

Design Principles for Improving Precision and Dexterity of Soft Robotic Hands

A dissertation presented

by

Clark Benjamin Teeple

to

The John A. Paulson School of Engineering and Applied Sciences

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Engineering Sciences

Harvard University

Cambridge, Massachusetts

November 2021

©2021 Clark Benjamin Teeple

All rights reserved.

Design Principles for Improving Precision and Dexterity of Soft Robotic Hands

Abstract

As robots move more and more into the real world, they need to be equipped to interact with a wide array of objects and environment features while maintaining a gentle touch. Soft robotic end effectors enable passive adaptability, but give up strength and precision in return. From simple grasping to in-hand manipulation, this thesis explores how the design of soft robotic hands influences a robot's overall manipulation capabilities. We begin with a foray into finger design, developing pneumatically-actuated soft fingers capable of robust precision grasping and power grasping via two independently-actuated bending segments. This new finger design forms the basis of a further investigation into the role of gripper compliance in grasping and manipulation of thin, flexible materials, where we find that vertical, lateral, and rotation compliance all play a role in minimizing damage. With a similar goal of gentle interaction, we then shift focus to dexterous in-hand manipulation, exploring how distributing controlled degrees of freedom into various parts of the hand contributes to overall dexterity. We first develop a dexterous soft finger design and soft hand platform capable of moving objects within the hand. Using this new hand platform, we show that active control of the arrangement of digits affects the categories of objects that can be successfully manipulated. We also show that controlling the interaction between objects and the palm (via the palm's frictional properties and location relative to fingers) enables greater grasp stability and expanded access to different motion primitives. Finally, we collect these results into a set of application-specific design principles which can be used to inform the design of soft hands with dexterity tuned for the particular application.

Contents

Title Page	i
Abstract	iii
Table of Contents	iv
Acknowledgments	viii
Dedication	x
1 Introduction	1
1.1 Gentle, Yet Precise Grasping in the Real-World	1
1.1.1 Motivation	1
1.1.2 Traditional Robotic Grasping	2
1.1.3 Compliance in Robotic Grasping Devices	4
1.1.4 Advances in Soft Robotic Hands for Grasping	5
1.2 Dexterous In-Hand Manipulation	7
1.2.1 Motivation	7
1.2.2 Traditional In-Hand Manipulation	7
1.2.3 Advances in Dexterous Soft Hands	8
1.3 Contributions and chapter organization	12
1.4 Previously published manuscripts	12
2 Multi-segment soft fingers for enhanced precision grasping	14
2.1 Introduction	14
2.1.1 Overview	15
2.2 High-Level Finger Design Principles	16
2.2.1 Fingertip Compliance	16
2.2.2 Fingertip Curvature	17
2.2.3 Precision Grasping with Soft Fingers	18
2.2.4 Power Grasping with Soft Fingers	21
2.3 Designing a Prototype Soft Hand	22
2.3.1 Designing Modular Two-Segment Fingers	22
2.3.2 Fabricating Two-Segment Fingers	23
2.3.3 Designing a Rigid Palm	25
2.4 Characterizing Fingers and Palm	26
2.4.1 Actuation Modes Replicate Finger Structure	26
2.4.2 Functional Evaluation of Fingers	27

2.4.3	Finding a Suitable Palm Angle	31
2.5	Characterizing Grasping Performance	32
2.5.1	Hand Placement	32
2.5.2	Robustness to External Forces	33
2.5.3	Fingertip Compliance During Grasping	35
2.5.4	Relative Segment Lengths	36
2.6	Results	36
2.6.1	Hand Placement Determines Grasp Type	37
2.6.2	Object Size Range	38
2.6.3	Rotational Stability	40
2.6.4	Robustness to External Forces	42
2.6.5	Grasp Stiffness	42
2.6.6	Relative Length of Finger Segments	44
2.6.7	Grasping Arbitrary Objects	44
2.7	Discussion	49
2.7.1	Two-Segment Fingers Enable Robust Pinch Grasps	49
2.7.2	Power Grasps Are Better-Performed With One Uniform Segment . .	49
2.7.3	Performance Tradeoffs Inform Design of Segment Lengths	50
2.7.4	Two Independently-Actuated Segments Enable Best Performance . .	51
2.7.5	Fingertip Compliance and Shape	51
2.7.6	Grasping Arbitrary Objects	52
2.7.7	Extension to Non-Planar Grasping	53
2.8	Conclusions and Future Work	54
3	The role of multi-dimensional compliance in fabric manipulation	56
3.1	Introduction	56
3.2	Conceptual Grasping Analysis	58
3.2.1	The Role of Compliance in Grasping	59
3.2.2	The Role of Compliance During Snags	60
3.2.3	The Role of Multi-Dimensional Compliance	61
3.3	Results	61
3.3.1	Robustness to Uncertainty During Grasping	62
3.3.2	Robustness During Non-Prehensile Manipulation	62
3.3.3	Handling Snags	65
3.3.4	Task-Relevant Demonstration	67
3.4	Discussion	69
3.5	Conclusions	70
4	Designing dexterous soft fingers for in-hand manipulation	71
4.1	Introduction	71
4.2	Design of a Dexterous Soft Hand	73
4.2.1	Task-Centric Performance Goals	73
4.2.2	High-Level Hand Design	73
4.2.3	Dexterous Finger Design	76
4.2.4	Hardware Fabrication	77

4.3	Characterization of the Hand Prototype	78
4.3.1	Actuation and Control	78
4.3.2	Finger Performance	78
4.3.3	Grasping Performance	81
4.4	In-Hand Manipulation Primitives	83
4.5	A Simple Finger Gait for Continuous Rotation	85
4.6	Demonstration of Real-World Manipulation	87
4.7	Conclusions and Future Work	89
5	Designing Digit Arrangement	91
5.1	Introduction	91
5.2	Related Work	92
5.3	Designing Digit Arrangement	93
5.4	Performance Metrics for In-Hand Manipulation	94
5.5	Large-Scale Design Study	95
5.6	Finger Control and Trajectory Design	96
5.6.1	Testing Procedure	98
5.6.2	Object Set	100
5.7	Results	101
5.8	Translation to Real-World Tasks	104
5.9	Discussion	105
5.9.1	Geometry Explains Performance Differences Between Digit Arrangements	105
5.9.2	Successful Digit Arrangements Depend on Finger Design	106
5.9.3	Limitations	106
5.10	Conclusions	107
6	Designing dexterous palm surfaces	109
6.1	Introduction	109
6.2	Controlling the Position of the Palm	110
6.2.1	Introduction	110
6.2.2	Achieving High-Quality In-Hand Manipulation	111
6.2.3	Design Exploration using Passive Palms	112
6.2.4	Design of the Actuated Palm	120
6.2.5	Results and Discussion	121
6.2.6	Conclusion	122
6.3	Controlling Palm-Object Interactions via Friction	123
6.3.1	Introduction	123
6.3.2	Grasping and In-Hand Manipulation Requirements	125
6.3.3	Design Parameters & Analysis	126
6.3.4	Results	131
6.3.5	Grasp Stability	136
6.3.6	Manipulation Tasks from Sequential Motion Primitives	136
6.3.7	Discussion	137
6.3.8	Conclusions	139

7	Conclusions and Future Work	141
7.1	Design Principles for Soft Robotic Hands	141
7.1.1	Grasping	141
7.1.2	In-Hand Manipulation	142
7.2	Further Exploration into Onboard Sensing	144
7.3	Further Exploration into Learning for Soft Manipulation	145
7.4	Conclusions	146
	Bibliography	148
	Appendix A Toward Soft Curvature and Contact Sensors for Grasping	165
A.1	Introduction	165
A.2	Sensor Design	167
A.2.1	Design Criteria	167
A.2.2	Waveguide Design for Soft Fingers	167
A.2.3	Deep-Sea Interface Design	168
A.3	Modeling optical losses as a function of deformation	168
A.3.1	Curvature Sensors	169
A.3.2	Contact (Normal) Force Sensors	170
A.4	Methods	171
A.4.1	Fabrication of Soft Waveguides	172
A.4.2	Fabrication of Integrated Curvature Sensors	173
A.4.3	Fabrication of Contact Force Sensors	174
A.4.4	Data Acquisition and Processing	174
A.5	Sensor characterization	175
A.5.1	Characterization of Discrete Sensors	175
A.5.2	Characterization of On-board Sensors	176
A.6	Evaluation of models and experiments	177
A.7	Characterization under deep-sea conditions	178
A.7.1	Temperature	178
A.7.2	Hydrostatic Pressure	179
A.8	Grasping Objects	180
A.9	Conclusions	181
	Appendix B Smooth pressure control for soft robots	182
	Appendix C Grasp Testing	184
C.1	Experimental setups for measuring grasp performance	184
C.1.1	Grasp testing	184
C.1.2	Robustness to external forces	184
C.2	Measuring finger curvature - experimental procedure	186
C.3	Measuring object motion during grasps	186
C.4	Hand Placement Tests	187

Acknowledgments

My sincerest gratitude goes out to everyone who has in some way helped get me here. So many people have either directly or indirectly contributed to my understanding of robotics, engineering, arts, and the world at-large. There's no way I can possibly include everyone, but I'll do the best I can.

First and foremost, thank you so, so, *so* much to my family. I thank my parents, Rod and Sandra Teeple, for their unwavering support and dedication. They lifted me up in innumerable ways, and shaped me into the fine young man I am today. I also thank my fiancé, Julia Ebert, whose love, support, and understanding through the PhD process has been unparalleled. Thank you also to my brother, Grant Teeple, whose passion for music has shaped my understanding of what's possible. Thank you all so much for getting me through this journey.

Next, I thank my PhD Committee, Rob Wood, Rob Howe, and Conor Walsh. In particular, thank you to my wonderful advisor, Rob Wood, who gave me the freedom to explore my research interests, while growing with me as I ventured into a field outside of the lab's expertise. Rob also sustained a dynamic, collaborative lab culture that is, frankly, unmatched, and is one of the primary reasons for my success during my time here. Thank you also to Rob Howe, who has always graciously lent his expertise in robotic hand design, and whose difficult questions have kept my work grounded in reality. Finally thank you to Conor Walsh for the helpful discussions about systems-level design, and mentorship during my teaching experience.

A lab is so much more than just a building, so now I thank everyone in the woodlab (and friends) community for keeping the culture so energizing, open, collaborative, and fun! I have sincerely enjoyed the countless discussions, advice, and random life chats throughout my time. In particular, thank you to Kait Becker for her wonderful mentorship over the years, and Moritz Graule for being a great partner in crime. I must also thank Emma, Siyi, and Becky for the camaraderie, because it's always better to do something in good company, PhD's included. Thank you to Qian Wan for helping me dip my feet into the world of robotic hand design (and band practice). Thank you to Sylvain Abondance for helping spearhead the in-hand manipulation work that has now expended to make up the majority of this dissertation, and Buse Aktas, Michelle Yuen, Grace Kim, Randy St. Louis, and Justin Werfel who have helped sustain this work. Thank you to Dan Bruder and Zeo Liu for bringing new perspectives to our "hand chats". Thank you also to Michael Bell and Daniel Vogt for their wealth of random fabrication knowledge, and Michelle Rosen for

Acknowledgments

showing me the ropes to our lab's community outreach. And of course, I must acknowledge my countless games of Foosball with so many other talented people.

Finally, thank you to all of the wonderful friends I have made over the course of the 5+ years I've spent at Harvard. Whether it was jamming out in the basement of Memorial Hall, planning events with the SEAS Graduate Council, doing community outreach, or just grabbing dinner, I'm so happy I had the opportunity to meet everyone here. I am truly thankful to everyone I have met for helping me keep up the excitement and drive to make it through this degree.

To all my families.

This dissertation is a body of proof
that success comes not only from standing on the shoulders of giants,
but also from leaning on the shoulders of those you love and trust.

1

Introduction

In the real world, robots must be capable of interacting in a robust and safe way with their environments, even in the presence of huge variation in conditions. From handling or cooking food, to setting the table, to cleaning up afterward, robots of the future could be expected to take care of a huge variety of household tasks, or even collaborate with humans to complete them. Outside the home, robots could be expected to stock delicate produce on store shelves, or perform labor-intensive maintenance and cleaning tasks such as cleaning restrooms, all while protecting the safety of nearby human shoppers or employees. Additionally, using human-centric tools while attempting any of these tasks will often require adjusting the grasp between the time the tool is acquired and when it is actually used. In all these applications, sources of variation arising from the object (size, shape, pose, compliance, frictional properties, etc.) and the surroundings (locations of obstacles, poses of surfaces such as walls or tables) still remain challenging for modern robots to adjust for. Thus, the ideal robotic manipulation system for the real world must be capable of robust, yet precise grasping, as well as dexterous in-hand manipulation.

1.1 Gentle, Yet Precise Grasping in the Real-World

1.1.1 Motivation

Robots are beginning to find their way into applications where gentle grasping of delicate objects is critical. Handling food items such as fruits, vegetables, or pastries requires gentle interaction at all stages to prevent damage [1]. Many activities of daily living (ADL) such as stocking a refrigerator, cooking food, or setting a table require maintaining gentle contact forces (typically below 10 N [2]). As robots become more capable, collaborating

with humans, operating in the lab, or even performing tasks in space requires a level of guaranteed safety such as maintaining low impact energy densities [3]. Even assembly and pick-and-place tasks can benefit from gentle end effectors when the environment is cluttered or uncontrolled [4].

To accomplish the aforementioned tasks, robots must accommodate a relatively large degree of uncertainty in the shape, size, and pose of objects in the environment [5]. To do this, humans and robots alike use a wide variety of grasps while performing manipulation tasks [6]. In most contexts, grasps can be broadly categorized into power and precision grasps as defined in the human grasp taxonomy of [7]. Power grasps emphasize stability, usually involve enveloping the object, and are often accomplished using multiple points of contact between the object and the surfaces of the fingers and palm. Precision grasps on the other hand, such as fingertip or pinch grasps, involve more focus on applying small forces and enabling such capabilities as dexterous manipulation. Thus, while power grasps are useful for securing larger, heavier objects, smaller objects tend to require the use of precision grasps, especially if they must be grasped from a surface.

1.1.2 Traditional Robotic Grasping

Traditional rigid robotic grippers tend to be relatively simple devices designed to provide reliable grasping of specific, well-defined parts in manufacturing settings, as shown in Figure 1.1. The most common type of gripper, the two-fingered parallel jaw gripper, usually involves two rigid fingers that move toward one another to produce a strong grasp. These grippers are rarely backdrivable, usually actuated by a motor and lead screw combination, or a pneumatic piston. As a consequence of their rigid construction, the traditional robotics pipeline (perception, planning, and execution) suffers from extreme sensitivity to uncertainty in the environment [4]. Even small errors in perception of object poses or world features can lead to catastrophic failures (i.e., failed grasps, damaged items) and safety concerns.

One approach to handling large environmental uncertainty of objects is to increase the number of actuated degrees of freedom in the hand and rely on complex control strategies to coordinate them. Several examples of highly articulated robotic hands exist (some are shown in Figure 1.1), and an overview is included as part of [8]. Most high-dimensional robotic hands are anthropomorphic in design, and usually include at least one actuator for




Figure 1.1: Rigid robotic hands are born out of industrial automation, with typical construction involving rigid links and joints designed for high-power or high-precision applications. Numerous grippers have been developed as products for the automation market, for example a) the Schunk Gripper, b) Barrett Hand [17], and c) Robotiq Grippers [18]. Research platforms with higher dexterity have also been developed, such as d) Utah/MIT Dexterous Hand [9], e) Allegro Hand [12], and f) Robonaut 2.0 Hand [11].

each finger joint. Examples include the Utah/MIT dexterous hand [9], the ShadowRobot Shadow hand [10], the Robonaut 2 hand [11], and the SimLab Allegro hand by [12]. Each of these hands have 16-20 actuators, and the Shadow hand is capable of being configured with up to 40 actuators. While highly dexterous, the complexity of control needed to coordinate these hands is usually large, and unnecessary for many grasping tasks.

Another approach to mitigating the effects of uncertainty on grasp success is to utilize machine learning to develop robust grasp planners trained on large object sets. For example, the DexNet project demonstrates that deep learning can yield a highly successful grasp planning system when trained using hundreds of simulated grasps on more than 10,000 objects [13, 14, 15]. DexNet was even used to develop an ambidextrous grasp planner [16]. However, these deep learning approaches must be retrained for each new gripper (on the order of 30,000 compute hours for DexNet 1.0) and make bold assumptions about the contact mechanics and rigidity of objects and the gripper. Furthermore, DexNet performs very poorly for some rigid object geometries with complex features such as protrusions and curved sides that would be simple or even helpful for humans to grasp [14]. Thus, bottom-up learning approaches could potentially be improved with better gripper designs.

1.1.3 Compliance in Robotic Grasping Devices

A paradigm shift toward under-actuated fingers with built-in compliance has emerged as a way to embody a robot hand with structural adaptability during grasping, without complex control. Structural compliance enables passive adaptation to object shapes without explicit knowledge of the object or environment, as shown in Figure 1.2. For example, the Robotiq two-finger gripper family uses a single actuator, but can still adapt between a parallel plate grasp and an enveloping power grasp [18]. Using carefully designed kinematics, joint limits, and joint compliance, this rigid hand adapts its grasp passively based on where and how the force vectors are applied to the plates. Achieving similar kinematic behavior, the Velo Gripper [19] utilizes tendon-driven fingers to passively adapt the grasp in a more compact mechanism. Furthermore, the Pisa/IIT Soft Hand [20] utilizes joint compliance and mechanical coupling between fingers as adaptive grasping synergies to achieve close to human performance with only four actuators.

While joint compliance in planar pin joints enables robust finger adaptation during grasping, three-dimensional compliance extends robustness to uncertainty. The SDM Hand [21], for example, uses compliant finger flexures as joints, allowing for small off-axis finger motions during grasping. With the addition of sensing and other design changes, the iRobot-Harvard-Yale (iHY) hand [22] achieves high passive compliance in actuated directions to enable robust power grasping, while retaining small off-axis compliance for precise fingertip control during precision grasping.

Building on the successes of simple, yet robust passive adaptation, others have chosen to focus on modulation of joint stiffness through additional actuators or impedance control. For example, the BarrettHand grasper [17] and the SRI Hand [23] use clutches in the joints to lock them in place. This enables passive adaptation to object shape, and strong grasps when the clutches are engaged. Conversely, the DLR Hand II [24] and recently CLASH [25] achieve fingertip stiffness modulation through impedance control and a clever differential drive mechanism. This allows on-the-fly stiffness control without additional actuators.

In an orthogonal approach to compliant grasping, local compliance at the fingertips is utilized through finger pads rather than compliant fingers or joints. For example, Maruyama *et al.* developed deformable fingertips that can interact gently with objects at first, then increase their stiffness as they deform [26]. More recently, McInroe *et al.* de-

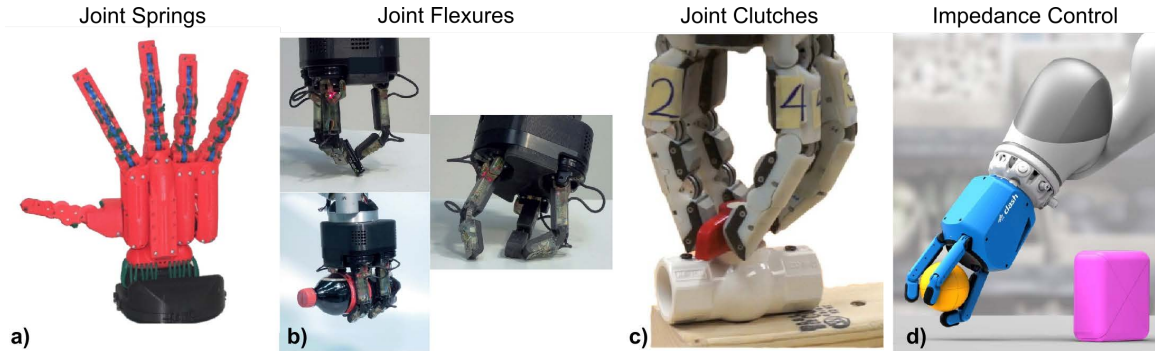


Figure 1.2: Structural compliance in fingers enables passive adaptation to object shapes without explicit knowledge of the object or environment. Selective compliance can be accomplished through mechanisms such as a) joint springs in the IIT/Pisa SoftHand [20], b) joint flexures in the iHY Hand [22], c) joint clutches in the SRI Hand [23], or d) impedance control in the CLASH hand [25].

veloped a similar soft fingertip that uses pneumatic actuation to apply forces to objects, while also being capable of measuring the fingertip’s complex deformation [27]. In addition, compliant fingertips have the potential to improve the stability of a grasp due to increased contact area and restoring forces according to Cutkosky [28]. However, compliant fingertips only ensure gentle interactions with objects directly at the fingertips.

1.1.4 Advances in Soft Robotic Hands for Grasping

Recently, more focus has been placed on building robotic hands that can safely and gently interact with their environments. This shift in application goals has given rise to soft robotic hands, where both the finger structure and contact surfaces are made of compliant materials [4]. Rubbers, fabrics, and foams are used to build actuators that minimize the risk of damage, especially when interacting with delicate targets [29, 30]. In addition, passive compliance of soft fingers reduces the control complexity required to robustly grasp objects [4, 31]. However, what soft robots gain in adaptability, they often lose in strength and precision [32].

Soft robotic hands or grippers are typically well-suited to grasp unknown, irregularly-shaped, or delicate objects. This can also translate to better handling of uncertainties in object pose that arise from vision and other sensory systems. For example, [33] demonstrate a soft gripper with a single pneumatic input capable of performing grasps on objects with minimal sensitivity to position errors. Brown *et al.* developed a universal jamming gripper

capable of grasping a wide array of object shapes [34]. In an interesting application, Galloway *et al.* designed and deployed a soft hand to perform sampling of delicate marine life in the deep sea [35], one of the most challenging environments to operate in. Furthermore, Deimel *et al.* built a dexterous soft hand, the RBO Hand II [36], capable of performing all but two grasps in the Feix taxonomy [37].

While most soft hands can achieve excellent power grasping due to passive compliance, they typically have trouble grasping small objects using precision grasps. During deep-sea exploration, the gripper built by Galloway *et al.* could withstand up to 17 N applied to the object during a power grasp with two opposing fingers, but relied on caging animals that were smaller than the minimum power grasping size [35]. Using a jamming gripper, enveloping grasps can be used for small objects, but can only be performed by pressing objects against a surface [34]. With more dexterous soft hands such as the RBO Hand II [36], the majority of successful grasps performed were power grasps. The RBO Hand 2 has six degrees of freedom, is capable of withstanding forces up to 8 N, but grasp stiffness was a main limitation due to the large finger compliance. Another soft hand of note comes from the prosthesis design community, where a 6-DOF hand was developed to mimic the shape of the human hand and finger kinematics [38]. This prosthetic hand can withstand forces up to 18 N, and can perform power grasps and precision grasps, but the main focus of the study was myoelectric control and tactile feedback using the hand [38]. In a final example, the BCL-13 Hand is capable of very robust power grasps, but with no characterization of precision grasp performance [39].

The focus of recent work in soft grasping has been on performing power grasps, without much emphasis on precision grasps. For example, the Pisa/IIT Soft Hand [20] makes use of postural “soft” synergies that describe principal components of hand motions over a set of grasping tasks. In another study, O'Brien *et al.* demonstrate how soft structures can be used to passively adapt between high force and high-speed operation modes [40]. However, precision grasps were not the focus of either of these studies. Other recent studies of soft finger design focused entirely on power grasps, using simulation [41] and experimentation [42]. Furthermore, one recent study has been presented by Vogt *et al.* where adding a passive extension to soft fingers enabled them to perform pinch grasps [43]. However, to date, precision grasping with soft fingers remains an open and relatively unexplored design space.

1.2 Dexterous In-Hand Manipulation

1.2.1 Motivation

As robots become more viable for use in the real world, we see an increased need for robots to perform dexterous manipulation to achieve their goals, while remaining safe when in proximity to humans. Many activities of daily living (ADL) require “gentle dexterity” when the arm is constrained (e.g., picking plates or food items from a cupboard, or loading dishes and cutlery into a dishwasher) [2]. In addition, food handling tasks, such as placing delicate produce at correct positions/orientations while packing a grocery bag or unloading into the fridge, require a delicate, dexterous touch. Finally, collaborating with humans during assembly tasks may involve transitioning the grasp to maneuver tools in the hand (e.g., grasping a tool from a box, then presenting it to a human with the handle exposed).

In all of these applications, fragile elements in the environments (objects, humans, or both) are commonplace, yet robots must navigate these manipulation challenges robustly and safely. In addition to the ever-present environmental uncertainty, target objects are often located in highly-constrained poses, or must be placed into new poses that are outside the workspace of the arm. Thus, a robotic solution must be capable of dexterous manipulation, while simultaneously remaining gentle and safe.

Another useful, if underappreciated, skill for real-world manipulation is the ability to adjust the grasp on an object without using external surfaces. Re-grasping an object to improve the grasp quality has been studied extensively [44, 45], but most examples require the robot to set down the target object on a tabletop before re-grasping. This workflow relies on reasonably accurate models of the object’s dynamics, environment geometry, and contact dynamics in addition to the robot’s own dynamics [46, 47]. Alternatively, if robots can manipulate the object within their hands, grasps can be adjusted without relying on external surfaces.

1.2.2 Traditional In-Hand Manipulation

Dexterous in-hand manipulation usually requires precise planning and control of finger motion based on models of the object and fingers when performed by rigid hands [48, 49]. This is due to complex contact interactions between the fingers and objects, and has led to a variety of studies on how to properly control object motion. Early work

applies control theory to the problem of dexterous manipulation, developing stable motions and finger gaits based on known finger kinematics [50, 45, 51]. On the design side, a variety of highly dexterous rigid robotic hands have been built, including the Utah/MIT Dexterous Hand [9], DLR hand [24], and Robonaut 2 hand [11], as discussed above in "Traditional Robotic Grasping". However, this high-dimensional hardware was notoriously difficult to control [52]. This led to a push toward compliant hands to handle uncertainty and reduce control complexity [53, 20, 22, 10].

In more-recent years, a focus on motion planning for in-hand manipulation has led to a host of advances in planning algorithms for various dexterous tasks. For example, in-hand regrasping [44, 48, 54], multi-finger manipulation [55], and finger gating [49, 56] have all been at the center of modern in-hand manipulation planning frameworks. However, the majority of work in planning revolves around rigid objects and traditional, rigid hands, limiting the resulting robustness of plans, and resulting in control policies that are sensitive to small variations in the environment [49].

A more modern, data-driven approach to in-hand manipulation has gained popularity in recent years. Some attempts to mitigate the complexity of these interactions using machine learning show incredible promise, but require extensive training on high performance computing systems. For example, OpenAI had great success in learning hand-eye coordination from scratch for a simple cube pivoting task [57] (using the Shadow Dexterous Hand [10]), but required over 90 years of simulated grasps and very precise dynamic model of the hardware involved. Later, OpenAI extended their results, learning to solve a Rubik's cube with the same hand, demonstrating that in-hand manipulation can be learned with enough computing power [58]. However, we can mitigate the need for planning and complex control for some in-hand manipulation tasks through targeted design of soft robotic hands.

1.2.3 Advances in Dexterous Soft Hands

Dexterity has a variety of definitions [52] ranging from the colloquial "skill in use of hands" [59] to the highly-technical "capability of changing the position and orientation of the manipulated object from a given reference configuration to a different one" [45]. Using the latter definition, we see dexterity from a task-centric perspective, where the level of dexterity is related to the range of poses an object can be manipulated to and from. We thus quantify dexterity in context with a set of manipulation tasks rather than looking

narrowly at the hand’s kinematics or dynamics. From here, we can subdivide the hand into parts, and evaluate the dexterity of these parts with respect to the tasks they need to perform (for example, fingers moving objects or the palm constraining objects, etc). In this way, we can discuss how the dexterity of each sub-component contributes to the overall dexterity of the hand, which allows a holistic view of robotic hand design.

Most recent advances in the design of compliant robotic hands for in-hand manipulation tasks focus on increasing finger dexterity, as shown in Figure 1.3. For example, the Shadow Dexterous Hand is a rigid-link mechanism with compliant joints and has 24 joints controlled by 20 actuators, with 17 controlled degrees of freedom distributed amongst the fingers and thumb (each of the four fingers has 3-DOF control, and the thumb has 5-DOF control) [10]. This hardware was designed with in-hand manipulation in mind, but is still difficult to control due to its high dimensionality. A similar control difficulty exists for the BCL-26 Hand, a completely soft mechanism with 26 actuated degrees of freedom, where 24 actuators are dedicated to finger motion (five DOF in each of the four fingers, and four DOF in the thumb) [60]. On the flip side, the RBO Hand 3 was designed for in-hand manipulation, but with a focus on minimizing complexity. The RBO Hand 3 is a fully-soft hand with only 16 controlled degrees of freedom, (two in each of the four fingers, five in the thumb, and one to control abduction-adduction of the fingers), but its capabilities are extremely promising [61].

In this dissertation, the goal is to design soft hands that have “just enough” dexterity to perform a set of desired tasks, as shown in Figure 1.4. One common theme that seems to pervade the dexterous hand design space is an anthropomorphic hand structure. As noted in [63], human hands are adapted to be an excellent general-purpose design due to their combination of passive compliance and proprioception. However, soft hand designs that are more “task-driven” could potentially achieve better performance with fewer actuated degrees of freedom. In addition, clever distribution of actuated structures in other parts of the hand can vastly simplify the finger dexterity needed to achieve various in-hand manipulation tasks. This conjecture is central to the hand design principles developed in this dissertation.





Figure 1.3: Recent advances in soft robotic hands have achieved increased dexterity, but at the cost of increased complexity. a) The Universal Gripper [34] and the “squishy fingers” gripper developed for deep-sea sampling [35] have low dexterity (grasping only) with only a single actuation input. In contrast, Zhao et al. [62] and the RBO Hand 2 [36] have higher dexterity (able to perform robust grasps and simple in-hand manipulation) via five independently controllable fingers, with the RBO hand using an additional thumb actuator. Finally, the Shadow Dexterous Hand [10], BCL-26 [60] and RBO Hand 3 [61] (not pictured) achieve high dexterity (capable of complex in-hand manipulation) through several actuated degrees of freedom per finger and additional actuators in other parts of the hand.

**Improving Precision
(for Gentle Grasping)**



Precision Grasping




Thin, Flexible Objects


**Improving Dexterity
(for In-Hand Manipulation)**



Dexterous Soft Fingers



Finger Arrangement



Dexterous Palms

Figure 1.4: The goal of this dissertation is to improve the precision of soft grippers while still achieving gentle grasping, and improve the dexterity of soft hands via thoughtful distribution of actuated degrees of freedom into various structures within the hand.

1.3 Contributions and chapter organization

This dissertation focuses on improving the precision of soft grippers while still achieving gentle grasping, and improving the dexterity of soft hands by exploring how to distribute actuated degrees of freedom throughout the hand. Chapter 1 introduces the field of soft robotic hand design, including relevant background on grasping and in-hand manipulation (*in fact, you are reading it right now*). Chapter 2 focuses on the design of soft robotic fingers for robust precision grasping, resulting in a soft finger design with two independent bending segments. Chapter 3 investigates the role of gripper compliance in successful grasping and manipulation of thin, flexible materials, utilizing the finger design developed in Chapter 2. Chapter 4 details the development of dexterous soft fingers and a dexterous hand platform capable of moving objects within the hand using several basic motion primitives. Chapter 5 shows how the arrangement of digits in a soft hand affects the quality of in-hand manipulation on different categories of objects. Chapter 6 demonstrates how controlling palm-object interactions (via the palm’s frictional properties and location relative to fingers) enables greater grasp stability and expanded access to different motion primitives. Finally, Chapter 7 concludes with a collection of application-specific design principles, closing thoughts on soft hand design, and a discussion of the outlook for future scientific exploration into dexterous soft hands.

1.4 Previously published manuscripts

Much of the content of this thesis has been published in the following manuscripts:

- C. B. Teeple, T. N. Koutros, M. A. Graule, and R. J. Wood, “Multi-segment soft robotic fingers enable robust precision grasping,” *International Journal of Robotics Research*, 2020
- S. Abondance, C. B. Teeple, and R. J. Wood, “A dexterous soft robotic hand for delicate in-hand manipulation,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5502–5509, 2020
- C. B. Teeple, G. R. Kim, M. A. Graule, and R. J. Wood, “An active palm enhances dexterity for soft robotic in-hand manipulation,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2021

- C. B. Teeple, R. C. St. Louis, M. A. Graule, and R. J. Wood, “The role of digit arrangement in soft robotic in-hand manipulation,” in *IEEE International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2021
- M. A. Graule, C. B. Teeple, T. P. McCarthy, R. C. St. Louis, G. R. Kim, and R. J. Wood, “Somo: Fast and accurate simulation of continuum robots in complex environments,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021
- C. B. Teeple, B. Aktaş, M. C. Yuen, G. R. Kim, R. D. Howe, and R. J. Wood, “Controlling palm-object interactions via friction for enhanced in-hand manipulation,” in *IEEE Robotics and Automation Letters*, IEEE, 2022 (In Review)
- C. B. Teeple, J. Werfel, and R. J. Wood, “Multi-dimensional compliance of soft grippers enables gentle interaction with thin, flexible objects,” in *2022 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2022 (In Review)

2

Multi-segment soft fingers for enhanced precision grasping

2.1 Introduction

Robotic grasping and manipulation often requires some form of online adaptation of grasp strategies. Complex contact interactions between the hand and objects make it challenging to perform grasping without perception, tactile feedback, or otherwise detailed information about the world. In addition, contact interactions can change dramatically depending on the mechanical properties of the object and fingers. Attributes of both structures, such as size, shape, compliance, and surface finish, all play a pivotal role in the stability and precision of the grasping process. Thus, in order to grasp a large range of objects, robots need the ability to adapt their grasps during run-time.

The focus of recent work in soft grasping has been on performing power grasps, without much emphasis on precision grasps. For example, the Pisa/IIT Soft Hand [20] makes use of postural “soft” synergies that describe principle components of hand motions over a set of grasping tasks. In another study, Obrien et al. [40] demonstrate how soft structures can be used to passively adapt between high force and high-speed operation modes. However, precision grasps were not the focus of either of these studies. Other recent studies of soft finger design focused entirely on power grasps, using simulation [41] and experimentation [42]. Furthermore, one recent study has been presented by Vogt et al. [43] where adding a passive extension to soft fingers enabled them to perform pinch grasps. However the extension would likely interfere with power grasp operation.

Additionally, the effect of adding multiple bending segments to soft fingers has not

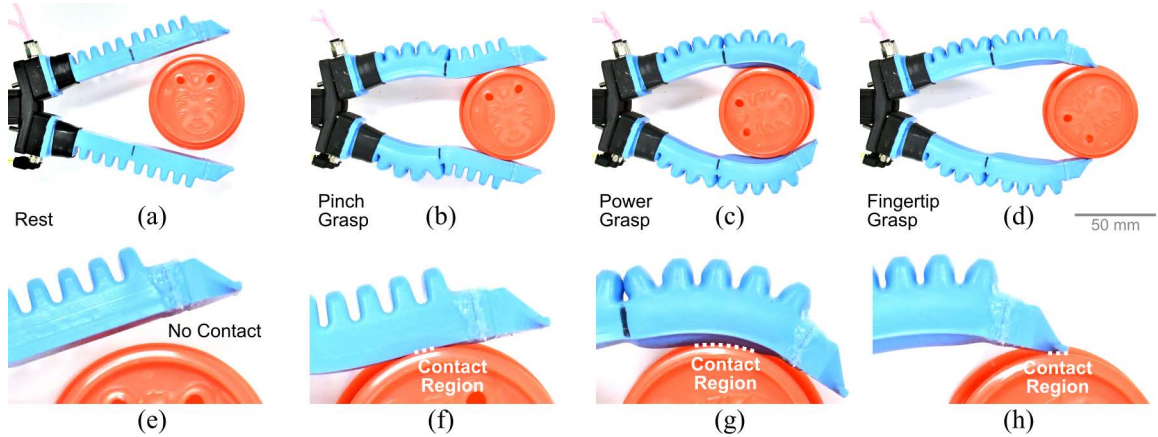


Figure 2.1: Our planar hand prototype with two co-planar segments per finger grasping a cup of 55 mm diameter. The hand is shown (a) at rest, (b) performing a pinch grasp with the side of the fingers, (c) performing a power grasp, and (d) performing a fingertip grasp. Close ups of fingers in (a), (b), (c), and (d) are shown in (e), (f), (g), and (h), respectively.

yet been explored for precision grasping. While Deimel et al. [71], Zhou et al. [39], and Zhou et al. [72] built soft fingers with more than one serial segment, all three studies focus on exploring the effect of their hand designs on power grasps. Zhou et al. [39] in particular find that the pullout force can be improved if fingers with two segments are actuated in a particular way. However, to date, precision grasping with soft fingers remains an open and relatively unexplored design space.

2.1.1 Overview

In this work, we show how simple, conceptual design rules can be used to design soft robotic fingers capable of excellent precision grasping without sacrificing power grasping performance. We present three main contributions: (1) a conceptual analysis of compliance and finger shape during grasping, which suggests that soft fingers should have at least two serial bending segments; (2) an empirical study of grasping performance comparing this two-segment finger design to uniformly-actuated fingers; and (3) experimental validation showing that fingers with two independently-actuated serial segments can achieve excellent precision and power grasps.

We first present a conceptual analysis of precision grasping with multi-segment soft fingers based on compliance and local finger shape. This high-level analysis suggests that

grasps are more stable when contact with the object occurs on the side of the finger (a pinch grasp) rather than the fingertip. In addition, achieving a pinch grasp with soft continuum fingers requires at least two independent bending segments each, but only requires actuation in the proximal segment. Furthermore, we explore the effect of the relative lengths of finger segments on pinch grasping performance.

Next, we empirically evaluate the grasping performance of the two-segment, proximally-actuated finger design compared to widely-used uniformly-actuated fingers. Performance is evaluated using several metrics: the acquisition region, object size range, rotational stability, and robustness to external forces. We confirm that the the proposed proximally-actuated finger design is capable of higher quality precision grasping than fingers with a uniformly-actuated design, and we show the tradeoff between power grasping strength and precision grasping capabilities as a function of segment length. However, power grasping is still best-performed with uniformly-actuated fingers. Thus, compromises in performance would need to be made if one single finger structure were to be chosen for each finger.

Finally, we show that adaptation between uniformly-actuated and proximally-actuated finger structures using two independently-actuated serial segments (as shown in Fig. 2.1) can achieve the best possible performance during both types of grasps, and can be implemented with only a limited increase in control complexity.

2.2 High-Level Finger Design Principles

To understand how the number of independent serially-linked segments in a soft finger affects its ability to perform robust grasps, we can build a conceptual argument around compliance and geometry at the contact point. Finger compliance directly affects fingertip motion when external forces are applied to the object. In addition, the shape of the finger and object at the contact point can be used to determine the sensitivity of object motion to fingertip motion. Combining knowledge of a finger’s compliance ellipse with the local fingertip shape provides useful insight into the stability of that grasp, which can ultimately be used to judge the quality of the finger design.

2.2.1 Fingertip Compliance

The compliance of any mechanical structure or linkage can be represented by the deflection of some point of interest in response to a unit force applied at all angles, resulting

in an elliptical region. This so-called compliance ellipse has been used to study human arm and finger function to visualize compliance in human extremities [73, 74, 75]. Compliance ellipses also form the basis of impedance control for robotic systems [76], where a desired endpoint stiffness can be achieved through joint control. In addition, Lim et al. [77] showed that designing the compliance ellipse at many points along the whole body of a mobile robot can improve the safety of human-robot interactions. Thus, it is natural to apply the same arguments to aid in the design of robotic fingers.

Designing fingers to achieve a desired tip-compliance ellipse (an ellipsoid in three dimensions) has proven to be a simple, yet effective method for highly under-actuated systems. For example, Anirban et al. [74] used impedance control to achieve a similar endpoint compliance ellipse as human fingers based on empirical measurements. In addition, Gravagne et al. [78] showed that the compliance ellipsoid of a continuum manipulator can be used to understand complex deflections under different tip loads.

Finally, Odhner et al. [22] used analysis of fingertip compliance when designing the flexure-based fingers of the iRobot-Harvard-Yale (iHY) hand, with the goal of aligning the major axis of the compliance ellipse normal to the fingertip surface. By examining the compliance ellipsoid of a set of generic soft fingers with multiple segments, we can understand how to best utilize control inputs for robust grasping.

2.2.2 Fingertip Curvature

Geometric analysis of how fingertip shape affects the rolling motion of an object can be used to gain some notion of the stability of a fingertip grasp. Cutkosky et al. [28] developed this analysis by investigating how several mechanical aspects of a finger affect the rotational stability of a grasp (i.e., how an infinitesimal rotation of the object affects stability). They found that the stability of a planar fingertip grasp increases as a function of both the radius of curvature and stiffness of the finger at the contact point. Furthermore, they found that with a sufficiently large radius of curvature, the grasp stability is infinitely stable regardless of finger stiffness. While this analysis assumes rigid fingertips, they note that the trends remain the same for soft fingertips.

Based on the insights from Cutkosky et al., Montana et al. [79] developed a description of grasp stability that agrees with this intuition. Under this framework, stability is increased with larger radii of curvature of both the object and fingertip. In addition,

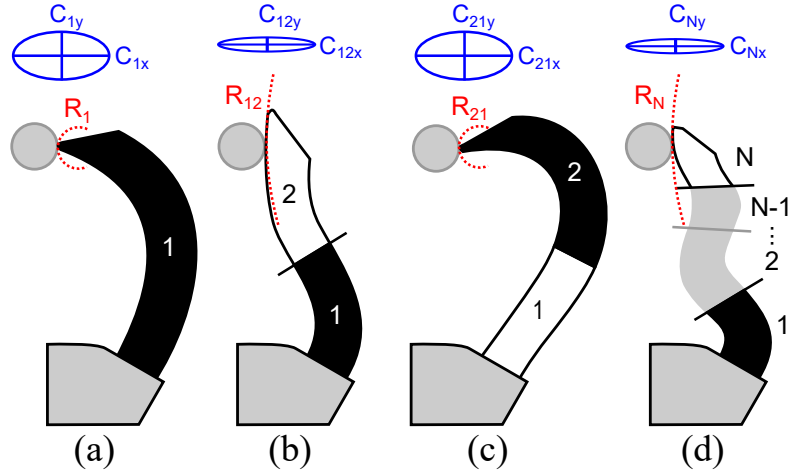


Figure 2.2: The shape of soft continuum fingers as they perform precision grasps with different numbers of bending segments. (a) One uniformly-actuated bending segment yields a fingertip grasp with small radius of curvature and larger anterior (C_y) compliance. (b) Two segments with only proximal actuation enable a pinch grasp with a much larger radius of curvature and smaller compliance. (c) Two segments with only distal actuation yields a fingertip grasp similar to that of a single uniformly-actuated bending segment. (d) More than two segments yields the ability to control fingertip orientation separately from position, with more options for actuation inputs to produce the desired pinch grasping configuration.

mechanical properties such as normal forces and viscoelasticity were found to only affect the stability of marginally-stable grasps. For example, increased viscoelasticity at the contact point was found to result in increased stability. This analysis again assumes perfect rolling contact, and looks at rotational stability, but is useful nonetheless to understand how soft fingers can be best utilized to perform robust precision grasps.

2.2.3 Precision Grasping with Soft Fingers

To design fingers that can perform high-quality precision grasping, the above analyses suggest the finger should have low compliance and small fingertip curvature at the contact point. Through examination of the local shape and compliance of soft fingers, we find that the placement of contact points on the side of the finger enables pinch grasping with a dramatic improvement in stability compared to fingertip grasps. The subsequent analysis is performed in the planar case ignoring the effects of gravity, but the resulting design rules can be extended to real-world grasping in a straightforward way (as presented in Section 2.7).

