Loeb and Carl Gans. *Electromyography for experimentalists*. The University of Chicago Press, Chicago, 1986.
- J McIntyre, FA Mussa-Ivaldi, and E Bizzi. The control of stable postures in the multijoint arm. *Experimental Brain Research*, 110(2):248–264, jul 1996. doi: 10.1007/BF00228556.
- David E Meyer, Richard A Abrams, Sylvan Kornblum, Charles E Wright, and J E Smith. Optimality in human motor performance: ideal control of rapid aimed movements. *Psychological review*, 95(3):340–370, jul 1988.
- L E Miller, P L van Kan, T Sinkjaer, T Andersen, G D Harris, and J C Houk. Correlation of primate red nucleus discharge with muscle activity during free-form arm movements. *The Journal of physiology*, 469:213–43, sep 1993.
- Theodore E. Milner. Contribution of geometry and joint stiffness to mechanical stability of the human arm. *Experimental Brain Research*, 143(4):515–519, 2002. doi: 10.1007/s00221-002-1049-1.
- Theodore E. Milner and David W. Franklin. Impedance control and internal model use during the initial stage of adaptation to novel dynamics in humans. *Journal of Physiology*, 567(2):651–664, sep 2005. doi: 10.1113/jphysiol.2005.090449.
- DW Moran and Andrew B. Schwartz. Motor cortical representation of speed and direction during reaching. *Journal of neurophysiology*, 82(5):2676–92, nov 1999.
- MM Morrow and LE Miller. Prediction of muscle activity by populations of sequentially recorded primary motor cortex neurons. *Journal of Neurophysiology*, 89(4):2279–88, apr 2003. doi: 10.1152/jn.00632.2002.
- R. B. Muir and R. N. Lemon. Corticospinal neurons with a special role in precision grip. *Brain Research*, 1983. doi: 10.1016/0006-8993(83)90635-2.
- F. A. Mussa-Ivaldi, Neville Hogan, and E. Bizzi. Neural, Mechanical, and Geometric Factors Subservicing Arm Posture in Humans. *Journal of Neuroscience*, 5(10):2732–2743, 1985.
- Isaac Newton. *Sir Isaac Newton’s mathematical principles of natural philosophy and his system of the world*. University of California Press, 1962.
- T R Nichols and J C Houk. Improvement in linearity and regulation of stiffness that results from actions of stretch reflex. *Journal of Neurophysiology*, 39(1):119–142, jan 1976. doi: 10.1152/jn.1976.39.1.119.

- Rieko Osu and Hiroaki Gomi. Multijoint Muscle Regulation Mechanisms Examined by Measured Human Arm Stiffness and EMG Signals. *Journal of Neurophysiology*, 81(4): 1458–1468, apr 1999. doi: 10.1152/jn.1999.81.4.1458.
- Rieko Osu, Etienne Burdet, David W. Franklin, Theodore E. Milner, and Mitsuo Kawato. Different Mechanisms Involved in Adaptation to Stable and Unstable Dynamics. *Journal of Neurophysiology*, 90(5):3255–3269, nov 2003. doi: 10.1152/jn.00073.2003.
- Sagi Perel. *Dynamic Functional Connectivity Between Cortex and Muscles*. PhD thesis, University of Pittsburgh, 2012.
- Eric J. Perreault, RF Kirsch, and PE Crago. Effects of voluntary force generation on the elastic components of endpoint stiffness. *Experimental brain research*, 141(3):312–23, dec 2001. doi: 10.1007/s002210100880.
- Eric J. Perreault, Robert F Kirsch, and Patrick E Crago. Voluntary Control of Static Endpoint Stiffness During Force Regulation Tasks. *Journal of Neurophysiology*, 87(6): 2808–2816, jun 2002. doi: 10.1152/jn.2002.87.6.2808.
- CG Phillips and R Porter. *Corticospinal Neurons: Their Role in Movement*. Academic Press, San Francisco, 1977.
- Davide Piovesan, Alberto Pierobon, Paul DiZio, and James R Lackner. Experimental measure of arm stiffness during single reaching movements with a time-frequency analysis. *Journal of Neurophysiology*, 110(10):2484–2496, nov 2013. doi: 10.1152/jn.01013.2012.
- A Polit and E Bizzi. Characteristics of motor programs underlying arm movements in monkeys. *Journal of neurophysiology*, 42(1 Pt 1):183–94, jan 1979.
- J. Andrew Pruszynski, Mohsen Omrani, and Stephen H. Scott. Goal-Dependent Modulation of Fast Feedback Responses in Primary Motor Cortex. *Journal of Neuroscience*, 34(13): 4608–4617, mar 2014. doi: 10.1523/JNEUROSCI.4520-13.2014.
- Peter M. H. Rack and D R Westbury. The short range stiffness of active mammalian muscle and its effect on mechanical properties. *Journal of Physiology*, 240(2):331–350, jul 1974. doi: 10.1113/jphysiol.1974.sp010613.
- Denis Rancourt and Neville Hogan. Stability in Force-Production Tasks. *Journal of Motor Behavior*, 33(2):193–204, jun 2001. doi: 10.1080/00222890109603150.
- JA Rathelot and Peter L. Strick. Muscle representation in the macaque motor cortex: an anatomical perspective. *Proceedings of the National Academy of Sciences of the United States of America*, 103(21):8257–62, may 2006. doi: 10.1073/pnas.0602933103.
- JA Rathelot and Peter L. Strick. Subdivisions of primary motor cortex based on cortico-motoneuronal cells. *Proceedings of the National Academy of Sciences of the United States of America*, 106(3):918–23, jan 2009. doi: 10.1073/pnas.0808362106.

