

All Models are Wrong, Simple Models Provide Insight: A Study of Human Manipulation

by

A. Michael West Jr.

B.S.

Mechanical Engineering, Yale University, 2018

S.M.

Mechanical Engineering, Massachusetts Institute of Technology, 2020

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

DOCTOR OF PHILOSOPHY

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2024

© 2024 A. Michael West Jr. This work is licensed under a [CC BY-SA 2.0](#) license.

The author hereby grants to MIT a nonexclusive, worldwide, irrevocable, royalty-free license to exercise any and all rights under copyright, including to reproduce, preserve, distribute and publicly display copies of the thesis, or release the thesis under an open-access license.

Authored by: A. Michael West Jr.
Department of Mechanical Engineering
April 30, 2024

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

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

All Models are Wrong, Simple Models Provide Insight: A Study of Human Manipulation

by

A. Michael West Jr.

Submitted to the Department of Mechanical Engineering
on May 3, 2024, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Humans possess a unique ability to manipulate tools, facilitated by the dexterity of our hands. Unfortunately, millions lose this capability annually due to conditions like limb amputation or Cerebral Vascular Accident. Robotic rehabilitation technologies, including prosthetics and exoskeletons, aim to restore motor function. However, understanding human control strategies is crucial for effective implementation. Unfortunately, studying humans is challenging due to their inherent complexity. Simple models can help untangle this complexity. This thesis delves into the role of simple models in analyzing human neural motor control and perception through the study of upper-limb motor control and hand manipulation.

Specifically, this thesis aims to develop a descriptive understanding of human manipulation in unimpaired subjects. To achieve this, we first delve into a common simplification of human hand manipulation—the presence of kinematic hand synergies—and stress the importance of studying functional hand manipulation beyond simple grasping. Secondly, we present findings from a motor control study that introduces a simple mathematical model, emphasizing mechanical impedance, that competently describes how humans manage physical interaction. Thirdly, we underscore the significance of mechanical impedance through a perceptual study which highlighted humans' robust ability to perceive limb stiffness. Lastly, we introduce a metric that can objectively quantify manipulation complexity, potentially broadening researchers' scope in studying human complex manipulation. The insights gained from this work have far-reaching implications, potentially enhancing existing robotic and rehabilitation technologies and guiding the development of new ones.

In a field dominated by large complex models fed by big data, this thesis highlights the value of conducting the basic science research required to uncover aspects of human motor control.

A video presentation of this thesis can be found at:
<https://www.youtube.com/watch?v=u2eCJHqEGww>

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

Acknowledgments

This thesis signifies the end of my formal academic journey, spanning 12 years of public elementary, middle, and high school education in California, followed by 4 years of undergraduate studies at Yale University, and finally, 6 years of graduate education at MIT. In total, my academic journey has spanned 22 consecutive years of education. This thesis marks the culmination of the final 6 years, resulting in the attainment of a PhD. As I reflect on the significance of this achievement, I must thank the individuals who have contributed to my academic journey and helped me reach this milestone.

First and foremost, I'd like to express my utmost gratitude to my advisor, Professor Neville Hogan. His expertise in controls and motor neuroscience is exceptional, and I am certain I wouldn't have acquired such profound knowledge under anyone else's guidance. Beyond his intellect, I deeply value his aptitude for teaching and mentoring students. Professor Hogan always demonstrated attentiveness, curiosity, and constructive feedback during our discussions about my research. I consistently left meetings with Professor Hogan more positive and energized. Working with and learning from him has been an amazing opportunity and incredible honor.

Moreover, thank you to my committee members Professor Harry Asada, Professor Alberto Rodriguez, and Professor Madhusudhan Venkadesan. Their perspectives, insights, and expertise were invaluable in conducting the research presented in this thesis.

I would also like to thank my labmates: Nicolas Arons, Nicole Attram, Max Burns, James Hermus, Meghan Huber, Johannes Lachner, Jongwoo Lee, Moses Nah, Erfan Shariari, Kaymie Shiozawa, Rika Sugimoto Dimitrova, Stephan Stansfield, Federico Tessari, Maggie Wang, and Kunpeng Yao. Throughout the years, they have all contributed to fostering a pleasant work environment. I have always been enthusiastic about collaborating with them, working through problems, and engaging in discussions. In particular, I want to extend special thanks to James, Meghan, and Federico, who have not only been invaluable labmates but also served as mentors,

providing me with both research and professional guidance. Additionally, I am grateful to Maggie and Nicole for affording me the opportunity to mentor them, which has enabled me to learn more about myself as a researcher and educator.

Prior to MIT, I was significantly influenced by several organizations during my undergraduate career. Specifically, I would like to express my gratitude to the National Society of Black Engineers (NSBE) for creating a platform for black engineers to convene and inspire one another. It was at a NSBE conference where I learned about what a PhD even is. NSBE has been instrumental in shaping my career aspirations, encapsulated by the organization's mission: *“to increase the number of culturally responsible black engineers who excel academically, succeed professionally, and positively impact the community.”* I am also grateful to the MIT Summer Research Program (MSRP) for equipping me with the necessary tools to confidently pursue a graduate degree at an institution like MIT. I extend a heartfelt acknowledgment to Former Dean Gloria Anglon, whose efforts bolstered my confidence in choosing to attend MIT. I can unequivocally state that without MSRP, I would not be here today.

Thank you to Noelle Wakefield and Maria Cervantes Gonzalez. They were always a pleasure to work with on various Office of Graduate Education initiatives. Working alongside them as a Pod Leader for MSRP has always been my favorite graduate school experience. Most importantly, I want to thank them for just being there. Their presence brought a familial feel to my experience at MIT.

Moreover, I extend my gratitude to the various Diversity, Equity, and Inclusion initiatives that have not only contributed to my academic development but also enriched me personally. At a time when some politicians have advocated for the demise of DEI, I want to underscore the positive impact these initiatives have had in nurturing the next generation of scientists. Indeed, part of their success is reflected in the scientific contributions presented in this thesis; without these programs which poured into me, I would not have been able to pour into my research as greatly as I did. I am thankful to the Institute on Teaching and Mentoring, the NextProf Nexus Future Faculty Workshop, the Georgia Tech Focus Fellows Program, Stanford

University's Postdoctoral Recruitment Initiative in Sciences and Medicine (PRISM), Purdue University's Trail Blazers in Engineering and Dr. Brooke Coley's BE GREAT retreat. As a minority PhD candidate in engineering, it's common to experience isolation, feeling as though advancing in your career requires leaving your community behind. These initiatives serve as a reminder of our collective journey towards success – we are not alone!

I want to thank the student groups I was a part of at MIT. Thank you to the Black Graduate Student Association (BGSA) and the University Center for Exemplary Mentoring (UCEM) for providing me with a community outside of the lab, especially during the trying times in 2020 following the murder of George Floyd. Thank you to the Academy of Courageous Minority Engineers (ACME) for holding me accountable as I pursued my goals. Finally, thank you to the MIT Club Rugby team for introducing a washed-up athlete to a whole new sport at the age of 22, for teaching me the game, tolerating my high tackles, and eventually electing me as your President and Captain. Thanks to these clubs, I spent my 6 years at MIT quite HAPPY! :)

Thank you to the lifelong friends I have made at MIT. Shout out to Dayo, Jide, Nick, Junior, Jadal, and Myles whom have made every moment unforgettable. Thank you for the endless laughs and always being down for my shenanigans.

Importantly, I want to thank my family. Thank you to my little brother, Amazi West. Since his birth, I've known that his eyes have been on me, aiming to follow in my footsteps. Knowing my importance in his life as a big brother has often given me the strength I needed to persevere through many challenges. Thank you to my parents. As I've entered adulthood and navigated life on my own, so many of the life lessons they taught me have become increasingly important. Without their mantras like "Work hard *then* play hard," "You have to be your own biggest fan if you want others to believe in you," and "Michael, you need to enact discipline — doing something you don't want to do in order to reach a goal," none of my accomplishments would have ever been possible. I love you guys.

I want to thank my fiancée, Dr. Caroline Ayinon. Firstly, I must credit her

for discovering and assisting me in applying to MSRP, the program that ultimately brought me to MIT. Secondly, it is thanks to her that I decided to pursue graduate studies at MIT. Finally, her unwavering love and support have been constants throughout my journey in graduate school. Thank you for motivating me to strive towards my goals. Thank you for reminding me to take a break when needed. Thank you for picking me up when I was down. Thank you for always reminding me of how great I am. Thank you for always taking the time to celebrate me. Thank you for loving me. I certainly could not have accomplished this without you by my side. I love you.

Lastly, I want to express gratitude to myself. This thesis represents tens of thousands of hours of dedicated research. Yet, what it does not capture are the countless mistakes, moments of self-doubt, and hours of frustration endured along the way. So, I want to finish with a moment of self-love. I want to thank me for doing the hard work. I want to thank me for giving myself grace when I made mistakes. I want to thank me for not giving up on myself during those moments of self-doubt. I want to thank me for taking the time to celebrate my wins and acknowledge my accomplishments. Above all, I want to thank me for being me, staying true to myself, and always showing myself love.

This research was performed in the Eric P. and Evelyn E. Newman Laboratory for Biomechanics and Human Rehabilitation at the Massachusetts Institute of Technology. It was supported in part by, the MIT Summer Research Program, the MIT Office of Graduate Education, Advanced Robotics for Manufacturing Institute (ARM-17-01-TA7), the Ford Foundation Fellowship, the Ben Gold Fellowship, the Takeda Fellowship, the Accenture Fellowship, and the Eric P. and Evelyn E. Newman Fund.

Members of the Committee

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

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

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

Professor Madhusudhan Venkadesan
Associate Professor of Mechanical Engineering & Materials Science
Yale University

Contents

1	Introduction	23
1.1	Current State-of-the-Art in Robotics	23
1.2	A Case for the Study of Human Motor Control	24
1.2.1	The Paradox of Human Performance	24
1.2.2	Rehabilitation Robotics	25
1.3	A Case for Simple Models	27
1.4	A Brief Literature Review	28
1.4.1	Synergies	29
1.4.2	Dynamic Primitives	32
1.5	Summary of Thesis	38
2	Reach-and-Grasp Synergies Differ from Manipulation Synergies	41
2.1	Introduction	41
2.2	Methods	45
2.2.1	Experimental Task	45
2.2.2	Data Acquisition	45
2.2.3	Data Analysis	48
2.3	Results	52
2.3.1	Emergence of Kinematic Hand Synergies	52
2.4	Comparison of Kinematic Hand Synergies	57
2.5	Discussion	68
2.5.1	Emergence of Kinematic Hand Synergies	68
2.5.2	Reach-and-Grasp vs. Manipulation Synergies	70

2.5.3	Limitations	71
2.5.4	Implications	73
2.5.5	Are synergies an epiphenomenon?	74
2.6	Conclusion	76
3	Dynamic Primitives Describe Limitations in Human Force Regulation during Motion	77
3.1	Introduction	78
3.2	Methods	80
3.2.1	Participants	80
3.2.2	Experimental Procedures	81
3.2.3	Experimental Design	83
3.2.4	Dependent Measures and Data Processing	83
3.2.5	Data Analysis and Statistics	84
3.3	Results	86
3.3.1	Performance Improvements	86
3.3.2	Existence of Motion-Dependent Force Errors	88
3.4	Additional Results	91
3.4.1	Performance Improvements	91
3.4.2	Existence of Motion-Dependent Force Errors	91
3.4.3	Differences in Force Error Magnitude	92
3.5	Discussion	92
3.5.1	Force Error	93
3.5.2	Dynamic Primitives	94
3.5.3	Limitations	95
3.5.4	Implications	96
3.6	Conclusions	97
4	Stiffness Can be Perceived via Visual Observation of Path Information	99
4.1	Introduction	100

4.2	Experiment 1: Methods	104
4.2.1	Ethics Statement	104
4.2.2	Experimental Protocol	104
4.2.3	Simulated Arm Motions	105
4.2.4	Task Instruction	111
4.2.5	Self-Reported Strategy for Stiffness Estimation	112
4.2.6	Statistical Analyses	112
4.3	Experiment 1: Results	114
4.3.1	Effects of Experimental Condition and Simulated Joint Stiffness on Stiffness Estimate	114
4.3.2	Effect of Experimental Condition on Motion Features Used to Estimate Stiffness	118
4.4	Experiment 2: Methods	118
4.4.1	Subjects	118
4.4.2	Experimental Protocol	120
4.4.3	Simulated Arm Motions	120
4.4.4	Task Instruction	120
4.4.5	Self-Reported Strategy for Stiffness Estimation	120
4.4.6	Statistical Analyses	120
4.5	Experiment 2: Results	121
4.5.1	Effects of Velocity Profile on Stiffness Estimate with the Same Endpoint Path	121
4.5.2	Self-Reported Stiffness Estimating Strategies	121
4.6	Discussion	124
4.7	Conclusions	133
5	A Synergy-based Complexity Index to Promote the Study of Complex Manipulation	135
5.1	Introduction	136
5.2	Methods	138

5.2.1	Singular Value Decomposition	138
5.2.2	Variance-Accounted-For based Complexity Index	139
5.2.3	Experimental Setup	140
5.2.4	Validation of the Variance-Accounted-For based Complexity Index	145
5.3	Results	147
5.3.1	Validation of the Variance-Accounted-For based Complexity Index	147
5.3.2	Comparison to Significant Number of Synergies	150
5.3.3	A Metric for Unsupervised Learning	150
5.3.4	Robustness to Data Pre-processing	155
5.4	Discussion	155
5.4.1	Limitations	157
5.4.2	Implications	157
5.5	Conclusions	158
6	Conclusions	161
6.1	Potential Future Work	164
6.1.1	Kinematic Hand Synergies	164
6.1.2	Visual Perception of Stiffness	166
6.1.3	Variance-Accounted-For based Complexity Index (VAF-CI)	167
6.2	Outlook	168
A	The Study of Complex Manipulation via Kinematic Hand Synergies: The Effects of Data Pre-Processing	171
A.1	Introduction	172
A.2	Methods	174
A.2.1	Theoretical Background	174
A.2.2	Numerical Validation	176
A.2.3	Piano Playing	177
A.3	Results	179

A.3.1	Numerical Validation	179
A.3.2	Piano Playing	180
A.4	Discussion	184
A.4.1	Effects of Centering the Data	184
A.4.2	Effects of Changing Data Variance	185
A.4.3	A Geometrical Understanding	186
A.4.4	Implications for Robotic Rehabilitation	186
A.5	Conclusion	187

List of Figures

1-1	The Author’s Research Methodology	28
1-2	Santello Hand Synergies [Santello et al., 1998]	31
1-3	1 DOF Mass-Spring-Damper System [West, 2020]	35
1-4	Soft synergy model [Bicchi et al., 2011]	37
2-1	The Mock Electrical Cabinet	46
2-2	Variance-accounted-for (VAF)	52
2-3	Number of Significant Synergies	53
2-4	The First Synergy.	55
2-5	The Second Synergy.	56
2-6	The Third Synergy	57
2-7	Reach-and-Grasp vs. Manipulation Cosine Similarities	59
2-8	Synergy Similitude	61
2-9	Reordered Reach-and-Grasp vs. Manipulation Cosine Similarities	63
2-10	Reordered Synergy Similitude	65
2-11	Reordered Manipulation Synergy Ranks	67
3-1	Direct Force Control vs Indirect Force Control Block Diagram	80
3-2	Human Force Regulation Experimental Setup	82
3-3	Human Force Regulation RMS Force Error across Trial	86
3-4	Human Force Regulation RMS Force Error across Block	87
3-5	Representative Subject Raw Force Data	89
3-6	Auto-correlation of Subjects’ Force	90

4-1	Visual Perception of Stiffness Experimental Setup	105
4-2	Arm Simulations' Endpoint Paths	108
4-3	Simulated Arm Motions	110
4-4	The Arm Simulations' Radius of Curvature and Tangential Velocity Relationship	111
4-5	Visual Perception of Stiffness Experiment 1 Stiffness Estimate Results	115
4-6	Visual Perception of Stiffness Experiment 1 Individual Subject Stiffness Estimate Results	116
4-7	Visual Perception of Stiffness Experiment 1 Coefficient of Determination	117
4-8	Visual Perception of Stiffness Experiment 1 Slope	117
4-9	Self-Reported Strategies in Experiment 1	119
4-10	Visual Perception of Stiffness Experiment 2 Stiffness Estimate Results	122
4-11	Visual Perception of Stiffness Experiment 2 Individual Subject Stiffness Estimate Results	123
4-12	Self-Reported Strategies in Experiment 2	123
4-13	RMS of Temporal Motion Features in Experiment 2	127
4-14	Change in Motion Related Features with Simulated Elbow Stiffness .	129
5-1	Wire Harnessing Experimental Setup	141
5-2	Piano Playing Experimental Setup	144
5-3	VAF of Each Synergy in the Reach-to-Grasp, Tool-Use, and Piano Playing Experiments	148
5-4	VAF-CI in the Reach-to-Grasp, Tool-Use, and Piano Playing Experiments	149
5-5	K-means Clustering Algorithm Results	152
5-6	Average VAF-CI of Individual Tasks	154
6-1	A Descriptive Understanding of Human Manipulation	162
A-1	Piano Playing Experimental Setup	177
A-2	Numerical Validation Data	180

A-3	Eigenvectors and VAF of the Numerical Validation Data	181
A-4	VAF of Each Synergy in the Piano Experiment	182
A-5	Comparison of Synergies across Different Pre-processing Types	183

List of Tables

2.1	Number of Statistically Similar Synergies	60
2.2	Number of Statistically Similar Synergies after Reordering	64
3.1	Comparison of Force Error between Experiments	92
4.1	Subject Self-Report Encoding Criteria	113
5.1	Number of Significant Synergies	149
A.1	Number of Significant Synergies in Piano Playing	182

Chapter 1

Introduction

“Everything should be made as simple as possible, but not simpler.”

–Albert Einstein

1.1 Current State-of-the-Art in Robotics

As of the writing of this thesis, the robotics community has entered a new era, with a keen interest in enhancing robot dexterity through advancements in artificial intelligence (AI) and machine learning (ML). Particularly, there’s been a notable shift towards the development of various imitation learning techniques, where robots strive to replicate human behavior [Hussein et al., 2017]. In such tasks, robots learn the mapping between observations and actions, commonly known as behavior cloning [Pomerleau, 1988]. This approach has enabled robots to acquire a range of skills beyond basic manipulation tasks like pick-and-place operations [Brohan et al., 2022, Padalkar et al., 2023] to more complex tasks involving bimanual mobile operation such as sauteing and serving a piece of shrimp [Fu et al., 2024]. Particularly noteworthy is the work by researchers at the Toyota Research Institute, who have leveraged advancements in generative AI to teach robots new manipulation abilities within a single afternoon [Chi et al., 2023, Chi et al., 2024]. A demonstration of their impressive work can be viewed on YouTube. While this work is quite impressive these

large learning models often suffer from a lack of interpretability, limiting the insight that can be derived.

1.2 A Case for the Study of Human Motor Control

1.2.1 The Paradox of Human Performance

Despite the rapid advancements in robotics led by the AI revolution, humans consistently outperform robots in tasks requiring physical interaction. This discrepancy underscores a fascinating dichotomy; humans, equipped with what might be considered inferior ‘wetware’ and ‘hardware’, excel at dexterously manipulating complex objects and simultaneously managing their own high degree-of-freedom bodies [Kandel et al., 2000a, Hogan and Sternad, 2012]. The limitations of current robotic technologies was highlighted in the DARPA Robotics Challenge. There, robots struggled with tasks trivial for humans, such as walking in sand, traversing a door, and turning a steering wheel [Banerjee et al., 2015, Knoedler et al., 2015]. These scenarios highlight the paradox of human performance: despite seemingly inferior biological mechanisms, humans possess an unparalleled ability to perform manipulation tasks that require physical interaction.

One of the most pronounced differences between human and robot ‘wetware’ and ‘hardware’ lies in the speed at which humans and robots operate and communicate. Human neural pathways and muscles operate at a significantly slower pace compared to computer processors and robotic actuators. For instance, human muscles contract at frequencies below 10 Hz, a stark contrast to robotic actuators which can achieve movements at frequencies well beyond 100 Hz [Kandel et al., 2000a, Hogan, 2017]. Additionally, human communication, facilitated by the transmission of neural signals, is approximately six orders of magnitude slower than robot communication, where electrons travel at speeds of 10^8 m/s compared to the 10^2 m/s speed of neuron signal transmission [Kandel et al., 2000a, Hogan, 2017]. Feedback delays present another stark contrast; the human trans-cortical loop incurs feedback delays greater

than 100 ms, whereas robotic systems can often boast feedback delays of less than 1 ms [Kandel et al., 2000a, Hogan, 2017]. These disparities highlight the mechanical and communicative efficiencies that robots possess over humans, emphasizing the paradoxical nature of human superiority in tasks involving physical interaction.

Another area where humans seemingly face a disadvantage is in managing the “curse of dimensionality” [Bellman, 1966]. This principle posits that control complexity increases exponentially with the addition of more degrees of freedom. Despite this, humans exhibit a remarkable ability to consistently manage their musculoskeletal system, which consists of over 200 joint degrees of freedom and 600 muscle actuators [Kandel et al., 2000a]. This capability suggests that humans employ highly efficient, yet currently not fully understood, strategies to overcome the challenges posed by high-dimensional control spaces.

This performance gap between human and robotic capabilities suggests that humans may employ simplifying control techniques to manage physical interactions – a concept that roboticists have yet to replicate successfully. The study of human motor control, learning, and perception is crucial in unraveling these underlying control mechanisms. By delving into the complexities of how humans achieve such dexterous manipulation, we can uncover insights that not only advance our understanding of human biology but also inform the development of more capable and dexterous robots. This thesis makes strides towards identifying these underlying, simplifying control techniques in an effort to bridge this gap.

1.2.2 Rehabilitation Robotics

The study of human motor control holds significant relevance for the field of rehabilitation robotics, including therapeutic robotics, assistive robots, exoskeletons, and prostheses. A hallmark of human performance is our remarkable ability to manipulate objects and tools with our hands, enabling us to interact effectively with the world around us. Unfortunately, each year approximately 15 million people worldwide and around 800,000 Americans suffer from loss of motor function due to cerebral vascular accident (CVA) (i.e. stroke) [Virani et al., 2020], with hand function

often being the most severely affected and experiencing the poorest recovery rates. Additionally, in the United States alone, there are approximately 185,000 amputations annually [Ziegler-Graham et al., 2008]. These stark statistics emphasize the pressing need for effective rehabilitation solutions. Consequently, extensive investments have been made in therapeutic devices (amounting to approximately \$53 billion [Tsao et al., 2022]) and prosthetic devices (valued between \$8-12 billion [Ziegler-Graham et al., 2008]) to address these challenges.

There has been demonstrated effectiveness and cost efficiency of robotic rehabilitation techniques over standard patient care in tasks like upper limb reaching [Fasoli et al., 2004, Krebs et al., 2003, Riener et al., 2005]. However, despite these promising developments the widespread adoption of such technologies remains limited, signaling a need for further innovation and refinement [Balasubramanian et al., 2010, Maciejasz et al., 2014]. Challenges hindering the success of rehabilitation robotics include high costs, the complexity of human-machine interfaces, and a general lack of customized treatment options that cater to individual patient needs [Calabrò et al., 2020, Qassim and Wan Hasan, 2020]. For review see, [Yue et al., 2017, Frolov et al., 2018, Li et al., 2017, Liu et al., 2022].

More importantly, a fundamental challenge in rehabilitation robotics is the lack of a comprehensive understanding of the control mechanisms humans use to produce movement [Rodgers et al., 2019]. This gap in knowledge raises a critical question: if we do not fully understand how humans naturally produce movement, how can we effectively relearn or restore lost motor functions? The study of human motor control serves to bridge this gap, providing the foundational knowledge necessary to develop robotic systems that can more accurately mimic or support human movements. By unraveling the complexities of how humans learn, control, and adapt movements, researchers can create rehabilitation robots that offer more natural and effective assistance or replacement for lost functions.

In conclusion, the intersection of human motor control study and rehabilitation robotics holds the potential to revolutionize the way we approach rehabilitation for motor impairments. By grounding the development of therapeutic robotics

in a thorough understanding of human motor functions, we pave the way for innovations that can significantly improve the quality of life for those affected by CVA, amputation, and other conditions that impair motor function. This thesis takes a step in that direction.

1.3 A Case for Simple Models

In the domain of robotics, the inclination towards leveraging large, complex models for understanding physical interactions has been primarily driven by the inherent challenges these interactions present. While these sophisticated models undoubtedly offer thrilling results [Chi et al., 2023, Chi et al., 2024, Fu et al., 2024], their complexity often obscures their interpretability, making them less accessible for intuitive understanding. This observation underscores the importance of embracing simpler models, particularly when the aim is to extract immediate, easily understandable insights.

The pursuit of simple models is further justified when considering the mechanisms of human motor control. It is reasonable to assume that humans do not engage in detailed, conscious modeling of their interactions with the environment during routine tasks. Taking the everyday activity of opening a door as an illustrative example, the mechanical requirement is to apply a force normal to the door's surface, a direction that dynamically changes with the door's movement. Despite the complexity implied by this interaction, the cognitive process for most individuals does not involve conscious calculation of these mechanics. Instead, humans intuitively grasp and pull the door handle in a direction that feels approximately correct. This process, largely automatic and devoid of explicit modeling, exemplifies the humans' innate ability to efficiently navigate complex physical interactions through intuitive understanding and execution.

The research methodology underlying the studies in this thesis, illustrated in Figure 1-1, follows a systematic approach to discerning the simplifying control mechanisms present in human motor tasks. It commences with a well-founded

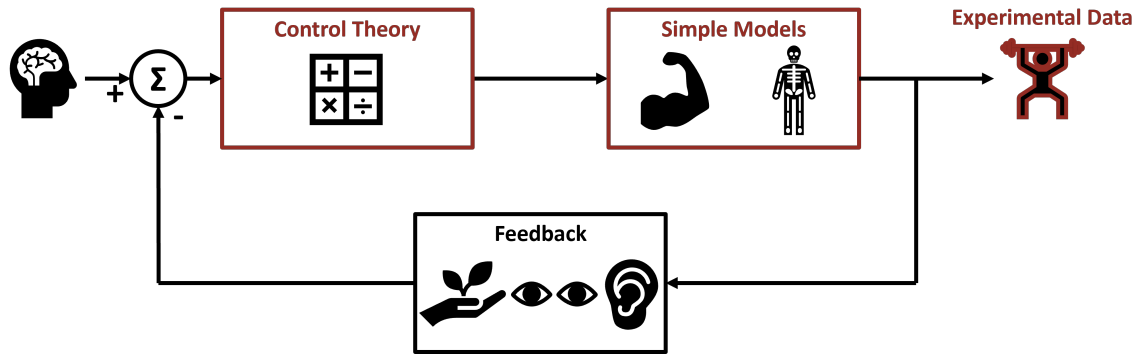


Figure 1-1: **The author’s research methodology** used throughout this thesis to uncover simplifying control mechanisms present in human motor tasks.

hypothesis, guiding the design of tailored behavioral studies aimed at investigating specific aspects of human motor control. These studies entail collecting data, primarily kinematics and occasionally kinetics, as participants engage in meticulously crafted experimental tasks. Subsequent analysis involves a thorough examination of the collected data to unveil underlying movement patterns and principles, accomplished through the development and application of simple yet effective models prioritizing clarity and interpretability. The final phase applies principles of control theory to further scrutinize observed behaviors, enabling inference of aspects of the internal human motor controller. By comparing modeled behavior with actual human performance, discrepancies and alignments are identified, shedding light on the internal processes governing motor control. This iterative process of hypothesis formulation, experimental investigation, modeling, and theoretical analysis forms the foundation of our approach to unraveling the simplifying controllers inherent in human motor tasks.

1.4 A Brief Literature Review

In this section, we will briefly highlight key discoveries in human motor control that have uncovered competent, simplifying models. Specifically, we will discuss synergies and dynamic primitives, as these concepts are important themes in this thesis.

1.4.1 Synergies

As mentioned earlier, a difficulty of managing physical interaction is the control of the complex, high degree-of-freedom system that is the human musculoskeletal system, which consists of 200+ joint degrees-of-freedom and 600+ muscle actuators [Kandel et al., 2000b, Bellman, 1966]. In tackling the complex nature of human motion, Bernstein formulated what is known as the *Degrees of Freedom Problem* which stated that [Bernstein, 1967, Bernstein, 2014]:

“It is clear that the basic difficulties for coordination consist precisely in the extreme abundance of degrees of freedom, with which the [nervous] centre is not at first in a position to deal.”

This quote highlights the complexity inherent in controlling movement, positing that such control might not involve managing each joint independently. Instead, he suggested that movements are orchestrated in a more coordinated and integrated manner, reducing the system’s overall complexity. This foundational idea paves the way for understanding how synergies can streamline hand movement control.

A synergy, in this context, is defined as a coordinated action whereby multiple degrees of freedom are regulated collectively to produce movement. Given the coordinates of the human action, θ , it is likely that many of these degrees of freedom are controlled simultaneously. If so, there exists a mapping such that

$$\theta = \phi(\delta) \tag{1.1}$$

where $\delta(t)$ is a command that has fewer dimensions than $\theta(t)$, and $\phi(\cdot)$ is a mapping between δ and θ . Direct access to δ is not immediately available. However, assuming it exists, the mapping is often found via dimensionality reduction techniques such as principal component analysis [Smith, 1988], singular value decomposition [Mason et al., 2001, West et al., 2023], and non-negative matrix factorization [D’Avella and Tresch, 2001, Tresch et al., 2006].

Depending on the domain in which the human data was collected, we can extract

neural synergies [Cheung and Seki, 2021, Santello et al., 2013], muscular synergies [Cheung et al., 2009, D’avella and Tresch, 2001, Weiss and Flanders, 2004, Ting and McKay, 2007, Yang et al., 2019], or kinematic synergies [Santello et al., 1998, Todorov and Ghahramani, 2004, Mason et al., 2001]. At the muscle level, synergies reflect the coordination among muscle activations, providing insight into how muscular efforts are pooled to facilitate movement. Similarly, at the neural level synergies represent the underlying neural commands that orchestrate these muscular actions. Overall, synergies are interpreted as a dimensionality reduction method adopted by humans to simplify control, fundamentally transforming a high-dimensional task into a lower-dimensional one. Together, these levels of analysis—kinematic, muscle, and neural—offer a multidimensional view of how synergies simplify the control of the neuromuscular system, providing a framework for both understanding natural movements and informing the design of prosthetic and rehabilitative technologies. By leveraging synergies, researchers and engineers can devise systems that more closely mimic natural hand movements, thereby enhancing the effectiveness of interventions aimed at restoring or replacing lost motor functions. For more comprehensive coverage, we direct the reader to [Santello et al., 2016].

This thesis focuses particularly on kinematic hand synergies, initially studied by Santello et al. [Santello et al., 1998]. In their study, subjects were asked to hold their right hand in a static posture, as if grasping 57 different household objects. Notably, the objects were neither physically nor virtually present; participants were required to imagine them. Principal component analysis revealed that the first two synergies explained over 80% of the variance in hand postures. The first synergy typically involved significant flexion and slight adduction of the four fingers, resembling a power grasp [Napier, 1956], while the second synergy entailed significant thumb inward rotation, index finger flexion, and moderate middle finger flexion, akin to a pinch grasp [Napier, 1956] (Figure 1-2).

Similar findings were reported by Weiss and Flanders [Weiss and Flanders, 2004], who observed analogous results when subjects spelled the alphabet using American Sign Language. Subsequent research extended the study of synergies to experiments




Figure 1-2: Postural synergies defined by the first two principal components. The hand posture at the center of the PC axes is the average of 57 hand postures for one subject. The postures to the right and left are for the postures for the maximum (max) and minimum (min) values of the first principal component (PC1), coefficients for the other principal components having been set to zero. The postures at the top and bottom are for the maximum and minimum values of the second principal component (PC2). Figure from Santello et al. [Santello et al., 1998]

involving object contact and temporal hand movement during reach-to-grasp tasks. Mason et al. [Mason et al., 2001] employed Singular Value Decomposition to analyze the temporal evolution of postural synergies throughout reach-to-grasp actions. Further analysis by Santello et al. [Santello et al., 2002] explored the temporal evolution of synergies across three experimental conditions: memory-guided movements, movements with virtual object representation, and movements with physically present objects. Remarkably, the first two synergies persisted across conditions, accounting for over 75% of variance in hand posture (Figure 1-2).

Overall, research has consistently demonstrated that a small number of synergies, primarily characterized by a power grasp, can explain a significant portion of variance across a diverse range of tasks. This suggest that synergies may be used as a simplifying controller to manage human motor control in the face of a high degree-of-freedom neuromuscular system.

However, much of the current work has focused on the kinematics of simple grasp and not on object manipulation or functional tool-use. If we aim to restore motor

function to those who have suffered from CVA or amputation, we must study hand manipulation during functional tool use. This prompts a question, do hand synergies differ during grasp and manipulation? Chapter 2 of this thesis reports an explicit comparison of grasp synergies and manipulation synergies.

1.4.2 Dynamic Primitives

Commonly, the central nervous system is believed to develop an “internal model” of an object’s dynamic behavior to facilitate control, enabling the generation of predictive forward-path control inputs for desired actions. The theory of dynamic primitives posits that motor behavior, whether involving physical interaction or not, is crafted using a limited set of primitive dynamic behaviors, acting as the fundamental “building blocks” of more intricate actions [Hogan and Sternad, 2012, Schaal and Sternad, 2001, Schaal, 2006, Degallier and Ijspeert, 2010]. In [Hogan and Sternad, 2012], researchers build upon mounting evidence indicating that sensorimotor control is based on a composition of primitive dynamic actions [Sternad et al., 2000, Thoroughman and Shadmehr, 2000, Flash and Hochner, 2005, Kargo and Giszter, 2008, Sternad, 2008, Sing et al., 2009, Degallier and Ijspeert, 2010, Dominici et al., 2011]. They propose that “human motor control is encoded solely in terms of *primitive dynamic actions*.” The authors identify three distinct primitives - submovements, oscillations, and mechanical impedances - to account for a wide range of human behaviors, emphasizing that encoding in terms of parameterized primitives is crucial for learning, performance, and retention of complex skills. The theory suggests that these primitives can be combined to produce observable forces and motions, possibly through a virtual trajectory composed of submovements and/or oscillations interacting with impedances.

Submovements

The first dynamic primitive we will discuss is submovements. Submovements are primitive dynamic elements of motor behavior; they resemble discrete movements

but may overlap in time [Hogan and Sternad, 2012]. They can be viewed as building blocks for more complex movements, assembled sequentially or in overlap to achieve fluid motion.

The concept of movements being composed of submovements is not new and has historical precedence in observations by [Woodworth, 1899] (for review, see [Meyer et al., 1988, Elliott et al., 2001]). However, more recent support for this theory stems from studies examining the stereotyped speed profiles of submovements in stroke patients [Krebs et al., 1999] and the development of infant reaching movements [Von Hofsten, 1991, Berthier, 1996]. These studies highlight that submovements are indeed a primitive element of human motor control. Moreover, studies have shown that such a primitive element is useful in facilitating the prediction of sensory consequences of actions and the detection of errors [Burdet and Milner, 1998].

Hogan and Sternad [Hogan and Sternad, 2012] further define submovements as attractors that describe smooth, sigmoidal transitions of a variable, with a stereotyped time profile indicative of a “trajectory attractor”. This definition emphasizes the control of trajectory rather than final position, challenging the notion of a point attractor as sufficient for describing the dynamic process of motor control [Bizzi et al., 1984, Won and Hogan, 1995]. The combination of submovements, through mechanisms like linear vector superposition, allows for the generation of complex movement trajectories, highlighting the flexibility and adaptability of this primitive-based approach for motor control [Flash and Henis, 1991].

Oscillations

Oscillations represent a dynamic primitive that facilitates rhythmic movement. While mathematically, an oscillation can be perceived as a combination of back-to-back submovements, rhythmic movement has a well-established presence in the literature and is phylogenetically ancient. Numerous studies have demonstrated the biological distinction between oscillations and submovements, as oscillatory behavior can be evoked and sustained with minimal intervention from the central nervous system [Brown, 1911, Brown, 1914, Grillner and Wallen, 1985, Schaal et al., 2004].

Furthermore, it has been observed that there is a limit to rhythmic actions; as the rhythmic period increases (to approximately 2 to 5 seconds), actions can no longer be perceived as rhythmic [Fraisse, 1984, James, 1890]. Additionally, experimental evidence suggests that executing smooth rhythmic motions very slowly is challenging for humans; instead, they tend to default to a motion composed of overlapping submovements [Park et al., 2017]. These limitations of rhythmic motion underscore the distinction between submovements and oscillations.

Mechanical Impedance

Submovements and oscillations may lay the groundwork for unconstrained motion, but contact and physical interaction are vital for the hallmark of human performance: object manipulation and tool utilization. In tasks necessitating physical interaction, the challenge lies in controlling both force and motion. One approach to address this challenge is impedance control, which governs the relationship between motion and force and is frequently modeled by the equation 1.2,

$$F = Z\{x_0 - x\} \tag{1.2}$$

where $Z\{\cdot\}$ is the mapping from a force to a motion (i.e. the impedance). Theoretically, this mapping may be any linear, or non-linear, operator. It was introduced by Hogan, [Hogan, 1985a, Hogan, 1985b, Hogan, 1985c] and has since been supported as a plausible description of human motor control.

To better understand impedance control, refer to the 1 DOF example in Figure 1-3. Figure 1-3 shows a mass-spring-damper system whose equation of motion, in the absence of external forces takes the form,

$$m\ddot{x} = k(x_0 - x) + b(\dot{x}_0 - \dot{x}) \tag{1.3}$$

and the control law takes the form,

$$F_{act} = k(x_0 - x) + b(\dot{x}_0 - \dot{x}) \tag{1.4}$$




Figure 1-3: A diagram of a 1 DOF mass-spring-damper system. All displacements and forces pointing rightward will be considered positive. Figure from [West, 2020].

where the control input, F_{act} (actuator force), is determined by the zero-force trajectory, x_0 , which represents the trajectory x would follow in the absence of external forces, and is linked to the actual trajectory, x , through a rigid body of mass, m , spring of stiffness, k , and damper of damping, b . In the Laplace domain, the forward path dynamics of this controller are,

$$\frac{X}{X_0} = \frac{bs + k}{ms^2 + bs + k} \quad (1.5)$$

Intuitively, one can think of x_0 as a planned motion that is “pulling” the actual motion, x , via a mass-spring-damper connection.

In the presence of an external force, F_{ext} , the disturbance response is modeled by equation 1.6,

$$Z = \frac{F_{ext}}{X} = ms^2 + bs + k \quad (1.6)$$

where Z is the impedance. This impedance characterizes the interactive behavior of the object (i.e. the relation between force and motion).

In the realm of human motor control, mechanical impedance emerges as a third category of dynamic primitives alongside submovements and oscillations. Humans can voluntarily modulate mechanical impedance through the co-contraction of antagonist

muscle groups, leading to an increase in net mechanical impedance due to the positive correlation between muscle impedance and muscle force [Nichols and Houk, 1976, Hoffer and Andreassen, 1981]. This mechanism highlights how mechanical impedance at a joint results from the combined contributions of each involved muscle group. Furthermore, mechanical impedance is subject to modulation by neural feedback mechanisms, with muscle spindles and Golgi tendon organs playing significant roles at the spinal level or above [Prochazka et al., 2000]. The adaptability of the feedback pathway gains underscores the complexity and flexibility of the system.

Hogan [Hogan, 1990] posits that mechanical impedance is particularly robust during contact and interaction with objects, influencing the dynamics of hand and arm movements and thereby altering behavioral output. This robustness, independent of object behavior or contact, underscores the qualification of mechanical impedance, driven by the neuromechanical system dynamics, as a dynamic primitive. This independence ensures behavior is not contingent on external interactions but on intrinsic system properties.

Remarkably, mechanical impedance, when integrated with skeletal inertia, permits the linear superposition of nonlinear impedances, enabling the control of interaction tasks across varying conditions [Hogan, 1985a, Toffin et al., 2003, Hogan and Buerger, 2018, Franklin et al., 2007]. Such capabilities indicate that modulating mechanical impedance is essential for managing tasks involving physical contact, solidifying its status as a foundational element of human motor behavior.

Mechanical impedance, when coupled with kinematic hand synergies acting as the zero-force trajectory (Eq. 1.2), plays a crucial role in managing grasping actions. This approach is exemplified in the soft synergy model, which emphasizes the importance of impedance modulation for ensuring grasp stability. Through this model, the dynamic equilibrium of the hand is maintained by two opposing force fields: one governed by the desired hand motion (a given synergy), and the other by the object, repelling the hand from penetrating it (Figure 1-4). The hand's ability to achieve a stable grasp depends on both modulating sufficient mechanical impedance and utilizing the initial synergies. Without precise control over these primary synergies,




Figure 1-4: The reference hand moves on the synergy manifold (a–d) and represents an attractor for the real hand (e–h), which is repelled by contact forces with the object. The resulting configuration is ultimately dictated by the hand mechanical impedance and control stiffness. Figure from Bicchi et al. [Bicchi et al., 2011]

achieving stability requires the engagement of numerous higher-order synergies. In sum, the soft synergy model serves as a simplified control model that highlights the interplay between mechanical impedance and kinematic synergies. Moreover, the soft synergy model has become a pivotal control strategy for numerous prosthetic and anthropomorphic robotic hands [Laffranchi et al., 2020, Varol et al., 2014, Brown and Asada, 2007, Catalano et al., 2014], underscoring its significance in both theoretical understanding and practical applications in robotics.

Additionally, using submovements as the reference trajectory of an impedance controlled robot has been used to successfully simplify robot control during tasks involving complex dynamics and interaction. Specifically, a controller based on dynamic primitives has been used successfully (in simulation) to control a 2 DOF arm to manipulate a dynamically complex whip with 50 DOF in a targeting task [Nah et al., 2020, Nah et al., 2021, Nah et al., 2023a]. Moreover, work has been done to demonstrate that dynamic primitives, renamed to elementary dynamic actions, facilitate physical robot control during contact-rich manipulation due to the embedded modularity [Nah et al., 2023b, Lachner et al., 2023]. This work reaffirms the transformative potential of simplified control models derived from the study of human motor control to advance robotic applications.

The research detailed in this literature review highlights the exciting prospect

of leveraging simplified controllers inspired by studies of human motor control for innovation in both motor neuroscience and robotics. This exploration of simplified control mechanisms, rooted in the innate efficiency of human motor strategies, forms the cornerstone of this thesis. Moreover, this endeavor not only deepens our understanding of the human body but also unlocks new possibilities in creating machines capable of extending, augmenting, and replicating human capabilities, marking a pivotal advancement in the evolution of human knowledge and mechanical innovation.

1.5 Summary of Thesis

The goal of this thesis is to provide a descriptive understanding of human manipulation in unimpaired subjects. To do so, it navigates the complex terrain of human motor control and its implications in robotics, driven by the quest to uncover the principles behind human dexterity. Each chapter builds upon the idea that uncovering simplified control mechanisms used by humans can significantly influence both theoretical knowledge and practical applications in rehabilitation and robotic design.

Starting off, Chapter 2 delves into the realm of kinematic synergies, shedding light on the subtle yet critical distinctions between the mechanisms underlying reach-and-grasp actions and those governing the manipulation of objects and tools. The chapter reveals that manipulation tasks involve a greater number and distinct subset of hand synergies compared to simple grasping. These results underscore the intricate nature of functional hand use and the need to specifically study *functional hand manipulation*, not just grasping.

In Chapter 3, our focus shifts to human force regulation, aiming to address an overlooked question from Chapter 2: how do humans manage force during physical interaction? Specifically, we inquire whether humans can manage force independently of motion. Contrary to longstanding assumptions in robotics, this study unveils that in humans control of force and motion are intertwined through mechanical impedance,

urging a reevaluation of control strategies in human-robot interactions.

In Chapter 4, we build upon Chapter 3 by investigating whether mechanical impedance is essential to the human internal model. To address this, we draw upon the well-cited relationship between human motor action and motor perception [Meltzoff and Moore, 1997, Grush, 2004, Bandura, 1986], asking the question: Can humans perceive mechanical impedance? Our findings reveal that an aspect of mechanical impedance, stiffness, can indeed be perceived, primarily through path information. These results not only offer insight into the mechanisms underlying human motion perception but also carry significant implications for the robotics and rehabilitation fields.

In Chapter 5, we aim to address a question subtly provoked in Chapter 2. While Chapter 2 established that synergies may simplify motor control during dexterous manipulation tasks, it also highlighted their task-specific nature. This observation suggests that we may have access to more synergies than degrees of freedom [Tessari et al., 2024], and that synergies may serve as a means to learn and store new behaviors in a library to be accessed during specific tasks. This raises the challenging question of how to ensure that the tasks we study represent the full spectrum of complex manipulation activities humans can perform with their hands. Chapter 5 seeks to answer this question by proposing a new metric: the Variance-Accounted-For based Complexity Index (VAF-CI). Validated across multiple dexterous tasks, ranging from basic reaching-to-grasp to piano playing, this novel metric can be leveraged to ensure that future studies of hand manipulation probe a wider range of the hands' capabilities, and related synergistic behavior.

In Chapter 6, the main findings of the work are reviewed and future work directions are discussed.

Not to be overlooked, in Appendix A, we delve into an important aspect of synergy decomposition: data pre-processing. The analysis presented therein confirms the efficacy of removing the mean as the appropriate data pre-processing step prior to synergy decomposition.

Together, these chapters leverage the study of human motor control, learning, and

perception through simple, easily interpretable models to weave a narrative that not only advances our understanding of human motor control but also lays the groundwork for innovations in robotics and rehabilitation. The aim is to restore or augment human hand function through more intuitive designs inspired by the natural efficiency of human movement.

As this thesis integrates peer-reviewed published papers (e.g., in Chapters 3, 4, and 5, and Appendix A), some redundancy may arise. This redundancy serves the purpose of reinforcing the context of each study, aiding in the comprehensive understanding of the research findings.

A video presentation of this thesis can be found at:
<https://www.youtube.com/watch?v=u2eCJHqEGww>

Chapter 2

Reach-and-Grasp Synergies Differ from Manipulation Synergies

In this chapter, we uncover new insights into kinematic synergies during functional human hand manipulation of objects and tools, through the study of wire harness installation. It emphasizes the nuanced distinctions between functional hand manipulation and simple grasping, revealing that manipulation tasks require a greater number and distinct subset of hand synergies compared to simple grasp actions. This research marks a significant step toward appreciating the intricacies of hand coordination in complex tasks beyond grasping.