2.2.3.1 Single Uniformly-Actuated Segment:

As a baseline, let us first explore precision grasping using fingers with a single uniform bending segment. As the fingers are actuated, they first have uniform curvature over the whole length, then contact with an object causes nonuniform curvature with lower curvature at the proximal ends. This decrease in proximal curvature is due to the long moment arm over which the contact force acts. The result is that during a grasp, the ends of the fingertips contact the objects, as shown in Figure 2.2a. Most examples of existing soft robotic fingers exhibit this behavior [35, 80, 36].

The placement of forces on the fingertips results in grasps with low stability. First consider the local shape of the fingertips. While the fingertip curvature is highly variable, many designers choose to use pointed fingertips to improve power grasping against surfaces [36, 80]. However, pointed fingertips result in very high curvature (small radius) at the contact point during fingertip grasps, causing large object deflections from relatively small fingertip deflections. Furthermore, according to Cutkosky et al. [28], the grasp stability is a function of the object's curvature. On top of this, the fingertip compliance in the axis normal to the palm's surface is usually relatively large, causing large fingertip deflection from relatively small forces on the object. While not studied in-depth here, fingertip design could be an interesting area of future development.

2.2.3.2 Actuated Proximal Segment, Passive Distal Segment

Now consider a finger that has two uniform bending segments of equal stiffness, but only the proximal segment (closest to the base) can be actuated. In this case, during grasping the passive distal segment can perform a passive backward bend in order to balance forces on the object. Zhou et al. [39] utilized this phenomenon to grasp objects larger than the opening width of the fingers. During a precision grasp, the passive bend allows the contact points to be moved from the fingertip to the inside edge of the finger to form a pinch grasp given an appropriate object position, as shown in Figure 2.2b.

Grasping with the sides of the finger results in more stable pinch grasps. Compared to the case of a fingertip grasp, the local geometry at the contact point has much higher curvature and lower compliance in the direction normal to the palm. In fact, the object contacts the finger on a flat surface of approximately zero curvature (infinite radius), meaning grasp stability should have low dependence on object curvature and hand place-

ment inaccuracies. Furthermore, since the finger compliance in the direction normal to the palm is much lower, grasp stability (and robustness) becomes mostly dependant on friction between the finger and object. This approximates grasping with a parallel-jaw gripper.

2.2.3.3 Passive Proximal Segment, Actuated Distal Segment

Next, consider flipping the two-segment configuration, where only the distal segment can be actuated. In this actuation scheme, the fingertip ends up touching the object and bends the passive proximal segment backward, as shown in Figure 2.2c. Similar to a single uniformly-actuated segment, contact at the fingertip yields poor grasp stability due to the large fingertip curvature and high compliance.

2.2.3.4 More Than Two Segments

In a final case, consider a finger that has more than two bending segments. During a pinch grasp, we assume the object will only touch the finger at a single point. Even with only three segments, the finger now has a family of input configurations that can place the contact points on the inside edge of the finger, as shown in Figure 2.2d. While an increase in the workspace of the finger would likely enable interesting functionality, we are focused on the two segment case since that is the minimal configuration where desirable pinch grasping behavior can occur with the sides of the fingers.

2.2.3.5 Relative Stiffness of Finger Segments

While it may be possible to enable the desired placement of contact points with only one bending segment of non-uniform stiffness, we restrict the focus of this work to uniform bending segments for simplicity. Prior work in this area from Knoop et al. [42] shows that non-uniform stiffness can be used to tune the contact pressure a soft finger applies at each point along its contact surface. However, the impact of these tuned contact-pressure profiles on grasping performance has not been evaluated in detail. In addition, there are likely inherent compromises in grasping performance when designing one mechanism to passively adapt between robust pinch grasping and strong power grasping. Rather than attempting to search a potentially large design space, we focused on using fairly simple sub-components (serially-linked uniform bending segments with equal stiffnesses) due to the simple relationships between actuation pressure and free curvature.

2.2.3.6 Relative Lengths of Finger Segments

To gain a high level conceptual understanding of how the relative lengths of finger segments affect precision grasping performance, we can abstract the two-segment fingers as two serially connected cantilever beams. Assuming a symmetric grasp, we can further simplify the grasp and look at only one finger. During a grasp, each finger segment can potentially have a single point load (from contact with the object) and an internal moment (from actuation pressure). With the proximal segment rigidly fixed to mechanical ground (the palm), the distal segment is joined serially with the proximal beam at the ‘connection point’. The following discussion is limited to a contact force applied only in the distal segment, which is the case in many successful pinch and power grasps.

First, consider the effect of lengthening the proximal segment. The deflection and bending angle at the ‘connection point’ increase with increasing actuation torque and length of the proximal actuator. This deflection and angle define the neutral position of the distal segment. Next, for an object with a fixed relative position and size, the position of the contact point on the distal segment is constrained. The deflection of the distal segment at the contact point therefore increases with increasing torque and length of the proximal actuator. Thus, for a constant actuation torque in the proximal segment, the contact force on the object increases as the proximal segment gets longer.

Similarly, the stiffness of the unactuated distal segment, and thus the contact force, increases with decreasing distal segment length. Furthermore, keeping the overall length constant, an increase in the length of one segment directly results in a decrease in length for the other segment.

Since these two effects are additive, a smaller distal length fraction generally enables higher contact forces on the object. Finally, higher contact forces lead to increased frictional forces assuming a constant friction coefficient between fingers and the object. Thus, we expect grasps to have increased robustness to external forces with proximally-actuated finger structure with decreasing distal length fraction.

2.2.4 Power Grasping with Soft Fingers

As detailed in Section 2.1, strong power grasping has been achieved in numerous studies, and can be robustly achieved using fingers composed of one uniformly-actuated bending segment. Therefore, to achieve the best possible grasping performance for power

and precision grasps using soft continuum fingers, we can see that two fundamentally-different finger structures are required. For the strongest possible power grasp, the obvious choice is the more-traditional finger design with one uniformly-actuated bending segment. Conversely, to achieve the most-stable precision grasping, fingers need at least two bending segments, but only the proximal segment needs to be actuated. In the following sections, we empirically investigate the tradeoffs in grasping performance that arise from each finger structure

2.3 Designing a Prototype Soft Hand

To illustrate the concepts explored in the previous section, we designed and built a soft robotic hand capable of interacting with objects in a plane. The hand consists of two fingers, each with two independent co-planar bending segments. The fingers are mounted on a rigid palm with some distance and angle between them, as shown in Figure 2.3. The following sections detail the design choices made and fabrication methods used to build a robust grasping system we can use to test our claims.

2.3.1 Designing Modular Two-Segment Fingers

Several criteria were taken into account when designing the fingers of our soft robotic grasping system. For simplicity, we limit the fingers to two serial bending segments. Two segments is the minimal configuration needed to enable placement of contact points on the side of the fingers, as discussed in the previous section. These segments should also have equal passive stiffness in order to approximate a single bending segment when equal actuation inputs are applied to both segments. In addition, the relative lengths of the two segments is a parameter of interest, so this should be easy to choose during construction. Finally, the fingertip shape should be consistent with other soft fingers designed for power grasping so as to preserve power grasping performance.

To address these design criteria, several key choices were made during the design process. First, the fingers utilize bellows-style pneumatic bending actuators as used in [35]. Each finger is split into two independently-actuated serial segments, as shown in Figure 2.4a, with the ability to control the relative lengths of the two segments at design-time. We keep each finger’s workspace free of tubing by routing pneumatic connections through one proximal hub on each finger. However, internal routing of pneumatic lines

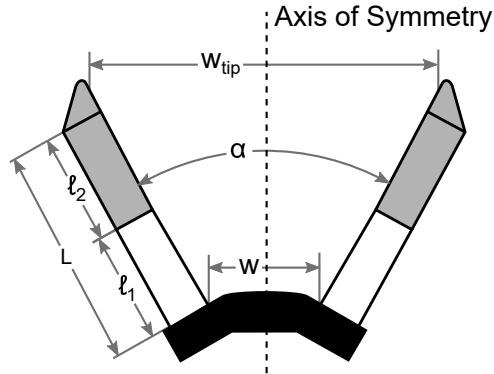


Figure 2.3: Schematic diagram of our planar soft-robotic hand prototype. Each finger has two bending segments of lengths ℓ_1 and ℓ_2 . Two fingers are mounted to a palm with a width of w and an angle of α between them.

required extra thickness in the proximal actuator wall. To ensure the bending stiffness of both segments is roughly equal, extra thickness in the proximal actuator walls was placed near the neutral bending axis. Finally, we used wedge-shaped fingertips since this shape is commonly employed for better power grasping against surfaces [80].

In addition to addressing explicit design goals, several aspects of the finger design space were held constant for simplicity. The most critical constant parameter is the overall length of the fingers, chosen to be on the order of 100 mm so as to be roughly the length of a large human finger. Differences in the segment lengths with respect to each other are controlled while keeping the overall finger length constant. In addition, we expect the stiffness of each segment (relative to actuation pressure) to contribute directly to the shapes that fingers form. However, while we would expect these fixed parameters to shift grasping performance and affect the magnitude of tradeoffs in the performance metrics, we would not expect them to fundamentally alter the results when comparing fingers with two actuated segments to fingers with one uniformly-actuated segment.

2.3.2 Fabricating Two-Segment Fingers

The fabrication process for our two-segment soft fingers involves a variation on the molding techniques discussed in [35], as well as coupling with rigid 3D printed hubs for fluid and structural connections. All molds were 3D printed on an Object Connex 500 printer (VeroClear material, Stratasys). All hubs were printed either on an Object Connex

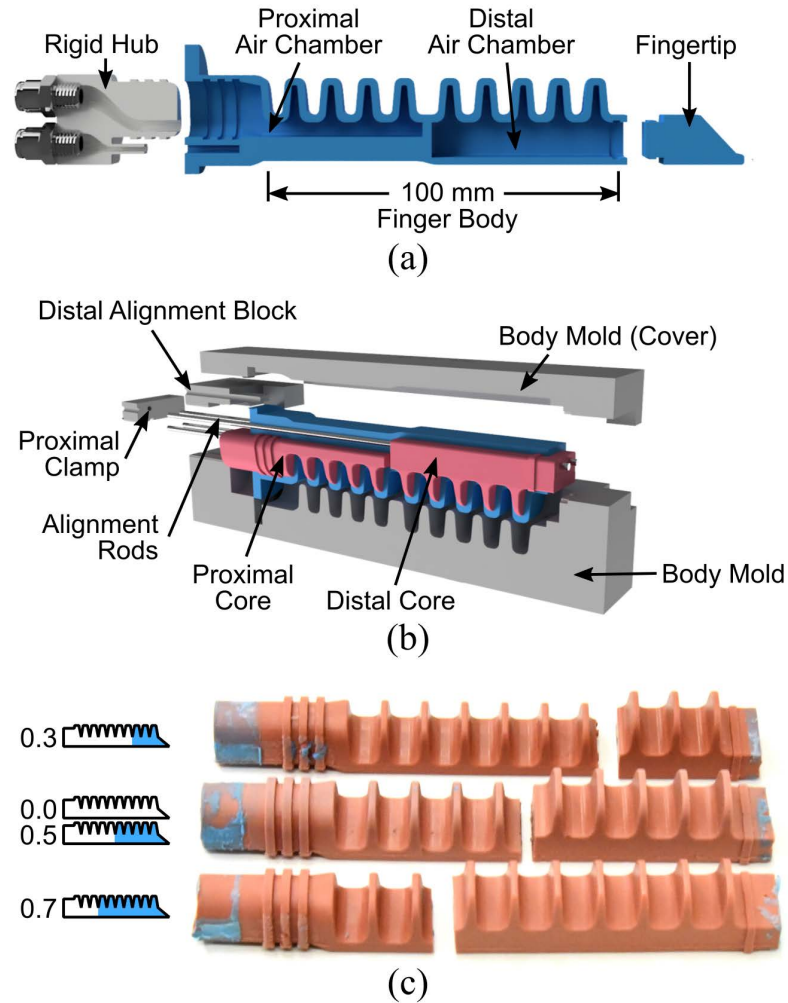


Figure 2.4: (a) Each soft finger is composed of a silicone skin and fingertip (both made out of Smooth-Sil 950), and a rigid hub at its base. The body of the finger contains two bellows-style pneumatic bending segments. (b) To fabricate a finger, we use a four-part mold made of 3D-printed resin (VeroClear RGD 810), as well as two soft silicone cores (Elastosil 4061, Wacker). (c) The soft cores can be trimmed to length to generate fingers with different segment lengths.

500 printer with VeroBlue material or on a Markforged Onyx One printer in Nylon with chopped carbon fiber (Onyx Material, Markforged).

The interior geometry of the finger is created using two soft-bodied cores. These cores are made using a typical molding process. First, Elasto-Sil M-4601 (Wacker) is mixed, then poured into both sides of the mold. Next, the mold halves are de-gassed in a vacuum chamber, and steel alignment rods (2mm diameter) are placed. The mold is then clamped together and placed in a 65 °C oven for three hours until fully cured.

The body of the finger is created by a four-part mold, as shown in Figure 2.4b, which is constant across all relative segment lengths. The relative lengths of the finger's two segments are instead chosen by adjusting the lengths of the soft cores before molding, as shown in Figure 2.4c. To build the body, the mold is filled halfway with Smooth-Sil 950 (Smooth-On Inc). After degassing, the proximal soft core (with alignment rods) is inserted into the mold, then fixed with the proximal clamping piece. Next, the distal alignment piece is placed, followed by the distal soft core. More silicone is subsequently poured to cover the cores completely. The mold is then clamped together between two aluminum plates and placed in a 65 °C oven for three hours until fully cured.

To plug the distal end of the finger, wedge-shaped fingertips are attached. The fingertip is created using Smooth-Sil 950 and the same basic molding process as the soft cores (without alignment rods). The piece is then attached to the distal segment using interlocking features and silicone adhesive (Silpoxy, Smooth-On, Inc). As discussed in Section 2.3.1, these pointed fingertips are commonly used for improved power grasping against surfaces [80].

Finally, a rigid hub is attached to the proximal side of the finger body to enable air delivery. The holes created by the alignment rods for the distal core also act as fluid channels to deliver air to the distal segment. Thus, both fluid connectors can be located in the proximal hub. The hub is fixed to the skin using interlocking features and Silpoxy. Once the adhesive is cured, heat-shrink tubing is wrapped around the proximal end of the finger to ensure no leaks.

2.3.3 Designing a Rigid Palm

For all subsequent testing and analysis, only grasping in a plane will be considered since it directly illustrates the benefits of including multiple bending segments in soft fingers.

To design a suitable rigid, planar palm, two competing design criteria were considered. Grasping objects of zero-width is only possible if the distal parts of the actuators come into contact. However, the resting position of the actuators should also have a wide opening angle to be able to grasp comparatively larger objects (see Figure 2.3).

Thus, two design parameters are free to be chosen: the width at the base, w , and the angle between fingers, α . Since the actual shape of the soft fingers can be complex during a grasp, several candidate palms were built and a single palm was chosen based on empirical testing (see Characterization of Fingers and Palm). All candidate palms were 3D printed on an Object Connex 500 printer with VeroBlue material, and the final design was printed on a Markforged Onyx One printer in Nylon with Onyx Material.

2.4 Characterizing Fingers and Palm

Characterizing the kinematic and mechanical properties of the fingers is critical before we can understand the grasping behavior of the hand as a whole. We first defined two actuation modes that allow the fingers to exhibit fundamentally different behavior while grasping. We then characterized the curvature and blocked force responses of individual segments under actuation pressure, and evaluate the maximum pressure before failure to determine a pressure operating point. We also measured the stiffness of each segment to confirm they are similar. Due to sufficient similarity between both segments, we can achieve finger motion similar to a single bending segment when equal pressures are applied to both segments. Finally, we use all of this information to design the rigid palm to be used for robust grasping.

2.4.1 Actuation Modes Replicate Finger Structure

To simplify the combinations of actuation inputs, we restrict our focus to two actuation modes that enable the fingers to replicate two fundamentally different finger structures, as shown in Figure 2.5. In the first actuation mode, (“*proximal-actuation*” mode), only the proximal segment is actuated while keeping the distal segment passive. In the second actuation mode, (“*uniform-actuation*” mode), both segments are driven with equal pressure so that the actuated region spans the entire length of the finger. This enables our prototype fingers to achieve the same behavior as fingers built with a single uniformly-actuated segment. Thus, using only pressure control, our two-segment fingers can be used




Figure 2.5: Fingers with two segments can be actuated a number of ways. A finger is shown (a) at rest, (b) with *proximal-only* actuation, (c) with only distal actuation, and (d) with equal pressure in both segments (*uniform-actuation* mode)

to investigate how grasping behavior differs depending on finger structure.

2.4.2 Functional Evaluation of Fingers

We performed a series of experiments on several fingers to determine the response of individual segments to input pressure. To control the pressures independently in each segment of the fingers, we used a custom pneumatic pressure control system with an accuracy of 1.4 kPa. For each of the four channels, the controller enables smooth control of output pressure around a setpoint, and execution of arbitrary pressure trajectories in real time. A more detailed description of our pressure control system can be found in Appendix B

We first recorded the change in curvature as a function of input pressure. We performed this experiment on the proximal and distal segment as well as for the whole actuator. The actuation pressure is increased from 0 kPa to 138 kPa in 13.8 kPa increments, and the resulting curvature is measured by hand from photographs. The average curvature ranges from 0/m (flat segment) to 32.24/m and 28.54/m for the proximal and distal segment respectively, as shown in Figure 2.6a. As the pressure reaches 100 kPa, the distal segment’s curvature does not show any further significant increase. Overall, the relationship between input pressure and curvature is fairly similar between segments, with a maximum of 22% difference in curvature occurring around 70 kPa.

Furthermore, to evaluate the limitations in actuation, both segments of three separate fingers were inflated until they failed by rupturing. The recorded burst pressures

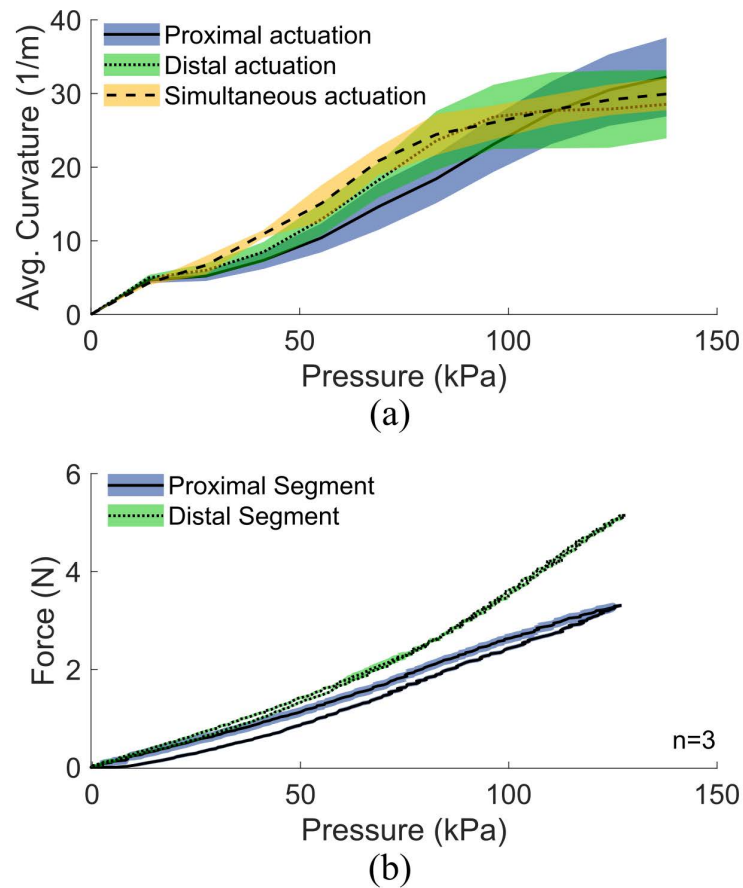


Figure 2.6: (a) The curvature of each finger segment as a function of applied actuation pressure is recorded every 13.8 kPa up to 138 kPa. (b) The blocked force as a function of applied actuation pressure shows only slight hysteresis over the 0 kPa to 138 kPa range. The mean and standard deviation of $n = 3$ trials is shown.

were (240 ± 14) kPa for the proximal segments, and (186 ± 20) kPa for the distal segments. Failures occurred in the bellows sections of both segments. To prevent structural failures during normal operation, we choose to use a maximum actuation pressure of 100 kPa.

Next, blocked force as a function of actuation pressure was measured for both segments using an Instron 5544A. Fingers are clamped in a vise and placed under the Instron, as shown in Figure 2.7b. The rigid hub is clamped when characterizing the proximal segment, and the finger itself is clamped when measuring the distal segment. Next, a thin plastic sheet is clamped in the jaws of an Instron machine to ensure a small contact point with the finger. The pressure is then applied and the resulting vertical blocked force is measured, as shown in Figure 2.6b. The difference in the slopes between the proximal and distal segments is likely due to the difference in the cross-section of the air chambers, which was discussed earlier in the “Designing Modular Two-Segment Fingers” section. Overall, the relationship between input pressure and blocked force has a similar shape for both segments.

Finally, the stiffness of each bending segment was characterized by applying small deflections at the tip of the segments and recording the resulting force. First a finger is clamped in a vise and placed under an Instron machine using the same procedure as the blocked force tests, as shown in Figure 2.7c. Next, a thin plastic sheet is clamped in the jaws of an Instron machine, and used to apply 10 mm of deflection to the tip of the segment. The resulting force is measured using a 10 N load cell. Given the linearity of the force-deflection curves, the stiffness is calculated as the slope of this line, as shown in Table 2.1.

For all of the four fingers characterized, the stiffness of the distal segment was within 33% of the proximal stiffness, with differences between fingers as low as 10%. This discrepancy in stiffness is caused partially by the mechanical design of the finger, since the wall of the proximal segment at the inside of the bend is thicker to accommodate the distal air supply channels.

To create the fairest comparison between finger structures with uniform actuation vs. proximal-only actuation, an important design goal was to ensure that each finger can achieve both structures through differences in actuation. The “proximally-actuated two segment” structure is trivial to implement by design. However, given the small magnitude of the difference in segment stiffness, combined with the similar bending and blocked force responses for both segments, we confirm our assumption that our fingers can behave like a single bending segment through *uniform-actuation*.




Figure 2.7: Characterization of stiffness and blocked force for each segment was performed on an Instron system. (a) The finger is clamped in a vise. (b) Blocked force is measured by applying input pressure and measuring the resulting force. (c) Segment stiffness is measured by deflecting the finger by 10 mm while measuring the force. In both cases, the finger presses against a thin plastic sheet.

Table 2.1: Bending stiffness of individual bending segments of fingers. The mean and standard deviation are reported for $n = 3$ trials of each sample.

Finger Sample	Stiffness (N/m)	
	Proximal	Distal
#1	84.2 ± 0.8	66.3 ± 0.2
#2	83.1 ± 0.4	92.4 ± 3.4
#3	75.3 ± 0.3	50.4 ± 0.2
#4	70.4 ± 1.0	57.9 ± 0.9
Average	78 ± 6	58 ± 8




Figure 2.8: Several palms with varying angles were evaluated at rest and with the proximal segment pressurized to 100 kPa. We chose an opening angle of 30 degrees, as this is the largest angle where the fingers achieve the nonzero distal contact area when actuated.

2.4.3 Finding a Suitable Palm Angle

Choosing an appropriate angle between fingers (palm angle α) and palm width, w , is critical to allow for robust pinch grasps while also maximizing the largest attainable object diameter. Since our soft fingers are limited to 100 kPa input pressures to prevent actuator failure, the resulting curvature is also limited. Thus, the geometry of the palm must ensure the fingers touch when only the proximal segment is actuated, while simultaneously achieving the widest possible distance between fingers at rest.

To maximize the potential for fingers to touch under *proximal-only* actuation, the palm width was chosen to be relatively short (15 mm). This short distance between the bases of each finger could potentially affect the overall robustness of power grasps. In fact, many other hand designs (including humans) utilize the palm as a contact surface during power grasps. However, our primary goal is to compare finger designs during both precision grasping and power grasping, so the palm width is less important.

To find the finger angle that meets these criteria under actuation constraints, we evaluated the contact area at the fingertips under a proximal actuation of 100 kPa on several prototype palms. Five palms were tested, ranging from an angle of 60 degrees to 20 degrees with a constant base width of 15 mm, as shown in Figure 2.8. Ultimately, an angle of 30 degrees between fingers was chosen due to the larger distal contact area under proximal-only actuation.

2.5 Characterizing Grasping Performance

For each pair of two fingers and the final palm design, we evaluated the effect of actuation modes on several relevant grasping metrics. We first evaluated the effect of hand placement (with respect to the object) on the type of grasp performed. These tests also enabled evaluation of the range of object diameters the hand can grasp. We then explored the hand’s robustness to external forces. Finally, we measured the finger compliance during grasping. All tests were performed for both actuation modes.

2.5.1 Hand Placement

For our soft fingers, the type of grasp performed is determined by the placement of the hand with respect to the object, and the actuation mode used. To evaluate this effect, we performed a series of grasp attempts on a set of cylindrical objects to determine the ranges of centering positions that cause power grasping, pinch grasping, or failure. In addition, these experiments also yield the region of acquisition for different objects along the axis normal to the palm.

To perform reliable grasping at precisely controlled positions, the hand was mounted to one of two Cartesian positioning systems: either a custom-built three-axis CNC gantry, or a UR5e 6DOF robot arm (Universal Robots, Denmark), each with a positioning accuracy better than 1 mm. For the gantry system, GCODE commands were used to command hand positions, while for the robot arm, MoveIt! [81] was used for motion planning. Robot Operating System (ROS, [82]) was used to coordinate motion and hand pressure control for both systems. To maintain a symmetric grasp, we actuate the homologous segments of both fingers with the same actuation signals, assuming symmetry between fingers. Pictures of both experimental setups are shown in Appendix C

The set of objects used in this study was chosen to reach both ends of the size spectrum that our soft hand can grasp. A set of eleven cylindrical objects ranging from 2.2mm to 116mm in diameter were chosen. Most of the objects belong to the Yale – Carnegie Mellon – Berkeley (YCB) object set [83], and a few extra objects were added to fill in gaps in the smaller size range, as shown in Figure 2.9. The actual objects used are discussed in Appendix C.

A typical test for a given object involves grasping an object with a known centering offset, then checking for relative motion, as shown in Figure 2.10. First, the object is

manually placed at a precise position on a low-friction table. Next, the hand approaches with some centering position, x_h , and attempts to grasp the object. The hand is then raised 10 mm to lift the object. Finally, the object is moved in a zig-zag pattern to determine if any relative motion between the hand and the object occurs. For a grasp to be considered successful, the object must remain in the same position before and after the zig-zag motion (thus, caging grasps are not considered a success in our testing).

For each successful grasp, the grasp is characterized as a power or precision grasp based on the number of contact points between the finger and object. Keeping consistent with our definitions of grasps found in Section 2.1, precision grasps involve one contact per finger, and power grasps involve more than one contact on at least one finger (at least three contacts total). Power grasps can also involve a large area of contact, which reduces to a line of contact between a finger and object in planar space.

The effect of centering position on grasp type was evaluated for both key actuation modes (*uniform* actuation vs. proximal-only) over the entire range of object sizes and centering positions. Centering positions ranged from 40 mm to 180 mm in 5 mm increments, as measured from the front of the palm to the center of the object. In addition, each object has a limit to how close it can be placed to the palm due to geometric constraints, so testing was restricted to positions that were geometrically reachable. Results of these experiments can be found in Figure 2.12. In addition, a summary of the range of object diameters capable of being grasped is shown in Figure 2.13.

2.5.2 Robustness to External Forces

The robustness of a grasp to applied external forces can be used as a metric to empirically evaluate the stability of a grasp. Grasp robustness, as defined by Ferrari et al. [84], is the maximum force on the object that the gripper can resist before the object is pulled out of the hand. A minimum is taken over all possible angles the force can be applied to find the “worst-case” situation. We measured the robustness empirically by grasping an object and measuring the force required to pull it out at different angles.

For this test we used a custom-built fixture to hold the hand at angles spanning from 0 (vertical) to 90 degrees (horizontal), as shown in Figure 3 in Appendix C. The fixture allows the position and orientation of the hand to be precisely controlled with respect to the target object. The object is then pulled vertically on a uniaxial testing machine (Instron



Figure 2.9: Objects used for grasping characterization. Objects are part of the YCB object set except the two bolts, the syringe, and the tube grommet.




Figure 2.10: Experimental setup for exploring grasp types over a range of centering offsets (x_h) and object diameters (D). (a) The object is grasped with some centering offset, then (b) the hand is moved in a zig-zag pattern to check if the grasp is successful.




Figure 2.11: Experimental setup for measuring finger stiffness while grasping, and associated schematic diagram. A single finger was actuated to perform half of a fingertip or pinch grasp on an object. Known deflections, δ , were applied to a finger at an angle θ , and the resulting force was measured. Forces applied at a desired angle were accomplished by mounting the hand at an angle, and using an Instron uniaxial testing machine to command deflections.

5544A), and the force is recorded simultaneously. Neglecting the effect of gravity, this setup is equivalent to pulling an object out of the grasp at the desired angle. As a target object, we chose an acrylic cylinder with a diameter of 50.8mm since it is in the middle of the object size range.

As mentioned in the previous section (Hand Placement), the position of the fingers with respect to the target object defines the type of grasp performed. Using this information, we chose to test two different object positions where precision grasping and power grasping occur: at the fingertips, and at the midpoint of the fingers. Grasping under *uniform* actuation, yields fingertip grasps when the object is placed at the fingertips, and power grasps when the object is placed at the midpoint of the fingers. Grasps under *proximal-only* actuation yeild pinch grasps for both object positions. We measured the force for each situation over a range of pulling angles (0, 15, 30, and 45 degrees).

2.5.3 Fingertip Compliance During Grasping

To estimate the stiffness of the finger at the contact point during a grasp, we measured the force generated by small deflections of the finger over several angles. The hand was first mounted at an angle near an Instron Uniaxial testing machine using the same fixture as was used to measure grasp robustness. Next, a 25.4 mm tube was positioned such that precision grasps could be performed using both modes of actuation (approximately

150mm from the center of the palm). Only one finger was actuated against the object, forming half of a fingertip or pinch grasp, as shown in Figure 2.11. Finally, small deflections ranging from 1 mm to 5 mm were applied and the resulting increase in force was measured. Three trials for each angle were performed.

To obtain the stiffness of the finger as a function of the angle at which force was applied, the slope of the force-deflection curve was found using linear regression for each trial. The range of angles tested includes 0, 15, 30, 60, and 90 degrees (as defined from the axis normal to the front of the palm). In addition, the small deflections of 1 mm to 5 mm were chosen to avoid slipping of the finger on the object. Results from these experiments are shown in Figure 2.16.

2.5.4 Relative Segment Lengths

To evaluate the grasping performance of fingers as a function of relative segment lengths, a subset of the tests presented above were conducted for fingers with 0.3 and 0.7 distal segment length fractions. Grasp success regions were evaluated with a subset of the objects (#2-56 bolt, 1/4-20 bolt, small marker, tube grommet, cup 1, cup 6, cup 10, and the pitcher), and robustness to external forces were measured as before. Combined with the more-detailed evaluation of fingers with equal segment lengths (0.5 distal length ratio), we aim to evaluate the trade-offs between various aspects of grasping performance as a function of segment lengths.

2.6 Results

The grasping performance of each set of fingers is evaluated on the basis of four metrics that can be used to compare the grasp quality and utility in a manipulation system. These metrics include a simplified estimate of the region of acquisition, an estimate of the range of object sizes that can be grasped, an estimate of the rotational stability, and the robustness to external forces on the object. In addition, the finger stiffness during grasping is characterized.

2.6.1 Hand Placement Determines Grasp Type

The region of acquisition describes how much error in hand position can be tolerated before it is unable to perform reliable grasps. As defined by Aukes et al. [85], the region of acquisition is the set of all hand positions (relative to a target object) where successful grasps can occur. We empirically measured a single axis of this region along the axis of symmetry for our hand (the axis normal to the palm). In addition, we tracked how the type of grasp is affected by hand placement, allowing the formation of regions of acquisition for each grasp type and each actuation mode. The results of these experiments are shown in Figure 2.12.

Overall, fingers with a passive distal segment (under *proximal* actuation) can perform precision grasps over a larger range of centering distances than with *uniform* actuation. The width of the success region for pinch grasping with *proximal* actuation is three to four times the width of the fingertip grasping region with *uniform* actuation for objects in the middle of the diameter range. Furthermore, the range of centering distances that yield stable grasps increases dramatically for objects under 40 mm in diameter. In fact, the pinch grasping region for fingers under *proximal* actuation includes some objects in the smaller range that could not be grasped with *uniform* actuation.

These results follow from the geometry of the fingers during a grasp. With a passive distal segment (*proximal* actuation), the object can be grasped anywhere along the distal segment, whereas contact points are limited to only the fingertips when using *uniform* actuation. Thus, we would expect a drastic increase in the size of the pinch-grasping region because the passive distal segment can grasp in positions that could only be caged using *uniform* actuation.

Conversely, fingers under *uniform* actuation can perform power grasps over a larger range of centering distances. The power grasping region for *proximal* actuation appears to shrink by approximately 70% on average for larger objects. This makes sense because much of the centering distances where power grasps are performed with *uniform* actuation result in pinch grasps with *proximal* actuation.

In addition, geometric considerations can explain the failure regions. For example, all grasps fail in the region beyond 160 mm centering distance, as this is beyond the reach of the fingers. In addition, some regions were unable to be tested due to geometric constraints. For objects 47 mm and larger, the empty region to the left of the power grasping region is

physically impossible to test.

One special case for our finger design involves fingertip grasps on small objects using *uniform* actuation. For small objects (25 mm and 16 mm), rolling instabilities on the fingertips can cause marginally stable grasps that snap to one side or the other. This forms grasps where the object is in contact with the fingertip of one finger and the back side of the other finger. An example of this type of grasp is shown in Figure 2.17 for fingers with distal segment length fraction of 0.0 grasping a tube grommet. For the purposes of this analysis, we consider these types of grasps as marginal failures, since the final object pose is not predictable.

Finally, it is important to note that for the small objects that could not be grasped by fingers under *uniform* actuation, failures involved rotational instability. All failures for objects smaller than 16 mm around an approximately 150 mm centering offset involved the fingertips applying force slightly off-center due to small differences in actuator performance. This slight off-center force balance caused the object to undergo large rotations, and thus large fingertip motions, eventually pushing the object out of the grasp.

2.6.2 Object Size Range

The range of object sizes capable of being stably grasped is another metric we can use to evaluate the effect of additional bending segments in our soft fingers. We can extract this metric from the results of the experiments performed in the previous section by identifying the largest and smallest objects that could be grasped for a given grasp type and actuation mode. A summary of the range of diameters that can be successfully grasped using each combination of actuation modes and grasp types are shown in Figure 2.13.

According to our measurements, the upper bound on object size is similar regardless of the actuation mode or type of grasp. Overall, the largest possible object that can be grasped is 116 mm in diameter. This makes sense because we expect the upper bound to be limited by hand geometry.

When operating with *uniform* actuation, a lower bound on object size exists. The smallest object that was successfully grasped is 16 mm in diameter, but this occurs only when performing a marginally-stable fingertip grasp. In this case, fingertip grasping is necessary because the diameter of the object is smaller than the minimum diameter that can be power grasped. The lower bound makes sense because fingertip grasp stability is a

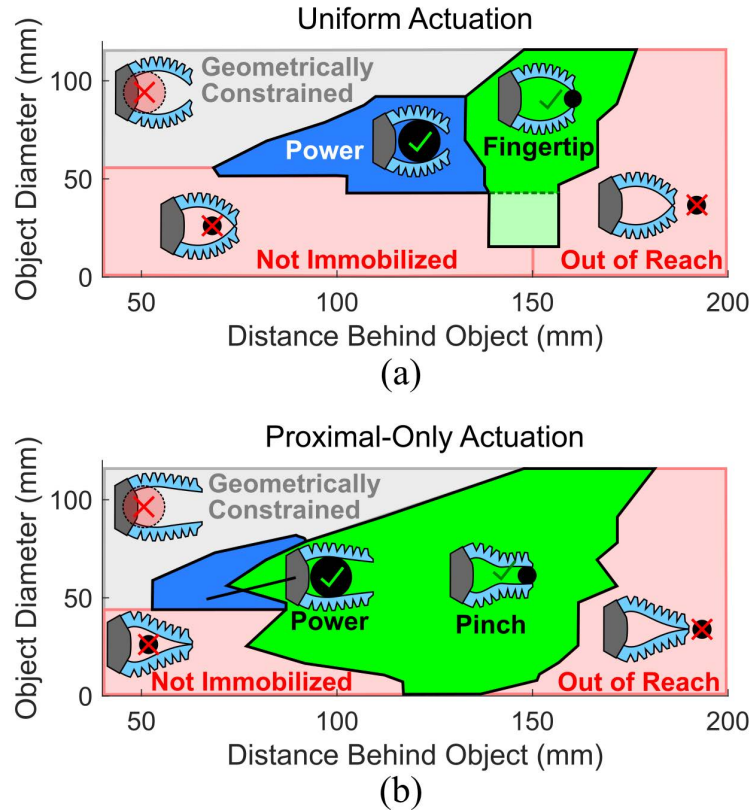


Figure 2.12: Hand position and object diameter determine the type of grasp produced upon actuation, as well as failure regions. (a) For *uniform* actuation, the power grasping region is large, but only marginally-stable grasps were observed for the 25 mm and 16 mm objects near a centering distance of 150 mm. (b) Under *proximal-only* actuation, the pinch grasping region encompasses a larger range of centering distances, and spans all the way to the thinnest object tested. All grasps were performed using fingers with equal-length segments.

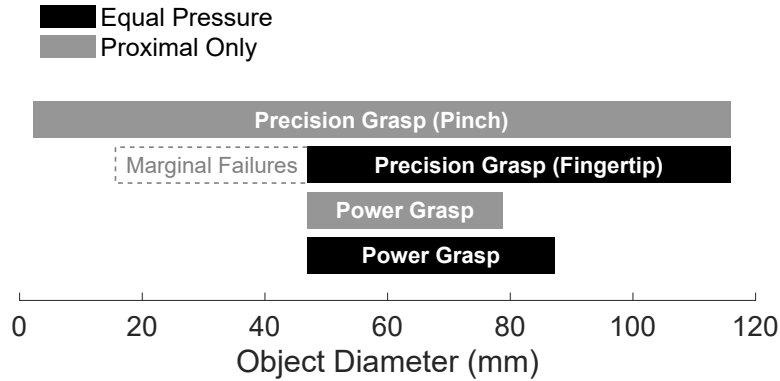


Figure 2.13: Summary of successfully-grasped objects for each actuation mode and grasp type using fingers with equal length segments. Under proximal-only actuation, the hand can grab smaller objects. The smallest grasped object is a #2-56 bolt. The dotted bar represents the size range if marginal failures were counted as successes.

function of the object and fingertip curvatures, making grasps on smaller objects unstable.

Alternatively, when operating with a passive distal segment (*proximal* actuation), the lower bound on object size appears to be arbitrarily small. Successful pinch grasps were capable of being performed on arbitrarily thin objects (such as a sheet of paper). This is due to the fact that the point of contact with objects is on the inside of the finger, and the passive distal segment can bend backward to achieve approximately zero curvature (flat plate). Thus, grasps are kinematically stable regardless of object diameter according to [28], and are instead limited by contact forces.

Overall, the results presented in the last two sections represent fundamental performance tradeoffs arising from finger structure. It is clear that having a passive distal segment (with the proximal segment actuated) is directly responsible for better functionality compared to fingers with one uniformly-actuated bending segment. Fingers with one uniform bending segment are unable to grasp objects below some minimum diameter due to rotational instability. However, fingers with passive distal segments can achieve stable pinch grasps on arbitrarily thin objects without affecting the maximum possible diameter.

2.6.3 Rotational Stability

Rotational stability plays a large role in the overall success of precision grasps. We can observe the effects of rotational stability (or instability) on grasps by tracking the

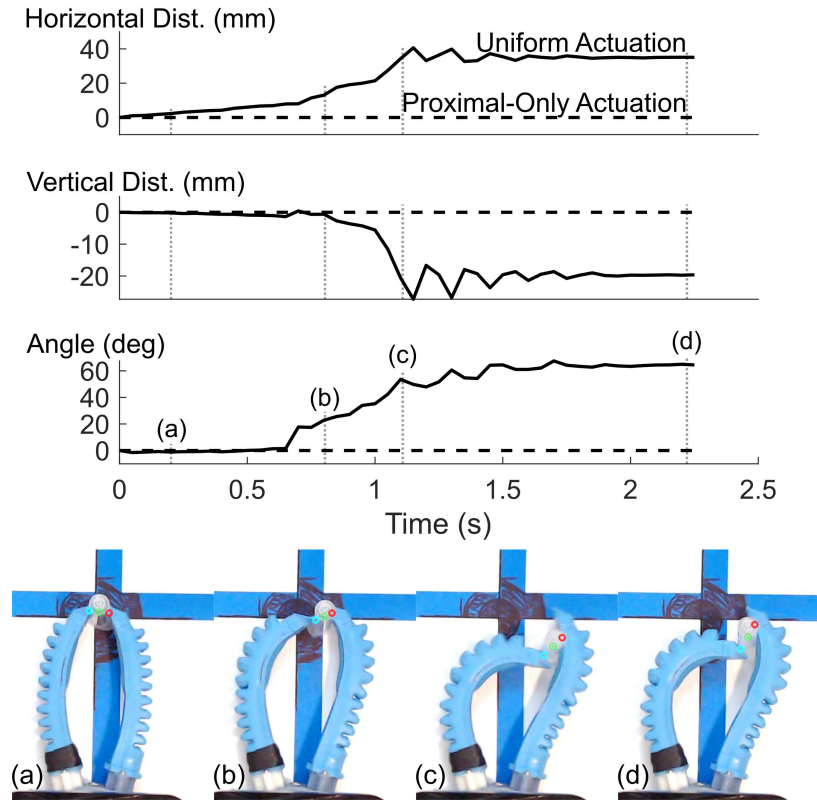


Figure 2.14: Rotational instability during a fingertip grasp causes large object rotation and translation. Pictures correspond to time points in the graphs as marked. The object grasped was a 16 mm syringe.

object’s pose over time. For both actuation modes, the position and orientation of a small object was tracked as a precision grasp was attempted. The positions of contact points were tracked manually for each frame in the videos using Tracker Video Analysis and Modeling Software [86] as described in Appendix C. The results are shown in Figure 2.14.

It can be clearly seen from the motion of the object that the rotational stability of a fingertip grasp under *uniform* actuation is lower than that of a pinch grasp under *proximal-only* actuation. During the grasp attempt, the object’s pose changes rapidly, with the angle of the object changing by 60 degrees from the starting position. This is due to rotation between the fingertip and the object. In contrast, the object neither moves nor rotates at all when a pinch grasp is performed with *proximal* actuation. This is because the object is grasped on the inside edge of the flat segment, so object rotation cannot occur unless the object slips with respect to the finger.

2.6.4 Robustness to External Forces

Robustness to external forces is critical for maintaining a grasp once it is successful. To understand how the grasp robustness differs for different actuation modes and grasp types, we look at the results of the pull force tests displayed in Figure 2.15.

It turns out that when a grasp occurs at the fingertip, grasping under *uniform* actuation shows a greater average resistance to forces applied to the object compared to pinch grasping with *proximal* actuation. However, when the cylinder is placed deeper in the pinch-grasp, the pull-out force is on average 96% higher than a fingertip grasp using *uniform* actuation. When it comes to a power grasp using *uniform* actuation, the required force is significantly higher (5.17 N on average over all angles). In all four situations, the angle (in this range) does not appear to be a consistent factor leading to any significant change in the pull force.

As before, these results using the two actuation modes are indicative of the performance of the two fundamental finger structures we are studying. These results indicate that if an object is large enough to be power-grasped, fingers with a single uniformly-actuated segment perform far better than fingers with two segments where the distal segment is passive. However, when an object is too small to be power grasped, fingertip grasping with single-segment uniformly-actuated fingers performs better than pinch grasping with passive distal segments. Finally, if the object is too small to be grasped with uniformly-actuated fingers, pinch grasps using fingers with a passive distal segment are more robust as the object is placed deeper in the grasp.