- GA Reina, DW Moran, and Andrew B. Schwartz. On the Relationship Between Joint Angular Velocity and Motor Cortical Discharge During Reaching. *Journal of Neurophysiology*, 92(February):2576–2589, 2001.
- Andrew B. Schwartz. Motor cortical activity during drawing movements: single-unit activity during sinusoid tracing. *Journal of Neurophysiology*, 68(2):528–541, aug 1992. doi: 10.1152/jn.1992.68.2.528.
- Andrew B. Schwartz. Cortical Representation of Drawing Direct Cortical Representation. *Science*, 265(July):540–542, 1994.
- Andrew B. Schwartz. Movement: How the Brain Communicates with the World. *Cell*, 164(6):1122–1135, 2016. doi: 10.1016/j.cell.2016.02.038.
- Stephen H. Scott. Optimal feedback control and the neural basis of volitional motor control. *Nature Reviews Neuroscience*, 5(7):532–545, jul 2004. doi: 10.1038/nrn1427.
- Stephen H. Scott, Tyler Cluff, Catherine R. Lowrey, and Tomohiko Takei. Feedback control during voluntary motor actions. *Current Opinion in Neurobiology*, 33:85–94, aug 2015. doi: 10.1016/j.conb.2015.03.006.
- Luc P J Selen, Peter J. Beek, and Jaap H. van Dieën. Impedance is modulated to meet accuracy demands during goal-directed arm movements. *Experimental Brain Research*, 172(1):129–138, jun 2006. doi: 10.1007/s00221-005-0320-7.
- Lauren E. Sergio and John F. Kalaska. Systematic Changes in Motor Cortex Cell Activity With Arm Posture During Directional Isometric Force Generation. *Journal of Neurophysiology*, 89(1):212–228, jan 2003. doi: 10.1152/jn.00016.2002.
- Lauren E. Sergio, C Hamel-Paquet, and John F. Kalaska. Motor Cortex Neural Correlates of Output Kinematics and Kinetics During Isometric-Force and Arm-Reaching Tasks. *Journal of Neurophysiology*, pages 2353–2378, 2005. doi: 10.1152/jn.00989.2004.
- Reza Shadmehr and John W. Krakauer. A computational neuroanatomy for motor control. *Experimental brain research*, 185(3):359–81, mar 2008. doi: 10.1007/s00221-008-1280-5.
- Reza Shadmehr and FA Mussa-Ivaldi. Adaptive Task of Dynamics during Learning of a Motor Task. *Journal of Neuroscience*, 74(May), 1994.
- C. D. Takahashi, R. A. Scheidt, and D. J. Reinkensmeyer. Impedance Control and Internal Model Formation When Reaching in a Randomly Varying Dynamical Environment. *Journal of Neurophysiology*, 86(2):1047–1051, aug 2001. doi: 10.1152/jn.2001.86.2.1047.
- CA Terzuolo, JF Soechting, and P Viviani. Studies on the control of some simple motor tasks. I. Relations between parameters of movements and EMG activities. *Brain Research*, 58(1):212–216, aug 1973. doi: 10.1016/0006-8993(73)90834-2.

- W T Thach. Correlation of neural discharge with pattern and force of muscular activity , joint position , and direction of intended next movement in motor cortex and cerebellu. *Journal of Neurophysiology*, 41(3), 1978.
- K A Thoroughman and Reza Shadmehr. Electromyographic correlates of learning an internal model of reaching movements. *Journal of Neuroscience*, 19(19):8573–8588, oct 1999.
- Emanuel Todorov. Direct cortical control of muscle activation in voluntary arm movements : a model. *Nature*, pages 391–398, 2000.
- Randy D Trumbower, Matthew A Krutky, Bing-Shiang Yang, and Eric J. Perreault. Use of Self-Selected Postures to Regulate Multi-Joint Stiffness During Unconstrained Tasks. *PLoS ONE*, 4(5):e5411, may 2009. doi: 10.1371/journal.pone.0005411.
- P Viviani and C A Terzuolo. Modeling of a simple motor task in man: intentional arrest of an ongoing movement. *Kybernetik*, 14(1):35–62, dec 1973.
- DanielM. Wolpert and Zoubin Ghahramani. Computational principles of movement neuroscience. *Nature Neuroscience*, 3(Supp):1212–1217, nov 2000. doi: 10.1038/81497.
- Robert Sessions Woodworth. Accuracy of voluntary movement. *Psychological Review: Monograph Supplements*, 3(3):i–114, 1899. doi: 10.1037/h0092992.