This chapter is an adapted version of a paper that was submitted to the Journal of Neurophysiology titled, “Kinematic Hand Synergies Differ Between Reach-and-Grasp and Functional Object Manipulation.” This work was done in collaboration with Neville Hogan.

2.1 Introduction

A hallmark of human behavior is the ability to use tools to help us manage our environment. For example, in reading this thesis, you have likely entered a car, driven to the office, opened a computer, and typed notes on the screen or written them out using a pen and paper. The steering wheel of the car, the keyboard of

the computer, and the pen are all examples of tools that are manipulated by the hand to facilitate human life. The dexterity of our hand allows us to manipulate a wide array of tools, a feat that not many other animals are capable of [Diogo et al., 2012, Kenward et al., 2005, Napier, 1956, Weir et al., 2002].

Unfortunately, many people worldwide suffer from injury that either inhibits or eliminates their hand dexterity. For example, 15 million people worldwide and 800,000 Americans suffer from loss of motor function due to Cerebral Vascular Accident (CVA) every year [Virani et al., 2020]. Moreover, hand function is often most severely affected and experiences the worst recovery rates post CVA. Additionally, 23 million people worldwide [McDonald et al., 2020] and 700,000 Americans [Ziegler-Graham et al., 2008] are currently living without their hands due to amputation; the number of Americans with limb amputations is expected to double by the year 2050 [Ziegler-Graham et al., 2008]. In a world designed for humans with dexterous hands, the loss of a functioning hand presents a great challenge, compromising the ability to perform activities of daily life.

In an effort to better rehabilitate those who have lost motor function post CVA and to design prosthetic hands for those who have lost their limbs, previous research has aimed to understand human hand motion. However, kinematic analysis of the human hand is challenging due to the 20+ degrees of freedom of the various joints [Kandel et al., 2000a]. Knowledge of human biomechanics and sensory-motor control may simplify this analysis. It is likely that many of these degrees of freedom are controlled in functional groups, based on biomechanical (and possibly neural) constraints. Commonly known as synergies, these groupings reduce the degrees of freedom in the human hand. Mathematically, if $\theta \in \mathbb{R}^n$ denotes the configuration (joint angles) of the hand, a kinematic synergy may be described by a mapping $\theta = \phi(\delta)$ where $\delta \in \mathbb{R}^m$ is a descending neural command with fewer dimensions than θ , i.e. $m < n$, and $\phi(\cdot)$ is a time-invariant mapping such that $\theta(t) = \phi(\delta(t))$. With this assumption, a linear approximation of the kinematic mapping $\phi(\cdot)$ may be found via Principal Component Analysis [Smith, 1988] or Singular Value Decomposition, two dimensionality reduction methods.

Inspired by the work of Bernstein [Bernstein, 1967, Bernstein, 2014], Santello et al. first studied kinematic synergies of the human hand in 1998; they found that two synergies accounted for more than 80% of the variance in hand posture during grasping of a set of 57 imagined objects [Santello et al., 1998]. Other researchers have found similar results in tasks that involve object contact [Santello et al., 2002], reach and grasp [Mason et al., 2001], object manipulation [Todorov and Ghahramani, 2004], and even the American sign language [Weiss and Flanders, 2004]. For review see [Santello et al., 2016].

This decrease in the hands’ operational degrees of freedom may enable simpler control. Experimentally, it affords researchers a means to simplify kinematic analysis. However, much of the current work has focused on the kinematics of simple grasp and not on object manipulation or functional tool-use. If we aim to restore motor function to those who have suffered from CVA or amputation, we must study hand manipulation during functional tool use. This prompts a question, do hand synergies differ during grasp and manipulation? This paper reports an explicit comparison of grasp synergies and manipulation synergies.

A study of human manipulation may also inform robotic manipulation [Mason, 2018]. In the robotics community, robots commonly pick and place objects [Correll et al., 2018, Eppner et al., 2018, Hernandez et al., 2017, Yu et al., 2016]. More recently, robots have demonstrated limited forms of object manipulation, such as finger gaits [Abondance et al., 2020, He and Ciocarlie, 2021, Khandate et al., 2021, Morgan et al., 2022, Rodriguez et al., 2012, Zeng et al., 2018] and object repositioning within grasp [Chavan-Dafle et al., 2019, Dafle et al., 2014, Hogan et al., 2018]. However, few robots exhibit the dexterity and versatility of humans. Thus, studying how humans manipulate tools during functional tasks may help inform the robotics community.

Robot control is further complicated during tasks that involve manipulation of deformable objects [Chi et al., 2022, Herguedas et al., 2019, Hopcroft et al., 1991, Sanchez et al., 2018, Yamakawa et al., 2011, Zhang et al., 2021]. The inability of robots to manipulate deformable objects is especially evident in the manufacturing process known as wire-harnessing. A wire harness is the assembly of electrical cables

used in machinery, such as aircraft, automobiles, and other systems with electrical components. The process of installing these wire harnesses into their respective electrical systems is known as wire-harnessing. Generally, on an assembly line, robots can often be used to quickly and accurately assemble a system of rigid parts. However, due to the non-rigidity of wire harnesses, automating this process is challenging [Sanchez et al., 2018]. Consequently, wire-harnessing is performed manually, creating a bottleneck in the assembly process. Humans’ ability to outperform robots in such a task is interesting as humans have much slower actuation, communication and computation than robots [Hogan, 2017, Hogan and Sternad, 2012]. This paradox motivates research to investigate how humans perform this task. To bridge the gap in our knowledge of how the human hand manipulates tools and complex objects, we studied wire-harnessing.

We aimed to test two hypotheses that compared the number of reduced degrees of freedom and the form of the reduced degrees of freedom:

1. The number of synergies underlying manipulation is the same as those underlying reach-and-grasp.
2. The identity of synergies underlying manipulation is different from those underlying reach-and-grasp.

To test these hypotheses, we measured human hand motion during two experiments, in the context of wire harness installation. In the first experiment, subjects reached for and grasped a tool or object commonly used in wire harness installation. In the second experiment, subjects manipulated those objects and tools to install a wire harness on a mock electrical cabinet. We found that more synergies were required to competently describe the manipulation data than the reach-and-grasp data. Moreover, upon comparing the kinematic hand synergies across reach-and-grasp and manipulation, we found that the difference was predominantly in the higher-order synergies. If we want to provide therapeutic technologies that restore hand function to those who have lost it, we must study and understand these differences; here, we provide a starting point.

2.2 Methods

2.2.1 Experimental Task

Seven adults, with no known history of neurological or musculoskeletal problems, participated in this study (3 women and 4 men, ages ranging from 18 to 28 years old). All subjects were right-handed. Participants were informed about the experimental procedure and agreed to sign a consent form. All procedures were approved by MIT's Institutional Review Board.

2.2.2 Data Acquisition

Each subject performed two different tasks. The first task, which will be referred to as the Reach-and-Grasp Experiment, involved reaching, grasping, and picking up a tool. Specifically, subjects were asked to reach for and grasp a pair of scissors, a Zip tie, a screwdriver, a wire harness with its branched ends tied, and a wire harness with its branched ends untied. They were specifically told to grasp the object as if they were going to use it. Subjects repeated this four times for each object. The second task was designed to emulate wire-harnessing in a manufacturing plant; subjects were required to use the tools from the first task to install a wire harness onto a mock electrical cabinet (Fig. 2-1). This task will be referred to as the Manipulation Experiment; it was performed in five distinct steps. At each step, subjects were verbally told what to do. They were also given a booklet with written instructions and a picture of the completed step. The steps were:

1. Zip tie the branched ends of the wire harness at three specified points.
2. Use the scissors to cut off the excess tips of the Zip ties.
3. Use the U-brackets, provided screws, and screwdriver to screw the wire harness into the top of the mock electrical cabinet.
4. Route the wire harness through the hooks on the mock electrical cabinet.
5. Plug the connector into the socket on the mock electrical cabinet.

These steps were chosen to simulate an assembly task that requires both dexterous manipulation of a deformable object and functional tool use.



Figure 2-1: after a subject had finished the Manipulation Experiment. On the bottom half of the wire harness are the zip ties that were used to tie the branched ends of the wire harness (Step 1 and 2). On the top half of the mock electrical cabinet are the u-brackets used to screw the wire harness on the mock electrical cabinet (Step 3). The red hooks were placed so that the subjects could route the wire harness (Step 4). A 3D printed socket appears above the right most red hook for subjects to plug the wire harness connector into (Step 5).

While subjects performed these tasks, we measured their hand posture using

a CyberGlove (CyberGlove; Virtual Technologies, Palo Alto, CA), a glove with embedded sensors that measure joint kinematics. Specifically, the flexion of the distal interphalangeal (DIP), proximal interphalangeal (PIP), and metacarpophalangeal (MCP) joints of the four fingers were measured. Additionally, the abduction (ABD) of the four fingers at the MCP joints was measured. At the thumb, the flexion at the MCP and interphalangeal (IP) joints, abduction (ABD) at the carpometacarpal joint, and rotation (ROT) about an axis passing through the trapeziometacarpal joint were measured. Lastly, palm arch (PA) and wrist (W) pitch and yaw were measured. Throughout both experiments, subjects wore a CyberGlove on each hand. However, only the data collected from the right hand is presented here. All subjects reported being right hand dominant. In all cases, the CyberGlove collected samples at ~ 200 Hz with a nominal angular resolution of $< 0.1^\circ$.

In the Reach-and-Grasp Experiment, subjects started with their hands flat on a table. After three seconds of data collection, subjects were verbally instructed to grasp the object placed 30cm in front of their hand. After 15 seconds, data collection ended. This was repeated 20 times (4 trials each for 5 different objects). In the second Mock Wire-Harnessing task (the Manipulation Experiment), subjects started with their hand in the same initial position, and the tools to be used were similarly placed 30 cm in front of them. After three seconds of data collection, subjects were verbally instructed to begin the specific step of the wire harness assembly task. Data collection terminated when subjects deemed they were done with the task and returned their hands to the initial position. All subjects were able to successfully complete the task. Notably, all subjects used their right hand as their primary tool-using hand.

To detect movement onset, we calculated average joint velocity, and then smoothed it using a sliding window of 5 adjacent samples. Then, movement was detected when the smoothed displacement was greater than 5% of its maximum value computed over the entire data set. Thus, in the Reach-and-Grasp Experiment, reach began at the instant movement was detected, and ended after the last movement was detected. After movement concluded, we assumed that the grasp phase had commenced. The grasp phase concluded with the end of data collection. Thus, the Reach-and-Grasp

Experiment was partitioned into two sets of data; the first set contained the time-sampled measurements of 23 joint angles during reach, and the second set contained the time-sampled measurements of 23 joint angles during grasp. Then, we re-partitioned each set of data to 100 bins.

In accordance with the work done by [Mason et al., 2001], we analyzed reach and grasp together. So, the two data sets were recombined to form the reach-and-grasp data of one subject during one trial, and the reach portion was evenly weighted with the grasp portion of the data set. In the Manipulation Experiment, we used the data after movement detection until the end of collection for a given step. Moreover, we interpolated the data in each step to 1000 bins as the Manipulation Experiments lasted longer than the Reach-and-Grasp Experiments.

2.2.3 Data Analysis

Emergence of Kinematic Hand Synergies

To extract kinematic hand synergies, we used the singular value decomposition (SVD) method introduced by [Mason et al., 2001] to analyze the evolving hand postures in both tasks. In the Reach-and-Grasp Experiment, the data for each subject was formed into matrix $X_{1,0} \in \mathbb{R}^{4000 \times 23}$, where the columns consisted of the 23 joint angles that were measured, and the 4000 rows stemmed from the 200 bins \times 5 objects \times 4 trials. In the Manipulation Experiment, we computed a data matrix, $X_{2,0}$, similarly – $X_{2,0} \in \mathbb{R}^{5000 \times 23}$, where the columns represent the 23 joint angles measured, and the rows represent the 1000 bins \times 5 steps. Prior to SVD, the data were centered; that is, the mean of each feature (column) was removed [West et al., 2023]. Singular value decomposition computes the left and right singular vectors of a matrix. So, $SVD(X_{i,j}) = U_{i,j} \Sigma_{i,j} V_{i,j}^T$ produces linear combinations of hand joint angles (i.e. kinematic synergies) in $V_{i,0} \in \mathbb{R}^{23 \times 23}$, their temporal evolutions in $U_{1,0} \in \mathbb{R}^{4000 \times 4000}$ or $U_{2,0} \in \mathbb{R}^{5000 \times 5000}$, and an associated variance measure, the singular values, on the principal diagonal of $\Sigma_{1,0} \in \mathbb{R}^{4000 \times 23}$ or $\Sigma_{2,0} \in \mathbb{R}^{5000 \times 23}$.

Additionally, to provide a direct comparison of reach-and-grasp of a particular tool

to manipulation of that specific tool, we computed data matrices for reach-and-grasp of the zip tie, scissors, and screwdriver ($X_{1,1}$, $X_{1,2}$, and $X_{1,3} \in \mathbb{R}^{800 \times 23}$) and steps 1, 2, and 3 of the manipulation experiment ($X_{2,1}$, $X_{2,2}$, and $X_{2,3}$). Since there was no need to equally weight the steps in the manipulation data, here we smoothed the joint position data by down-sampling to 60Hz, rather than 1000 bins. Then, synergies were extracted using the SVD method described above.

To obtain a measure of the variance-accounted-for (VAF) in the data by a given kinematic hand synergy, we used the singular values along the diagonal of matrix $\Sigma_{i,j}$. Specifically, we report VAF as a percentage obtained by dividing the square of a singular value by the total sum of the squares of the singular values in $\Sigma_{i,j}$.

To test if manipulation and reach-and-grasp can be described by the same number of synergies, we must quantify the number of significant synergies. Thus, we followed the method of [Lambert-Shirzad and Van Der Loos, 2017]. Specifically, we report the number of synergies required to achieve at least 90% VAF and where inclusion of another subsequent synergy did not add an additional 5% VAF. To determine if there was an effect of tool or experiment on the number of significant synergies, these values were submitted to a 2 (experiment) \times 4 (tool) repeated-measures ANOVA.

Comparison of Kinematic Hand Synergies

While the above analysis compares the size of the synergy space, it is also insightful to compare the form of the synergy space. Thus, to compare the calculated kinematic synergies, $V_{i,j}$, across experiments, we computed the product of each individual subject’s reach-and-grasp synergies, $V_{1,j}$, and their manipulation synergies, $V_{2,j}$. This resulted in a matrix consisting of the cosine similarities, $C(i, j)$, between synergies where i denotes the i^{th} reach-and-grasp synergy and j denotes the j^{th} manipulation synergy. Here, we report the magnitude of these cosine similarities, ranging from 0 to 1.

$$C = \cos(\theta) = |V_1^T V_2| \quad (2.1)$$

If synergies were the same irrespective of experiment (i.e. the data span the same hyperspace), we would expect a matrix of cosine similarities with ones on the diagonal and zeros otherwise. Contrarily, if the reach-and-grasp synergies differed from the manipulation synergies (i.e. $V_{1,j} \neq V_{2,j}$), the resulting matrix, $C \in \mathbb{R}^{m \times m}$, would be asymmetric. In our report of the cosine similarities, in line with a plethora of other studies [Cheung et al., 2009, Cheung and Seki, 2021, D’avella and Tresch, 2001, Tresch et al., 2006] we highlight the cosine similarities that were greater than 0.9. Moreover, the values along the diagonal of this matrix are referred to here as the synergy similitude – a measure of the likeness between a subject’s reach-and-grasp synergies, $V_{1,j}$, and their manipulation synergies, $V_{2,j}$.

Ideally, we would aim to test the hypothesis that two synergy vectors are the same, $V_{1,j} = V_{2,j}$, which would result in a cosine similarity of 1, $C = 1$. However, cosine similarity is upper bounded by unity. Consequently, its distribution violates the implicit assumption of standard statistical test (i.e. approximate normality).

Instead, to determine if subjects’ reach-and-grasp synergies differed from their manipulation synergies we tested the hypothesis that the two synergy vectors were orthogonal, $C_{i,j} = 0$. To do so, we generated a distribution of the manipulation synergy by randomizing the order of the features in the synergy vector and computing the cosine similarity between the reach-and-grasp synergy and this random manipulation synergy. This was then repeated 1000 times to create a distribution. Then, we calculated where the observed cosine similarity fell within this distribution. It is important to note that here we computed $C = \cos(\theta) = V_1^T \cdot V_2$, ranging from -1 to 1; thus, the null distribution was centered about 0 and had a standard deviation less than unity. Finally, the reported p-value indicates the probability of obtaining a cosine similarity as extreme as the observed cosine similarity if there were no relationship (i.e. the reach-and-grasp and manipulation synergies were orthogonal, $C_{i,j} = 0$).

To determine if there was an effect of synergy-order on synergy similitude, we computed a linear regression between the two across all subjects. A significant effect was indicated if the slope was different from zero, i.e., the 95% confidence interval of the slopes did not include zero.

Recognizing that the default synergy order, based on variance-accounted-for, is subject to random variation due to experimental ‘error’ and ‘noise’, we reordered the manipulation synergies to continue the above analysis on similitude. Specifically, we recorded the columns of the manipulation synergy matrix, $V_{2,j}$, such that upon computing the cosine similarity matrix, C , the sum of the similitude (i.e., the sum of the diagonal values of C) was maximized. Subsequently, we followed the above procedures to determine if subjects’ reach-and-grasp synergies differed from their manipulation synergies and we tested if there was an effect of synergy-order on synergy similitude. Finally, we report the new manipulation synergies’ rank based on the new synergy matrix’s column order that maximized similitude.

All analyses were conducted in MATLAB 2022b.

2.3 Results

2.3.1 Emergence of Kinematic Hand Synergies

Variance-Accounted-For

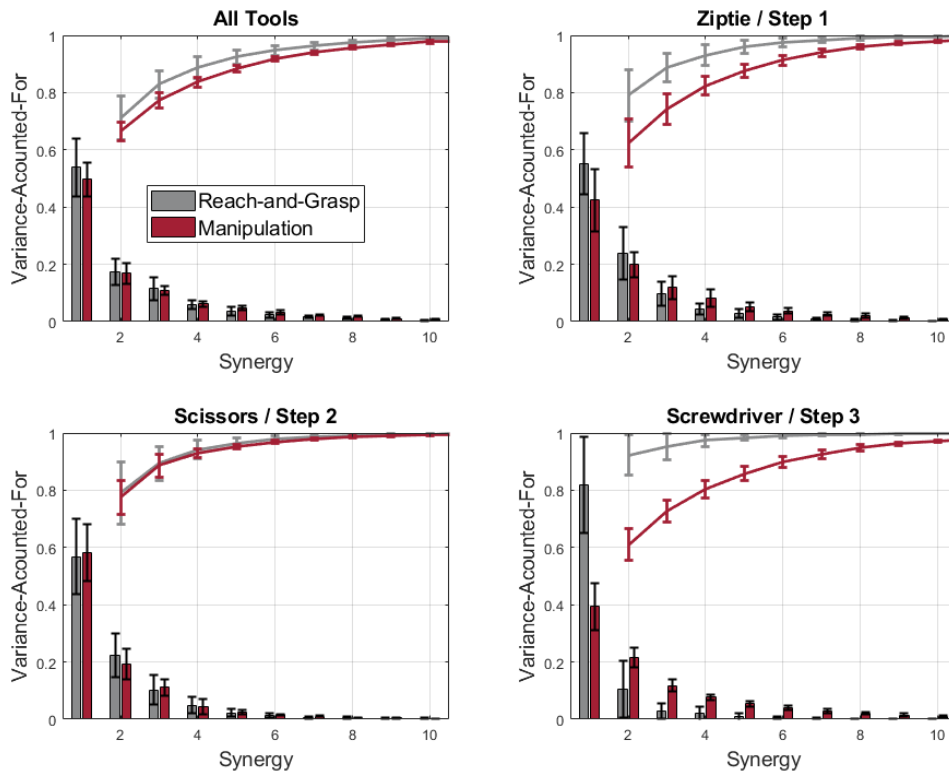


Figure 2-2: Variance-accounted-for (VAF) by each synergy averaged across subject for (Top Left) all tools, (Top Right) zip tie, (Bottom Left) scissors, and (Bottom Right) the screwdriver. Lines show the cumulative sum of the VAF. Error bars are ± 1 standard deviation.

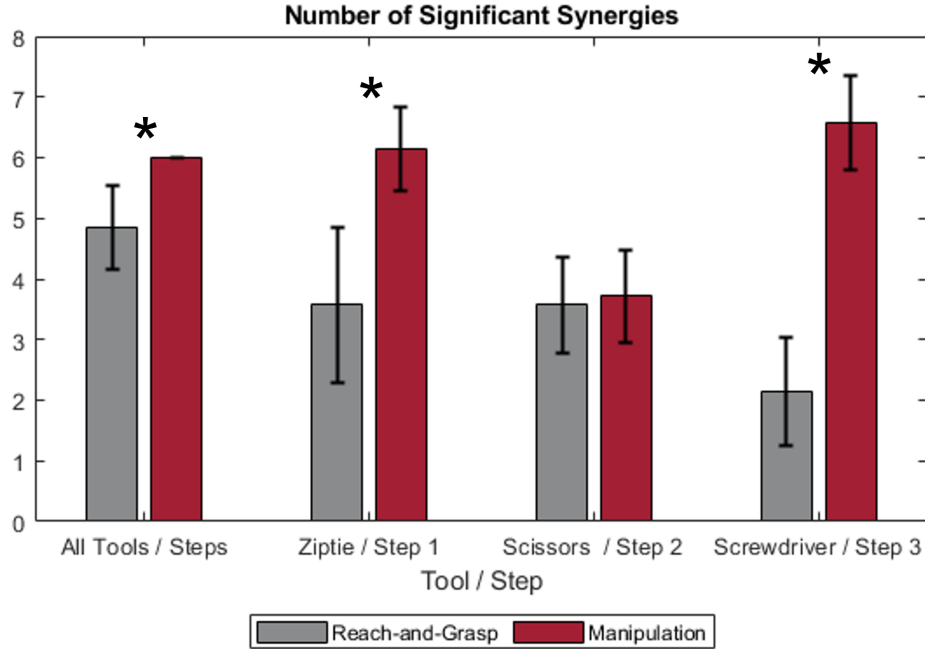


Figure 2-3: Number of significant synergies in reach-and-grasp or manipulation of each tool averaged across subjects. Error bars are ± 1 standard deviation. An asterisk (*) indicates a significant difference between the number of reach-and-grasp and manipulation synergies.

In general, manipulation required more synergies than reach-and-grasp. Fig. 2-2 shows the average variance accounted for by each of the synergies and their cumulative sum. Moreover, Fig. 2-3 shows the number of significant synergies for each tool and experiment. These values range from 1 synergy, in the case of one subject's reach-and-grasp of the screwdriver, to 8 synergies, in the case of one subject's manipulation of the screwdriver. Nonetheless, this analysis demonstrates that fewer synergies were needed than the full 23 DOF joint space of the hand.

A two-way (2 experiment \times 4 tool) ANOVA on number of significant synergies showed that there was an effect of experiment ($p = 1.07e-10$) and tool ($p = 3.96e-06$) on the number of synergies required to describe the data. However, there was a significant interaction ($F_{3,48} = 18.86$, $p = 3.62e-08$).

Post-hoc t-tests showed that the number of significant synergies computed across all tools ($M = 5.43$, $SD = 0.76$) was statistically greater than that of the scissor ($M =$

3.64, $SD = 0.75$)($p = 2.37e - 06$) and the screwdriver ($M = 4.34$, $SD = 2.44$)($p = 0.0051$). Additionally, the significant number of synergies computed across the zip tie ($M = 4.86$, $SD = 1.66$) was statistically greater than that of the scissor ($p = 0.0013$). Moreover, the significant number of synergies in the manipulation experiment ($M = 5.61$, $SD = 1.29$) was statistically greater than that of the reach-and-grasp experiment ($M = 3.54$, $SD = 1.32$)($p = 8.78e - 13$). The number of significant synergies was statistically greater in the manipulation experiment than in the reach-and-grasp experiment when computed across all tools ($p = 8.93e - 04$), the zip tie ($p = 5.14e - 04$) and the screwdriver ($p = 4.44e - 07$) (Bonferroni adjusted $\alpha_{Bonferroni} = 0.05/4 = 0.0125$). However, no effect was observed when comparing scissor reach-and-grasp to scissor manipulation ($p = 0.74$). To summarize, the number of necessary synergistic degrees-of-freedom was higher during manipulation than reach-and-grasp, except in the case of scissors.

Kinematic Hand Synergies

As mentioned above, we used singular value decomposition to compute the synergies in our data set. The first reach-and-grasp and manipulation synergies computed for each subject are shown in Fig. 2-4. Across all datasets, the first synergy was dominated by flexion of the MCP and PIP joints. These movements are akin to the power grasp that has been reported throughout the previous literature [Napier, 1956, Santello et al., 1998, Santello et al., 2016]. This motion was similar across subjects. The second synergy was dominated by thumb rotation (Fig. 2-5). Here, there appeared to be greater variation across subjects. The third (Fig. 2-6) and higher order synergies had such great variation among subjects that description of a common physiological movement risks being specious.

Synergy 1

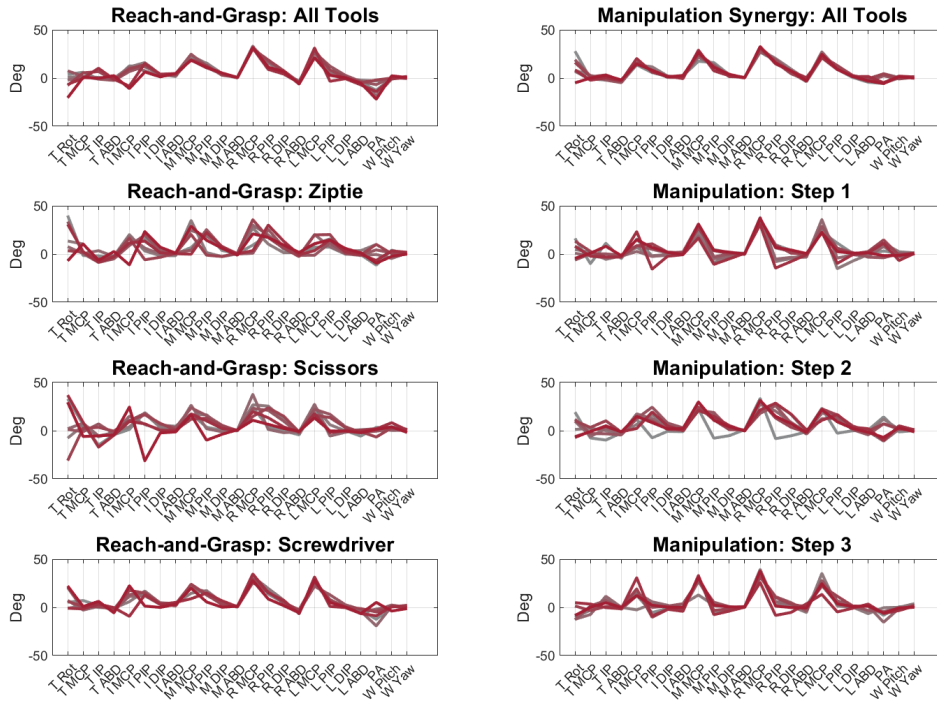


Figure 2-4: The first synergy. Each line denotes a different subject. Each row denotes synergies computed across a different data set (i.e. j in $V_{i,j}$) while each column denotes synergies computed across either reach-and-grasp, $V_{1,j}$, or manipulation, $V_{2,j}$.

Synergy 2

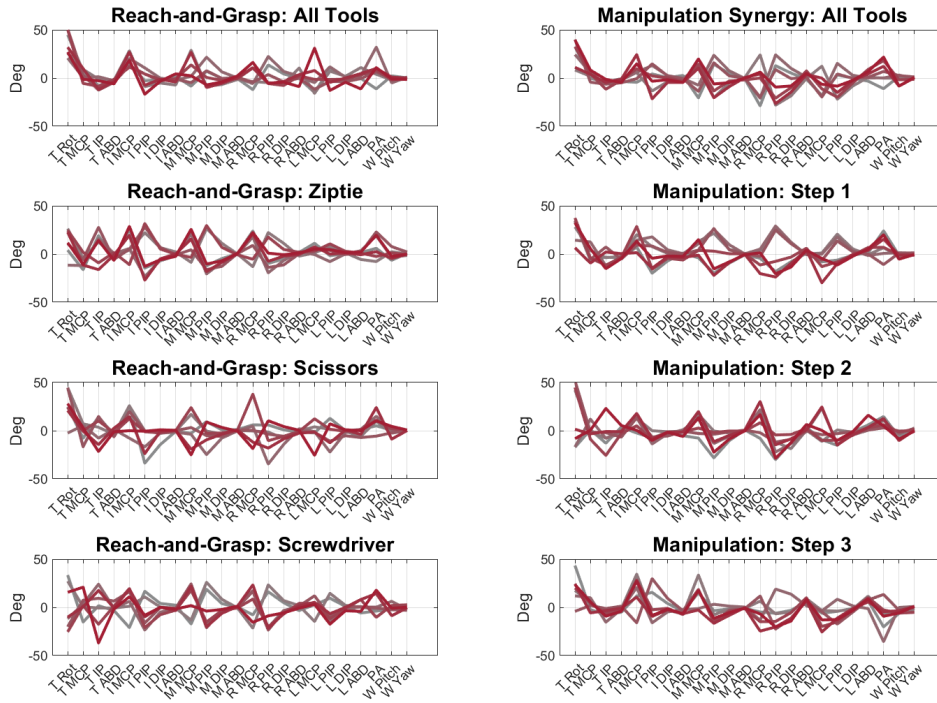


Figure 2-5: The second synergy. Each line denotes a different subject. Each row denotes synergies computed across a different data set (i.e. j in $V_{i,j}$) while each column denotes synergies computed across either reach-and-grasp, $V_{1,j}$, or manipulation, $V_{2,j}$.

Synergy 3

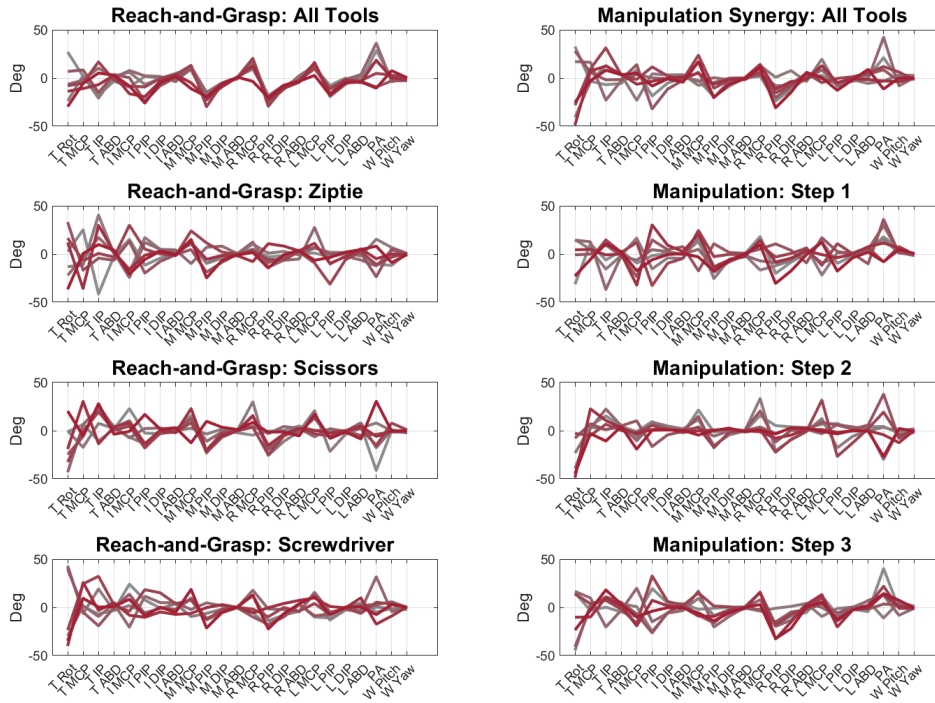


Figure 2-6: The third synergy. Each line denotes a different subject. Each row denotes synergies computed across a different data set (i.e. j in $V_{i,j}$) while each column denotes synergies computed across either reach-and-grasp, $V_{1,j}$, or manipulation, $V_{2,j}$.

2.4 Comparison of Kinematic Hand Synergies

While Fig. 2-2 and 2-3 compare the size of the synergy subspaces they do not compare their form; thus, to compare the form of the synergies across the two experiments, we computed a cosine similarity matrix; its magnitude for each subject is shown in Fig. 2-7. Each row denotes a comparison across a different subset of tools and each column denotes a different subject. Additionally, the right column in the figure shows the average of all the subjects' similarity matrices. Moreover, we highlight cosine similarities that are greater than 0.9 with color. Specifically, in the all-tools case, we observe a cosine similarity of greater than 0.9 in 3 of the 7 subject's first synergy comparison; in the scissors case we observe a cosine similarity greater than

0.9 in 2 of the 7 subjects first synergy comparison and 1 of the 7 subject's third synergy comparison. In the Zip tie and screwdriver similarity synergy comparisons we do not observe a cosine similarity greater than 0.9. Moreover, despite the evident differences between subjects, the average across subjects reflects common features seen in all subjects. Notably, the cosine similarity between the first synergy of each experiment is large while that of the subsequent synergies decreases.

Cosine Similarity ($|C|$)

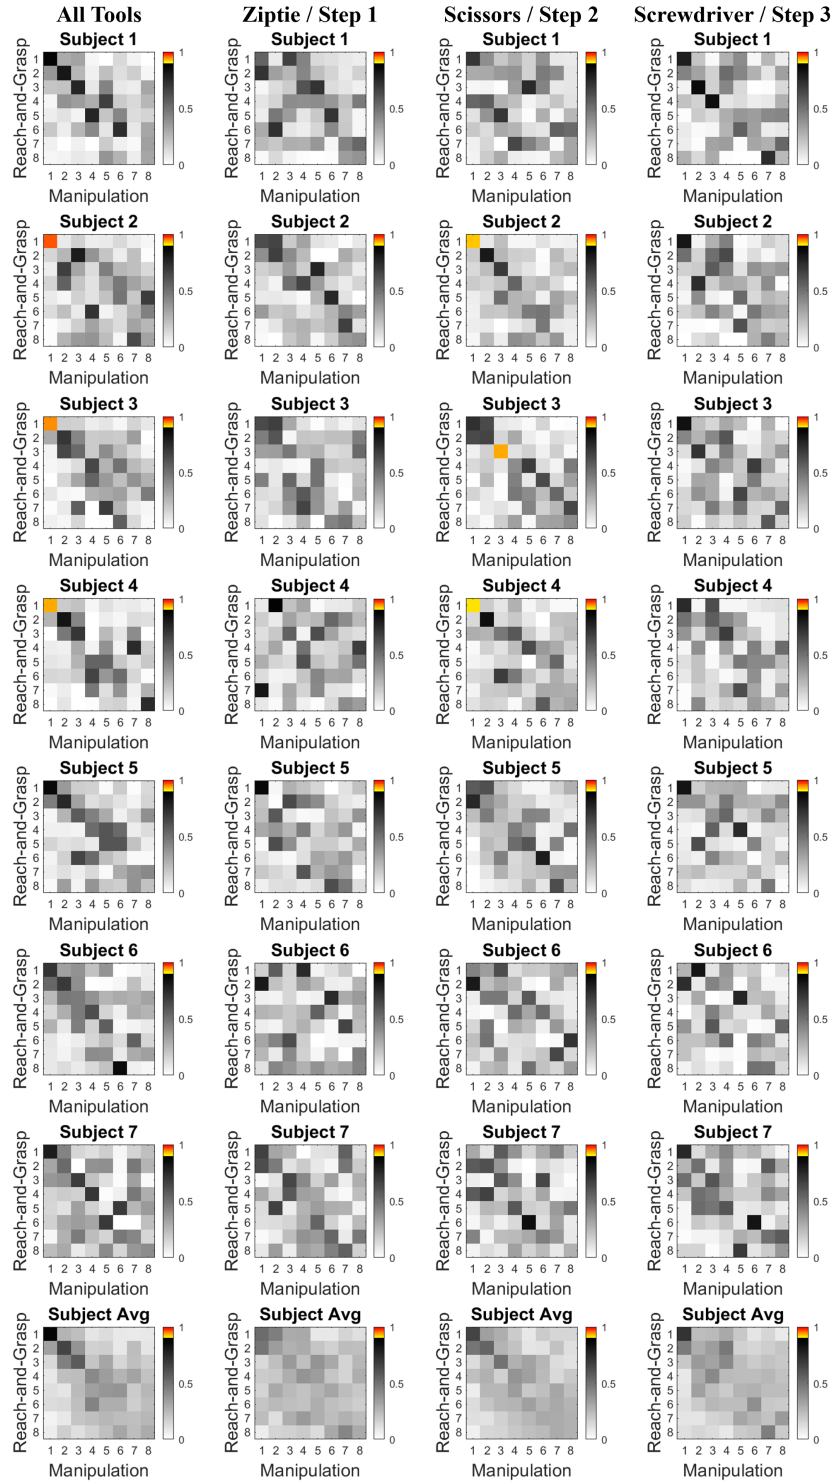


Figure 2-7: Cosine similarities between the first 8 synergies for the reach-and-grasp and manipulation tasks for all subjects. Cosine similarities greater than 0.9 are highlighted in color. Each row denotes a different subject; the last row shows the average across all subjects. Each column denotes comparison across a different subset of tools.

Number of Statistically Similar Synergies								
	Syn 1	Syn 2	Syn 3	Syn 4	Syn 5	Syn 6	Syn 7	Syn 8
All Tools	7	6	6	4	3	2	0	0
Zip tie / Step 1	5	3	3	3	1	1	0	0
Scissors / Step 2	5	4	5	2	0	0	0	0
Screwdriver / Step 3	6	1	3	0	2	3	0	0

Table 2.1: The number of statistically similar synergies upon comparing a subjects’ reach-and-grasp synergy (syn) to their manipulation synergy (syn) of the same order. The total number of subjects was 7. For brevity, only the first 8 synergies are shown. No subsequent synergies were statistically similar.

Additionally, to determine if subjects’ reach-and-grasp synergies differed from their manipulation synergies we tested the hypothesis that the two synergy vectors were different, $C_{i,j} = 0$. Two synergy pairs that reject this hypothesis suggest that the reach-and-grasp synergy and manipulation synergy span a similar subspace. Table 2.1 reports the number of statistically similar synergy pairs of the same order from the 7 subjects. That is, Table 2.1 tallies the number of statistically similar synergy pairs on the diagonal of the cosine similarity matrix presented in Fig. 2-7. In general, as synergy order increased, the number of statistically similar synergies decreased.

To determine if there was an effect of synergy-order on synergy similitude, we computed a linear regression between the synergy similarity and synergy order across all subjects. Note, only the synergies that were deemed significant (Fig. 2-3) were included. This regression is shown in Fig. 2-8. In all cases, the regression slopes were negative and significantly different from zero (All Tools: $p = 2.1e - 07$; Zip tie / Step 1: $p = 0.0039$; Scissors / Step 2: $p = 0.0055$; Screwdriver / Step 3: $p = 0.0022$). Note that while the coefficient of determination of these linear regressions were in the medium to low range (All Tools: $R^2 = 0.49$; Zip tie / Step 1: $R^2 = 0.19$; Scissors / Step 2: $R^2 = 0.27$; Screwdriver / Step 3: $R^2 = 0.19$), all of the regression models were significant when compared to a constant model; hence, the significant slope. Thus, we conclude that synergy similarity decreased with synergy-order.

Synergy Similitude

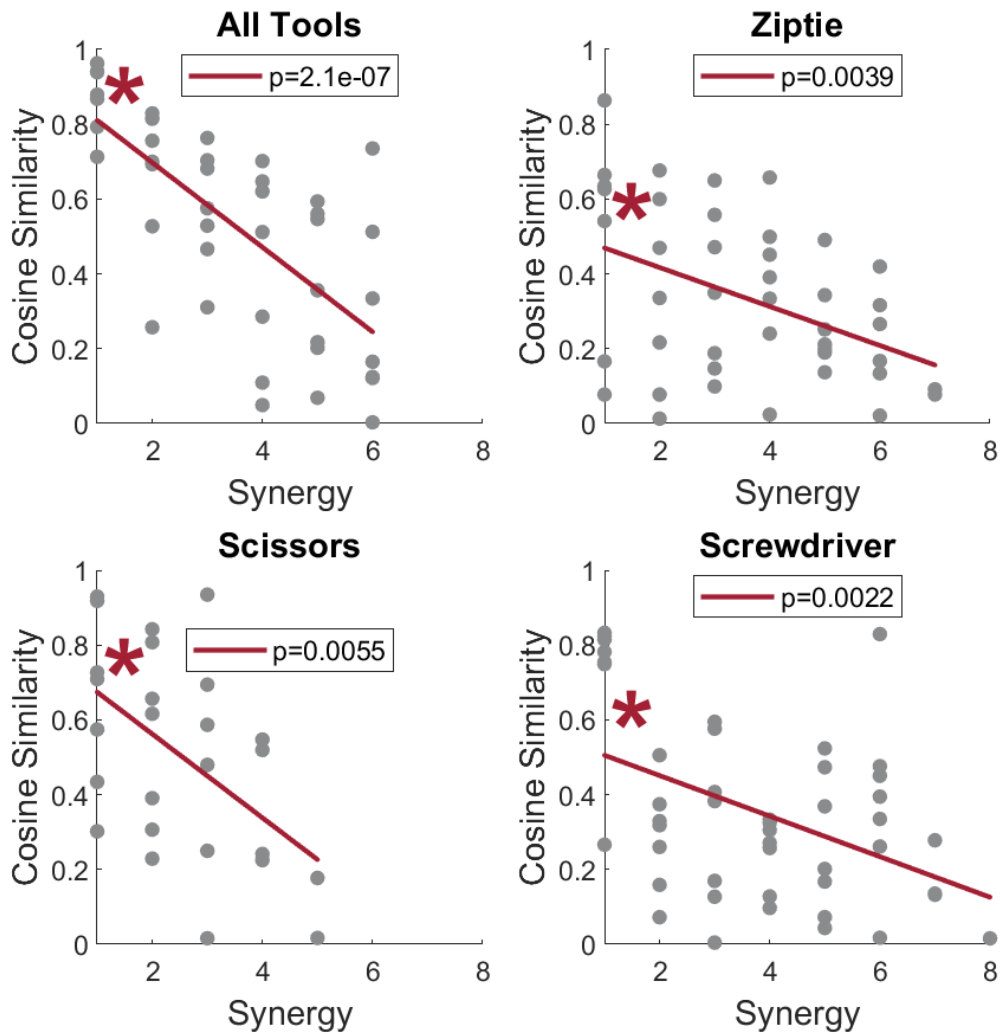


Figure 2-8: Synergy similitude across participants computed across each tool. The asterisk (*) indicates a significant slope.

The default synergy order produced by SVD was based on the amount of variance a given synergy accounted for in the data; as synergy order increased the variance-accounted-for (VAF) decreased. Our proposed similitude measure focused on the cosine similarity of synergies with the same order and was thus susceptible to synergy order. This highlights that the default synergy order does not exactly allow us to compare the most similar synergies across reach-and-grasp and manipulation. To combat this, we reordered the manipulation synergies to maximize synergy similitude

when compared to the reach-and-grasp synergies. These comparisons are presented in Fig. 2-9. There, we highlight cosine similarities that are greater than 0.9 with color. As in Fig. 2-7, we still observe in the all-tools case, a cosine similarity of greater than 0.9 in 3 of the 7 subject's first synergy comparison and in the scissors case a cosine similarity greater than 0.9 in 2 of the 7 subjects first synergy comparison and 1 of the 7 subject's third synergy comparison. In the zip tie and screwdriver similarity synergy comparisons we do not observe a cosine similarity greater than 0.9. Notably, the number of statistically similar synergies is quantified in Table 2.2. There, we observed a higher coincidence of statistically similar synergies on the diagonal. As expected, we see an increase of statistically similar synergies in Table 2.2 when compared to Table 2.1.

Reordered Cosine Similarity (|C|)

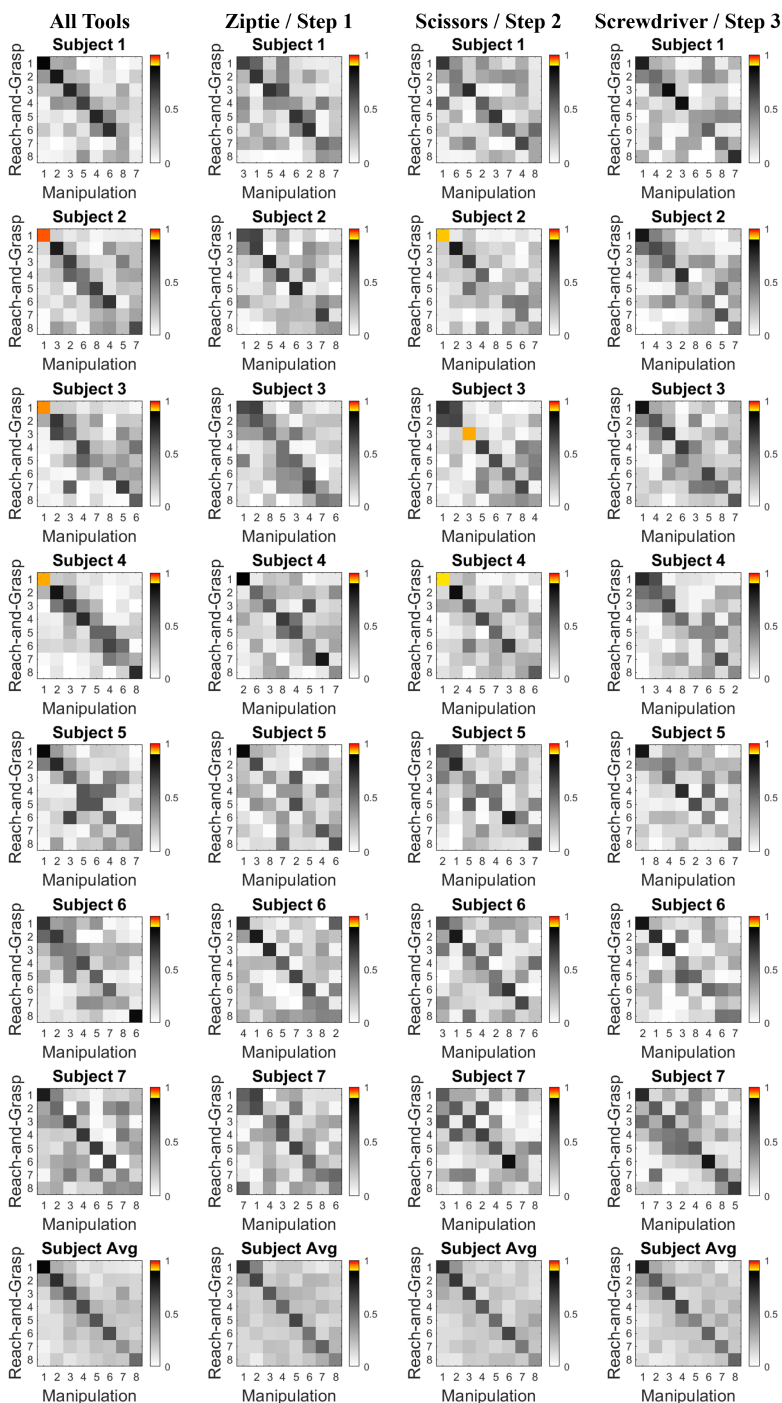


Figure 2-9: All subjects' first 8 synergy cosine similarities between the reach-and-grasp and manipulation synergies upon reordering the columns of the manipulation synergy matrix to maximize similitude. Note, the label of the x-axis refers to the original synergy rank based on the variance-accounted-for in the manipulation data. Cosine similarities greater than 0.9 are highlighted in color. Each row denotes a different subject; the last row shows the average across all subjects. Each column denotes comparison across a different subset of tools.