2.6.5 Grasp Stiffness

Overall, the finger stiffness during grasping was higher with two-segment fingers than with single-segment fingers. The results of these experiments are shown in Figure 2.16. At 0 degrees (aligned with the axis of symmetry), the stiffness of the finger under *uniform* actuation is 410 ± 30 N/m. Meanwhile, the pinch-grasping stiffness for fingers under *proximal-only* actuation is 1200 ± 110 N/m. For angles of 0, 15, and 30 degrees, the finger with a passive distal segment had a stiffness on average 2.9 times that of the single-segment finger. The variability in stiffness measurements was on the order of 8-9% at 0 degrees, and less than 2% at all other angles. Interestingly, stiffnesses were much lower for both finger designs at more extreme angles of 60 and 90 degrees.

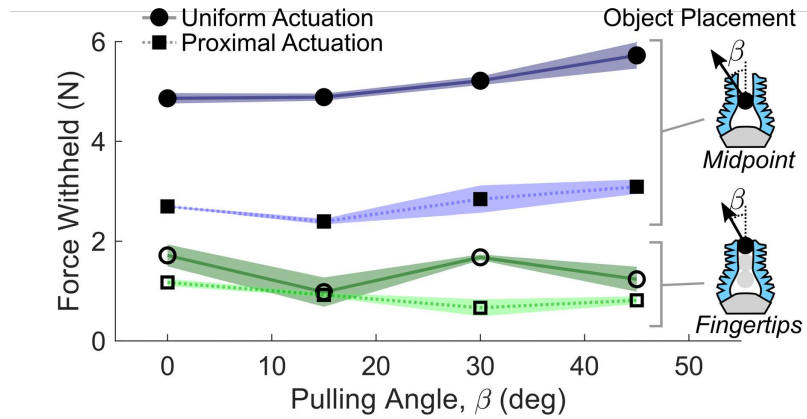


Figure 2.15: The average minimum force to pull out a 50.8 mm cylinder for different grasps using fingers with equal-length segments. Error bars represent one standard deviation over $n = 3$ runs.

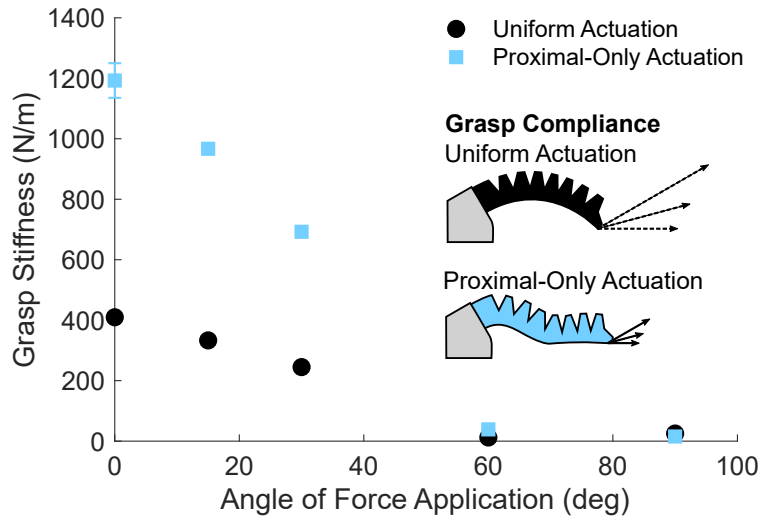


Figure 2.16: Stiffness of the fingers as a function of the force application angle. The stiffness when performing a pinch grasp using *proximal* actuation is roughly three times larger than the fingertip grasps using *uniform* actuation. The diagram shows the deflection for a constant force applied at angles of 0, 15, and 30 degrees. Magnitude of the vectors are exaggerated for clarity, using an equivalent of 15 N applied force. Error bars represent twice the standard deviation over $n = 3$ trials.

2.6.6 Relative Length of Finger Segments

To evaluate the effect of relative segment lengths on grasping performance, we performed hand placement and grasp robustness tests on two additional length ratios (0.3 and 0.7 distal length fraction) under *proximal-only* actuation. The results of these tests can be directly compared with the original fingers with equal segment lengths (0.5 distal length fraction) under *proximal-only* actuation. In addition, fingers with equal segment lengths can simulate the performance of similar fingers with one uniformly-actuated bending segment (0.0 distal length fraction) when actuated with *uniform* pressure. Representative samples of the resulting grasps for some of the objects tested are shown in Fig. 2.17.

From a geometric standpoint, the region of grasp success transitions from mostly power grasps to entirely precision grasps as the relative length-fraction of the distal segment (distal segment fraction) increases, as shown in Fig. 2.18. With a distal segment fraction of 0.0, power grasping occurs when the palm is between 70 mm and 135 mm behind the object, and precision grasping occurs between 135 mm and 170 mm. By contrast, fingers with a distal segment ratio of 0.7 exhibit no power grasping region, and precision grasps occur between 60 mm and 165 mm behind the object. In addition, increasing the distal segment length increases the range of successful hand positions for smaller objects.

Our experiments evaluating the forces during grasping show that the grasp robustness generally decreases as the distal segment length fraction increases, as shown in Fig. 2.19. Robustness was tested with objects placed at the midpoint of the fingers and at the fingertips. In both cases, the minimum force withheld appears to decrease as a function of the distal segment fraction. This relationship appears distinct when the object is placed at the midpoint of the fingers, while the robustness of fingertip grasps is similar for all distal segment fractions tested. In addition, grasps on the object at the midpoint of the fingers are overall stronger than with the object placed at the fingertips. This makes sense due to the larger lever arm over which contact forces are transmitted, and the shorter actuated proximal segment as the distal segment length increases.

2.6.7 Grasping Arbitrary Objects

In addition to cylindrical objects, we tested grasps on rectangular prisms of varying thicknesses from 2 mm (cardboard sheet) to 60 mm (a square box). In all cases, prisms were placed with the sides parallel to the axis of symmetry of the hand. Representative samples

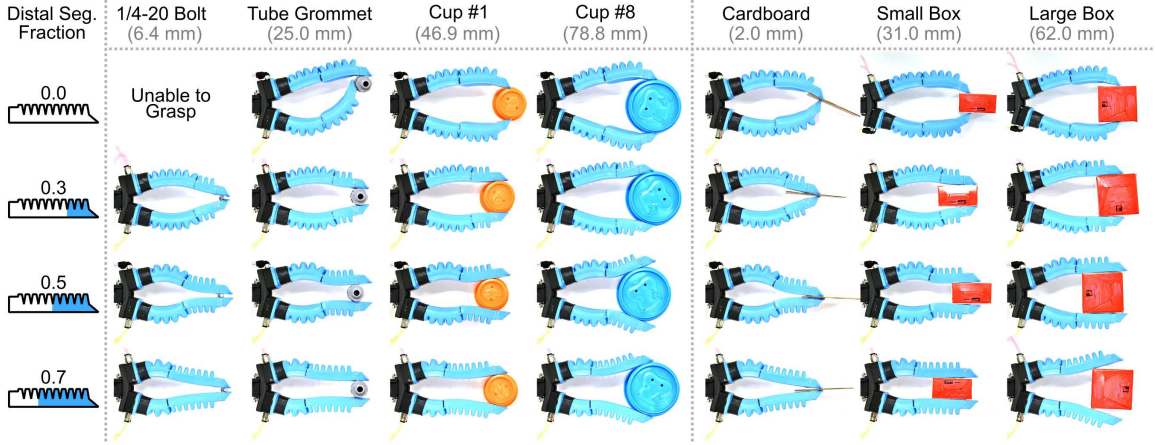


Figure 2.17: Precision grasps performed on cylinders and rectangular prisms of a variety of widths. The general shape of the fingers during grasping is relatively constant for each object as the distal segment length fraction decreases from 0.7 to 0.3. However, a fundamental difference in shape occurs for 0.0 (no distal segment) compared to fingers with distal segments

of the resulting grasps are shown in Fig. 2.17. As expected, the shapes of fingers during grasping are not substantially different from grasps performed on cylindrical objects of similar width. In addition, the results of hand position tests for this set of rectangular prisms are shown in Fig. 2.20.

In addition to similar finger shapes during grasps, the stability of precision grasps on thin objects (judged by examining the amount of extraneous object motion during the initial grasp) is still increased when using fingers with a passive distal segment as compared to no distal segment. For example, grasping a thin plate with no distal segment causes large object rotation due to the small radius of curvature of the fingertips. Conversely, grasping with fingers that have passive distal segments causes minimal object motion. This is the same trend as was observed with cylinders.

While the benefit of a passive distal segment on grasp stability is similar for prisms and cylinders, the success region for fingers with no distal segment is substantially improved for small objects. Using fingers with no distal segment, a 2 mm thick cardboard sheet is easily grasped. This is due to the aid of a third contact point at one of the fingers. In addition, the radius of curvature of the cardboard sheet near the fingertips is large (essentially infinite), and the distance between contacts is small, leading to higher rotational stability compared to a cylinder with a similar diameter.

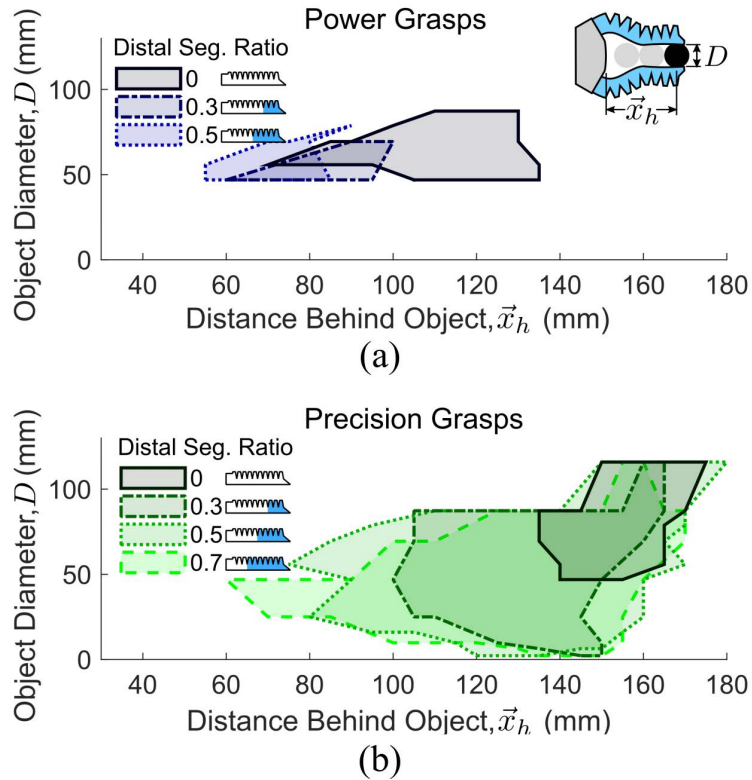


Figure 2.18: The relative lengths of finger segments affect which grasp types are successful under proximal-only actuation. a) The power grasping region appears to shrink as the length of the distal finger segments increases (relative to overall finger length). In fact, a distal segment ratio of 0.7 exhibits no power grasping region at all. b) Conversely, the precision grasping region appears to expand as the distal segment length increases. Shaded regions represent successful grasps.

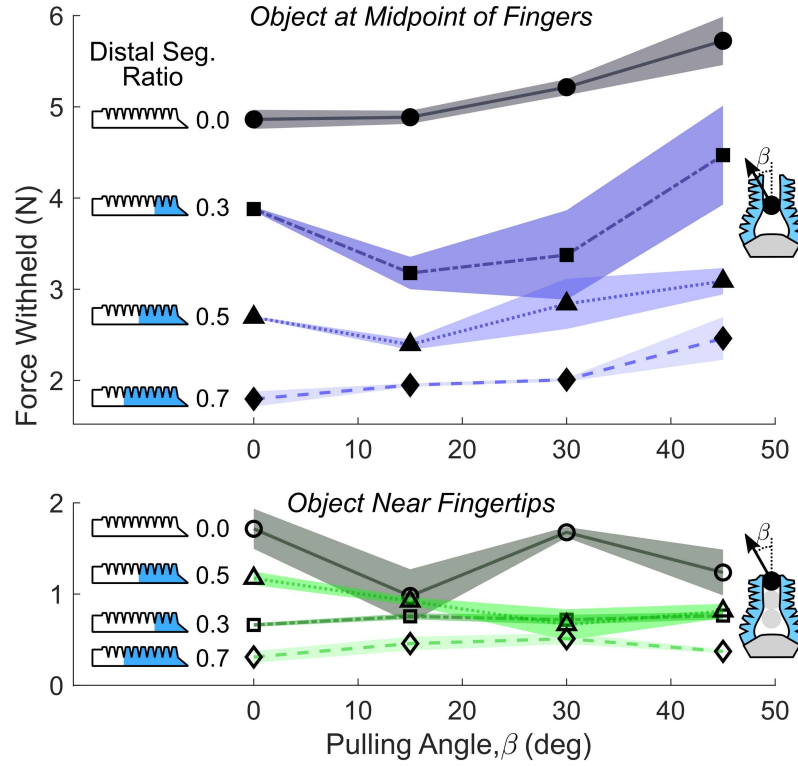


Figure 2.19: The relative lengths of finger segments affect the grasp robustness under proximal-only actuation for two different object placements. With the object at the midpoint of the fingers, grasps are overall stronger (more robust) with a length ratio of 0.0 providing the most robust grasps. With the object near the fingertips, grasps are overall weaker (less robust) with the length ratio of 0.7 being the least robust. In both cases, the minimum force withheld generally decreases as the distal length increases. Grasps were performed on a 50.4 mm cylinder, and shaded regions represent the standard deviation in the force withheld over $n = 3$ trials.

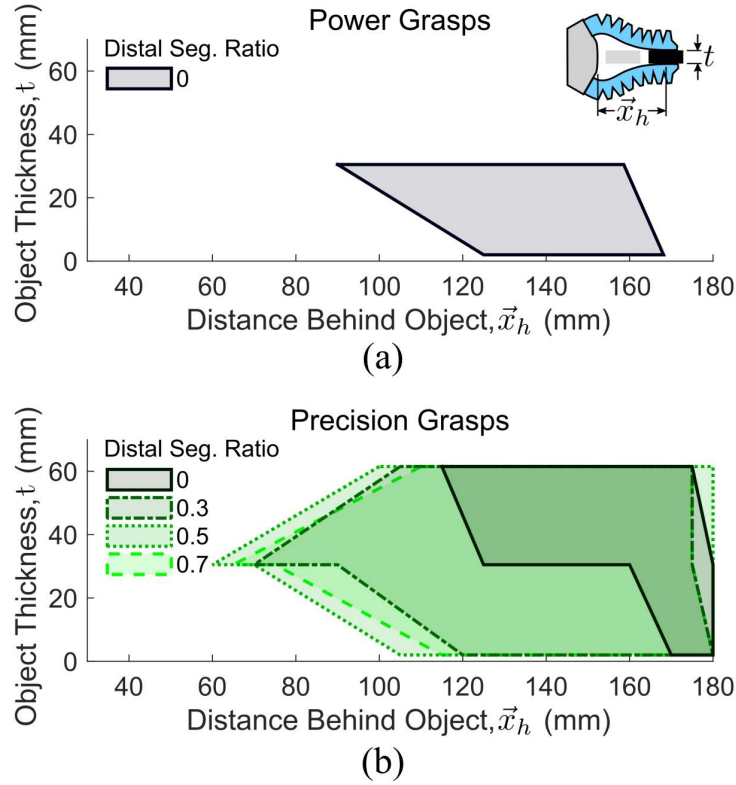


Figure 2.20: Grasping rectangular prisms does not substantially change grasping performance compared to cylindrical objects. a) Due to object geometry, only fingers with a distal segment ratio of 0 was capable of power grasps on the rectangular prisms tested. b) The range of centering distances that result in successful precision grasps is larger for smaller objects compared to similarly-sized cylinders.

The other main difference in grasping rectangular prisms vs. cylinders is that precision grasps are successful over a much larger range of hand centering positions, as shown in Fig. 2.20. This is because the precision grasping region for rectangular shapes extends for the entire length of the object, whereas precision grasps are often not successful for cylinders until the fingertips pass the midpoint of the object. The opposing sides of rectangular prisms are best grasped by parallel forces from the fingers, which can be generated robustly by fingers that have a distal segment (distal segment fractions greater than 0.0). In this way, these fingers behave similar to a parallel-jaw gripper.

2.7 Discussion

In this section, we analyze the results of the grasp performance characterization from above. In addition to confirming the proposed high-level design principles, we also confirm the reasoning behind them, and discuss performance tradeoffs that arise. We then generalize our results to arbitrary objects, since we expect trends to remain similar regardless of object geometry. Finally, we discuss how our results can be generalized to grasping in three dimensions.

2.7.1 Two-Segment Fingers Enable Robust Pinch Grasps

Based on three of the four grasping metrics, the precision grasping capabilities of a soft finger can be drastically improved using two bending segments with only the proximal segment actuated. For our prototype soft fingers, the only way to successfully grasp small objects below 16 mm in diameter is to perform a pinch grasp using *proximally-actuated* fingers. In addition, for the entire range of objects, a pinch grasp using a two-segment finger design had a much larger range of centering positions where successful grasps could be performed. The larger region of success for pinch grasping makes the hand much less sensitive to positioning errors compared to using fingertip grasps with one-segment fingers. The rotational stability of pinch grasps is also higher due to the much smaller finger curvature at the contact points. Finally, while the grasp robustness of pinch grasps was lower than for fingertip grasps, the difference was small compared to the magnitude of forces applied.

Overall, the benefit of designing soft fingers with a passive distal segment (or similar compliant mechanism at the fingertip) is clear. Using soft fingers with one uniform segment, a hand can only perform precision grasps with the fingertips. However, the two-segment structure enables pinch grasping, which has a higher utility than fingertip grasping when fingers are highly compliant.

2.7.2 Power Grasps Are Better-Performed With One Uniform Segment

Based on the grasp robustness measurements, power grasping is better-performed with a single uniformly-actuated segment compared to two-segment fingers with passive distal segments. First, the power grasping region is very small when using a two-segment structure with a passive distal segment compared to single uniformly-actuated segment, so

there are less opportunities to perform a power grasp. In most cases, power grasping is simply not possible with our two-segment finger design.

In addition, the robustness of power grasping with fingers that have one uniformly-actuated segment (or two uniform segments actuated together) is much higher than that of pinch grasping with a passive distal segment, even in the best case. The best pinch grasping performance occurs when the object is deeper in the grasp. However, the minimum pull-out force for a power grasp using a single uniformly-actuated segment was approximately 150% higher than for a pinch grasp with a passive distal segment. Thus, for our soft hand, the best power grasping performance requires a single uniformly-actuated bending segment.

2.7.3 Performance Tradeoffs Inform Design of Segment Lengths

While it is clear that fingers with passive distal segments enable more stable precision grasping, our exploration of the relative lengths of finger segments suggests that a fundamental tradeoff exists between grasp robustness and precision grasp stability. Some optimal ratio of segment lengths exists, however the solution is likely dependent on the detailed design of the fingers and task requirements, and would indubitably require compromises in performance. An understating of tradeoffs in performance space is therefore critical during the design process.

From our investigation, it is clear that precision grasping is necessary to grasp smaller objects, but this can only be achieved robustly with continuum fingers if a passive distal segment (or similar compliant fingertip structure) is employed. As shown in Fig. 2.18, fingers with no distal segment are unable to produce stable grasps on smaller objects, while even a small passive distal segment (0.3 distal length fraction) enables grasping arbitrarily small objects. Furthermore, the size of the precision grasping region increases as the distal length ratio increases, as shown in Fig. 2.18. This means that grasps can be performed over a larger range of hand positions on smaller objects with a larger distal segment fraction.

Conversely, power grasping is clearly the strongest grasping mode for larger objects, but the robustness (ability to withstand pull-out forces) of power grasps diminishes as the passive distal segment is lengthened. In fact, the power grasp robustness is highest when the finger has no distal segment (distal length ratio of 0.0). However, with a short distal segment (length fraction of 0.3), the minimum force withheld by power grasps is 35% smaller than with no distal segment.

From this investigation, we can draw the conclusion that some local optimum exists where the grasping region extends down to arbitrarily small objects while minimally affecting power grasp robustness. However, a more thorough understanding of the shape of this performance space would be needed before an optimal ratio of segment lengths can be determined for any particular application.

2.7.4 Two Independently-Actuated Segments Enable Best Performance

Based on the discussion thus far, it is clear that a hand with soft, continuum fingers can only achieve the best possible grasping performance during both power grasps and pinch grasps by using two different finger structures. The most robust power grasps occur using fingers with a single uniformly-actuated bending segment. On the flip side, the most-successful precision grasps occur during pinch grasping, which requires two bending segments with only the proximal segment actuated.

Using two actuated segments in each finger enables on-the-fly adaptation between both desired finger structures with only a small increase in control complexity. As shown in this chapter, fingers with two independently-actuated segments can replicate the performance of both fundamental soft finger structures. Thus, through control of both finger segments, we can achieve the best grasping performance of both structures.

2.7.5 Fingertip Compliance and Shape

According to our experimental investigation, we have corroborated the conceptual design analysis presented in Section 2.2. Fingertip compliance and local shape clearly play a role in the stability of grasps performed by soft robotic fingers. In addition, these features were tracked for both the one-segment and two-segment finger structures, and can be used to explain trends in stability.

First, the stability of fingertip grasps is very low when using a single bending segment, since all of the failed attempts to grasp smaller objects were caused by fingertip rolling instabilities. The rolling instabilities are caused by the extremely small radius of curvature of the fingertips used in our prototype system. Conversely, the range of successfully grasped diameters was extended down to zero when using pinch grasps with two bending segments. This is due to increased rotational stability gained when grasping with flat contacts on the side of the finger.




Figure 2.21: Online adaptation between pinch grasps and power grasps is essential for grasping tasks involving objects of small and large sizes. This is illustrated by the task of placing a bolt of 6 mm diameter into a large cup of 80 mm, then moving the cup. (a-d) Due to its small size, a pinch grasp must be performed on the bolt to move it to the cup. (e-f) After releasing the bolt, the hand performs a power grasp on the cup. (g-j) Finally, the cup is moved to a new position and released.

Next, the compliance of the fingers during a pinch grasp is smaller in the axis normal to the palm’s surface compared to a fingertip grasp. Thus, not only are rolling instabilities mitigated, but the amount of finger deflection per unit force on the object is smaller.

Overall, our experiments clearly demonstrate that consideration of simple design attributes can be used to explain why two independently-controlled bending segments (or similar compliant fingertip structures) in soft robotic fingers are fundamentally necessary. Through increased rotational stability due to low curvature at the contact point, pinch grasps enable soft fingers to grasp smaller objects, and are more robust against small perturbations. Additionally, two-segment fingers can control object motion to a higher degree through increased stiffness. All of these capabilities can lead to important functions in real use cases.

2.7.6 Grasping Arbitrary Objects

In the real world, robots need to be capable of interacting with a variety of object shapes. Rarely do robots encounter perfect cylinders or prisms in two dimensions. To

understand the changes in performance when grasping arbitrary objects, we can turn back to the stability analysis of [28].

For any arbitrary object, we can break down a grasp on that object into local object curvature near the contacts and the distance between contact points. In our experiments with cylindrical objects, the object curvature and distance between contacts are coupled due to geometry. However, in general these two parameters are decoupled. This decoupling leads to potentially far better grasping performance if the object is thin and flat at the contact points, and far worse performance if the object is thick and rounded at the contacts.

Based on this analysis, we would expect the stability of any grasp to increase as the object curvature decreases (as the sides of the object become flatter). Coupled with the ability to passively enable parallel finger segments, we would expect the increase in grasp stability for soft fingers with two segments vs. one segment to be even more drastic. In our study, we confirm this trend by the fact that grasps could be performed on arbitrarily-thin rectangular prisms using fingers with no distal segment, whereas grasps on cylinders with the same finger structure had a lower bound on object size that could be successfully grasped.

Overall, the design principles laid out in this work for how to utilize multi-segment continuum fingers for pinch grasping appear to generalize to a wide variety of object shapes. While complicated shapes may change the magnitudes of the trends found in this study, we expect the general trends to remain similar. As we showed in our comparison of cylinders vs. rectangular prisms, the shapes of fingers during grasping were similar for similarly-sized objects. In addition, a general increase in the stability of precision grasps when using passive distal segments also remained true regardless of object shape. Based on this evaluation, we expect the fundamental design principles discussed in this chapter to remain similar for arbitrary objects.

2.7.7 Extension to Non-Planar Grasping

In addition to arbitrary objects, real-life manipulation tasks involve moving objects in six dimensions (translations and rotations about all three axes). While our analysis and experimental validation of finger design was performed for the planar case, we can extend our results directly to real-world conditions in two important cases.

One simple extension involves using a planar hand to perform antipodal grasps,

but moving the object in 3D space. Antipodal grasps are commonly performed in robotic manipulation. If the off-axis stiffness of soft fingers is high enough, grasping can be performed in a plane with minimal fingertip deviation, and translations and rotations in 3D space become trivial, as shown in Figure 2.21 and the Supplemental Video. In this case, the finger design rules presented in this study can be directly used to build a robust antipodal gripper.

Finally, we also expect our results to apply to grippers with radial symmetry, similar to the work in Chapter 4. With fingers arranged radially, pinch grasps on axisymmetric target objects would be functionally similar to pinch grasps performed by our planar hand. This becomes useful in our subsequent hand designs in Chapter 4, and especially relevant for our exploration of digit arrangement in Chapter 5. Overall, we expect the analysis and empirical results of this study to remain structurally similar for object motion in 3D space.

2.8 Conclusions and Future Work

We explored how simple design rules can produce soft robotic fingers capable of excellent precision grasping without sacrificing power grasping performance. We presented and validated a conceptual analysis of grasping using soft fingers with multiple serially-linked bending segments. Through this analysis and an extensive empirical investigation, we showed that designing different finger structures for each grasp type clearly outperforms any single finger structure. We found that pinch grasps have increased stability compared to fingertip grasps, and achieving pinch grasps requires fingers with at least two bending segments, though only the proximal segment needs to be actuated. Further, we showed that robust power grasping requires fingers with one uniformly-actuated bending segment.

Based on this investigation, we showed that fingers with two independently actuated segments can gain the best functionality of both finger structures through a small increase in control complexity. Finally, we demonstrate the necessity of online choice between power and precision grasps during a pick-and place operation, and discuss extensions of our work to arbitrary objects and three-dimensional manipulation.

Nonetheless, there remains potential for future work exploring the capabilities of these multi-segmented soft fingers. The benefits of intrinsic compliance become most-relevant when soft actuators are dealing with fragile objects, or targets with complex morphologies. Furthermore, the design concepts presented here could be further refined by

exploring finger designs with non-uniform stiffness. Extending hands beyond planar configurations to explore three-dimensional grasps using soft, two-segment fingers is also a natural next step.

In addition to the mechanical design of two-segment soft fingers, mathematical models and sensory feedback could enable finer tuning of grasps, or even in-hand manipulation. Models of finger deformation under contact could be used as a design tool, or to aide in the generation of grasping strategies. On-board shape estimation could provide insight into how local deformations lead to successful grasps. Contact sensing could improve the sensitivity of pinch grasps, and potentially enable success estimation without the need for external vision systems. Overall, the work in this chapter sets the stage for high-quality grasping using soft robotic hands.

3

The role of multi-dimensional compliance in fabric manipulation

3.1 Introduction

Grasping and manipulating thin, flexible objects (fabric, tape, bags, etc.) is an essential skill for robots to achieve in the home, in built settings, and more-remote environments. Assistive tasks such as folding clothes, using towels to clean messes, making a bed, and handling some foods (tortillas, pizza dough, pastry sheets, etc.) are all aspirational tasks for home-based robots. In commercial settings, robots could be used for cleaning tasks requiring handling of towels or similar implements, applying tape for packaging, etc. Looking into the future, robots could also help with autonomous protection of habitats or vehicles in remote or dangerous environments. For example, using fire blankets to extinguish a fire or applying an adhesive patch to seal leaks in undersea or extraterrestrial settings will require robots to robustly handle thin flexible materials. However, handling such materials is still challenging for modern robotic manipulators.

Several advances in robotic handling of thin, flexible objects in recent years aim to enhance the capabilities of traditional rigid robotic systems in the realms of perception, planning, and learning. For example, representing the pose of a piece of fabric or a garment presents fundamental challenges to traditional pipelines, requiring new methods for efficient geometric representations [87]. In grasp planning, some earlier work also looks specifically at grasp planning for hemmed fabrics [88], and more recently for handling folded fabrics [89]. For task planning, the task of making a bed was investigated using deep learning to determine where to grasp a bed sheet in order to unfold it to a desired state [90]. Finally,




Figure 3.1: Three-axis compliance enables soft grippers to gently interact with thin, flexible objects and the surfaces they rest on. Gentle interaction is especially important for delicate materials, like the damp paper towel shown here. During the grasp, the vertical compliance of a soft gripper allows for large positional uncertainty; during a snag (when the swatch is caught), the soft gripper’s lateral compliance increases the time over which forces are applied, and rotational compliance decreases the maximum tensile forces applied, preventing damage. By contrast, the rigid gripper tears the swatch.

attempts to transfer learned behaviors to real hardware were reasonably successful for the task of hanging a sheet of fabric on a hanger [91]. However, the main limitation in the latter study was the fact the combination of a rigid robot arm and rigid gripper could only handle very small vertical uncertainty when grasping before failure.

At a higher level, Borras et al. developed a comprehensive framework for grasping and manipulating fabric that attempts to abstract away the specific gripper morphology [92]. They present a taxonomy of grasps commonly used in fabric manipulation, and demonstrate how specialized grasps can make manipulation of fabrics easier. The grasp taxonomy includes point, line, and surface contacts, and also includes the environment (tables, surfaces) as a source of extrinsic dexterity. While this framework nicely enables high-level reasoning

about which types of grasps to use for a given task, our work is focused on understanding how to enable each of these grasps in a robust way through hardware design.

There has also been a recent push toward the low-level design of end effectors specifically for grasping fabric and other thin materials. A review of gripping devices for commercial fabric handling shows many purpose-built grippers used in manufacturing, while most grippers used in fabric handling research are simple two-fingered grippers [93]. Some grippers use alternative means of grasping that are particularly well-suited for fabrics and other thin materials, such as electrostatic attraction [94] and micro-needles [95]. Several grippers have also been designed based on a study of how humans grasp fabric, where key motions are extracted and robotic finger kinematics are optimized to replicate human motion [96, 97, 98]. Conversely, a top-down approach to gripper design has also been taken, designing morphology from a task-centric perspective, and leading to key innovations such as a large base for supporting fabric, and a variable-friction surface on the fingertip for sliding against substrates [99]. However, the precise role of compliance for interaction with thin, flexible objects has not been fully explored, and mitigating snags (where the swatch is caught and tensile forces are applied) has not been considered.

In this chapter, we discuss the role of gripper compliance in successful, safe grasping and manipulation of thin, flexible materials. We show that for a planar gripper, vertical, lateral, and rotational compliance each contribute to preventing damage to the material. We demonstrate these benefits through an empirical case study comparing a rigid gripper to a soft gripper. We evaluated the level of vertical uncertainty each gripper can handle for prehensile and non-prehensile manipulation, and the forces and displacements incurred during snags. Finally, we demonstrated the integrated utility of three-axis gripper compliance for grasping and manipulating delicate materials, as shown in Figure 3.1.

3.2 Conceptual Grasping Analysis

Successful handling of thin, flexible objects requires both grasping and manipulating swatches of material such that no damage occurs. Many tasks require grasping swatches that initially lie flat on a surface; thus grippers must be able to gently handle contact with those surfaces even in the presence of positional uncertainty. Once a swatch is grasped, the robot must also take care not to apply large tensile forces to it, even when the material unexpectedly snags. In this section, we develop a conceptual analysis of how compliance

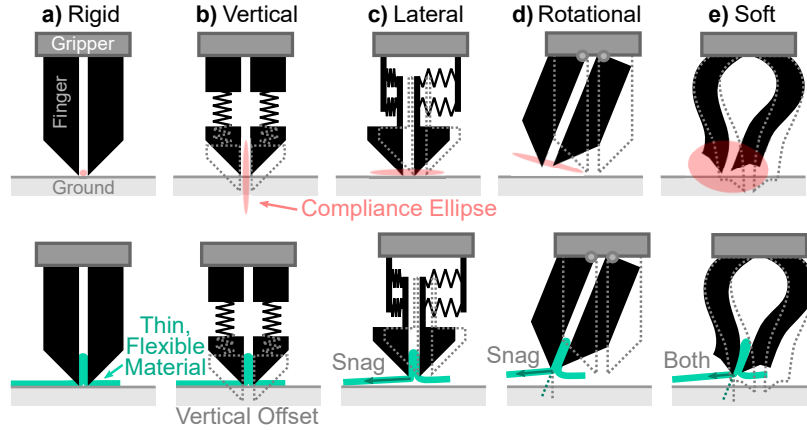


Figure 3.2: Grasp compliance in three axes is critical for grasping thin objects from a surface, while also enabling snag-resistance. Five conceptual finger designs are shown which explore each axis. a) With zero compliance, the gripper is rigid. b) With only vertical compliance, the gripper passively handles vertical uncertainty. c) With only lateral compliance, the gripper has graceful snag protection. d) With only rotational compliance, the tensile force applied during a snag is redirected away from the lateral direction. e) With compliance in three axes, (commonplace for soft robotic grippers), the robot can natively handle vertical uncertainty and snags.

affects grasping and manipulation success, and discuss key performance metrics related to the risk of damage to swatches.

3.2.1 The Role of Compliance in Grasping

Grasping thin, flexible objects from a surface on which they rest (e.g., a tabletop) usually requires a robot to interact with that surface. Since the objects are thin, achieving point-to-point or line-to-line grasps (as defined by Borras et al. [92]) requires a robot to interact with the surface through the object, or at least operate in very close proximity to the surface. Furthermore, utilizing the resting surface plane as part of an extrinsic grasp involving plane-to-point (or -line or -plane) contact explicitly relies on the robot interacting with the surface. Thus, robustness to uncertainty in the surface’s vertical position relative to the gripper is critical.

To achieve high robustness to vertical positioning error, large vertical compliance (low stiffness) in the gripper can be utilized. Such a gripper can be pressed into the table with minimal vertical force applied to the object (Figure 3.2b), while a rigid gripper (Figure 3.2a) must be placed with high precision to avoid applying large forces. Vertical compliance

achieves a similar effect as hybrid position/force control, where force control is desired in the vertical axis ([100, 101]). In addition, using the table surface as part of the gripper via extrinsic dexterity is simplified by leveraging vertical compliance with that surface [102].

For a successful grasp of a thin, flexible object, we can define several damage-related performance metrics. Normal forces must be low enough to prevent denting, creasing, or wearing through the material during the grasping process, and shear forces must be low enough to prevent tearing. To consider both of these failure modes, we can define the range of acceptable uncertainty (normal to the surface) in the hand position before forces become large enough to damage a swatch. This allows us to take into account the fact that interaction forces between the table, swatch, and gripper are related to positional error.

3.2.2 The Role of Compliance During Snags

Once a thin swatch is grasped, the robot must prevent damage to it if it gets snagged, i.e., caught on or under other surfaces such that a tensile force is applied. Detecting and recovering from snags is especially important for delicate sheets like tissue paper, garments, dough, or pastry crusts. External sensing such as vision provides little information about the stress state of thin swatches. When tensile forces are applied, elastic swatches may deform by a measurable amount, but inextensible materials will incur very little deformation even under large loads. To detect snags, a robot would then need to use force or torque sensing in the arm or in fingers to detect such forces, requiring high sensitivity and bandwidth to react to snags before catastrophic damage occurs.

To achieve high robustness to unexpected snags, gripper designs can utilize both high lateral compliance and high rotational compliance. Lateral compliance can be used to increase the time over which forces are applied (assuming constant wrist motion) by allowing the grasp to translate relative to the arm during a snag event (Figure 3.2c), giving the robot more time to detect and react to the snag. Additionally, rotational compliance enables the grasp to rotate relative to the arm during a snag event (Figure 3.2d), allowing tensile forces to be applied closer to parallel with the gripper. This change in direction causes the grasp to fail at a lower snag force than it would without rotation.

3.2.3 The Role of Multi-Dimensional Compliance

Based on the previous two conceptual analyses, a gripper with three-axis compliance should be robust to both vertical uncertainties and unexpected snags. One class of grippers with 3-axis (and often fully 6-dimensional) compliance is soft robotic grippers, where fingers are made entirely of soft materials. This soft construction allows such grippers to be extremely robust to a wide range of uncertain conditions in their environments, including the positions, sizes, and shapes of target objects, as well as those of obstacles and ground planes [4, 36, 64, 22, 21, 35, 34]. As such, soft grippers are particularly well-suited for handling thin, flexible objects in a gentle way, as indicated by Figure 3.2e.

For a given material, the compliance ellipse (planar representation of compliance) of a soft gripper should be tuned to ensure correct force thresholds are maintained for minimal damage or wear from pinching or snagging. For example, an elastic fabric can sustain large normal forces (i.e., won't dent) and large shear forces during snags, so the lower bound on compliance is small in all directions. Conversely, thin tissue paper cannot sustain large shear forces, so large vertical compliance is necessary to prevent high friction forces with the table when grasping, and high lateral and rotational compliance allows large deformation of fingers and lower snag force limits to gracefully handle snags.

3.3 Results

We demonstrate the benefits of multi-dimensional compliance for handling thin, flexible materials by performing a series of empirical investigations comparing a rigid gripper (zero compliance) to a gripper with fully soft fingers (2D compliance). The rigid gripper used was a PhantomX Parallel AX-12 (Trossen Robotics) parallel jaw gripper based on a Dynamixel AX-12 servo (ROBOTIS) with vertical and lateral grasp stiffnesses of 19800 ± 300 N/m and 5200 ± 200 N/m, respectively (mean and standard deviation of $n = 3$ trials). The soft gripper is a custom, two-fingered, pneumatic gripper from [64] designed for pinch grasping, with vertical and lateral grasp stiffnesses of 2580 ± 20 N/m and 260 ± 10 N/m, respectively. All experiments were performed by mounting each gripper to the wrist of a UR5e robot arm (Universal Robots). Robot Operating System (ROS) was used to coordinate robot and gripper motion, as well as capture video and gripper poses using AprilTags [103] where applicable.

3.3.1 Robustness to Uncertainty During Grasping

To test the effect of vertical compliance on grasp success in the presence of uncertainty, grasps on a 1 mm thick swatch of woven cotton were performed with both grippers over a range of known vertical centering offsets. With the position where the fingertips just touch the table’s surface set as the 0 mm reference point, grasps were performed for offsets ranging from 4 mm (above the table) to -40 mm (below the table), with increments of 1 mm for positive offsets and 2 mm for negative offsets. A grasp is considered successful if the arm can pick up the swatch and drag it along the table. In these experiments, the rigid gripper was operated at maximum actuation strength (with a pull-out force of 7.4 ± 0.8 N), while the soft gripper was operated at 70% of its maximum actuation strength (with a pull-out force of 1.82 ± 0.02 N) to preserve its lifespan. Pullout forces are reported using the mean and standard deviation of $n = 3$ trials.

The results of this study demonstrate that the vertical compliance of a soft gripper greatly increases the range of allowable vertical uncertainty that can be handled compared to a rigid gripper, as shown in Figure 3.5. For the rigid gripper, the region where successful grasps occur is extremely small (2-3 mm total), as shown in Figure 3.3. The rigid gripper successfully grasps the swatch at offsets between $+0$ mm and $+1$ mm, but fails to grasp at $+2$ mm offset above the table. The rigid construction of the gripper also prevents the arm from moving lower than a 0 mm offset (into the table) without applying large-enough forces to trip the arm’s emergency stop condition. Conversely, the soft gripper can handle large vertical offsets (30 mm total) with no decrease in performance, as shown in Figure 3.4. The soft gripper can successfully grasp the swatch between 0 mm and -30 mm below the initial position, since the fingers can compress and bend to adapt to the vertical position offset.

3.3.2 Robustness During Non-Prehensile Manipulation

In addition to a top-down grasp, we demonstrate that vertical compliance is essential for non-prehensile manipulation (e.g., planar sliding on the tabletop surface) of thin, flexible materials. With the same 1 mm thick cotton swatch and same range of centering offsets as in the previous experiments, the grippers were commanded to press down on the swatch, then slide it on the table to a final pose. A grasp is considered successful if the fingers maintained contact with the swatch for the entire duration of the sliding phase.




Figure 3.3: For a rigid gripper attempting to grasp a 1 mm-thick swatch of fabric from a table, the region of allowable vertical uncertainty in the gripper's position is extremely small (between 0 mm and + 2 mm). Example grasp sequences are shown with fingers a) 2 mm above the table and b) touching the table. Due to the rigid construction of the gripper, the arm cannot move lower than the position shown in b).




Figure 3.4: For a soft gripper with 2D compliance attempting to grasp a 1 mm-thick swatch of fabric from a table, the region of allowable vertical uncertainty in the gripper’s position is very large (30 mm). Example grasp sequences are shown with fingers a) 2 mm above the table, b) just pressing the table, and c) pressed 30 mm below “just-pressing”.




Figure 3.5: The range of vertical uncertainty that each gripper can withstand is shown for both the grasping and non-prehensile manipulation tasks. The soft gripper successfully manipulates the sample under large vertical displacements, while the rigid gripper is only successful for a small range of displacements. Bars represent regions of success for $n = 1$ trial, and error bars represent uncertainty due to the resolution of positions tested.

We find that the vertical compliance of a soft gripper again greatly increases the range of allowable vertical uncertainty that can be handled compared to a rigid gripper (Figs. 3.5, 3.6). For the rigid gripper, successful grasps occurred only in a small window of vertical offsets: a +2 mm vertical offset leads to failure via loss of contact with the object, and an offset of -2 mm leads to failure by pressing forces large enough to trigger the arm’s emergency stop condition. For the soft gripper, contact is maintained with the swatch over a wide range (50 mm) of vertical offsets.

3.3.3 Handling Snags

We investigate how both grippers compare when snagging on materials with different properties. Grasps were performed with a 0 mm vertical offset on three swatches, with relevant properties detailed in Table 3.1. To ensure snagging behaviour is fairly compared between the two grippers, the soft gripper was operated at 100% of its pneumatic limit of 193 kPa with a grip strength of 3.17 ± 0.01 N, and the rigid gripper’s strength was matched as closely as possible (3.6 ± 0.2 N at 15% motor torque). Due to friction limits in the rigid gripper’s mechanism, 15% torque is the lowest actuation setting that still resulted in grasping motion.

To demonstrate the effect of lateral and rotational compliance on the handling of thin, flexible materials, both grippers were tested in a simulated snag scenario, as shown in Figure 3.8. The robot was commanded to grasp one side of a swatch, then lift 5 mm and attempt to move laterally by 30 cm at a speed of 0.10 m/s. The other side of the swatch




Figure 3.6: Visual comparison of the rigid vs. soft gripper performing non-prehensile (extrinsic) manipulation of a piece of cloth by sliding it on a tabletop). a) The rigid gripper is successful at moving the swatch when perfectly positioned, but c) a small vertical offset of just 2 mm results in failure to contact the swatch. b,d) Conversely, the soft gripper is successful in manipulating the swatch with large vertical offsets.

Table 3.1: Properties of swatches used in snag experiments

Swatch Type	Thickness (mm)	Young's Modulus (MPa)
Elastic (spandex)	0.69 ± 0.02	0.21 ± 0.053
Woven (cotton)	0.41 ± 0.08	12.0 ± 0.16
Woven (cotton), Folded	3.03 ± 0.01	0.56 ± 0.046

Values are reported as the mean \pm standard deviation of $n = 4$ trials. Thickness is measured using ASTM standard D1777 [104] with 0.16 kPa preload

was manually clamped to the edge of the table, causing a snag when the robot attempted to move the swatch. The poses of the wrist and fingertips were recorded during these tests using AprilTags [103], viewed by a world-mounted webcam at a framerate of 30 Hz, and the forces applied by the gripper to the swatch were recorded by the built-in force/torque sensor in the UR5e arm at a rate of 500 Hz. Three trials were performed for each swatch type.

The results of these tests (Fig. 3.7) demonstrate that the lateral compliance of the soft gripper leads to tensile forces applied over much longer time spans than for a rigid gripper. This is due to the large lateral grasp displacement during a snag: For all swatches, the displacement of the grasp during snags was large (>30 mm) for the soft gripper, and negligible (<1 mm) for the rigid gripper. In addition, forces were applied to inextensible (woven) swatches over approximately twice as long a period with the soft gripper as with the rigid gripper in both cases. The elastic swatch, however, saw forces applied over the

same amount of time for both grippers ($p > 0.5$), due to that swatch’s high compliance (Table 3.1).

For context, if the arm were to manipulate the woven swatch at more realistic speeds, such as 2 m/s ($20\times$ our testing speed), the rigid gripper would take 12 ms to reach the maximal snag force, which is only enough time for six force measurements to be made with our robot’s force/torque wrist sensor. The soft gripper would take 25 ms, allowing for 12 measurements to be made. With a more compliant soft gripper, this time could be lengthened as needed to enable the robot to detect and react in time to prevent damage.

Our results also show that rotational compliance results in lower force thresholds before grasp failure, even when grip strength is held constant. As shown in Fig. 3.7, the maximal snag force before grasp failure was consistently lower for the soft gripper for all swatches tested. The largest difference appeared for the folded swatch, where the rigid gripper applied more than twice the lateral force of the soft gripper before failure, despite having the same nominal grip strength.

3.3.4 Task-Relevant Demonstration

To demonstrate the benefits of the planar compliance of soft grippers in a real-world task, we consider an example of table-cleaning in the restaurant industry. Robot cleaners are likely to interact with delicate objects such as wet napkins or paper towels, which could be caught underneath cutlery or crockery. In our demonstration, a damp paper towel is caught under a heavy bowl (Fig. 3.1). The rigid gripper tears the towel slightly during grasping due to high normal forces, then tears it badly when pulling laterally due to the large tensile force induced in the towel. By contrast, the soft gripper successfully grasps the towel even under a small vertical offset without causing damage, then passively rotates when the towel snags, releasing the grasp before any permanent damage occurs.

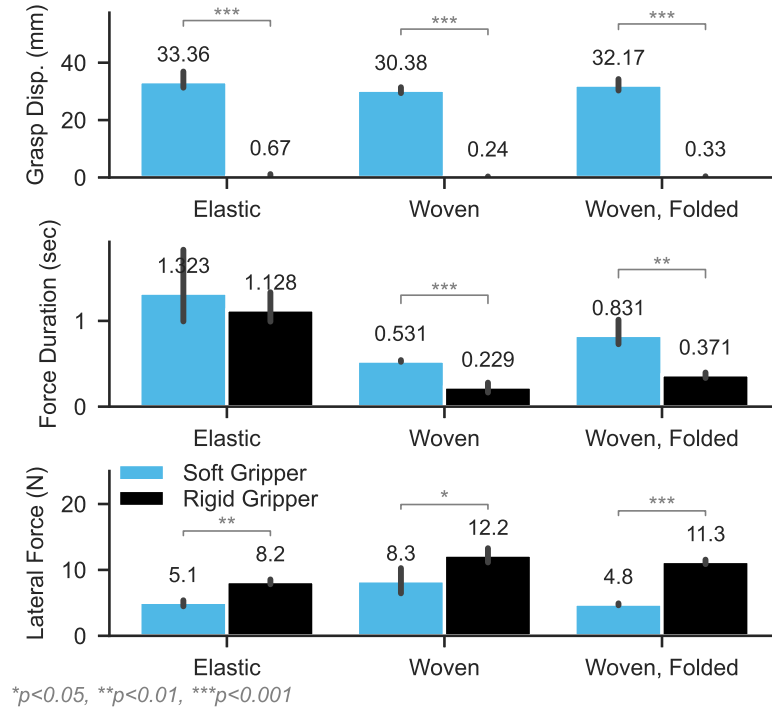


Figure 3.7: The lateral and rotational compliance of the soft gripper enable it to gracefully handle snags. During a snag with both elastic and woven materials, lateral grasp displacement, duration of applied lateral force, and the lateral force on the swatch are shown. Bars and errorbars represent the mean and standard deviation for $n = 3$ trials, and statistical significance is calculated using t-tests.




Figure 3.8: Visual comparison of rigid and soft grippers in the presence of a lateral snag (one side of the fabric swatch is fixed to the table). The lateral grasp displacement as well the grasping angle during the snag are negligible for the rigid gripper, but large for the soft gripper, enabling graceful grasp failure during snags without damaging the swatch.

3.4 Discussion

Our results demonstrate that the same feature (multi-dimensional compliance) that enables soft robots to gently interact with commonly-studied objects also has important benefits for handling thin, flexible objects. For both prehensile grasps and non-prehensile manipulation, the gripper inherently operates on or near tabletop surfaces. Vertical compliance in a gripper enables gentle, force-limited interactions with the table, even with large vertical position error. This can be particularly useful in situations where visual perception is difficult or unreliable. In addition, once a swatch is grasped, an unexpected snag could occur, which causes tensile forces to be applied to the object by the robot. We showed that rotational compliance in the gripper decreases the maximal tensile force applied, leading to passive force-limited grasps. Additionally, lateral compliance increases the time over which these forces are applied, which directly reduces the sensing bandwidth required to successfully detect snags before damage occurs in the swatch.

While only two values of compliance were tested in this study, we expect the results to apply to gripper designs with intermediate compliance as well. If the three directions of compliance could be fully decoupled and tuned independently per an application’s specifications, finer-grained control over the exact forces applied to a swatch could be achieved. Through characterization of the maximal expected error in perception for a given robotic system, the vertical compliance of the gripper can be tuned to provide safe interaction with tabletop surfaces without sacrificing precision. With proper characterization of the maximum allowable tensile forces before damage occurs to a given material, the rotational compliance of the gripper could be tuned to set an upper limit on snag forces, and the lateral compliance could be tuned such that snags can be detected with onboard sensors quickly enough for the robot to react.

Finally, we note that compliance is not necessarily required to be implemented at the finger level. While this is the most common implementation in the soft robotics space, a variety of alternative approaches exist that could enable grasp compliance without finger compliance. For example, a gripper with rigid fingers connected to a highly compliant wrist would obtain the same performance benefits of the soft gripper in this chapter with respect to grasp robustness and snag resistance. However, the rigid components could still cause damage to materials during grasping or rapid motion, and the robot would also lose the benefits of finger softness when manipulating other delicate objects.

3.5 Conclusions

Here we have demonstrated that a soft robotic approach to gripper design increases a robot's ability to safely grasp and manipulate thin, flexible objects, as well as facilitating detection and recovery from unexpected snags. We discussed the critical roles that planar compliance plays in successful handling of thin materials, and demonstrated these with hardware experiments. We showed that vertical compliance enables graceful handling of vertical uncertainty as well as limiting potentially damaging forces applied to objects while pressing against surfaces. We also showed that lateral and rotational compliance can prevent damage to objects during unexpected snags by decreasing the maximal tensile forces applied to the object and increasing the time window over which forces are applied. Overall, our results demonstrate that a soft robotic gripper with planar compliance can achieve all of these benefits through passive means.

These results give rise to a number of future directions in studying robotic handling of thin, flexible materials. One promising area involves using variable-stiffness actuators, where dynamic control of both vertical and lateral stiffness could enable a gripper to adapt for materials of different fragility. Another possible area of interest is the use of onboard sensors in the fingers to directly detect snags early via finger deformation. Finally, compliant grippers could enable higher success in bi-manual manipulation tasks, mitigating any potential snag forces between hands.

4

Designing dexterous soft fingers for in-hand manipulation

4.1 Introduction

In recent years, there has been a sizable push toward utilizing compliant robotic hands that passively adapt to uncertainty in the environment [4]. Soft-bodied hands enable robots to grasp objects of varying shape, size, and pose without explicit knowledge of those properties [22, 36, 72, 60]. Furthermore, passive compliance enables robots to safely interact with delicate target objects or other fragile features in the environment [35, 105, 106].

Dexterous in-hand manipulation usually requires precise planning and control of finger motion based on models of the object and fingers when performed by rigid hands [48, 49, 50]. This is due to complex contact interactions between the fingers and object, as well as minimal passive adaptation to object variation. Some attempts to mitigate the complexity of these interactions using machine learning show incredible promise, but require extensive training on high performance computing systems [57]. However, we can mitigate the need for planning and complex control for some in-hand manipulation tasks through targeted design of a soft robotic hand with dexterous fingers.

In this chapter, I show that soft robotic hands can robustly perform in-hand manipulation in the presence of uncertainty. I first design a soft hand prototype with dexterous fingers capable of fingertip motion conducive to several basic motion primitives. Through empirical validation, I show that these dexterous fingers enable the desired object motion primitives within the hand while still maintaining strong grasping capabilities. I then examine a simple, heuristic finger gait which enables continuous object rotation for a wide variety




Figure 4.1: Our soft, dexterous hand prototype is capable of performing real-world manipulation tasks within the hand. a)-c) The hand is shown unscrewing the cap of an empty plastic jar using a heuristic finger gait for rotation. d)-f) The hand can also impart planar translations to objects. The empty jar was lightly taped to the ground to enable un-capping.

of object shapes and sizes. Finally, I demonstrate the utility of in-hand manipulation using a dexterous soft hand in three real-world cases: unscrewing the cap of a jar, orienting food items for packaging, and gravity compensation during grasping.

4.2 Design of a Dexterous Soft Hand

4.2.1 Task-Centric Performance Goals

In this work, we focus on three motion primitives in which the object moves while maintaining a grasp: rotation about one axis, and translation in two axes of a plane, as displayed in Figure 4.2. These primitives were chosen based on their utility in a set of target tasks drawn from a variety of application areas. These areas include performing ADL in the home, assembling and packaging delicate items such as pastries [1], and picking/handling produce [25].

Our first desired motion primitive is rotation about an object’s central axis. Rotation is useful for a variety of activities of daily life including unscrewing bottle caps, turning dials and knobs, and reorienting non-axisymmetric objects. Rotations can also enable tool use during assembly, adjustment of items during packing tasks, and twisting fruits and vegetables to pick them.

The second and third target motion primitives are translations in the plane perpendicular to the object’s central axis. Planar translations are useful for fine, local adjustments during packing tasks, as well as picking produce by shifting side-to-side. Furthermore, translations can be used to compensate for finger deflection caused by gravitational forces on the object.

4.2.2 High-Level Hand Design

To build a soft hand capable of achieving the three motion primitives of interest, several design decisions were made based on previous successful hand designs with compliant fingers. The high-level structure of the hand was designed with several compliant digits placed around a flat palm. The axis normal to the palm was set to be parallel to the object’s central axis so the rotations and translations occur in the reference frame attached to the palm. To simplify object motions, the fingers and palm were designed such that objects only contact fingers at the fingertips. We also assume fingertips have rolling contact




Figure 4.2: The necessary fingertip motions for our soft hand are conceptually derived from desired object motion primitives. Based on the motion of typical soft bending actuators, the desired fingertip motions can be achieved using two parallel bending actuators




Figure 4.3: The two high-level hand designs considered in this work are the a) four-finger cross design, and b) more-traditional anthropomorphic design.

with objects without slipping.

Additionally, two high-level configurations were considered for finger arrangement: four-finger cross, and anthropomorphic, as displayed in Figure 4.3. The four-finger cross configuration has two sets of antipodal finger pairs, similar to several existing hands [53, 107, 49]. Conversely, the anthropomorphic configuration is based on the designs of several soft hands [36, 72, 60].

4.2.2.1 Four-Finger Cross Finger Arrangement

The four-finger cross design features two pairs of opposing fingers located perpendicular to each other, as shown in Figure 4.3a. With this design, the hand can always maintain static equilibrium with objects using one or both of the antipodal finger pairs [53, 107]. In addition, we can sequence the two sets of antipodal pairs to achieve our three target motion primitives in a straightforward way. Objects can be rotated about the axis perpendicular to the palm’s surface using one set of antipodal pairs at a time. Objects can also be translated in the plane perpendicular to the palm’s surface using simple combinations of translations in orthogonal axes (one axis from each antipodal pair).

While the four-finger cross design would likely work well for objects with aspect ratios of $\sim 1:1:1$, manipulating objects with more-extreme ratios (such as a pencil) may be challenging. This design has a roughly cylindrical space between fingers, so objects with extreme aspect ratios like pencils ($\sim 1:1:27$) may not fit into this volume at certain orientations. However, these challenges can be partially mitigated if unused fingers are able to retract out of the way during manipulation.

4.2.2.2 Anthropomorphic Finger Arrangement

The more-traditional anthropomorphic hand design uses three aligned fingers and an opposing “thumb”, as shown in Figure 4.3b. This configuration is similar to other soft hands capable of gentle grasping and manipulation [36, 72, 60]. With this configuration, objects can be rotated about the axis perpendicular to the palm’s surface in two ways: rolling contact with the thumb and middle finger, or antagonistically by pressing with one of the outside fingers while grasping the object with the thumb and the middle finger [72]. Translations in a plane (perpendicular to the palm’s surface) can be achieved with the thumb and the middle finger for one direction, and side-to-side motions of the fingers for other directions.

While the anthropomorphic design would likely work well for objects with more-extreme aspect ratios, manipulating objects with closer to equal aspect ratios may be challenging. Utilizing all four fingers during a grasp would be challenging since applying forces with the outer two fingers along their main axes could cause force imbalances with the thumb. In addition, the control strategy for rotation may change drastically depending on object shape.

4.2.2.3 Choosing one design

Considering these two hand configurations, we ultimately chose the 4-finger cross design, as it appears to simplify the control of our three target motion primitives. Under open-loop pressure control, this hand design makes it easier to control translation compared to the anthropomorphic hand, where more than two fingers are implicated for the same translation. Finally, many common objects in the home have aspect ratios of $\sim 1:1:1$ [83], making the four-finger cross the best choice due to symmetry between pairs of fingers.

4.2.3 Dexterous Finger Design

In order to design dexterous soft-bodied fingers, we used an object-centric approach to extract how the fingertips should move based on our desired object motion primitives, as shown in Figure 4.2. To produce object translations, fingertips need to translate in the primary grasping axis (in/out toward the center of the palm) and side-to-side (perpendicular to the grasping axis). Side-to-side motion is also necessary to perform rotatory motion of the grasped object. In addition to fingertip motion, the fingers must be strong enough to impart that motion onto real objects as well as sustain strong grasps along the grasping axis.

To achieve compliant, strong, and lightweight fingers, we created pneumatic bending actuators with two side-by-side air chambers separated with a central wall. Each chamber is based on a typical bellows actuator design, similar to those found in [35]. This configuration of actuators enables strong grasping when both chambers are actuated with equal pressure, while also offering ample side-to-side motion when actuated differentially. In addition, the side-to-side motion can be amplified by using vacuum to increase the pressure differential between chambers. A more-detailed explanation of how actuation pressures affect fingertip motion can be found in [65].

We carried out several design iterations in an attempt to maximize the strength, range of motion, and burst pressure of these fingers to increase the reliability of this system. Several design parameters were modified including the number and size of bellows, wall thickness, and fingertip shape. The most recent version presented in Figure 4.4a has 14 short bellows with a wall thickness of 2mm. We chose a bending segment of 100mm and a width of 22.5mm to maintain consistency with similar designs such as [64, 35]. Soft finger pads were added to each fingertip to increase compliance with grasped objects [27, 28].




Figure 4.4: The design and fabrication of our two-chamber dexterous fingers. a) Our finger design utilizes two side-by-side bellows actuators. b)-d) The fabrication process for these fingers involves injecting silicone elastomer into 3D-printed molds with soft cores.

4.2.4 Hardware Fabrication

The fabrication process for our soft fingers is based on the process found in [35]: a rigid mold creates the outer geometry of the finger, and a soft silicone core creates the internal geometry. Before building a finger, two soft cores are made by injecting True Skin[®]10 silicone into a 3D printed mold (VeroClear, Stratasys), then curing in an oven for one hour. To build a finger, the two soft cores are de-molded and secured in the main mold using two square rods, as shown in Figure 4.4b. Next, all parts of the mold are clamped together with screws, and Smooth-Sil[™]945 silicone (45 Shore-A hardness) is injected using a custom injection system, as shown in 4.4c. When the material is cured, the two soft cores are removed, resulting in the two side-by-side air chambers. The two holes created by the rod are sealed with a small amount of silicone to prevent air leakage. The soft finger pads are created using the same injection process but with a softer silicone (Ecoflex 00-30[™]) with a 30 Shore-00 hardness. The injection process ensures that the fingers are made repeatably, with minimal variation between each batch.

To complete a newly molded finger, first a 3D printed hub with ribs is fitted with two pneumatic fittings enabling easy swapping of pneumatic control lines, and two screw posts for mounting on the scaffold. This hub is then glued into the main body of the finger

with Sil-PoxyTM. A Kevlar thread binding and heat shrinking tubing is fixed around the base of the finger to strengthen the pneumatic connection between the rigid hub and the finger body. The completed finger can be seen in Figure 4.4d. All of the molds and the rigid hubs were 3D printed on an Object Connex 500 printer (VeroClear, VeroWhite and VeroBlue, Stratasys).

Once four fingers are completed, they are attached to a rigid scaffold to arrange them into two perpendicular antipodal finger pairs. The finger mounting angle was chosen based on the fingers' range of motion. An angle of 40° between the two fingers in each pair permits grasping a wide variety of objects without losing force before making contact with the object. Since the scaffold has the same mechanical interface as the fingers and palm, it is easy to swap new fingers and palms as needed. The scaffold is also designed to be easily attached to a robotic arm (UR5e) with a standard mounting flange.

Finally, a palm is placed in the middle of the scaffold, with a flat surface of 70 mm diameter located at a distance of 40 mm from the contact point of the fingertip. This palm height ensures that objects are close enough to the fingertips when resting on the palm. Both the scaffold and rigid palm were 3D printed on a Markforged Onyx One printer (Onyx Material, Markforged).

4.3 Characterization of the Hand Prototype

4.3.1 Actuation and Control

To control pressures independently on both sides of each finger, we used a custom pneumatic pressure control system with eight independent channels [108]. Based on the system used in [64], the controller enables execution of arbitrary pressure trajectories in real time with an accuracy of 1.4 kPa. Detailed performance specifications can be found in Appendix B

4.3.2 Finger Performance

Before characterizing the hand, we first characterized fingers individually to understand the range of motion, strength, and operating pressures. These metrics subsequently inform how fingers can be used for grasping and manipulation within the hand.

To evaluate the workspace of our soft, 2DOF fingers, we applied sweeps of actuation

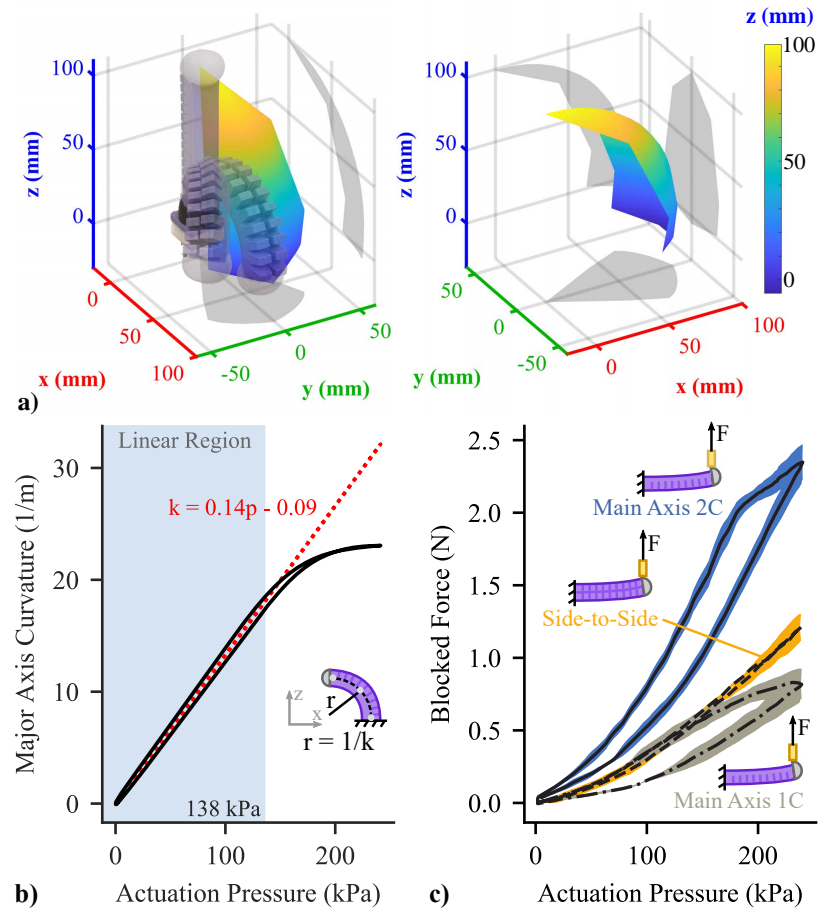


Figure 4.5: Validation of finger performance. a) The workspace of our fingers enables ample side-to-side motion during a grasp. Two views of the fingertip workspace surface are shown. b) The curvature in the grasping axis vs. actuation pressure is roughly linear in the 0 to 138 kPa range. c) The maximum blocked force of each finger is 2.3 N in the grasping axis, and 1.2 N side-to-side.

Table 4.1: Unactuated Stiffness of individual fingers in each axis.

	Unactuated Stiffness (N/m)	
	Grasping Axis	Side-to-Side
Finger 1	6.09 ± 0.03	28.23 ± 0.20
Finger 2	6.23 ± 0.07	28.90 ± 0.11
Finger 3	6.04 ± 0.02	30.02 ± 0.16
Total	6.12 ± 0.09	29.05 ± 0.8

* The mean and standard deviation are reported for $n = 3$ trials for each sample.

pressure/vacuum to both chambers of the finger (-34 kPa to 240 kPa), and measured the free deflection of the fingertip. Using a Vicon motion capture arena, a finger’s workspace was traced out in 3D, and the resulting surface is shown in Figure 4.5a. The fingertips have a range of approximately 80 mm in the major grasping axis, and ± 30 mm of lateral range with minimal change in grasping displacement.

In addition to the 3D workspace, the grasping axis curvature in response to actuation pressure can help us understand how to control these fingers. To measure the curvature of fingers along the main grasping axis, we applied controlled pressures equally to both actuator chambers, and measured the resulting displacement of three points along the length of a finger. Using MATLAB to synchronize the motion tracking data to actuation pressure, we pressurized the finger from 0 to 250 kPa. Then, the curvature was computed from the marker positions using a least squares fit for a circle. The results are shown in Figure 4.5b. The curvature linearly increases with pressure from 0 to 138 kPa, after which minimal change in curvature occurs. This suggests an operating point of 138 kPa during grasping is a reasonable starting point.

To evaluate the strength of our soft fingers, we began by measuring the blocked force during actuation. Individual fingers were placed in a custom fixture beneath the crosshead of an Instron universal testing machine with a 10 N load cell. The fingers were actuated from 0 to 250 kPa, and the resulting vertical blocked force was measured, as shown in Figure 4.5c. Each test was repeated three times for three different fingers. The maximum strength in the grasping axis is 2.3 N when both air chambers are pressurized. However, the strength when only one air chamber is pressurized is 0.8 N in the grasping axis and 1.2 N side-to-side, indicating a tradeoff in grasping forces vs. side-to-side motion.

Another useful performance metric is the unactuated stiffness, which can help us understand finger behavior under load. To measure the stiffness, we used a custom fixture

to mount fingers as cantilevers underneath the crosshead of the Instron. We then used the Instron to deflect the finger by 10 mm along the x and y axes of the finger while measuring the resulting force. We then found the slope of the load-displacement curve for each sample, which results in the average stiffness shown in Table 4.1. It is notable that our finger design is approximately 4.7 times stiffer in the side-to-side axis than the grasping axis, which can potentially help maintain stable grasps when fingertip forces are not perpendicular to the grasping axis.

Finally, we obtained a burst pressure of 310 kPa for both air chambers. Actuation pressure was slowly increased at approximately 10 kPa per second until the fingers failed by rupturing. Combined with the results of our curvature evaluation, we chose 240 kPa as the maximum operating pressure and 138 kPa as the nominal grasping pressure.

4.3.3 Grasping Performance

With a better understanding of the performance of individual fingers, we now focus on the grasping performance of the whole four-fingered hand. Through an analysis of the grip strength and grasp stiffness, we determined the range of object masses that can be sustained in a grasp, as well as how hand orientation affects this.

To measure the grip strength of our soft hand as well as the effect of soft finger pads on grasping performance, we performed a series of grasps on a cylinder of diameter 76.2 mm, similar to [35]. The cylinder was attached to the crosshead of our Instron machine, and the hand was fixed on a structure below the cylinder. The cylinder was then pulled vertically and the maximum vertical force from the fingers was measured. This process was repeated three times for each hand configuration. With this setup, we evaluated the effect of the actuation pressure, soft finger pads, and the use of two vs. four fingers on grasp strength, as displayed in Figure 4.6a.

Based on our experiments, the hand has the highest grip strength when using four fingers with soft finger pads actuated at the highest pressure. With a pressure of 173 kPa, the hand achieves a grip strength of 4.7N, demonstrating that the hand is able to grasp objects up to 0.47kg when four fingers are actuated with soft finger pads. Conversely, the best possible performance without finger pads was only 0.21N (grasping a 0.21kg object), showing that adding soft pads significantly increases grip strength. Finally, our experiments show that grip strength is roughly two times higher when grasping with four fingers compared to

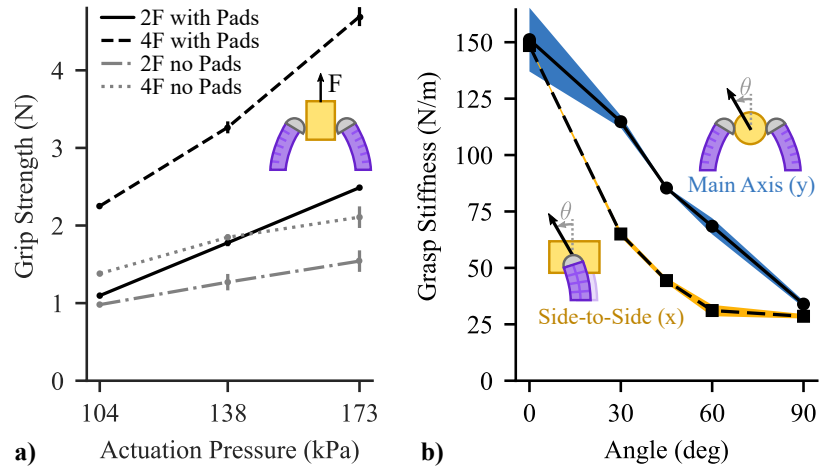


Figure 4.6: Characterization of grasping performance. a) The grip strength increases by a factor of two when compliant finger pads are used compared to harder fingertips. b) The grasp stiffness decreases roughly linearly with the angle along the main axis, but much faster side-to-side. In both graphs, the mean and standard deviation are reported for $n = 3$ trials for each sample.

two fingers.

In addition to grasping strength, the hand must also be able to resist external forces applied off-axis. To evaluate the hand’s ability to resist off-axis loading, we measured the stiffness of two-finger grasps as a function of the angle. For these measurements, a cylinder of 25 mm diameter was fixed to the crosshead of the Instron, and the hand was mounted at precise angles to the cylinder using a custom fixture. To measure stiffness, a two-finger antipodal grasp was performed on the cylinder (with an actuation pressure of 138 kPa), then the cylinder was deflected by 10 mm using the Instron and the resulting vertical force was measured. This process was repeated three times for angles ranging from 0 to 90 degrees about the x and y axes.

The resulting grasp stiffness for angles about both axes is presented in Figure 4.6b. For angles about the major axis of the cylinder (x -axis), the grasp stiffness decreases roughly linearly with angle, similar to the fingers in [64]. In addition, the stiffness decreases much quicker for angles about the y -axis (side-to-side with respect to the antipodal grasp). This discrepancy is likely due to the added stiffness of each finger in the grasping direction when actuated with air pressure. Finally, the worst-case loading condition involves forces perpendicular to the palm (90°) where the stiffness is lowest (30–35 N/m). In practical use,

this suggests a 50 g object held perpendicular to gravity would deflect by 17 mm.

4.4 In-Hand Manipulation Primitives

To evaluate the in-hand manipulation performance of our prototype hand, we empirically tested all three motion primitives using a small standard set of objects, as shown in Figure 4.7a. To keep consistent with standard object sets, two rigid cylinders were taken from the YCB Object set (47 mm and 88 mm diameters), and a rigid box of dimensions 60 mm \times 60 mm \times 60 mm was included from [64]. In addition to these rigid objects, a compliant box made of memory foam (50 mm \times 50 mm \times 70 mm) was used. Finally, three fragile real-world objects were tested: two muffins (70 mm and 40 mm diameters), and a broccoli crown (approximately 110 mm diameter). The approximate dimensions, masses, and materials are shown in Table 4.2.

Table 4.2: The set of objects used to evaluate manipulation

Object	Dimensions (mm)	Mass (g)	Material
Cylinder (Small)*	$\varnothing 47 \times 60$	35	Rigid Plastic
Cylinder (Large)*	$\varnothing 88 \times 78$	46	Rigid Plastic
Cardboard Box †	60 \times 60 \times 60	13	Cardboard
Foam Box	50 \times 50 \times 70	13	Memory Foam
Muffin (Small)	$\varnothing 40 \times 40$	30	Muffin
Muffin (Large)	$\varnothing 70 \times 60$	147	Muffin
Broccoli Crown	$\varnothing 110 \times 90$	230	Broccoli

* Items are part of the YCB Object Set [83]. † Items are included in [64].

To evaluate the spatial range of object motions our hand is capable of achieving, we performed sweeps over all three motion primitives. Using heuristically-designed pressure trajectories (sequences of hand-tuned pressure waypoints), finger motion was commanded slow enough to assume quasi-static conditions. The motion primitives were carried out with objects starting at rest on a transparent acrylic ground plane, and object motion was captured using a camera mounted below. Using Tracker Video Analysis software [86], we measured the positions of two known markers on each object, enabling the object’s planar pose to be calculated using Python.

The planar “workspace” for each object is shown in Figure 4.7b and 4.7d. The region of reachable positions in the x - y plane is similar for all objects, with a roughly circular shape of 50 mm diameter. This region is likely limited by the size of each finger’s workspace,

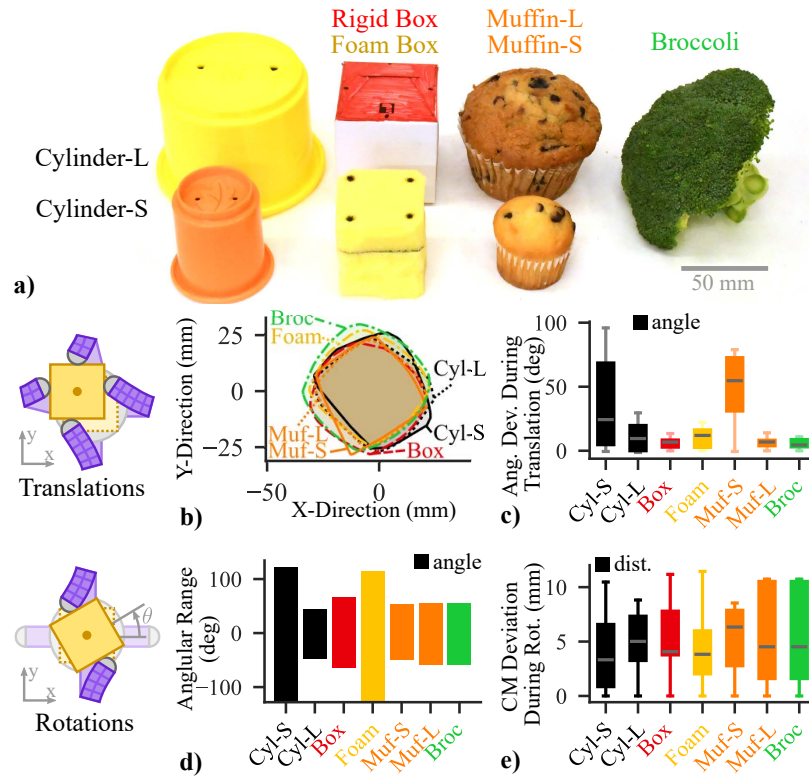


Figure 4.7: a) Our soft hand can perform the three object motion primitives for which it was designed on a variety of objects. b) The space of reachable object translations is similar for all objects tested. c) The undesired angular motion of the object during one translation cycle is somewhat large. d) The range of reachable object rotations is highly dependent on object size. e) The extra translational motion of the object during one rotation cycle is within 12 mm for all objects.

with ± 30 mm of side-to-side deflection. Conversely, the range of reachable object rotations is inversely dependent on object size, with smaller objects undergoing larger rotations. This also meets expectations given the range of side-to-side deflection each finger can achieve. Finally, while knowledge of object size could be used to calculate the necessary finger motion to achieve a given rotation angle, this relationship is highly nonlinear.

When commanding pure motion primitives on the object, we also measured the object's undesired off-axis motion, as shown in Figure 4.7c and 4.7e. During the sweep of reachable translations, smaller objects tended to incur larger undesired angular deflection. This is likely due to rotational instabilities that occur when grasping objects of smaller diameter (higher curvature) [28]. On the flip side, the uncontrolled lateral translation during attempts to achieve pure object rotation were small compared to object dimensions (within 12 mm for all objects).

The effect of irregularities in object shape can be derived from the results of individual objects. Pure translation appears to be relatively invariant to object size, so objects with protrusions or extreme aspect ratios could still be manipulated predictably. Conversely, pure rotation is very dependent on the distance between fingertips, so we expect that objects with lobes or extreme aspect ratios would experience rotation similar to simpler objects of the same diameter for any given set of contact points.

4.5 A Simple Finger Gait for Continuous Rotation

While simple combinations of our desired motion primitives enable considerable object motion, these primitives have limited spatial range if the finger maintains contact with the object. However, if contact with the object can be broken and reformed once the fingers reach the edge of their workspace, a finger gait can be developed to move objects an arbitrary amount. Our hand design with two pairs of antipodal grasping fingers can perform such finger gaits stably within the hand, and without placing the object onto an external surface during the reset period.

Focusing on continuous object rotation, we can utilize the compliance of our soft fingers to design one simple finger gait that enables continuous rotation of a wide variety of objects. The gait involves a two-part cycle with two parameters based on an extension of our simple rotation primitive, as shown in Figure 4.8. First, the object is grasped with one pair of antipodal fingers and rotated to the edge of the fingertip's workspace while

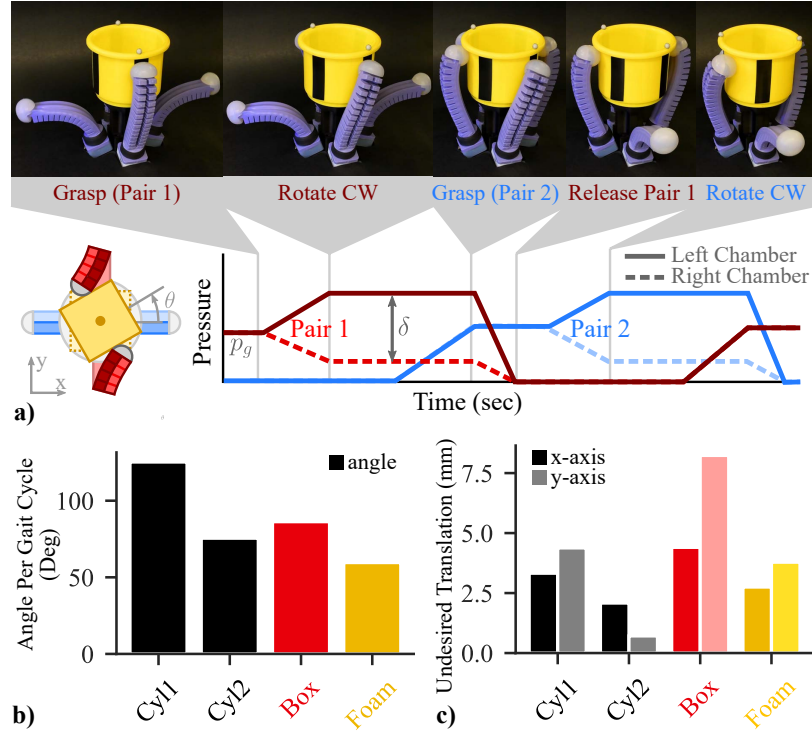


Figure 4.8: Using a simple heuristic finger gait, our soft hand is capable of continuously rotating a variety of object shapes and sizes. a) The rotation gait is defined via actuation pressures, and photographs of key events in the gait are displayed. Unless noted, $p_g = 138$ kPa and $\delta = 34$ kPa b) The angular displacement per gait cycle is reported for $n = 20$ cycles per object (mean and standard deviation). c) The undesired lateral object motion is similar in magnitude to that seen during the rotation primitive. † The palm was raised by 10 mm and $\delta = 28$ kPa. †† The palm was raised by 10 mm and $p_g = 172$ kPa.

maintaining contact. Next, the other two fingers form an antipodal grasp on the object while the first grasping pair holds position. From here, the first grasping pair releases while attempting to maintain the fingers' lateral positions, then the pair is reset to its resting position. Finally, the second part of the cycle is identical to the first part except the roles of the grasping pairs are flipped. The gait can be defined entirely in actuation pressure space, with a nominal grasping pressure, p_g , and a deviation from that pressure, δ , which differentially actuates the two fingers to produce rotation.

We evaluated the utility of this finger gait by testing its performance on our previously defined set of objects. The results of these tests are shown in Figure 4.8. For each object, we performed the finger gait for 20 gait cycles while capturing the resulting object motion using a Vicon motion capture system. The average number of cycles to rotate each object are shown in Figure 4.8, along with the average extraneous lateral motion.

Not only is this class of finger gaits simple, the gait is also robust to variation in object shape and size. In fact, for this set of experiments, the exact same actuation pressure trajectories were used for all objects, which still resulted in successful, continuous object rotation for all 20 gait cycles. We suspect the mechanism that this simple finger gait exploits is the fact that controlling actuation pressure indirectly controls contact forces on the object. Using pressure control on a soft finger results in controlling finger motion until the finger is blocked, then acts as contact force control on the object [109]. This means that we can achieve robust continuous object rotations without fingertip sensors, and without re-planning fingertip trajectories in response to object size and shape differences. Moreover, this continuous object rotation is useful when a robotic arm is constrained by the environment or reaches joint limits. This continuous motion can also be advantageous in arms that do not have rotational wrists.

4.6 Demonstration of Real-World Manipulation

To showcase the utility of our soft dexterous hand, three manipulation tasks were performed using a UR5e 6-DOF robot arm. While visual perception is essential to any full manipulation system [5], no visual perception was used as it is outside the scope of this study. In these demonstrations, objects are manually placed, and the arm moves our dexterous hand to predefined poses without explicit knowledge of the object's pose, size, or mass properties (unless noted). Full motion sequences for all of these demonstrations can




Figure 4.9: Using the heuristic finger gait developed earlier, our soft hand can gently manipulate delicate objects. a)-b) The hand rotates a cupcake to place a design in a desired orientation (forward). c) The hand then grasps the cupcake. d)-f) Finally, cupcakes are packed in a carton with the help of a robot arm. The full sequence is shown in the Supplemental Video [65].

be found in the Supplementary Video of [65].

In the first task, our soft hand is used to un-screw the cap of a plastic jar (as shown in Figure 4.1), a common activity of daily life. Using our simple finger gait for continuous rotation, the cap is successfully removed and picked up by the arm. Then the x - y translation primitives are combined to move the cap in mid-air without the use of an external surface. The cap is then replaced, and the whole jar is grasped and translated in mid-air.

In addition to rigid objects, soft hands can gently manipulate delicate objects such as fruits, vegetables, or pastries [1]. For example, the task of packing cupcakes into a container is easily performed using our dexterous soft hand, as demonstrated in Figure 4.9. Furthermore, we can utilize our simple finger gait to rotate the cupcakes prior to packing, placing a design forward in the container.

Finally, the dexterity of the two-chamber fingers in our hand can be used to com-

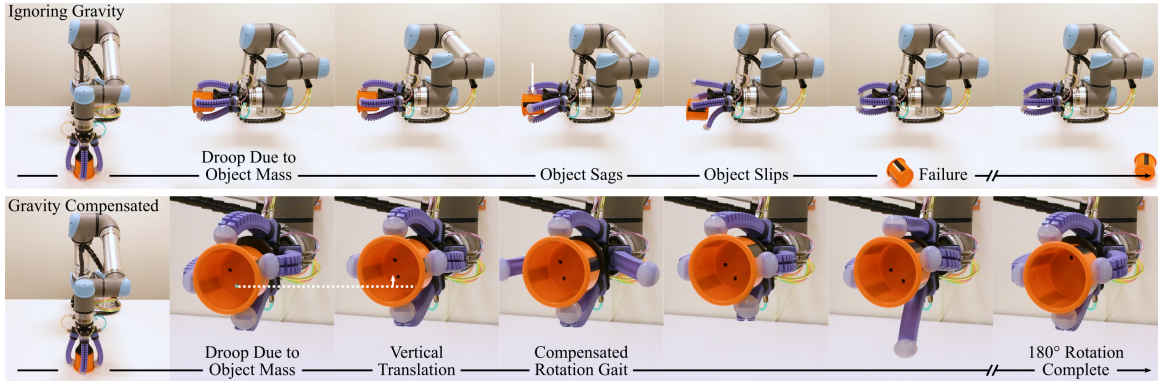


Figure 4.10: By combining motion primitives (translation + rotation), the hand can adjust the object’s position to compensate for finger droop caused by gravity. The full, annotated motion sequence can be found in the Supplementary Video of [65]

compensate for finger droop caused by the weight of grasped objects. Finger deflections during a grasp are usually large for soft robotic hands, including ours, relative to the mass of typical objects due to high finger compliance. Given knowledge of the hand’s orientation and approximate object mass, a translation primitive (simple pressure offsets) can be applied to each section of the rotation gait to shift the fingertip position vertically and recover from finger droop, as shown in Figure 4.10.

Compensating for finger droop is often crucial to the success of the rotation task for our soft hand. As demonstrated in Figure 4.10, when the fingertips are not shifted to account for droop, the object is dropped during the rotation gait attempt. However, when the gait is gravity compensated, the object’s position is closer to the center of the palm during the rotation, and stable rotation is achieved. Thus, designing fingers and a hand with the capability to translate objects enables soft hands to manipulate objects in a much wider variety of orientations.

4.7 Conclusions and Future Work

In summary, we demonstrated that a soft robotic hand is capable of robust in-hand manipulation of delicate objects without knowledge of the precise position, shape, or size of those objects. Through a conceptual analysis of desired object motion, we designed a soft hand with four soft, dexterous fingers capable of moving objects in a plane. Using simple control, our soft hand can achieve three object motion primitives (translation and

rotation in a plane), which can be combined in a straightforward way. In addition, using a simple heuristic finger gait, the hand can achieve continuous rotation of objects. Finally, we demonstrated three real-world tasks where our dexterous soft hand utilizes in-hand manipulation to maneuver objects in a gentle way.

In future work, there is endless potential to dive deeper into soft manipulation. We show in Chapter 6 that careful consideration of palm design can enhance the in-hand manipulation performance. Manipulation of grasped objects could be further improved using mathematical models of finger motion, and on-board sensing of finger shape or contact forces. With these tools, closed-loop control of contact forces could provide even more-robust manipulation in cases where passive compliance fails. Finally, dexterous soft robotic hands could be applied to a variety of exciting applications, from in-home assistive robots to bimanual manipulation, since soft robotic hands can interact safely with themselves and the surrounding environment.

5

Designing Digit Arrangement

5.1 Introduction

Common robotic hand designs fall into two main categories: anthropomorphic and “task-driven”. Human hands incorporate a combination of passive compliance and proprioception to simultaneously achieve high degrees of dexterity and gentleness, making them an excellent general-purpose design [63]. In robotics, anthropomorphic hands attempt to mimic the form of human hands, usually involving several digits arranged in a line with a single opposing ‘thumb’ [11, 10, 20, 12, 110]. Alternatively, “task-driven” hand designs typically descend from simpler grippers, with the goal of performing well for a specific set of tasks [17, 53, 22, 39, 25, 65]. By virtue of this “task-driven” design process, these hands have ideal dexterity for their applications, but tend to show limited performance outside of those domains.