Number of Statistically Similar Synergies								
	Syn 1	Syn 2	Syn 3	Syn 4	Syn 5	Syn 6	Syn 7	Syn 8
All Tools	7	7	7	6	7	5	0	0
Zip tie / Step 1	7	7	7	4	5	5	2	0
Scissors / Step 2	6	6	6	4	2	0	0	0
Screwdriver / Step 3	7	7	7	4	6	6	1	1

Table 2.2: The number of statistically similar synergies upon reordering the columns of the manipulation synergy (syn) matrix to maximize similitude. The total number of subjects was 7. For brevity, only the first 8 synergies are shown. No subsequent synergies were statistically similar.

Fig. 2-10 compares the synergy similitude after reordering the manipulation synergies. To determine if there was an effect of synergy-order on synergy similitude, we computed a linear regression between the synergy similarity and the significant synergies across all subjects. This regression is shown in Fig. 2-10. Unlike Fig. 2-8, the regression slopes were negative and significantly different from zero only for the All-Tools and Screwdriver / Step 3 datasets but not the Zip tie / Step 1 and Scissors / Step 2 datasets (All Tools: $p = 1.0e - 05$; Zip tie/ Step 1: $p = 0.05$; Scissors / Step 2: $p = 0.42$; Screwdriver / Step 3: $p = 1.7e - 03$). Note, here the coefficient of determination of these linear regressions were in the medium range for the data whose model did present a significant slope and in the very low range, $R^2 < 0.01$, for the data whose model did not present a significant slope (All Tools: $R^2 = 0.49$; Zip tie / Step 1: $R^2 = 0.09$; Scissors / Step 2: $R^2 = 0.026$; Screwdriver / Step 3: $R^2 = 0.28$). This may indicate that there was similarity between the reach-and-grasp and manipulation synergies, especially in the handling of the Zip tie and scissors. However, in the all-tools and screwdriver data these results suggest that synergies clearly present in manipulation may not be as prevalent in reach-and-grasp.

Reordered Synergy Similitude

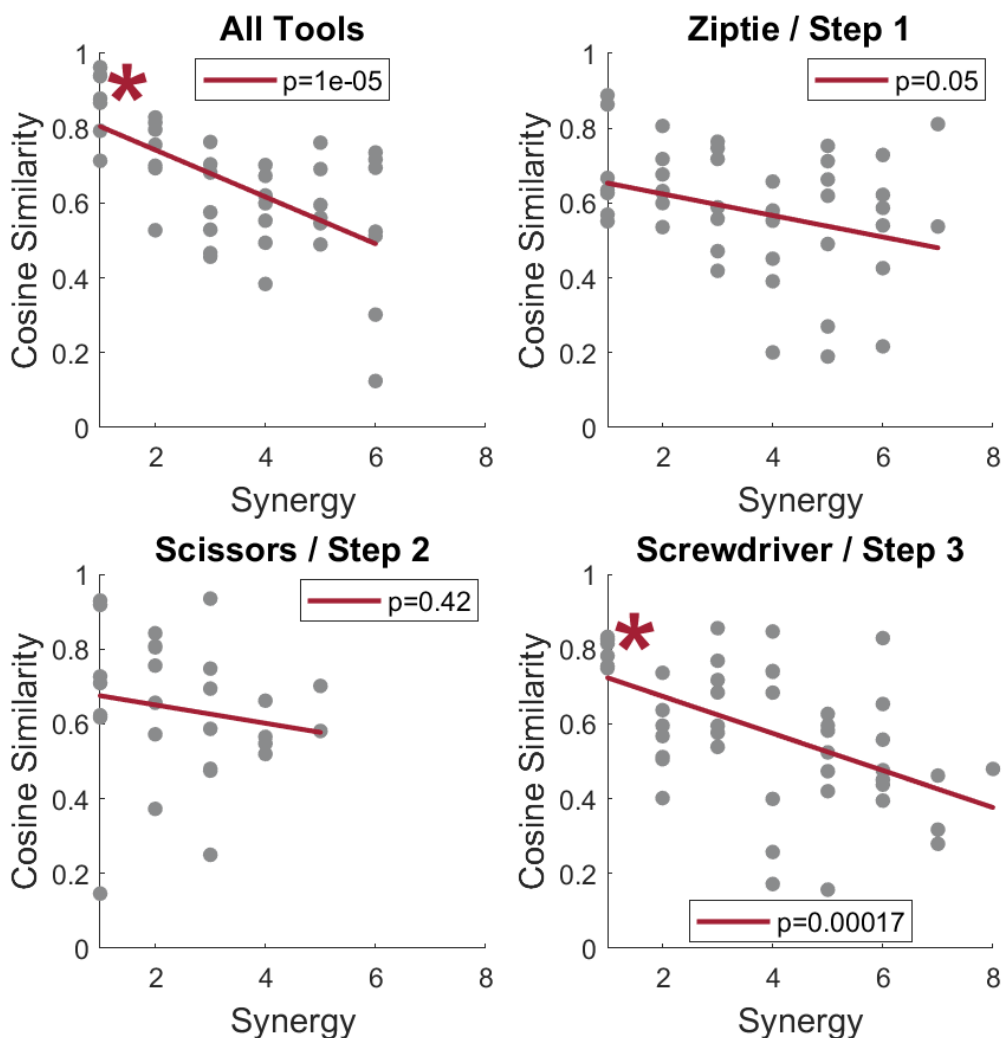


Figure 2-10: Synergy similitude across participants computed across each tool, upon reordering the columns of the manipulation synergy matrix to maximize similitude. The asterisk (*) indicates a significant slope.

Fig. 2-11. provides insight into how reordering the manipulation synergies led to the increased incidence of statistically similar synergies (Table 2.2). Specifically, Fig. 2-11 shows the manipulation synergy ranks averaged across subjects after the manipulation synergies were reordered to maximize similitude; a dashed black line demonstrates no change in synergy order. Values above this line demonstrate a decrease in synergy-order and values below this line demonstrate an increase in

synergy-order. That is, values above the dashed black line suggests that synergies that were less prevalent in reach-and-grasp became more prevalent in manipulation and values below the dashed black line suggests that synergies that were more prevalent in reach-and-grasp became less prevalent in manipulation. In Fig. 2-11 it is readily seen that higher-order synergies often became lower order synergies while lower-order synergies often became higher-order synergies, especially in the screwdriver data. This suggests that synergies that were deemed insignificant in reach-and-grasp were important in manipulation.

Reordered Manipulation Synergy Ranks

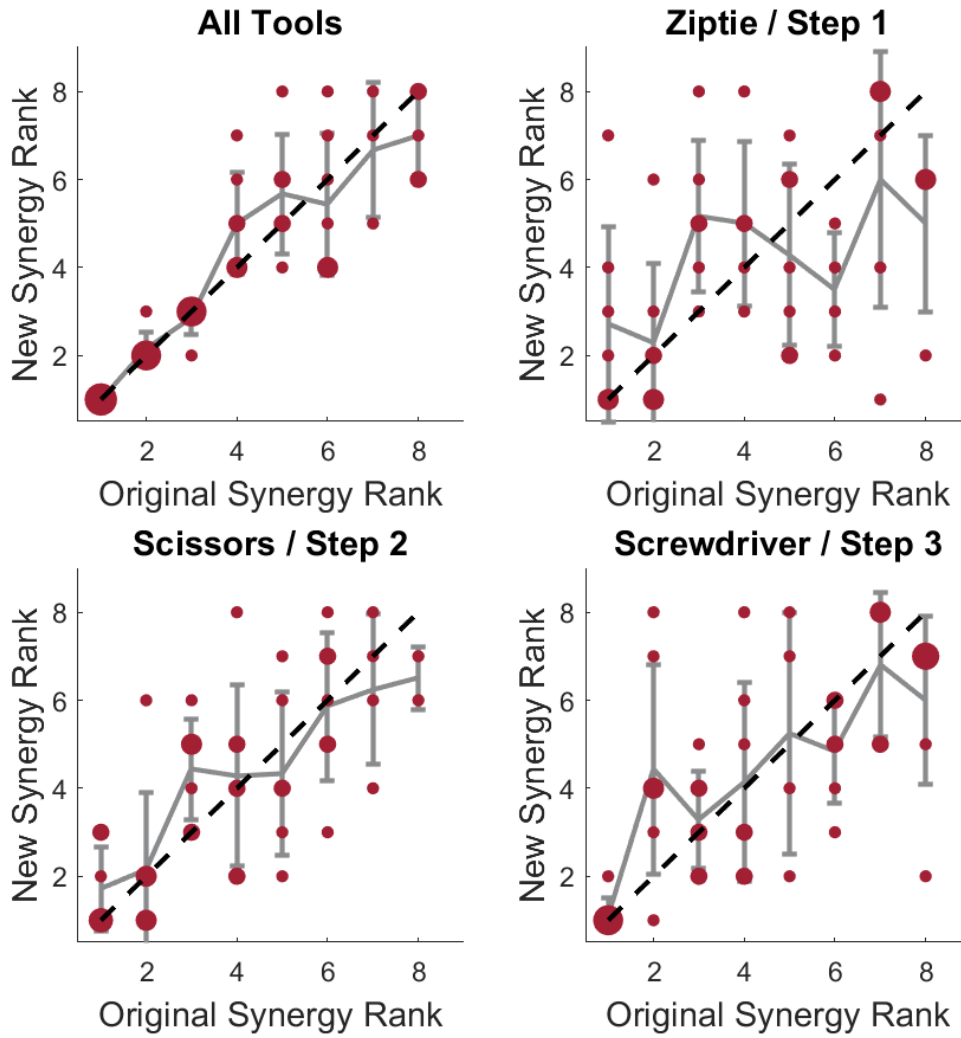


Figure 2-11: Manipulation synergy ranks after the manipulation synergies were reordered to maximize similitude upon comparing to the reach-and-grasp synergies. The gray line denotes the average and error bars denote the standard deviation. The red dots are the data; larger dots represent a repetition of the data point. The dashed black line is the direct relation between the old and new rank and represents no change in order. Values above the dashed black line suggests that synergies that were less prevalent in reach-and-grasp became more prevalent in manipulation and values below the dash black line suggests that synergies that were more prevalent in reach-and-grasp became less prevalent in manipulation.

2.5 Discussion

In this study, we tested two hypotheses: (1) the number of synergies underlying manipulation is the same as those underlying reach-and-grasp, and (2) the identity of synergies underlying manipulation is different from those underlying reach-and-grasp. We found that more synergies were required to competently describe the manipulation data than the reach-and-grasp data. Moreover, upon comparing the kinematic hand synergies across reach-and-grasp and manipulation, we found that the difference was predominantly in the higher-order synergies. However, after reordering the manipulation synergies to maximize similitude, we found that synergies clearly present in manipulation may not be as prevalent in reach-and-grasp.

2.5.1 Emergence of Kinematic Hand Synergies

Our analysis did confirm a reduction of the hand’s operational degrees of freedom. The first 3 synergies accounted for greater than 75% of the variance in both experimental tasks (Fig. 2-2). This value is consistent with what has been reported in prior literature [Mason et al., 2001, Santello et al., 1998, Santello et al., 2016, Todorov and Ghahramani, 2004, Weiss and Flanders, 2004]. Thus, we conclude that this simplified kinematic analysis can continue to be useful to understand aspects of motor control of the human hand performing a functional task. Consistent with much of the synergy work first introduced by Santello et al. [Santello et al., 1998], the first synergy, in both of the experimental tasks presented here, was dominated by flexion of the MCP and PIP joints (Fig. 2-4). This result has been observed across several experiments that studied hand posture in grasping, both with and without contact [Mason et al., 2001, Santello et al., 2002]; and even including signing in American sign language [Weiss and Flanders, 2004]. This first synergy is generally equated to Napier’s power grasp. A power grasp is when flexed fingers and the palm clamp down on an object while the thumb applies counter pressure [Napier, 1956]. Typical examples of a power grasp include gripping a hammer or holding a bottle. The presence of this synergy in these experiments is consistent with prior studies.

While a reduction of the hand’s operational degrees of freedom was supported by our analysis in both experiments, interestingly, more synergies were required to account for the same variance in the Manipulation Experiment than in the Reach-and-Grasp Experiment (Fig. 2-3). This difference was statistically significant when the synergies were extracted from all tools, the Zip tie alone, and the screwdriver alone. However, this was not the case when synergies were extracted from the scissors. We speculate that this is because, when using a pair of scissors, the finger and thumb motion is physically constrained by the object. With the understanding that increasing variability in motion increases the number of significant synergies, it is plausible that the reason we saw no statistical difference in the number of significant synergies between reach-and-grasp of the scissors and manipulation of the scissors is because the object constrained the hand’s motion. We want to highlight that the intuitiveness of this result reassured us that despite the challenge of noisy data and complicated processing, our analysis didn’t detect differences when none were to be expected.

Unlike the scissors, in the other tools studied here, more synergies were required to account for the same variance in the Manipulation Experiment than in the Reach-and-Grasp Experiment (Fig. 2-3). For instance, to use a screwdriver, one must repeatedly grip and turn the object; this motion might also require some form of finger gaiting. Similarly, to manipulate the zip tie will not require a power-grasp (synergy 1) but will require more fine-grained finger manipulation. Nonetheless, we saw the greatest difference in the number of significant synergies when comparing screwdriver reach-and-grasp to manipulation as screwing requires many re-grasps and finger gaiting. We also saw a similar difference when comparing zip tie reach-and-grasp to manipulation. Together, these results suggest that a greater variety of hand postures was used when manipulating a tool or object than when reaching for and grasping it, unless that tool greatly constrained hand motion. This may be because manipulation requires more variable hand kinematics than reach-and-grasp, which also suggests that high-order synergies may be more important in manipulation than in reach-and-grasp. Intuitively, this would be consistent with the need for high dexterity in tasks that

require object manipulation.

2.5.2 Reach-and-Grasp vs. Manipulation Synergies

This study compared synergies across reach-and-grasp and manipulation (Fig. 2-7). Particular focus was centered on the diagonal values of the similarity matrices – the synergy similitudes. These values are shown for the significant synergies in Fig. 2-8. There it is seen that synergy similarity decreased with synergy order. Since these higher-order synergies differed between the reach-and-grasp and the manipulation experiments, current knowledge of higher-order kinematic hand synergies cannot be applied reliably when the hand is performing a functional task. Considering that higher-order synergies become significant during manipulation (Fig. 2-3), it is important that we investigate these differences.

It is important to note that higher-order synergies are often dismissed because they do not explain a significant portion of the variance in the data, and their minimal contribution is typically considered indistinguishable from noise (i.e., they are near the “noise floor”). In our data, the first synergy, irrespective of task, consistently represented a power grasp. However, the higher-order synergies, which usually contribute substantially to the overall action, exhibited significant differences between reach-and-grasp and manipulation tasks. Typically dismissed as experimental error or noise, these higher-order synergies have obscured differences between reach-and-grasp and manipulation. In this study, we have elucidated these distinctions by refraining from dismissing higher-order synergies. Note that investigating the diagonal values of the similarity matrices in Fig. 2-7 assumed that synergies of the same order should be the same. However, this need not be true. In fact, in Fig. 2-7 several off-diagonal synergy comparisons present statistically similar synergies. To truly test similarity between reach-and-grasp and manipulation synergies, we developed a comparison between the synergies with the highest cosine similarity. Specifically, in Fig. 2-9 we reordered the manipulation synergies to maximize the diagonal cosine similarities. There, we see an increased incidence of statistically similar synergies on the diagonal (Table 2.2). Fig. 2-11 quantifies this. We occasionally observed that

higher order synergies increased in rank (became less important) while lower-order synergies decreased in rank (became more important), especially in screwdriver use. This suggests that synergies that were deemed insignificant in the reach-and-grasp case were important in manipulation. Understanding the hand kinematics during the grasp of a screwdriver and the use of a screwdriver may help explain this. The screwdriver is a cylindrical object that can be picked up using a power grasp, which is akin to synergy 1. However, to screw requires rotating the tool, which, in turn, requires more finger-gaiting motions. Thus, in this task it would make sense that higher-order reach-and-grasp synergies became more important during this functional tool use, while synergy 1 became less important. This is what is shown in Fig. 2-11. This result implies that we should further investigate the typically-ignored higher-order reach-and-grasp synergies as they may be important for manipulation even though they are closer to the “noise floor”.

2.5.3 Limitations

Because reach was inherently incorporated into the manipulation experiment, we could not fundamentally de-couple the two groups of synergies. It is true that reach was incorporated in the data used to extract both groups of synergies. However, in the first experiment the reach portion of each trial lasted 4.75s on average while the second experiment lasted 92.73s on average. Thus, when computing the synergies from the manipulation data they were not heavily influenced by the reach portion of the motion. Moreover, our results showed that reach-and-grasp and manipulation synergies differed. The fact that reach motion was incorporated in the data used to extract both groups of synergies would make the groups of synergies more similar, not more different.

Singular Value Decomposition was used to extract the kinematic hand synergies. This is only one of many dimensionality reduction technique and suffers flaws; however, it is the most commonly used in kinematic data [Santello et al., 2016]. As a data-driven approach, it is susceptible to data pre-processing effects. We aimed to mitigate this by following the guidelines presented in [West et al., 2023]. Namely,

we centered the data by removing the mean; this ensured that we did not obtain a dominant first synergy that only served to center the data and thereby reduced the number of significant synergies [West et al., 2023]. Moreover, we chose to bin the data in a manner that was least likely to change the variance within the data; doing so would have resulted in a synergy indicating related DOFs but without accurate quantification of their co-variation [West et al., 2023]. While it is possible that our results may have been influenced by the data pre-processing and processing methods, reasonable assumptions consistent with prior literature were made in the synergy extraction methods. This will be discussed in length in Appendix A.

Ideally, we aimed to test the hypothesis that two synergy vectors are the same, $V_{1,j} \neq V_{2,j}$. To compare the reach-and-grasp and manipulation synergies, we computed their cosine similarity. This metric ranges from -1 to 1. However, two synergies with a cosine similarity of -1 and 1 will span the same subspace. Therefore, we reported the magnitude of the cosine similarity ranging from 0 to 1. However, cosine similarity is upper bounded by unity, and consequently, its distribution violates the implicit assumption of standard statistical tests (i.e., approximate normality). Thus, we tested the null hypothesis that the two synergy vectors were orthogonal, $C_{i,j} = 0$. It is important to emphasize that falsification of this hypothesis strictly tells us that the two compared synergy vectors do not span completely orthogonal subspaces. Thus, the reported similar synergies in Tables 1 and 2 do not necessarily span the same subspace; rather, they do not span orthogonal subspaces. Due to the boundedness of the cosine similarity, the authors cannot test and comment on whether the compared synergy vectors are statistically equal. Thus, in line with other literature [Cheung et al., 2009, Cheung and Seki, 2021, D’avella and Tresch, 2001, Tresch et al., 2006] the authors have highlighted cosine similarity greater than 0.9.

Furthermore, it’s important to note that cosine similarity, particularly in high-dimensional spaces, comes with known limitations. For instance, it normalizes vectors based on their magnitudes, focusing primarily on direction rather than magnitude, which can be problematic in scenarios where magnitude is significant.

Additionally, sparse vectors, which are common in high-dimensional datasets, can lead to artificially high similarity scores because only non-zero dimensions are considered. Another challenge is that cosine similarity is undefined for zero vectors. Most notably, in very high-dimensional spaces, the distinction between different directions can become less meaningful, especially if the data contains a lot of noise. This is partly due to the curse of dimensionality, where distances in high-dimensional spaces do not always behave as intuitively expected, potentially reducing the effectiveness of cosine similarity in distinguishing between truly similar and dissimilar items [Houle et al., 2010]. Nonetheless, it was employed in this study because it remains one of the most widely used methods for multidimensional comparisons [Cheung et al., 2009, Cheung and Seki, 2021, D'avella and Tresch, 2001, Tresch et al., 2006, West et al., 2023, West, 2020, Lahitani et al., 2016, Luo et al., 2018]. Future research should aim to develop a metric that can compare high-dimensional subspaces without encountering these limitations.

In this study, muscle dynamics and muscle synergies were not studied. As muscles are the human's actuators, one could argue that we cannot make conclusions about the human motor controller. However, considerable insight has been gained from studying purely kinematic hand motion [Bicchi et al., 2011, Brown and Asada, 2007, Gabbicini et al., 2011, Gabbicini et al., 2013, Gioioso et al., 2013, Jarque-Bou et al., 2019, Mason et al., 2001, Santello et al., 1998, Santello et al., 2002, Santello et al., 2016, Todorov and Ghahramani, 2004, Weiss and Flanders, 2004]. Moreover, studies of adaptation to visual and/or mechanical perturbation have shown that kinematics dominate human motion planning [Flanagan and Rao, 1995, Shadmehr and Mussa-Ivaldi, 1994]. Studying muscle dynamics during manipulation is important and should be pursued further; however, one should not ignore the insights that can be gained from kinematics.

2.5.4 Implications

Synergies have been used as a basis for rehabilitation post Cerebral Vascular Accident [di Luzio et al., 2022, Singh et al., 2018]. Moreover, they have led to the

design of several hand rehabilitation devices [Iandolo et al., 2019, di Luzio et al., 2022, Maciejasz et al., 2014, Sierotowicz et al., 2022, Ueki et al., 2008]. The study reported here demonstrated that there are differences between the synergies used for reach-and-grasp and for tool-use and object manipulation. Thus, it may be important to incorporate these differences in functional rehabilitation tasks intended to help patients recover activities of daily living, which often require tool-use. For upper limb reaching, the advancement of robotic rehabilitation techniques has demonstrated greater effectiveness and better cost efficiency than standard patient care [Fasoli et al., 2004, Krebs et al., 2003, Riener et al., 2005, Rodgers et al., 2019]. Successful deployment of upper-limb robotic rehabilitation has been attributed to studies that have explored innate human motor function, similar to what was done here. Ultimately, if we want to restore hand function to those who have lost it, kinematic hand synergies can be useful; however, they must be studied in the context of activities of daily living, which often require functional tool-use.

The results shown here also have implications for prosthetic design. Many prostheses have used synergies to simplify design [Fani et al., 2016, Fu and Santello, 2018, Furuji et al., 2019, Laffranchi et al., 2020, Li et al., 2014, Matrone et al., 2012, Santello et al., 2016]. However, those designs have relied on the use of lower-order grasp synergies. As demonstrated by the work reported here, higher-order grasp synergies become important for tool-use. Thus, the previously used low-order grasp synergies may not encompass movements required for tool-use.

2.5.5 Are synergies an epiphenomenon?

The concept of synergies has a venerable history in the motor neuroscience and robotics literatures (for review see [Santello et al., 2016]). Confining attention to kinematic synergies, as described above they may be defined by a time-invariant map $\phi(\cdot)$ from a lower-dimensional control input $\delta \in \mathbb{R}^m$ to a higher-dimensional configuration space $\theta \in \mathbb{R}^n$, where $n > m$, such that $\theta(t) = \phi(\delta(t))$. Numerical methods such as singular value decomposition may be applied to experimental data to identify a linear approximation to the map $\phi(\cdot)$. Is this concept insightful?

For example, any single action necessarily follows a path $p(\cdot)$ in configuration space. A highly-skilled or well-learned action is likely to be repeatable such that every replication of the action follows the same path in configuration space. That path defines the coordinate of a one-dimensional (albeit likely nonlinear) subspace to which the task is confined. Displacement along that path $s \in \mathbb{R}^1$ defines a single coordinate. Its time-history $s(t)$ may serve as a control input to define the configuration-space trajectory $\theta(t) = p(s(t))$ used for the task. Therefore $p(\cdot)$ defines a synergy. By this reasoning, synergies are a necessary consequence of skill.

Moreover, distinct actions likely follow distinct paths $q(\cdot)$, $r(\cdot)$ etc. and each of these also defines a synergy. As a result, we may predict that synergies will vary with subject, task, and object, consistent with our findings and those of previous researchers [Furuya et al., 2011, Jarque-Bou et al., 2019, Santello et al., 1998, Santello et al., 2002, Weiss and Flanders, 2004]. Thus, we may conclude that the dimensionality of a single highly-learned action should be one. Imperfect execution or noise might lead to an apparent dimensionality greater than one. Those additional dimensions may be useful, e.g. to absorb inevitable motor variability, but any additional dimensions do not describe the task.

Moreover, the linearization implicit in the use of numerical methods such as singular value decomposition may also lead to an apparent dimensionality greater than one. This may motivate the application of nonlinear approaches in future work. Nonetheless, this reasoning does not challenge the biological reality of synergies. Learning a skill is loosely analogous to solving an optimization problem and is costly in time and physical and mental effort. Storing, and retrieving, learned patterns in the form of synergies may be a way to avoid the need to re-optimize on every execution. But there's no free lunch: we should expect that learning a task that requires a configuration-space path orthogonal to $p(\cdot), q(\cdot), r(\cdot)$ etc. will be much harder than learning a task that interpolates between these subspaces. That, too, has been reported [Golub et al., 2018, Sadtler et al., 2014].

A collection of skilled behaviors need not use all regions of configuration space but may be confined to a subspace spanned by the synergies $p(\cdot), q(\cdot), r(\cdot)$ etc. The key to

identifying this skilled subspace is to study tasks that require the true versatility of human manipulation—such as wire-harnessing.

2.6 Conclusion

This study sought evidence of kinematic hand synergies during wire harness installation – a task that involves manipulation of complex tools and objects. Human hand motion was measured during two experiments: (1) reaching for and grasping a tool or object and (2) manipulating those objects. In both experiments, a reduction of the operational degrees of freedom was observed (i.e. synergies). Moreover, we found that manipulation of a tool generally required more significant synergies than grasp of that same tool. Nonetheless, consistent with previous literature on kinematic hand synergies, the first synergy participated in power grasp, and the second synergy was dominated by thumb rotation. However, upon comparing reach-and-grasp with manipulation, we found that synergy similarity decreased with synergy-order. Considering that higher-order synergies become significant during manipulation, it is important that we investigate these differences; this study serves as a point of entry to understanding them. Investigation of the human hand during functional object manipulation may lead to better prosthetic hand design, and hand rehabilitation techniques. If we want to restore hand function to those who have lost it we must study hand manipulation beyond grasping.

Chapter 3

Dynamic Primitives Describe Limitations in Human Force Regulation during Motion

In Chapter 2, we discussed a simplifying control technique to plan motion during dexterous tasks. However, that chapter overlooked a key aspect of physical interaction: force modulation. We aim to address this aspect in the present chapter. Specifically, we inquire: How do humans control force? Are they capable of managing force independent of motion? Contrary to longstanding assumptions in robotics, this study reveals that humans do not modulate force and motion independently; instead, they are intertwined through mechanical impedance.

This chapter is an adapted version of a paper published in IEEE Robotics and Automation Letters titled, “Dynamic Primitives Limit Human Force Regulation during Motion” [West et al., 2022a]. This work was done in collaboration with James Hermus, Meghan Huber, Pauline Maurice, Neville Hogan, and Dagmar Sternad.

A video summary of this chapter can be found at <https://www.youtube.com/watch?v=RRfWe9KVSDY>

3.1 Introduction

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

During physical interaction, bidirectional forces between the actor and the environment critically affect the behavior of the coupled system. If humans regulate motion during free reaching, a simple extension of this idea to contact tasks may be to regulate both force and motion. In robotics, hybrid control allows for simultaneous and independent control of both motion and force in complementary subsets of the workspace [Mason, 1981, Raibert and Craig, 1981]. In human motor control, it is yet unresolved whether humans can control force independent of motion.

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

On the other hand, it has been shown that the central nervous system (CNS) contains a controller that modulates the coupling of force and motion [Kolesnikov et al., 2011, Piovesan et al., 2019]. Other studies demonstrated that humans modulate the relation between motion and force during upper limb reaching in unstable force

fields [Burdet et al., 2001, Milner and Franklin, 2005, Osu et al., 2003, Takahashi et al., 2001]. Additionally, our own previous research showed that exerted force depended on the velocity profile when grasping and following a robot manipulandum. Specifically, participants were asked to trace the motion of a robot manipulandum without exerting force as it moved on an elliptical path with varying velocity profiles [Maurice et al., 2018]. If force can be controlled independent of motion, the velocity profile should not matter; however, it did.

This study aimed to examine human control of physical interaction that could resolve these seemingly contradictory results. We conducted an experiment in which participants physically interacted with a motion-controlled robot to test whether humans could regulate force independent of motion. We refer to this independent control as “direct force control” (Fig. 3-1a). Explicitly, direct force control applies an actual force as a function of only a planned force. This function is an operator that may be dynamic and nonlinear. If participants can regulate force independent of motion, direct force control can be accepted as a plausible schema for human physical interaction. Conversely, if humans are unable to decouple force from motion, an alternative hypothesis is “indirect force control.” With indirect force control, a planned force $f_p(t)$ may still exist in the forward path, but an impedance term $Z\{\cdot\}$ is needed to relate the difference between input motion $x_0(t)$ and actual motion $x(t)$ to the output force $f(t)$ (Fig. 3-1b). The core feature of indirect force control is that force depends on motion.

The direct force control hypothesis leads to a testable prediction: Errors in contact force will be independent of motion. Thus in this experiment, participants were instructed to apply a specific constant force on a robot manipulandum in its direction of motion as it moved along an elliptical path. To give participants the best opportunity to complete the task, the robot moved with a velocity profile that matched human movement preferences, i.e., angular velocity scaled with curvature with a power of $2/3$ [Maurice et al., 2018] [Zago et al., 2018]. Despite visual feedback and some practice, errors in exerted force persisted and were dependent on motion, suggesting (1) a coupling of force and motion, and (2) the existence of an underlying

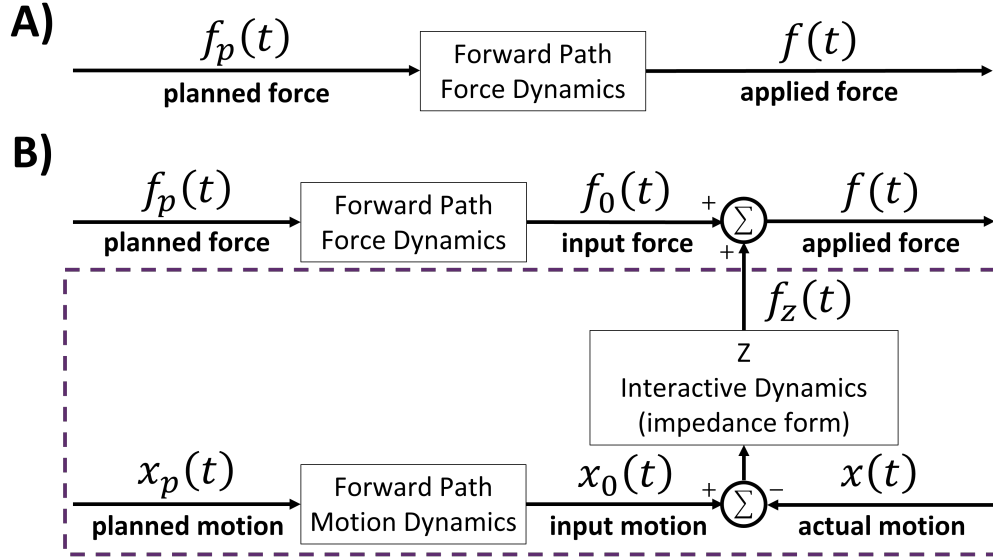


Figure 3-1: (a) **Direct force control.** Applied force is a function of a planned force that is independent of motion. (b) **Indirect force control.** Applied force is a function of a planned force but also depends on motion. Further details are in the Discussion section.

structure in the feedforward motion planning signal. Additional analysis of previous data from [Maurice et al., 2018] further validated the current results. In sum, this work showed that interactive dynamics are significant and of particular concern in (1) quantification of human performance and in (2) physical human-robot interaction.

3.2 Methods

3.2.1 Participants

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

3.2.2 Experimental Procedures

Task and Instructions

Participants were instructed to hold the handle of a moving robotic manipulandum (HapticMaster) [Linde et al., 2002] and apply a constant 5N force in the direction of robot motion (i.e., tangential direction) as it traversed an elliptical trajectory in a horizontal plane (Fig. 3-2a).

Participants performed the experiment seated and held the robot through a vertical handle which could pivot around its vertical axis; the pivot decoupled the robot end-effector orientation and the participant’s wrist orientation (Fig. 3-2b). Participants were positioned such that when holding the robot handle at the 270° position of the ellipse (Fig. 3-2d), the right upper arm hung downward slightly away from the torso. This position aligned the forearm with the minor axis of the ellipse. The robot height was adjusted such that the forearm was approximately parallel to the ground. This resulted in an angle between the upper arm and forearm of approximately 90° (Fig. 3-2a-b).

Visual Display

Participants sat approximately 2.2m in front of a projection screen (height: 1.8m, width: 2.4m). In conditions with visual feedback, two horizontal bars appeared on the screen (Fig. 3-2c). A red horizontal bar moved vertically to indicate the tangential force (averaged over 80ms) applied by the participant onto the robot; the stationary white bar indicated the desired tangential force of 5N. Otherwise, the screen was black.

Control of Robot Motion

The robot handle was commanded to move counterclockwise along an elliptical path (major axis = 30cm, minor axis = 10cm) on a horizontal plane with a period of 3s (Fig. 3-2d). The velocity profile of the robot handle followed the so-called 2/3 power-law relation [Maurice et al., 2018] [Zago et al., 2018] between path curvature

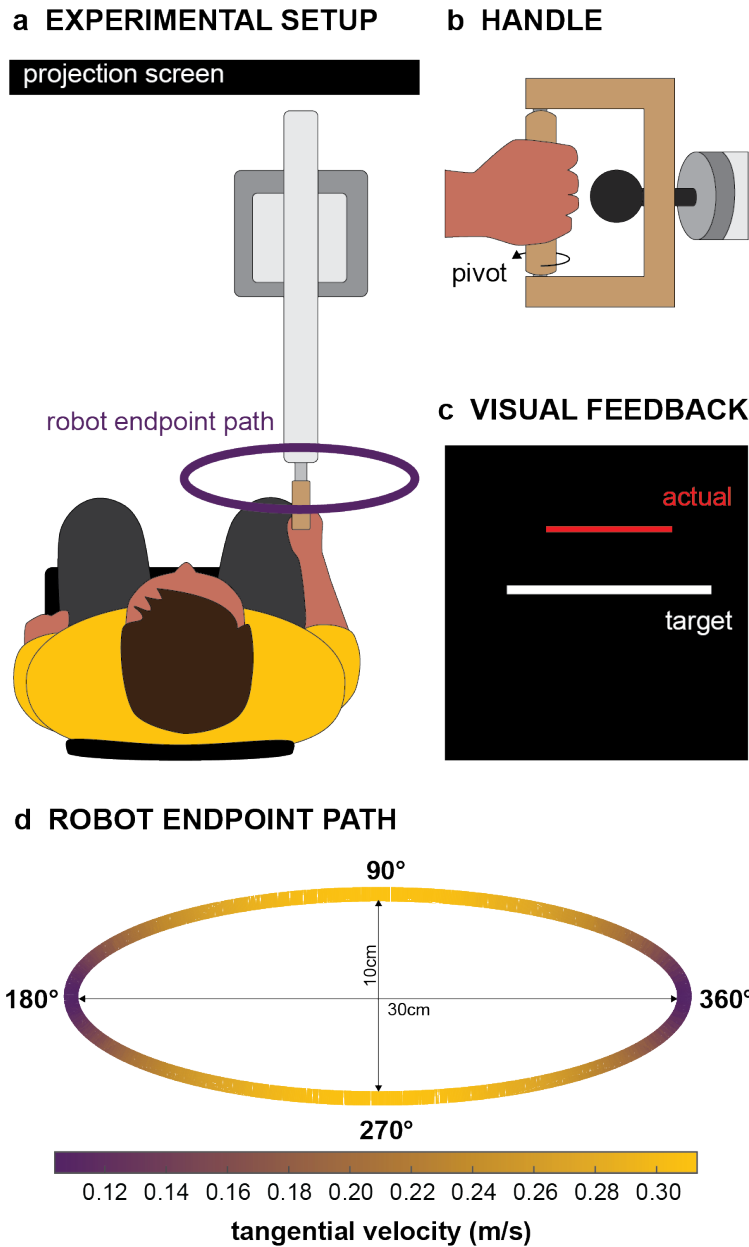


Figure 3-2: **Experimental setup.** (a) Top-down view of the experimental setup. Participants were instructed to hold the handle of a moving robotic manipulandum and apply a constant force in the direction of the robot’s motion. The elliptical path of the robot endpoint is displayed on the figure for clarity. However, participants did not see any visual display of the robot path. (b) Robot handle used to decouple human wrist and robot end-effector orientations. (c) To provide visual feedback, the projection screen displayed a stationary white bar indicating the target tangential force of 5N and a moving red bar indicating the current applied tangential force. Visual feedback was given during Block 1V and Block 2V. Otherwise, the screen was black. (d) Elliptical trajectory of the robot endpoint (i.e., handle) in the horizontal plane. The robot manipulandum moved counterclockwise and followed a velocity profile that was in accordance with the two-thirds power law [Zago et al., 2018]. Tangential velocity is shown by color.

and angular velocity (Fig. 3-2d), decreasing in highly curved portions and increasing in less curved portions. The position of the robot handle was controlled with a Cartesian PD controller; a high proportional gain was used such that deviation from the desired trajectory was negligible. The desired position of the robot was updated at 700Hz, and an internal force control loop ran at 2kHz.

3.2.3 Experimental Design

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

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

3.2.4 Dependent Measures and Data Processing

The force that participants applied to the robot handle was measured at ~ 560 Hz with a 3 DoF force sensor mounted at the robot end effector. In each trial, the

tangential component of the force applied by the human to the robot was calculated and resampled as a function of robot position along the elliptical path at a resolution of 1° .

Angular position along the elliptical path was defined using the eccentric anomaly¹, E , such that $E = \text{atan}(ay/bx)$, where a and b were half the length of the major and minor axes of the ellipse, respectively. x and y were the magnitudes of the position vector in the direction of the major and minor axes in Cartesian space, respectively.

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

3.2.5 Data Analysis and Statistics

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

Performance Improvements

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

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

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

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

Existence of Motion-Dependent Force Errors

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

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




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

3.3 Results

3.3.1 Performance Improvements

Change in RMS Force Error Within Blocks

Inspection of the grouped time series of force error revealed that subjects showed a consistent decline of the force error in the first part of Block 1V (Fig. 3-3a). The iteratively computed linear regressions between the average RMS force error across participants and trial for the last 15, 14, 13, and so forth trials identified that the force error values in Block 1V plateaued when calculated over the last 9 trials (i.e., from trial 7 to 15). From trial 7 onwards the regression slopes did not differ from zero. As this initial drop of error seemed to be a result of familiarization, subsequent analyses only examined the last 9 trials of all four blocks to evaluate the errors reached in each condition.

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

To statistically evaluate whether visual feedback and practice had a significant effect on the force error, a 2 (block) x 2 (feedback) repeated-measures ANOVA was conducted. The force error revealed a significant interaction ($F_{1,10} = 15.74, p =$

$2.66e - 03$) as the mean RMS error decreased from Block 1V ($M = 3.63N, SD = 1.12N$) to Block 1NV ($M = 3.30N, SD = 1.12N$) and increased from Block 2V ($M = 3.08N, SD = 1.32N$) to Block 2NV ($M = 3.86N, SD = 1.96N$) (Fig. 3-4a). However, neither the main effect of block ($p = 0.99$), nor the main effect of feedback ($p = 0.60$) were statistically significant. Recall, all subjects completed the experiment in order: Block 1V, 1NV, 2V, 2NV. Thus, the increase in mean RMS of force in Block 2 was likely the result of cognitive or physical fatigue as the experiment was quite long.




Figure 3-4: Mean RMS of force error from the last 9 trials for (a) all experimental blocks in the main experiment, and (b) Experiment 0N. Yellow dots depict individual participants. Error bars depict ± 2 standard errors of the mean. The mean RMS of force error was significantly higher in all experimental blocks of the main experiment compared to Experiment 0N (Table I).

3.3.2 Existence of Motion-Dependent Force Errors

Given this indifference to feedback and practice, the time series of force error were inspected. As illustrated by the raw force data shown for a representative participant in Fig. 3-5, force error was periodic with pronounced peaks at multiples of 180° in all blocks. The means of each cluster identified in the maximum lag data of trials in Blocks 1V, 1NV, 2V, and 2NV were 179.3° and 359.3° (Fig. 3-6a). The average maximum auto-correlation coefficient was 0.43 ($SD = 0.13$). Trials with maximum lags of 360° indicate that the peaks in force error at half and full cycle were different, while the maximum lag at 180° indicates that the two peaks in force error were similar. Analyses of individual participants revealed that five subjects showed higher force applied at 180° and six subjects showed higher force applied at 360° . Taken together, these results indicate that force error strongly depended on the phase of the oscillatory robot motion.

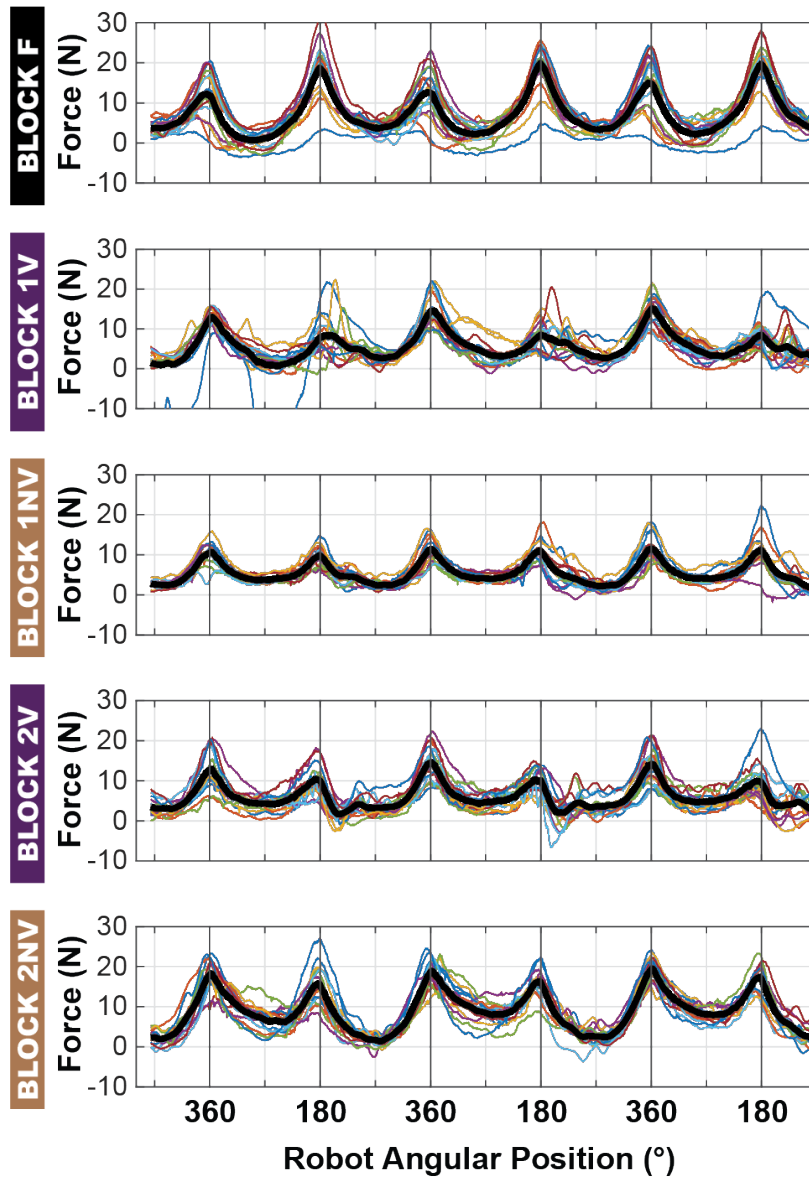


Figure 3-5: **Raw force data in each block for a representative participant.** For each of the 5 blocks in the main experiment, plots of tangential force over robot angular position are shown for the last 3 cycles of every trial. The tangential force was resampled as a function of robot position along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

Given this pronounced periodicity in the experimental blocks that specified 5N force, we also examined whether this periodicity was present spontaneously. The same autocorrelation analyses were run on the trials of the familiarization block (Block F). Fig. 3-6b shows two clusters with mean values of 179.9° and 359.5° . The average

maximum auto-correlation coefficient was 0.53 ($SD = 0.12$). As illustrated in Fig. 3-6b, these results give strong evidence for a spontaneous coupling of motion and force.

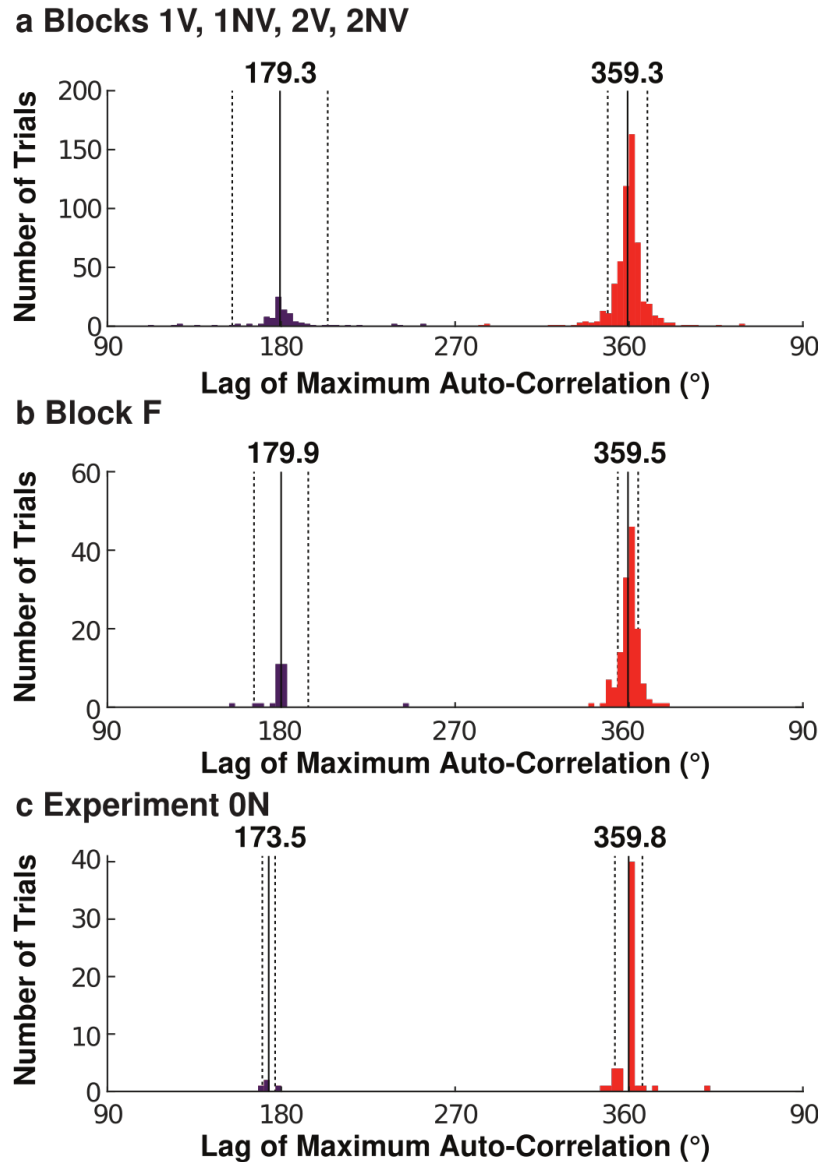


Figure 3-6: Evidence of motion-dependent periodic force errors for all trials in (a) Blocks 1V, 1NV, 2V, and 2NV of the main experiment, (b) Block F of the main experiment (c) Experiment 0N. Histogram of lags of maximum autocorrelation (referred to as maximum lag) values measured in units of robot angular position. In all conditions, two clusters were identified. The solid lines indicate the mean of each cluster, and the dashed lines depict ± 1 standard deviation of each cluster.

3.4 Additional Results

To further validate the existence of motion-dependent force errors, similar analyses were performed on data collected in a previously published study of human-robot interaction by Maurice et al. (see [Maurice et al., 2018] for full experimental details). The objective of that study was to examine how humans adapt to different velocity patterns in elliptic planar robot movements, identical to the ones used in this study. However, instead of being instructed to apply a constant 5N force in the tangential direction, participants were instructed to minimize the total force magnitude applied to the robot end effector (i.e., apply 0N total force). Participants ($N = 6$) performed 10 trials, where each trial consisted of 4 cycles as in the present study. No visual feedback was provided to participants at any point in the experiment. This dataset is referred to as Experiment 0N.

To allow comparison with the results of the main experiment, force error was defined as the difference between the actual tangential force and the target tangential force of 0N. Otherwise, all data processing methods and dependent measures were identical.

3.4.1 Performance Improvements

The linear regression between the average RMS force error calculated across participants and trial numbers was not statistically significant, indicating no evidence of change in task performance with practice (Fig. 3-3b).

3.4.2 Existence of Motion-Dependent Force Errors

Two clusters were identified in the maximum lag data. The means of each cluster were 173.5° and 359.8° (Fig. 3-6c), and the average maximum auto-correlation coefficient was 0.41 ($SD = 0.11$). Despite the difference in task instruction, these results were strikingly similar to those in the main experiment. Most importantly, they indicate that the force errors were also motion-dependent.

	<i>t</i> statistic (df = 15)	<i>p</i> value
Block 1V - Experiment 0N	4.40	$5.21e - 04$
Block 1NV - Experiment 0N	3.79	$1.77e - 03$
Block 2V - Experiment 0N	2.95	$9.87e - 03$
Block 2NV - Experiment 0N	2.02	$8.62e - 03$

Table 3.1: Statistical results comparing mean RMS of force error between main experiment and experiment 0N

3.4.3 Differences in Force Error Magnitude

While the force patterns were dependent on motion in both experiments, the magnitude of the force errors differed. Independent *t*-tests compared the mean RMS force errors from each block of the main experiment with those of Experiment 0N. Note that in all cases, the mean RMS force error was calculated over the last 9 trials for consistency. As summarized in Table I and depicted in Fig. 3-4a-b, the mean RMS force error was significantly higher in all blocks of the main experiment (Bonferroni adjusted $\alpha = 0.05/4 = 0.0125$). Despite more practice and added visual feedback, participants who aimed to apply 5N force on the moving robot performed considerably worse than those instructed to minimize total force applied.

3.5 Discussion

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

and motion are coupled as schematically shown in Fig. 3-1b.

3.5.1 Force Error

In the main (5N) experiment subjects were given visual feedback of their tangential force in two of the blocks (Fig. 3-2c). In contrast to static tasks, where visual feedback enables subjects to apply a constant force quite accurately [Massey et al., 1992], the elliptic motion of the robot manipulandum in this study significantly compromised the subjects' ability to regulate force. Subjects did not eliminate residual errors, which varied periodically with motion.

Interestingly, the overall magnitude of force errors was significantly lower when the target force was lower (Fig. 3-4). There are several plausible explanations why this occurred. One possibility is that greater force applied induced higher noise (i.e., signal dependent noise) [Osu et al., 2004]. Another possibility is that greater force applied induced higher hand impedance [Lipps et al., 2020], which would amplify any errors between the input and actual trajectories. This would provide further support for indirect force control (Fig. 3-1b).

Nonetheless, production of actual force $f(t)$ that equals input force $f_0(t)$ is possible using the indirect force control strategy of $f(t) = f_0(t) + Z\{x_0(t) - x(t)\}$ when $Z\{x_0(t) - x(t)\} = 0$ (Fig. 3-1b). This can be achieved in one of two ways: (1) zero interaction dynamics and/or (2) a simultaneous prediction of the input² trajectory $x_0(t)$ that matches the actual trajectory $x(t)$. Thus, it is critical to note that if motion-dependent force errors were not observed, it would be impossible to distinguish between the direct and indirect force control strategies. However, the force error we observed was dominated by motion dependency (Fig. 3-5). Specifically, the force error was periodic with maximum auto-correlation at lag corresponding to the 180° and 360° ellipse positions (Fig. 3-6).

These motion-dependent force errors were also observed in both the familiarization Block F of the main experiment, where subjects were instructed to apply a constant

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

tangential force, (Fig. 3-6b) and Experiment 0N (Fig. 3-6c), where subjects were instructed to apply zero force. In both, subjects did not receive any visual feedback. Despite some practice with and without visual feedback, the motion dependency of the applied force persisted throughout the main experiment (Fig. 3-5). This robust observation suggests an underlying structure in humans’ ability to regulate force during motion that limits the performance of this task.

3.5.2 Dynamic Primitives

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

Dynamic primitives do not preclude arbitrary patterns of force production. A sufficiently accurate internal model might be used to compute both $f_0(t)$ and the corresponding $x_0(t)$ (Fig. 3-1). However, if the parameters of oscillatory primitives used to plan the motion were limited, the time-course of force production would also be limited. Periodic force errors in our experimental results suggest that the controller appears to be content with “good-enough” performance, which can be obtained using a limited set of “primitive” oscillations and a sufficiently low mechanical impedance. Thus motor behavior constructed by dynamic primitives may result in performance limitations – such as the observed imperfect, periodic force regulation reported here (Fig. 3-5).

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

3.5.3 Limitations

In the main experiment, participants experienced the task for approximately one hour (300 cycles). It is possible that participants could learn to better regulate their force with additional practice (e.g., over multiple days). However, investigation of extensive practice was not the goal of our work. Humans regularly perform a variety of novel forceful interaction tasks with ease and apparently without requiring long-term practice. In fact, task performance slightly worsened at the end of practice in the main experiment, possibly indicating that fatigue set in. Hence, this study aimed to identify the performance that might be expected from intuitive and spontaneous human-robot interaction.

Force errors might also be ascribed to poor perception of the robot’s motion. However, the motion slowly ($\sim 0.33\text{Hz}$) followed a large elliptical path of 66.8cm in circumference. Additionally, if errors in the perception of the robot’s motion led to force errors, we would expect to see differences in error between the blocks that did and did not have visual feedback. Fig. 3-4 demonstrates that this was not observed. Motion dependent deviations from the instructed force were persistent throughout

the entire main experiment (Fig. 3-5, & 3-6a-b).

It is also possible that there may have been too much cognitive demand from mapping the vertical feedback display to the horizontal force. While this argument cannot be directly refuted from the results reported here, it is unlikely to account for our main result. Fig. 3-6b shows that subjects force error was motion dependent even before visual feedback had been provided. In short, the position dependence of force error was consistent throughout the main experiment, despite the presence of visual feedback.

3.5.4 Implications

Understanding the preferred control strategy employed by humans may guide the design of robot controllers to manage physical interaction. A roboticist may draw upon the proposed “building blocks” to program a simple controller to achieve a complex task [Ijspeert et al., 2002, Stulp et al., 2012, Peters and Schaal, 2006]. For example, a controller based on dynamic primitives has been used successfully (in simulation) to control a 2 DOF arm to manipulate a dynamically complex whip with 50 DOF in a targeting task [Nah et al., 2020]. Furthermore, the human body has a large number of redundant degrees of freedom. Kinematic redundancy has commonly been viewed as a difficult challenge to overcome, especially if control is performed via conventional optimization-based techniques. However, redundant degrees of freedom may be controlled by superposition of mechanical impedance primitives. Remarkably, unlike optimization-based methods, as the number of redundant degrees of freedom increased, control based on the superposition of impedance primitives improved; in effect, with greater redundancy control became easier [Hermus et al., 2021].

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

amputation prostheses, assistive exoskeletons, robot-aided rehabilitation, and physical human-robot collaboration.

3.6 Conclusions

In this work, we scrutinized a pervasive assumption: force and motion can be controlled independently (an idea referred to here as direct force control). To examine this assumption, subjects were asked to apply a constant force on a robot manipulandum that moved along an elliptical path with a speed profile consistent with the preferred pattern of human motion (the two-thirds power law). Results showed that subjects were unable to control force accurately during motion, despite some practice and the presence of visual feedback; errors in force were periodic in response to the periodic motion of the robot. These results point towards an indirect force control formulation (Fig. 3-1b), in which commanded motion acts through mechanical impedance to evoke force. Furthermore, the periodic pattern of path-dependent force errors was consistent with commanded motion composed of oscillatory primitives. Taken together, these findings suggest that a relatively simple mathematical model combining dynamic motion primitives with mechanical impedance, as an additional primitive, is competent to describe how humans control contact and physical interaction. A quantitative model is especially important for designing devices that physically collaborate with humans.

Chapter 4

Stiffness Can be Perceived via Visual Observation of Path Information

In Chapter 3, we established that humans cannot independently manage force and motion. Rather, they are coupled via mechanical impedance. The next key question becomes: Is impedance essential to humans' internal model? One way to tease this out is to draw upon the well-cited human motor action motor perception relationship [Meltzoff and Moore, 1997, Grush, 2004, Bandura, 1986] and ask: Can humans perceive mechanical impedance without access to force information? This question is addressed in this chapter, where we demonstrate that an aspect of mechanical impedance, stiffness, can be perceived without contact. Moreover, we illustrate that stiffness estimates are made primarily through inferring path information. These results not only provide insight into the mechanisms underpinning human motion perception but also carry significant implications for the robotics and rehabilitation fields.

This chapter is an adapted version of a paper published in PLOS Computational Biology titled, "Role of path information in visual perception of joint stiffness" [West et al., 2022b]. This work was done in collaboration with Meghan Huber and Neville Hogan.

A video summary of this chapter can be found at: <https://www.youtube.com/watch?v=-bbXeqr7b0>

4.1 Introduction

When observing another human, we can only see their overt behavior (e.g., motion); we cannot perceive the underlying neural signals that generate the behavior. Yet humans have an astonishing ability to extract latent information from visually observing the movement of others (for review, see [Blake and Shiffrar, 2007]). This impressive ability has been best shown in a plethora of point light animation studies, in which subjects were shown the motion of only a small subset of points on the body. Even from such sparse motion information, humans can easily determine intention from arm movement [Pollick et al., 2001], distinguish emotion from patterns in dancing [Walk and Homan, 1984], and identify individuals from gait patterns [Loula et al., 2005]. In the context of motor learning, Mattar and Gribble [Mattar and Gribble, 2005] showed that subjects who observed the motion of other humans reaching in an unseen force field subsequently performed better when reaching in a similar unseen force field. That study demonstrated humans’ ability to learn about novel force environments solely based on visual observation of kinematics during physical interaction. As a whole, these studies demonstrate that humans can determine latent information from visual observation of motion; extrapolating, it may be possible that humans can interpret underlying control signals used to generate motion too.

Aligned with these results, in previous work we found that humans could infer dynamic properties, specifically joint stiffness, from multi-joint limb motion [Huber et al., 2017, Huber et al., 2019]. In the prior study, subjects observed the motion of a simulated stick-figure, two-link planar arm on a computer screen and then estimated its stiffness on a numeric scale. To mimic aspects of human neuromotor control, the arm movement was driven by the superimposition of a hand-space mechanical impedance controller and a joint-space mechanical impedance [Hogan, 2017], where mechanical impedance is characterized mathematically by the dynamic relation between motion and a resultant force [Hogan, 1985a, Hogan, 1985b, Hogan, 1985c]. Results showed that subjects’ stiffness estimates positively correlated with

the joint stiffness values used in the control policy, indicating that they could estimate changes in joint stiffness. Remarkably, this was possible without force information or explicit knowledge of the underlying limb controller. It is impossible to unequivocally quantify features such as limb stiffness from motion alone. Thus, humans must have used prior knowledge to estimate latent information from visual observation of motion. To estimate limb stiffness, for instance, their prior knowledge had to be congruent with the relation between limb stiffness and motion produced by the control policy used to drive the simulated limb. However, there are still open questions regarding the form of the prior knowledge used and how it was acquired.

One possibility is that shared resources are used for action execution and action perception [Meltzoff and Moore, 1997, Grush, 2004, Bandura, 1986]. The existence of mirror neurons—neurons in the premotor cortex that respond both when a subject performs an action and observes another person perform that same action—supports this notion [Gallese and Goldman, 1998]. Behavioral studies showing that infants [Meltzoff, 2002] and adults [Wolpert et al., 2011] learn from observing and imitating the motor behavior of others also indicates a strong link between action and perception. Importantly, this relation appears to be bi-directional [Hecht et al., 2001] that is, both transfer of knowledge from perception to action and from action to perception exists. Considering this relation, studying how humans perceive movement can inform how humans produce movement.

The possibility that humans used knowledge of their own neuromotor system to perform the visual-perception-of-stiffness task suggests that the controller used to simulate the arm motions was an adequate approximation of neuromotor control of upper limb movements. In the previous experiments, the simulated arm movements were produced using a motor controller built on dynamic primitives as proposed by Sternad and Hogan [Hogan, 2017, Hogan and Sternad, 2012]. The authors proposed that encoding behavior in terms of primitive actions can simplify the control of the otherwise complex neuromuscular system. Specifically, the prior simulations were a composition of two mechanical impedances, one referenced to a nominal joint configuration, the other referenced to an underlying oscillatory motion. Humans’

ability to estimate stiffness from purely visual information indicates a relation between motor action and motor perception. Moreover, it demonstrates that some of the ‘primitive’ elements believed to underlie motor production (in this case stiffness) also figure prominently in sensory processing and perception.

While the composition of motor behavior with dynamic primitives may serve as a competent model of motor generation and perception, how motor behavior is represented in the nervous system remains an open question. Extensive research shows that parameters such as motion and force are correlated with neural activity, but we have limited understanding of how they are integrated to generate motor commands or interpret motor behavior (for review, see [Branco et al., 2019]). Furthermore, there is evidence to show that aspects of mechanical impedance are also correlated with neural activity [Haruno et al., 2012]. However, it is unclear how, or even whether, mechanical impedance is encoded in the nervous system. For instance, it is possible that the prior knowledge subjects used to estimate stiffness in our previous work included other aspects of mechanical impedance such as velocity-dependent force (damping). The goal of the study presented here was to further probe the prior knowledge humans use to estimate changes in limb stiffness from visual observation. Specifically, we investigated the role of temporal information in the visual perception of stiffness.

Highly regular patterns exist in the temporal aspects of human motor behavior. When humans generate motion, it is commonly observed that their tangential hand velocity changes logarithmically with the radius of curvature of its path, the so-called $1/3$ power law (also commonly referred to as the $2/3$ power law, depending on formulation) [Huh and Sejnowski, 2015, Viviani and Flash, 1995]. Furthermore, removing this temporal pattern worsens motion perception. For instance, humans perceive motions that follow the $1/3$ power law to be more natural [Bidet-Ildei et al., 2006] and uniform [Viviani and Stucchi, 1992]. They can also anticipate the motion of a system more accurately when it follows the $1/3$ power law [Kandel et al., 2000b]. Additionally, Dayan et al. [Dayan et al., 2007] found stronger and more extensive neural responses, especially in motor-related areas, when humans perceived motion

that followed the $1/3$ power-law compared to motion that did not. Maurice et al. [Maurice et al., 2018] also showed that humans can control physical interaction with a robot better when its velocity profile follows the $1/3$ power law. Moreover, both velocity magnitude (i.e., speed) and direction are also well-correlated with the motor cortical activity of the brain [Moran and Schwartz, 1999]. Thus, temporal information appears to play a key role in biological motion perception and understanding.

Given the important role of temporal patterns in the generation and perception of human motor behavior, we hypothesized that changing the velocity profile of the simulated arm without changing its path would hinder subjects' ability to estimate limb stiffness from visual observation. Conversely, if changing the velocity profile had no effect, it would suggest that temporal information is at least subordinate to path information. Here, in Experiment 1, subjects viewed simulated arm motions in which the joint paths varied with simulated stiffness as determined by the dynamic simulation; however, the simulated velocity profiles were modified in three different ways. Counter to our hypothesis, the presence and type of velocity manipulation did not significantly affect subjects' ability to estimate limb stiffness. Given these results, we conducted a follow-up experiment to determine if temporal information had *any* influence on the perception of stiffness. Thus, in Experiment 2, subjects observed simulated arm motions that all followed the same joint path but with different velocity profiles. Subjects perceived simulations with a veridical velocity profile to be less stiff than that of a non-veridical velocity profile; but there was no difference in subjects' stiffness estimates of simulations that followed a non-veridical velocity profile. Together, these results suggest that while temporal information may influence humans' stiffness perception, path information is the predominant factor used by humans to visually estimate changes in limb stiffness. These observations provide further insight into humans' representation of motor behavior and how humans interpret and learn from the motor actions of others.

4.2 Experiment 1: Methods

4.2.1 Ethics Statement

A total of 30 subjects (15 males and 15 females with a mean age of 25.5 ± 5.6 years) took part in Experiment 1 (10 in each of the 3 experimental conditions). Subjects had a variety of educational backgrounds, and none had any prior experience with the experimental task. All subjects gave informed written consent before the experiment. The experimental protocol was reviewed and approved by the Institutional Review Board of the Massachusetts Institute of Technology (MIT IRB Protocol #1508154608). Data of an additional 10 subjects collected as part of a prior study (Experiment 2; [Huber et al., 2019]), referred to here as the *original* condition, were used for comparison in the statistical analyses of Experiment 1 in the present study.

4.2.2 Experimental Protocol

In each trial, subjects were instructed to observe a stick-figure display of a two-link planar arm move along a closed path for 20s and then estimate its stiffness on a numeric scale from 1 (“least stiff”) to 7 (“most stiff”) (Fig. 4-1). During this time, participants were not restricted from moving their body during the experiment, and they often mimicked the movement using their arm. Note that the endpoint path was not explicitly displayed. After submitting their estimate, subjects self-initiated the next trial. S1 Video demonstrates the display that subjects interacted with during these experiments.




Figure 4-1: **Set up for Experiments 1 and 2.** In each trial, subjects were instructed to observe a stick-figure display of a two-link planar arm move along a closed path for 20s and then estimate its stiffness on a numeric scale from 1 (“least stiff”) to 7 (“most stiff”).

Each subject performed 30 trials. Subjects were shown six unique arm motions, each of which was simulated with a different value of elbow stiffness, repeated five times in a blocked manner. The order was randomized within each block. After completing the experiment, subjects were asked to provide a written description of their strategy for estimating “arm stiffness”. The whole experiment lasted approximately 20 minutes.

A custom MATLAB program (The Mathworks, Natick, MA) was used to simulate and display the arm motions and to record subjects’ stiffness estimates. Participants sat approximately 20” in front of a laptop screen (12” w \times 7” h) on which the arm motions were displayed. On the screen, the length of each arm link was \sim 1”.

4.2.3 Simulated Arm Motions

The arm was modelled as a two-link planar manipulator moving in a vertical plane and was driven with a controller comprising two attractors, inspired by the proposal

that human motor behavior is composed of dynamic primitives [Hogan, 2017, Hogan and Sternad, 2012]. The first attractor was a combination of an oscillatory primitive with mechanical impedance in endpoint coordinates, acting to pull the endpoint along a circular path, and the second was a combination of a fixed-point primitive with mechanical impedance in joint coordinates, acting to pull it to a nominal joint configuration.

The dynamics of this model were described as:

$$M(q)\ddot{q} + C(q, \dot{q})\dot{q} + g(q) = \tau \quad (4.1)$$

where $q, \dot{q}, \ddot{q} \in \mathbb{R}^{2 \times 1}$ are the joint angular positions, velocities, and accelerations, respectively, $M(q) \in \mathbb{R}^{2 \times 2}$ is the inertia matrix, $C(q, \dot{q}) \in \mathbb{R}^{2 \times 2}$ are the Coriolis and centrifugal terms, $g(q) \in \mathbb{R}^{2 \times 1}$ are the gravitational terms, and $\tau \in \mathbb{R}^{2 \times 1}$ are the controller joint torques. The length, mass, center of mass, and moment of inertia parameters for the two links were chosen to match the forearm and upper arm of an average, male human as described in [Zatsiorsky, 2002]. The controller joint torques τ were determined by

$$\tau = J(q)^T K_x (x_r - x) - J(q)^T B_x \dot{x} + K_q (q_r - q) \quad (4.2)$$

$$x_r = \begin{pmatrix} 0.1 \cos\left(\frac{20\pi t}{3}\right) \\ 0.1 \sin\left(\frac{20\pi t}{3}\right) \end{pmatrix}, \quad q_r = \begin{pmatrix} \frac{\pi}{4} \\ \frac{\pi}{4} \end{pmatrix} \quad (4.3)$$

$$K_x = \begin{pmatrix} 500 & 0 \\ 0 & 500 \end{pmatrix}, \quad B_x = \begin{pmatrix} 10 & 0 \\ 0 & 10 \end{pmatrix}, \quad K_q = \begin{pmatrix} 0 & 0 \\ 0 & E \end{pmatrix} \quad (4.4)$$

where $x, \dot{x} \in \mathbb{R}^{2 \times 1}$ were the endpoint (i.e., hand) positions and velocities, respectively, $J(q) \in \mathbb{R}^{2 \times 2}$ was the Jacobian matrix, x_r was the reference endpoint position, which followed a circular path, q_r was the reference joint configuration which was constant, K_x and B_x were the endpoint stiffness and damping matrices, respectively, K_q was the joint stiffness matrix, and $E \in \mathbb{R}_{\geq 0}$ was the value in the

joint stiffness matrix corresponding to the elbow joint. The six unique arm motions were generated by setting $E = \{0, 10, 20, 30, 40, 50\} Nm/rad$ (Fig. 4-2). The range of elbow stiffness values used was akin to those reported in human studies [Bennett et al., 1992, Lacquaniti et al., 1982, Mussa-Ivaldi et al., 1985]. The dynamic simulation of the arm used in this study was identical to that used in Experiment 2 of our previous study [Huber et al., 2019].




Figure 4-2: **The six endpoint motions of the simulated arm in Experiment 1.** Elbow stiffness (E) was varied. During the experiments, subjects only saw the moving limb and were not shown the endpoint traces displayed here. The distance between the centroids of the orbital endpoint paths, starting from $E = 0 \text{ Nm/rad}$ to $E = 10 \text{ Nm/rad}$, were 0.24", 0.15", 0.09", 0.05", 0.03", respectively, when measured on screen.

While all subjects in Experiment 1 observed the arm move so that it followed the same six endpoint paths (Fig. 4-2), the velocity profiles along those paths varied across four experimental conditions (Fig. 4-3, S1 Video and S2 Video). Subjects

in the original condition saw the joints move with a velocity profile governed by the aforementioned dynamic simulation. In the constant condition, the tangential velocity of the endpoint, v , was set to a constant value of 0.185 m/s for all six motions (Fig. 4-4). In the inverse condition, tangential velocity varied as a power of the endpoint path's radius of curvature $R(t)$:

$$v(t) = KR(t)^{-1/3} \quad (4.5)$$

where K was the velocity gain tuned to match the period of arm motions shown in the *constant condition* (Fig. 4-4). In the *inverse condition*, the speed of the endpoint increased with the curvature of its path. Note that this is the inverse of the typical power law relation between speed and curvature observed in human motor behavior (for review see [Zago et al., 2018]). The velocity manipulations implemented in the constant and inverse conditions were chosen for this study because they have been previously proven to affect both motion perception and production [Dayan et al., 2007, Maurice et al., 2018]. In the *variable condition*, the relation between tangential velocity and radius of curvature changed in each condition. The differential equation of the model dynamics was solved using the ODE45 function in MATLAB to obtain a time vector at 1ms resolution. A new time vector was generated to maintain constant tangential velocity for the endpoint path produced from $E = 50 \text{ Nm/rad}$ (*group1*). This newly generated time vector was used to generate the arm motion with six different endpoint paths. The position vectors differed for each endpoint path, but the time vector stayed the same. As a result, the relation between speed and curvature differed across endpoint paths and did not follow a power law (Fig. 4-4). A video of the different arm simulations subjects observed can be found in S2 Video.

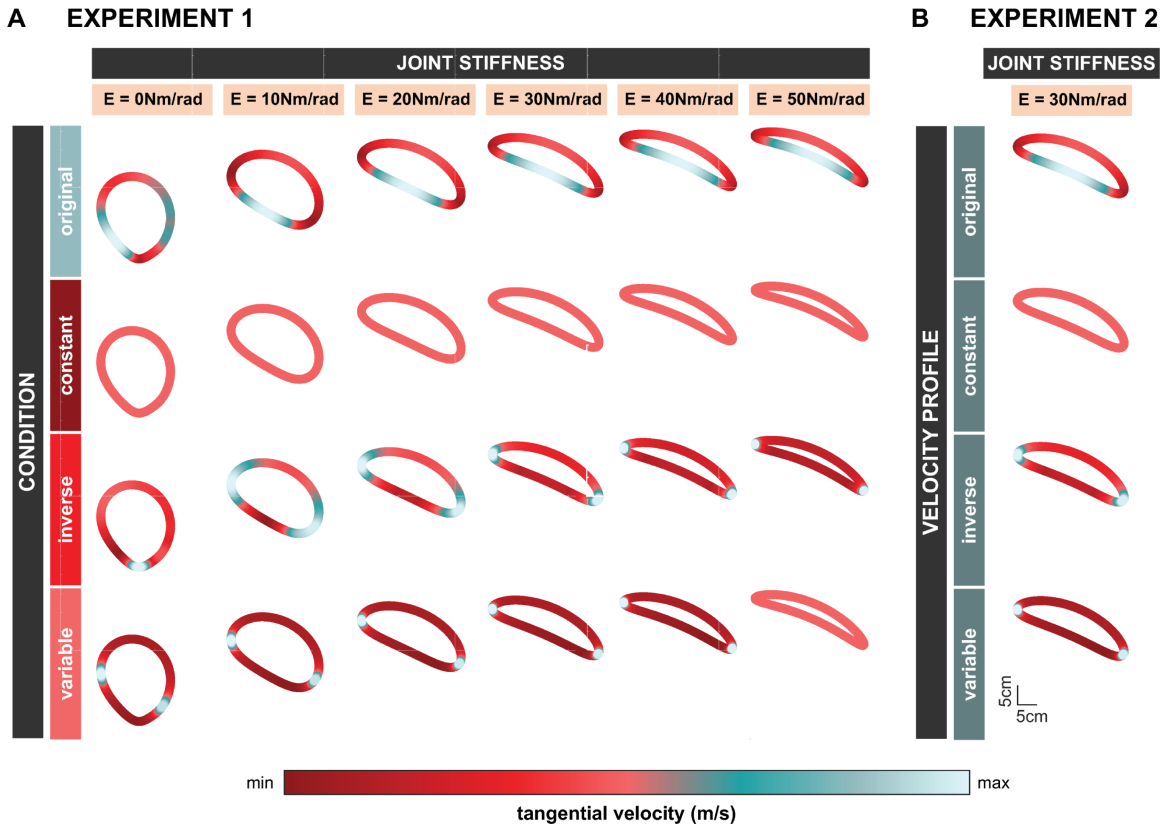


Figure 4-3: **(A) Simulated arm motions in Experiment 1.** The simulated endpoint velocity profiles for each endpoint path across all conditions in Experiment 1 are shown. In the original condition, the velocity profiles followed from the dynamic simulation. In the constant condition, the velocity profiles were manipulated to be constant. In the inverse condition, the velocity profiles were manipulated to follow the inverse of the veridical power-law relation. In the variable condition, the relation between tangential velocity and radius of curvature changed with each simulated stiffness; the velocity profiles did not have a simple velocity-curvature power-law relation (see Fig. 4-4). **(B) Simulated arm motions in Experiment 2.** In Experiment 2, the endpoint path (and simulated stiffness) remained the same across the four different arm motions shown to participants, while the temporal pattern along the path differed. The endpoint path was generated with $E = 30\text{ Nm/rad}$; the velocity profiles along the path were chosen from the four conditions in Experiment 1 and are referred to as original, constant, inverse, and variable, accordingly.

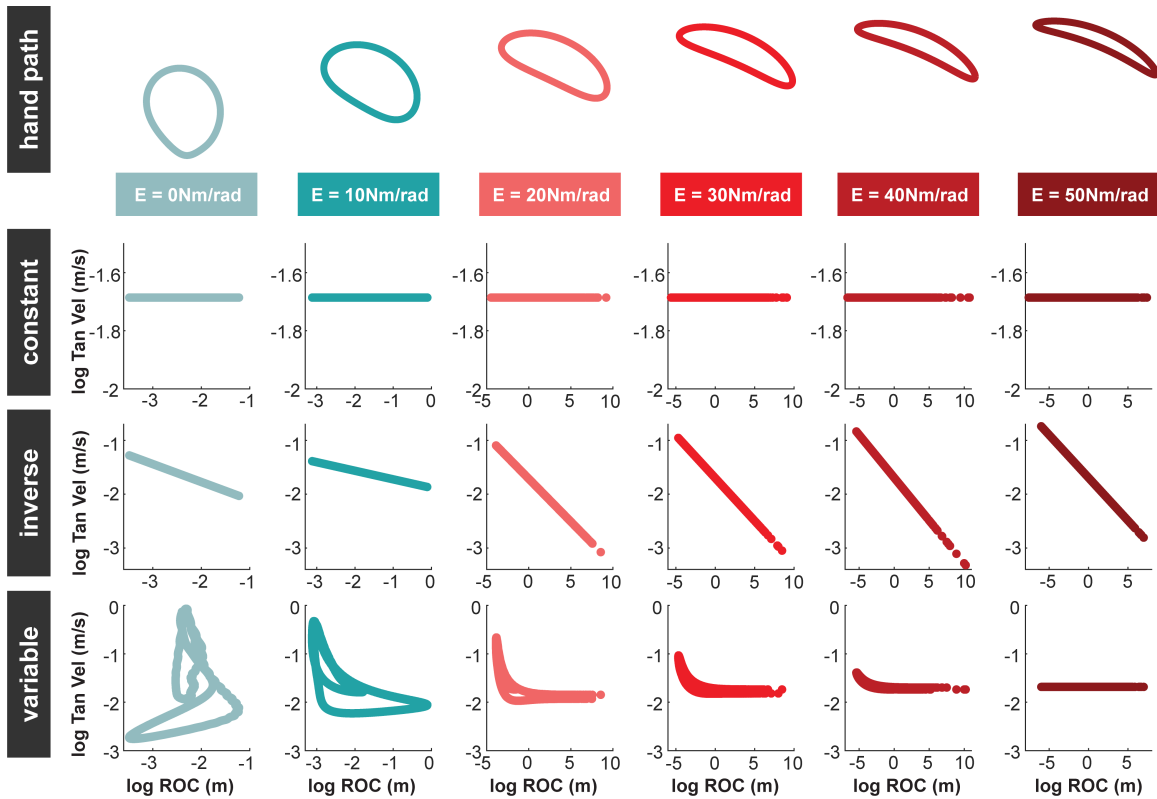


Figure 4-4: **The relation between radius of curvature (ROC) and tangential velocity (Tan Vel) for each endpoint path in the constant, inverse, and variable conditions of Experiment 1.** In the *constant condition*, this relation was constant for all endpoint paths. In the *inverse condition*, tangential velocity increased with radius of curvature for all endpoint paths. In the *constant* and *inverse conditions*, the relation was the same for all endpoints path within each condition. In the *variable condition*, this relation differed for each endpoint path. Specifically, the time vector generated to maintain constant tangential velocity for the endpoint path produced from the $E = 50 \text{ Nm/rad}$ simulation (*constant condition*) was applied to all six different endpoint paths.

While the period of arm motion increased slightly across the six motions displayed in the *constant* and *inverse conditions* (from 3.33 to 3.68 seconds), it was constant in the *original* (3.33s) and *variable* (3.68s) *conditions*.

4.2.4 Task Instruction

We intentionally did not provide subjects with any details regarding the underlying controller. Prior to the start of the experiment, subjects were not presented with examples of “more” and “less” stiff arm motions. They also did not receive feedback

regarding the accuracy of their estimates at any point during the experiment. If a subject was unsure of what the term stiffness meant, they received the following definition: “*Stiffness is the extent to which an object resists deformation or deflection in response to an applied force. A stiffer object has higher resistance to deflections than a less stiff object.*” In experiment 1, only 8 out of 30 subjects requested and were provided with a definition of stiffness.

4.2.5 Self-Reported Strategy for Stiffness Estimation

After subjects had viewed all the simulations, they were asked to write down their strategy for estimating stiffness. Subjects were not told that they would have to self-report their strategy before the experiment to prevent potential interference of conscious strategies. Each subject’s response was manually codified with four binary motion features of interest (joint motion, endpoint motion, path information, and temporal information) to quantitatively assess what type of motion information subjects used to estimate stiffness. The criteria used to set the value of each binary feature as either ‘yes’ or ‘no’ are presented in Table 4.1.

4.2.6 Statistical Analyses

To assess whether the estimates were similar across the four experimental conditions, we conducted a 4 (condition) \times 6 (joint stiffness) \times 5 (block) analysis of variance (ANOVA) of the arm stiffness estimate. ‘Condition’ was a between-subjects factor, and ‘joint stiffness’ and ‘block’ were within-subject factors. The Greenhouse-Geisser correction was applied to the within-subject factors.

In addition, a linear model of stiffness estimate as a function of joint stiffness was fit to the data for each subject. The coefficient of determination (R^2) was calculated for each subject. The R^2 value represents the fraction of the overall variance of the dependent variable (i.e., stiffness estimate) that can be accounted for by variability of the independent measure (i.e., simulated joint stiffness). It served as a performance measure of each subject’s ability to estimate changes in stiffness. R^2 values closer to

<u>Feature</u>	<u>Value</u>	<u>Criteria</u>
Path Information	'yes'	Use of the following words or phrases: -“distance” -“displacement” -“range of motion” -“angle”
	'no'	Did not meet criteria for a 'yes'
Temporal Information	'yes'	Use of the following words or phrases: -“speed” -“rate” -“acceleration” -“jerk” -“smooth”
	'no'	Did not meet criteria for a 'yes'
Joint Motion	'yes'	Use of the following words or phrases: -“joint” -“shoulder” -“elbow” -“angle” Hand-drawn picture of the arm with a pointer to at least one of the joints
	'no'	Did not meet criteria for a 'yes'
Endpoint Motion	'yes'	Use of the following words or phrases: -“endpoint” -“hand” Hand-drawn picture of the arm with a pointer to at least one of the joints
	'no'	Did not meet criteria for a 'yes'

Table 4.1: **Criteria used to encode the type of information subjects reported using to estimate stiffness in Experiments 1 and 2.**

1 indicated better fit of a linear model, and hence, better performance. A one-way ANOVA of the R^2 values with ‘condition’ as a between-subjects factor was conducted. This analysis tested whether the amount of unmodeled variability in subjects’ stiffness estimates differed across experimental conditions. A one-way ANOVA of the fitted slope values with ‘condition’ as a between-subjects factor was also conducted.

Binomial regression was conducted to assess the effect of experiment on the likelihood that subjects reported using path information, temporal information, joint motion, and endpoint motion to estimate stiffness. Linear regression was also

conducted to further investigate whether the reported use of the aforementioned motion features could predict subjects' abilities as quantified by the R^2 values of the stiffness estimate linear fits. Data from the *original condition* was not included in regression analyses since, in that experiment, subjects did not self-report their strategy in writing.

In all statistical tests, the significance level was set to $p = 0.05$. Statistical analyses were performed using SPSS Statistics for Windows, Version 24.0 (IBM Corporation, Armonk, NY).

4.3 Experiment 1: Results

4.3.1 Effects of Experimental Condition and Simulated Joint Stiffness on Stiffness Estimate

A three-way ANOVA revealed a significant effect of simulated joint stiffness [$F(1.64, 59.184) = 85.22, p < 0.001$]. Across all experimental conditions, subjects increased their stiffness estimate with the simulated joint stiffness used to generate the arm motion paths (Fig. 4-5). However, the remaining effects and interactions were not significant [condition: $F(3, 36) = 0.64, p = 0.59$; block: $F(2.67, 96.11) = 0.35, p = 0.77$; simulated joint stiffness \times block: $F(11.40, 410.36) = 1.21, p = 0.28$; simulated joint stiffness \times condition: $F(4.93, 59.18) = 1.22, p = 0.31$; block \times condition: $F(8.01, 96.11) = 0.22, p = 0.99$; simulated joint stiffness \times condition \times block: $F(34.20, 410.36) = 0.95, p = 0.55$].

Fig. 4-6 shows each individual subject's stiffness estimates for every motion path simulated with a different joint stiffness value, along with the linear model fit to each subject's data and the corresponding R^2 value. The better the linear model fit, the better the subject's ability to estimate changes in stiffness. Subjects across all experiments varied in their ability to estimate stiffness from motion, as indicated by the overall variation of R^2 values ($M = 0.55, SD = 0.25$; Fig. 4-7A–4-7D). A one-way ANOVA revealed that there was no significant effect of condition on the




Figure 4-5: **Experiment 1 stiffness estimate results.** In all four experimental conditions, there was a significant positive effect of joint stiffness on the arm stiffness estimates. Solid lines show the average arm stiffness estimate across subjects within each condition. The dashed line shows the average stiffness estimate across subjects in all conditions. Error bars represent ± 2 standard errors of the mean.

values [$F(3, 36) = 1.04, p = 0.39$] (Fig. 4-7E). A one-way ANOVA of the slopes of the linear fits similarly revealed no significant effect of condition [$F(3, 36) = 0.90, p = 0.45$] (Fig. 4-8A–4-8E). Across all experimental conditions, the average slope was 0.56 ($SD = 0.33$). These results indicate that the speed profile of the arm motions had no statistically detectable influence on subjects' ability to estimate stiffness.

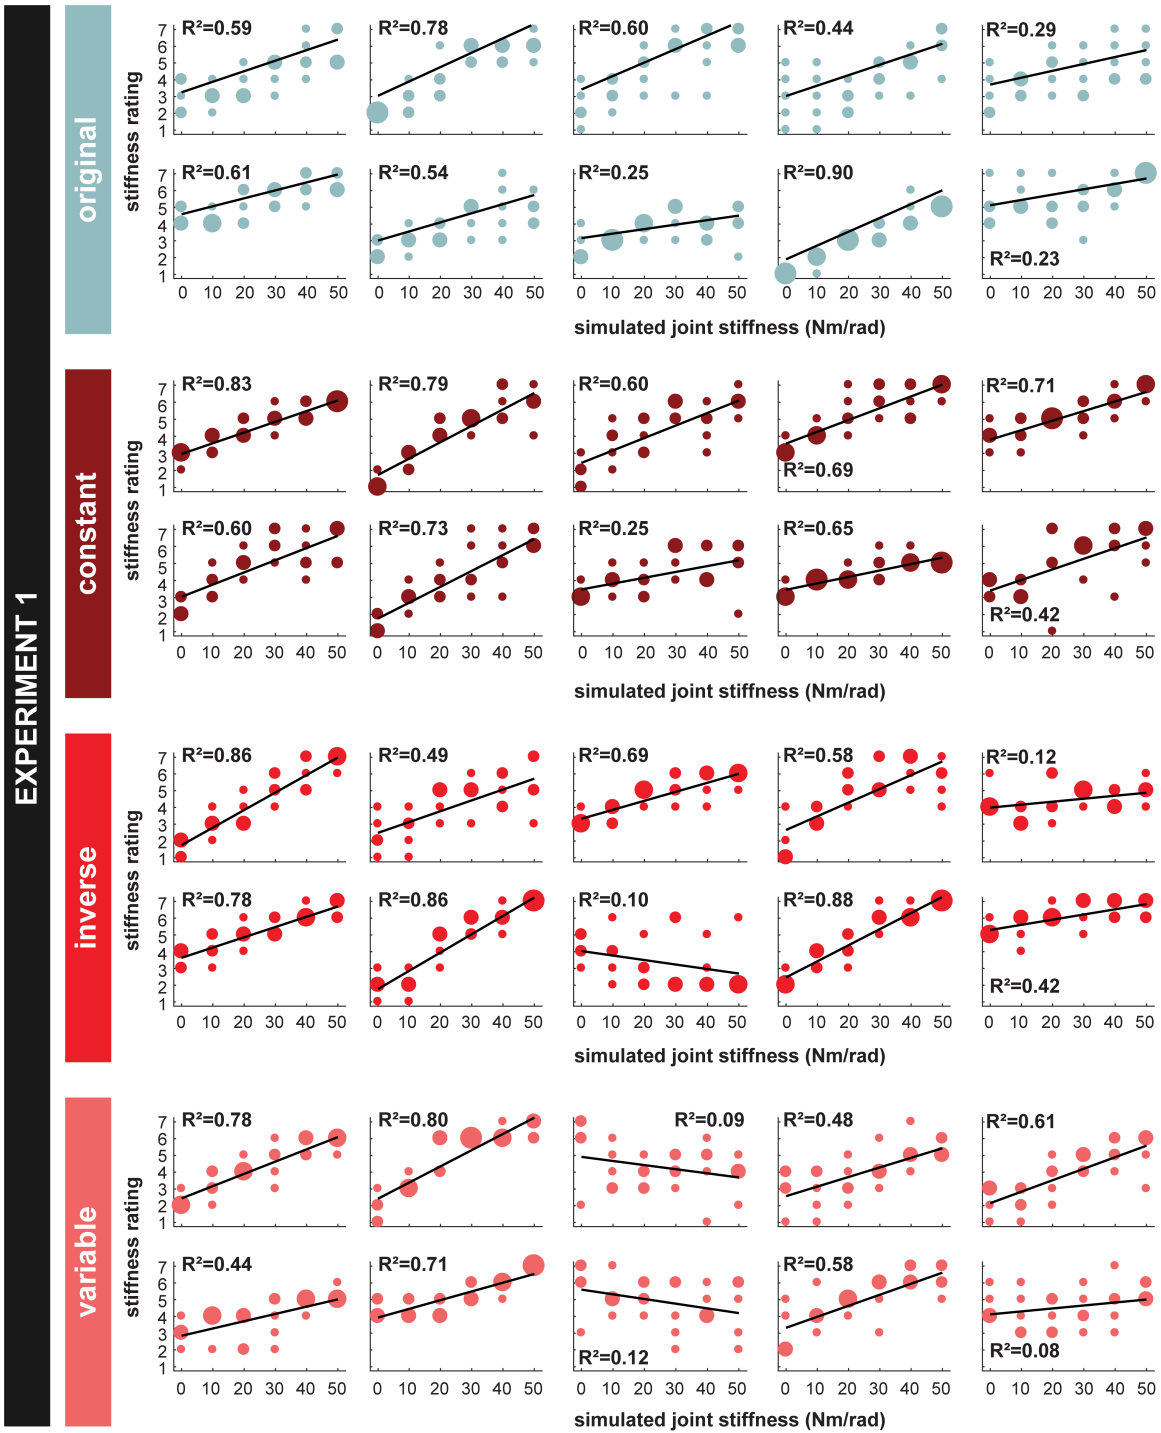


Figure 4-6: **Experiment 1 stiffness estimates for each individual subject.** All the individual subjects' stiffness estimates across simulated joint stiffness in all four conditions. Larger dots indicate greater response frequency. The black lines represent a linear fit of each subject's data. The coefficient of determination, R^2 , is also reported.