In addition to dexterity, many of our target applications require a level of gentle interaction and safety which can most-easily be achieved using soft robots. Robots made with soft materials and structures have limited force output by design, which makes them inherently safe even in the event of a power outage or errors in control systems [4]. Over the past two decades, a large variety of soft grippers have emerged, with finger designs and digit arrangements that produce high-quality grasps in delicate situations [32, 71, 35, 64, 105]. Most of these grippers utilize several digits arranged in a radially-symmetric pattern to envelop target objects, enabling strong power grasps even in uncertain conditions. However, these soft grippers typically lack the dexterity for in-hand manipulation.

Most high-dexterity soft robotic hands use anthropomorphic finger arrangements,




Figure 5.1: In this study, we are interested in the effect of finger arrangement on in-hand manipulation capabilities. a) Based on the 2-DOF fingers from [65], we controlled the resting pose of two fingers in the hand along an ellipse to smoothly transition between “4 finger cross” and “anthropomorphic” arrangements. Using simulations and b) physical prototypes at five discrete finger arrangements, we explore the effect of finger arrangement on manipulation of a wide variety of objects.

even though the kinematics of typical soft fingers are very different from human fingers. The RBO Hand 2 [36], BCL-13 [72], and BCL-26 [60] all use anthropomorphic finger arrangements despite also using pneumatic bending actuators as fingers. Conversely, the hand developed in [65] uses two orthogonal pairs of opposing fingers (“4-Finger Cross” arrangement), which decreases the complexity of mapping control inputs to object motions, but limits the sizes of objects that can be manipulated. All of these hands demonstrate excellent grasping and in-hand manipulation capabilities, but the question still remains: How does the arrangement of digits in a soft robotic hand affect in-hand manipulation if all other factors are identical?

5.2 Related Work

A small number of studies have investigated the effect of other design parameters on in-hand manipulation. In one such study, Feiz et. al found that the precision manipulation workspace of human hands becomes more restricted as the number of digits involved increases, but the range of controllable axes increases [111]. They find that using two digits for in-hand manipulation enables the largest workspace, but three digits may be more useful

to ensure objects can be moved in more directions. In addition, the effect of finger design and number of degrees of freedom on dexterous manipulation has been studied [60, 112], but usually in isolation from whole-hand design.

The effect of digit arrangement on grasping has also been studied in several application-driven cases. For robotic grasping, the ability to use either an antipodal grasp (two fingers opposing) or power grasp (three or more fingers in a circle) was found to be extremely useful to expand the range of graspable object sizes and shapes [22] compared to just one of those configurations. Several commercial grippers include mechanisms to switch between these two digit configurations on the fly including the Barrett hand [17], iHY hand [22], and Robotiq 3-Finger Adaptive gripper [113]. However, these two digit configurations are designed for grasping, and may not directly transfer to in-hand manipulation. Additionally, the effect of digit arrangement on in-hand manipulation has not yet been thoroughly studied for robotic hands.

5.3 Designing Digit Arrangement

In this study, we employ an empirical approach to study the effects of finger arrangement on in-hand manipulation performance. Our study utilizes the dexterous soft hand platform developed in [65] which has four modular dexterous fingers, each with 2-DOF, where the finger arrangement is fully adaptable. In prior work, this hand used a “4-finger cross” digit configuration with fingers arranged with radial symmetry, enabling an intuitive mapping between finger control inputs and object motions. The downside of this configuration, however, is difficulty handling objects with high aspect ratios, or with dimensions larger than the hand [65]. To combat these limitations, we propose a more “anthropomorphic” finger configuration with three fingers arranged along a line, with a “thumb” opposing them, as shown in Figure 5.1. This enables cylindrical grasps of objects much longer than the hand.

For the initial design phase, we used a simulated version of the dexterous hand inside of the SoMo simulation framework [68] for quick design iterations. In simulation, we define the digit arrangement with a variable, c_a , which enables the hand to morph between the “4-finger cross” ($c_a = 0$) and “anthropomorphic” ($c_a = 1$) designs by controlling the position and orientation of the two outside fingers around an elliptical shape. Since the finger pose can be controlled directly in simulation, we enable a fine-grained evaluation of

finger arrangement designs.

To validate the simulation results, we leveraged the modularity of our physical prototype hand to study five finger arrangements spanning the space from “4-finger cross” to “anthropomorphic” ($c_a = \{0.00, 0.25, 0.50, 0.75, 1.00\}$). For each finger arrangement, we designed a static scaffold to hold the fingers in the correct locations. This enables a coarser study of finger arrangements.

5.4 Performance Metrics for In-Hand Manipulation

We propose a task-based approach for quantifying in-hand manipulation performance that extends beyond the typical metrics used in this space. Most performance metrics for in-hand manipulation in the literature focus on the presence or absence of finger motions or specific capabilities [63, 60] rather than on quantifying overall performance on actual tasks. Instead, we introduce a set of task-based performance metrics that do not depend upon any particular hand/finger morphology and can be easily obtained from experiments. For any given hand design, we can define a set of motion primitive tasks (moving the object) that the hand should be capable of achieving. The hand can then be commanded to perform each motion primitive task while the object’s resulting motion is measured. Our performance metrics emerge from a comparison of the object’s deviation from the expected trajectory over many repetitions, as shown in Figure 5.2.

The performance metrics introduced in this study quantify the object’s motion along a desired axis as well as off-axis during a manipulation primitive. Looking in the desired axis, the *range of motion* describes the maximum position or orientation achieved by the hand, and *repeatability* describes the spread in the range of motion over multiple task




Figure 5.2: We introduce three task-based performance metrics for in-hand manipulation. The *range of motion*, *repeatability*, and *off-axis deviation* measure the performance of a given motion primitive task.

repetitions. The *off-axis deviation* describes the amount of undesired motion in all other axes (L_2 norm over the other five pose dimensions normalized by the maximum translational and rotational ranges of motion). Taken together, these metrics describe how well a hand can impart a set of motion primitives onto an object, and yields a practical method to evaluate the in-hand manipulation performance of any particular hand design.

In this study, the hand prototype used is only capable of controlling object motion in the plane parallel to the palm. We can thus define three motion primitive tasks in these cardinal directions: lateral translation in the x and y -axes, and rotation about the z -axis. In addition, the hand is controlled open-loop without explicit knowledge of object properties (as discussed in the “Finger Control” section), so exact object poses are not possible to prescribe. Instead, we can re-define our motion primitive tasks to involve moving objects along desired axes (with the goal of pure motion in one axis, and zero off-axis motion). The “*range of motion*” then becomes an observed measure of the amount of motion along the desired axis rather than a measure of pose accuracy with respect to a desired object pose. This simplification enables a comprehensive study of digit arrangements without any controller design.

5.5 Large-Scale Design Study

We performed a large-scale design study to understand how finger arrangements affect in-hand manipulation of a variety of objects. This study is performed exhaustively in simulations, and validated with a lower resolution using our physical prototype hand. The parameter space spans six key variables, as explained in later sections:

1. **Finger arrangement**, c_a
2. **Motion primitive task**
3. Trajectory design (discussed in next section)
4. Cross-sectional shape of object (square or circular)
5. **Characteristic dimension of object’s cross section**
6. **Aspect ratio of object** (perpendicular to cross section)

While finger arrangement, motion primitive task, and the object’s cross sectional size and aspect ratio are primary independent variables in this study, we also know that the shape of the object and the design of the actuation signals play a large role. The shape of the object (curvature and convexity) can have a large effect on the stability of grasps even before in-hand manipulation begins [28]. In addition, the design of actuation trajectories is a critical factor in successful in-hand manipulation performance. To understand these effects, we performed a preliminary study of actuation signals before the larger design study.

5.6 Finger Control and Trajectory Design

To control the hand (in both simulation and hardware), we developed two sets of heuristic pressure trajectories: one set for each major finger arrangement, where each set enables the three critical object motion primitives. These pressure trajectories were developed for a nominal object (60 mm cube) in simulation through trial-and-error (taking actuation limits into account), and applied to the physical hand using a simple linear conversion, as discussed in [68]. The trajectories are developed for open-loop motion primitives with no aid from a perception system or any other state information. To evaluate the ranges of finger arrangements where each trajectory is valid, we performed each motion primitive study twice, applying both trajectories for every finger arrangement design.

Before an exhaustive search of the parameter space, we first explored the effects of trajectory design in depth for a nominal object. As shown in Figure 5.3, there exists a clear trade off in performance when manipulating a 60mm cube with these two trajectory designs. Specifically, finger arrangements close to the *4-finger cross* design have higher performance when using the trajectory designed for the *4-finger cross* configuration, and vice versa for *anthropomorphic* designs. This is shown by the crossing pattern in the left column of Figure 5.3.

If we take the best-performing trajectory at every finger arrangement (largest range of motion), we obtain the right-column of Figure 5.3. These graphs show a gradual degradation in performance as each designed trajectory moves away from its reference finger arrangement, indicating that the combination of these two trajectories is robust to changes in finger arrangement. Out of the best-performing trajectories, the lowest range of motion generally occurred towards the middle of the arrangement space (near $c_a = 0.5$). For this object, the loss in range of motion due to imperfect trajectory design is 15% for x -translation,

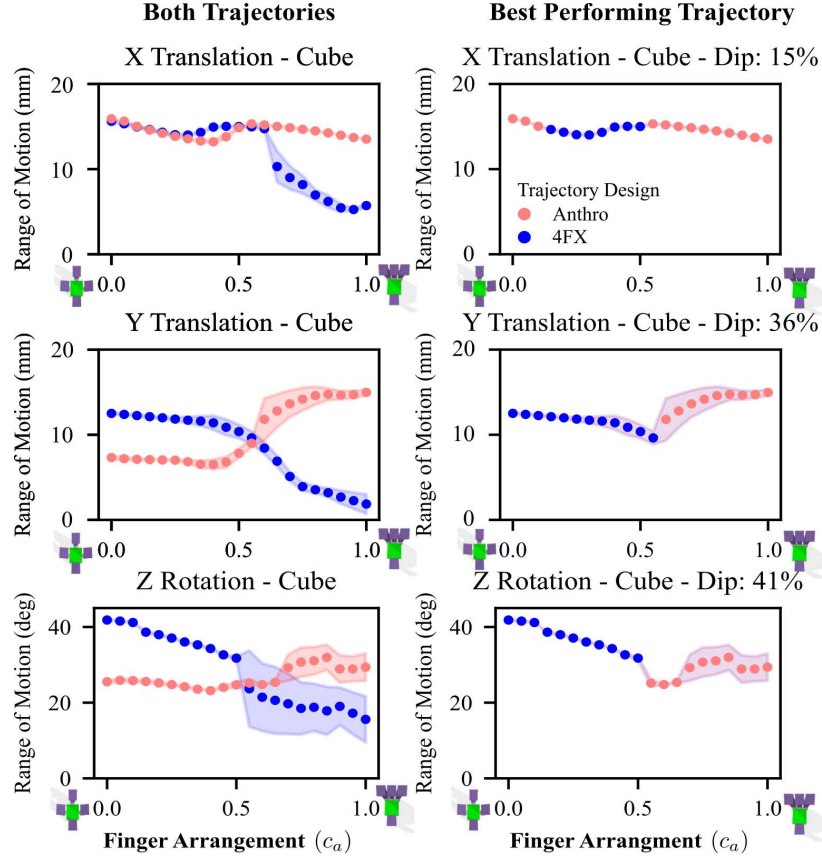


Figure 5.3: The trade off in performance as a function of the trajectory design is shown for a 60 mm cube in simulation. The left column shows the range of motion (lines) and repeatability (shaded regions around lines) of each manipulation primitive is shown with both trajectories. The right column shows the superposition of the best-performing trajectory at each finger arrangement, which exhibits only gradual degradation in the range of motion as finger arrangements move toward the $c_a = 0.5$. At worst, we lose 41% of the range of motion when performing a rotation about the z -axis at $c_a = 0.6$.




Figure 5.4: Our task-based approach to evaluating in-hand manipulation starts with performing a set of primitive motions while recording the object’s pose. An example of a rotation primitive performed on an object of 80 mm width and 0.5:0.5:1 aspect ratio with $c_a = 0.75$ (close to anthropomorphic) is shown a) in real hardware, and b) in simulation. c) The simulated object set consists of five widths (20 mm to 100 mm), five aspect ratios (0.25:0.25:1 to 1:1:4), and two shapes (box and cylinder). d) The real object set consists of several standard-sized everyday objects.

36% for y -translation, and 41% for z -rotation primitives. These percentages represent the percent decrease from the maximum achieved motion across the best-performing trajectory. While these performance losses are significant, we do not expect any macro-scale conclusions to be affected to a large degree by these heuristic trajectory designs.

5.6.1 Testing Procedure

The task-based performance metrics used in this study require a standardized testing procedure that works for simulations and real hardware. For a given object, finger arrangement, desired primitive motion, and trajectory design, the manipulation task is

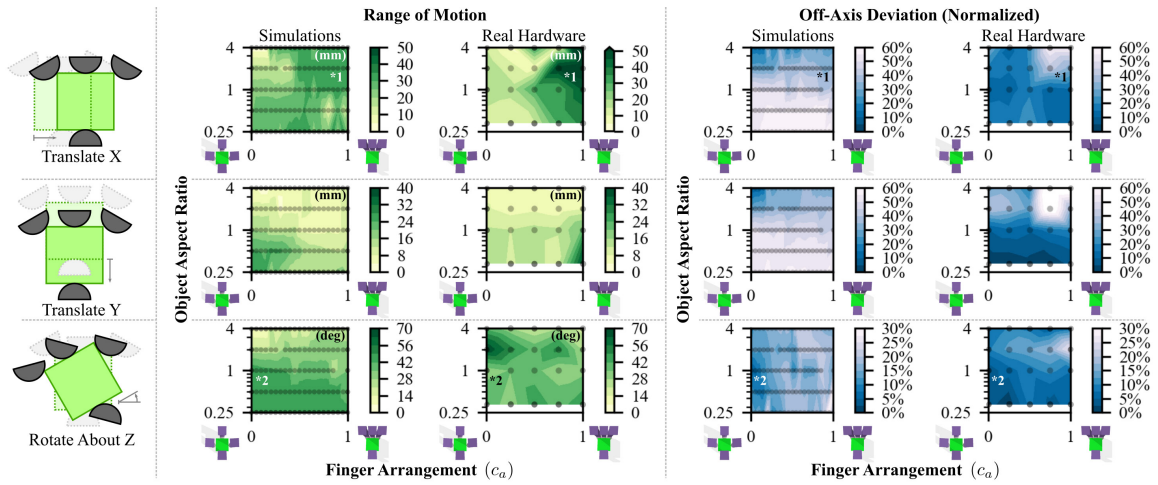


Figure 5.5: Each of the three motion primitive tasks is best-performed with different finger arrangements. Overall, “anthropomorphic” finger arrangements ($c_a = 0.75\text{--}1.0$) are good at performing translations in the x -direction, with a 30-50 mm range of motion, marked by “*1” in the plots. The “4-finger cross” arrangements ($c_a = 0.0\text{--}0.25$) are excellent at rotations about the z -axis, with 40-70° range of motion and 10-20% off-axis deviations, marked by “*2” in the plots. All finger arrangements struggled with y -translation of high-aspect ratio objects (above 2.0). All of these trends appear in real hardware and simulations, with reasonable agreement between the two for x -translation and z -rotation. Points on plots represent values of parameter space tested. Each point is an average over $n = 10$ trials in simulation, and $n = 4$ trials for real experiments. Contours are linear interpolations.

executed several times while measuring the object’s 6D pose, as shown in Fig. 5.4. For ease of testing, objects are initially set on the ground with the hand approaching from directly above.

Simulations and real-life tests were preformed according to similar procedures, with customized steps for each. In the simulations, actuation torque signals are applied to the fingers while object’s pose is logged, enabling performance metrics to be calculated directly from the data. However, a few simulations became numerically unstable; these points were excluded from the results. For real-life experiments, actuation signals are converted to actuation pressure signals and executed with a real-time pressure controller [108] (same as in [65]). In addition, objects are fitted with April Tags [103] for pose tracking, and manipulated on a clear acrylic table with a camera viewing the scene from below. All real hardware experiments are controlled using Robot Operating System (ROS, [82]), and performance metrics are calculated using the pose estimates.

5.6.2 Object Set

To test the effects of the size and aspect ratio of target objects on in-hand manipulation, we chose a suitable set of simple objects which span this space. For the simulations, we used a procedurally-generated set of geometrically simple objects, where we control the cross-sectional size and the aspect ratio with only two unique parameters. We chose two basic geometric primitives (cylinders and boxes), five cross-sectional widths ranging from 20 mm (fingertip diameter) to 100 mm (largest object to fit inside the resting hand), and five aspect ratios ranging from 0.25:0.25:1 to 1:1:4. Overall, $2 \times 5 \times 5 = 50$ simulated objects were generated using this discretization, as shown in Figure 5.4c.

For real-world testing, we used a sparser set of everyday objects and food items. To validate our simulation results, we used a combination of packing boxes for various retail items (raspberry pi camera, jello, and pneumatic connectors), and also a plastic jar and tube. These objects roughly span the space of sizes and aspect ratios of the simulations (as shown in Figure 5.4d). Finally, in a real-world demonstration, we used two delicate pastries (sweet bun and cupcake) as shown in Figure 5.7.

Several object properties were ignored in this study, but remain important for future explorations. Since we are interested in the geometric properties of hand design, the mass and mass distributions of objects were not varied, however the mass of objects

can drastically affect real-world tasks. In simulations, all objects had a constant mass of 0.010 kg to minimize friction effects with the ground. We also designed the fingers to ensure that friction is high between the fingertips and the object, and chose testing surfaces to ensure sufficiently low friction between the object and the ground. Rolling contact between the fingertips and the object (forces always within the friction cone) during in-hand manipulation is also assumed, though not always true. While these simplifications restrict our testing environment, we do not expect them to affect our conclusions.

5.7 Results

Based on our experiments, we find that the finger arrangement of a soft hand affects not only the in-hand manipulation performance as a function of object sizes and aspect ratios, but also the overall success of motion primitive tasks. These results are summarized in Figures 5.5 and 5.6.

To obtain relevant slices of the 6D parameter space, the data is condensed by taking averages or maximums over some of the dimensions. Since two trajectory designs were tested for each experiment, we can collapse this dimension by taking the best-performing result (largest range of motion) at each point. Additionally, in the simulations, we took the mean value over the two object shapes (box and cylinder), but the physical tests only report results from boxes since the boxes spanned the entire space of aspect ratios. From here, slices of the parameter space are taken at relevant values of the object’s size and aspect ratio.

Figure 5.5 shows that each of the three motion primitive tasks is best-performed within separate ranges of finger arrangements. Taking a slice through the parameter space at mid-sized objects (60 mm), we see that “anthropomorphic” finger arrangements ($c_a = 0.75$ – 1.0) can perform high quality translations of high-aspect ratio objects in the x -direction, with a large range of motion (30-50 mm), but also large off-axis deviations ($\sim 40\%$). Conversely, the “4-finger cross” arrangements ($c_a = 0.0$ – 0.25) are excellent at rotations about the z -axis, with large range of motion (40 – 70°) and small off-axis deviations (10-20%). All finger arrangements struggled with y -translation of high-aspect ratio objects (above 2.0) with less than 16 mm of motion. These results are summarized in Table 5.1. Overall, we see the same large-scale trends in real hardware and simulations, with the range of motions matching well. However, the off-axis deviation for translational motions have op-

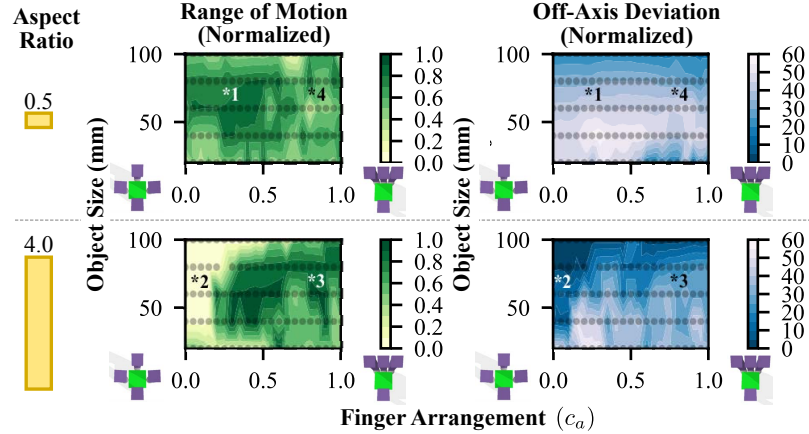


Figure 5.6: Object aspect ratio directly affects in-hand manipulation performance for all motion primitives. “4-finger cross” finger arrangements excel at manipulating larger objects with small aspect ratios (indicated by “*1” in the plots), and have extreme difficulty with high-aspect ratio objects (“*2”). Anthropomorphic finger arrangements perform well with high aspect ratio objects regardless of size (“*3”), but do poorly with objects of lower aspect ratio, (“*4”). Points on the plots represent values in the parameter space tested. Each point is an average over $n = 10$ trials in simulation. Contours are linear interpolations.

Table 5.1: Overall performance for each motion primitive

Finger Arr.	Max Range of Motion			Max Off-Axis Dev.		
	x-tran	y-tran	z-rot	x-tran	y-tran	z-rot
4-Finger X	20 mm	24 mm	70°	25%	40%	20%
Anthro.	50 mm	12 mm	45°	40%	60%	30%

* Best-performing finger arrangements are marked in **bold**.

posite trends in hardware compared to simulations, likely due to limitations of the modeling framework when implementing high-friction fingertip surfaces.

Figure 5.6 shows that in-hand manipulation performance of each major family of finger arrangements is tied directly to the aspect ratio of the object. Taking an average over all primitive motions (with ranges of motion normalized by the maximum values for each primitive), we can see that the “4-finger cross” arrangement ($c_a = 0.0$ – 0.25) and even the intermediate arrangements ($c_a = 0.25$ – 0.6) successfully manipulate large, thin objects. For an aspect ratio of 0.5, the “4-finger cross” has a large range of motion (80% of maximum) with moderate off-axis deviation (20–30%) for objects larger than 50 mm. In addition, the “4-finger cross” arrangements produce extremely poor performance with high-aspect ratio objects (essentially no motion produced). On the other hand, “anthropomorphic” finger ar-

Table 5.2: Overall performance as a function of object aspect ratio

Finger Arr.	Max Range of Motion		Max Off-Axis Dev.	
	aspect: 0.5	4.0	0.5	4.0
4-Finger X	80%	0%	30%	60%
$c_a = 0.50$	80%	100%	60%	40%
Anthro.	40%	100%	40%	30%

* Best-performing finger arrangements are marked in **bold**.

rangements ($c_a = 0.75$ – 1.0) perform well (50-100% of maximum range of motion) with high aspect ratio objects regardless of size, but do poorly (20-40% of maximum range of motion) with objects of lower aspect ratio. Lastly, there appears to be a local performance maximum for finger arrangements between $c_a = 0.25$ and $c_a = 0.5$, with good performance (above 60% of maximum range of motion) over both aspect ratios. These results are summarized in Table 5.2.

Finally, while repeatability was evaluated, no significant trends existed. In the physical experiments, the standard deviation of the amplitude of motions never exceeded 24% of the maximum range of motion for over all primitive motions, and the simulations yielded a standard deviation of at most 28% of the maximum range of motion. This larger variance in the motion of objects is likely due to variation in friction forces with the ground as well as slippage between the fingertips and the object.




Figure 5.7: When manipulating two delicate pastries, two different finger arrangements are required to effectively perform fine pose adjustments. Due to the high aspect ratio of the sweet bun, the “anthropomorphic” finger arrangement performs the best translation in the x -direction, with the “4-finger cross” arrangement damaging the bun (marked with “*”). Similarly, due to the lower aspect ratio of the cupcake, the “4-finger cross” arrangement performs the most-effective rotation about the z -axis, with the “anthropomorphic” arrangement inducing large off-axis motion (marked with “†”).

5.8 Translation to Real-World Tasks

The results from our study with simple objects and simple motion primitives can be extended to pick-and-place operations on delicate real-world objects. In Figure 5.7, two different finger arrangements are required to effectively perform fine pose adjustments on two pastries with different aspect ratios (sweet bun and cupcake). For each of the two major finger arrangements, we attempt to translate the bun in the x -direction before releasing, and rotate the cupcake about its z -axis to reveal a fiducial marker. For these demonstrations, the hand is fixed to a UR5e 6-DOF robot arm (Universal Robots), and poses are heuristically-designed without the aid of a perception system. The hand and robot are coordinated using ROS [82].

To translate the high-aspect ratio sweet bun (1:1:3.3, 30 mm×30 mm×100 mm) in the x -direction, the “anthropomorphic” finger arrangement performs best, translating the bun without damage. However, the “4-finger cross” finger arrangement damaged the bun (applied a large stress from the fingertip) while grasping due to limitations in the maximum dimension that can fit within the fingers when retracted. The bun was damaged further when attempting to perform the translation primitive. These results are expected given the trends we see in our design study.

To rotate the cupcake (aspect ratio of 1:1:1, 50 mm×50 mm×50 mm) about its z -axis, the “4-finger cross” arrangement performs best, with a stable grasp and stable, consistent, well-controlled rotation within the hand. However, the “anthropomorphic” arrangement is unable to perform a stable grasp without a large (60 deg.) initial rotation, and rotation primitives had much larger off-axis motion during the motion. These observations also match the results of the design study.

5.9 Discussion

The trends presented in the design study and real-world demonstrations can be explained by a simple analysis of geometry and grasping forces. Based on this analysis, we extract several design rules for high-performance in-hand manipulation using soft hands. Our results also suggest that on-the-fly control of digit arrangement would enable a robot to maximize in-hand manipulation performance with arbitrary objects. These results can then be extended to other finger designs, actuation mechanisms, and combinations of degrees of freedom.

5.9.1 Geometry Explains Performance Differences Between Digit Arrangements

Static grasp stability can be used to explain why certain finger arrangements are better suited for specific tasks. When grasping, the “4-finger cross” arrangements apply contact forces to the object in an axisymmetric way, thus excelling at axisymmetric motion primitives such as z -axis rotations. Conversely, the “anthropomorphic” arrangements apply contact forces to the object with several fingers on one side balanced by an opposing thumb, leading to high-quality linear motion primitives.

A simple analysis of geometric features of each finger arrangement can help explain the strong effect of object aspect ratio on in-hand manipulation performance over all motion primitive tasks. While all finger arrangements have practical size limits on graspable objects, “anthropomorphic” finger arrangements have an effectively infinite size range parallel to the direction of the fingers. This liberates “anthropomorphic” hands to grasp high-aspect ratio objects with ease, and in-hand manipulation performance is high as a consequence. However, the opposing nature of the forces produced during grasping make stabilizing objects with smaller aspect ratios challenging for “anthropomorphic” hands.

Conversely, “4-finger cross” arrangements have object size limits in all directions, making it difficult to grasp high aspect ratio objects (and thus difficult to manipulate), but enabling more-stable grasps on low-aspect ratio objects.

This simple analysis reveals an important result: to maximize open-loop in-hand manipulation performance for a large variety of objects and tasks, active control of digit arrangement is necessary. Achieving the best of both worlds without additional control or planning can be achieved if a hand can adapt between these two finger arrangements with a 1-DOF mechanism. This result is consistent with findings in previous studies on hand design for grasping, where a 1-DOF mechanism is sufficient to achieve precision and power grasping [22].

5.9.2 Successful Digit Arrangements Depend on Finger Design

In all of these results, passive compliance plays a primary role in enabling successful in-hand manipulation. Bending compliance in the fingers is responsible for the ability to grasp and manipulate a wide range of object sizes without explicit knowledge of any properties of objects. Additional compliance in the fingertips helps ensure stability while manipulating objects due to the large, high-friction contact areas. Compliance also limits contact pressure on delicate objects (such as pastries) by ensuring contact forces are small and distributed over large areas.

Finally, some of the motion primitives included in this study are only possible due to the 2-DOF finger design we used. If 1-DOF fingers were used (where only the primary bending axis is controlled), then the “anthropomorphic” finger arrangement could not perform controlled x -translations since the fingertips only point in the y -direction. Similarly, the “4-finger cross” arrangement could not perform controlled z -rotations, as the fingers only point toward the center of the object. The 2-DOF fingers used in this study were specifically designed to combat these potential limitations, but the total number of controlled degrees of freedom could potentially be reduced in future iterations (i.e., mixing “simple” and “dexterous” fingers together in one hand).

5.9.3 Limitations

A number of design decisions limit the scope of the results presented above. The biggest limitation is that all in-hand manipulation in this study was performed open-loop,

and is highly-dependent on the exact finger trajectories used. Trajectories were designed specifically for the two major finger arrangements tested, but not for any intermediate finger arrangements. Without any knowledge of object sizes, on-board sensing in the fingers, or modeling, the performance of finger arrangements in the middle of the scale (near $c_a = 0.5$) could potentially be improved if trajectories were specifically designed for each finger arrangement. Furthermore, if a modeling framework were implemented, finger input trajectories could then be planned directly rather than designed by a human. Given a motion planning system for the fingers, we would expect overall improved performance, but the trends seen in our study would likely still hold.

The other main limitation of this study is our choice of the two main finger arrangements. These two arrangements are an excellent starting point, but this study does not involve any first-principles analysis or design. An first-principles analysis of the in-hand manipulation problem could yield a superior finger arrangement that falls outside of these two specific designs (for example, different number of digits, different mounting angles, etc). However, we believe the conclusions drawn in this study would still apply to finger designs falling within the two families of designs studied (anthropomorphic and axisymmetric).

5.10 Conclusions

In summary, we demonstrated that in-hand manipulation performance is directly tied to the arrangement of digits in a soft robotic hand, and that on-the-fly control of digit arrangement is necessary to achieve the best performance for arbitrary objects. Through a large-scale design study, we found that certain motion primitives are best-accomplished with different ranges of finger arrangements, with “anthropomorphic” arrangements performing well with x -translations, and “4-finger cross” arrangements excelling with z -rotations. We also found that the aspect ratio of the object affects in-hand manipulation performance over all motion primitive tasks, where high aspect ratio objects are best-handled with anthropomorphic designs, and “4-finger cross” arrangements performing better with low aspect ratio objects. Finally, we demonstrated that these findings extend to a real-world manipulation task where gentle in-hand manipulation is desired.

In future work, we plan to develop a soft robotic hand with one additional degree of freedom to control the digit arrangement on-the-fly. We are also interested in developing a modeling and motion planning framework for soft fingers to directly plan finger motion to

achieve desired object motions. Given such a modeling framework, new finger arrangements could be explored beyond those a human designer can manifest. Another additional avenue would be to explore how to distribute the total number of controlled degrees of freedom such that we maximize the benefits of passive compliance while still enabling useful in-hand manipulation. Finally, integrating soft dexterous hands into more real-world testing scenarios will help push forward the development of safe, highly-capable soft hands suitable for real applications in human-centric environments.

6

Designing dexterous palm surfaces

6.1 Introduction

Substantial advances in gripper and hand design have occurred over the past four decades, leading to the development of many robotic hands optimized for grasping. For all of these hands, the design of the palm (if one is present) is a high-friction and often compliant material which aides the fingers in securely grasping the object. For example, rigid robotic hands (such as the Barrett Hand [17] and Robonaut hand [11]) have demonstrated strong and precise grasps, but require sophisticated motion planning when uncertainty exists in perception of the objects or the environment. Compliant hands and grippers (such as SDM, iHY, CLASH, RBO Hand 2) are mechanically robust to uncertainty in the environment, so they tend to perform better on a large variety of objects, however, often at the expense of strength or precision [21, 22, 36, 35, 25, 39, 60]. Jamming-based grippers (such as the universal gripper [34]) add a new and interesting dimension to grasping, since they are highly compliant during the initial grasp, but stiffen substantially when jammed to provide strong grasps. In some gripper designs, the palm is also omitted entirely in favor of finger compliance [21, 19, 35, 64].

In addition to the passive palm designs of the aforementioned hands, many recent works have developed actuated palm structures which help improve grasp robustness. Meng, *et al.* developed a soft, telescopic palm device designed to absorb impacts with objects [114]. Jamming-based palm designs have also been employed to easily conform around objects to further secure them in a grasp with the fingers [115, 116]. The shape of the palm has also been explored, where Capsi-Morales, *et al.* used a motion-synergistic approach to design a

palm that can change concavity to increase the workspace of the Pisa/IIT SoftHand [117]. Several actuated palm designs in the literature also focus on controlling the base positions of fingers rather than utilizing the palm itself as a control surface [118, 119, 120, 121]. While these studies yield important results for grasping, the role of the palm changes when performing in-hand manipulation.

For in-hand manipulation, the goal is no longer to secure the object, but to control its motion relative to the hand. Most recent advances in hand design for in-hand manipulation focus on increasing finger dexterity [10, 60, 65, 61, 122]. However, the palm’s surface takes on new functions during in-hand manipulation: a convenient surface to support the object while the fingers maneuver it. Some work has been done investigating the effect of palm design on in-hand manipulation [66, 60, 123], but this area is still ripe for exploration.

6.2 Controlling the Position of the Palm

6.2.1 Introduction

The palm can contribute to hand dexterity through controlled shape or position relative to the fingers. The importance of the palm is especially evident in human hands during grasping and in-hand manipulation. For the dominant human hand, the palm provides a higher force contribution than other anatomic areas to the hand’s total grip strength [124]. In addition, palm surfaces are sometimes used in robotic grasping to provide stability to objects during power grasps [22, 36, 72] or act as convenient surfaces to support objects during in-hand manipulation [65]. Utilizing an active palm as a movable/deformable contact surface has been demonstrated in a soft hand, but can have limited utility when fingers are extremely dexterous. For example, the BCL-26 Hand has 26 controlled DOF’s, with three DOF’s in the the palm [60]. However, some evidence suggests that humans only utilize five to six DOF’s during grasping [125], and seven to nine DOF’s during in-hand manipulation tasks [126]. Understanding how to control all 26 DOF’s proves challenging, and may not be necessary from a strict task-based perspective.

In this section, we demonstrate a soft robotic hand with an actuated palm capable of increasing in-hand manipulation capabilities by more effectively utilizing limited finger dexterity. We first discuss our soft robotic hand platform and important palm design parameters, choosing to focus on the height and diameter of the palm. Next, we used a




Figure 6.1: Our 1-DOF actuated palm can adjust its surface position and diameter to enable in-hand manipulation of a variety of object sizes while preventing interference between the palm and fingers.

combination of physical and virtual experiments to explore the effects of the palm design on in-hand manipulation performance. The results of the design exploration suggest that no single palm design can enable successful in-hand manipulation for all objects. Based on these results, we propose an actuated palm design with one degree of freedom which can change its height and diameter to accommodate a larger range of object sizes. Finally, we demonstrate the palm’s utility during manipulation primitives and a complex in-hand manipulation task inspired by stocking display shelves.

6.2.2 Achieving High-Quality In-Hand Manipulation

Achieving precise, repeatable, and sufficiently fast in-hand manipulation with a soft robotic hand is challenging due to finger compliance. A few observations can help in designing a system that meets these requirements. Treating in-hand manipulation as a similar problem to multi-limb manipulation, we can observe that the object has its largest range of motion if it is in contact with the end of the manipulator as opposed to the middle. This same rule applies for soft fingers with limited controlled degrees of freedom: a larger dexterous workspace is achieved if objects are manipulated by the fingertips.

Another key insight is found when observing the role of the palm as a mechanical

stabilizer for the object. This stabilizing effect creates a “local mechanical ground”, giving the hand a foundation to bear the weight of heavy objects. Additionally, the palm constrains the object’s pose, which removes uncertainty during in-hand manipulation by reducing the possible number of DOFs in the object’s position.

Taken together, these two insights imply that objects are best-manipulated when in contact with the fingertips, and when a palm or external surface provides a stable platform. Thus, we want our palm to be in contact with the object whenever possible. Furthermore, the effect of gravity on a grasped object’s rotational stability is minimized when the fingertips make contact close to the object’s center of mass.

Finally, the soft robotic hand platform used in this study was demonstrated previously in Abondance et. al [65]. It can achieve planar in-hand manipulation of a variety of objects using a symmetric array of four fingers, each with two degrees of freedom. In our previous work, this hand had a simple passive palm, which produced an inability to support small objects at the correct height to be grasped by the fingertips.

Our goal is to achieve these two basic design constraints in our soft hand platform for a variety of object widths and heights, thereby enhancing the hand’s manipulation capabilities. To achieve this, we could either 1) add one additional DOF per finger (4 DOF’s total in our hand platform) to enable the fingertips to move in and out with respect to the palm, or 2) add 1DOF in a dexterous palm that can move its surface in and out with respect to the fingers. In this study, we chose to investigate the role of the palm for in-hand manipulation.

6.2.3 Design Exploration using Passive Palms

To understand the effects of the palm’s physical dimensions on the overall in-hand manipulation performance of the hand, we utilized a new simulation platform for soft robots (SoMo, [68]) to perform a large-scale series of experiments on a simulated hand, then validated the observed trends on physical hardware. Examples of these experiments are shown in Figure 6.2. Overall, the simulations captured the trends in manipulation performance of the hardware platform, and results matched for key phenomena. These experiments validate our intuition about the coupling of palm diameter and height to avoid interfering with finger motion.

To evaluate the grasping and manipulation capabilities of the simulated and actual




Figure 6.2: To evaluate the effect of various design parameters on in-hand manipulation performance, we used a simple analysis of the finger’s workspace, and performed large parameter sweeps using the SoMo framework [68]. a) A schematic is shown of the palm design space compared to the finger workspace. b) We tested several simple objects of varying shapes, sizes and aspect ratios, and c) included three basic in-hand manipulation primitives. d) Simulations were validated using several isolated experiments on real hardware.

hardware as a function of the palm design, we defined four metrics related to task performance. The first two metrics are binary checks during each manipulation task. “Contact success” is true when the fingertips make contact with the object before manipulating, and “manipulation success” is true if the hand can perform the manipulation primitive for an extended time without dropping the object. These two metrics directly inform palm design: “contact success” acts as a filter to ensure a wide range of objects can be grasped, and “manipulation success” refines the design space to ensure in-hand manipulation is possible. Next, for all successful manipulation tasks and objects used, we defined two more metrics related to the object’s dimensions: the minimum object height and width successfully manipulated. These four performance metrics enable use to quantitatively evaluate palm design in a task-relevant way.

6.2.3.1 Palm Design

When designing a palm for our soft hand platform, several physical properties of the palm’s surface were considered, but ultimately only the height and diameter were explored in-depth due to their specific utility for in-hand manipulation. The height of the palm’s surface is critical to manipulation performance, as it determines where the fingers make contact with the object. The diameter of the palm is also important for maintaining the range of object sizes that can be manipulated, as it should not interfere with the fingers. Two additional design parameters were considered (the palm’s friction and concavity), but initial testing indicated that the height and diameter were most-promising as long as the palm’s surface has sufficiently low friction.

6.2.3.2 Experimental Setup

To ensure no interference between the fingers and palm, we can examine a simple profile of the finger’s inner surface during a grasp of an arbitrarily thin object, as shown in Figure 6.2a. This profile shape produces an estimate of the maximum palm diameter for a given height. Palm designs within this boundary will not interfere with with fingers during grasping. However, this analysis does not hold during in-hand manipulation tasks due to 3D kinematics of the fingers, so a more comprehensive set of experiments is necessary to fully explore the design space.

To understand the design space in the context of in-hand manipulation, we utilized

a new simulation framework (SoMo, [68]) to rapidly explore the effects of these design parameters. In this framework, soft fingers are modeled as hyper-redundant rigid-body serial manipulators with compliant joints. For these simulations, we used the mass, stiffness, and blocked force measurements of the soft fingers from [65] to calibrate the simulated fingers. The calibration from bending beams to discretized links with spring joints is described in [68], and results in actuator-level accuracy in blocked force of 0.28 N, and task-level accuracy of 9% for a complex finger gait [68].

With a sufficiently accurate simulation framework, we explored our palm design space through a series of simulated experiments. Over a variety of objects, the hand was commanded to perform three planar in-hand manipulation primitives using heuristically-designed actuation torque trajectories. While co-design of hand morphology and control would produce a fairer assessment of the design space, this is intractable without a detailed dynamic model, and is thus outside the scope of this chapter. Additionally, this design study was performed for the case where the palm is most-utilized (hand is placed with the palm facing vertically), but results apply beyond this simple case. The object set consisted of simple geometries with sizes that span typical household objects: three spheres (20, 60, and 100 mm diameters), nine cylinders of varying aspect ratios (3x3 grid of height and diameter with 20, 60, and 100 mm), and nine boxes (3x3 grid of height and width with 20, 60, and 100 mm). All objects had a fixed mass of 50 g, which is similar to many common household objects, such as those found in the YCB object set [83]. The three in-hand manipulation primitives were the same as in [65]: translation and rotation in a plane perpendicular to the palm’s surface. For all experiments, we recorded the 3D pose of the target object, contact points between the hand and object, and the finger actuation signals. The results of these experiments are shown in Figure 6.4.

In addition to the higher density parameter sweeps in simulation, we validated these results with sparser experiments on physical hardware. We built an array of circular palms with two diameters (70 mm and 85 mm) and two heights (60 mm and 80 mm). All palms were 3D printed on FDM Printers using PLA (using a CR-10s printer, Creality). The palms are interchangeable on the soft hand platform. The object set used consisted of three boxes (23, 30.5, and 61.5 mm in thickness), two cylinders (Cups 1 and 10 from the YCB object set [83], with diameters of 47 mm and 87 mm respectively), and one wooden sphere 63.5 mm in diameter. Using a top-down camera and April Tags [103] to track the object’s pose, we commanded motion primitives using the same heuristic actuation torque

trajectories, but converted to actuation pressures for the physical hardware. Actuation pressure signals were generated with a custom real-time pressure controller (same as in [65]). The results of these experiments are shown in Figure 6.3.

6.2.3.3 Results of Design Exploration

The results of our extensive simulated experiments suggest that a single palm height and diameter is not capable of manipulating all object sizes, which is supported by the experiments in real hardware. These findings suggests that a palm that decreases in diameter as a function of height can achieve manipulation performance for all object sizes.

The trends in bulk performance as a function of the palm design and object width appear to match between simulation and real hardware. In Figure 6.3, the contact success rates and manipulation success rates are shown for both systems as a function of object width. Several key real-world phenomena are preserved in simulation, shown in the callout pictures. For the large, tall palm (85 mm diameter, 80 mm height), we see that small objects cannot be grasped by the fingers, leading to zero success rate. Conversely, larger objects can be grasped and manipulated with high success. On the other end of the design spectrum, the short, small palm (70 mm diameter, 60 mm height) could grasp and manipulate all object widths can be grasped with relatively high success (above 50%), but the fingertips touch far from the center of mass of these objects. Simulations were only performed at three object widths (20, 60, and 100 mm) due to the high dimensionality of the sweeps, but the results still agree well with the hardware.

From here, the in-hand manipulation performance of the large-scale simulations over all objects and tasks as a function of the palm design are displayed in Figure 6.4. Using our performance metrics defined earlier, the results for each palm design (height and diameter) are averaged over all 21 objects and three manipulation tasks. First, Figure 6.4a shows the "contact success rate" which describes the ratio of trials where the fingers successfully contact the object at the beginning of the manipulation task. Next, we defined the "manipulation success rate" as the fraction of runs where the manipulation task is completed without the object falling out of the hand, as shown in Figure 6.4b. Finally, for all trials that successfully completed 10 task repetitions, the shortest object manipulated at each set of palm parameters is shown in Figure 6.4c, and the smallest object widths are shown in Figure 6.4d.