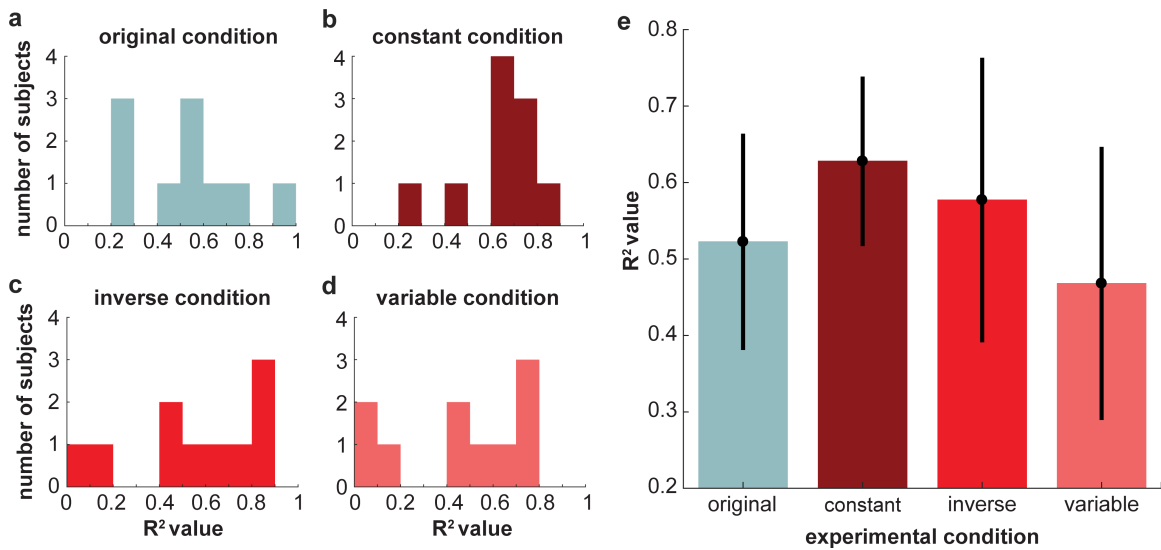


Figure 4-7: **Histograms of the coefficient of determination, R^2** in (a) the original condition, (b) the constant condition, (c) the inverse condition, and (d) the variable condition of Experiment 1. (e) Average coefficient of determination, R^2 , in each experimental condition. Error bars show ± 2 standard errors of the mean. There was no significant effect of experimental condition on the R^2 values.

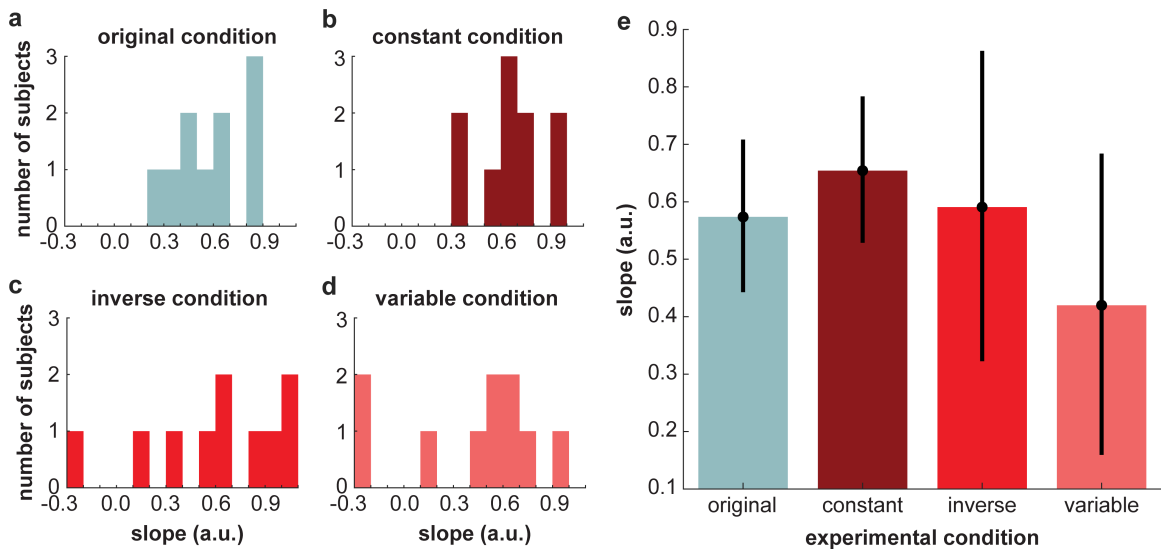


Figure 4-8: **Histograms of the slopes** from the fitted linear models in (a) the original condition, (b) the constant condition, (c) the inverse condition, and (d) the variable condition of Experiment 1. (e) Average slope in each experimental condition. Error bars show ± 2 standard errors of the mean. There was no significant effect of experimental condition on the slope values.

4.3.2 Effect of Experimental Condition on Motion Features Used to Estimate Stiffness

To further examine the motion features subjects used to estimate stiffness, analyses of the self-reported strategies were conducted. As seen in Fig. 4-9A, more subjects reported using path information ($N = 18$) compared to temporal information ($N = 8$) to estimate stiffness, and more subjects reported using joint motion ($N = 16$) compared to endpoint motion ($N = 2$). Most subjects ($N = 10$) reported using both path information and joint motion Fig. 4-9B). However, a similarly large number of subjects ($N = 8$) did not report using any of the four features.

Results of binomial regression found that experiment did not significantly affect the likelihood that subjects reported using any of the four motion features [path information: $\chi^2(2) = 0.84, p = 0.66$; temporal information: $\chi^2(2) = 2.62, p = 0.27$; joint motion: $\chi^2(2) = 1.08, p = 0.58$; endpoint motion: $\chi^2(2) = 1.69, p = 0.43$] (Fig. 4-9B). Results of the linear regression also found that none of the four motion features extracted from subjects were significant predictors of a subject's ability to estimate stiffness (i.e., R^2 value) [$F(4.29) = 1.90, p = 0.94$] (Fig. 4-9C).

4.4 Experiment 2: Methods

The results of Experiment 1 showed that manipulation of temporal patterns did not preclude subjects' ability to identify changes in simulated joint stiffness. It is still possible, however, that temporal patterns influence the magnitude of humans' stiffness estimates. Experiment 2 directly tested this possibility.

4.4.1 Subjects

Ten subjects took part in Experiment 2 (5 males and 5 females with a mean age of 25.6 ± 1.7 years). As in Experiment 1, subjects had a variety of educational backgrounds, and none had any prior experience with the experimental task. All subjects gave informed written consent before the experiment. The experimental protocol was

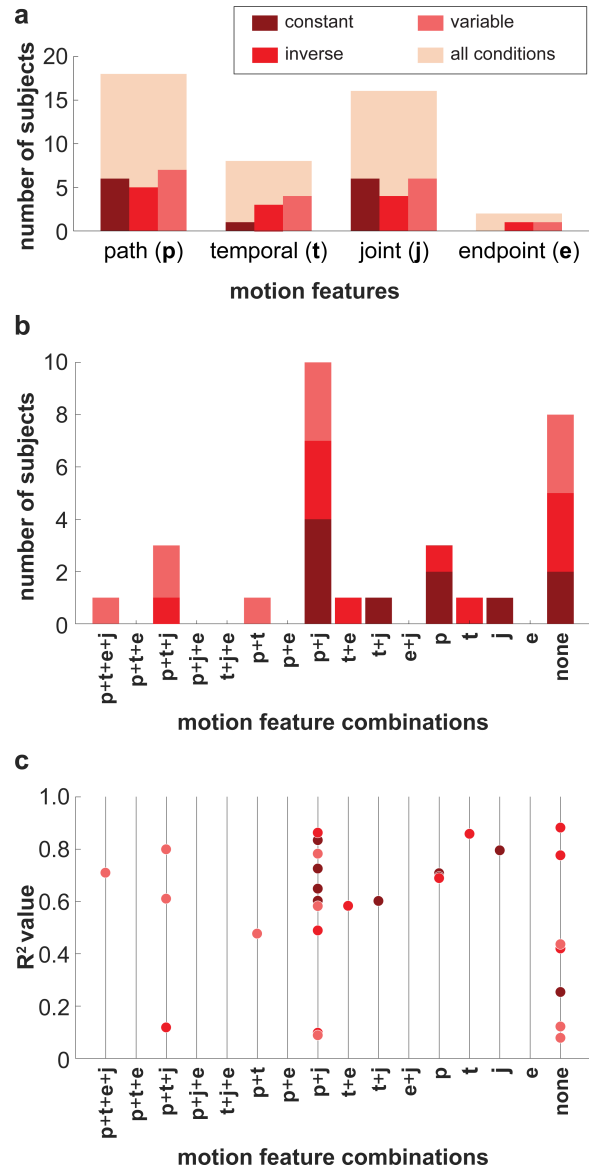


Figure 4-9: **Self-reported strategies in Experiment 1.**(a) Histogram of motion features that subjects reported using to estimate stiffness. (b) Stacked histogram of the sixteen possible motion feature combinations that subjects could have reported using to estimate stiffness (p: path information, t: trajectory information, j: joint motion, e: endpoint motion). A subject who reported using path information (p) and joint motion (j), but not trajectory information (t) and endpoint motion (e), for example, would be counted under the combination ‘p+j’. (c) R^2 values for each motion feature combination that subjects reported using to estimate stiffness. Each circle represents an individual subject. Color differentiates subjects based on the condition they participated in.

reviewed and approved by the Institutional Review Board of the Massachusetts Institute of Technology.

4.4.2 Experimental Protocol

The experimental protocol was identical to that of Experiment 1, except that subjects in Experiment 2 were shown four different arm motions, repeated five times in a blocked manner. Each subject performed 20 trials.

4.4.3 Simulated Arm Motions

The endpoint path (and simulated stiffness) remained the same across the four different arm motions, while the temporal pattern along the path differed. As shown in Fig. 4-3B, the endpoint path was generated with an elbow stiffness of $E = 30 \text{ Nm/rad}$; the velocity profiles along the path were chosen from the four conditions in Experiment 1 and are referred to as *original*, *constant*, *inverse*, and *variable*, accordingly.

4.4.4 Task Instruction

Task instruction was the same as Experiment 1. In Experiment 2, three subjects requested and were provided with a definition of stiffness.

4.4.5 Self-Reported Strategy for Stiffness Estimation

As in Experiment 1, after observing all simulations, subjects were asked to write down their strategy for estimating stiffness. Again, subjects were not told that they would have to self-report their strategy prior to the experiment. The same criteria as before (Table 4.1) were used to codify subjects' self-reported strategy.

4.4.6 Statistical Analyses

To assess the effect of velocity profile on stiffness estimate, we conducted a 4 (velocity profile) \times 5 (block) ANOVA on the arm stiffness estimate. 'Velocity profile' and 'block' were within-subject factors. The Greenhouse-Geisser correction was applied

to the within-subject factors. Planned comparisons using paired t-tests were made to probe the within-subject effect of ‘velocity profile’.

4.5 Experiment 2: Results

4.5.1 Effects of Velocity Profile on Stiffness Estimate with the Same Endpoint Path

A two-way ANOVA revealed a significant effect of velocity profile on stiffness estimate [$F(1.37, 12.31) = 5, p = 0.029$] (Fig. 4-10). Paired tests revealed that stiffness estimates were significantly lower for original velocity profile compared to the three other velocity profiles ($ps < 0.03$). There was no statistical difference in stiffness estimate across the other three velocity profiles ($ps > 0.18$). The remaining effects and interactions were not significant [block: $F(2.50, 22.53) = 1.08, p = 0.37$; velocity profile x block: $F(2.69, 24.20) = 0.56, p = 0.63$]. Fig. 4-11 shows each individual subject’s stiffness estimates for each velocity profile.

4.5.2 Self-Reported Stiffness Estimating Strategies

As in Experiment 2, subjects were asked to write down their strategy for estimating stiffness after having viewed all the simulations. Contrary to the results of Experiment 1, in Experiment 2 more subjects reported using temporal information ($N = 8$) as opposed to path information ($N = 6$) to estimate stiffness (Fig. 4-12A). However, this was to be expected given that path information did not change across motions in Experiment 2. Moreover, subjects in Experiment 2 rarely reported using either joint motion ($N = 1$) or endpoint motion ($N = 1$) (Fig. 4-12A). Rather, many subjects (40%) reported using both path and temporal information (Fig. 4-12B).

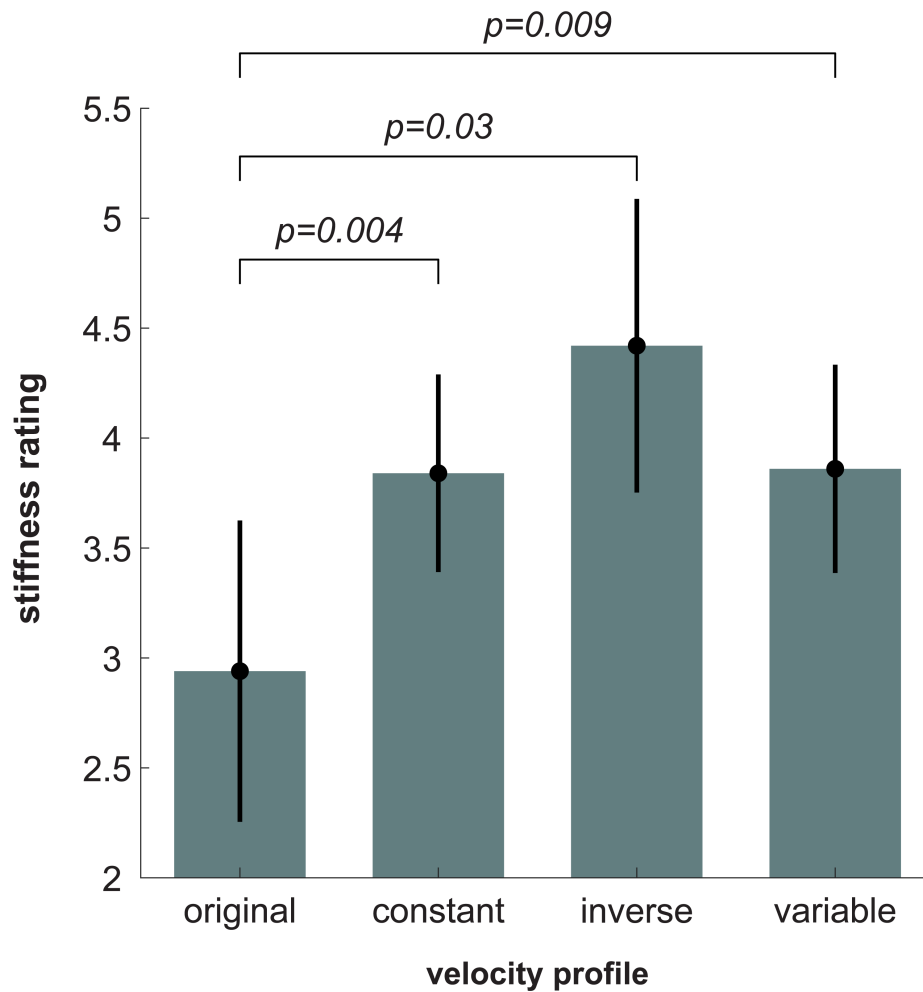


Figure 4-10: **Experiment 2 stiffness estimate results.** A significant effect of velocity profile on stiffness estimate was observed. Stiffness estimates for the original velocity profiles were statistically smaller than for the manipulated velocity profiles. There was no statistical difference in stiffness estimate between the three manipulated velocity profiles. Bars show the average arm stiffness estimate across subjects for each velocity profile. Error bars represent ± 2 standard errors of the mean.

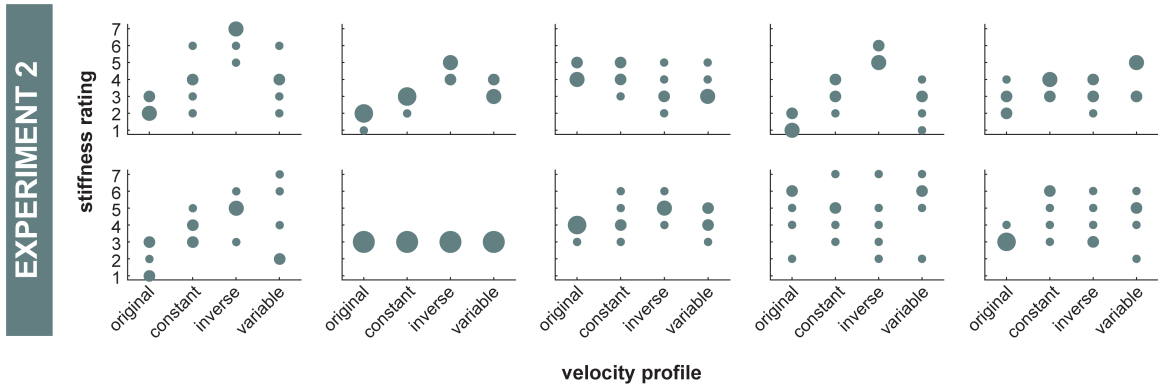


Figure 4-11: **Experiment 2 stiffness estimates for each individual subject.** All the individual subjects' stiffness estimates across the four velocity profiles. Larger dots indicate greater response frequency.

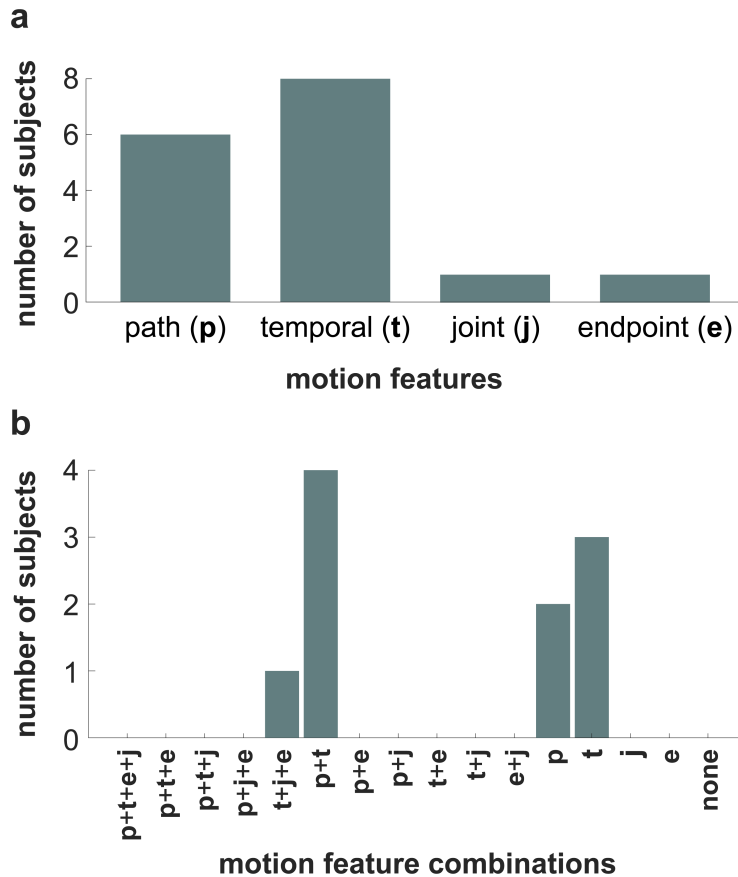


Figure 4-12: **Self-reported strategies in Experiment 2.** (a) Histogram of motion features that subjects reported using to estimate stiffness. (b) Histogram of the sixteen possible motion feature combinations that subjects could have reported using to estimate stiffness (p: path information, t: trajectory information, j: joint motion, e: endpoint motion).

4.6 Discussion

Given the observed relation between motor action and perception and the dependence of both on velocity, we hypothesized that manipulating the velocity profile would hinder subjects' ability to estimate stiffness. To test this hypothesis, we simulated arm motion with three new conditions. In each condition, we distorted the velocity profiles of a simulated arm by manipulating the velocity-curvature relation of its endpoint while keeping the paths the same. In all conditions, we found that despite the manipulation of temporal information subjects still increased their arm stiffness estimate with simulated elbow stiffness. Moreover, we did not observe a statistically significant difference in subjects' ability to estimate limb stiffness across the experimental conditions, including the original condition with veridical velocity profiles. These results emphasize the robustness of our prior finding that humans can estimate changes in limb stiffness from visually observing its motion [Huber et al., 2019, Huber et al., 2017]. While the results of Experiment 1 demonstrated that manipulation of temporal patterns did not preclude subjects' ability to identify changes in simulated joint stiffness, Experiment 2 directly tested the possibility that temporal patterns influence the *magnitude* of humans' stiffness estimates. Specifically, subjects observed 20 arm simulations that followed the same path (i.e., simulated stiffness), but had 4 different velocity profiles. Analysis demonstrated that subjects estimated simulations with the original velocity profile to be significantly less stiff than simulations with a manipulated velocity profile. Given these results, we conclude that temporal information does influence visual perception of stiffness. However, path information is the predominant factor used by humans to visually estimate changes in limb stiffness. These results provide further insight into how humans interpret and learn from the motor actions of others.

Considerable research has shown that temporal information plays a key role in both the generation and perception of biological motion. And yet, our results show that changing the temporal patterns of the arm motion did not render subjects unable to estimate limb stiffness. However, it is possible that the Likert scale used

to quantify subjects' stiffness estimates was not sensitive enough to capture small differences across experiments. Discrete numeric scales are vulnerable to quantization error, and the variance of quantization error is inversely proportional to the number of estimate options (i.e., the lower the number of estimate options, the higher the noise) [Bennett, 1948]. Additionally, the Likert scale is susceptible to response bias, which occurs whenever a person responds systematically on some basis other than what the items were specifically designed to measure [Bertram, 2007]. For instance, it is possible that subjects could have different interpretations of the term "stiffness," leading to increased variability in the strategies used for producing stiffness estimates. If subjects asked what the term "stiffness" meant, they were given the following definition: "Stiffness is the extent to which an object resists deformation or deflection in response to an applied force. A stiffer object has higher resistance to deflections than a less stiff object." As mentioned earlier, only eleven subjects asked for and were given this definition; However, omitting those subjects from the statistical analysis did not change the significance (or lack of significance) of the results. It is important to emphasize that if our definition of stiffness had been given to all participants, an assessment of whether the instruction potentially biased participants' performance could not have been made. This motivated our decision to only provide an instruction if explicitly asked. Another common form of response bias is central tendency bias, which occurs when subjects avoid selecting the most extreme results. As seen in Fig. 4-6, not all subjects used the full range of the Likert scale, indicating that they were affected by such bias. Some subjects even explicitly reported looking for the most and least stiff simulations; however, those simulations never came. Since central tendency bias effectively reduces the number of estimate options, it further increases measurement noise for a given subject. Nonetheless, we used the Likert scale to quantify subjects' stiffness estimates to allow comparison of our experimental results with those of our prior studies [Huber et al., 2019, Huber et al., 2017]. Furthermore, despite the presence of measurement noise, the results show that the Likert scale measurements were still fine enough to capture the effect of simulated joint stiffness on stiffness estimate in Experiment 1. Moreover, it allowed us to determine the effect

of velocity profile on the magnitude of subjects' stiffness estimates in Experiment 2. We conclude that, if the experimental procedure did affect subjects' stiffness estimating ability, the effect was minimal.

In Experiment 2, where subjects observed simulations of the same endpoint path produced from the same simulated joint stiffness, subjects estimated the stiffness of the veridical velocity profile to be less than that of the other three non-veridical velocity profiles. If subjects solely used path information to estimate stiffness, we would expect subjects to have chosen the exact same stiffness estimate. While one subject did exhibit this behavior (Fig. 4-11), that was not the trend for all subjects (Fig. 4-10). Upon comparing Figs 10 and 13, it is seen that subjects' stiffness estimates may be related to the acceleration or jerk of the simulated motion. Thus, participants possibly used one of these temporal metrics to estimate changes in stiffness, as the path-based metrics typically used were not available. Consistent with these results, it has been shown that minimizing mean-squared jerk yields the widely-observed $1/3$ power law relation between speed and curvature [Huh and Sejnowski, 2015, Viviani and Flash, 1995, Schaal and Sternad, 2001, Flash and Hogan, 1985], and moreover motions that follow the $1/3$ (veridical) power law are perceived to be more natural [Bidet-Ildei et al., 2006] and uniform [Viviani and Stucchi, 1992]. Here, subjects perceived these more 'natural' and 'uniform' movements to be less stiff. Throughout the literature (for review see [Zago et al., 2018]), it has often been shown that velocity encodes information about biological vs non-biological motion. Given that minimizing jerk leads to a velocity profile commonly observed in biological motion, subjects may have adopted this heuristic when estimating stiffness (Fig. 4-13). As a result, subjects rated biological velocity profiles to be less stiff than non-biological velocity profiles (Fig. 4-10). While participants had an overall cognitive bias to estimate smoother (i.e., less jerky), veridical motions as less stiff, this bias was likely not learned within the time course of the experiment. There was no statistically significant effect of trial nor a statistically significant interaction between trial and velocity profile on the stiffness ratings in Experiment 2. Ultimately, while this temporal information may have an effect on subjects' ability to estimate stiffness,

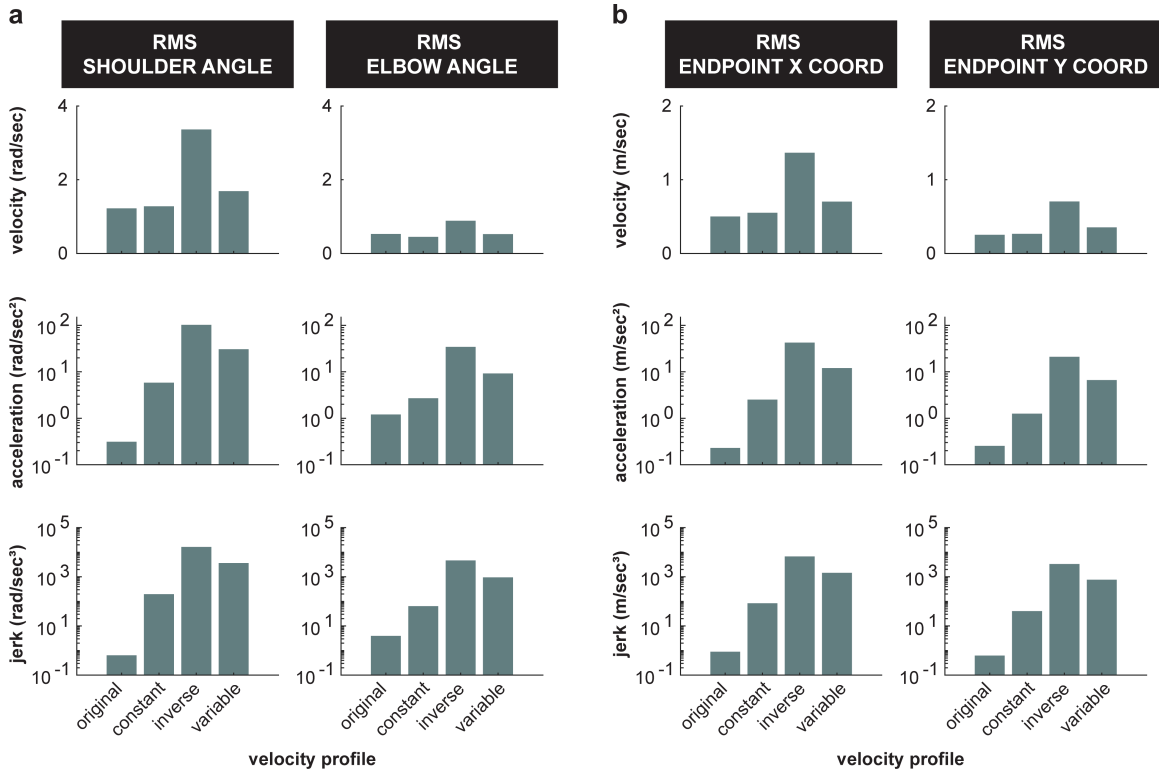


Figure 4-13: **RMS of temporal motion features in Experiment 2.** All 4 velocity profiles' RMS of velocity (rad/s), acceleration (rad/s^2), and jerk (rad/s^3) for 1 cycle of motion in (a) joint coordinates, and (b) endpoint coordinates.

it is still subordinate to path information.

Given that stiffness is the relationship between force and motion and no information related to force was provided, participants must have used some set of path-based heuristics to estimate changes in stiffness. It is important to emphasize, however, that we cannot claim what path-based heuristic(s) participants used as there are a number of possibilities (e.g., endpoint path oblongness, endpoint path area, endpoint path mean curvature, joint motion relative phase, shoulder range of motion, or elbow range of motion) (Fig. 4-14). Inherently, our findings are consistent with the idea of attribute substitution [Morewedge and Kahneman, 2010]. Attribute substitution occurs when people unconsciously make a judgement about a target attribute based on simpler and/or more accessible substitute attributes [Morewedge and Kahneman, 2010]. The process of attribute substitution reflects the assumptions or cognitive biases that an individual holds about the relation

between the target and substitute attribute(s). In our study, participants used some motion attribute(s) (substitute attributes) to estimate changes in stiffness (target attribute). Here, Experiment 1 (Fig. 4-5) demonstrates that overall participants held a “correct” bias or assumption about how changes in motion related to changes in joint stiffness; moreover, this bias was relatively consistent across individuals (Fig. 4-6). Importantly, this assumption of the relation between motion and stiffness was unaffected by manipulations of temporal information. However, Experiment 2 shows that participants still held a bias or assumption about how temporal attributes related to stiffness when differences in path were not present. Thus, we conclude with the novel finding that path-related attributes predominate over time-related attributes when humans estimate stiffness. Future work will investigate what specific path-related attributes are used.

Subjects’ self-reported estimating strategies further support the finding that the use of path information predominated over temporal information when subjects visually estimated changes in limb stiffness. To estimate arm stiffness, in Experiment 1, more subjects reported using motion features that contained path information than those that contained temporal information. Moreover, velocity profile did not significantly influence whether subjects reported using path or temporal information (Fig. 4-9A). It is possible that subjects only chose to ignore temporal information because it was not veridical. On the other hand, in Experiment 2, subjects were forced to use temporal information as the simulated endpoint paths were indistinguishable. Thus, contrary to Experiment 1 where few subjects ($\sim 23\%$) reported using temporal information (Fig. 4-9A), in Experiment 2 many subjects (80%) reported using temporal information (Fig. 4-12A). Interestingly, of the 30 subjects in Experiment 1, only 1 ($\sim 3\%$) reported using the words “jerk” or “smooth”; however, 4 of the 10 subjects (40%) in Experiment 2, reported using the words “jerk” or “smooth”. These results, together with Fig. 4-10 and 4-13, suggest it is likely that subjects did use temporal information to estimate stiffness when differences in path information were removed. Even so, the fact that velocity profile manipulations did not hinder subjects’ ability to estimate stiffness in Experiment 1 indicates that the influence of temporal




Figure 4-14: **Change in motion related features with simulated elbow stiffness** (adapted from [Huber et al., 2019]) (A) Endpoint path oblongness was defined as the ratio of the eigenvalues of the covariance matrix of the endpoint path in Cartesian coordinates. (B) Endpoint path area was defined as the area enclosed by the end-point path in Cartesian coordinates. (C) Endpoint path mean curvature was defined as the average inverse of the radius of curvature (see Methods). (D) Joint motion relative phase was defined as the relative phase between the elbow and shoulder motion and was calculated using cross-correlation analysis. (E) Shoulder range of motion (ROM). (F) Elbow range of motion (ROM).

information is at least subordinate to path information.

Many subjects reported strategies that did not clearly include any of the four motion features (Fig. 4-9B). Examples of such strategies reported interest in “flexibility of the object,” “rigidity of the motion,” and “how constrained the object’s motion appeared”. To avoid codifying subjects’ open-ended, written responses, we could have asked subjects to choose from a pre-determined list of motion features. While this may have simplified the coding of their estimating strategies, it would have

also increased the chances of response bias (e.g., choosing features from the list for the sake of it). Moreover, Table 4.1 is not an exhaustive list of the motion heuristics subjects could have used. Fig. 4-14 highlights other path-based motion heuristics that were related to the joint stiffnesses of the observed simulations. This study did not aim to determine which path-based heuristics subjects used to estimate stiffness. Rather, we examined subjects' use of path and temporal information; the features chosen in Table 4.1 helped us to do so. Furthermore, it is also important to be mindful that no matter how subjects reported their estimating strategy, subjects could have reported features that they did not actually use (or omitted reporting features they did use). An inability to clearly articulate what strategy they used could also signal that estimating stiffness using solely observation of motion is an implicit process rather than a conscious one. This could explain why the combination of reported motion features did not affect the stiffness estimating performance (Fig. 4-9C).

Nonetheless, the fact that subjects could successfully estimate changes in limb stiffness is remarkable, and even more so is the fact that this ability is robust to manipulations of temporal information (**Experiment 1**). The same limb motion can be generated with an infinite number of stiffness values, making it fundamentally impossible to unambiguously estimate stiffness from motion alone. This begs the question, how are humans able to estimate limb stiffness from just the visual observation of motion? The results from individual subjects indicated that they did not achieve task success by simply guessing how changes in motion related to changes in stiffness. If they did, we would expect the relation between their stiffness estimates and the simulated joint stiffness values to be significantly negative for some subjects. However, this was not the case. While 3 out of 40 subjects, in Experiment 1, did have a negative slope in the linear model fit to their estimating data, the corresponding R^2 values were very low (< 0.12) (Fig. 4-6). Excluding these subjects from the statistical analyses did not change the significance or non-significance of any results. Moreover, calculation of the Pearson's correlation coefficient for each of these three subjects confirmed that there was no significant, linear correlation between stiffness estimates and simulated joint stiffness ($p > 0.05$). If these subjects did guess, it

was a fruitless strategy. Instead, subjects must have drawn on prior knowledge to successfully perform the experimental task. Furthermore, that prior knowledge had a similar relation between limb stiffness and motion as the one resulting from the control policy used to drive the simulated limb.

Extensive evidence from behavioral and neuroimaging studies shows a strong coupling between action generation and perception processes in the human nervous system [Hecht et al., 2001, Casile and Giese, 2006, Calvo-Merino et al., 2006, Avenanti et al., 2013]. Thus, one possible explanation is that subjects may have used implicit knowledge of their own neuromotor control system to successfully estimate the changing stiffness of the simulated arm. In fact, a study of adaptation to a force field during reaching by a proprioceptively deafferented patient found that an internal representation of limb dynamics could be updated by visual information alone [Sarlegna et al., 2010]. As mentioned previously, the controller used to drive the simulated arm intentionally emulated aspects of human neuromotor control, based on the proposal that motor actions are built from dynamic primitives using the compositionality of mechanical impedance [Hogan, 2017, Hogan and Sternad, 2012]. As described in this theoretical framework, composing motor actions from motion primitives alone is insufficient. Inclusion of mechanical impedance, minimally defined as stiffness, is also required to control interaction or even the prospect of interaction. Studies of kinematically constrained movements corroborate this notion [Won and Hogan, 1995, Hermus et al., 2020b]. For instance, an oscillatory motion primitive combined with a mechanical impedance is a competent model of human interaction control during crank-turning (i.e., moving with a circular kinematic constraint) [Hermus et al., 2020b] and during human-robot physical interaction [West et al., 2022a]. This model offers a plausible explanation of how humans manage physical interaction with their high degree of skeletal redundancy [Hogan, 2017, Hogan and Sternad, 2012]. Outside the domain of human motor control, this computational model has proven to be an effective control strategy for kinematically redundant robots [Hermus et al., 2021]. The fact that subjects could estimate stiffness from observing motion alone further supports the role of impedance in the generation

of motor actions. Moreover, the fact that subjects could do this without explicit knowledge of the underlying controller suggests that the internal representation used by humans to interpret the motion of others is consistent with the form of the computational model proposed to describe the generation of motion. However, further research is needed to determine the degree of congruency between subjects' motion generation and perception needed to perform this experimental task. Nevertheless, the importance of the agreement between the simulated controller and subjects' motor perception should not be undervalued.

Whatever internal model or process was used to perform the task, regardless of its exact form, it must have incorporated stiffness. The novel contribution of the study presented here is that the representation of stiffness in this model is based primarily on path information. At first blush, this appears to contradict the extensive evidence indicating the important role of temporal patterns in biological motion perception. One feasible explanation for this apparent discrepancy is that here, subjects were asked to judge a mechanical property of a simulated arm, not anticipate or predict its kinematic behavior. Smoothness is a measure of predictability, and prior work has shown that maximizing smoothness (i.e., minimizing mean-squared jerk) yields a widely-observed power law relation between speed and curvature [Huh and Sejnowski, 2015, Viviani and Flash, 1995, Schaal and Sternad, 2001, Flash and Hogan, 1985]. Hence, manipulating the temporal pattern of an object's motion to eliminate the power law relation between speed and curvature decreases its predictability. While this would likely affect a human's ability to track or anticipate the object, explaining the results of studies such as Kandel et al. [Kandel et al., 2000b] and Maurice et al. [Maurice et al., 2018], it would not necessarily impair their judgement of mechanical properties, especially if such estimations are made based on the geometric features of movement.

Another plausible reason why distorting temporal patterns did not inhibit subjects' ability to estimate joint stiffness was the presence of multiple moving bodies. Here, subjects were presented with motion of a two-link arm, and more subjects reported using joint motion ($N = 16$) rather than endpoint motion ($N = 2$). In

prior studies where modulation of the speed-curvature power law relation reportedly affected motion perception, subjects were only shown a single moving body, or part of its trajectory, which is analogous to displaying only endpoint motion [Viviani and Stucchi, 1992, Kandel et al., 2000b, Dayan et al., 2007]. It is possible that spatial structure may be prioritized over temporal structure when multiple bodies are moving. For instance, using the point light animation paradigm with a walking figure, Hirai and Hiraki [Hirai and Hiraki, 2006] found that manipulation of spatial structure (i.e., randomizing the start point of a particular point light, while keeping its velocity vector the same) had a substantially greater effect on the brain’s response in the occipitotemporal regions, than manipulation of the temporal structure (i.e., scrambling the frames in the point light animation).

Ultimately, the fact that our results suggest that path information (i.e., geometry) is the predominant motion characteristic in humans’ internal representation of stiffness is consistent with the foundation of mechanics. Only with the postulates of geometry derived by Euclid could the likes of Newton, Lagrange, and Hamilton describe the motion of objects and deduce their own postulates regarding the forces that cause it (for review see [Dugas, 1988]). Geometry is fundamental to mechanics. Therefore, it is reasonable that internal models used to interpret dynamics during action perception are predominantly based on geometry.

4.7 Conclusions

To conclude, here we showed that humans can correctly infer changes in limb stiffness from nontrivial changes in multi-joint limb motion. This result was robust despite manipulations of the arm’s endpoint velocity profile. However, when path information was indistinguishable, veridical velocity profiles were perceived as less stiff than non-veridical velocity profiles. While other researchers have shown that manipulation of the velocity profile may hinder humans’ ability to anticipate kinematics, we show that it does not hinder humans’ ability to estimate stiffness. These observations suggest that stiffness, or more generally mechanical impedance, of the limb is encoded

in an internal model used to perform this task. Moreover, these results provide new insight into how humans interpret the motor actions of others and suggest that path, not trajectory, information is more important to humans when estimating stiffness from motion. This exploration of how humans extract latent features of neuromotor control from kinematics provides new insight into how humans interpret the motor actions and interactions of others.

Chapter 5

A Synergy-based Complexity Index to Promote the Study of Complex Manipulation

In Chapter 2, we established that synergies may simplify motor control during dexterous manipulation tasks. However, we highlighted that these synergies are task-specific, suggesting that we may have access to more synergies than degrees of freedom. Moreover, these synergies may aid in learning new behaviors, stored in a library to be utilized during specific tasks [Tessari et al., 2024]. This raises a daunting research question: How can we ensure that the tasks we study represent the entire spectrum of complex manipulation activities humans can perform with their hands? In Chapter 5, we aim to answer this question by proposing a new metric: the Variance-Accounted-For based Complexity Index (VAF-CI). This innovative metric can broaden the scope of future studies on hand manipulation, ensuring a comprehensive exploration of the hands' capabilities and related synergistic behavior.

This chapter is an adapted version of a paper that was submitted IEEE RAS EMBS 10th International Conference on Biomedical Robotics and Biomechatronics (BioRob 2024), “The Study of Dexterous Hand Manipulation: A Synergy-Based Complexity Index.” This work was done in collaboration with Federico Tessari and Neville Hogan.

5.1 Introduction

The study of robotic grippers, hands and exoskeletons is of great academic and industrial interest due to the need for an intuitive, multi-functional and seamless physical-interaction between machines and their external environment [Billard and Kragic, 2019]. Multiple studies, spanning more than a century, tried to tackle the design of dexterous end-effectors inspired by the superb manipulation abilities of the human hand [Schlesinger, 1919, Jacobsen et al., 1986, Bicchi and Kumar, 2000, Catalano et al., 2014, Piazza et al., 2019, Laffranchi et al., 2020]. However, manipulation proves to be a challenging problem not only from a design perspective, but also from a control standpoint [Billard and Kragic, 2019, Mason, 2018].

The “classical” robotic control techniques—based on rigid body dynamics—provide a sound starting point but struggle when faced with the strong non-linearities of contact-rich dynamics [Murray et al., 1994]. Two main approaches have emerged. On one hand, several studies tried to solve dexterous manipulation by exploiting learning-based methods [Li et al., 2018, Nagabandi et al., 2020, Kumar et al., 2016, Kroemer et al., 2021, Andrychowicz et al., 2020]. A complementary approach studies how humans exploit their end-effectors—i.e. their hands—to perform object manipulation. The extraordinary manipulation abilities of humans can provide both inspiration for and intuition about how to achieve dexterous robotic manipulation.

One aspect that has facilitated the study of human manipulation is the investigation of hand synergies [Santello et al., 1998, Mason et al., 2001, Todorov and Ghahramani, 2004]. A synergy can be defined as a coordinated action of an ensemble of features. Such features can be described at different levels such as muscular [D’avella and Tresch, 2001, Ting and McKay, 2007, Weiss and Flanders, 2004] (typically measured through EMGs) or kinematic [Santello et al., 1998, Mason et al., 2001, Todorov and Ghahramani, 2004, Weiss and Flanders, 2004, Santello et al., 2002] (typically recorded through joint trajectories or limb motions). A leading theory in human neuromotor control supports the hypothesis that humans exploit synergies

to control their dimensionally complex neuromusculoskeletal system [Santello et al., 2016]. In other words, synergies are a lower dimensional operating space that simplifies the control of a larger set of controllable features.

Human hand synergies have been investigated to simplify robot mechanical design and control [He and Ciocarlie, 2022, Ciocarlie and Allen, 2009]. Most studies investigated the role of synergies in simple and common everyday tasks, such as grasping and tool handling [Santello et al., 1998, Mason et al., 2001]. However, the human hand is capable of a wide array of tasks of various degree of complexity (e.g., from grasping cylindrical objects such as a water bottle to a complex manipulation task such as playing the piano).

With this in mind, one large question remains unanswered: “How can we ensure that the tasks analyzed represent the whole spectrum of complex manipulation activities that humans can perform with their hands?”. The answer to this question may lead to a more accurate description of how humans exploit synergies to control their hands during dexterous manipulation. Such comprehension could be leveraged to improve both the design of exoskeletons, prostheses, and general robotic end-effectors as well as the execution of existing synergy-based control strategies.

In this study, we present a new metric to quantify the complexity of hand manipulation tasks. This metric can ensure that the study of hand manipulation incorporates tasks that represent the breadth of the hands’ dexterity. To validate this metric, we analyzed three different tasks involving various degrees of hand dexterity: (1) a simple human manipulation task i.e., reaching to grasp a tool, (2) a more advanced occupational task incorporating tool-use, and (3) a widely recognized advanced manipulation exercise: piano playing.

For each of these tasks, the analysis focused on kinematic synergies of the hand and, in particular, quantification of their complexity based on an index extracted from the cumulative variance-accounted-for (VAF) by all synergies. Here, we introduce the VAF-based complexity index (VAF-CI). The efficacy of this index is statistically validated and its robustness is investigated through a sensitivity analysis. Furthermore, a supervised machine learning method, based on k-means clustering, is

developed as an unbiased verification of the experimental methodology. The results of this study can be used to ensure that future studies of hand manipulation probe a wider range of the hands’ capabilities, and related synergistic behavior.

5.2 Methods

5.2.1 Singular Value Decomposition

Given a data matrix $X \in \mathbb{R}^{n \times m}$, where n denotes the number of observations (e.g., the time evolution of a degree of freedom) and m denotes the number of features (e.g., the number of analyzed DoFs) for a given task, synergies can be extracted using the Singular Value Decomposition (SVD) algorithm:

$$X = U \cdot S \cdot V^T \quad (5.1)$$

where $U \in \mathbb{R}^{n \times n}$ is an orthonormal matrix whose columns denote the temporal evolution of a given synergy, $S \in \mathbb{R}^{n \times m}$ is diagonal matrix of singular values which give an estimate of the VAF in a given synergy, and $V \in \mathbb{R}^{m \times m}$ is an orthonormal matrix whose columns represent a synergy. For a geometrical interpretation of this, see [West et al., 2023].

The VAF enables us to understand how each extracted synergy represents the variability in the original data set. The VAF of a given (i.e., the i^{th}) synergy can be computed using:

$$VAF(i) = \frac{\sigma_i^2}{\sum \sigma_i^2} \quad for \ i = 1, \dots, m \quad (5.2)$$

where σ_i is the singular value on the i^{th} diagonal element of S . Throughout this paper, the VAF is reported as a decimal ranging from 0 to 1. Intuitively, a synergy with a larger VAF will account for a greater portion of the data variability.

5.2.2 Variance-Accounted-For based Complexity Index

In this work we propose a VAF-based complexity index (VAF-CI). Given a certain experimental task, the VAF-CI can be computed using the following equation based on the trapezoidal integral rule:

$$VAF-CI = 1 - \frac{1}{N_{DoF}} \sum_{i=1}^{N_{DoF}-1} \frac{1}{2} \left(cVAF(i) + cVAF(i+1) \right) \quad (5.3)$$

where $cVAF(i) = \sum^i VAF(i)$ is the cumulative variance-accounted-for at the i -th synergy and N_{DoF} is the total number of degrees-of-freedom. Importantly, in the SVD, the total number of synergies coincides with the total number of degrees-of-freedom. Note that the VAF-CI has been formulated as a metric that can range from 0 to 1.

Mathematically, the VAF-CI is the normalized area above the cumulative VAF curve (See Fig. 5-3). As this area increases, so does the VAF-CI; for this area to increase, the cumulative VAF-curve must decrease. Conceptually, the VAF-based complexity index increases when more synergies are needed to describe the variance in the data. However, from a physical standpoint, this metric has an upper limit of 0.5. If each degree-of-freedom acts independently, then the contribution of each synergy will be weighted equally. In turn, the cumulative variance-accounted-for will approach a straight line whose slope is $\frac{1}{N_{DoF}}$. The normalized area above this curve will approach a maximum at 0.5.

In this work, we tested the hypothesis that as a task increases in complexity, more individual DoFs are needed to perform the task, and this could be uniquely quantified by the VAF-CI. It is important to emphasize that, here, we define complexity by the necessity of individuated motion. A task that enables joint coordination is considered less complex, while a task that requires joint individuation is considered more complex. In sum, the proposed VAF-based complexity index can be used as an intuitive measure of task complexity without the need to choose the number of important synergies, e.g. based on an arbitrary level of variance-accounted-for.

5.2.3 Experimental Setup

To test the VAF-CI, we needed to obtain kinematic hand data for a wide range of tasks with widely recognized different levels of complexity. Specifically we considered the three following manipulation tasks: (1) *reach-to-grasp*, as the most basic manipulation task – every able-bodied human is capable of it, (2) *tool-use*, as a more advanced manipulation – most but not all able-bodied humans are capable of it, and (3) *piano playing*, as one of the most advanced manipulation tasks – only a small percentage of specifically-trained able-bodied humans are capable of it. We want to emphasize that piano playing is a highly dynamic task that requires temporal coordination, joint individuation, and the application of varying forces in a goal-oriented, time specific manner.

Wire Harnessing

The initial two tasks involving manipulation (*reach-to-grasp* and *tool-use*) were conducted within the framework of wire-harnessing. Figure 5-1 shows the mock-up electrical cabinet used for the wire-harnessing experiment.




Figure 5-1: **Wire harnessing experimental setup**: top image shows the use of a screw-driver (Step 3), bottom image shows the threading of the wire harness (Step 4).

The study enlisted seven (three women and four men) right-handed participants, ages ranging from 18 to 28 years, all of whom had no documented history of neurological or musculoskeletal issues. All subjects received a comprehensive briefing on the experimental procedures and signed a consent form approved by MIT's Institutional Review Board.

Each participant executed two distinct tasks. The first task, denoted as the '*reach-to-grasp* experiment', entailed reaching out, grasping, and lifting various tools. Specifically, participants were instructed to reach for and grip items such as scissors, zip ties, a screwdriver, a wire harness with its branches secured, and a wire harness with its branches unsecured, as if they intended to use them. Each subject repeated this action four times for each object. The concatenation of each of these trials, for a given subject, produced one data matrix, X . The second task was structured to replicate wire-harness assembly in a manufacturing setting. Participants were

tasked with utilizing the tools from the reach-to-grasp experiment to install a wire harness onto a simulated electrical cabinet (Fig. 5-1). This task, termed the ‘*tool-use experiment*’, was divided into five distinct steps. At each stage, participants received verbal instructions and were provided with a booklet containing written directions and a depiction of the completed step. One subject’s completion of a given step produced one data matrix X . The steps were as follows: (1) Attach zip ties to the branched ends of the wire harness at three specified locations, (2) employ the scissors to trim excess tips of the zip ties, (3) utilize U-brackets, screws, and a screwdriver to fasten the wire harness securely to the top of the mock electrical cabinet, (4) thread the wire harness through hooks, (5) connect the wire harness into the socket.

Throughout the task performance, we recorded hand postures using a CyberGlove (CyberGlove; Virtual Technologies, Palo Alto, CA), a glove equipped with embedded sensors designed to capture joint kinematics. We recorded 20 joint angles. Specifically, measurements included flexion of the distal interphalangeal (DIP), proximal interphalangeal (PIP), and metacarpophalangeal (MCP) joints of the four fingers. Additionally, abduction (ABD) of the four fingers at the MCP joints was recorded. For the thumb, measurements included flexion at the MCP and interphalangeal (IP) joints, abduction (ABD) at the carpometacarpal joint, and rotation (ROT) around an axis intersecting the trapeziometacarpal joint. During both experiments, participants wore a CyberGlove on each hand, although only data from the right hand are presented in this report. In all instances, the CyberGlove collected samples at approximately 200 Hz, with a nominal angular resolution of less than 0.1 degrees.

In total, we obtained 10 (5 reach-to-grasp tasks + 5 manipulation tasks) sets of data, $X \in \mathbb{R}^{n \times 20}$ (where $m = 20$ joint DoFs), for each of the 7 subjects. For greater experimental details, see Chapter 2.

The piano playing study enlisted five right-handed participants (three men and two women) with no documented history of neurological or musculoskeletal disorders. Their ages ranged from 19 to 36 years and they had an average piano experience of 20.0 ± 8.3 years. All subjects were right-handed, received a comprehensive briefing

on the experimental procedures, and signed a consent form approved by MIT's Institutional Review Board.

Each participant was asked to play any piece they could perform from a series of 9 piano pieces from the classical repertoire, ranging from beginner to advanced level. No subject could play the advanced pieces; thus, they are not reported here: (Beginner) J. S. Bach, *Prelude in C Major, BWV 846*; F. Chopin, *Prelude Op. 28 No. 5*; E. Satie, *Gnossienne No. 1*, (Intermediate) W. A. Mozart, *Sonata No. 8 in A minor, K310, I movement*; L. v. Beethoven, *Sonata No. 23 in F minor "Appassionata", Op. 57, III movement*; C. Debussy, *12 Etudes, No. 5, Pour Les Octaves*.

Afterwards, subjects were invited to play any additional piano piece in their personal repertoire. Each subject was instructed to perform the piece in a unique session without interruptions. This was done to prioritize continuity of motion over accuracy and precision of the piece execution. On average each subject played 6.5 ± 1.9 pieces.

During the execution of every piece subjects hand movements were recorded using two cameras (Sony RX0 II) mounted above the two extremities of the keyboard (Roland FP-10) as shown in Figure 5-2. The cameras' recordings were analogically synchronized using two dedicated Sony Camera Control Boxes connected to an Ethernet router. The Sony related acquisition software was used and data were saved at a sampling frequency of 60 *Hz*.

Piano Playing




Figure 5-2: Piano playing experimental setup.

The collected data were then post-processed using the MediaPipe Hands marker-less motion tracking algorithm. MediaPipe Hands is a high prediction quality and lightweight machine learning model. It consists of two modules: (i) a palm detector that places an oriented hand bounding box and (ii) a hand landmark model that operates on the cropped bounding box [Zhang et al., 2020]. This base model detects 21 key-points: one on the wrist and four additional points per finger. A custom algorithm, developed in [Wang, 2023], was used to transform the extracted key-points into 20 joint rotations – the same as those collected by the CyberGlove described above.

The left camera (Fig. 5-2) was used to extract the right hand middle, ring and little fingers' motion while the right camera took care of the right hand thumb and index

fingers. The camera configuration and finger selection were heuristically determined. Only right hand data were analyzed for this report.

5.2.4 Validation of the Variance-Accounted-For based Complexity Index

Given three sets of tasks with widely accepted different levels of manipulation complexity, we tested the proposed VAF-CI. First the data were centered, in accordance with [West et al., 2023]. Then, after SVD (Eq. 5.1), the cumulative VAF (Eq. 5.2) and subsequent VAF-CI (Eq. 5.3) were computed. Finally, we compared the VAF-CIs of the reach-to-grasp, tool-use, and piano playing trials using a 3 (task) x 1 ANOVA.

Comparison to Significant Number of Synergies

To emphasize the robustness of the novel VAF-CI, we compared our metric to the current standard for estimating the number of significant synergies. The literature has used a plethora of methods to define the number of significant synergies [Santello et al., 1998, Todorov and Ghahramani, 2004]. Typically, researchers will choose the significant number of synergies using an arbitrarily chosen amount of cumulative VAF. For instance, some have defined the number of significant synergies as the number of synergies required to achieve at least 90% cumulative VAF and where inclusion of another subsequent synergy did not add an additional 5% cumulative VAF [Todorov and Ghahramani, 2004]. Here, we computed the number of significant synergies for each task considering at least 90%, 80%, or 70% cumulative VAF and where inclusion of another subsequent synergy did not add an additional 5% cumulative VAF. For brevity, each of these three metrics will be referred to as 90 & 5, 80 & 5, and 70 & 5 respectively. For each of the 3 thresholds used to define the significant number of synergies, we performed three one-way ANOVAs (one per each task) to determine if there was an effect of manipulation task on the number of significant synergies.

A Metric for Unsupervised Learning

As an additional test, we submitted the computed VAF-CIs to an unsupervised learning algorithm, namely k-means clustering [Lloyd, 1982], to test whether this metric can be used to classify various manipulation tasks. The k-means clustering algorithm was implemented with the *kmeans* function in MATLAB. Moreover, it was initialized with k centroids whose values were randomly sampled from the input VAF-CI values. To ensure that a local minimum was reached, the algorithm was updated with online updates – each sample was individually reassigned to a new cluster if it reduced the sum of point-to-centroid distances. Finally, in order to find a global minimum this algorithm was repeated 5 times with different randomly-sampled initializations. The repetition with the lowest sum of point-to-centroid distances is reported.

To identify the number of clusters, we used the elbow method – identifying where adding an additional number of clusters no longer significantly decreases the sum of point-to-centroid-distances. For the case of 3 chosen clusters, we also report the confusion matrix comparing the true and predicted tasks. Moreover, we report the algorithm’s accuracy – defined as the number of true positives (correct classifications) for each group divided by the total number of sample tasks.

Robustness to Data Pre-processing

One might argue that the time needed to complete the various measured tasks could significantly impact the variance in the data, thereby introducing bias to the variance-accounted-for based complexity index. To address this concern, we applied two pre-processing techniques before Singular Value Decomposition. Specifically, we downsampled the data using either a specific frequency (60Hz) or a set number of bins (800). These values were chosen to be close to those of the trial with the lowest sampling frequency or number of bins.

As done above, we conducted SVD (Eq. 5.1) on the down-sampled data and calculated the cumulative VAF (Eq. 5.2) and subsequent VAF-CI (Eq. 5.3). Finally,

we compared the VAF-CIs of the reach-to-grasp, tool-use, and piano playing trials using a 3 (task) x 1 ANOVA. Also, we report the accuracy of the unsupervised learning algorithm on these data sets when 3 clusters were identified.

5.3 Results

5.3.1 Validation of the Variance-Accounted-For based Complexity Index

With kinematic hand data from the three sets of tasks (reach-to-grasp, tool-use, and piano playing), we used Eqs. 5.1 and 5.2 to compute the variance-accounted-for by each synergy. The average and standard deviation across trials are shown in Fig. 5-3. Additionally, the cumulative VAFs are shown averaged across trials for each set of tasks. The cumulative VAF for a given synergy is generally lower in tool-use than in reach-to-grasp, and lower in piano than in tool-use.

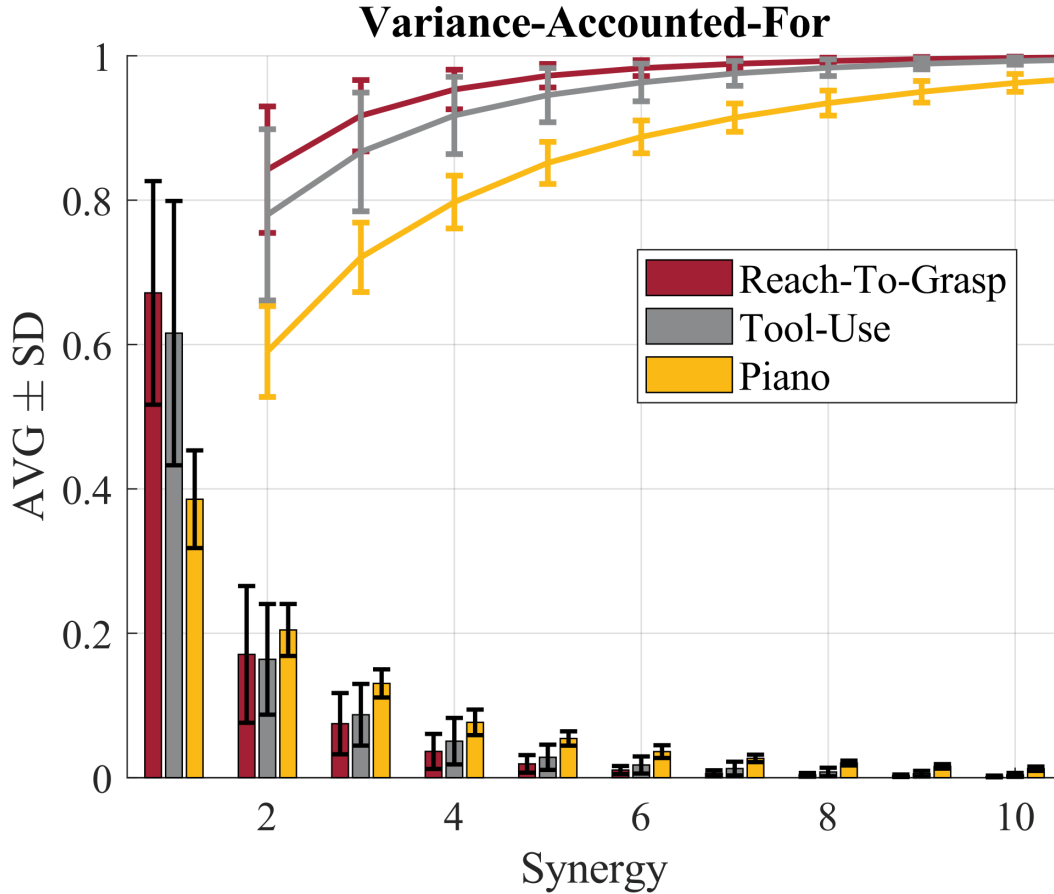


Figure 5-3: **VAF of each synergy in the reach-to-grasp, tool-use, and piano playing experiments** averaged across subject. The cumulative sum of the VAF is denoted by the lines. Errorbars are $\pm 1SD$. To save space, only the first 10 of 20 synergies are shown.

With the cumulative VAF, our proposed VAF-based complexity index was computed using Eq. 5.3 for each trial. The averages and standard deviations across tasks are reported in Fig. 5-4. A one-way ANOVA revealed a significant effect of task on VAF-CI ($F_{2,101} = 125.75$, $p = 3.86e - 20$). Post-hoc t-tests (Bonferroni corrected $\alpha = 0.05/3 = 0.0167$) revealed that the tool-use VAF-CI was significantly greater than that of reach-to-grasp ($p = 0.0061$) while the piano playing VAF-CI was significantly greater than that of tool-use ($p = 1.27e - 20$) and reach-to-grasp ($p = 2.26e - 27$).

	90 & 5	80 & 5	70 & 5
Reach-to-Grasp	3.03 ± 1.01	2.83 ± 0.92	2.83 ± 0.92
Tool-Use	3.94 ± 1.35	3.34 ± 0.94	3.34 ± 0.94
Piano	6.97 ± 0.87	4.74 ± 0.71	4.68 ± 0.68

Table 5.1: **Number of Significant Synergies** for the reach-to-grasp, tool-use, and piano playing experiments to achieve at least 90%, 80%, or 70% cumulative VAF and where inclusion of another subsequent synergy did not add an additional 5% VAF. The averages and standard deviations are reported.




Figure 5-4: **VAF-CI in the reach-to-grasp, tool-use, and piano playing experiments** averaged across subject. Errorbars are $\pm 1SD$. * represents $p \leq 0.05$, ** represents $p \leq 0.01$, and *** represents $p \leq 0.001$

5.3.2 Comparison to Significant Number of Synergies

Table 5.1 reports the number of significant synergies in the reach-to-grasp, tool-use, and piano-playing experiments for the three aforementioned thresholds of cumulative VAF.

Moreover, for each of the 3 metrics used to define the significant number of synergies, we performed a one-way 3 (tasks) x 1 ANOVA to determine if there was an effect of manipulation task on the number of significant synergies. The one-way ANOVA on the 90 & 5 significant synergies returned a significant effect of task ($F_{2,101} = 121.27$, $p = 1.42e - 27$). Post-hoc t-tests (Bonferroni corrected $\alpha = 0.05/3 = 0.0167$) revealed that the tool-use had significantly more synergies than the reach-to-grasp ($p = 0.0021$) while, piano had significantly more synergies than tool-use and reach-to-grasp ($p = 1.48e - 19$ and $4.74e - 27$ respectively).

In comparison, the ANOVA on the 80 & 5 and 70 & 5 significant synergies demonstrated a significant effect of task in both cases ($F_{2,101} = 44.76$, $p = 1.20e - 14$ and $F_{2,101} = 42.51$, $p = 4.04e - 14$ respectively). However, post-hoc t-tests (Bonferroni corrected $\alpha = 0.05/3 = 0.0167$) revealed that while piano playing exhibited a greater number of significant synergies than tool-use ($p = 3.78e - 09$ and $1.13e - 08$, respectively) and reach-to-grasp ($p = 1.81e - 14$ and $5.35e - 14$, respectively) tool-use did not exhibit more significant synergies than reach-to-grasp in both cases ($p = 0.038$ and 0.036 , respectively for 80 & 5 and 70 & 5). This demonstrated that arbitrarily choosing a threshold to define which synergies are significant can lead to different results, both in the number of synergies as well as how much they discriminate between manipulation tasks. This method of measuring complexity could lead to arbitrarily biased results.

5.3.3 A Metric for Unsupervised Learning

To demonstrate that VAF-CI can be used to classify various manipulation tasks, we submitted these values to an unsupervised learning algorithm, k-means clustering. Fig. 5-5a shows the sum of point-to-centroid distances as a function of the number

of clusters. Intuitively, as this value increases to the exact number of samples the sum of point-to-centroid distances will tend to zero. Importantly, Fig. 5-5a shows that after ~ 3 clusters, adding an additional cluster did not greatly reduce the sum of point-to-centroid distances; there is an “elbow” around 3 clusters. That is consistent with the three different manipulation tasks investigated in this study.

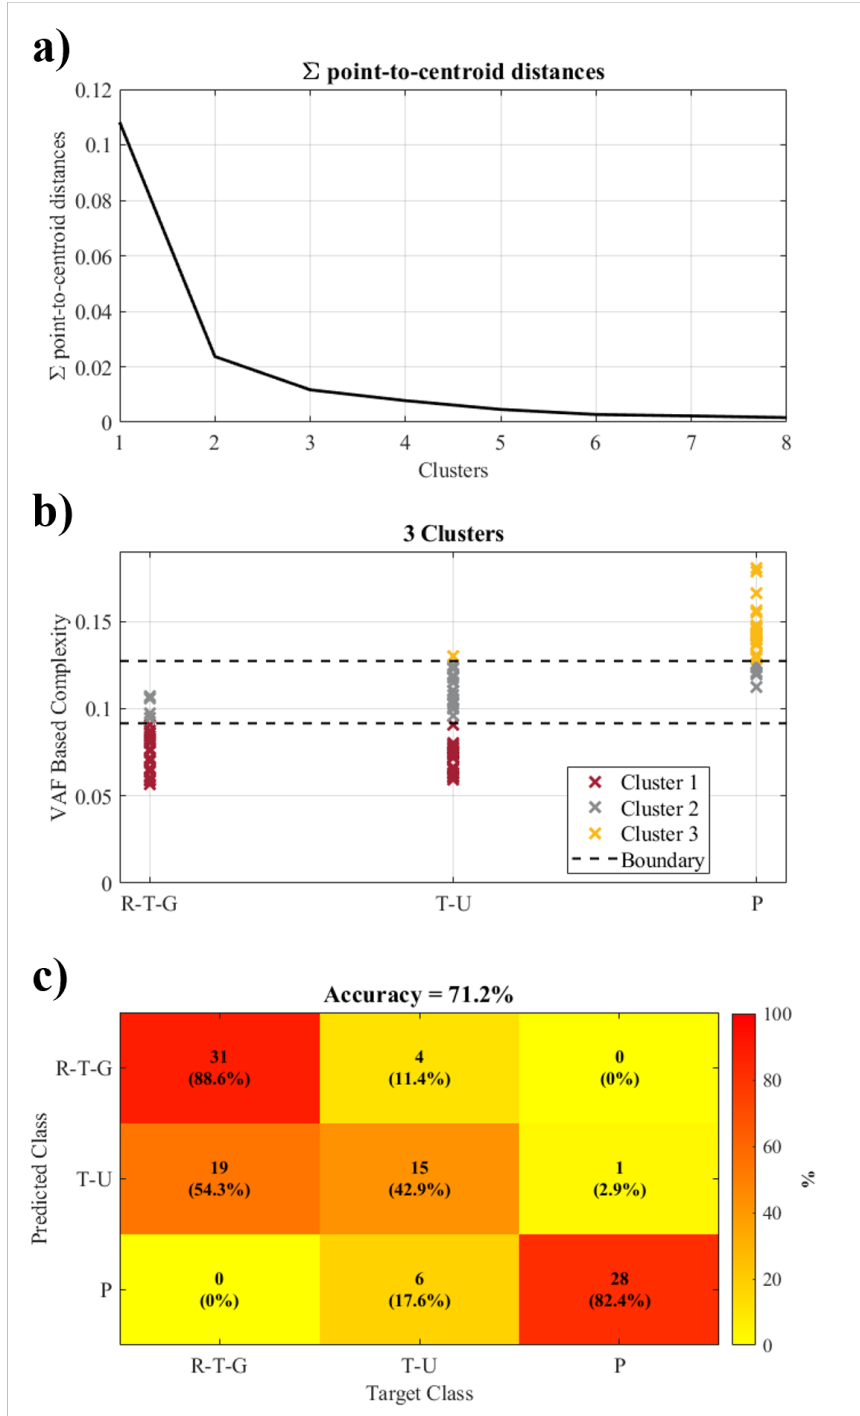


Figure 5-5: **K-means clustering algorithm results** (a) The sum of point-to-centroid distances as a function of the number of clusters. At ~ 3 clusters, adding another cluster did not greatly reduce the sum of point-to-centroid distances. (b) The identified clusters of the algorithm for 3 chosen clusters and the associated (c) confusion matrix. X's denoted the VAF-CI for each of the samples in the 3 tasks: reach-to-grasp ('R-T-G'), tool-use ('T-U'), and piano playing ('P'). Different colors denote the various identified clusters while the dashed lines denote the boundaries between the identified clusters.

Fig. 5-5b shows the results of the k-means clustering algorithm for three identified clusters. Fig. 5-5c shows the confusion matrix between the identified clusters and the originally considered manipulation tasks. The algorithm correctly identified the task using the VAF-CI with an accuracy of 71.2%. Notably, it classified the reach-to-grasp task with an accuracy of 88.6% and the piano playing task with an accuracy of 82.4%. The tool-use task, whose VAF-CI was generally less than that of the piano-playing task and generally greater than that of the reach-to-grasp task, had a classification accuracy of 42.9%; the tool-use tasks were often classified as a reach-to-grasp task (54.3%).

Fig. 5-6a shows the VAF-CI of the individual tasks that comprise the reach-to-grasp (R-T-G) and tool-use (T-U) experiments. Additionally, the boundaries between clusters identified in Fig. 5-5b are displayed. It is seen that some of the simpler reach-to-grasp tasks, such as plugging in (Step 5) and routing (Step 4) the wire harness, have lower VAF-CIs. On the other hand, the tool use tasks involving more dexterity - using a zip tie and a screwdriver - were classified in the middle cluster.

Fig. 5-6b displays the VAF-CI of the individual piano pieces that were played as part of the piano experiment. Three pieces of generally accepted levels of easy (J.S. Bach, *Prelude in C major, BWV 846* [Bach, 2011]), medium (W.A. Mozart, *Piano Sonata No. 8 in A minor, K. 310*), and advanced (F. Chopin, *Étude Op. 25, No. 11* [Chopin and Palmer, 1995]) difficulty are highlighted. VAF-CI correctly orders these pieces based on complexity, demonstrating that even within a particular manipulation task (i.e., piano playing), VAF-CI can capture the varying complexity.

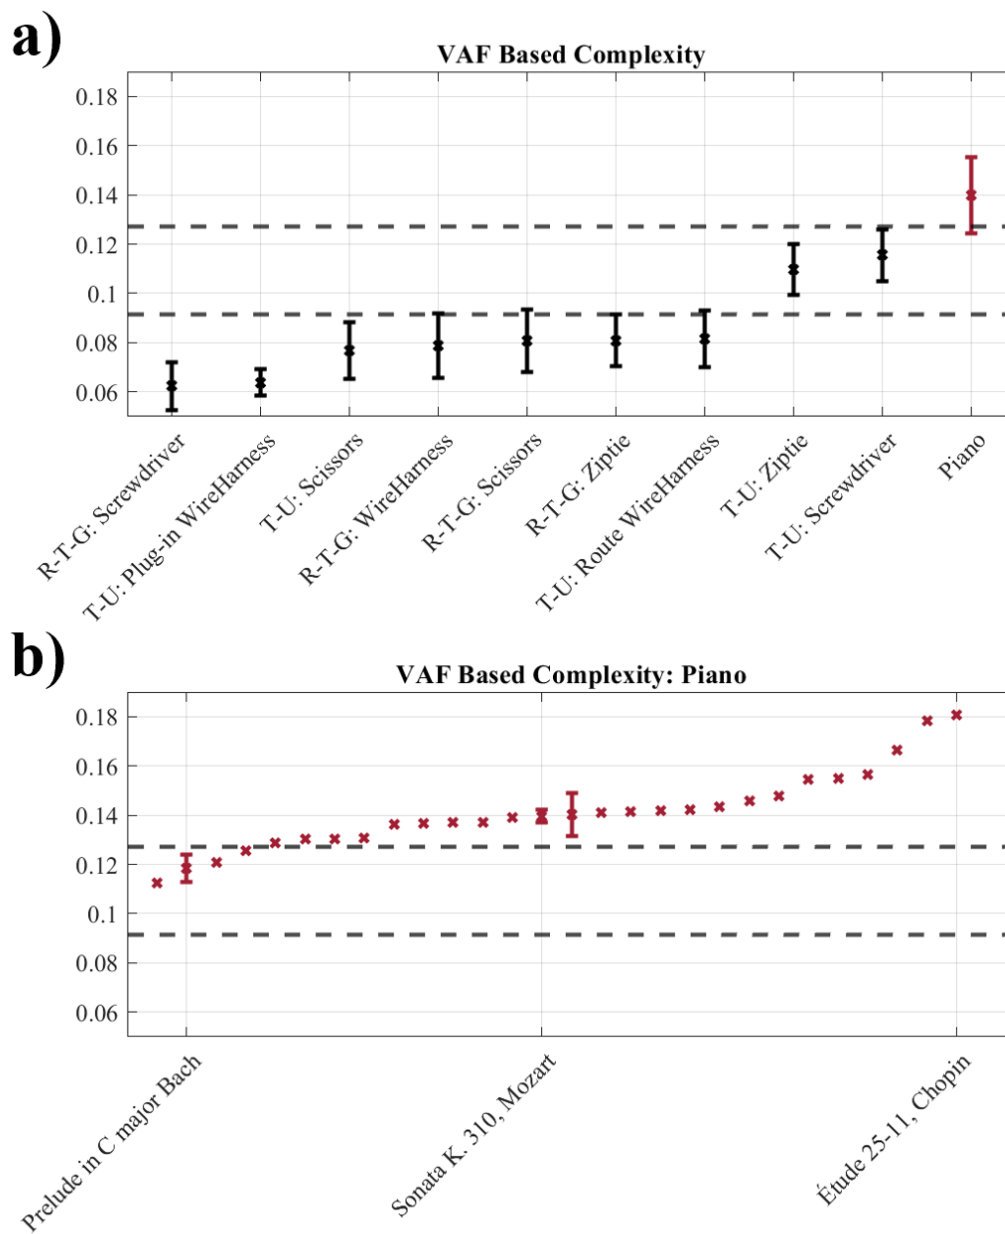


Figure 5-6: **Average VAF-CI of individual tasks.** (a) The VAF-CI for the distinct tasks in the reach-to-grasp (R-T-G) experiment and tool-use (T-U) experiment are shown in increasing order in black, while the piano piece VAF-CI is shown in red. (b) The VAF-CI of individual pieces played in the ‘*piano*’ experiment are shown in increasing order in red. Pieces of generally accepted levels of easy (J.S. Bach, *Prelude in C major, BWV 846*), medium (W.A. Mozart, *Piano Sonata No. 8 in A minor, K. 310*), and advanced (F. Chopin, *Étude Op. 25, No. 11*) difficulty are highlighted. Error bars are $\pm 1SD$. Note, some piano pieces do not have error bars as they were only played by one subject. Dashed lines denote the boundaries between the identified clusters as in Fig. 5-5b.

5.3.4 Robustness to Data Pre-processing

As mentioned in Section 5.2.4, we sought to assess the metric’s robustness to data length and pre-processing. Two pre-processing techniques were employed: downsampling to (1) 800 bins and (2) 60 Hz. In the 800 bins data the VAF-CI was 0.076 ± 0.013 , 0.089 ± 0.022 , and 0.12 ± 0.011 for the reach-to-grasp, tool-use, and piano playing data respectively. Similarly, in the 60 Hz data the VAF-CI was 0.076 ± 0.013 , 0.089 ± 0.022 , and 0.12 ± 0.012 for the reach-to-grasp, tool-use, and piano playing data respectively. Moreover, in both data sets a one-way ANOVA revealed a significant effect of task on VAF-CI (800 bins: $F_{2,101} = 69.22$, $p = 1.17e - 19$ and 60 Hz: $F_{2,101} = 69.63$, $p = 9.83e - 20$). In both data sets, post-hoc t-tests (Bonferroni corrected $\alpha = 0.05/3 = 0.0167$) revealed that the tool-use VAF-CI was significantly greater than that of reach-to-grasp (800 bins: $p = 0.0030$ and 60 Hz: $p = 0.0035$) while the piano playing VAF-CI was significantly greater than that of tool-use (800 bins: $p = 3.64e - 12$ 60 Hz: $p = 2.63e - 12$) and reach-to-grasp (800 bins: $p = 1.40e - 19$ 60 Hz: $p = 1.25e - 19$). Additionally the k-means clustering algorithm, with three identified clusters, had an accuracy of 61.5% and 60.6% for the 800 bins data and 60 Hz data, respectively. Together, these results demonstrate that VAF-CI is able to distinguish between and classify the complexity of the 3 chosen tasks despite pre-processing methods that vary the length of the data array.

5.4 Discussion

In this study, we proposed and tested a novel metric for identifying complexity of a manipulation task, termed variance-accounted-for based complexity index (VAF-CI). The VAF-CI is obtained from computing the area above (Eq. 5.3) the cumulative VAF-curve (Eq. 5.2) obtained after Singular Value Decomposition of the kinematic hand joint data (Eq. 5.1). We tested the VAF-CI with respect to three different manipulation tasks of distinct levels of complexity: reach-to-grasp, tool-use, and piano playing. We found a significant effect of manipulation task on VAF-CI (Fig. 5-3 and 5-4), validating the index’s ability to measure task complexity.

Importantly, the VAF-CI is a non-arbitrary metric to quantify task complexity. Previously, researchers classified complexity using the number of synergies needed to achieve a certain cumulative VAF [Santello et al., 1998, Todorov and Ghahramani, 2004]. As demonstrated by Table 5.1, this method is sensitive to the user-defined cumulative VAF cut-off. In contrast, VAF-CI does not require such choices.

Additionally, we demonstrated that VAF-CI, can be used by an unsupervised learning algorithm to identify different manipulation tasks (Fig. 5-5b and 5-5c). K-means clustering was able to accurately identify the 3 different manipulation tasks with an accuracy of 71.2%. Insight can be gleaned from observing how each tool of the wire harnessing task was classified (Fig. 5-6a). The algorithm correctly classified tool-use tasks with an accuracy of only 42.9%. Further inspection revealed that for all 7 subjects, the algorithm classified Step 5, a task involving reaching, grasping and plugging-in an electrical socket, as reach-to-grasp. Given this simple action, this result seems reasonable. Moreover, the algorithm classified scissors tool-use as reach-to-grasp in 6 of the 7 subjects. This may occur because the motion of the joints in the hand is physically constrained by the scissors, leading to motion composed of fewer synergies. Screwdriver reach-to-grasp was always correctly classified, while screwdriver tool-use was classified as tool-use in 6 of the 7 subjects and piano playing in the other. This emphasizes the clear difference in the manipulation complexity required to grasp a screwdriver versus to use it. Notably, this difference was captured by the VAF-CI. Moreover, three of the six piano pieces classified as tool-use were when subjects played J. S. Bach, *Prelude in C Major, BWV 846* – a piece of beginner level [Bach, 2011]. Additionally, Fig. 5-6b demonstrates that even within piano playing there was a spectrum of manipulation complexity and VAF-CI captured this; notoriously complex pieces, such as F. Chopin, *Étude Op. 25, No. 11* [Chopin and Palmer, 1995], were identified as the most complex. Importantly, k-means clustering correctly identified the extremity of piano playing with an accuracy of 82.4% and reach-to-grasp with an accuracy of 88.6%; it never classified piano playing as reach-to-grasp nor reach-to-grasp as piano playing. This emphasizes the robustness of our metric for identifying manipulation tasks without any prior knowledge.

5.4.1 Limitations

It is important to emphasize here that manipulation complexity can be composed of many different factors – such as tactile, visual and proprioceptive complexity, and force, torque or stiffness modulation. A limitation of the variance-accounted-for based complexity index is that it only takes into account one factor: joint coordination, or lack-there-of. Thus, in this study, when referring to complexity, VAF-CI defines complexity as increasing solely with the required individuated joint motion. Nonetheless, it is still a novel and continuous metric that can non-arbitrarily quantify the complexity of different tasks. As such, it is superior to the common practice of arbitrarily identifying an integer number of significant synergies.

One might interpret VAF-CI as a measure of variability. With this interpretation, greater variability in the task will inherently result in a higher VAF-CI. This means that when a task can be performed in multiple ways, making it generally ‘easier’ to perform, the complexity index increases. Conversely, if a task requires very precise actions with minimal variability, the index will likely decrease. However, this interpretation only holds if you set up your data matrix prior to SVD in a manner such that you have observed a subject do a repetitive task. In this case, if a subject uses different strategies each time during a repetitive task, the VAF-CI will likely yield higher values. This highlights that VAF-CI should be used to differentiate between the complexities of different non-repetitive tasks. Nevertheless, this insight presents an intriguing avenue for future research on the VAF-CI metric: determining its efficacy in quantifying user skill for a particular, repetitive manipulation task. We hypothesize that tasks requiring greater skill and precision would yield lower VAF-CI values among experts compared to novices.

5.4.2 Implications

A major challenge in manipulation of dexterous robotic hands is the “curse of dimensionality” [Bellman, 1958]. Reducing the control space to a subset of synergies is one solution to this problem [Santello et al., 2016]. This has led to the design

of many synergy-based dexterous robot hands [Catalano et al., 2014, Laffranchi et al., 2020, Brown and Asada, 2007, Varol et al., 2014, Matrone et al., 2012, Furui et al., 2019]. However, these designs often ignore the fact that synergies may stem from biomechanics as opposed to ease of neuromotor control. In fact, Todorov et al. [Todorov and Ghahramani, 2004] found that—when asked to individually actuate each joint in the hand—subjects required significantly more synergies (~ 8.7), representing an upper limit to finger individuation and thus manipulation complexity. Contrarily, grasping studies generally report 2-3 significant synergies [Santello et al., 1998, Mason et al., 2001, Weiss and Flanders, 2004, Santello et al., 2002], representing the lower end of finger individuation and thus manipulation complexity. The study presented here introduced a synergy-based analysis of piano playing, a complex manipulation task that requires years of training. Future study of piano playing could provide an upper limit of manipulation complexity during a functional task. Moreover, the distinct difference between reach-to-grasp, tool-use, and piano playing demonstrates the need to study a wide array of manipulation tasks if we truly want to replicate human dexterity. Future work should investigate the synergies that emerge from complex manipulation tasks as building blocks for manipulation.

Perhaps the greatest value of the proposed VAF-CI is as a metric for bio-inspired learning from demonstration. One proposed solution to multi-finger dexterous manipulation is to learn control from human demonstration [Qin et al., 2022]. However, as previously motivated, if we want our robots to perform dexterous manipulation tasks, we should train them using a rich set of manipulation data. As VAF-CI proved to be a reliable, continuous and robust indicator of manipulation complexity, it can be used to ensure that a spectrum of manipulation tasks is presented to a dexterous manipulator during learning from demonstration.

5.5 Conclusions

Here, we developed a novel complexity index, VAF-CI, able to objectively quantify the complexity of a manipulation task. After measuring the kinematic behavior

of the hand joints in three different manipulation tasks of increasing complexity (reaching-to-grasp, tool-use and piano playing), the proposed index identified a clear distinction between these groups. Moreover, it could be used by an unsupervised learning algorithm to reasonably identify 3 different classes of manipulation. This new metric could support the improvement of learning algorithms to achieve robot dexterity from human demonstration and human bio-inspiration. Future work should investigate human dexterity during reasonably complex manipulation tasks quantified by the VAF-CI.

Chapter 6

Conclusions

The goal of this thesis was to provide a descriptive understanding of human manipulation in unimpaired subjects. To achieve this, human subject experiments were conducted to uncover simplified control mechanisms used by humans (Figure 1-1). A summary of these results can be visualized in Figure 6-1. The icons on the right side of the figure highlight the perception of stiffness study (Chapter 4) and force control study (Chapter 3). The former led to the indirect force control formulation (adapted from Figure 3-1), while the latter highlighted that mechanical impedance was essential to humans' internal model. The icon of the mock electrical cabinet (adapted from Figure 5-1) emphasizes that synergistic motion may be used to simplify the planning of hand motor function during tool use and object manipulation, as discussed in Chapter 2. However, it is important to note that these synergies were task-dependent, suggesting the existence of a “synergy library” in which task-dependent synergies are stored for various tasks [Tessari et al., 2024]. This raised the question: how can we ensure that studied tasks represent the full spectrum of human hand manipulation capabilities? The icon above the “Synergy Library” block denotes VAF-CI, a metric that can be used to answer this question. In summary, Figure 6-1 highlights how each chapter in this thesis contributes to the development of a descriptive understanding of human manipulation in unimpaired subjects. Below, we will recap in greater detail the contributions of this thesis.

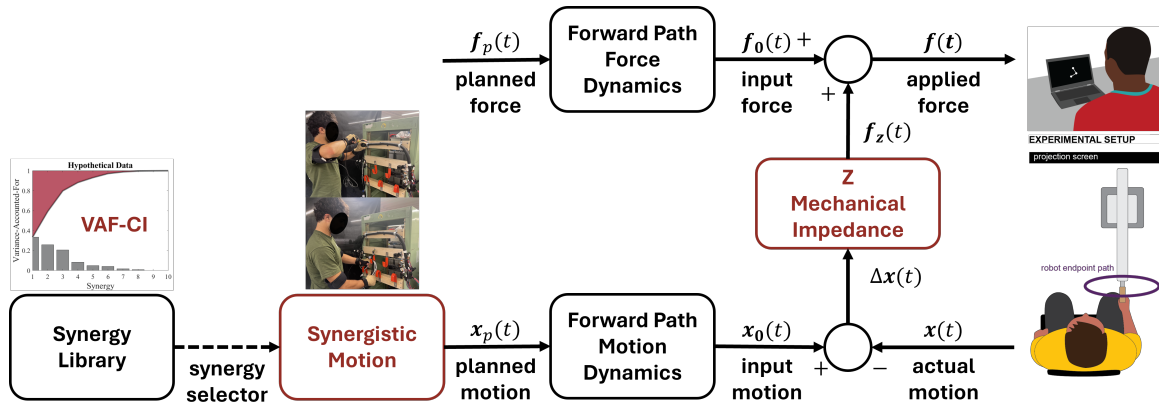


Figure 6-1: **A descriptive understanding of human manipulation.** Here we highlight the contributions of this thesis that have pushed towards a descriptive understanding of human manipulation in unimpaired subjects. Each icon highlights the particular study that contributed to a finding in the developed block diagram.

In Chapter 2, our aim was to uncover evidence of kinematic hand synergies during wire harness installation, a task involving the manipulation of complex tools and objects. Our findings revealed a reduction in operational degrees of freedom, indicating the presence of synergies during both reaching-and-grasping and manipulation tasks. However, manipulation of a tool generally necessitated more significant synergies than grasping the same tool. Moreover, manipulation often required a different subset of synergies than reach-and-grasp. These results underscore the nuanced complexities of functional hand use beyond basic grasping actions, highlighting the importance of studying the human hand during functional manipulation tasks beyond grasping.

Transitioning to Chapter 3, our focus shifted to human force regulation, addressing an overlooked question from Chapter 2: how do humans manage force during physical interaction? Contrary to longstanding assumptions in robotics, this study revealed that the modulation of force is coupled to motion via mechanical impedance. Moreover, in this work, we developed a simple mathematical model competent to describe how humans control physical interaction and their limitations. These findings have significant implications for the design of devices collaborating with humans, emphasizing the importance of a quantitative model in understanding human upper-limb force control.

Chapter 4, building upon Chapter 3, delved into whether mechanical impedance is essential to the human internal model. Using a perceptual study, we demonstrated that humans can infer changes in limb stiffness from multi-joint limb motion. It should be emphasized that this is quite a feat; stiffness is the relation between motion and force, yet subjects were able to perceive stiffness using only motion information. Notably, we identified that human perception of stiffness is dominated by path information. These observations provide new insights into how humans interpret motor actions and interactions, offering implications for fields such as robotics and rehabilitation.

Finally, in Chapter 5, our aim was to address a question subtly provoked in Chapter 2. While Chapter 2 established the potential of synergies to simplify motor control during manipulation tasks, it also highlighted their task-specific nature, raising the challenging question of ensuring that studied tasks represent the full spectrum of human hand manipulation capabilities. Here, we proposed the Variance-Accounted-For based Complexity Index (VAF-CI), a novel metric capable of objectively quantifying the complexity of manipulation tasks. Validated across various dexterous tasks, this index can ensure that future studies probe a wider range of hand capabilities and synergistic behavior, offering new pathways for advancing robot dexterity via bio-inspiration.

The insights in this thesis stemmed from human subject experimentation, dedicated to uncovering the simple control mechanisms humans use for motor function. This approach was not random but driven by the belief that understanding these strategies requires direct human observation. Our research was driven by a commitment to developing simple models of these mechanisms. Preferring simplicity over complexity was not just a methodological choice but a strategy to ensure our studies provided immediate, understandable insights. Simple models offer clarity that complex systems often obscure, making it easier for both researchers and practitioners to grasp human motor control principles and apply them to robotic systems and rehabilitation techniques. This focus on simple, interpretable models distinguishes our work, bridging the gap between human motor function complexity and the practical

needs of robotics and rehabilitation, ensuring the insights we have uncovered will have tangible impacts on these fields.

6.1 Potential Future Work

While this thesis has made notable strides toward achieving a descriptive understanding of human manipulation in unimpaired subjects, thorough research typically uncovers more questions than answers, and this study is no different. The insights gained have shed light on numerous areas primed for deeper exploration. In the pursuit of furthering this field, presented here are several potential avenues for future research.

6.1.1 Kinematic Hand Synergies

In the literature, there has been a persistent question regarding the neural basis of synergies [Santello et al., 2016]. The key question is, are synergies the result of physical constraints and biomechanics, or neuromotor control? The answer to this questions speaks to whether or not synergies are in fact used as a simplifying control technique. In Chapter 2, we analyzed kinematic data without considering constraints imposed by the tools or the external environment, thus hindering our ability to address this question. One potential method to address this issue involves a more refined analysis, such as projecting measured variations onto the tangent plane of the constraints. This approach could provide a more accurate estimation of biological synergies, rather than those influenced by mechanical constraints. For instance, if a subject is required to slide their fingertip along a surface, the surface's shape would impact the recorded finger kinematics and resultant singular value distribution. However, projecting the measured variation onto the plane tangent to the constraint could yield a better estimate of intrinsic motor control strategies, isolating them from extrinsic task-specific adaptations. In fact, similar work has been conducted in [Sharma, 2023]. However, it's analysis could be extended to the functional manipulation task presented here. Future research could focus

on developing analytical techniques that incorporate these considerations, thereby improving the fidelity of synergy analysis in complex, constrained manipulation tasks.

Moreover, exploring kinematic hand synergies further offers a promising avenue for future work, particularly through the development of a “Manipulation Taxonomy.” This taxonomy aims to distill the essential synergies, or canonical motions, underpinning dexterity in complex manipulation tasks into a set of simple models. By identifying these core synergies, researchers can advance our understanding of hand manipulation, providing valuable insights for both theoretical exploration and practical application in robotic and prosthetic design.

We envision that the proposed future work can be pursued through two main research branches. The first branch focuses on human subject experiments in virtual environments, where participants control an anthropomorphic robot in a simulated environment that mirrors the subject’s real hand movements. By systematically varying the subset of synergies available to the simulated hand, this approach aims to pinpoint which synergies are vital for manipulation. This work not only seeks to establish a Manipulation Taxonomy but also evaluates the efficacy and limitations of synergy-based control strategies in complex tasks. Such investigations can enhance the design and functionality of prosthetic devices, ensuring they more closely emulate the dexterous capabilities of the human hand.

Simultaneously, the second research branch leverages the power of machine learning to investigate how machines can emulate human dexterity with minimal control inputs. Employing reinforcement learning within simulated environments, an anthropomorphic hand will endeavor to accomplish manipulation tasks by adopting a synergy-based control strategy. This approach facilitates the exploration of which synergies AI utilizes in learning to manipulate objects, providing valuable insights into efficient robot learning processes for achieving dexterous manipulation. These dual research trajectories not only contribute to the development of a Manipulation Taxonomy but also shed light on avenues to instill robots with human-like dexterity, bridging the gap between human motor skills and robotic capabilities.

6.1.2 Visual Perception of Stiffness

The visual perception of stiffness study presented in Chapter 4 opens up a whole new line of research questions. Some might include:

- Can humans perceive damping from visual observation of motion?
- Can we extend the visual observation of motion task to a three link manipulator?
- Can humans perceive stiffness during the kinematic simulation of an anthropomorphic dynamic walker?
- How do humans perceive stiffness during 3-dimensional motion?

The last question in this list can be addressed through a VR-based perceptual experiment. Specifically, we will explore whether and how subjects can gauge the impedance of a two-link arm moving within a 3D environment. These studies will not only test the boundaries of human perceptual abilities but also serve as a foundation for developing a software tool that applies computer vision and machine learning algorithms to estimate limb impedance non-invasively. Looking further ahead, the long-term goal of this research is to create a comprehensive tool that can accurately measure limb dynamics and muscle activity without direct physical contact. This tool will be rigorously validated through human subject experiments focusing on upper-limb kinematics and dynamics. The development of such a non-invasive impedance measurement tool holds the promise of wide-ranging applications, from monitoring muscle health in the elderly to optimizing athletic training regimens. By advancing our understanding of visual perception of stiffness and translating these insights into practical tools, we can pave the way for innovative solutions that enhance rehabilitation, and personal health monitoring.

6.1.3 Variance-Accounted-For based Complexity Index (VAF-CI)

The full potential of the Variance-Accounted-For based Complexity Index (VAF-CI) metric remains largely untapped. Its expansion to musculature data holds promise for uncovering insightful observations regarding movement pattern optimization in skilled manipulators. We hypothesize that with extensive practice, highly skilled individuals refine their movements to follow patterns engaging fewer synergies, thereby minimizing musculature effort. This optimization may represent an efficient strategy for reducing both cognitive and physical demands associated with dexterous manipulation.

A forthcoming research endeavor could entail capturing hand kinematics from skilled and novice dexterous manipulators, such as comparing an experienced attending surgeon to a first-year resident during simulated surgical procedures. The objective is to quantify movement complexity and discern differences in synergy utilization between these groups using the VAF-CI and synergy-based analysis. This investigation aims to reveal how experience and skill level influence hand movement coordination. By elucidating movement strategies contributing to surgical proficiency, this research could streamline surgical education, potentially minimizing errors and enhancing patient outcomes.