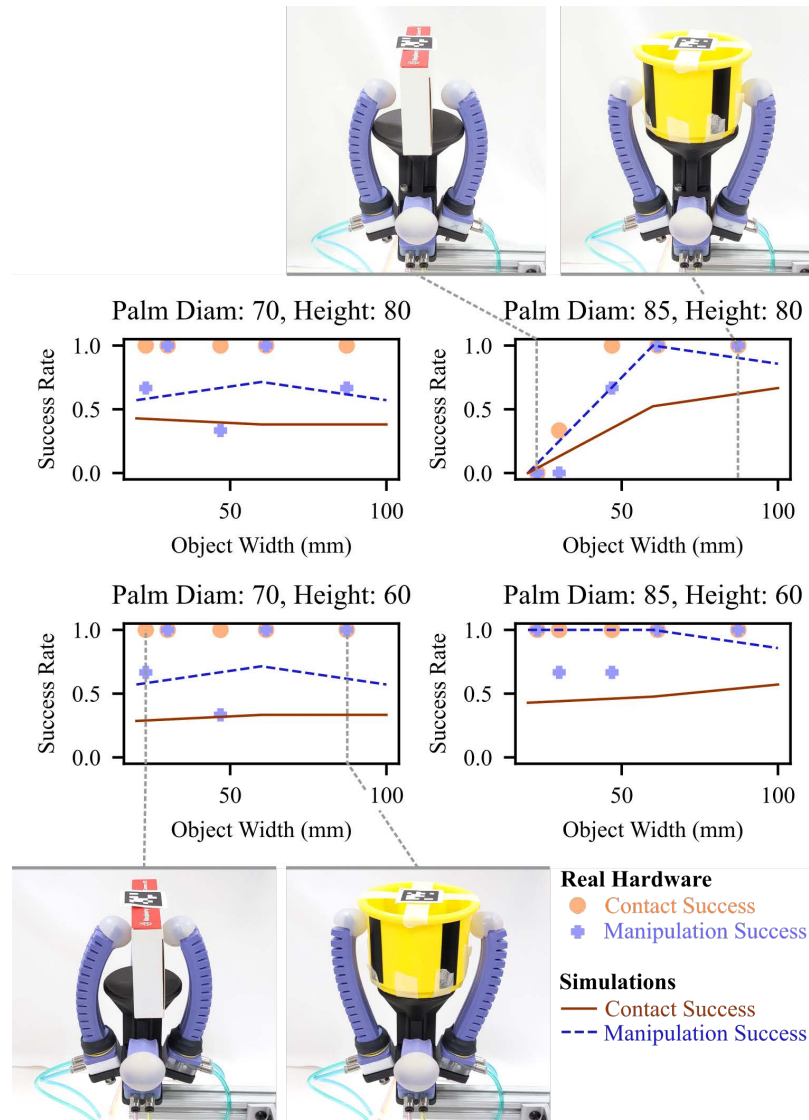


Figure 6.3: The main trends in manipulation performance from the simulated sweeps agree with the performance of real hardware. For the tall, large palm (85 mm diameter, 80 mm height), small objects cannot be grasped by the fingers leading to zero success rate, whereas larger objects can be grasped and manipulated with high success. Conversely, for the short, small palm (70 mm diameter, 60 mm height), all object widths can be grasped, but the fingertips touch far from the center of mass. *Each point from real hardware represents the average over the three primitive motions. Each point in simulation represents the average of 21 runs, (three tasks and seven objects), and are shown as lines for graphical clarity.*

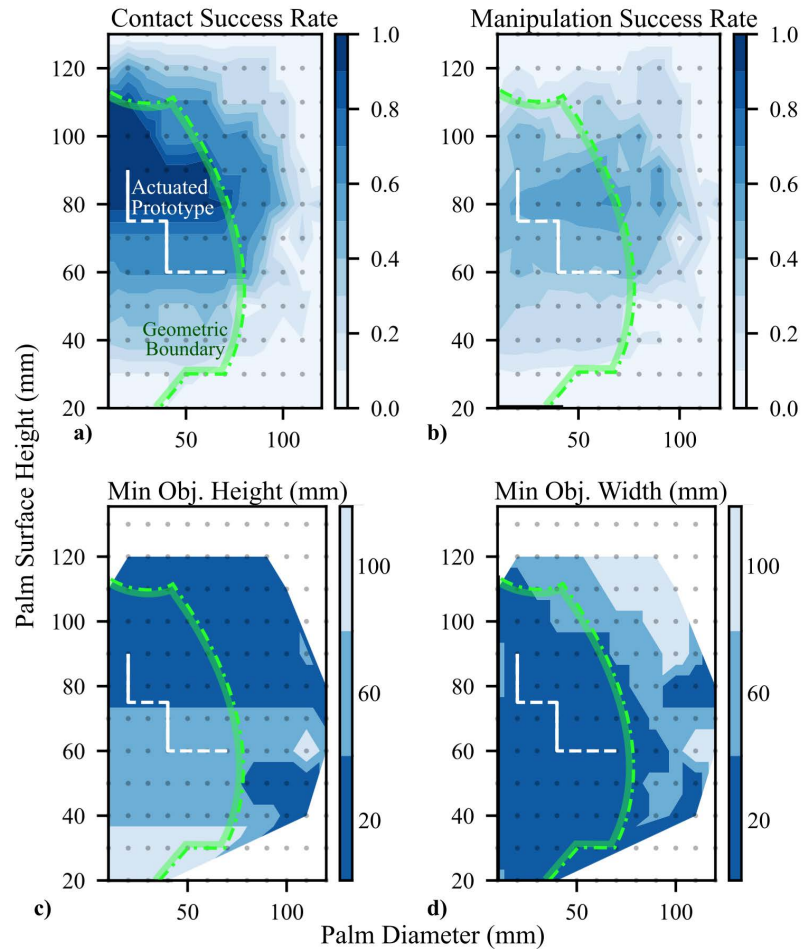


Figure 6.4: The palm height and diameter exhibit tradeoffs in grasping and manipulation performance over a variety of objects and manipulation tasks. a) The ratio of trials where the fingers successfully contact the object before manipulating shows a clear maximum near smaller-diameter, taller palms. b) The manipulation success rate reaches a maximum near taller palms of moderate diameter. c) The minimum object heights that were successfully manipulated is dependent only on the palm height and d) the shape of the boundary where the hand successfully manipulates thin objects is similar to the shape of our soft fingers (geometric boundary). Our proposed actuated palm with discrete sections changes diameter as a function of height.




Figure 6.5: Our actuated palm design expands the range of object sizes that can be successfully manipulated by controlling the object’s vertical position such that the fingertips touch closer to the objects’ centers. a) A cut-away view of the actuated palm mechanism shows how the diameter of the surface decreases as a function of height. b-e) The palm can lift smaller-sized objects (such as a lime) into place for the fingers to grasp, and f-i) manipulation tasks can be performed on a variety of object sizes.

The results illustrate a clear region of design space that produces successful in-hand manipulation. There appears to be a local maximum in contact success rate near palms of smaller diameter, but larger height (near 40 mm diameter, 90 mm height), as indicated in Figure 6.4a. In addition, Figure 6.4b clearly shows that manipulation success is relatively high near similar palm heights, but toward the middle of the palm diameter range. However, these first two metrics only tell part of the story. Taking object size into account, we can clearly see that palms with lower heights lose their ability to manipulate short objects, as shown in Figure 6.4c. On the flip side, thin objects can only be manipulated by smaller, shorter palms as shown in Figure 6.4d. Furthermore, the shape of the boundary between successful manipulation of thin objects and failure is similar to the shape of our soft fingers when actuated. This makes sense when considering that the source of failures for large-diameter, taller palms is interference between the palm and the fingers.

Based on these results, we can see that no single palm design is capable of highly reliable manipulation for all object sizes, so we propose instead to design an actuated palm capable of adjusting palm height and diameter simultaneously. With one DOF, our proposed palm can transition between a large-diameter, low height palm design, as well as a small-diameter, tall palm. Considering several practical design factors discussed in the next section, we chose to use a discretized approach capable of three different palm diameters

which change as a function of height to traverse the design-space, as shown with dotted lines in Figure 6.4.

6.2.4 Design of the Actuated Palm

Based on the parameter exploration in the previous study, we designed a single palm with one DOF that achieves useful motion perpendicular to the palm’s surface, as shown in Figure 6.1. With a discretized telescoping mechanism in which each palm level has a different diameter-to-height ratio, our palm has the ability to move continuously along the axial direction without causing interference with finger motion. The palm utilizes a rigid mechanism, but is a stepping stone toward our ultimate goal of a soft-bodied palm.

The main constraint informed by our design exploration was determining the maximum palm diameter we could accommodate as a function of the palm’s surface height. To make the best use of one additional DOF, the palm diameter should decrease as a function of the palm height to ensure the palm does not interfere with finger motion. This relationship should follow the shape of the geometric boundary of the soft fingers when grasping an extremely thin object.

Our final palm actuation design utilizes a telescoping mechanism to achieve the necessary discretized changes in diameter as a function of the palm’s height, as shown in Figure 6.5a. Only three height levels were chosen for the purpose of demonstrating functionality, and all palm levels were 3D printed on an FDM Printer using PLA (using a CR-10s printer, Creality). The second and third levels are held at their maximum heights by a parallel spring system to isolate their motion, utilizing four springs for the second level to accommodate larger objects with greater mass. To control the height of the palm assembly, a rigid wire connects the third level palm with a servo below. With this one connection, the servo can pull the third level palm down to touch the second level, then continue to pull both the second and third level until they are flush with the base (Figure 6.5b).

With one degree of freedom, this palm can achieve a spectrum of palm designs through actuation as shown in Figure 6.5. The dimensions of layers were limited by the mechanism itself, however they fall into a reasonably-successful region. The first level of the palm (base) has a height of 60 mm and diameter of 70 mm. The second and third levels can travel vertically by 15 mm each. The diameter of the third layer was restricted by the size

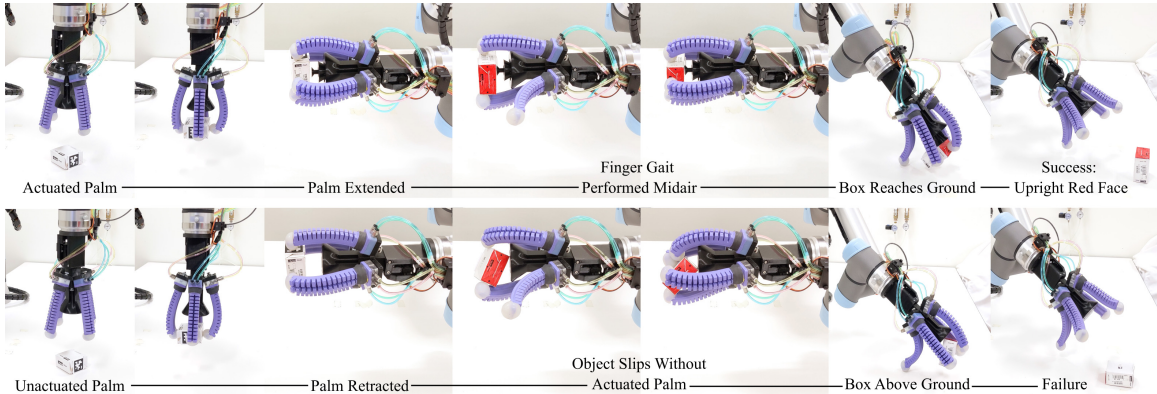


Figure 6.6: Having the ability to actuate the palm provides our soft dexterous hand increased stability when grasping smaller objects, which is useful for real-world tasks such as stocking shelves. This is especially evident for in-hand manipulation tasks performed in midair. In this example, the hand rotates the box using a finger gait developed in [65], with the goal of placing a particular face (in red) upright. If the actuated palm extends to support the object (top row), the task is completed successfully. However, if the palm is not extended (bottom row), the object slips sideways, and ultimately the task fails.

of the fingertips of our soft dexterous hand (20 mm). We chose a diameter of 40 mm for the second level to prevent collisions with the fingers. With these various palm diameters and heights, we can ensure our palm has the largest possible supporting surface at any given height while ensuring no interference with finger motion. While the design is directly tied to the geometry and kinematics of the fingers, we expect this design process to remain valid for a variety of soft finger and hand designs.

6.2.5 Results and Discussion

Overall, the actuated palm design translates to real gains in grasping and manipulation performance. With some limited knowledge of the object’s dimensions, we can adjust the palm’s height to allow manipulation to occur at the fingertips while ensuring no interference with finger motion. When left retracted at a height of 60 mm, the palm supports larger objects (such as the grapefruit or a large muffin in Figure 6.5) at the same height as the original hand in [65]. With the palm extended to its maximum height (80 mm), it can now support smaller objects such as a grape or small muffin. Furthermore, we can clearly see the benefit of the actuated palm’s ability to adjust the height of the object on-the-fly to support small objects (such as a lime in Figure 6.5b-e) at the correct height to be manipulated by

the fingertips.

To demonstrate the benefits provided by our actuated palm in a more-realistic task, we set up a manipulation task derived from the application of stocking shelves. A common task in shelf stocking involves placing items in a specific orientation that displays graphics towards the customer. Using a UR5e (Universal Robots) 6-DOF robot arm we performed a similar task manipulating the object within the hand in midair with the support of our actuated palm, as shown in Figure 6.6.

The goal of the manipulation task is to pick up a short square pneumatic fittings box (Parker), and rotate it such that the red side faces upwards when placed on the table. This box does not touch the retracted palm when picked up from the table. Next, two gait cycles of a finger gait for continuous rotation are executed by the hand. The gait compensates for the weight of the object and fingers (see [65]), but no other tuning is required. If the palm remains retracted, the finger gait fails, resulting in undesired motion of the box. Thus, when the robot attempts to place the box on the table, it drops the box instead. Alternatively, if the palm extends (increases in height) before the rotation begins, the box’s motion during manipulation is constrained by contact with the palm. This results in predictable object motion, which ultimately enables the robot to successfully place the box on the table with the correct side up.

6.2.6 Conclusion

In summary, we demonstrated an active palm that can change its height and diameter to enable high-quality in-hand manipulation for a large range of object sizes. The actuated palm provides stabilizing contact to objects, and helps control the axial position of objects with respect to the fingertips. We leveraged a combination of physical and virtual experiments to explore the effects of the palm’s design space on in-hand manipulation performance. The results of these experiments indicate that the palm’s diameter must decrease as a function of its height to ensure finger motion is not impeded. Using these results, we designed an actuated palm mechanism that controls the palm height and diameter relationship using one actuated degree of freedom. We demonstrate that the actuated palm enables successful manipulation of larger range of object widths and heights. Finally, the actuated palm enables the hand to perform a complex in-hand manipulation task in mid-air.

In future work, we expect that studying the role of the palm in soft manipu-

lation could yield a plethora of opportunities to improve the design and performance of soft hands. Expanding the palm’s design space to include more physical properties such as surface concavity, friction, and compliance could lead to insights about suitable fully soft palm designs. Further investigating palm design for hands with different finger arrangements, such as anthropomorphic, could also yield an understanding of the palm’s role more-generally in grasping and manipulation. Finally, an actuated palm such as the one proposed in this study could enable soft robotic hands to perform complex in-hand manipulation tasks in the real-world, such as activities of daily living in the home, using hand tools, or stocking display shelves.

6.3 Controlling Palm-Object Interactions via Friction

6.3.1 Introduction

Taking a step back from the palm design presented in the previous section, we note that in addition to the position of the palm’s surface, the contact interactions of that surface with the object are directly responsible for grasping and in-hand manipulation success. Importantly, for the same finger forces applied to the object, the palm’s surface can determine whether an object slides or tips on the palm. If these surface interactions could be actively controlled, then the robot could have more control over grasp stability, as well as greater capabilities for manipulating the object within the hand.

Examining the palm designs of common dexterous, compliant hands reveals a roughly flat surface for supporting objects, and either a high-friction material for grasp stability and pivoting (e.g., Shadow Hand [10], BCL-26 [60], RBO Hand 3 [61], and Pagoli, *et al.* [123]) or low-friction material to enable sliding [65, 66]. Palm compliance is also a key design factor to enable gentle interaction with delicate objects.

In this section, we show that active control of friction forces between objects and the palm of a soft robotic hand can be used to enhance in-hand manipulation capabilities and improve grasp stability. Through a simple analysis, we show how two key design parameters (friction coefficient and preload) directly control an object’s slipping vs. tipping behavior when finger forces are applied. We also show how mechanical compliance can be used to reduce the finger forces required to tip objects. We then demonstrate these concepts using our soft robotic hand platform by building two physical prototype palms which can actively




Figure 6.7: Varying the amount of friction force applied to the object by the palm (by utilizing either a *variable-friction* palm or a *variable-preload* palm) enables different in-hand manipulation primitives (sliding vs. tipping). (a) The *variable-friction* palm has a soft high-friction membrane underneath a rigid low-friction porous surface. Pneumatic actuation exposes the high-friction film, increasing the coefficient of friction. (b) The *variable-preload* palm has a compliant open-cell foam housed inside a thermoplastic elastomer pouch. With the object held securely by the fingers, pneumatic actuation raises or lowers the surface, thereby modulating the normal force.

control either the friction coefficient or normal force on the object, as shown in Fig. 6.7. We found that active control of both the *variable-friction* palm and the *variable-preload* palm allows the hand to translate and pivot the object given identical finger actuation, whereas a passive palm design could only achieve one of these manipulation modes. We show that the active palm designs also contribute substantially to grasp stability. Finally, we demonstrate that through sequential motion primitive operations, the active palms enable the soft hand to perform real-world manipulation tasks.

6.3.2 Grasping and In-Hand Manipulation Requirements

In this section, we identify two main capabilities that our robotic hand prototype should be capable of: 1) in-hand manipulation in the plane of the palm *and* out-of-plane, as shown in Fig. 6.7, and 2) robust, stable grasping. The ability to control five axes of object motion (lateral translations, lateral rotations, and rotations about the axis normal to the palm), allows the hand to move objects to any orientation through sequential primitive motions. Simultaneously, the hand must still maintain stable grasps to ensure it can withstand external forces (such as gravity or acceleration due to arm motions). Here we discuss how these goals lead to palm design requirements.

Our in-hand manipulation goal is to achieve both planar and out-of plane object motion. To achieve planar motion, the palm must act as a surface on which the fingers can slide the object, enabling x - y translations and z -rotations (as defined in Fig. 6.8). To achieve out-of-plane manipulation, we can utilize pivoting to rotate the object about a contact point or line on the palm without slipping (achieving rotation about x and y axes with some coupled translation). To enable both sliding and pivoting, these two competing requirements must be met either by increasing the dexterity of the hand (fingers, palm, or both), or utilizing extrinsic dexterity (such as gravitational forces, object dynamics, or environment constraints [46]). While increasing the number of controlled degrees of freedom in each finger could enable switching between planar and out-of-plane object motion[112], we show that the same effect can be achieved with a single DOF in the palm.

For strong grasping, the goal is to secure the object within the hand while resisting external forces. With compliant fingers, external forces will usually result in object motion, but finger compliance maintains the grasp even under large deformations. The failure criterion is governed by friction with some number of “virtual fingers”, of which the palm can

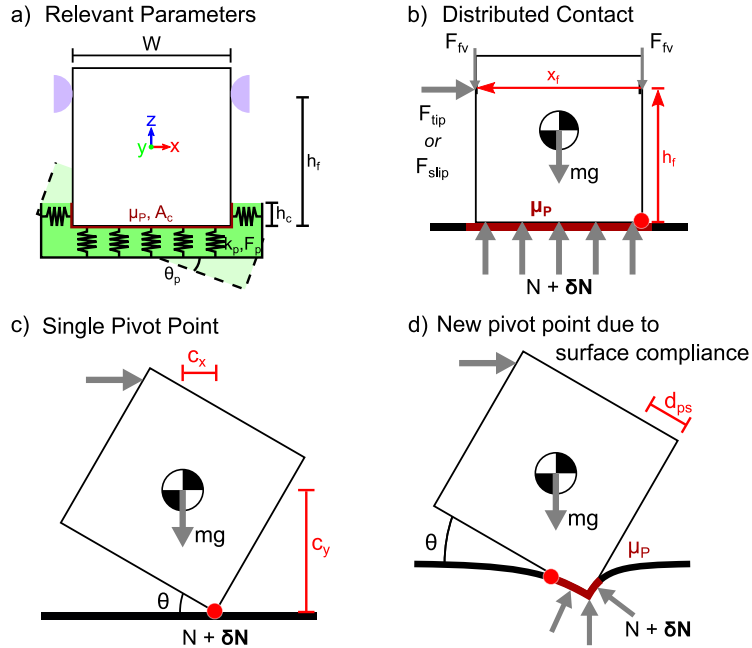


Figure 6.8: a) Schematic of physical properties to consider when designing a palm for a robotic hand. All properties affect grasping and/or in-hand manipulation performance, as detailed in Table 6.1. Both palm designs exploit contact switching between sliding and fixed contact to control whether an object translates laterally or rotates about a pivot point. b-c) For a flat surface, the friction coefficient directly determines whether an object will tip or slip. d) For a compliant palm surface, the palm provides a new lateral surface constraint and new pivot point when the object sinks into the surface.

act as one or more. The palm’s contribution to grasp stability occurs at small deformations, where it helps resist moments and lateral loads initially. Thus, our goal is to build a palm that can help resist the largest forces possible before the object displaces.

Toward the goal of enabling both grasping and in-hand manipulation capabilities with a gentle touch, we explore how the palm’s surface can be actively modulated for on-the-fly control of contact constraints on the object. In addition, palm compliance can enable gentle interaction with objects and make our in-hand manipulation goals easier to achieve with strength-limited fingers (as discussed below).

6.3.3 Design Parameters & Analysis

Table 6.1: Physical Properties of the palm and considerations for an active degree of freedom

Phenomenon	Physical Property	Practical Considerations for Active DOF	
		Pros	Cons
Surface Interactions	Friction coefficient	Directly determines tipping vs. slipping behavior when palm is flat	Depends on the pairing of objects and palm surfaces, so designs must consider a range of values.
	Surface roughness	Determines friction coefficient	Resulting friction coeff. is unpredictable depending on material choice.
	Contact area	Useful if location of contact could be controlled	Difficult to reliably control for a large variety of object geometries
Local Geometric Constraints (Conformability)	Conformal depth	Surface can act as a wall	Depends on the elasticity, vertical stiffness, and object mass.
	Stiffness	Can affect the force required to tip object	Behavior depends on elasticity
	Elasticity	Determines how the depth of confirmation is affected by external loads	Mechanisms with low elasticity and low vertical stiffness may be difficult to reset
	Preload	Palm can control the normal force on the object during grasping, thereby controlling the critical magnitude of friction forces	Strength of preload force must be matched to finger strength/stiffness.
Surface Location	Position (w.r.t. fingers)	Requires only a simple linear mechanism, as shown in [66, 114]	Width of the palm must decrease as a function of height to ensure no collisions with the fingers
	Orientation	Controls the direction of the surface normal, which could be used to direct object motion	Exact effect is unpredictable for a large variety of objects.

* Properties selected for further study are marked in *green-bold*.

There are a variety of physical properties that could be considered when designing a palm, but some of them are more straightforward than others to control. We present a list of options along with practical considerations for controlling them with an actuated mechanism. We consider three high-level phenomena that are relevant for the design of palms: surface interactions, local geometric constraints (i.e., conformability), and surface location. Fig. 6.8a illustrates the relevant physical properties associated with each phenomenon.

To identify the most-relevant physical properties to explore, we need to consider our goal of developing a palm with one actuated degree of freedom that enables strong grasping and switching of in-hand manipulation modes. Table 6.1 details the benefits and drawbacks of each property when considering it for use as an actuated degree of freedom. Furthermore, a simple analysis of tipping vs. slipping of an object on the palm’s surface illuminates two main relevant parameters, both of which are related to frictional contact with the object: the friction coefficient and preload.

6.3.3.1 Analysis of tipping vs. slipping

To understand how the friction and stiffness of the palm affect in-hand manipulation, we can look to a simple static analysis of the hand-object system at key time points. In both cases, designs that control either of these properties will exploit contact mode switching between sliding and fixed contact. Determining conditions for contact mode switching to occur has been studied extensively as an essential part of motion planing for pivoting of objects [127, 54]. The mode switching directly controls whether an object translates laterally or rotates about a pivot point when finger forces are applied. Fig. 6.8b-d illustrate how both parameters affect this contact mode switching behavior.

For a palm with a rigid, flat surface with friction coefficient μ_p , we analyze whether the object will tip or slip. This analysis consists of two conditions to check: the force required for the object to slip, F_{slip} (obtained by balancing forces), and the force required to tip, F_{tip} (obtained with a moment balance about the pivot point). The condition with lower force determines which motion will occur. For a palm surface orthogonal to gravity, we analyze the point at which the object just starts to slip or tip, resulting in:

$$F_{slip} = (N + \delta N)\mu_p, \quad N = 2F_v + mg \quad (6.1)$$

$$F_{tip} = \frac{mgc_x + F_v x_f}{h_f} \quad (6.2)$$

where N is the normal force from the palm, δN is an additional preload applied by the palm, F_{fv} is the vertical reaction force applied by the fingers, m is the object's mass, g is the acceleration due to gravity, c_x is the horizontal distance from the pivot point to the center of mass, h_f is the vertical height of the fingers with respect to the palm, and x_f is the horizontal distance between the finger and the pivot. For any given object (constant m, c_x) and hand (constant h_f, x_f), the friction coefficient μ_p and preload δN are the only controllable parameters.

Our initial design goal is to actively control the palm's surface to select whether objects slide or pivot within the hand for a wide range of objects. Examining (6.2) shows that for a given hand (constant h_f and x_f), the object mass and size determines F_{tip} , while (6.1) indicates that only the mass affects F_{slip} . To maximize the potential for pivoting over a large range of object masses, mass distributions, and sizes using the two controllable parameters, we must maximize F_{slip} by maximizing μ_p or δN . Conversely, to maximize the potential for objects to slip, we must minimize F_{slip} by minimizing μ_p or δN . Thus, only a binary mechanism is required, where the difference between high and low states is maximized.

For a palm with a compliant surface, the previous tip/slip analysis applies for objects that are lightweight compared to the palm's stiffness. Such low-mass objects sink into the palm a negligible amount under their own weight at rest. However, when objects are in the tipping mode and the palm's stiffness is suitably matched to the object's mass, the palm's compliance allows the object to sink into it. This provides a new surface constraint and shifts the pivot point toward the center of mass (as shown in Fig. 6.8c-d), which plays an interesting role: decreasing the overall finger force required to tip the object. To analyze this effect, we can perform a static moment balance for an object of width W pivoting on a soft surface at angle θ such that the pivot point shifts toward the center of mass by a distance d_{ps} . This results in (6.3):

$$\begin{aligned}
 F_{tip} &= mgX_c/Y_f, & (6.3) \\
 X_c &= c_x \cos\theta - c_y \sin\theta - d_{ps} \cos\theta, \text{ and} \\
 Y_f &= W \sin\theta + h_f \cos\theta - d_{ps} \sin\theta.
 \end{aligned}$$

where c_y is the vertical distance from the pivot point to the center of mass. The pivot shift distance (d_{ps}) is related to the stiffness of the palm, the mass of the object, the contact area,

and the pivot angle. However, for this analysis we assume it is a small constant distance that can be prescribed for simplicity. Thus, we show that a compliant palm decreases the finger force required to tip an object via a small shift in the pivot point. For a typical square object with a pivot adjustment distance (d_{ps}) of 10% of the object’s width, the required finger force to tip the object decreases by 20% compared to if the palm were rigid ($d_{ps} = 0$), which could be extremely beneficial for hands with limited finger strength.

6.3.3.2 Final Palm Designs

We developed two pneumatically-driven palm designs, with each design exploiting one of the two critical design parameters, as shown in Fig. 6.9. The *variable-friction* palm is based on the mechanism in [128] to actively control the friction coefficient of the palm’s surface, and a compliant, *variable-preload* palm is based on elastic, compressible foam. The principles of operation of both designs are shown in Fig. 6.9.

Our *variable-friction* palm controls the friction coefficient of the palm’s surface via a high-friction membrane that inflates out of holes in a low-friction rigid frame. At rest, the membrane (Ecoflex 00-50, Smooth-On) is well-below the palm’s surface, so objects can only make contact with the low-friction polystyrene top layer (1/16” th., McMaster-Carr). As the membrane inflates, it pushes through small holes in the top layer, creating a distributed surface with a high friction coefficient. By controlling the pneumatic pressure, we can control the effective friction coefficient of the palm’s surface, as shown in Fig. 6.9a.

To measure the friction coefficient, a cardboard box (60g, 60 mm cube) is placed on the palm, then pulled laterally at 0.1 mm/s via a string connected to an Instron uniaxial testing machine. The maximum force sustained prior to slip is divided by the object’s mass to calculate the friction coefficient. Interestingly, for the *variable-friction* palm in the “low-friction” regime (where the elastomer film does not yet protrude from the top layer), slight increases in pressure result in a bowed palm surface. This reduces the contact area between the object and palm, resulting in a lower measured coefficient of friction. We note that the specific values of the friction coefficient will vary based upon the object and palm, but the results presented in Fig. 6.9a are representative of a typical object. For the *variable-friction* palm, the friction coefficient changes by roughly $2\times$, between 0.55 ± 0.03 and 1.2 ± 0.06 .

Our second design, the *variable-preload* palm, can modulate the normal force on objects after they are grasped via inflation or deflation of a foam-filled pouch. The operating

principle is similar to existing vacuum-driven, foam-based soft actuators [129, 130]. A cylinder of highly-compliant open cell foam is sealed inside a pouch made of a low-friction thermoplastic elastomer film (Stretchlon 200, Airtech Intl.). At its “zero-preload” state, the pouch is placed under a small negative air pressure (-3 kPa) to partially compress the foam. During a grasp, the fingers hold the object in place against the palm, and the palm’s air pressure can then be modulated up or down to relax (expand) or further compress the foam, as shown in Fig. 6.9b. This small change in the position of the palm’s surface controls the palm’s preload on the object by increasing or decreasing the normal force between the palm and object. To measure the preload, the palm was first retracted by applying -4.1 kPa, then the actuation pressure was raised in 0.7 kPa increments and the resulting force was measured by the Instron machine. Across this range of pressures, the (blocked) preload force varies from 1.2 ± 0.03 N and 6.0 ± 0.2 N.

6.3.4 Results

To evaluate the performance of the two palm designs during grasping and in-hand manipulation, we performed a series of simple tests and measurements using our existing soft, dexterous hand platform [65]. This hand consists of four radially-oriented fingers, each with two orthogonal degrees of freedom. In prior work, this hand was capable of only planar in-hand manipulation, but could utilize a telescoping palm mechanism to control the object’s position in the axis normal to the palm [66]. However, this hand was still incapable of rotating an object about its lateral axes, and its grasp stability was relatively low for most objects.

In this study, we replace the original palm with our new, friction-controlled palms. Using the air pressure controller setup from [65], the fingers and palm are commanded using air pressure trajectories to control grasping and lateral motions, and palm actuation states. In all experiments, the hand is operated open-loop with no external vision system or on-board sensing in the fingers or palm.

6.3.4.1 In-Hand Manipulation

To evaluate the effectiveness of the two palm mechanisms to control object motion, the hand was commanded to grasp an object, set the palm’s actuation state, then shift the grasp location while the object’s resulting motion was observed, as shown in

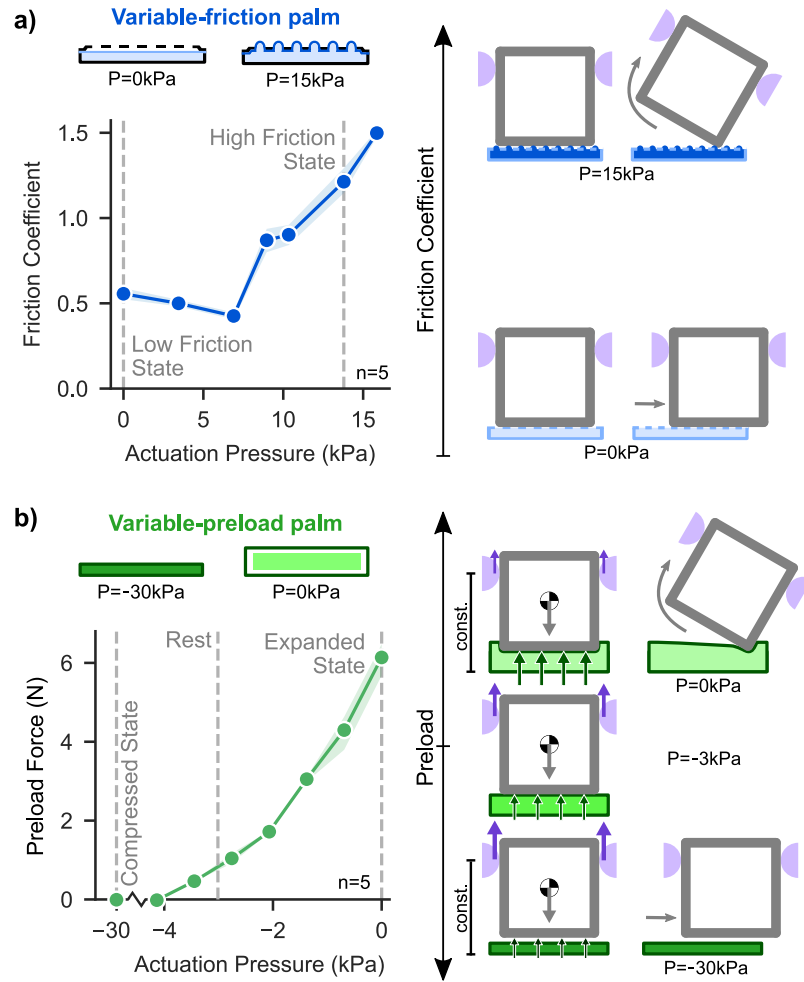


Figure 6.9: The two pneumatically-driven palm designs can be used to control friction forces on the object. a) The *variable-friction* palm uses a positive air pressure signal to control the friction coefficient, μ_p between the palm and object. b) The *variable-preload* palm uses a negative air pressure signal after grasping to control the normal force on the object (applying a preload, δN), which changes the load distribution between the palm and the fingers. Controlling the tip/slip condition for the largest range of objects is achieved by switching between two extreme states for both palm mechanisms.

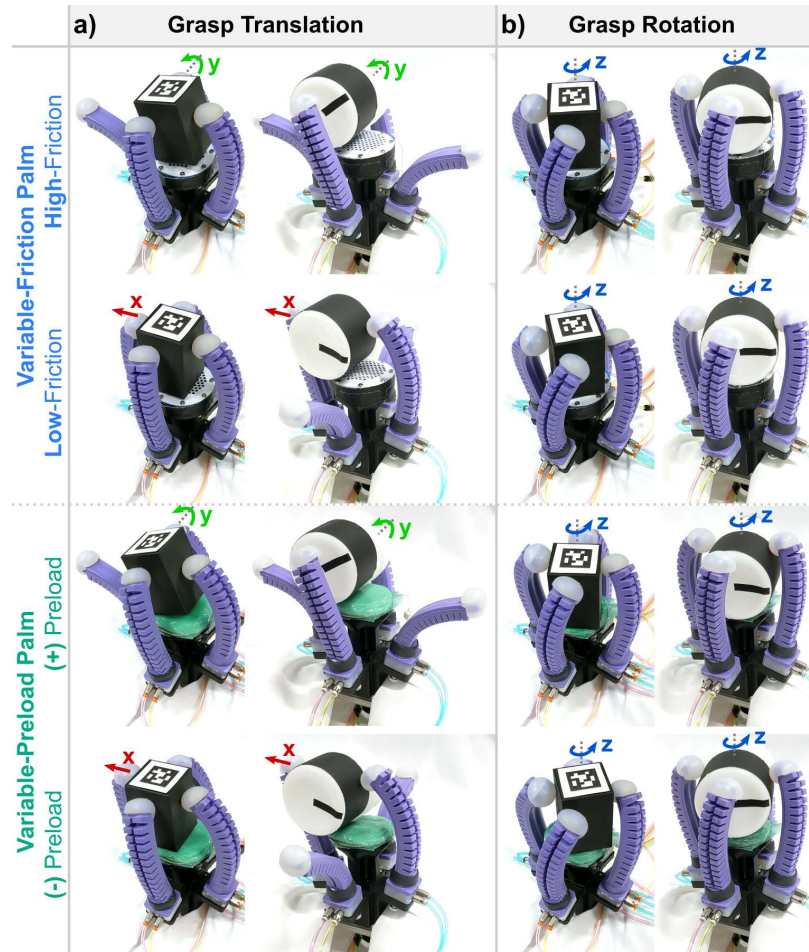


Figure 6.10: The actuation of our two palm mechanisms allows the hand to control whether objects slide or tip when finger forces are applied, and thus the direction of object motion. a) As the grasp is translated in the x -direction, both palm designs control whether the object pivots or rolls about the y -axis, or slides in the x -direction. Due to the hand's symmetry, the same capabilities exist for grasp translations in the y -direction. b) Grasp rotations about the z -axis result in object rotations, but the resulting angle depends on the palm's actuation state (high-friction and + preload states result in lower range of motion). Results are collected via visual inspection.

Fig. 6.10. Objects of rectangular and circular cross-section were tested: a 3D-printed box (76 g, 50 mm×50 mm×75 mm), and a plastic jar (43 g, 100 mm×100 mm). After grasping the object, the grasp is translated in the x -axis, and rotated about the z -axis, and the resulting object motion is collected by inspection.

Observing the motion of the object undergoing the same finger actuations, both palm mechanisms can directly control the sliding vs. tipping behavior of objects in the hand. For the *variable-friction* palm, the low-friction state produces lateral translation of the object when the grasp is translated laterally, while the high-friction state results in pivoting (or rolling) about the axis orthogonal to finger motion. For the *variable-preload* palm, a negative preload (where the palm retracts to reduce the normal force on the object) produces lateral object translations when the grasp is moved laterally, and a positive preload results in pivoting. Grasp rotation about the z -axis results in z -axis object rotation for both actuation states of both palms.

Looking at the behavior as a whole, we see that modifying the palm’s frictional characteristics with the object (either through the friction coefficient, or the magnitude of normal forces) enables switching between sliding or pivoting/rolling. If the palm is implemented as a single, passive design, the hand only has access to one of these manipulation modes (chosen at design-time). Thus, an actuated palm design, such as the two we present in this section, provides run-time control of in-hand manipulation modes, expanding the dexterity of the hand via only one degree of freedom, as shown in Fig. 6.10.

An added advantage of the *variable-preload* palm is that the compliance of the foam reduces the tipping force required by the fingers to rotate the object out-of-plane. This is especially important when trying to manipulate heavier objects with low finger forces, or when trying to keep forces low during manipulation of delicate objects. We demonstrate this by trying to tip a heavy object using the *variable-preload* palm. When the palm has a positive pre-load, the fingers can easily tip the object on the compliantly conforming palm. When the palm has a negative pre-load, our soft fingers (which can apply a maximum force of 2N) are not strong enough to tip the same object, as demonstrated in the Supplementary Video. This demonstrates the importance of compliance in a palm for in-hand-manipulation tasks that require gentle dexterity.

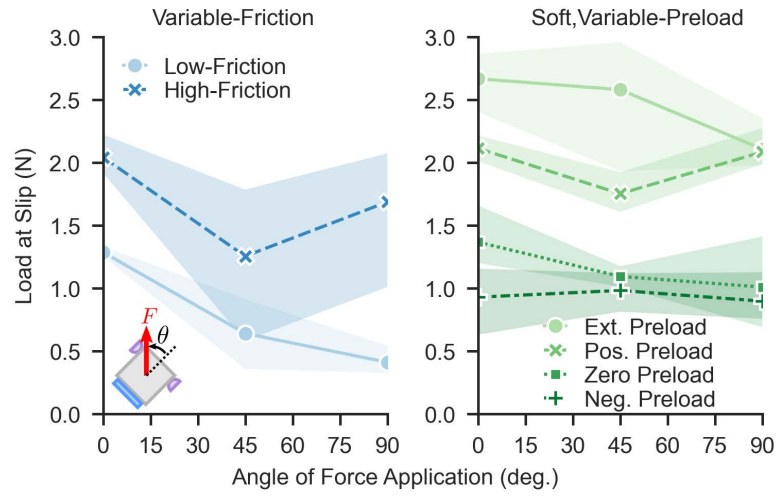


Figure 6.11: The “initial grasp stability” is displayed for each actuation mode of each palm design (mean ± 1 standard deviation of $n = 3$ trials). We can see that for the *variable-friction* palm, the high-friction state has consistently higher grasp stability ($\sim 2\times$) compared to the low-friction state. Additionally, for the *variable-preload* palm, the grasp stability increases as a function of the preload, and an external preload on the object results in a ($\sim 2.5\times$) increase in stability compared to resting on the surface.

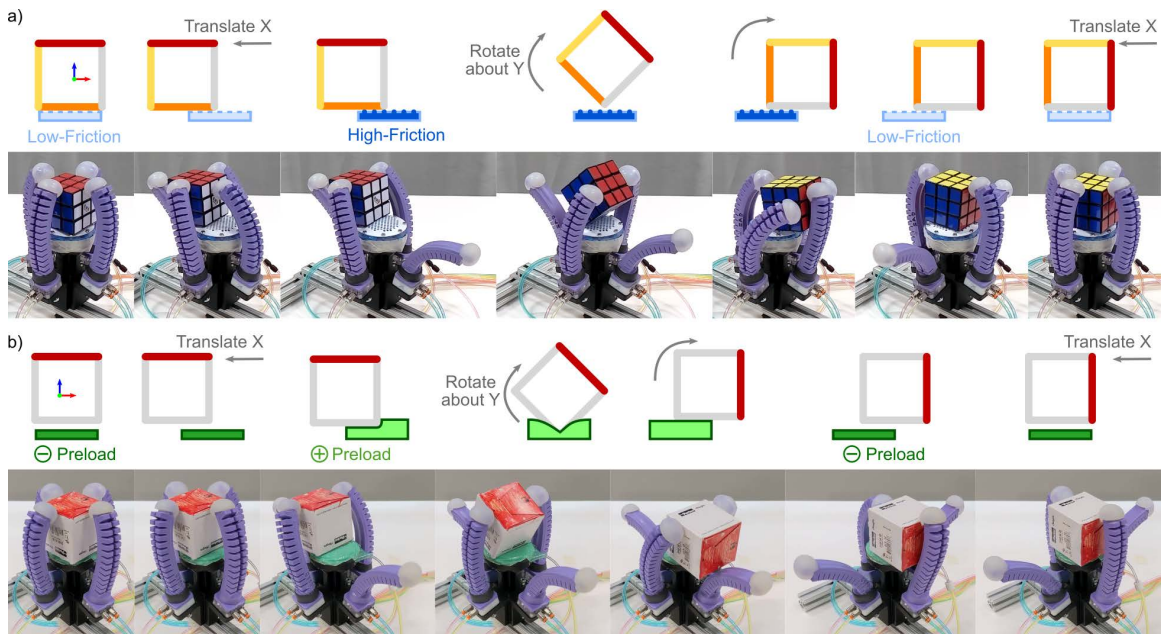


Figure 6.12: Controlling the palm’s surface enables the soft hand to perform arbitrary pose shifts on an object through combinations of sliding and tipping, using either the variable-friction palm (a) or the variable-preload palm (b).

6.3.5 Grasp Stability

The effect of the palm on grasp stability is most directly seen in the initial grasp stability, when the object is secured in a grasp. The initial grasp stability is related to the external load (applied to the object) when the object first slips from its initial pose. To get a simple estimate of how the palm contributes the initial grasp stability, we measured the initial slipping force as a function of the angle at which force is applied. These tests were conducted per the procedure detailed in [65]. The object (a 60 mm×60 mm×60 mm cardboard box, as first seen in [65]) is mounted in the jaws of an Instron universal testing machine with a 2 kN load cell, and hand is mounted to a fixture where its orientation relative to the object is controlled. We then applied a constant extension ramp of 1 mm/s and measured the resulting force.