In the future, our overarching objective is to develop tools that streamline the acquisition of complex motor skills. We aim to create a sophisticated tool capable of autonomously monitoring surgical proficiency using VAF-CI and providing immediate, synergy-driven feedback. Such a tool has the potential to significantly transform surgical training methodologies.

Beyond surgery, insights gleaned from expert manipulators could aid individuals, such as stroke patients, in regaining lost motor skills. This approach holds promise for improving rehabilitation processes by providing targeted, synergy-focused guidance to restore dexterity. Furthermore, a deeper, model-based comprehension of dexterous human activities, spanning surgery, sports, and artistic performance, could inspire

advancements in robotic automation, enhancing automated systems' ability to replicate or augment human movements. Through this multifaceted exploration, the research aims to expand our understanding of human motor control and dexterity, with far-reaching implications for enhancing both human and robotic capabilities.

6.2 Outlook

In a world increasingly dominated by the narratives of artificial intelligence (AI) and machine learning (ML), the enduring value of simple models underscores the significance of foundational science in unraveling complex systems. In the pursuit of ever more intricate computational models, there remains a critical need for scientists and engineers dedicated to the fundamental science behind human motor control and its practical applications. Simple models, characterized by their clarity and interpretability, serve as a vital complement to the complexity of AI and ML, offering a crucial link between theoretical insights and real-world implementations. Embracing simplicity represents a deliberate approach to understanding the intricacies of human motor function and designing robotic systems that emulate the grace and efficiency of natural human movements. It is within this realm that engineers, focusing on the foundational aspects of motor control, can make significant strides, leveraging simple models to uncover the principles that enable humans to interact seamlessly with their surroundings.

This perspective fundamentally reshapes how we approach the study of human motor control and its integration into robotics. If humans can navigate their environment effectively through simplified, intuitive models, then robotics has the opportunity to adopt a similar approach. By prioritizing simple, interpretable models, researchers and engineers can develop robotic systems that not only comprehend but also adapt to the nuances of human interaction. These systems would possess the ability to execute complex tasks without the need for elaborate computational modeling, mirroring the efficiency and adaptability of human motor function.

Furthermore, the emphasis on simple models illuminates the potential for these

models to serve as a foundation for more advanced learning and adaptation in robots. Just as humans learn and refine their motor skills over time, robots equipped with simple, interpretable models could be designed to evolve and adapt their interactions based on experience, achieving a level of dexterity and responsiveness akin to human capabilities.

In conclusion, the case for simple models in studying human motor control and developing robotic systems is compelling. It aligns with the inherent efficiency of human motor strategies and offers a pathway toward creating robots that can more naturally and effectively interact with their environment and with humans. This approach not only facilitates a deeper understanding of the underlying principles of human motor control but also advances the field of robotics towards more intuitive, accessible, and adaptable systems.

A video presentation of this thesis can be found at:
<https://www.youtube.com/watch?v=u2eCJHqEGww>

Appendix A

The Study of Complex Manipulation via Kinematic Hand Synergies: The Effects of Data Pre-Processing

The exploration of kinematic hand synergies, employing matrix decomposition techniques such as singular value decomposition, as undertaken in Chapters 2 and 5, suggests that humans potentially regulate a predetermined subspace of motions during manipulation tasks, commonly referred to as synergies. However, the manner in which data is pre-processed significantly influences the quantitative conclusions drawn about these synergies. This chapter aims to elucidate the impact of data pre-processing on the analysis of hand synergies by examining both numerical simulation and real kinematic data from a complex manipulation task, namely piano playing. The findings indicate that centering the data, achieved by removing the mean, emerges as the most suitable pre-processing technique for studying kinematic hand synergies,

This chapter is adapted from a paper titled, “The Study of Complex Manipulation via Kinematic Hand Synergies: The Effects of Data Pre-Processing” [West et al., 2023] published in the 2023 International Consortium on Rehabilitation Robotics (ICORR). This work was done in collaboration with Federico Tessari and Neville Hogan.

A video summary of this chapter can be found at:
<https://www.youtube.com/watch?v=J1Xp5Ekl6JQ>

A.1 Introduction

The human hand is the end-effector that serves humans in interaction with themselves, the external environment, and other individuals [Klopsteg et al., 1955]. It is an anatomically complex system comprised of 27 bones and 30 muscles [Kandel et al., 2013] that are controlled to perform highly dexterous manipulation tasks [Napier, 1956, Diogo et al., 2012]. All in all, the hand presents more than 20 degrees-of-freedom (DOFs) [Kandel et al., 2013], thus providing a highly redundant control space to interact with the environment which, for a rigid object, has at most 6 DOFs.

An average of 800,000 individuals in the US suffer from Cerebral Vascular Accident (CVA) [Virani et al., 2020] every year, while roughly 700,000 upper-limb hand amputees currently live in the United States [McDonald et al., 2020, Ziegler-Graham et al., 2008]. Losing hand functionality poses a great barrier for these individuals in a society designed for humans with dexterous hands. This has a major detrimental impact on their quality of life as they lose their capability to perform numerous activities of daily living. To improve the available rehabilitation and assistive technologies for impaired individuals (e.g., hand prostheses and rehabilitative exoskeletons), it appears necessary to deepen the understanding of how humans control and move their hands.

Among the many research questions about the human hand, there is growing interest in comprehending how humans seamlessly move a highly redundant system without particular apparent effort. Santello et al. were the first to suggest that kinematic synergies of the hand (i.e., linear combinations of joint motions) could provide a reasonable reduction of dimensionality, decreasing the number of effective DOFs to be controlled [Santello et al., 1998]. Specifically, they observed that most (> 80%) of the joint kinematic variance during grasping could be described by only two synergies. Subsequent studies confirmed similar results [Santello et al., 2002, Santello et al., 2016, Todorov and Ghahramani, 2004, Weiss and Flanders, 2004, Mason et al., 2001].

From a mathematical point of view, these researchers approached the study

of variance by means of matrix decomposition techniques. Specifically, they took advantage of Singular Value Decomposition (SVD) [Stewart, 1993] and Principal Component Analysis (PCA) [Pearson, 1901]. These methods allow the decomposition of a non-square matrix - containing the collected kinematic data of the hand - into singular vectors (or principal components) and associated singular values (or principal component scores) to investigate the distribution of variance. However, pre-processing of the hand kinematic data before performing the decomposition plays a crucial role. For example, [Santello et al., 1998] centered each joint’s data around its mean value, [Todorov and Ghahramani, 2004] considered two cases: joint data re-scaled to contain values between 0 and 1, and joint data re-scaled to have unit variance, while [Mason et al., 2001] used the raw data binned to fixed increments of time between start and end of a trial. These pre-processing approaches led to similar conclusions in terms of dimensionality reduction, but different quantitative estimates of the effective amount of this reduction. While some have explored how data structure affects synergy extraction [Zhao et al., 2022, Lambert-Shirzad and Van Der Loos, 2017] and others have investigated how electromyography pre-processing changes the resultant muscle synergies [Kieliba et al., 2018], to our knowledge no-one has explored how pre-processing affects kinematic hand synergies.

In this work, we present an analysis of the effect of the most common pre-processing techniques on the dimensionality reduction problem. First, to provide intuitive understanding of the effect of the different pre-processing methods, we show a numerical example for a simulated 2 degree-of-freedom (DOF) system. Then to validate the intuition gained from the 2 DOF system, we apply the same approach to real hand kinematic data obtained from a subject performing a complex hand manipulation task: piano playing.

Studying piano playing has two intrinsic advantages. First, it is widely recognized to be a complex manipulation task [Chen et al., 1997] that requires specific training and years of practice [Krings et al., 2000]. Thus, it should exhibit a more refined distribution of variance than a simpler task such as object grasping. Second, playing piano requires limited range-of-motion trajectories of the joints, thanks to the

constraint imposed by the piano keyboard. Joint trajectories are generally curved, but over small ranges of motion, a linear approximation may be satisfactory. Therefore, piano playing is an ideal task to evaluate linear mathematical tools such as Singular Value Decomposition. Moreover, the current literature [Santello et al., 2002, Santello et al., 2016, Todorov and Ghahramani, 2004, Weiss and Flanders, 2004, Mason et al., 2001] largely studies grasping of common household items, which approximates a lower bound on the repertoire of the hands’ capabilities. However, if we want to restore activities of daily living that involve manipulation it is important to specifically study hand manipulation. Studying piano playing can help approximate an upper bound on the hands’ capabilities.

The results presented in this study help to shed light on the quantification - and correct interpretation - of the synergistic behavior of the human hand. These results have implications for the design of assistive and rehabilitative technologies.

A.2 Methods

A.2.1 Theoretical Background

Consider a data matrix $X \in \mathbb{R}^{n \times m}$, where n represents the number of observations (e.g., the time evolution of a joint angle or degree of freedom) and m represents the number of features (e.g., the number of analyzed DOFs) for a given task. The main goal in the study of kinematic hand synergies is to mathematically process i.e., decompose, the matrix X to identify patterns (synergies) that could be related to a reduction of the dimensionality of the control problem. This is a proposed solution to the infamous “curse of dimensionality”, which states that higher degree-of-freedom systems are more difficult to control [Bellman, 1966]. Dimensionality reduction is often found via SVD.

In SVD, the matrix X is decomposed into the product of an orthonormal matrix U , a matrix with values only on its main diagonal S , and another orthonormal matrix

V :

$$X = U \cdot S \cdot V^T \quad (\text{A.1})$$

where $U \in \mathbb{R}^{n \times n}$, $S \in \mathbb{R}^{n \times m}$ and $V \in \mathbb{R}^{m \times m}$. From a geometric perspective, these operations could be interpreted as the succession of a rotation U , a stretch S and another rotation V^T . Moreover, the principal diagonal elements of S , also known as singular values σ_i , are related to the eigenvalues λ_i of the square matrix $X^T \cdot X$ by the following relationship: $\sigma_i^2 = \lambda_i$ for $i = 1, \dots, m$.

Mathematically, SVD provides information about the matrix X 's distribution over the hyperspace $\mathbb{R}^{n \times m}$; of particular interest is the data distribution over the feature space \mathbb{R}^m . The matrix V^T can be used to obtain the principal directions in the feature space, while the singular values - obtained from S - return an estimate of the magnitude of variance projected onto each principal direction.

To obtain a measure of the data variance-accounted-for (VAF) by a given synergy, we used the singular values along the principal diagonal of matrix S . Specifically, we report VAF as a decimal, ranging from 0 to 1 obtained by dividing the square of a singular value by the total sum of the squares of the singular values: $VAF(i) = \sigma_i^2 / \sum \sigma_i^2$ for $i = 1, \dots, m$.

Pre-processing operations performed on the data matrix, X , will, of course, affect the singular value decomposition and the resultant estimation of the related VAF for a given analyzed motion. The important question is how and how much.

In the work presented here, we considered four different pre-processing operations that have previously been used in the literature for kinematic hand synergy decomposition: (1) no pre-processing [Mason et al., 2001], (2) removing the mean [Santello et al., 1998], (3) z-score [Todorov and Ghahramani, 2004], and (4) range 0-1 [Todorov and Ghahramani, 2004]. Mathematically, the latter three approaches result in three new data matrices:

$$X_{rm} = X - \text{mean}(X) \quad (\text{A.2})$$

$$X_{zs} = \frac{X - \text{mean}(X)}{\text{std}(X)} \quad (\text{A.3})$$

$$X_{r01} = \frac{X - \min(X)}{\max(X - \min(X))} \quad (\text{A.4})$$

where X_{rm} represents a data matrix with the mean removed, X_{zs} represents a z-scored data matrix, and X_{r01} represents a data matrix scaled between 0 and 1. Excluding trivial cases in which the data matrix features $x(1 : n, m)$ already have either zero mean, unit variance or a 0-1 range, it is evident that the singular values and related orthogonal directions $v_i \in V^T$ will change due to pre-processing.

A.2.2 Numerical Validation

A set of numerical simulations was performed in Matlab 2022b (Mathworks, USA) to observe the changes in singular value decomposition due to the different pre-processing methods. To provide an intuitive understanding, a 2 DOF system ($m = 2$) was considered to facilitate graphical interpretation of the numerical simulations. The two DOFs were linearly correlated with the addition of external noise:

$$x_2(i) = k_1 x_1(i) + k_0 + \epsilon(i) \quad i = 1, \dots, N \quad (\text{A.5})$$

where $x_1(t), x_2(t)$ represent two features of the system, k_1, k_0 are the linear correlation parameters between these features, $\epsilon(t)$ is an external disturbance, and N represents the number of simulated points. The results presented in Section A.3.1 used the following simulation parameters: $k_1 = 0.50$, $k_0 = 3$, and $N = 101$. SVD was performed on the data matrix $X = [x_1, x_2] \in \mathbb{R}^{N \times 2}$ and the results of all the proposed pre-processing techniques.



Figure A-1: **Experimental setup** of one subject playing the piano while kinematic data of their right hand was recorded with a CyberGlove.

A.2.3 Piano Playing

To validate the intuition provided by the numerical example, a preliminary experimental validation on real human hand kinematic data was performed with a single subject¹ performing a complex manipulation task: playing the piano (Fig. A-1). He was informed about the experimental procedure and agreed to sign a consent form. All procedures were approved by MIT's Institutional Review Board.

Piano playing was selected due to its evident complexity and high level of required skill. This was expected to address a common weakness of synergy analysis via SVD: most of the data variance is 'lumped' into the first few synergies², compromising the reliability of numerical differences between the higher-order synergies. Piano-playing was expected to evoke a wider variability of the extracted synergies, thus avoiding the effect of cumulative lumping of VAF into the first few synergies. This would provide a better understanding of the effect of different pre-processing methods.

The subject separately performed a set of 7 piano pieces (Chopin: Nocturne

¹The subject was a 31 year old, right-handed skilled piano playing male who had 10 years of trained piano practice and 20 years of overall piano experience.

²Note, in the piano playing study the angles of the 23 joints of the hand comprise the feature space. Thus, a reported synergy is a vector (hyperspace) representing a common hand configuration which is described by angles of the joints in the hand.

Op. 9 No. 2; Chopin: Waltz Op. 64 No.1; Giovanni Allevi: Come Sei Veramente; Hans Zimmer: Oogway Ascends; W.A. Mozart: Turkish March; Christopher Norton: A Whimsy; Paul Desmond: Take Five) while wearing a CyberGlove (CyberGlove; Virtual Technologies, Palo Alto, CA), a glove with embedded sensors that measure joint kinematics (Fig. A-1). Specifically, the flexion of the distal interphalangeal (DIP), proximal interphalangeal (PIP), and metacarpophalangeal (MCP) joints of the four fingers were measured. Additionally, the abduction (ABD) of the four fingers at the MCP joints was measured. At the thumb, the flexion at the MCP and interphalangeal (IP) joints, abduction (ABD) at the carpometacarpal joint, and rotation (ROT) about an axis passing through the trapeziometacarpal joint were measured. Lastly, palm arch (PA) and wrist (W) pitch and yaw were measured. The subject wore the CyberGlove on his right hand. The CyberGlove collected samples at ~ 200 Hz with a nominal angular resolution of $< 0.1^\circ$.

Synergies (i.e., linear combinations of the joint DOFs) and their VAF were extracted using the SVD algorithm presented above. To identify the number of significant synergies, we followed the method of [Lambert-Shirzad and Van Der Loos, 2017]. Specifically, we report the number of synergies required to achieve at least 90% VAF and where inclusion of another subsequent synergy did not add an additional 5% VAF. To determine if there was an effect of pre-processing on the number of significant synergies, these values were submitted to a 4 (pre-processing type) x 1 repeated-measures ANOVA.

Moreover, we compared the calculated kinematic hand synergies, V , across pre-processing types. To do so, we computed the product of each individual piece’s synergies identified by a pre-processing type with those of a different pre-processing type:

$$C = |\cos(\theta)| = |V_{preprocessing1}^T \cdot V_{preprocessing2}| \quad (\text{A.6})$$

This resulted in a matrix consisting of the cosine similarities, $C(i, j)$, between synergies where i denotes the i^{th} synergy of pre-processing 1 and j denotes the j^{th} synergy of pre-processing 2. Here, we report the magnitude of these cosine similarities,

ranging from 0 to 1. If synergies were the same irrespective of pre-processing (i.e. the data spans the same hyperspace), we would expect a matrix of cosine similarities with ones on the diagonal and zeros otherwise. In all other cases (i.e. $V_{preprocessing1} \neq V_{preprocessing2}$), the resulting matrix, $C \in \mathbb{R}^{m \times m}$, will be an asymmetric square matrix.

All data processing and statistical analyses were performed using custom scripts in MATLAB. The significance level for statistical tests was $\alpha = 0.05$.

A.3 Results

A.3.1 Numerical Validation

In Fig. A-2 we present singular vector decomposition of numerical data as described in Section A.2.2. There the raw and pre-processed data can be seen. It is seen that in the raw data case, the first eigenvector does not appropriately represent the slope of the data (Fig. A-2a).

The coefficients of the first and second eigenvectors are shown in Fig A-3a-b. Given that the slope of the data was 0.50 (i.e., $k_1 = 0.50$), we expect the ratio of the coefficients of the first eigenvector to equal ~ 0.50 with some error due to the added noise (i.e., $V(2,1)/V(1,1) \sim 0.50$). The observed ratio was 1.177, 0.550, 1.000, and 0.952 in the raw, mean removed, z-score, and range 0-1 data, respectively. Note, the mean values of the raw data are 6.022 and 5.013; their ratio is 1.201. Additionally, the mean values of the range 0-1 data are 0.486 and 0.503; their ratio is 0.967. This indicates when the data is not centered, the first synergy is directed towards the mean of the data.

Fig A-3c shows the VAF of each eigenvector. The VAF of the first eigenvector was 0.995, 0.968, 0.957, and 0.989 in the raw, mean removed, z-score, and range 0-1 data, respectively. Evidently, when the data was not pre-processed, the second synergy was considered to negligibly contribute to the variation of the data. While it may appear to be an advantage to have a system adequately described by fewer synergies, it is

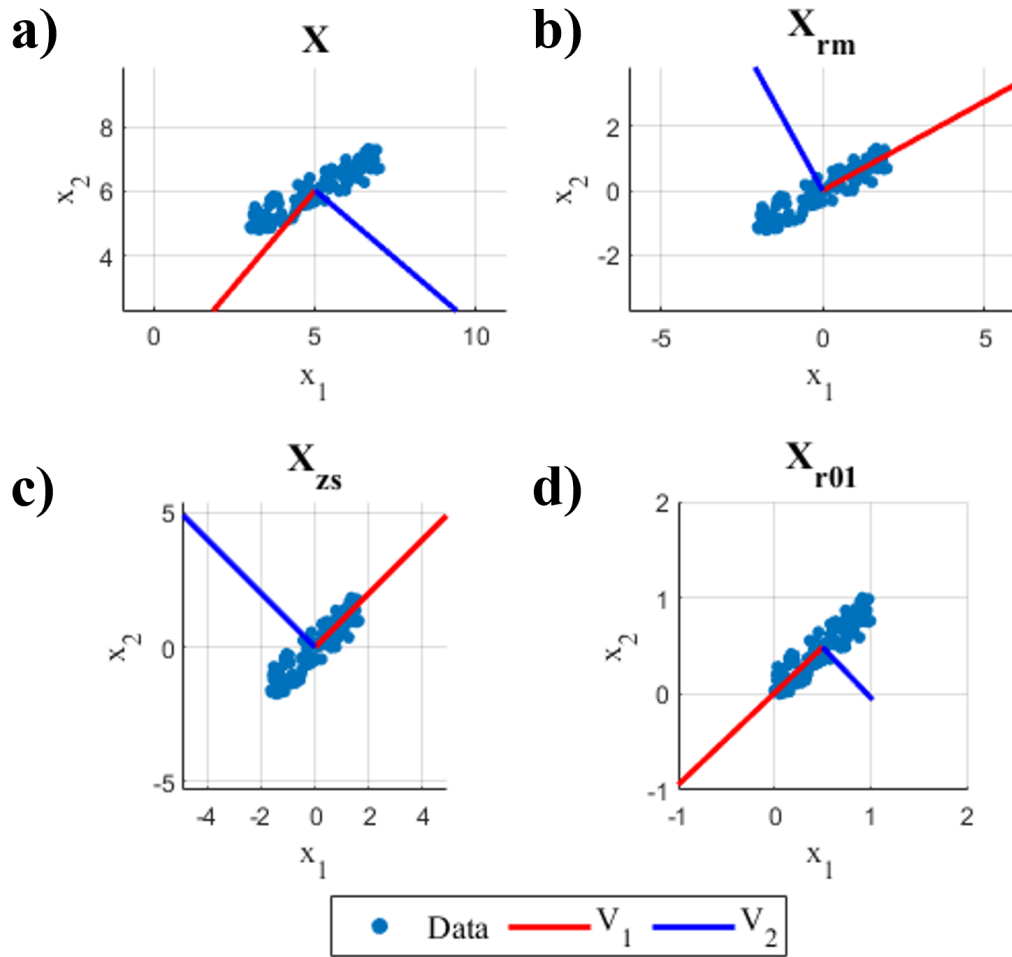


Figure A-2: **Numerical validation data** of the various pre-processing types are shown by the light blue dots. Additionally, the computed 1st and 2nd eigenvectors (i.e. synergies) are shown by the red and blue lines, respectively. (a) denotes the raw data while (b), (c) and (d) denote the data after the pre-processing described by Eq. 2, 3, and 4, respectively.

important to recall that in the raw data case the first synergy was incorrect; rather than representing the co-variation in the data it reflected the data mean.

A.3.2 Piano Playing

Fig. A-4 demonstrates the VAF in the piano experiments averaged across pieces for each pre-processing type. It is seen that, on average, two synergies achieve 0.90 VAF when data is not pre-processed or set to range from 0 to 1. However when the data

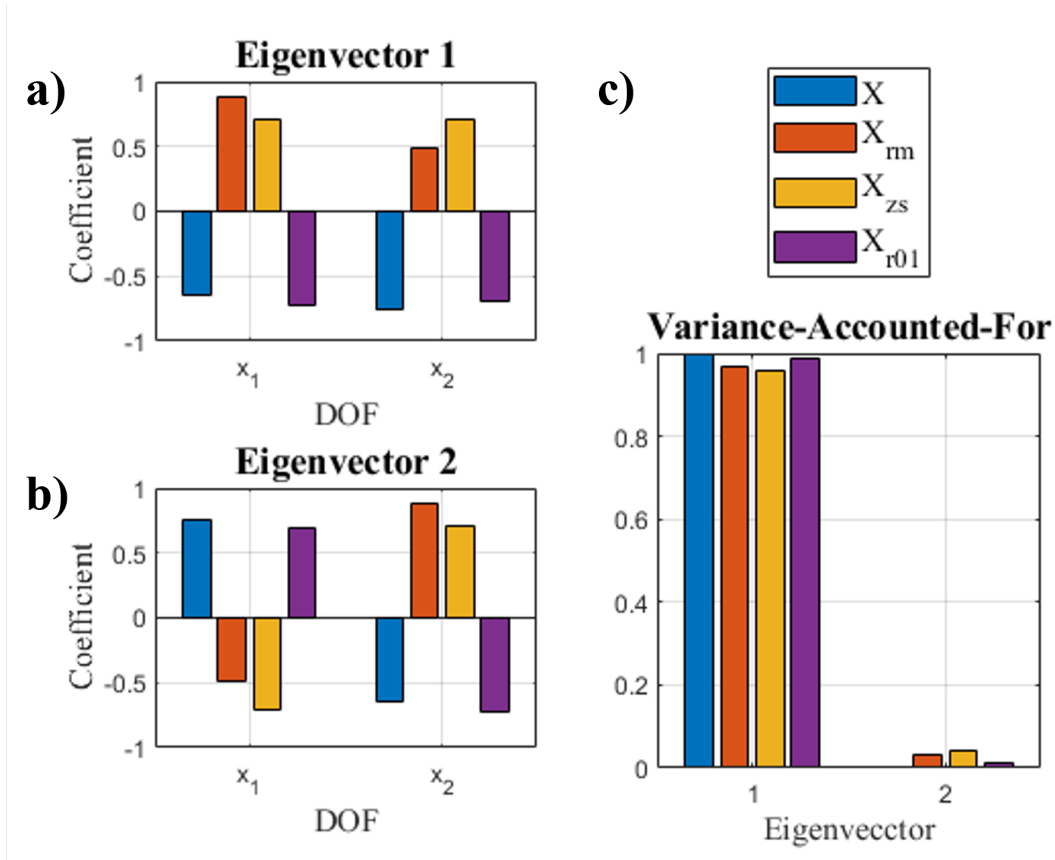


Figure A-3: **Eigenvectors and VAF of the numerical validation data.** Different colors denote different pre-processing types.

has the mean removed or is z-scored, 6 and 9 synergies (respectively) are needed to reach 0.90 VAF.

We quantified the number of significant synergies based on each data pre-processing type. The average and standard deviation is reported in Table A.1. A one-way ANOVA revealed a significant effect of pre-processing on the number of significant synergies ($F_{3,24} = 94.47$, $p = 1.99e - 13$). Post-hoc t-tests (Bonferroni corrected $\alpha = 0.05/6 = 0.0083$) revealed that the range 0-1 data had a statistically fewer number of significant synergies than the raw data ($p = 3.67e - 04$), data with the mean removed ($p = 1.23e - 08$), and z-score data ($p = 3.35e - 08$). Moreover, the raw data had a statistically fewer number of significant synergies than the data with the mean removed ($p = 3.70e - 10$), and z-score data ($p = 2.28e - 07$). Further, the data with the mean removed had a statistically fewer number of significant synergies

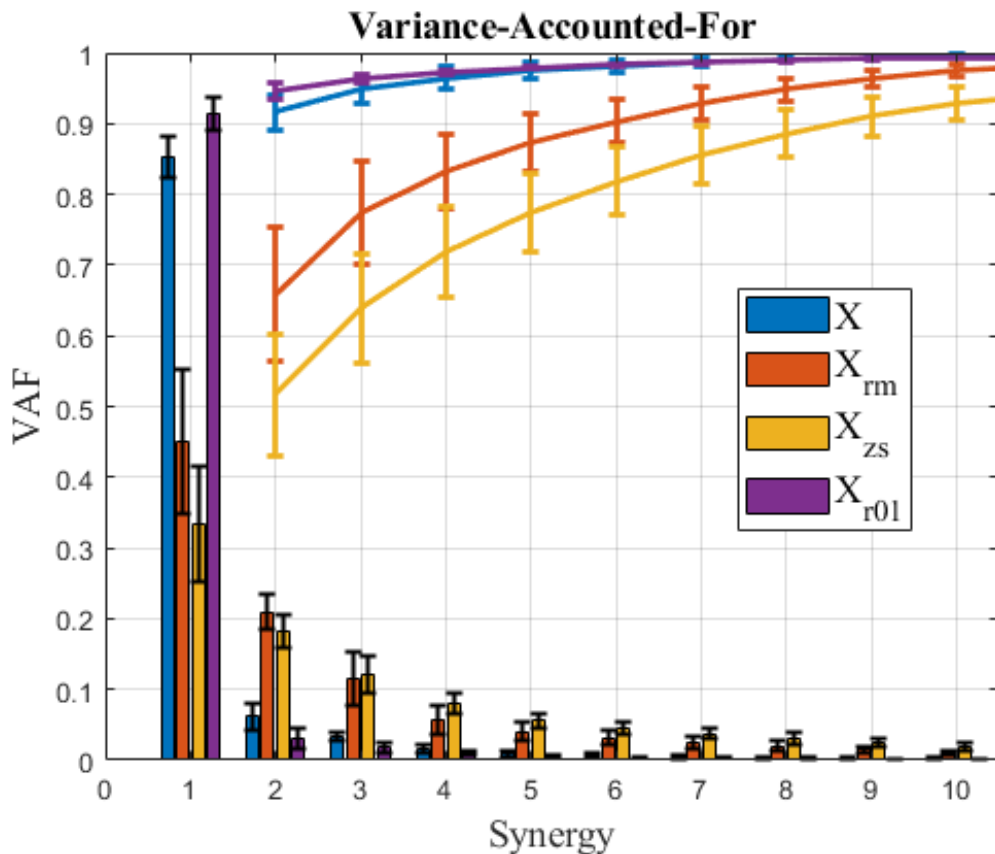


Figure A-4: **VAF of each synergy in the piano experiment** averaged across piece. Different colors denote different pre-processing types. Lines show the cumulative sum of the VAF. Errorbars are $\pm 1SD$. Due to space constraints, only the first 10 synergies are shown.

	X	X_{rm}	X_{zs}	X_{r01}
Avg	2.29	6.14	9.00	1.14
SD	0.49	0.90	1.63	0.38

Table A.1: **Number of significant synergies in the piano playing study** based on data pre-processing.

than the z-score data ($p = 0.0016$).

To compare the synergies across pre-processing type, we computed the magnitude of the similarity coefficients between the matrices of synergies. The 6 comparisons, stemming from the pairwise combinations of the 4 pre-processing types (i.e., $\binom{4}{2} = 6$), are reported in Fig. A-5 for a representative piece. If two types of pre-processing methods lead to similar synergies we would observe a matrix of ones on the diagonal

Chopin: Waltz Op. 64 No.1




Figure A-5: **Comparison of synergies across different pre-processing types** for one representative piece. The matrix of cosine similarities between synergies of different pre-processing types is presented for the Chopin Waltz Op. 64 No. 1. A black value denotes a cosine similarity of 1 and a white value denotes a cosine similarity of 0. Due to space constraints, comparisons of only the first 10 synergies, ordered by decreasing VAF, are shown.

and zeros otherwise. Clearly this is not the case. In fact, when comparing the raw data synergies with the ones obtained from removing the mean (Fig. A-5a), we observe that there is a large similarity in the subdiagonal of the matrix of cosine similarities. High similarity is also seen in the superdiagonal when we compare the z-score data synergies and the range 0-1 synergies (Fig. A-5f). These relations were also observed in the pieces not reported here.

A.4 Discussion

A.4.1 Effects of Centering the Data

The mean removed and z-scored data are centered while the raw and the range 0-1 data are not. In Fig A-3c, it is seen that the first synergy of non-centered data has a greater VAF than the centered data. This was also observed in the piano study's resultant synergies (Fig. A-4). Due to the large first synergy VAF, the non-centered data had statistically fewer significant synergies than the centered data (Table A.1).

In the validation data, Fig A-2a showed that the first eigenvector did not align well with the expected principal direction of the data. Recall that the validation data was set up using Eq. A.5, where the slope, k_1 , was 0.50. Thus, we expected the ratio of the coefficients of the first eigenvector (i.e., $V(2,1)/V(1,1)$) to be 0.50. This was not the case for the raw data – $V(2,1)/V(1,1) = 1.18$. Fig A-2a suggests that this first synergy was used to reach the center of the data, resulting in an eigenvector whose coefficients did not tell us how our data co-varied. Moreover, this relationship was also observed in the piano data. The cosine similarity matrix was computed to compare synergies across pre-processing types (Fig. A-5). If the synergies were similar, we would have observed a matrix of ones on the diagonal and zeros otherwise. Figure A-5a shows a matrix that contains high similarity on the subdiagonal when comparing the raw data synergies to the ones obtained from removing the mean. Given this in conjunction with the observation made in Fig A-2a, we hypothesize that the first synergy in piano playing is similarly used to center the data, while the

subsequent synergies inform us how the joints of the hand co-vary. Moreover in Fig. A-5e, the cosine similarity matrix between the z-score data synergies and the range 0-1 data synergies has high similarity in the superdiagonal. Again, the range 0-1 data is not centered, suggesting that the first synergy centers the data, while the subsequent synergies inform us how the joints of the hand co-vary. In sum, not centering the data leads to a first-synergy-dominated behavior that is not representative of the principal directions of motion.

A.4.2 Effects of Changing Data Variance

The z-scored and range 0-1 data change the variance of each DOF while the raw data and the data with the mean removed do not. Specifically, z-scoring forces each DOF to unit variance. As such, we are unable to determine how much individual DOFs co-vary with one another. This is represented by the eigenvectors of the data presented in Sections A.2.2 and A.3.1. We set up these data using Eq. A.5, where the slope, k_1 , was 0.50. Thus, we expected the ratio of the coefficients of the first eigenvector (i.e., $V(2,1)/V(1,1)$) to be 0.50. In the z-scored and range 0-1 data this ratio was 1.00 and 0.95, respectively. Because the variance of the data was changed, information about how joint motions changed in relation to one another was lost. Thus, the ratio of the first eigenvector coefficients was constrained to 1.00, due to z-scoring, as opposed to the expected value of 0.50.

Applying this understanding to kinematic hand synergies, we conclude that reducing the variance of the data would lead us to a synergy that does not describe how one joint varies with another, rather just that the two are related. For example, let's say x_1 and x_2 in Eq. A.5 represent flexion/extension of the index PIP and DIP. Z-scoring the data would result in a synergy that whenever the PIP flexes (or extends) the DIP also flexes (or extends) the same amount. However, based on the data, this is not true; whenever the PIP flexes (or extends), the DIP flexes (or extends) half that amount. To conclude, changing the variance of the data during pre-processing would result in a synergy that does not describe how the several DOFs co-vary.

In this study, we do not have a physiological ground truth of the synergies used in

the piano playing task as we do not have access to humans' internal model. However, we can derive a reasonable understanding of the dimensionality reduction provided by the linear combinations of hand movements using our knowledge of linear algebra and the intuition provided by the numerical example in Section A.3.1. In sum, not centering the data during pre-processing will result in a dominant first synergy that aims to center the data (Fig. A-2a and A-2d), leading to a greatly reduced significant number of synergies. Moreover, changing the variance of the data during pre-processing will lead to a synergy that demonstrates that certain DOFs do, indeed, co-vary but it will not quantify how much (i.e., the ratio between them). Thus, if you are interested in conducting SVD to both accurately estimate the number of significant synergies and how each DOF co-varies with another in a given synergy, you should center your data but not change its variance. This is consistent with the pre-processing step of removing the mean.

A.4.3 A Geometrical Understanding

Understanding that Singular Value Decomposition decomposes a data matrix, X , into a product of a rotation matrix, U , a (diagonal) stretch matrix, S and another rotation matrix, V^T provides insight into how pre-processing methods may affect the decomposed synergies. Specifically, a pure rotation of the data matrix, X , would lead to a pure rotation of the eigenvectors. Moreover, a heterogeneous stretch of each DOF in X will lead to heterogeneous changes of the eigenvalues, resulting in different VAFs. In the z-score X_{zs} and range 0-1 X_{r01} , transformations presented here, each feature (i.e., joint DOF) is separately scaled; that is a heterogeneous stretch. Thus, the VAF of each feature will be accordingly altered.

A.4.4 Implications for Robotic Rehabilitation

The design of rehabilitative devices for the hand is still in search of the best trade-off between number of actuated DOFs, and device function, appearance, and comfort (e.g., minimized weight and size). In this scenario, prosthetic hand designs that

exploit fewer actuators to move multiple joints represent a growing portion of the research prototypes [Laffranchi et al., 2020, Varol et al., 2014, Brown and Asada, 2007, Catalano et al., 2014]. Moreover, clinicians have used synergies as a basis for rehabilitation post Cerebral Vascular Accident [di Luzio et al., 2022, Singh et al., 2018]. Thus, the understanding of how pre-processing affects synergy decomposition presented here can better inform the selection of the quantity and kinematic coupling of synergies that assistive and prosthetic hand devices should implement.

A.5 Conclusion

This work investigated the role of data pre-processing on the extraction of kinematic hand synergies. Using numerical simulation and human hand kinematic data of a subject performing playing the piano, we showed that removing the mean appears to be the best approach to minimize error in the interpretation of computed synergies and their related VAF. This understanding may inform the design of devices that replicate and rehabilitate the human hand. Future work will aim to expand the study of shoulder, arm and hand kinematics during piano playing to uncover the role of synergy decomposition in this complex manipulation task.

Bibliography

- [Abondance et al., 2020] Abondance, S., Teeple, C. B., and Wood, R. J. (2020). A Dexterous Soft Robotic Hand for Delicate In-Hand Manipulation. *IEEE Robotics and Automation Letters*, 5(4):5502–5509.
- [Andrychowicz et al., 2020] Andrychowicz, O. A. M., Baker, B., Chociej, M., Józefowicz, R., McGrew, B., Pachocki, J., Petron, A., Plappert, M., Powell, G., Ray, A., Schneider, J., Sidor, S., Tobin, J., Welinder, P., Weng, L., and Zaremba, W. (2020). Learning dexterous in-hand manipulation. *International Journal of Robotics Research*, 39:3–20.
- [Avenanti et al., 2013] Avenanti, A., Candidi, M., and Urgesi, C. (2013). Vicarious motor activation during action perception: Beyond correlational evidence. *Frontiers in Human Neuroscience*, 7:185.
- [Bach, 2011] Bach, J. S. (2011). *Bach for beginners*. Boosey & Hawkes Music, London, England.
- [Balasubramanian et al., 2010] Balasubramanian, S., Klein, J., and Burdet, E. (2010). Robot-assisted rehabilitation of hand function.
- [Bandura, 1986] Bandura, A. (1986). *Social foundations of thought and action*. Englewood Cliffs, NJ.
- [Banerjee et al., 2015] Banerjee, N., Long, X., Du, R., Polido, F., Feng, S., Atkeson, C. G., Gennert, M., and Padir, T. (2015). Human-supervised control of the ATLAS humanoid robot for traversing doors. In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, pages 722–729. IEEE.
- [Bellman, 1958] Bellman, R. (1958). Dynamic programming and stochastic control processes. *Information and Control*, 1(3):228–239.
- [Bellman, 1966] Bellman, R. (1966). Dynamic Programming. *Science*, 153(3731):34–37.
- [Bennett et al., 1992] Bennett, D. J., Hollerbach, J. M., Xu, Y., and Hunter, I. W. (1992). Time-varying stiffness of human elbow joint during cyclic voluntary movement. *Experimental Brain Research*, 88(2):433–442.

- [Bennett, 1948] Bennett, W. R. (1948). Spectra of Quantized Signals. *Bell System Technical Journal*, 27(3):446–472.
- [Bernstein, 1967] Bernstein, N. (1967). The coordination and regulation of movements.
- [Bernstein, 2014] Bernstein, N. A. (2014). *Dexterity and Its Development*.
- [Berthier, 1996] Berthier, N. E. (1996). Learning to reach: A mathematical model. *Developmental Psychology*, 32(5):811–823.
- [Bertram, 2007] Bertram, D. (2007). Likert scales. *academia.edu*.
- [Bicchi et al., 2011] Bicchi, A., Gabbicini, M., and Santello, M. (2011). Modelling natural and artificial hands with synergies. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1581):3153–3161.
- [Bicchi and Kumar, 2000] Bicchi, A. and Kumar, V. (2000). Robotic grasping and contact: A review. *Proceedings-IEEE International Conference on Robotics and Automation*, 1:348–353.
- [Bidet-Ildei et al., 2006] Bidet-Ildei, C., Orliaguet, J.-P., Sokolov, A. N., and Pavlova, M. (2006). Perception of elliptic biological motion. *Perception*, 35(8):1137–47.
- [Billard and Kragic, 2019] Billard, A. and Kragic, D. (2019). Trends and challenges in robot manipulation. *Science*, 364(6446).
- [Bizzi et al., 1984] Bizzi, E., Accornero, N., Chapple, W., and Hogan, N. (1984). Posture control and trajectory formation during arm movement. *Journal of Neuroscience*, 4(11):2738–2744.
- [Blake and Shiffrar, 2007] Blake, R. and Shiffrar, M. (2007). Perception of Human Motion. *Annual Review of Psychology*, 58(1):47–73.
- [Branco et al., 2019] Branco, M. P., de Boer, L. M., Ramsey, N. F., and Vansteensel, M. J. (2019). Encoding of kinetic and kinematic movement parameters in the sensorimotor cortex: A Brain-Computer Interface perspective.
- [Brohan et al., 2022] Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Dabis, J., Finn, C., Gopalakrishnan, K., Hausman, K., Herzog, A., Hsu, J., et al. (2022). Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*.
- [Brown and Asada, 2007] Brown, C. Y. and Asada, H. H. (2007). Inter-finger coordination and postural synergies in robot hands via mechanical implementation of principal components analysis. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2877–2882. IEEE.

- [Brown, 1911] Brown, T. G. (1911). The intrinsic factors in the act of progression in the mammal. *Proceedings of the Royal Society of London. Series B, Containing Papers of a Biological Character*, 84(572):308–319.
- [Brown, 1914] Brown, T. G. (1914). On the nature of the fundamental activity of the nervous centres; together with an analysis of the conditioning of rhythmic activity in progression, and a theory of the evolution of function in the nervous system. *The Journal of Physiology*, 48(1):18–46.
- [Burdet and Milner, 1998] Burdet, E. and Milner, T. E. (1998). Quantization of human motions and learning of accurate movements. *Biological cybernetics*, 78(4):307–318.
- [Burdet et al., 2001] Burdet, E., Osu, R., Franklin, D. W., Milner, T. E., and Kawato, M. (2001). The central nervous system stabilizes unstable dynamics by learning optimal impedance. 414(6862):446–449.
- [Calabrò et al., 2020] Calabrò, R. S., Müller-Eising, C., Diliberti, M. L., Manuli, A., Parrinello, F., Rao, G., Barone, V., and Civello, T. (2020). Who will pay for robotic rehabilitation? The growing need for a cost-effectiveness analysis. *Innovations in Clinical Neuroscience*, 17(10-12).
- [Calvo-Merino et al., 2006] Calvo-Merino, B., Grèzes, J., Glaser, D. E., Passingham, R. E., and Haggard, P. (2006). Seeing or Doing? Influence of Visual and Motor Familiarity in Action Observation. *Current Biology*, 6(19):1905–10.
- [Campos and Calado, 2009] Campos, F. M. and Calado, J. M. (2009). Approaches to human arm movement control—A review. *Annual Reviews in Control*, 33(1):69–77.
- [Casadio et al., 2015] Casadio, M., Pressman, A., and Mussa-Ivaldi, F. A. (2015). Learning to push and learning to move: The adaptive control of contact forces. *Frontiers in Computational Neuroscience*, 9(November):1–17.
- [Casile and Giese, 2006] Casile, A. and Giese, M. A. (2006). Nonvisual motor training influences biological motion perception. *Current Biology*, 16(1):69–74.
- [Catalano et al., 2014] Catalano, M. G., Grioli, G., Farnioli, E., Serio, A., Piazza, C., and Bicchi, A. (2014). Adaptive synergies for the design and control of the pisa/iit soft-hand. *International Journal of Robotics Research*, 33:768–782.
- [Chavan-Daffe et al., 2019] Chavan-Daffe, N., Holladay, R., and Rodriguez, A. (2019). Planar in-hand manipulation via motion cones. <https://doi.org/10.1177/0278364919880257>, 39(2-3):163–182.
- [Chen et al., 1997] Chen, R., Gerloff, C., Hallett, M., and Cohen, L. G. (1997). Involvement of the ipsilateral motor cortex in finger movements of different complexities. *Annals of Neurology*, 41:247–254.

- [Cheung et al., 2009] Cheung, V. C., Piron, L., Agostini, M., Silvoni, S., Turolla, A., and Bizzi, E. (2009). Stability of muscle synergies for voluntary actions after cortical stroke in humans. *Proceedings of the National Academy of Sciences of the United States of America*, 106(46).
- [Cheung and Seki, 2021] Cheung, V. C. and Seki, K. (2021). Approaches to revealing the neural basis of muscle synergies: A review and a critique. *Journal of Neurophysiology*, 125(5):1580–1597.
- [Chi et al., 2022] Chi, C., Burchfiel, B., Cousineau, E., Feng, S., and Song, S. (2022). Iterative Residual Policy: for Goal-Conditioned Dynamic Manipulation of Deformable Objects.
- [Chi et al., 2023] Chi, C., Feng, S., Du, Y., Xu, Z., Cousineau, E., Burchfiel, B., and Song, S. (2023). Diffusion policy: Visuomotor policy learning via action diffusion. *arXiv preprint arXiv:2303.04137*.
- [Chi et al., 2024] Chi, C., Xu, Z., Pan, C., Cousineau, E., Burchfiel, B., Feng, S., Tedrake, R., and Song, S. (2024). Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots. *arXiv preprint arXiv:2402.10329*.
- [Chib et al., 2009] Chib, V. S., Krutky, M. A., Lynch, K. M., and Mussa-Ivaldi, F. A. (2009). The separate neural control of hand movements and contact forces. *Journal of Neuroscience*, 29(12):3939–3947.
- [Chopin and Palmer, 1995] Chopin, F. and Palmer, W. (1995). *Chopin*. Alfred Publishing Company, Van Nuys, CA.
- [Ciocarlie and Allen, 2009] Ciocarlie, M. T. and Allen, P. K. (2009). Hand Posture Subspaces for Dexterous Robotic Grasping. *The International Journal of Robotics Research*, 28(7):851–867.
- [Correll et al., 2018] Correll, N., Bekris, K. E., Berenson, D., Brock, O., Causo, A., Hauser, K., Okada, K., Rodriguez, A., Romano, J. M., and Wurman, P. R. (2018). Analysis and observations from the first Amazon picking challenge. *IEEE Transactions on Automation Science and Engineering*, 15(1):172–188.
- [Dafle et al., 2014] Dafle, N. C., Rodriguez, A., Paolini, R., Tang, B., Srinivasa, S. S., Erdmann, M., Mason, M. T., Lundberg, I., Staab, H., and Fuhlbrigge, T. (2014). Extrinsic dexterity: In-hand manipulation with external forces. *Proceedings - IEEE International Conference on Robotics and Automation*, pages 1578–1585.
- [D’avella and Tresch, 2001] D’avella, A. and Tresch, M. C. (2001). Modularity in the motor system: decomposition of muscle patterns as combinations of time-varying synergies. *Advances in Neural Information Processing Systems*, 14.
- [Dayan et al., 2007] Dayan, E., Casile, A., Levit-Binnun, N., Giese, M. A., Hendler, T., and Flash, T. (2007). Neural representations of kinematic laws of motion:

- Evidence for action-perception coupling. *Proceedings of the National Academy of Sciences of the United States of America*, 104(51):20582–20587.
- [Degallier and Ijspeert, 2010] Degallier, S. and Ijspeert, A. (2010). Modeling discrete and rhythmic movements through motor primitives: A review.
- [di Luzio et al., 2022] di Luzio, F. S., Cordella, F., Bravi, M., Santacaterina, F., Bressi, F., Sterzi, S., and Zollo, L. (2022). Modification of Hand Muscular Synergies in Stroke Patients after Robot-Aided Rehabilitation. *Applied Sciences 2022, Vol. 12, Page 3146*, 12(6):3146.
- [Diogo et al., 2012] Diogo, R., Richmond, B. G., and Wood, B. (2012). Evolution and homologies of primate and modern human hand and forearm muscles, with notes on thumb movements and tool use. *Journal of Human Evolution*, 63(1):64–78.
- [Dominici et al., 2011] Dominici, N., Ivanenko, Y. P., Cappellini, G., D’Avella, A., Mondì, V., Cicchese, M., Fabiano, A., Silei, T., Di Paolo, A., Giannini, C., Poppele, R. E., and Lacquaniti, F. (2011). Locomotor primitives in newborn babies and their development. *Science*, 334(6058):997–999.
- [Dugas, 1988] Dugas, R. (1988). A History of Mechanics.
- [Elliott et al., 2001] Elliott, D., Chua, R., and Helsen, W. F. (2001). A century later: Woodworth’s (1899) two-component model of goal-directed aiming. *Psychological Bulletin*, 127(3):342–357.
- [Eppner et al., 2018] Eppner, C., Höfer, S., Jonschkowski, R., Martín-Martín, R., Sieverling, A., Wall, V., and Brock, O. (2018). Four aspects of building robotic systems: lessons from the Amazon Picking Challenge 2015. *Autonomous Robots*, 42(7):1459–1475.
- [Fani et al., 2016] Fani, S., Bianchi, M., Jain, S., Neto, J. S. P., Boege, S., Grioli, G., Bicchi, A., and Santello, M. (2016). Assessment of myoelectric controller performance and kinematic behavior of a novel soft synergy-inspired robotic hand for prosthetic applications. *Frontiers in Neurorobotics*, 10.
- [Fasoli et al., 2004] Fasoli, S. E., Krebs, H. I., and Hogan, N. (2004). Robotic technology and stroke rehabilitation: Translating research into practice.
- [Flanagan and Rao, 1995] Flanagan, J. R. and Rao, A. K. (1995). Trajectory adaptation to a nonlinear visuomotor transformation: Evidence of motion planning in visually perceived space. *Journal of Neurophysiology*, 74(5):2174–2178.
- [Flash and Henis, 1991] Flash, T. and Henis, E. (1991). Arm trajectory modifications during reaching towards visual targets. *Journal of Cognitive Neuroscience*, 3(3):220–230.
- [Flash and Hochner, 2005] Flash, T. and Hochner, B. (2005). Motor primitives in vertebrates and invertebrates.