Fig. 6.11 shows the grasp stability for both palm designs as a function of the actuation mode and force application angle. We can see that for the *variable-friction* palm, the high-friction state yields consistently higher grasp stability ($\sim 2\times$ at best) compared to the low-friction state. For the *variable-preload* palm, the preload on the object has a similar effect on stability, where a negative preload leads to lower stability compared to zero preload, and positive preload leads to approximately $2\times$ higher grasp stability compared to zero preload. In addition, when the object is pressed into the palm before grasping by an external force, the hand achieves a further increase in grasp stability ($\sim 2.5\times$) compared to zero preload. Looking at the results overall, we see that both palms can modulate their grasp stability as a function of actuation, but the compliant construction of the *variable-preload* palm also allows it to make use of external preload forces if they are available.

6.3.6 Manipulation Tasks from Sequential Motion Primitives

By chaining together in-hand manipulation motion primitives discussed previously, objects can be translated and rotated to any desired orientation. Fig. 6.12 illustrates the sequence of palm and finger actuation required to perform a full 90° out-of-plane rotation of a box. Briefly, this requires the object to be slid so that the edge of rotation is shifted close to the center of the finger workspace, followed by a tipping motion. We leverage the palm’s abilities to change friction or preload against the fingers to trigger whether slipping or tipping occurs with equivalent finger actuation. The objects used in this demonstration are a Rubik’s Cube, a cardboard box (60 g, 60 mm×60 mm ×60 mm), and a cylindrical

plastic jar (43 g, 100 mm×100 mm). Videos of these demonstrations can be found in the Supplementary Video.

To further demonstrate the utility of the grasping and in-hand manipulation capabilities afforded by control of the palm’s friction, we show a food preparation task, as shown in Fig. 6.13. The goal of the task is to present a sandwich bun with the sliced side out such that a worker can fill it with sandwich toppings, then place it on a plate. The successful task is shown in the Supplementary Video. To accomplish this task, a UR5e robot arm (Universal Robots) is outfitted with our soft hand and the *variable-friction* palm. The arm grasps the bun, engages its high-friction state for maximal grasp stability, then flips the hand upside down, allowing the bun to rest on the palm’s surface. Next, the hand utilizes the sequence of primitive motions detailed above to slide and pivot the bun 90° placing a the pre-sliced side of the bread upward. The bun is loaded with toppings by a worker, then pivoted 90° to place the top of the bun facing the palm. Finally, the completed sandwich is placed on a plate, right side up.

6.3.7 Discussion

Our results demonstrate that active control of the friction force between the palm of a robotic hand and the object it interacts with has a direct effect on in-hand manipulation modes and grasp stability. In this section, we discuss how these results can inform future hand designs, what additional factors (beyond isolated palm design) should be considered, and how these results could be generalized.

Active control of the friction coefficient or preload between the palm and object can play a large role in determining the direction of in-hand object motion (sliding vs. pivoting) when finger forces are applied. This is accomplished by controlling which constraints on object motion the palm provides through contact with the object. This is especially relevant for robotic hands with limited finger dexterity since, as demonstrated here, a single actuated DOF in the palm can be used instead of added dexterity in each finger.

Controlling the frictional characteristics of the palm also allows for control of grasp stability, but the upper bound is limited for both palm designs. For controlling the friction coefficient, the grasp stability improvements are limited by the available material choice (elastomers tend to have the highest coefficients already). Through control of the preload, the gains in grasp stability are limited by the strength (and friction coefficient) of the fingers,



Figure 6.13: The actuated palm enables a real-world grasping and manipulation task: Helping a worker make a submarine sandwich.

since the palm presses the object against the fingers when applying positive preload. Thus, the range of acceptable preload forces must be matched to the finger strengths.

In addition to the two main design parameters, we showed that building a palm with a compliant surface has several real, measurable benefits for grasping and in-hand manipulation. A compliant palm can substantially decrease the finger forces required to tip heavier objects, as analyzed and demonstrated in this study. We also showed that palm compliance can be used to take advantage of external preloads (i.e., via the environment) before grasping an object with the fingers to increase grasp stability. Both of these benefits are especially useful for soft robotic hands whose fingers are often weak compared to the weight of many common household objects.

In addition to the palm's design, a large factor in the design process should be the target set of objects. For example, for a desired set of objects, the palm's materials and actuation modes must be chosen such that the range of friction coefficients and/or preloads between the palm and the object set are different enough to enable the desired contact mode switching behavior. Furthermore, while we used rectangular and circular objects in this study to demonstrate the palm's impact on in-hand manipulation, we expect our analysis and results to remain pertinent even for objects with more interesting geometries. However, the precise motion of objects in the hand will be affected by the object's center-of-mass and position of the fingers on the object.

Finally, further exploring the effect of palm-object interactions on in-hand manipulation in a more quantitative manner could yield insights into the complexity of these interactions. For example, measuring the manipulability of the object in the hand as a function of the palm's applied constraints could indicate which physical mechanisms dominate. This could also shed more light into the coupling of the hand's performance to its pose and directions of external forces or constraints. Future palm designs could then be developed based on a more bottom-up approach.

6.3.8 Conclusions

In this paper, we demonstrated how control of the friction force between an object and the palm of a soft robotic hand dictates slipping and tipping behaviors, and thus in-plane and out-of-plane motions, during manipulation. Through a theoretical analysis, we showed that tip vs. slip behavior can be determined by the friction coefficient and the

preload of the object on the palm, given the same object and finger force. Furthermore, we show in this analysis that a compliant palm lowers the force required to tip an object, a note of particular interest for soft robotic hands that are often underpowered. With these design concepts in mind, we fabricated two pneumatically-driven palms, a *variable-friction* and a *variable-preload* palm. We experimentally verified that by actuating the palms, the same finger motion can result in either sliding or tipping, generating both in-plane and out-of-plane motion primitives. Additionally, we showed that the grasp stability can also be varied via the actuation state of the palms. Finally, through a sequence of grasping and in-hand manipulation motion primitives, we demonstrated that controlling the palm-object interaction enables a robot to stably grasp an object, then re-orient it in 5-DOF.

The area of palm design for robotic hands has great potential for future innovations. Most immediately, we are interested in a more in-depth study of the palm design concepts presented in this paper as they pertain to real application tasks. For example, choosing the range of friction coefficients or preloads could enable control of in-hand manipulation of a wider range of objects or tasks. Continuous control of the actuation space (rather than binary) and spatially modulating the friction forces could enable a finer-grained approach to controlling contact interactions. We are also interested in further exploring how both compliance and plasticity of the palm can benefit grasping and in-hand manipulation. Looking further into the future, this work could be expanded upon to achieve greater control of objects within the hand. Combining the two palm designs presented in this paper with a variable height mechanism would enable the hand to control finger placement on the object, potentially enabling 6-DOF control of the object's pose. With careful engineering, active palm surfaces could be a gateway to highly-dexterous in-hand manipulation with relatively simple hands.

7

Conclusions and Future Work

In this chapter, I discuss how desired applications can affect the design requirements for soft robotic grippers and hands, as well as some potential future explorations based on the work presented in this dissertation.

7.1 Design Principles for Soft Robotic Hands

The main driver of robotic system design is usually the specific applications in which robots are to be deployed. Even within common aspirational applications of soft robots, such as agriculture, food handling, or assistance with household tasks, functional requirements for end effectors can vary dramatically. In this section, I discuss how some common types of applications drive the design of soft robotic grippers and hands in terms of the necessary level of dexterity, the arrangement of DOFs, and actuation requirements, as well as limitations of various design approaches.

7.1.1 Grasping

The goal of all robotic grippers is to secure objects such that they can be moved by a manipulation system (arm, gantry, etc.). This necessarily requires grippers to restrain objects against forces and torques applied externally to the robot such as those from inertia, gravity, or environmental constraints. Grasping objects can often be accomplished with relatively low dexterity. However, the applications presented here demand that a gripper's overall success depends not only on successful control of the object, but also on minimizing damage, and maximizing success under large uncertainty and larger variability. Thus, narrowing down the specific item set that a soft gripper is designed to interact with is a

critical first step.

For picking large and/or heavy items, a power grasp is usually the most effective grasp [7] to use, which can be accomplished using very simple soft grippers. As demonstrated by a host of prior work, power grasps can be easily achieved with compliant grippers using simple, 1-DOF mechanisms [21]. If objects are delicate, an extremely gentle touch can be achieved through minimizing contact pressure on the object via highly-compliant materials, high numbers of contact points, or both [35, 105, 131]. In Sinatra et al. 2019, we showed that with appropriate tuning of actuator stiffness, a soft gripper could perform a secure power grasp on a jellyfish, all without damaging the organism [105]. Successful grippers can also be designed that forgo the traditional digit-centric design paradigm in favor of a more stochastic grasping approach with a large number of distributed contact interactions [131].

If the desired item set includes handling thinner or smaller objects, precision grasps must be used to grasp these object, especially when the objects rest on surfaces. I showed in Chapter 2 that soft robotic fingers can still achieve pinch grasps using a 1-DOF mechanism, but the finger design must include a passive distal segment for greater grasp stability [64]. I also illuminated a fundamental design tradeoff between pinch grasping and power grasping capabilities, which indicated that a 2-DOF gripper could achieve the best performance for both grasp types. In addition, I showed in Chapter 3 that for thin, flexible objects, 3D finger compliance enables robust precision grasping and protects objects from damage [70].

7.1.2 In-Hand Manipulation

For in-hand manipulation, the goal is usually to maintain control of the object's motion relative to the hand. In this work, we treat in-hand manipulation as a series of quasi-static grasp transitions which directly lead to object motion. As such, different real-world tasks require different object motion primitives for success. Taking the desired primitive motions into account, one can choose how to distribute degrees of freedom amongst fingers and other hand features.

Many tasks in the real world, such as twisting dials and knobs, performing fine position adjustments to objects while packing them, and even writing with a pen, can benefit from in-hand manipulation, but do not require finger-gaiting. For such tasks, a hand with two dexterous fingers (at least 2-DOF each) can accomplish the necessary motion

primitives via quasi-planar motion. With the necessary friction coefficient at the fingertips, simple motion can be performed, but is subject to the combined workspace of the fingers. Furthermore, fingers cannot necessarily break contact with the object, since both fingers are critical for the object to remain grasped. Thus, if the robot needs to move the object over a larger range of motion, the object must be set down and regrasped.

To handle tasks where a large range of in-hand motion is required, such as quickly re-orienting items to be scanned and bagged at a grocery store, or stocking items on shelves with labels facing out, finger gaits can be used to extend the range of achievable object motion beyond the limits of the fingers. When fingers are free to make and break contact with the object, at least three digits are required to ensure two points of contact are maintained at all times [50, 111]. In many hands, more than three fingers are used in an effort to provide more ways to design finger gaits [53, 107, 49]. In this dissertation, we determined that four fingers are enough to compensate for the limited lateral workspace of individual fingers, while reducing complexity as much as possible.

After determining the number of fingers, the next important factor is the required dexterity of each finger. The arrangement of degrees of freedom in each finger plays a large role in which directions the hand can manipulate objects, especially if the fingers contain the only actuated degrees of freedom in the hand. To enable planar motions perpendicular to the palm, fingers with two parallel degrees of freedom (in the grasping and side-to-side directions) can achieve robust results [65]. However, if the task requires objects to move orthogonal to the palm, then fingers need at least two serially-linked degrees of freedom in the sagittal plane. Finally, if planar motion, vertical motion, and pivoting are required, then each finger requires 3 DOF (two in the sagittal plane, and one off-axis). With at least three fingers, each with three DOF, the hand could achieve 6-DOF control over object motion, but with nine or more total degrees of freedom.

From here, we can reduce the necessary number of actuated degrees of freedom in the fingers by instead enabling actuation of the finger arrangement and palm surface, as discussed in Chapters 5 and 6. For tasks where only planar object motion is necessary, such as performing fine adjustments during packing, using dials and knobs, or simple regrasping, only a passive palm is necessary, and 1-DOF control of the finger arrangement yields high-quality manipulation of a variety of objects. For more complex tasks, such as re-orienting items, 1-DOF control of the palm's frictional properties enables pivoting and sliding of objects with fingers which only have one actuated degree of freedom in the sagittal plane

rather than two (a reduction of 1 DOF per finger).

Overall, the design of compliant or fully soft hands must be fundamentally based on some desired set of tasks. As discussed throughout this dissertation, the necessary arrangement of degrees of freedom within the fingers, finger arrangement, and palm surface are highly dependent on the set of tasks for which the hand is designed. In addition, the ways in which compliance can be exploited are also affected greatly by conditions of a chosen application. For in-hand manipulation, soft hands are well-suited for handling uncertainty in initial conditions of objects, but will inherently lack the necessary precision for fine manipulation such as writing with a pen, or assembling structures with tweezers. Thus, soft hands should be used in applications where safety and adaptability are more important than precision or repeatably. While one could embark on an effort to make a “general-purpose” soft hand capable of a wide variety of tasks, a well-defined description of “how general” is always necessary.

7.2 Further Exploration into Onboard Sensing

The role of sensing in robotic grasping and in-hand manipulation has been debated in the literature and PhD dissertations for decades, but a clear answer has still yet to emerge. Humans make extensive use of rich tactile sensing and proprioception during both grasping and dexterous manipulation, yet this dissertation demonstrates that robust grasping and dexterous manipulation can be achieved even without these senses. While it may be easy to consider this proof that on-board sensing in the hand itself is unnecessary for soft robots, I believe the exact opposite has actually been demonstrated.

The precision required to perform repeatable in-hand manipulation with limited *a-priori* knowledge about the object is difficult to achieve without sensing. In most of the demonstrations shown in this dissertation, robot and hand motions are designed by human operators for success on a limited set of objects. However, these hand-tuned trajectories remain sensitive to perturbations in the object’s size, shape, mass, surface texture, friction coefficient, and pose.

Vision-based sensing can provide shape and pose information about objects, but cannot provide explicit estimates of other attributes. For example, deep learning has been successfully applied to object pose estimation, provided the system has an offline training period with rich camera views of each new object [132]. For objects that have not been cat-

aloped, the size and shape can be estimated reasonably well using simultaneous localization and mapping (SLAM) [133] or data from RGB-D cameras [134, 135]. However, estimating the mass, surface texture, or frictional properties requires the robot to interact with the object [136].

On-board sensing of finger shape, and/or contact forces can provide important realtime feedback that could be used for estimation of grasp quality [137], object pose [138, 139, 140], and physical properties of the object and environment [141, 142]. Appendix A details some early work toward sensor development, where I integrated soft curvature and contact sensors into existing soft fingers for use in grasping delicate organisms in the harsh conditions of the deep sea [143]. While the sensors developed in this work are insufficient for measuring minute changes in contact and shape during in-hand manipulation, there has been a huge amount of recent development toward high-quality, low-cost soft sensors for robotics applications [62, 144, 145, 146, 147].

With reliable, high-quality sensors at-the-ready, their placement on fingers and hands is the next important question. Wall et al. developed a technique to empirically determine how sensor placement leads to information gain, which enables identification of which sensors contribute most to shape estimates of soft fingers [148, 149]. In addition, learning-based approaches have been extremely successful in estimating the shape, contact location, and contact forces on soft fingers from relatively simple sensing modalities [150].

Overall, while the work in this dissertation demonstrates the capability of soft robotic hands to grasp and maneuver objects robustly in the presence of uncertainty, the addition of on-board sensing and integration of this information into feedback control is still necessary for truly robust performance. Furthermore, on-board sensing is critical for determining grasp quality, and especially important for making on-the-fly decisions about grasp transitions. A wide variety of soft sensing modalities are now reaching maturity for use in soft robots, so understanding how to best sense the objects and environment that a soft hand encounters is a critical direction for future exploration.

7.3 Further Exploration into Learning for Soft Manipulation

A huge amount of work is currently focusing on designing, modeling, and controlling highly-dexterous robots, yet robots still interact with the real, unstructured world using the dexterity of a toddler. Data-driven methods such as deep learning and reinforcement

learning have become more and more prevalent in the robotics world, applied to difficult, high-dimensional problems with high success. In most cases, machine learning models are trained in simulated worlds, which are usually far from reality. This leads to the fundamental challenge of the sim2real problem: transferring learned behavior from simulation to physical hardware in a robust, intelligent way.

Soft robots present a unique way to deal with discrepancies between the simulated environment and real life: the mechanical compliance of soft robots, usually designed for robust interactions with the world with minimal sensing, may also provide a level of robustness during sim2real transfer. The potential advantage is that learning could occur in low-fidelity simulations (saving time and compute resources) while still resulting in reasonable behavior on the real system.

In a first attempt to demonstrate the robustness of sim2real transfer for soft robot systems, we showed that a control policy for in-hand manipulation developed via reinforcement learning in our SoMo simulation environment could be transferred directly to hardware with no hand-tuning, resulting in very similar object behavior [151]. This policy was trained in a simulated environment built using the SoMo simulation framework [68], where soft fingers are approximated using a series of rigid links and spring-loaded joints. While SoMo does not capture material compliance, the framework approximates compliant structures with reasonable accuracy, and can be easily calibrated to simulate physical hardware. Even with these simplifications, we were able to successfully train a soft hand to perform a desired object motion (rotate a cube) in SoMoGym, then perform the task in real hardware by converting the actuation signals using a simple calibration function. These results are a promising indicator that more-complex tasks and behaviors can be learned in simulation, then transferred to hardware with relative ease.

7.4 Conclusions

The goal of my work has always been to develop some intuition about how to design soft hands that have “just enough” dexterity as necessary for a given task, while minimizing complexity. In this dissertation, I showed how this simple design philosophy leads to hand designs that, while tied to specific applications, can achieve precise grasping and robust, dexterous manipulation that meets the needs of their applications. Looking into the future, there are myriad opportunities to improve the capabilities of soft grippers

and hands through sensing and control. I believe that applying this same mentality (determine the necessary system complexity required for a task, and choose the least-complex implementation strategy) is key to success moving forward. The world deserves safe, gentle, dexterous robots, so we must strive to give robots what they deserve: end effectors with the precision and dexterity to function in the real world, but with the inherent safety to work by our sides.

Bibliography

- [1] Soft Robotics Inc., “5 reasons to automate your bakery operations,” *White Paper*, 2019.
- [2] P. Maciejasz, J. Eschweiler, K. Gerlach-Hahn, A. Jansen-Troy, and S. Leonhardt, “A survey on robotic devices for upper limb rehabilitation,” *Journal of neuroengineering and rehabilitation*, vol. 11, no. 1, pp. 1–29, 2014.
- [3] S. Robla-Gómez, V. M. Becerra, J. R. Llata, E. Gonzalez-Sarabia, C. Torre-Ferrero, and J. Perez-Oria, “Working together: A review on safe human-robot collaboration in industrial environments,” *IEEE Access*, vol. 5, pp. 26754–26773, 2017.
- [4] D. Rus and M. T. Tolley, “Design, fabrication and control of soft robots,” *Nature*, vol. 521, no. 7553, p. 467, 2015.
- [5] M. T. Mason, “Toward robotic manipulation,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 1–28, 2018.
- [6] I. M. Bullock, J. Z. Zheng, S. De La Rosa, C. Guertler, and A. M. Dollar, “Grasp frequency and usage in daily household and machine shop tasks,” *IEEE transactions on haptics*, vol. 6, no. 3, pp. 296–308, 2013.
- [7] M. R. Cutkosky, “On grasp choice, grasp models, and the design of hands for manufacturing tasks,” *IEEE Transactions on robotics and automation*, vol. 5, no. 3, pp. 269–279, 1989.
- [8] J. Amend and H. Lipson, “The JamHand: Dexterous Manipulation with Minimal Actuation,” *Soft Robotics*, vol. 4, no. 1, p. soro.2016.0037, 2017.

- [9] S. Jacobsen, E. Iversen, D. Knutti, R. Johnson, and K. Biggers, “Design of the utah/mit dextrous hand,” in *Proceedings. 1986 IEEE International Conference on Robotics and Automation*, vol. 3, pp. 1520–1532, IEEE, 1986.
- [10] A. Kochan, “Shadow delivers first hand,” *Industrial robot: an international journal*, vol. 32, no. 1, pp. 15–16, 2005.
- [11] L. B. Bridgwater, C. Ihrke, M. A. Diftler, M. E. Abdallah, N. A. Radford, J. Rogers, S. Yayathi, R. S. Askew, and D. M. Linn, “The Robonaut 2 hand-designed to do work with tools,” in *2012 IEEE International Conference on Robotics and Automation*, pp. 3425–3430, IEEE, 2012.
- [12] J.-H. Bae, S.-W. Park, J.-H. Park, M.-H. Baeg, D. Kim, and S.-R. Oh, “Development of a low cost anthropomorphic robot hand with high capability,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 4776–4782, IEEE, 2012.
- [13] J. Mahler, F. T. Pokorny, B. Hou, M. Roderick, M. Laskey, M. Aubry, K. Kohlhoff, T. Kröger, J. Kuffner, and K. Goldberg, “Dex-net 1.0: A cloud-based network of 3d objects for robust grasp planning using a multi-armed bandit model with correlated rewards,” in *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1957–1964, IEEE, 2016.
- [14] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, “Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics,” *arXiv preprint arXiv:1703.09312*, 2017.
- [15] J. Mahler, M. Matl, X. Liu, A. Li, D. Gealy, and K. Goldberg, “Dex-net 3.0: Computing robust vacuum suction grasp targets in point clouds using a new analytic model and deep learning,” in *2018 IEEE International Conference on robotics and automation (ICRA)*, pp. 5620–5627, IEEE, 2018.
- [16] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, “Learning ambidextrous robot grasping policies,” *Science Robotics*, vol. 4, no. 26, p. eaau4984, 2019.

- [17] W. Townsend, “The Barrett Hand grasper – programmably flexible part handling and assembly,” *Industrial Robot: an international journal*, vol. 27, no. 3, pp. 181–188, 2000.
- [18] Robotiq, “2f-85 and 2f-140 grippers,” *Online datasheet*, 2019.
- [19] M. Ciocarlie, F. M. Hicks, R. Holmberg, J. Hawke, M. Schlicht, J. Gee, S. Stanford, and R. Bahadur, “The Velo gripper: A versatile single-actuator design for enveloping, parallel and fingertip grasps,” *International Journal of Robotics Research*, vol. 33, no. 5, pp. 753–767, 2014.
- [20] M. G. Catalano, G. Grioli, E. Farnioli, A. Serio, C. Piazza, and A. Bicchi, “Adaptive synergies for the design and control of the pisa/iit soft hand,” *The International Journal of Robotics Research*, vol. 33, no. 5, pp. 768–782, 2014.
- [21] A. M. Dollar and R. D. Howe, “The highly adaptive SDM hand: Design and performance evaluation,” *International Journal of Robotics Research*, vol. 29, no. 5, pp. 585–597, 2010.
- [22] L. U. Odhner, L. P. Jentoft, M. R. Claffee, N. Corson, Y. Tenzer, R. R. Ma, M. Buehler, R. Kohout, R. D. Howe, and A. M. Dollar, “A compliant, underactuated hand for robust manipulation,” *The International Journal of Robotics Research*, vol. 33, no. 5, pp. 736–752, 2014.
- [23] D. M. Aukes, B. Heyneman, J. Ulmen, H. Stuart, M. R. Cutkosky, S. Kim, P. Garcia, and A. Edsinger, “Design and testing of a selectively compliant underactuated hand,” *The International Journal of Robotics Research*, vol. 33, no. 5, pp. 721–735, 2014.
- [24] J. Butterfaß, M. Grebenstein, H. Liu, and G. Hirzinger, “Dlr-hand ii: Next generation of a dextrous robot hand,” in *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No. 01CH37164)*, vol. 1, pp. 109–114, IEEE, 2001.
- [25] W. Friedl, H. Höppner, F. Schmidt, M. A. Roa, and M. Grebenstein, “Clash: Compliant low cost antagonistic servo hands,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6469–6476, IEEE, 2018.

- [26] R. Maruyama, T. Watanabe, and M. Uchida, “Delicate grasping by robotic gripper with incompressible fluid-based deformable fingertips,” *IEEE International Conference on Intelligent Robots and Systems*, pp. 5469–5474, 2013.
- [27] B. W. McInroe, C. L. Chen, K. Y. Goldberg, R. Bajcsy, and R. S. Fearing, “Towards a soft fingertip with integrated sensing and actuation,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6437–6444, Oct 2018.
- [28] M. R. Cutkosky and P. K. Wright, “Friction, Stability and the Design of Robotic Fingers,” *The International Journal of Robotics Research*, vol. 5, no. 4, pp. 20–37, 1986.
- [29] J. Hughes, U. Culha, F. Giardina, F. Guenther, A. Rosendo, and F. Iida, “Soft manipulators and grippers: a review,” *Frontiers in Robotics and AI*, vol. 3, p. 69, 2016.
- [30] C. Majidi, “Soft robotics: a perspective—current trends and prospects for the future,” *Soft Robotics*, vol. 1, no. 1, pp. 5–11, 2014.
- [31] P. Polygerinos, N. Correll, S. A. Morin, B. Mosadegh, C. D. Onal, K. Petersen, M. Cianchetti, M. T. Tolley, and R. F. Shepherd, “Soft robotics: Review of fluid-driven intrinsically soft devices; manufacturing, sensing, control, and applications in human-robot interaction,” *Advanced Engineering Materials*, vol. 19, no. 12, p. 1700016, 2017.
- [32] J. Shintake, V. Cacucciolo, D. Floreano, and H. Shea, “Soft robotic grippers,” *Advanced Materials*, p. 1707035, 2018.
- [33] F. Ilievski, A. D. Mazzeo, R. F. Shepherd, X. Chen, and G. M. Whitesides, “Soft robotics for chemists,” *Angewandte Chemie*, vol. 123, no. 8, pp. 1930–1935, 2011.
- [34] E. Brown, N. Rodenberg, J. Amend, A. Mozeika, E. Steltz, M. R. Zakin, H. Lipson, and H. M. Jaeger, “Universal robotic gripper based on the jamming of granular material,” *Proceedings of the National Academy of Sciences*, vol. 107, no. 44, pp. 18809–18814, 2010.
- [35] K. C. Galloway, K. P. Becker, B. Phillips, J. Kirby, S. Licht, D. Tchernov, R. J. Wood, and D. F. Gruber, “Soft Robotic Grippers for Biological Sampling on Deep Reefs,” *Soft Robotics*, vol. 3, no. 1, p. soro.2015.0019, 2016.

- [36] R. Deimel and O. Brock, “A novel type of compliant and underactuated robotic hand for dexterous grasping,” *The International Journal of Robotics Research*, vol. 35, no. 1-3, pp. 161–185, 2016.
- [37] T. Feix, R. Pawlik, H.-B. Schmedmayer, J. Romero, and D. Kragic, “A comprehensive grasp taxonomy,” in *Robotics, science and systems: workshop on understanding the human hand for advancing robotic manipulation*, vol. 2, pp. 2–3, 2009.
- [38] G. Gu, N. Zhang, H. Xu, S. Lin, Y. Yu, G. Chai, L. Ge, H. Yang, Q. Shao, X. Sheng, *et al.*, “A soft neuroprosthetic hand providing simultaneous myoelectric control and tactile feedback,” *Nature Biomedical Engineering*, pp. 1–10, 2021.
- [39] J. Zhou, S. Chen, and Z. Wang, “A soft-robotic gripper with enhanced object adaptation and grasping reliability,” *IEEE Robotics and Automation Letters*, vol. 2, pp. 2287–2293, Oct 2017.
- [40] K. W. O’Brien, P. A. Xu, D. J. Levine, C. A. Aubin, H.-J. Yang, M. F. Xiao, L. W. Wiesner, and R. F. Shepherd, “Elastomeric passive transmission for autonomous force-velocity adaptation applied to 3d-printed prosthetics,” *Science Robotics*, vol. 3, no. 23, p. eaau5543, 2018.
- [41] R. Deimel, P. Irmisch, V. Wall, and O. Brock, “Automated co-design of soft hand morphology and control strategy for grasping,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1213–1218, IEEE, 2017.
- [42] E. Knoop, M. Bächer, V. Wall, R. Deimel, O. Brock, and P. Beardsley, “Handshakiness: Benchmarking for human-robot hand interactions,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4982–4989, IEEE, 2017.
- [43] D. M. Vogt, K. P. Becker, B. T. Phillips, M. A. Graule, R. D. Rotjan, T. M. Shank, E. E. Cordes, R. J. Wood, and D. F. Gruber, “Shipboard design and fabrication of custom 3d-printed soft robotic manipulators for the investigation of delicate deep-sea organisms,” *PloS one*, vol. 13, no. 8, p. e0200386, 2018.
- [44] A. Fernandez, J. P. Gazeau, S. Zegloul, and S. Lahouar, “Regrasping objects during

- manipulation tasks by combining genetic algorithms and finger gaiting,” *Meccanica*, vol. 47, p. 939–950, Apr 2012.
- [45] A. Bicchi, “Hands for dexterous manipulation and robust grasping: a difficult road toward simplicity,” *IEEE Transactions on Robotics and Automation*, vol. 16, p. 652–662, Dec 2000.
- [46] N. C. Daffe, A. Rodriguez, R. Paolini, B. Tang, S. S. Srinivasa, M. Erdmann, M. T. Mason, I. Lundberg, H. Staab, and T. Fuhlbrigge, “Extrinsic dexterity: In-hand manipulation with external forces,” in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, p. 1578–1585, IEEE, May 2014.
- [47] N. Furukawa, A. Namiki, S. Taku, and M. Ishikawa, “Dynamic regrasping using a high-speed multifingered hand and a high-speed vision system,” in *2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, p. 181–187, IEEE, 2006.
- [48] B. Sundaralingam and T. Hermans, “Geometric in-hand regrasp planning: Alternating optimization of finger gaits and in-grasp manipulation,” in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 231–238, May 2018.
- [49] Y. Fan, W. Gao, W. Chen, and M. Tomizuka, “Real-time finger gaits planning for dexterous manipulation,” *IFAC-PapersOnLine*, vol. 50, p. 12765–12772, Jul 2017.
- [50] L. Han and J. Trinkle, “Dextrous manipulation by rolling and finger gaiting,” in *Proceedings. 1998 IEEE International Conference on Robotics and Automation (Cat. No.98CH36146)*, vol. 1, p. 730–735, IEEE, 1998.
- [51] M. Huber and R. Grupen, “Robust finger gaits from closed-loop controllers,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2, pp. 1578–1584 vol.2, 2002.
- [52] R. R. Ma and A. M. Dollar, “On dexterity and dexterous manipulation,” in *2011 15th International Conference on Advanced Robotics (ICAR)*, pp. 1–7, IEEE, 2011.
- [53] R. R. Ma and A. M. Dollar, “An underactuated hand for efficient finger-gaiting-based dexterous manipulation,” in *2014 IEEE International Conference on Robotics and Biomimetics (ROBIO 2014)*, p. 2214–2219, IEEE, Dec 2014.

- [54] F. R. Hogan, J. Ballester, S. Dong, and A. Rodriguez, “Tactile dexterity: Manipulation primitives with tactile feedback,” in *2020 IEEE international conference on robotics and automation (ICRA)*, pp. 8863–8869, IEEE, 2020.
- [55] J. Xu, T. J. Koo, and Z. Li, “Finger gaits planning for multifingered manipulation,” in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2932–2937, 2007.
- [56] M. Pfanne, M. Chalon, F. Stulp, H. Ritter, and A. Albu-Schäffer, “Object-level impedance control for dexterous in-hand manipulation,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2987–2994, 2020.
- [57] O. M. Andrychowicz, B. Baker, M. Chociej, R. Józefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, and et al., “Learning dexterous in-hand manipulation,” *The International Journal of Robotics Research*, vol. 39, p. 3–20, Jan 2020.
- [58] I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas, *et al.*, “Solving rubik’s cube with a robot hand,” *arXiv preprint arXiv:1910.07113*, 2019.
- [59] R. Sturges, “A quantification of machine dexterity applied to an assembly task,” *The International Journal of Robotics Research*, vol. 9, no. 3, pp. 49–62, 1990.
- [60] J. Zhou, X. Chen, U. Chang, J.-T. Lu, C. C. Y. Leung, Y. Chen, Y. Hu, and Z. Wang, “A soft-robotic approach to anthropomorphic robotic hand dexterity,” *IEEE Access*, vol. 7, p. 101483–101495, 2019.
- [61] A. Bhatt, A. Sieler, S. Puhlmann, and O. Brock, “Surprisingly robust in-hand manipulation: An empirical study-supplementary material,” in *Robotics Science and Systems Conference (RSS)*, 2021.
- [62] H. Zhao, K. O’Brien, S. Li, and R. Shepherd, “Optoelectronically Innervated Soft Prosthetic Hand via Stretchable Optical Waveguides,” *Science Robotics*, vol. 7529, no. December, pp. 1–10, 2016.
- [63] J. Zhou, Y. Chen, D. C. F. Li, Y. Gao, Y. Li, S. S. Cheng, F. Chen, and Y. Liu, “50 benchmarks for anthropomorphic hand function-based dexterity classification and

- kinematics-based hand design,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 9159–9165, 2020.
- [64] C. B. Teeple, T. N. Koutros, M. A. Graule, and R. J. Wood, “Multi-segment soft robotic fingers enable robust precision grasping,” *International Journal of Robotics Research*, 2020.
- [65] S. Abondance, C. B. Teeple, and R. J. Wood, “A dexterous soft robotic hand for delicate in-hand manipulation,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5502–5509, 2020.
- [66] C. B. Teeple, G. R. Kim, M. A. Graule, and R. J. Wood, “An active palm enhances dexterity for soft robotic in-hand manipulation,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2021.
- [67] C. B. Teeple, R. C. St. Louis, M. A. Graule, and R. J. Wood, “The role of digit arrangement in soft robotic in-hand manipulation,” in *IEEE International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2021.
- [68] M. A. Graule, C. B. Teeple, T. P. McCarthy, R. C. St. Louis, G. R. Kim, and R. J. Wood, “Somo: Fast and accurate simulation of continuum robots in complex environments,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [69] C. B. Teeple, B. Aktaş, M. C. Yuen, G. R. Kim, R. D. Howe, and R. J. Wood, “Controlling palm-object interactions via friction for enhanced in-hand manipulation,” in *IEEE Robotics and Automation Letters*, IEEE, 2022 (In Review).
- [70] C. B. Teeple, J. Werfel, and R. J. Wood, “Multi-dimensional compliance of soft grippers enables gentle interaction with thin, flexible objects,” in *2022 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2022 (In Review).
- [71] R. Deimel and O. Brock, “A compliant hand based on a novel pneumatic actuator,” in *2013 IEEE International Conference on Robotics and Automation*, pp. 2047–2053, IEEE, 2013.

- [72] J. Zhou, J. Yi, X. Chen, Z. Liu, and Z. Wang, “Bcl-13: A 13-dof soft robotic hand for dexterous grasping and in-hand manipulation,” *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3379–3386, 2018.
- [73] F. Mussa-Ivaldi, N. Hogan, and E. Bizzi, “Neural, mechanical, and geometric factors subserving arm posture in humans,” *The Journal of Neuroscience*, vol. 5, no. 10, pp. 2732–2743, 1985.
- [74] A. De and U. Tasch, “A two-dof manipulator with adjustable compliance capabilities and comparison with the human finger,” *Journal of Robotic Systems*, vol. 13, no. 1, pp. 25–34, 1996.
- [75] A. Z. Hajian and R. D. Howe, “Identification of the Mechanical Impedance at the Human Finger Tip,” *Journal of Biomechanical Engineering*, vol. 119, pp. 109–114, 1997.
- [76] N. Hogan, “Impedance Control: An Approach to Manipulation: Part II-Implementation,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 107, pp. 8–16, 03 1985.
- [77] H. O. Lim and K. Tanie, “Human safety mechanisms of human-friendly robots: Passive viscoelastic trunk and passively movable base,” *International Journal of Robotics Research*, vol. 19, no. 4, pp. 307–335, 2000.
- [78] I. A. Gravagne and I. D. Walker, “Manipulability, force, and compliance analysis for planar continuum manipulators,” *IEEE Transactions on Robotics and Automation*, vol. 18, pp. 263–273, June 2002.
- [79] D. J. Montana, “Contact Stability for Two-Fingered Grasps,” *IEEE Transactions on Robotics and Automation*, vol. 8, no. 4, pp. 421–430, 1992.
- [80] J. Morrow, H.-s. Shin, C. Phillips-Grafflin, S.-H. Jang, J. Torrey, R. Larkins, S. Dang, Y.-l. Park, and D. Berenson, “Improving Soft Pneumatic Actuator fingers through integration of soft sensors, position and force control, and rigid fingernails,” *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5024–5031, 2016.

- [81] S. Chitta, I. Sucan, and S. Cousins, “Moveit! [ros topics],” *IEEE Robotics Automation Magazine*, vol. 19, pp. 18–19, March 2012.
- [82] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, “Ros: an open-source robot operating system,” in *ICRA Workshop on Open Source Software*, 2009.
- [83] B. Calli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar, “Benchmarking in manipulation research: Using the yale-cmu-berkeley object and model set,” *IEEE Robotics & Automation Magazine*, vol. 22, no. 3, pp. 36–52, 2015.
- [84] C. Ferrari and J. Canny, “Planning optimal grasps,” in *Proceedings 1992 IEEE International Conference on Robotics and Automation*, pp. 2290–2295, IEEE, 1992.
- [85] D. M. Aukes and M. R. Cutkosky, “Simulation-based tools for evaluating underactuated hand designs,” *2013 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2067–2073, May 2013.
- [86] Tracker, “Tracker video analysis and modeling software, version 5.0.7,” *Computer Software*, Mar 2019.
- [87] C. Chi and S. Song, “Garmentnets: Category-level pose estimation for garments via canonical space shape completion,” *arXiv preprint arXiv:2104.05177*, 2021.
- [88] M. Shibata, T. Ota, Y. Endo, and S. Hirai, “Handling of hemmed fabrics by a single-armed robot,” in *2008 IEEE International Conference on Automation Science and Engineering*, pp. 882–887, IEEE, 2008.
- [89] Y. Moriya, D. Tanaka, K. Yamazaki, and K. Takeshita, “A method of picking up a folded fabric product by a single-armed robot,” *ROBOMECH Journal*, vol. 5, no. 1, pp. 1–12, 2018.
- [90] D. Seita, N. Jamali, M. Laskey, R. Berenstein, A. K. Tanwani, P. Baskaran, S. Iba, J. Canny, and K. Goldberg, “Robot bed-making: Deep transfer learning using depth sensing of deformable fabric,” *arXiv preprint arXiv:1809.09810*, vol. 26, 2018.
- [91] J. Matas, S. James, and A. J. Davison, “Sim-to-real reinforcement learning for deformable object manipulation,” in *Conference on Robot Learning*, pp. 734–743, PMLR, 2018.

- [92] J. Borràs, G. Alenyà, and C. Torras, “A grasping-centered analysis for cloth manipulation,” *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 924–936, 2020.
- [93] P. Koustoumpardis and N. Aspragathos, “A review of gripping devices for fabric handling,” *hand*, vol. 19, p. 20, 2004.
- [94] B. Sun and X. Zhang, “A new electrostatic gripper for flexible handling of fabrics in automated garment manufacturing,” in *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*, pp. 879–884, IEEE, 2019.
- [95] S. Ku, J. Myeong, H.-Y. Kim, and Y.-L. Park, “Delicate fabric handling using a soft robotic gripper with embedded microneedles,” *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4852–4858, 2020.
- [96] P. N. Koustoumpardis, K. X. Nastos, and N. A. Aspragathos, “Underactuated 3-finger robotic gripper for grasping fabrics,” in *2014 23rd International Conference on Robotics in Alpe-Adria-Danube Region (RAAD)*, pp. 1–8, IEEE, 2014.
- [97] P. N. Koustoumpardis, S. Smyrnis, and N. A. Aspragathos, “A 3-finger robotic gripper for grasping fabrics based on cams-followers mechanism,” in *International Conference on Robotics in Alpe-Adria Danube Region*, pp. 612–620, Springer, 2017.
- [98] M. J. Thuy-Hong-Loan Le, A. Landini, M. Zoppi, D. Zlatanov, and R. Molino, “On the development of a specialized flexible gripper for garment handling,” *Journal of Automation and Control Engineering Vol.*, vol. 1, no. 3, 2013.
- [99] S. Donaire, J. Borràs, G. Alenyà, and C. Torras, “A versatile gripper for cloth manipulation,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6520–6527, 2020.
- [100] M. H. Raibert and J. J. Craig, “Hybrid Position/Force Control of Manipulators,” *Journal of Dynamic Systems, Measurement, and Control*, vol. 103, pp. 126–133, 06 1981.
- [101] B. Aktaş and R. D. Howe, “Flexure mechanisms with variable stiffness and damping using layer jamming,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 7616–7621, IEEE, 2019.

- [102] C. Eppner, R. Deimel, J. Alvarez-Ruiz, M. Maertens, and O. Brock, “Exploitation of environmental constraints in human and robotic grasping,” *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 1021–1038, 2015.
- [103] J. Wang and E. Olson, “AprilTag 2: Efficient and robust fiducial detection,” in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4193–4198, IEEE, oct 2016.
- [104] ASTM, “Test method for thickness of textile materials,” *ASTM International*, 2019.
- [105] N. R. Sinatra, C. B. Teeple, D. M. Vogt, K. K. Parker, D. F. Gruber, and R. J. Wood, “Ultragentle manipulation of delicate structures using a soft robotic gripper,” *Science Robotics*, vol. 4, no. 33, 2019.
- [106] B. Shih, D. Drotman, C. Christianson, Z. Huo, R. White, H. I. Christensen, and M. T. Tolley, “Custom soft robotic gripper sensor skins for haptic object visualization,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, p. 494–501, IEEE, Sep 2017.
- [107] M. Higashimori, H. Jeong, I. Ishii, M. Kaneko, A. Namiki, and M. Ishikawa, “A new four-fingered robot hand with dual turning mechanism,” in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, p. 2679–2684, IEEE, 2005.
- [108] C. B. Teeple, “Ctrl-p, v2.1,” *Computer Software*, Feb. 2020.
- [109] P. Paoletti, G. W. Jones, and L. Mahadevan, “Grasping with a soft glove: intrinsic impedance control in pneumatic actuators,” *Journal of The Royal Society Interface*, vol. 14, p. 20160867, Mar 2017.
- [110] L. Tian, H. Li, Q. Wang, X. Du, J. Tao, J. S. Chong, N. M. Thalmann, and J. Zheng, “Towards complex and continuous manipulation: A gesture based anthropomorphic robotic hand design,” *arXiv preprint arXiv:2012.10981*, 2020.
- [111] T. Feix, I. Bullock, Y. Gloumakov, and A. Dollar, “Effect of number of digits on human precision manipulation workspaces,” *IEEE Transactions on Haptics*, 2020.

- [112] M. Liarokapis and A. M. Dollar, “Deriving dexterous, in-hand manipulation primitives for adaptive robot hands,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1951–1958, IEEE, 2017.
- [113] Robotiq, “3-finger adaptive robot gripper,” *Online datasheet*, 2021.
- [114] J. Meng, L. Gerez, J. Chapman, and M. Liarokapis, “A tendon-driven, preloaded, pneumatically actuated, soft robotic gripper with a telescopic palm,” *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*, pp. 476–481, 2020.
- [115] J. Lee, J. Kim, S. Park, D. Hwang, and S. Yang, “Soft robotic palm with tunable stiffness using dual-layered particle jamming mechanism,” *IEEE/ASME Transactions on Mechatronics*, 2021.
- [116] Y. Li, Y. Wei, Y. Yang, and Y. Chen, “A novel versatile robotic palm inspired by human hand,” *Engineering Research Express*, vol. 1, no. 1, p. 015008, 2019.
- [117] P. Capsi-Morales, G. Grioli, C. Piazza, A. Bicchi, and M. G. Catalano, “Exploring the role of palm concavity and adaptability in soft synergistic robotic hands,” *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4703–4710, 2020.
- [118] A. Yamaguchi, K. Takemura, S. Yokota, and K. Edamura, “A robot hand using electro-conjugate fluid: Grasping experiment with balloon actuators inducing a palm motion of robot hand,” *Sensors and Actuators A: Physical*, vol. 174, pp. 181–188, 2012.
- [119] V. Subramaniam, S. Jain, J. Agarwal, and P. Valdivia y Alvarado, “Design and characterization of a hybrid soft gripper with active palm pose control,” *The International Journal of Robotics Research*, vol. 39, no. 14, pp. 1668–1685, 2020.
- [120] Y. Sun, Q. Zhang, and X. Chen, “Design and analysis of a flexible robotic hand with soft fingers and a changeable palm,” *Advanced Robotics*, vol. 34, no. 16, pp. 1041–1054, 2020.
- [121] H. Wang, F. J. Abu-Dakka, T. N. Le, V. Kyrki, and H. Xu, “A novel soft robotic hand design with human-inspired soft palm: Achieving a great diversity of grasps,” *IEEE Robotics & Automation Magazine*, vol. 28, no. 2, pp. 37–49, 2021.