- [Flash and Hogan, 1985] Flash, T. and Hogan, N. (1985). The coordination of arm movements: an experimentally confirmed mathematical model. *The Journal of neuroscience*, 5(7):1688–1703.
- [Fraisse, 1984] Fraisse, P. (1984). Perception and Estimation of Time. *Annual Review of Psychology*, 35(1):1–37.
- [Franklin et al., 2007] Franklin, D. W., Liaw, G., Milner, T. E., Osu, R., Burdet, E., and Kawato, M. (2007). Endpoint stiffness of the arm is directionally tuned to instability in the environment. *Journal of Neuroscience*, 27(29):7705–7716.
- [Frolov et al., 2018] Frolov, A. A., Kozlovskaya, I. B., Biryukova, E. V., and Bobrov, P. D. (2018). Use of Robotic Devices in Post-Stroke Rehabilitation. *Neuroscience and Behavioral Physiology*, 48(9):1053–1066.
- [Fu and Santello, 2018] Fu, Q. and Santello, M. (2018). Improving fine control of grasping force during hand-object interactions for a soft synergy-inspired myoelectric prosthetic hand. *Frontiers in Neurorobotics*, 11(JAN).
- [Fu et al., 2024] Fu, Z., Zhao, T. Z., and Finn, C. (2024). Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. *arXiv preprint arXiv:2401.02117*.
- [Furui et al., 2019] Furui, A., Eto, S., Nakagaki, K., Shimada, K., Nakamura, G., Masuda, A., Chin, T., and Tsuji, T. (2019). A myoelectric prosthetic hand with muscle synergy-based motion determination and impedance model-based biomimetic control. *Science Robotics*, 4(31):467.
- [Furuya et al., 2011] Furuya, S., Flanders, M., and Soechting, J. F. (2011). Hand kinematics of piano playing. *Journal of Neurophysiology*, 106(6):2849–2864.
- [Gabbicini et al., 2011] Gabbicini, M., Bicchi, A., Prattichizzo, D., and Malvezzi, M. (2011). On the role of hand synergies in the optimal choice of grasping forces. *Autonomous Robots*, 31(2-3):235–252.
- [Gabbicini et al., 2013] Gabbicini, M., Stillfried, G., Marino, H., and Bianchi, M. (2013). A data-driven kinematic model of the human hand with soft-tissue artifact compensation mechanism for grasp synergy analysis. In *IEEE International Conference on Intelligent Robots and Systems*, pages 3738–3745.
- [Gallese and Goldman, 1998] Gallese, V. and Goldman, A. (1998). Mirror neurons and the simulation theory of mind-reading.
- [Georgopoulos et al., 1992] Georgopoulos, A. P., Ashe, J., Smyrnis, N., and Taira, M. (1992). The Motor Cortex and the Coding of Force. *Science*, 256(5064):1692–1695.
- [Gioioso et al., 2013] Gioioso, G., Salvietti, G., Malvezzi, M., and Prattichizzo, D. (2013). Mapping Synergies From Human to Robotic Hands With Dissimilar Kinematics: An Approach in the Object Domain. *IEEE Transactions on Robotics*, 29(4):825–837.

- [Golub et al., 2018] Golub, M. D., Sadtler, P. T., Oby, E. R., Quick, K. M., Ryu, S. I., Tyler-Kabara, E. C., Batista, A. P., Chase, S. M., and Yu, B. M. (2018). Learning by neural reassociation. *Nature Neuroscience*, 21(4).
- [Grillner and Wallen, 1985] Grillner, S. and Wallen, P. (1985). Central Pattern Generators for Locomotion, with Special Reference to Vertebrates. *Annual Review of Neuroscience*, 8(1):233–261.
- [Grush, 2004] Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27(3):377–396.
- [Gulletta et al., 2020] Gulletta, G., Erlhagen, W., and Bicho, E. (2020). Human-like arm motion generation: A review.
- [Hamel-Pâquet et al., 2006] Hamel-Pâquet, C., Sergio, L. E., and Kalaska, J. F. (2006). Parietal Area 5 Activity Does Not Reflect the Differential Time-Course of Motor Output Kinetics During Arm-Reaching and Isometric-Force Tasks. <https://doi.org/10.1152/jn.00789.2005>, 95(6):3353–3370.
- [Haruno et al., 2012] Haruno, M., Ganesh, G., Burdet, E., and Kawato, M. (2012). Differential neural correlates of reciprocal activation and cocontraction control in dorsal and ventral premotor cortices. *Journal of Neurophysiology*, 107(1):126–133.
- [He and Ciocarlie, 2021] He, Z. and Ciocarlie, M. (2021). Discovering Synergies for Robot Manipulation with Multi-Task Reinforcement Learning.
- [He and Ciocarlie, 2022] He, Z. and Ciocarlie, M. (2022). Discovering synergies for robot manipulation with multi-task reinforcement learning. *Proceedings - IEEE International Conference on Robotics and Automation*, pages 2714–2721.
- [Hecht et al., 2001] Hecht, H., Vogt, S., and Prinz, W. (2001). Motor learning enhances perceptual judgment: A case for action-perception transfer. *Psychological Research*, 65(1):3–14.
- [Herguedas et al., 2019] Herguedas, R., López-Nicolás, G., Aragiúes, R., and Sagiúes, C. (2019). Survey on multi-robot manipulation of deformable objects. In *IEEE International Conference on Emerging Technologies and Factory Automation, ETFA*, volume 2019-Septe, pages 977–984. Institute of Electrical and Electronics Engineers Inc.
- [Hermus et al., 2021] Hermus, J., Lachner, J., Verdi, D., and Hogan, N. (2021). Exploiting Redundancy to Facilitate Physical Interaction. *IEEE Transactions on Robotics*, pages 1–17.
- [Hermus et al., 2020a] Hermus, J., Sternad, D., and Hogan, N. (2020a). Evidence for Dynamic Primitives in a Constrained Motion Task.

- [Hermus et al., 2020b] Hermus, J. R., Doeringer, J., Sternad, D., and Hogan, N. (2020b). Separating Neural Influences from Peripheral Mechanics: The Speed-Curvature Relation in Mechanically-Constrained Actions. *Journal of Neurophysiology*.
- [Hernandez et al., 2017] Hernandez, C., Bharatheesha, M., Ko, W., Gaiser, H., Tan, J., van Deurzen, K., de Vries, M., Van Mil, B., van Egmond, J., Burger, R., Morariu, M., Ju, J., Gerrmann, X., Ensing, R., Van Frankenhuyzen, J., and Wisse, M. (2017). Team delft’s robot winner of the amazon picking challenge 2016. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9776 LNAI:613–624.
- [Hirai and Hiraki, 2006] Hirai, M. and Hiraki, K. (2006). The relative importance of spatial versus temporal structure in the perception of biological motion: An event-related potential study. *Cognition*, 99(1):B15–B29.
- [Hoffer and Andreassen, 1981] Hoffer, J. and Andreassen, S. (1981). Regulation of soleus muscle stiffness in premammillary cats: intrinsic and reflex components. *Journal of neurophysiology*, 45(2):267–285.
- [Hogan et al., 2018] Hogan, F. R., Bauza, M., Canal, O., Donlon, E., and Rodriguez, A. (2018). Tactile Regrasp: Grasp Adjustments via Simulated Tactile Transformations. *IEEE International Conference on Intelligent Robots and Systems*, pages 2963–2970.
- [Hogan, 1985a] Hogan, N. (1985a). Impedance control: An approach to manipulation: Part I-theory. *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME*, 107(1):1–7.
- [Hogan, 1985b] Hogan, N. (1985b). Impedance control: An approach to manipulation: Part II-implementation. *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME*, 107(1):8–16.
- [Hogan, 1985c] Hogan, N. (1985c). Impedance control: An approach to manipulation: Part III-applications. *Journal of Dynamic Systems, Measurement and Control, Transactions of the ASME*, 107(1):17–24.
- [Hogan, 1990] Hogan, N. (1990). Mechanical impedance of single-and multi-articular systems. In *Multiple muscle systems: Biomechanics and movement organization*, pages 149–164. Springer.
- [Hogan, 2017] Hogan, N. (2017). Physical interaction via dynamic primitives. In Laumond, J., Mansard, N., and Lasserre, J., editors, *Geometric and Numerical Foundations of Movements.*, pages 269–299. Springer International Publishing AG.
- [Hogan and Buerger, 2018] Hogan, N. and Buerger, S. P. (2018). Impedance and interaction control. In *Robotics and automation handbook*, pages 375–398. CRC press.

- [Hogan and Sternad, 2012] Hogan, N. and Sternad, D. (2012). Dynamic primitives of motor behavior.
- [Hopcroft et al., 1991] Hopcroft, J. E., Kearney, J. K., and Kraftt, D. B. (1991). A Case Study of Flexible Object Manipulation. *The International Journal of Robotics Research*, 10(1):41–50.
- [Houle et al., 2010] Houle, M. E., Kriegel, H.-P., Kröger, P., Schubert, E., and Zimek, A. (2010). Can shared-neighbor distances defeat the curse of dimensionality? In *Scientific and Statistical Database Management: 22nd International Conference, SSDBM 2010, Heidelberg, Germany, June 30–July 2, 2010. Proceedings 22*, pages 482–500. Springer.
- [Huber et al., 2017] Huber, M. E., Folinus, C., and Hogan, N. (2017). Visual perception of limb stiffness. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3049–3055. IEEE.
- [Huber et al., 2019] Huber, M. E., Folinus, C., and Hogan, N. (2019). Visual perception of joint stiffness from multijoint motion. *Journal of neurophysiology*, 122(1):51–59.
- [Huh and Sejnowski, 2015] Huh, D. and Sejnowski, T. J. (2015). Spectrum of power laws for curved hand movements. *Proceedings of the National Academy of Sciences of the United States of America*, 112(29):E3950–E3958.
- [Hussein et al., 2017] Hussein, A., Gaber, M. M., Elyan, E., and Jayne, C. (2017). Imitation learning: A survey of learning methods. *ACM Computing Surveys (CSUR)*, 50(2):1–35.
- [Iandolo et al., 2019] Iandolo, R., Marini, F., Semprini, M., Laffranchi, M., Mugnosso, M., Cherif, A., De Michieli, L., Chiappalone, M., and Zenzeri, J. (2019). Perspectives and challenges in robotic neurorehabilitation.
- [Ijspeert et al., 2002] Ijspeert, A. J., Nakanishi, J., and Schaal, S. (2002). Movement imitation with nonlinear dynamical systems in humanoid robots. *Proceedings - IEEE International Conference on Robotics and Automation*, 2:1398–1403.
- [Jacobsen et al., 1986] Jacobsen, S., Iversen, E., Knutti, D., Johnson, R., and Biggers, K. (1986). Design of the utah/m.i.t. dextrous hand. pages 1520–1532. Institute of Electrical and Electronics Engineers (IEEE).
- [James, 1890] James, W. (1890). *The principles of psychology*, volume 2. Henry Holt and Company, NY.
- [Jarque-Bou et al., 2019] Jarque-Bou, N. J., Scano, A., Atzori, M., and Müller, H. (2019). Kinematic synergies of hand grasps: A comprehensive study on a large publicly available dataset. *Journal of NeuroEngineering and Rehabilitation*, 16(1):1–14.

- [Kandel et al., 2000a] Kandel, E., Schwartz, J., Jessell, T., Siegelbaum, S., and Hudspeth, A. (2000a). *Principles of neural science*. The McGraw-Hill Companies, fifth edition.
- [Kandel et al., 2013] Kandel, E., Schwartz, J., Jessell, T., Siegelbaum, S., and Hudspeth, A. J. (2013). *Principles of neural science*. Principles of Neural Science. McGraw-Hill, 5th ed. edition.
- [Kandel et al., 2000b] Kandel, S., Orliaguet, J. P., and Viviani, P. (2000b). Perceptual anticipation in handwriting: The role of implicit motor competence. *Perception and Psychophysics*, 62(4):706–716.
- [Kargo and Giszter, 2008] Kargo, W. J. and Giszter, S. F. (2008). Individual premotor drive pulses, not time-varying synergies, are the units of adjustment for limb trajectories constructed in spinal cord. *Journal of Neuroscience*, 28(10):2409–2425.
- [Kenward et al., 2005] Kenward, B., Weir, A. A., Rutz, C., and Kacelnik, A. (2005). Behavioural ecology: Tool manufacture by naive juvenile crows. *Nature*, 433(7022):121.
- [Kepler, 1609] Kepler, J. (1609). *Astronomia Nova*.
- [Khandate et al., 2021] Khandate, G., Haas-Heger, M., and Ciocarlie, M. (2021). On the Feasibility of Learning Finger-gaiting In-hand Manipulation with Intrinsic Sensing. *2022 International Conference on Robotics and Automation (ICRA)*, pages 2752–2758.
- [Kieliba et al., 2018] Kieliba, P., Tropea, P., Pirondini, E., Coscia, M., Micera, S., and Artoni, F. (2018). How are Muscle Synergies Affected by Electromyography Pre-Processing? *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 26(4):882–893.
- [Klopsteg et al., 1955] Klopsteg, P. E., Alldredge, R. H., Eberhart, H. D., Mcmath, R. R., Mitchell, C. L., Taylor, C. L., Thompson, T. C., Wilson, P. D., Allen, R. S., Inman, V. T., Wilson, A. B., and Fleer, B. (1955). Artificial limbs advisory committee on artificial limbs executive secretary.
- [Knoedler et al., 2015] Knoedler, K., Dimitrov, V., Conn, D., Gennert, M. A., and Padir, T. (2015). Towards supervisory control of humanoid robots for driving vehicles during disaster response missions. In *2015 IEEE International Conference on Technologies for Practical Robot Applications (TePRA)*, pages 1–6. IEEE.
- [Koeppen et al., 2017] Koeppen, R., Huber, M., Sternad, D., and Hogan, V. (2017). Controlling physical interactions: Humans do not minimize muscle effort. *ASME 2017 Dynamic Systems and Control Conference*.

- [Kolesnikov et al., 2011] Kolesnikov, M., Piovesan, D., Lynch, K. M., and Mussa-Ivaldi, F. A. (2011). On force regulation strategies in predictable environments. *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, pages 4076–4081.
- [Krebs et al., 1999] Krebs, H. I., Aisen, M. L., Volpe, B. T., and Hogan, N. (1999). Quantization of continuous arm movements in humans with brain injury. *Proceedings of the National Academy of Sciences of the United States of America*, 96(8):4645–4649.
- [Krebs et al., 2003] Krebs, H. I., Palazzolo, J. J., Dipietro, L., Ferraro, M., Krol, J., Rannekleiv, K., Volpe, B. T., and Hogan, N. (2003). Rehabilitation robotics: Performance-based progressive robot-assisted therapy. *Autonomous Robots*, 15(1):7–20.
- [Krings et al., 2000] Krings, T., Töpper, R., Foltys, H., Erberich, S., Sparing, R., Willmes, K., and Thron, A. (2000). Cortical activation patterns during complex motor tasks in piano players and control subjects. a functional magnetic resonance imaging study. *Neuroscience Letters*, 278:189–193.
- [Kroemer et al., 2021] Kroemer, O., Niekum, S., and Konidaris, G. (2021). A review of robot learning for manipulation. *The Journal of Machine Learning Research*, 22:1–82.
- [Kumar et al., 2016] Kumar, V., Todorov, E., and Levine, S. (2016). Optimal control with learned local models: Application to dexterous manipulation. *Proceedings - IEEE International Conference on Robotics and Automation*, 2016-June:378–383.
- [Lachner et al., 2023] Lachner, J., Nah, M. C., Tessari, F., and Hogan, N. (2023). Elementary dynamic actions: key structures for contact-rich manipulation. In *IROS 2023 Workshop on Leveraging Models for Contact-Rich Manipulation*.
- [Lacquaniti et al., 1982] Lacquaniti, F., Licata, F., and Soechting, J. F. (1982). The mechanical behavior of the human forearm in response to transient perturbations. *Biological Cybernetics*, 44(1):35–46.
- [Laffranchi et al., 2020] Laffranchi, M., Boccardo, N., Traverso, S., Lombardi, L., Canepa, M., Lince, A., Semprini, M., Saglia, J. A., Naceri, A., Sacchetti, R., Gruppioni, E., and De Michieli, L. (2020). The Hannes hand prosthesis replicates the key biological properties of the human hand. *Science Robotics*, 5(46):467.
- [Lahitani et al., 2016] Lahitani, A. R., Permanasari, A. E., and Setiawan, N. A. (2016). Cosine similarity to determine similarity measure: Study case in online essay assessment. In *2016 4th International conference on cyber and IT service management*, pages 1–6. IEEE.
- [Lambert-Shirzad and Van Der Loos, 2017] Lambert-Shirzad, N. and Van Der Loos, H. F. (2017). Data sample size needed for analysis of kinematic and muscle

- synergies in healthy and stroke populations. *IEEE International Conference on Rehabilitation Robotics*, pages 777–782.
- [Li et al., 2017] Li, B., Li, G., Sun, Y., Jiang, G., Kong, J., and Jiang, D. (2017). A review of rehabilitation robot. *Proceedings - 2017 32nd Youth Academic Annual Conference of Chinese Association of Automation, YAC 2017*, pages 907–911.
- [Li et al., 2014] Li, S., Sheng, X., Liu, H., and Zhu, X. (2014). Design of a myoelectric prosthetic hand implementing postural synergy mechanically. *Industrial Robot*, 41(5):447–455.
- [Li et al., 2018] Li, Y., Ganesh, G., Jarrasse, N., Haddadin, S., Albu-Schaeffer, A., and Burdet, E. (2018). Force, impedance, and trajectory learning for contact tooling and haptic identification. *IEEE Transactions on Robotics*, 34(5):1170–1182.
- [Linde et al., 2002] Linde, R. V. D., Lammertse, P., Frederiksen, E., and B (2002). The HapticMaster, a new high-performance haptic interface. *Proc. EuroHaptic, Edinburgh, UK*, pages 1–5.
- [Lipps et al., 2020] Lipps, D. B., Baillargeon, E. M., Ludvig, D., and Perreault, E. J. (2020). Quantifying the Multidimensional Impedance of the Shoulder During Volitional Contractions. *Annals of Biomedical Engineering*, 48(9):2354–2369.
- [Liu et al., 2022] Liu, C., Lu, J., Yang, H., and Guo, K. (2022). Current State of Robotics in Hand Rehabilitation after Stroke: A Systematic Review.
- [Lloyd, 1982] Lloyd, S. P. (1982). Least Squares Quantization in PCM. *IEEE Transactions on Information Theory*, 28(2).
- [Loula et al., 2005] Loula, F., Prasad, S., Harber, K., and Shiffrar, M. (2005). Recognizing people from their movement. *Journal of Experimental Psychology: Human Perception and Performance*, 31(1):210–220.
- [Luo et al., 2018] Luo, C., Zhan, J., Xue, X., Wang, L., Ren, R., and Yang, Q. (2018). Cosine normalization: Using cosine similarity instead of dot product in neural networks. In *Artificial Neural Networks and Machine Learning–ICANN 2018: 27th International Conference on Artificial Neural Networks, Rhodes, Greece, October 4–7, 2018, Proceedings, Part I 27*, pages 382–391. Springer.
- [Maciejasz et al., 2014] Maciejasz, P., Eschweiler, J., Gerlach-Hahn, K., Jansen-Troy, A., and Leonhardt, S. (2014). A survey on robotic devices for upper limb rehabilitation.
- [Mason et al., 2001] Mason, C. R., Gomez, J. E., and Ebner, T. J. (2001). Hand Synergies During Reach-to-Grasp. *Journal of Neurophysiology*, 86(6):2896–2910.
- [Mason, 1981] Mason, M. T. (1981). Compliance and Force Control for Computer Controlled Manipulators. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(6):418–432.

- [Mason, 2018] Mason, M. T. (2018). Toward Robotic Manipulation. <https://doi.org/10.1146/annurev-control-060117-104848>, 1(1):1–28.
- [Massey et al., 1992] Massey, J. T., Lurito, J. T., Pellizzer, G., and Georgopoulos, A. P. (1992). Three-dimensional drawings in isometric conditions: relation between geometry and kinematics. *Experimental Brain Research* 1992 88:3, 88(3):685–690.
- [Matrone et al., 2012] Matrone, G. C., Cipriani, C., Carrozza, M. C., and Magenes, G. (2012). Real-time myoelectric control of a multi-fingered hand prosthesis using principal components analysis. *Journal of NeuroEngineering and Rehabilitation*, 9(1):1–13.
- [Mattar and Gribble, 2005] Mattar, A. A. and Gribble, P. L. (2005). Motor learning by observing. *Neuron*, 46(1):153–160.
- [Maurice et al., 2018] Maurice, P., Huber, M. E., Hogan, N., and Sternad, D. (2018). Velocity-Curvature Patterns Limit Human-Robot Physical Interaction. *IEEE Robotics and Automation Letters*, 3(1):249–256.
- [McDonald et al., 2020] McDonald, C. L., Westcott-McCoy, S., Weaver, M. R., Haagsma, J., and Kartin, D. (2020). Global prevalence of traumatic non-fatal limb amputation. *Prosthetics and Orthotics International*.
- [Meltzoff and Moore, 1997] Meltzoff, A. N. and Moore, M. K. (1997). Explaining facial imitation: A theoretical model. *Infant and Child Development*, 6(3-4):179–192.
- [Meltzoff, 2002] Meltzoff, A. N. W. P. (2002). *The Imitative Mind: Development, Evolution and Brain Bases (Cambridge Studies in Cognitive and Perceptual Development)*. Cambridge University Press.
- [Meyer et al., 1988] Meyer, D. E., Abrams, R. A., Kornblum, S., Wright, C. E., and Smith, J. E. K. (1988). Optimality in human motor performance: Ideal control of rapid aimed movements. *Psychological Review*, 95(3):340–370.
- [Milner and Franklin, 2005] Milner, T. E. and Franklin, D. W. (2005). Impedance control and internal model use during the initial stage of adaptation to novel dynamics in humans. *The Journal of Physiology*, 567(2):651–664.
- [Moran and Schwartz, 1999] Moran, D. W. and Schwartz, A. B. (1999). Motor cortical representation of speed and direction during reaching. *Journal of Neurophysiology*, 82(5):2676–2692.
- [Morewedge and Kahneman, 2010] Morewedge, C. K. and Kahneman, D. (2010). Associative processes in intuitive judgment. *Trends in Cognitive Sciences*, 14(10):435–440.

- [Morgan et al., 2022] Morgan, A. S., Hang, K., Wen, B., Bekris, K., and Dollar, A. M. (2022). Complex In-Hand Manipulation Via Compliance-Enabled Finger Gaitting and Multi-Modal Planning. *IEEE Robotics and Automation Letters*, 7(2):4821–4828.
- [Murray et al., 1994] Murray, R. M., Li, Z., and Sastry, S. S. (1994). A mathematical introduction to robotic manipulation. *A Mathematical Introduction to Robotic Manipulation*, pages 1–456.
- [Mussa-Ivaldi et al., 1985] Mussa-Ivaldi, F. a., Hogan, N., and Bizzi, E. (1985). Neural, mechanical, and geometric factors subserving arm posture in humans. *Journal of Neuroscience*, 5(10):2732–2743.
- [Nagabandi et al., 2020] Nagabandi, A., Konolige, K., Levine, S., Kumar, V., and Brain, G. (2020). Deep dynamics models for learning dexterous manipulation.
- [Nah et al., 2020] Nah, M. C., Krotov, A., Russo, M., Sternad, D., and Hogan, N. (2020). Dynamic Primitives Facilitate Manipulating a Whip. *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics*, 2020-Novem:685–691.
- [Nah et al., 2021] Nah, M. C., Krotov, A., Russo, M., Sternad, D., and Hogan, N. (2021). Manipulating a Whip in 3D via Dynamic Primitives. In *IEEE International Conference on Intelligent Robots and Systems*, pages 2803–2808. Institute of Electrical and Electronics Engineers Inc.
- [Nah et al., 2023a] Nah, M. C., Krotov, A., Russo, M., Sternad, D., and Hogan, N. (2023a). Learning to manipulate a whip with simple primitive actions – A simulation study. *iScience*, 26(8).
- [Nah et al., 2023b] Nah, M. C., Lachner, J., Tessari, F., and Hogan, N. (2023b). Kinematic Modularity of Elementary Dynamic Actions.
- [Napier, 1956] Napier, J. R. (1956). THE PREHENSILE MOVEMENTS OF THE HUMAN HAND. *The Journal of Bone and Joint Surgery*, 38(4):902–913.
- [Nichols and Houk, 1976] Nichols, T. and Houk, J. (1976). Improvement in linearity and regulation of stiffness that results from actions of stretch reflex. *journal of Neurophysiology*, 39(1):119–142.
- [Ohta et al., 2004] Ohta, K., Svinin, M. M., Luo, Z., Hosoe, S., and Laboissière, R. (2004). Optimal trajectory formation of constrained human arm reaching movements. *Biological Cybernetics*, 91(1):23–36.
- [Osu et al., 2003] Osu, R., Burdet, E., Franklin, D. W., Milner, T. E., and Kawato, M. (2003). Different Mechanisms Involved in Adaptation to Stable and Unstable Dynamics. <https://doi.org/10.1152/jn.00073.2003>, 90(5):3255–3269.

- [Osu et al., 2004] Osu, R., Kamimura, N., Iwasaki, H., Nakano, E., Harris, C. M., Wada, Y., and Kawato, M. (2004). Optimal impedance control for task achievement in the presence of signal-dependent noise. *Journal of Neurophysiology*, 92(2):1199–1215.
- [Padalkar et al., 2023] Padalkar, A., Pooley, A., Jain, A., Bewley, A., Herzog, A., Irpan, A., Khazatsky, A., Rai, A., Singh, A., Brohan, A., et al. (2023). Open x-embodiment: Robotic learning datasets and rt-x models. *arXiv preprint arXiv:2310.08864*.
- [Park et al., 2017] Park, S. W., Marino, H., Charles, S. K., Sternad, D., and Hogan, N. (2017). Moving slowly is hard for humans: Limitations of dynamic primitives. *Journal of Neurophysiology*, 118(1):69–83.
- [Pearson, 1901] Pearson, K. (1901). Principal components analysis. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 6(2):559.
- [Peters and Schaal, 2006] Peters, J. and Schaal, S. (2006). Policy gradient methods for robotics. *IEEE International Conference on Intelligent Robots and Systems*, pages 2219–2225.
- [Piazza et al., 2019] Piazza, C., Grioli, G., Catalano, M., and Bicchi, A. (2019). A century of robotic hands. <https://doi.org/10.1146/annurev-control-060117-105003>, 2:1–32.
- [Piovesan et al., 2019] Piovesan, D., Kolesnikov, M., Lynch, K., and Mussa-Ivaldi, F. A. (2019). The Concurrent Control of Motion and Contact Force in the Presence of Predictable Disturbances. *Journal of Mechanisms and Robotics*, 11(6).
- [Pollick et al., 2001] Pollick, F. E., Paterson, H. M., Bruderlin, A., and Sanford, A. J. (2001). Perceiving affect from arm movement. *Cognition*, 82(2).
- [Pomerleau, 1988] Pomerleau, D. A. (1988). Alvin: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1.
- [Prochazka et al., 2000] Prochazka, A., Clarac, F., Loeb, G. E., Rothwell, J. C., and Wolpaw, J. R. (2000). What do reflex and voluntary mean? modern views on an ancient debate. *Experimental brain research*, 130:417–432.
- [Qassim and Wan Hasan, 2020] Qassim, H. M. and Wan Hasan, W. Z. (2020). A review on upper limb rehabilitation robots.
- [Qin et al., 2022] Qin, Y., Su, H., and Wang, X. (2022). From one hand to multiple hands: Imitation learning for dexterous manipulation from single-camera teleoperation.
- [Raibert and Craig, 1981] Raibert, M. H. and Craig, J. J. (1981). Hybrid position/force control of manipulators. *Journal of Dynamic Systems, Measurement, and Control*, 102(127):126–133.

- [Riener et al., 2005] Riener, R., Nef, T., and Colombo, G. (2005). Robot-aided neurorehabilitation of the upper extremities.
- [Rodgers et al., 2019] Rodgers, H., Bosomworth, H., Krebs, H. I., van Wijck, F., Howel, D., Wilson, N., Aird, L., Alvarado, N., Andole, S., Cohen, D. L., Dawson, J., Fernandez-Garcia, C., Finch, T., Ford, G. A., Francis, R., Hogg, S., Hughes, N., Price, C. I., Ternent, L., Turner, D. L., Vale, L., Wilkes, S., and Shaw, L. (2019). Robot assisted training for the upper limb after stroke (RATULS): a multicentre randomised controlled trial. *The Lancet*, 394(10192):51–62.
- [Rodriguez et al., 2012] Rodriguez, A., Mason, M. T., and Ferry, S. (2012). From caging to grasping. *International Journal of Robotics Research*, 31(7):886–900.
- [Russell and Hogan, 1989] Russell, D. and Hogan, N. (1989). Dealing with constraints: a biomechanical approach. In *Images of the Twenty-First Century. Proceedings of the Annual International Engineering in Medicine and Biology Society*, pages 892–893. IEEE.
- [Sadtler et al., 2014] Sadtler, P. T., Quick, K. M., Golub, M. D., Chase, S. M., Ryu, S. I., Tyler-Kabara, E. C., Yu, B. M., and Batista, A. P. (2014). Neural constraints on learning. *Nature*, 512(7515).
- [Sanchez et al., 2018] Sanchez, J., Corrales, J.-A., Bouzgarrou, B.-C., and Mezouar, Y. (2018). Robotic manipulation and sensing of deformable objects in domestic and industrial applications: a survey. *The International Journal of Robotics Research*, 37(7):688–716.
- [Santello et al., 2013] Santello, M., Baud-Bovy, G., and Jörntell, H. (2013). Neural bases of hand synergies. *Frontiers in Computational Neuroscience*, 0(MAR):23.
- [Santello et al., 2016] Santello, M., Bianchi, M., Gabiccini, M., Ricciardi, E., Salvietti, G., Prattichizzo, D., Ernst, M., Moscatelli, A., Jörntell, H., Kappers, A. M., Kyriakopoulos, K., Albu-Schäffer, A., Castellini, C., and Bicchi, A. (2016). Hand synergies: Integration of robotics and neuroscience for understanding the control of biological and artificial hands.
- [Santello et al., 1998] Santello, M., Flanders, M., and Soechting, J. F. (1998). Postural hand synergies for tool use. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 18(23):10105–15.
- [Santello et al., 2002] Santello, M., Flanders, M., and Soechting, J. F. (2002). Patterns of hand motion during grasping and the influence of sensory guidance. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 22(4):1426–35.
- [Sarlegna et al., 2010] Sarlegna, F. R., Malfait, N., Bringoux, L., Bourdin, C., and Vercher, J. L. (2010). Force-field adaptation without proprioception: Can vision be used to model limb dynamics? *Neuropsychologia*, 48(1):60–67.

- [Schaal, 2006] Schaal, S. (2006). Dynamic Movement Primitives -A Framework for Motor Control in Humans and Humanoid Robotics. *Adaptive Motion of Animals and Machines*, pages 261–280.
- [Schaal and Sternad, 2001] Schaal, S. and Sternad, D. (2001). Origins and violations of the 2/3 power law in rhythmic three-dimensional arm movements. *136(1)*:60–72.
- [Schaal et al., 2004] Schaal, S., Sternad, D., Osu, R., and Kawato, M. (2004). Rhythmic arm movement is not discrete. *Nature Neuroscience*, *7(10)*:1136–1143.
- [Schlesinger, 1919] Schlesinger, G. (1919). Der mechanische aufbau der künstlichen glieder. *Ersatzglieder und Arbeitshilfen*, pages 321–661.
- [Sergio and Kalaska, 1998] Sergio, L. E. and Kalaska, J. F. (1998). Changes in the Temporal Pattern of Primary Motor Cortex Activity in a Directional Isometric Force Versus Limb Movement Task. <https://doi.org/10.1152/jn.1998.80.3.1577>, *80(3)*:1577–1583.
- [Shadmehr and Mussa-Ivaldi, 1994] Shadmehr, R. and Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience*, *14(5)*:3208–3224.
- [Sharma, 2023] Sharma, N. (2023). The minimum intervention principle of optimal control relates the uncontrolled manifold to muscle synergies. *bioRxiv*, pages 2023–08.
- [Sierotowicz et al., 2022] Sierotowicz, M., Lotti, N., Nell, L., Missiroli, F., Alicea, R., Zhang, X., Xiloyannis, M., Rupp, R., Papp, E., Krzywinski, J., Castellini, C., and Masia, L. (2022). EMG-Driven Machine Learning Control of a Soft Glove for Grasping Assistance and Rehabilitation. *IEEE Robotics and Automation Letters*, *7(2)*:1566–1573.
- [Sing et al., 2009] Sing, G. C., Joiner, W. M., Nanayakkara, T., Brayonov, J. B., and Smith, M. A. (2009). Primitives for Motor Adaptation Reflect Correlated Neural Tuning to Position and Velocity. *Neuron*, *64(4)*:575–589.
- [Singh et al., 2018] Singh, R. E., Iqbal, K., White, G., and Hutchinson, T. E. (2018). A Systematic Review on Muscle Synergies: From Building Blocks of Motor Behavior to a Neurorehabilitation Tool. *Applied bionics and biomechanics*, 2018.
- [Slotine and Asada, 1992] Slotine, J.-J. E. and Asada, H. (1992). *Robot Analysis and Control*. John Wiley & Sons, Inc., New York, 1st edition.
- [Smith, 1988] Smith, L. (1988). A tutorial on Principal Components Analysis.
- [Sternad, 2008] Sternad, D. (2008). Towards a unified theory of rhythmic and discrete movements-behavioral, modeling and imaging results. *Understanding Complex Systems*, 2008:105–133.

- [Sternad et al., 2000] Sternad, D., Dean, W. J., and Schaal, S. (2000). Interaction of rhythmic and discrete pattern generators in single-joint movements. *Human Movement Science*, 19(4):627–664.
- [Stewart, 1993] Stewart, G. W. (1993). On the early history of the singular value decomposition. *SIAM Review*, 35(4):551–566.
- [Stulp et al., 2012] Stulp, F., Theodorou, E. A., and Schaal, S. (2012). Reinforcement learning with sequences of motion primitives for robust manipulation. *IEEE Transactions on Robotics*, 28(6):1360–1370.
- [Takahashi et al., 2001] Takahashi, C. D., Scheidt, R. A., and Reinkensmeyer, D. J. (2001). Impedance Control and Internal Model Formation When Reaching in a Randomly Varying Dynamical Environment. <https://doi.org/10.1152/jn.2001.86.2.1047>, 86(2):1047–1051.
- [Tessari et al., 2024] Tessari, F., West Jr, A. M., and Hogan, N. (2024). De coordinatione motus humani: The synergy expansion hypothesis. *bioRxiv*, pages 2024–04.
- [Thoroughman and Shadmehr, 2000] Thoroughman, K. A. and Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature*, 407(6805):742–747.
- [Ting and McKay, 2007] Ting, L. H. and McKay, J. L. (2007). Neuromechanics of muscle synergies for posture and movement. *Current Opinion in Neurobiology*, 17:622–628.
- [Todorov and Ghahramani, 2004] Todorov, E. and Ghahramani, Z. (2004). Analysis of the synergies underlying complex hand manipulation. *Annual International Conference of the IEEE Engineering in Medicine and Biology - Proceedings*, 26 VI:4637–4640.
- [Toffin et al., 2003] Toffin, D., McIntyre, J., Droulez, J., Kemeny, A., and Berthoz, A. (2003). Perception and reproduction of force direction in the horizontal plane. *Journal of neurophysiology*, 90(5):3040–3053.
- [Tresch et al., 2006] Tresch, M. C., Cheung, V. C., and D’Avella, A. (2006). Matrix factorization algorithms for the identification of muscle synergies: Evaluation on simulated and experimental data sets. *Journal of Neurophysiology*, 95(4):2199–2212.
- [Tsao et al., 2022] Tsao, C. W., Aday, A. W., Almarzooq, Z. I., Alonso, A., Beaton, A. Z., Bittencourt, M. S., Boehme, A. K., Buxton, A. E., Carson, A. P., Commodore-Mensah, Y., Elkind, M. S. V., Evenson, K. R., Eze-Nliam, C., Ferguson, J. F., Generoso, G., Ho, J. E., Kalani, R., Khan, S. S., Kissela, B. M., Knutson, K. L., Levine, D. A., Lewis, T., Liu, J., Loop, M. S., Ma, J., Mussolino, M. E., Navaneethan, S. D., Perak, A. M., Poudel, R., Rezk-Hanna, M., Roth,

- G. A., Schroeder, E. B., Shah, S. H., Thacker, E. L., VanWagner, L. B., Virani, S. S., Voecks, J. H., Wang, N.-Y., Yaffe, K., and Martin, S. S. (2022). Heart disease and stroke statistics-2022 update: A report from the american heart association. *Circulation*, 145(8):e153–e639.
- [Ueki et al., 2008] Ueki, S., Nishimoto, Y., Abe, M., Kawasaki, H., Ito, S., Ishigure, Y., Mizumoto, J., and Ojika, T. (2008). Development of virtual reality exercise of hand motion assist robot for rehabilitation therapy by patient self-motion control. In *Proceedings of the 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS'08 - "Personalized Healthcare through Technology"*, pages 4282–4285. IEEE Computer Society.
- [Varol et al., 2014] Varol, H. A., Dalley, S. A., Wiste, T. E., and Goldfarb, M. (2014). *Biomimicry and the Design of Multigrasp Transradial Protheses*, pages 431–451. Springer International Publishing, Cham.
- [Virani et al., 2020] Virani, S. S., Alonso, A., Benjamin, E. J., Bittencourt, M. S., Callaway, C. W., Carson, A. P., Chamberlain, A. M., Chang, A. R., Cheng, S., Delling, F. N., Djousse, L., Elkind, M. S., Ferguson, J. F., Fornage, M., Khan, S. S., Kissela, B. M., Knutson, K. L., Kwan, T. W., Lackland, D. T., Lewis, T. T., Lichtman, J. H., Longenecker, C. T., Loop, M. S., Lutsey, P. L., Martin, S. S., Matsushita, K., Moran, A. E., Mussolino, M. E., Perak, A. M., Rosamond, W. D., Roth, G. A., Sampson, U. K., Satou, G. M., Schroeder, E. B., Shah, S. H., Shay, C. M., Spartano, N. L., Stokes, A., Tirschwell, D. L., VanWagner, L. B., Tsao, C. W., Wong, S. S., and Heard, D. G. (2020). Heart disease and stroke statistics—2020 update: A report from the American Heart Association.
- [Viviani and Flash, 1995] Viviani, P. and Flash, T. (1995). Minimum-Jerk, Two-Thirds Power Law, and Isochrony: Converging Approaches to Movement Planning. *Journal of Experimental Psychology: Human Perception and Performance*, 21(1):32–53.
- [Viviani and Stucchi, 1992] Viviani, P. and Stucchi, N. (1992). Biological Movements Look Uniform: Evidence of Motor-Perceptual Interactions. *Journal of Experimental Psychology: Human Perception and Performance*, 18(3):603–623.
- [Von Hofsten, 1991] Von Hofsten, C. (1991). Structuring of early reaching movements: A longitudinal study. *Journal of Motor Behavior*, 23(4):280–292.
- [Walk and Homan, 1984] Walk, R. D. and Homan, C. P. (1984). Emotion and dance in dynamic light displays. *Bulletin of the Psychonomic Society*, 22(5):437–440.
- [Wang, 2023] Wang, M. (2023). Non-invasive vision-based measurement of hand kinematics and interaction.
- [Weir et al., 2002] Weir, A. A., Chappell, J., and Kacelnik, A. (2002). Shaping of hooks in new caledonian crows.

- [Weiss and Flanders, 2004] Weiss, E. J. and Flanders, M. (2004). Muscular and Postural Synergies of the Human Hand. *Journal of Neurophysiology*, 92(1):523–535.
- [West, 2020] West, A. M. (2020). Towards a non-invasive measurement of human motion, force, and impedance during a complex physical-interaction task : wire-harnessing.
- [West et al., 2022a] West, A. M., Hermus, J., Huber, M. E., Maurice, P., Sternad, D., and Hogan, N. (2022a). Dynamic Primitives Limit Human Force Regulation During Motion. *IEEE Robotics and Automation Letters*, 7(2):2391–2398.
- [West et al., 2022b] West, A. M., Huber, M. E., and Hogan, N. (2022b). Role of path information in visual perception of joint stiffness. *PLoS Computational Biology*, 18(11):e1010729.
- [West et al., 2023] West, A. M., Tessari, F., and Hogan, N. (2023). The Study of Complex Manipulation via Kinematic Hand Synergies: The Effects of Data Pre-Processing. *IEEE ... International Conference on Rehabilitation Robotics : [proceedings]*, 2023:1–6.
- [Wolpert et al., 2011] Wolpert, D. M., Diedrichsen, J., and Flanagan, J. R. (2011). Principles of sensorimotor learning.
- [Won and Hogan, 1995] Won, J. and Hogan, N. (1995). Stability properties of human reaching movements. 107:125–136.
- [Woodworth, 1899] Woodworth, R. S. (1899). The Accuracy of Voluntary Movements. *The Psychological Review*, 3(13):1–114.
- [Yamakawa et al., 2011] Yamakawa, Y., Namiki, A., and Ishikawa, M. (2011). Dynamic Manipulation of a Cloth by High-speed Robot System using High-speed Visual Feedback. *IFAC Proceedings Volumes*, 44(1):8076–8081.
- [Yang et al., 2019] Yang, N., An, Q., Kogami, H., Yamakawa, H., Tamura, Y., Takahashi, K., Kinomoto, M., Yamasaki, H., Itkonen, M., Shibata-Alnajjar, F., Shimoda, S., Hattori, N., Fujii, T., Otomune, H., Miyai, I., Yamashita, A., and Asama, H. (2019). Temporal Features of Muscle Synergies in Sit-to-Stand Motion Reflect the Motor Impairment of Post-Stroke Patients. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 27(10):2118–2127.
- [Yu et al., 2016] Yu, K.-T., Fazeli, N., Chavan-Daffe, N., Taylor, O., Donlon, E., Lankenau, G. D., and Rodriguez, A. (2016). A Summary of Team MIT’s Approach to the Amazon Picking Challenge 2015.
- [Yue et al., 2017] Yue, Z., Zhang, X., and Wang, J. (2017). Hand Rehabilitation Robotics on Poststroke Motor Recovery.

- [Zago et al., 2018] Zago, M., Matic, A., Flash, T., Gomez-Marin, A., and Lacquaniti, F. (2018). The speed-curvature power law of movements: a reappraisal. *Experimental Brain Research*, 236(1):69–82.
- [Zatsiorsky, 2002] Zatsiorsky, V. M. (2002). *Kinetics of human motion*. Human Kinetics, Champaign.
- [Zeng et al., 2018] Zeng, A., Song, S., Welker, S., Lee, J., Rodriguez, A., and Funkhouser, T. (2018). Learning Synergies between Pushing and Grasping with Self-Supervised Deep Reinforcement Learning. In *IEEE International Conference on Intelligent Robots and Systems*.
- [Zhang et al., 2020] Zhang, F., Bazarevsky, V., Vakunov, A., Tkachenka, A., Sung, G., Chang, C.-L., and Grundmann, M. (2020). MediaPipe Hands: On-device Real-time Hand Tracking.
- [Zhang et al., 2021] Zhang, H., Ichnowski, J., Seita, D., Wang, J., Huang, H., and Goldberg, K. (2021). Robots of the Lost Arc: Self-Supervised Learning to Dynamically Manipulate Fixed-Endpoint Cables. *Proceedings - IEEE International Conference on Robotics and Automation*, 2021-May:4560–4567.
- [Zhao et al., 2022] Zhao, K., Zhang, Z., Wen, H., and Scano, A. (2022). Number of trials and data structure affect the number and components of muscle synergies in upper-limb reaching movements. *Physiological Measurement*, 43(10):105008.
- [Ziegler-Graham et al., 2008] Ziegler-Graham, K., MacKenzie, E. J., Ephraim, P. L., Travison, T. G., and Brookmeyer, R. (2008). Estimating the Prevalence of Limb Loss in the United States: 2005 to 2050. *Archives of Physical Medicine and Rehabilitation*, 89(3):422–429.