- [122] A. J. Spiers, B. Calli, and A. M. Dollar, “Variable-friction finger surfaces to enable within-hand manipulation via gripping and sliding,” *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4116–4123, 2018.
- [123] A. Pagoli, F. Chapelle, J. A. Corrales, Y. Mezouar, and Y. Lapusta, “A soft robotic gripper with an active palm and reconfigurable fingers for fully dexterous in-hand manipulation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7706–7713, 2021.
- [124] A. Cai, I. Pingel, D. Lorz, J. Beier, R. Horch, and A. Arkudas, “Force distribution of a cylindrical grip differs between dominant and nondominant hand in healthy subjects,” *Archives of Orthopaedic and Trauma Surgery*, vol. 138, no. 9, pp. 1323–1331, 2018.
- [125] M. Santello, M. Flanders, and J. F. Soechting, “Postural hand synergies for tool use,” *Journal of neuroscience*, vol. 18, no. 23, pp. 10105–10115, 1998.
- [126] E. Todorov and Z. Ghahramani, “Analysis of the synergies underlying complex hand manipulation,” in *The 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 2, pp. 4637–4640, IEEE, 2004.
- [127] A. Holladay, R. Paolini, and M. T. Mason, “A general framework for open-loop pivoting,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3675–3681, IEEE, 2015.
- [128] K. P. Becker, N. W. Bartlett, M. J. Malley, P. M. Kjeer, and R. J. Wood, “Tunable friction through constrained inflation of an elastomeric membrane,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4352–4357, IEEE, 2017.
- [129] M. A. Robertson and J. Paik, “New soft robots really suck: Vacuum-powered systems empower diverse capabilities,” *Science Robotics*, vol. 2, no. 9, 2017.
- [130] Y. Yamada and T. Nakamura, “Laminated foam-based soft actuator for actuatable flexible structure,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4359–4364, Nov. 2019. ISSN: 2153-0866.
- [131] K. P. Becker, C. B. Teeple, N. Charles, L. Mahadevan, and R. J. Wood, “Grasping via entanglement,” *In Review*, 2022.

- [132] J. Tremblay, T. To, B. Sundaralingam, Y. Xiang, D. Fox, and S. Birchfield, “Deep object pose estimation for semantic robotic grasping of household objects,” *arXiv preprint arXiv:1809.10790*, 2018.
- [133] D. Gálvez-López, M. Salas, J. D. Tardós, and J. Montiel, “Real-time monocular object slam,” *Robotics and Autonomous Systems*, vol. 75, pp. 435–449, 2016.
- [134] L. Bo, X. Ren, and D. Fox, “Unsupervised feature learning for rgb-d based object recognition,” in *Experimental robotics*, pp. 387–402, Springer, 2013.
- [135] A. Eitel, J. T. Springenberg, L. Spinello, M. Riedmiller, and W. Burgard, “Multimodal deep learning for robust rgb-d object recognition,” in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 681–687, IEEE, 2015.
- [136] T.-H. Pham, A. Kheddar, A. Qammaz, and A. A. Argyros, “Towards force sensing from vision: Observing hand-object interactions to infer manipulation forces,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2810–2819, 2015.
- [137] R. Calandra, A. Owens, M. Upadhyaya, W. Yuan, J. Lin, E. H. Adelson, and S. Levine, “The feeling of success: Does touch sensing help predict grasp outcomes?,” *arXiv preprint arXiv:1710.05512*, 2017.
- [138] R. Li, R. Platt, W. Yuan, A. ten Pas, N. Roscup, M. A. Srinivasan, and E. Adelson, “Localization and manipulation of small parts using gelsight tactile sensing,” in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3988–3993, IEEE, 2014.
- [139] J. Bimbo, P. Kormushev, K. Althoefer, and H. Liu, “Global estimation of an object’s pose using tactile sensing,” *Advanced Robotics*, vol. 29, no. 5, pp. 363–374, 2015.
- [140] J. Bimbo, S. Luo, K. Althoefer, and H. Liu, “In-hand object pose estimation using covariance-based tactile to geometry matching,” *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 570–577, 2016.
- [141] Z. Kappassov, J.-A. Corrales, and V. Perdereau, “Tactile sensing in dexterous robot hands,” *Robotics and Autonomous Systems*, vol. 74, pp. 195–220, 2015.

- [142] S. Luo, J. Bimbo, R. Dahiya, and H. Liu, “Robotic tactile perception of object properties: A review,” *Mechatronics*, vol. 48, pp. 54–67, 2017.
- [143] C. B. Teeple, K. P. Becker, and R. J. Wood, “Soft curvature and contact force sensors for deep-sea grasping via soft optical waveguides,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1621–1627, Oct 2018.
- [144] L. Viry, A. Levi, M. Totaro, A. Mondini, V. Mattoli, B. Mazzolai, and L. Beccai, “Flexible three-axial force sensor for soft and highly sensitive artificial touch,” *Advanced Materials*, vol. 26, pp. 2659–2664, feb 2014.
- [145] R. K. Kramer, “Soft electronics for soft robotics,” in *Micro and Nanotechnology Sensors, Systems, and Applications VII* (T. George, A. K. Dutta, and M. S. Islam, eds.), SPIE-Intl Soc Optical Eng, 2015.
- [146] J. C. Yeo, H. K. Yap, W. Xi, Z. Wang, C.-H. Yeow, and C. T. Lim, “Flexible and stretchable strain sensing actuator for wearable soft robotic applications,” *Advanced Materials Technologies*, vol. 1, p. 1600018, may 2016.
- [147] G. Zöllner, V. Wall, and O. Brock, “Acoustic sensing for soft pneumatic actuators,” in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 6986–6991, IEEE, 2018.
- [148] V. Wall, G. Zöllner, and O. Brock, “A method for sensorizing soft actuators and its application to the rbo hand 2,” in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4965–4970, IEEE, 2017.
- [149] V. Wall and O. Brock, “Multi-task sensorization of soft actuators using prior knowledge,” in *2019 International Conference on Robotics and Automation (ICRA)*, pp. 9416–9421, IEEE, 2019.
- [150] G. Zöllner, V. Wall, and O. Brock, “Active acoustic contact sensing for soft pneumatic actuators,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7966–7972, IEEE, 2020.
- [151] M. A. Graule, T. P. McCarthy, C. B. Teeple, J. Werfel, and R. J. Wood, “Somogym: A toolkit for developing and evaluating controllers and reinforcement learning algorithms for soft robots,” in *IEEE Robotics and Automation Letters*, IEEE, 2022 (In Review).

- [152] Q. Wan and R. D. Howe, “Modeling the Effects of Contact Sensor Resolution on Grasp Success,” *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1933–1940, 2018.
- [153] L. P. Jentoft, Q. Wan, and R. D. Howe, “Limits to compliance and the role of tactile sensing in grasping,” in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pp. 6394–6399, IEEE, 2014.
- [154] Q. Wan, R. P. Adams, and R. D. Howe, “Variability and predictability in tactile sensing during grasping,” *IEEE International Conference on Robotics and Automation*, vol. 2016-June, pp. 158–164, 2016.
- [155] C. To, T. L. Hellebrekers, and Y.-L. Park, “Highly stretchable optical sensors for pressure, strain, and curvature measurement,” in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, sep 2015.
- [156] D. Vogt, Y.-L. Park, and R. J. Wood, “A soft multi-axis force sensor,” in *2012 IEEE Sensors*, Institute of Electrical and Electronics Engineers (IEEE), oct 2012.
- [157] H. Yousef, M. Boukallel, and K. Althoefer, “Tactile sensing for dexterous in-hand manipulation in robotics—a review,” *Sensors and Actuators A: Physical*, vol. 167, no. 2, pp. 171 – 187, 2011. Solid-State Sensors, Actuators and Microsystems Workshop.
- [158] B. Nie, R. Li, J. Cao, J. D. Brandt, and T. Pan, “Flexible transparent iontronic film for interfacial capacitive pressure sensing,” *Advanced Materials*, vol. 27, pp. 6055–6062, sep 2015.
- [159] G. P. Agrawal, *Fiber-optic communication systems*. Wiley, 3rd ed., 2010.
- [160] A. Yariv and P. Yeh, *Photonics: optical electronics in modern communications*. Oxford University Press, 6 ed., 2007.



Toward Soft Curvature and Contact Sensors for Grasping

A.1 Introduction

Gentle grasping has become particularly relevant for recovering samples of deep-sea organisms. Minimizing damage to organisms while taking samples is critical for biologists to accurately study morphology and DNA expression. Thus, in a recent effort toward gentle deep sea sampling, soft robotic grippers have been used as tools on deep-sea remotely operated vehicles (ROVs) to cause less damage to animals than traditional rigid grippers [35].

Soft robots have been shown to interact more-gently with objects in their environments using minimal control effort compared to rigid robots [4]. However, while compliance allows soft robots to passively adapt their shape to complex or uncertain objects, compliance also introduces uncertainty in the robot's position and applied forces. Uncertainties in actuator positioning, contact force direction, and gripping force make the grasp quality difficult to predict [152, 7]. These uncertainties are magnified when soft robots are teleoperated on ROVs due to limited state feedback provided to the operator.

While gentle deep-sea grasping has been demonstrated, sampling tasks are often still cumbersome and lengthy due to poor estimation of arm, hand, and soft finger positioning. To enhance the robustness of grasping operations to actuator and environmental uncertainty, many studies of under-actuated rigid hands have used on-board proprioception and contact force sensing in gripper control strategies [153, 154, 22]. Using these two pieces




Figure A.1: We outfitted a soft robotic gripper platform developed in our lab with curvature and contact sensors. (a) This platform has been used on a variety of deep-sea sampling expeditions to grasp delicate corals and other organisms. (b) In this work, we use soft optical waveguides as sensors. (c) These sensors are implemented onboard the fingers of the gripper, with curvature sensors integrated directly into the structure.

of sensor information in the control of soft actuators, we can improve the robustness of grasps and enable quicker, more effective sampling of deep-sea organisms.

To implement shape and contact force sensing on soft robots operating in the deep-sea, we chose to use soft optical waveguides. Intentionally lossy, soft waveguides have been used as strain sensing elements in other works through gold-plated elastomer channels [155], and urethane rubber cores surrounded by a silicone cladding [62]. Sensing is accomplished by emitting light into the guide, then measuring the intensity loss of transmitted light as a result of stretching, compressing, or bending.

Compared to other soft sensing modalities, lossy optical waveguides made of elastomers are the most robust for use in the deep-sea, where temperatures reach as low as 2°C and hydrostatic pressures as high as 9000 psi at the ocean floor. The low coefficient of thermal expansion and incompressibility of elastomers makes them insensitive to changes in temperature and hydrostatic pressure. Furthermore, optical properties of elastomers such as refractive index change very little with temperature. Finally, while optical couplings add some complexity, the sensors can be fabricated without local electrical parts to dramatically simplify the waterproofing process necessary for the deep-sea.

By contrast, electrical soft sensors used in soft robots, including resistive and capacitive sensors, are less-suitable for use in the deep-sea. Resistive sensors rely on changes in resistance via deformation of channels filled with conductive liquids (such as liquid metal or ionic liquids), piezoresistive materials, or conductive textiles [145, 156, 146]. These sensors are simple in design, but suffer from thermal drift [157]. In addition, liquid resistive materials such as eGaIn or ionic liquids freeze below 10 – 18 °C, rendering them unusable in the deep-sea. Capacitive sensors are usually implemented with similar materials as resistive sensors [144, 158], and tend to be more stable in response to temperature changes, but can be quite sensitive to electrical noise [157].

A.2 Sensor Design

A.2.1 Design Criteria

In addition to the necessity for invariance to environmental changes mentioned above, a primary objective for our soft sensing elements is the ability to distinguish between bending and external contact forces. This distinction can be made by mechanically decoupling the sensitivity of sensing elements to different modes of deformation through geometry and material selection.

A.2.2 Waveguide Design for Soft Fingers

Taking the above criteria into account, we designed the structure of our soft waveguide sensors to enable simple integration with our existing deep-sea soft gripper platform. Waveguides were designed to fit into the 2.5 mm-thick skin of the fingers (bellows-style bending actuators) at the inside of the bend, as shown in Figure A.1. The waveguides are patterned in loops that begin and end on the proximal side of the finger for ease of optical connections. The smallest reliable waveguide diameter was 1 mm due to fabrication limitations (discussed later), so only two waveguides could be placed into the volume of the flat face of each finger. Further fabrication refinement may enable the incorporation of more waveguides.

To achieve total internal reflection in our waveguides, the core material for both sensors is ClearFlex 30 (Smooth-On Inc.). Clear-Flex was chosen for its high refractive index of 1.488 to maximize the difference between the indices of the core and silicone cladding

(approximately 1.38-1.41). This material combination promotes total internal reflection for angles of incidence up to approximately $18^\circ - 22^\circ$.

Materials for the sensor body (which also acts as the optical cladding) are chosen using mechanical stiffness matching [30] by taking into account the stiffnesses of actuators and target objects to be grasped. To design a bending sensor that is relatively insensitive to contact forces, we seek to maximize the cladding material stiffness while being limited by the stiffness of an actuator. Thus, it is convenient to use the same material as the finger actuators, Smooth-Sil 950 (Smooth-On Inc.). To design a maximally sensitive contact force sensor, we used Ecoflex 00-30 (Smooth-On Inc.) as the cladding material because of its extremely low stiffness (27 times smaller elastic modulus than Smooth-Sil 950). Practically, Ecoflex 00-30 is also the lowest-stiffness silicone we can use before our fabrication method fails.

Finally, we amplify the effect of distributed contact force on the finger by including rigid indenters attached across each waveguide. These indenters amplify contact pressure from objects directly onto the fibers, as shown in Fig. A.1.

A.2.3 Deep-Sea Interface Design

We used commercially available plastic optical fiber in a modular design to allow the soft sensing elements to interface with typical ROV systems. The plastic optical fibers transmit light between the soft sensing elements and optoelectronics, allowing all electronics to be located proximally in a waterproof and depth-proof container. In addition, modular quick optical disconnects enable easy swapping of actuators and sensors in the field.

A.3 Modeling optical losses as a function of deformation

To understand how design decisions affect sensor performance, we used simple models for geometric optical losses in our soft optical fibers as well as material deformation and stress. The effect of normal strain on the optical losses in our soft waveguides is much more complex, so the relationship between compression and optical intensity loss is not modeled.

Optical waveguides have several loss mechanisms due to the complexity of light transmission through a waveguide, as shown in Fig. A.2, however many of them can be ignored for modeling purposes. We are only interested in the intensity loss incurred as a




Figure A.2: Large optical waveguides have several important loss mechanisms, some of which can be ignored when modeling intensity loss in our sensors. Absorption, volume scattering, and microbending will be ignored because they remain roughly constant regardless of deformation.




Figure A.3: Schematic diagrams of optical phenomena in soft waveguides during (a) bending and (b) normal force application, including key dimensions and light loss modes.

function of deformation, so we define the intensity loss I_{loss} in decibels as:

$$I_{loss} = 10 \log_{10} \left(\frac{I_0}{I} \right) \quad (\text{A.1})$$

where I is the measured intensity, and I_0 is the baseline intensity under no deformation.

A.3.1 Curvature Sensors

To construct a model for curvature losses in our soft optical fibers, we can focus on macrobending and surface roughness as the primary loss mechanisms. The effect of macrobending has been well-studied for circular optical fibers. Our waveguides have substantially rounded corners, so we assume an approximately-circular cross-section. As such,

the attenuation coefficient, α_B , per unit length in a fiber undergoing a bend with curvature κ can be written as [159]:

$$\alpha_B = C_{b,1} \exp\left(-\frac{n_{core}^2 - n_{clad}^2}{a} \frac{1}{\kappa}\right) \quad (\text{A.2})$$

where n_{core} and n_{clad} are the refractive indices of the core and cladding, a is the radius of the waveguide, and $C_{b,1}$ is a constant. Using Beer's law, we can find the total power around a bend of length L to be $I = I_0 \exp(\alpha_B L)$. We can rearrange Beer's law into the form in Eq. A.1 to obtain the loss per unit length around the bend in dB/cm:

$$I_{loss,mb} = C_{b,2} \alpha_B \quad (\text{A.3})$$

where $C_{b,2}$ is a constant that transforms the units to decibels and the base of the exponent to base 10.

The role of surface roughness in bending losses has not been studied thoroughly in the literature, but plays a critical role in the bending behavior of soft waveguides in this work and others [62]. The bending response is highly dependent on the direction of curvature due to differences in surface roughness on opposing surfaces of the guide, (a by-product of the fabrication process as discussed in the Fabrication Section).

We can gain an intuition for how surface roughness affects the relationship between bending and optical intensity loss by examining the simplified problem of a planar waveguide (see Fig. A.3a). At rest, light rays propagate with an equal number of reflections on both surfaces of the guide. However, during bending, light begins to reflect more often on the outer surface (dashed blue rays in Fig. A.3a). Thus, with a rough outer surface, substantial bending losses occur at much smaller curvatures (larger radii) than pure macrobending. When bent in the opposite direction (smooth outer surface), macrobending takes over (a phenomenon seen in [62] as well).

A.3.2 Contact (Normal) Force Sensors

To construct a model for losses in our sensors as a function of input force on a rigid indenter, we need a material model and an optical loss model. In this work, we use the neo-Hookean material model, but only develop an intuition for the optical losses as a function of compression due to the complexity of the problem.

We begin with a cross section of the soft optical sensor with an applied force from a flat indenter, as shown in Fig. A.3b. The indenter is displaced into the sensor, which compresses the sensor body (cladding and core), inducing a local change in cross-section of the fiber and a small bend at either edge.

Under uniaxial compression, we relate the axial engineering strain, ϵ , exhibited in the sensor body to the applied axial engineering stress, σ_{zz} , using the Neo-Hookean material model for large deformations, a common material model for elastomers. This relationship takes the form:

$$\sigma_{zz} = C_1 \left(\lambda(\epsilon) - \frac{1}{\lambda(\epsilon)^2} \right) \quad (\text{A.4})$$

where C_1 is a material constant and λ is the principle stretch. Rather than implement this material model for the core and both cladding layers in series, we assume the ratio of core to cladding stiffness is large enough to ignore the effect of cladding compression. Thus, we approximated C_1 based on the Young's modulus and Poisson's ratio of the core (roughly 0.5 MPa and 0.47 respectively).

Based on a high-level analysis, it is likely that the majority of optical intensity losses come from the small bends produced on either side of the indenter, not optical mode de-coupling. Due to the large dimensions of our waveguides (1 mm side lengths), there exist on the order of 10^7 optical transmission modes. A rigorous analysis would require a sum of mode overlap integrals over all modes [160]. However, only the highest-order modes with the lowest power densities would be attenuated due to dimension changes in the waveguide. Thus, mode de-coupling is likely not a major source of losses. The small bends on either side of the indenter are the only remaining loss mechanism, so we would expect the relationship between compressive strain and optical loss to remain roughly similar regardless of indenter size.

A.4 Methods

We designed and fabricated soft waveguide sensor arrays (2×1) both as discrete units and integrated into the structure of our existing soft finger actuators. We then characterized the light loss in these sensors as a function of curvature and local normal force.




Figure A.4: Fabrication process for soft robotic finger actuators with integrated soft waveguides. (a) The finger actuator is molded per [35], partially cured, then (b) part of the mold is removed exposing channels that create three walls of the waveguide cladding. (c) A separate rubber graft is wet-bonded over the open channels per [156]. (d) Once fully cured, the channels are injected with the core material. (e) Once fully cured, the proximal end is trimmed and (f) glued into an alignment clip. Plastic optical fibers are glued into the other side of the clip. (g) Finger actuator with integrated optical curvature sensor array. (h) Discrete soft optical sensor arrays.

A.4.1 Fabrication of Soft Waveguides

The fundamental building block in both curvature and contact sensors is a soft waveguide, developed and fabricated using molding methods similar to those presented in [62] and [156]. Waveguides are molded by creating channels in the cladding material, then filling the channels with core material. In the case of discrete sensors, this cladding material is a strip of rubber only slightly thicker than the optical core (as seen in Fig. A.4h). When integrated into the structure of a soft finger (as is the case for our curvature sensors), the cladding material is the skin of the flat face of the finger (as seen in Fig. A.4b).

A.4.2 Fabrication of Integrated Curvature Sensors

The methods below (and in Figure A.4) focus on the fabrication of curvature sensors integrated into the structure of bellows-style soft finger actuators. Due to fabrication limitations, contact sensors are molded as discrete sensors, then adhered to the flat face of actuators. The process modification for discrete units is described later in this section.

To create a finger with integrated waveguides, a mold of our typical bellows-style bending actuator was modified to include grooves (with a square cross-section) that later become waveguides. The mold also accommodates alignment clips for plastic optical fibers that interface with the soft waveguides. The molds were 3D printed on a Stratasys Objet 30 printer with polyjet VeroBlue and VeroClear material, then baked overnight in an oven at 60 °C before use.

Next, the bellows actuator is molded from Smooth-Sil 950 and placed in a pressure chamber according to the procedure described in [35], as shown in Fig. A.4a. The actuator is removed from the pressure chamber before it has finished curing (4 hours from silicone mix time), and the channel-side face is removed to expose the flat side of the finger (Fig. A.4b). From here, a flat graft of Smooth-Sil 950 is wet-bonded to the exposed flat face of the finger using the procedure found in [156], creating enclosed channels (Fig. 4c). Everything is then cured overnight at room temperature. The resulting roughness on the channel walls formed by the printed mold are much rougher than the grafted side of the channel (atomically smooth) because the new layer of rubber is cured in-place without any surface contact.

After the finger is fully cured, the channels are cut open at the proximal end of the finger, and the core of the optical waveguides is created. Clear-Flex 30 (Smooth-On, Inc.) is degassed and injected into each of the cladding channels until it comes out the other side (Fig. A.4d). The assembly is then allowed to cure at room temperature overnight.

With the soft optical waveguides fully formed, rigid plastic optical fibers are aligned and adhered. First the proximal section of the sensor is cut with a razor blade to ensure the a clean optical surface (Fig. A.4e). Then, 9-inch lengths of plastic optical fiber (Industrial Fiberoptics, 1 mm core diameter) are aligned with the soft waveguides using a custom-designed 3D printed alignment clip (Markforged, Onyx). The clip is adhered to the finger by silicone adhesive (Silpoxy, Smooth-On), and a strong physical and optical bond between the soft urethane cores and the plastic fibers is made with a cyanoacrylate glue (Loctite 401). Finally, the entire assembly is placed under a 1 kg mass until all adhesives are cured.

To finish the fingers with integrated sensors, the opening to the finger is plumbed with a custom 3D printed adapter and pneumatic hardware, heat shrink tubing is placed around the base of the finger (and alignment clip), and open-cell memory foam is attached to the flat face of the finger. All of these steps follow the procedure described in [35].

To create a discrete 2×1 sensor array, the finger mold is replaced with a flat, rectangular mold to generate three sides of the channel. All subsequent steps (applying the flat graft, cutting the proximal end, attaching plastic optical fibers via alignment clip) remain the same.

A.4.3 Fabrication of Contact Force Sensors

To make contact sensors, Ecoflex 00-30 is used as the cladding, while the core remains Clear-Flex 30. Since the cladding material is different from the material used to make the finger, direct integration would pose significant fabrication challenges. Instead, discrete contact sensor arrays are adhered to the flat face of a finger using silicone adhesive (Silpoxy). In addition, 3D printed rigid indenters (Markforged, Onyx) are embedded into the memory foam with the pointed edge directly overtop of the soft waveguides, as seen in Fig. A.1c.

A.4.4 Data Acquisition and Processing

To acquire light intensity signals from our sensors, red LEDs designed to couple with 1mm fiber (SP000063802, Broadcom Limited) provide a light source, and custom designed acrylic housings allow digital light intensity sensor chips (LTR-329ALS-01, Lite-on Inc.) to couple with the fibers. We used a microcontroller (Arduino Nano, 16 MHz) to interface digitally with the light intensity sensors. Ranges were chosen on a case-by-case basis to maximize the resolution without saturating. All data were captured at 10 Hz, which is more than sufficient for our fingers operating underwater at frequencies much less than 1 Hz.

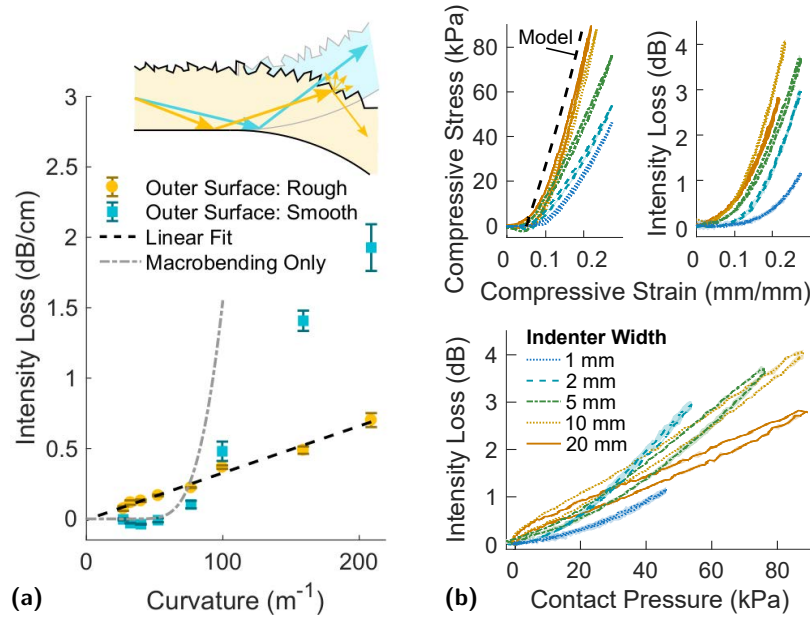


Figure A.5: The sensitivity of discrete curvature and contact force sensors show similar trends to our models and intuition. (a) The intensity loss as function of curvature follows a linear relationship when the outer surface is rough, and agrees in shape with the macrobending model when the outer surface is smooth. (b) Models for stress as a function of strain (upper left) are validated by experimental data, and the intensity loss vs. strain curves (upper right) are roughly invariant to indenter width, validating our intuition. These two relationships combine to form the desired calibration curve (bottom). *Data points in (a) represent the mean of $n = 20$ trials, and curves in (b) represent the mean of $n = 3$ trials.*

A.5 Sensor characterization

A.5.1 Characterization of Discrete Sensors

Characterization procedures were performed on each type of discrete sensor based on its intended use (curvature and normal force). These calibration procedures are intended to approximate field conditions while also isolating the effects of bending and local normal force. While the variation in sensitivity among sensors was not quantified, large variations in the baseline intensity (within approximately 50%) were noted, stemming from variation in soft fiber alignment.

The effect of bending on light loss in our sensors was investigated by manually bending standalone sensors around cylinders ranging from 7.9 mm to 76 mm in diameter,

as exhibited in the supplementary video. Under typical actuation pressures (< 25 psi), the fingers exhibit curvatures of at most 52 m^{-1} (corresponding to a diameter of 38 mm). For diameters larger than 25 mm, the cylinders were attached to the tip of the sensor body with tape and rolled toward the base. For smaller diameters, the sensors were bent 180 degrees around stationary cylinders.

Calibration curves are based on the intensity loss (as defined in Eq. A.1) per unit length around the bend that occurs at the point of maximum bend length for each cylinder. This corresponds to the peak intensity loss during the rolling procedure, producing calibration curves as shown in Fig. A.5a.

The effect of applied normal force on light loss in our sensors was characterized by pressing flat indenters into the sensor using an Instron uniaxial material testing machine. The indenters were 3D printed (Polyjet VeroWhite) with widths ranging from 1 mm to 20 mm, and a length of 20mm, spanning the width of the sensors. The sample (2.2 mm thick) was compressed by 0.6 mm and released at a strain rate of 0.05 mm/sec while recording the resulting axial force and light intensity (see supplementary video).

The calibration curve for force sensors is defined using the contact pressure (force divided by estimated indenter area) vs. intensity loss measurements averaged over three trials for each indenter, as shown in Fig. A.5b. It should be noted that contact sensors are also sensitive to curvature due to fabrication limitations, however this effect can be compensated by simultaneously measuring curvature using the curvature sensors.

A.5.2 Characterization of On-board Sensors

Curvature and contact sensors implemented on-board each soft finger were characterized in response to finger actuation pressure, curvature, and contact force (if applicable). The effect of actuation pressure on intensity loss was characterized for both types of sensors by blocking the finger's bending motion and applying pressure up to 24 psi in steps of 2 psi. The effect of curvature was characterized for both types of sensors using the same procedure as for standalone curvature sensors without actuating the finger. The effect of contact force on contact sensors was characterized by manually pressing the embedded indenter against a load cell.

The calibration function used to describe intensity loss I_{loss} as a function of actuation pressure p_a , curvature κ , and contact force f_c on a finger is derived from empirical

observations:

$$I_{loss} = a_1 \exp(a_2 p_a) + a_3 \kappa + a_4 f_c \quad (\text{A.5})$$

where a_i are calibration constants. For each finger, parameters a_1 and a_2 were fit using a nonlinear least squares regression on data from blocked actuation tests. a_3 and a_4 were characterized using the datasets from curvature and contact tests respectively.

A.6 Evaluation of models and experiments

Experimental characterization of discrete curvature and contact sensors show similar trends to our models, as shown in Fig. A.5. The curvature response of the curvature sensors agree in shape with pure macrobending when the outer surface of the bend is smooth, but show a strong linear relationship when the outer surface is rough. In addition, the relationship between compressive strain and stress in contact sensors agrees with the Neo-Hookean model. Finally, The optical loss as a function of compressive strain appears to be loosely invariant to indenter width for indenters wider than 1 mm, confirming our intuition that the majority of light loss is likely due to the small bends on the edges of the indenters.

We can use trends predicted by our models to estimate how changes in design parameters might affect sensor performance. For example, to increase the linearity of curvature the sensors, we should decrease the effects of macrobending and increase the effects of surface roughness. The macrobending model suggests a dependence of $\exp(-a^{-1})$ on the channel size and a dependence of $\exp(-n_{core}^2 - n_{clad}^2)$ on the refractive indices. Thus, we might consider decreasing the channel size a , or increasing the difference in refractive index between the core and cladding to lower the losses from macrobending. We could also consider introducing small surface roughness to linearize the response. Similarly, the force model suggests a direct dependence on the material parameter C_1 in eqn. A.4 which is proportional to elastic modulus, and inversely dependent on area. Therefore, to increase sensitivity to applied force, we could minimize the elastic modulus or the area of the indenter.

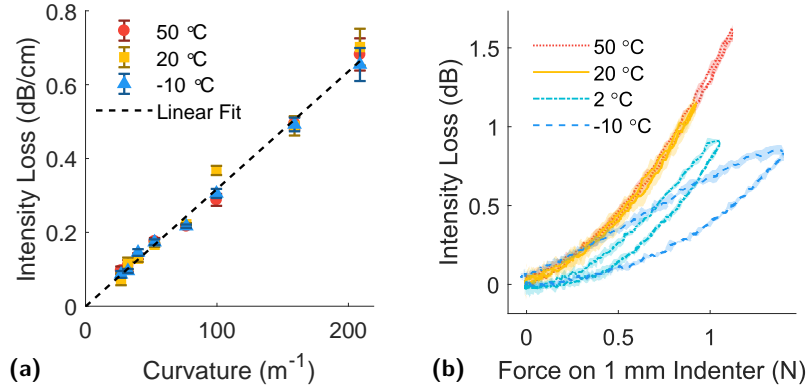


Figure A.6: The bending sensitivity of our soft fiber-optic sensors is invariant to operating temperature, while normal force sensitivity decreases with decreasing temperature. (a) Invariance of curvature sensors is illustrated by an overlap in calibration data for all curvatures measured. (b) However, temperature dependence of force sensors is illustrated by calibration tests with a 1 mm flat indenter. *Curvature points represent the mean of $n = 20$ repetitions for a single sensor. Force curves represent the mean of $n = 3$ trials. Error bars/regions represent one standard deviation from the mean.*

A.7 Characterization under deep-sea conditions

To determine the extent to which our soft optical sensors are affected by deep-sea conditions, we tested performance under simulated environments. Operating temperature and hydrostatic pressure effects were explored separately.

A.7.1 Temperature

The effect of operating temperature was evaluated using discrete curvature sensors and contact sensors at different ambient temperatures. Sensors were equilibrated for 15 minutes on a heating/cooling plate (Teca, Model AHP-301CPV) set to a desired temperature ($-10^{\circ}C$, $20^{\circ}C$, or $50^{\circ}C$). After equilibration, either a curvature or indentation calibration procedure was performed directly on the temperature plate. Comparisons of calibration curves for one curvature and one contact force sensor at temperatures from $-10^{\circ}C$, to $50^{\circ}C$ are shown in Figure A.6.

Based on these temperature-controlled results, we have determined that the bending sensitivity of our soft fiber-optic sensors is invariant to operating temperature within the measured range. This makes sense because our curvature sensors are thin sheets that are

sensitive to geometric optical losses, but they only incur small deformation in the material.

However, the normal force sensitivity of our sensors decreases as temperatures drop below room temperature. This makes sense because optical properties of materials do not change very quickly with temperature, but the elastic moduli of elastomers usually increase as temperatures decrease. In addition, the force sensors show a dramatic increase in hysteresis as the temperature drops below 0°C. This also makes sense due to the viscoelasticity of the polyurethane core material, which could be improved by material choice.

Finally, while the quasi-static sensitivity of our sensors exhibit the relationships described, the dynamic characteristics such as bandwidth are likely very temperature-dependent. These effects were not characterized in this work because deep-sea sampling procedures performed *in-situ* are slow enough to assume quasi-static conditions.

A.7.2 Hydrostatic Pressure

The effect of hydrostatic pressure on the sensitivity of our curvature sensors was characterized using a finger actuator with integrated curvature sensor array as well as a second discrete sensor array. Both were placed into a vessel pressurized between 0 psi and 4000 psi (2700m equivalent depth) with 4°C tap water using a pump and back-flow valve, as shown in the supplementary video. Optical sensing circuitry was waterproofed and connected to a pressure-rated electrical passthrough so sensor readings could be collected from outside the vessel.

Under high pressure, the finger was actuated and the resulting combined sensitivity to internal pressure and curvature was found. At each increment of high pressure, the finger’s internal pressure was first allowed to equilibrate, then a hand pump was used to pressurize the finger to a gauge pressure of 21 psi (set by an internal check valve), then the finger was vented back to the vessel pressure. An example of sensor readings during these actuation cycles is shown in Fig. A.7. The sensitivity of the sensor at each pressure was calculated as $S_{pres} = I_{loss,max}/P_{max}$, where $I_{loss,max}$ is the maximum intensity loss over one actuation cycle, and P_{max} is 21 psi.

Based on these tests, the sensitivities of our integrated curvature sensors are reasonably invariant to the combined effects of low temperatures and hydrostatic pressure beyond 1000 psi (700 m equivalent). This implies that curvature calibration done anywhere in this pressure range would be valid for the rest of the pressure range tested.

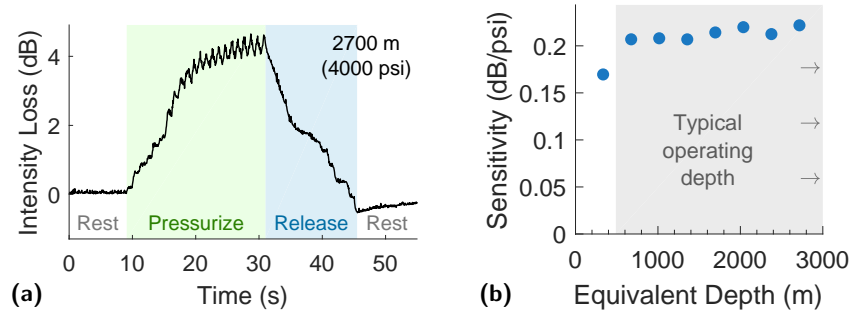


Figure A.7: The sensitivity of integrated curvature sensors does not change appreciably under typical deep-sea hydrostatic pressures. Pressure testing was performed in a pressure vessel at 4°C with the ability to actuate a finger. (a) Sensor readings are shown from a typical actuation cycle from 0 psi up to 21 psi and back down to 0 psi gauge pressure, with the large steps caused by successive pumps of the hand pump. (b) shows the sensitivity of the integrated curvature sensor as a function of equivalent depth.

A.8 Grasping Objects

To demonstrate how our integrated sensors can provide informative feedback during grasping, we performed grip-pull tests with a rigid cylinder, and underwater grasps of a compliant sphere. In the grip test, two fingers were mounted to the base of the Instron: Finger #1 with two integrated curvature sensors, and finger #2 with one functioning contact sensor. The fingers were actuated to grasp a cylinder of 1 in diameter, which was then pulled out of the grasp. The resulting curvature and contact force estimates are shown in A.8, where finger #2’s contact force estimate is curvature-corrected using the curvature estimates from finger #1. As shown, the curvature estimates provide higher-fidelity information about the grasp than actuation pressure alone. In addition, we performed underwater pick-and-release operations on an extremely compliant sphere to demonstrate successful sensor function underwater, as can be seen in the supplementary video.

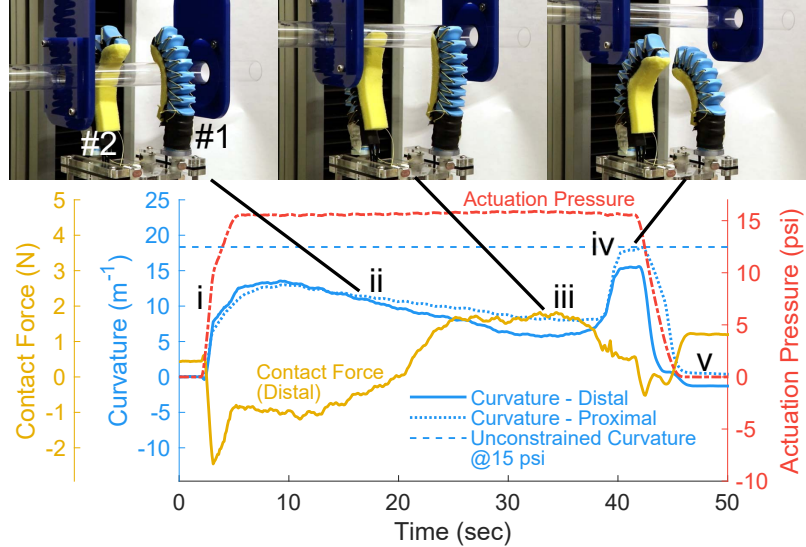


Figure A.8: Grasping a cylinder shows key points where curvature and contact force sensors provide more information than internal pressure alone. At *i*, the fingers are pressurized to 15 psi, but blocked by the cylinder. The cylinder is then moved upward, forcing fingers to straighten (*ii*). The cylinder continues upward until it reaches the fingertips (*iii*) and eventually slips out of the grasp, allowing the fingers to reach their unconstrained curved positions (*iv*). Finally, the fingers are de-pressurized (*v*).

A.9 Conclusions

In summary, we have designed and integrated soft optical waveguide sensors on-board the fingers of a soft robotic hand that enable proprioception and contact force measurements. We validated the function of the sensors in deep-sea conditions, and demonstrated their ability to provide useful state feedback during grasping in both air and water.

In future work, the sensitivity of contact sensors should be increased through improved modeling, the density of sensors in each finger should be increased by improving fabrication techniques, and other types of sensors could be explored such as multi-axis contact sensing. Application of these sensors in deep-sea soft manipulators will enable improved sampling of organisms through enhanced state feedback, and can be extended to enable grasp planning and other high-level autonomous behavior.

B

Smooth pressure control for soft robots

Smooth, high-bandwidth control of multiple pneumatic pressure signals was important for all of the work presented in this dissertation, but critical for my work on in-hand manipulation using soft hands. While a host of commercial options exist for control of air and gas pressure in industrial applications, these tend to be expensive, slow, and bulky. To this end, I developed a custom pressure control solution (“Ctrl-P”) that is scalable between 1 and 10 independent channels, has adaptable control range (depending on sensor choice), is compatible with most robotic systems (via ROS [82]), and is open source [108]. This system has been used extensively in most of my publications [64, 65, 66, 67, 69, 70], and the most-recent iteration is shown in Figure B.1.

For smooth, realtime control of numerous independent pressure signals, the open-source “Ctrl-P (Control Pressure) project” [108] was born. Based on a set of inexpensive proportional valves, industry-standard pressure sensors, a high-quality microcontroller, and cheap drive electronics, the hardware of this system is highly-capable, but reasonably priced. The majority of the development has been in software design, building both reliable, realtime firmware, and reliable, easy-to-use host-side control (via ROS [82]). This combination enables execution of arbitrary pressure trajectories (between -210 kPa **and** 275 kPa) on up to 10 channels in real time with an accuracy of 1.4 kPa and response time of 0.2 s. Additionally, ROS compatibility has enabled high-level coordination between pneumatic hands, our robot arm (UR5e), and camera systems for motion tracking via AprilTags [103]. Since this system is being actively developed at the time of writing, documentation about both the hardware and software is available on GitHub (https://github.com/cbteple/pressure_control_interface) [108].




Figure B.1: The custom pressure control system used throughout this body of work is part of an ongoing, open-source project [108]. The system can control pneumatic pressure between -210 kPa and 275 kPa with an accuracy of 1.4 kPa and a response time of 0.2 s.



Grasp Testing

This appendix includes the procedures, objects, and experiments used in Chapter 2.

C.1 Experimental setups for measuring grasp performance

C.1.1 Grasp testing

To measure the hand's region of acquisition and object size range for which grasping is successful, the experimental setup is shown in Fig. C.1

C.1.2 Robustness to external forces

To measure the robustness to external forces for each hand configuration, we used the experimental setup shown in Fig. C.2. With the hand mounted at a prescribed angle and grasping an object, a uniaxial testing machine was used to pull the object out of the grasp while simultaneously measuring the force applied to the object. Neglecting gravity, this setup is equivalent to pulling the object out of the grasp at the prescribed angle. This procedure was repeated three times for each angle tested.




Figure C.1: The experimental setups used to measure the hand's grasping performance consists of a three-axis positioning system and a tabletop. The positioning systems used were a) a custom-built three-axis CNC gantry, and b) a UR5e 6DOF robot arm (Universal Robots, Denmark)), each with a positioning accuracy of better than 1 mm. The object is placed precisely on the table by hand using a placement grid with radial tick marks every 5 mm.




Figure C.2: The experimental setup used to measure the hand's robustness to external forces over a variety of angles. The hand was mounted at prescribed angles using the fixture shown. With a grasp performed on the object, an Instron uniaxial testing machine was used to pull the object vertically. This is equivalent to pulling the object out at the prescribed angle.

C.2 Measuring finger curvature - experimental procedure

The curvature of three different actuators was measured using PowerPoint from the Microsoft Office Suite and captures at different pressure inputs. The captures were taken for all three actuation modes: proximal, distal and simultaneous. In each case, the input pressure is gradually increased every 13.8 kPa up to 138 kPa. To measure the curvature in a consistent way for all structures, we draw a red circle following the backbone as seen in Figure C.3 . The latter is formed during the fabrication when the external molds are clamped together. The calibration was done by using a fixed and known distance at rest (the first four knuckles) of 38.1 mm.




Figure C.3: (a) The actuator is at rest, without any pressure applied to neither section.(b) The proximal part of the actuator is pressurized to 110.3 kPa. (c) The distal segment’s input pressure is 138 kPa. (d) The simultaneous actuation shown here is achieved with a pressure of 96.5 kPa in both sections.

C.3 Measuring object motion during grasps

We captured a top-down video of each grasp, then used Tracker Video Analysis and Modeling Software (<https://physlets.org/tracker>) to track contact points, as shown in Fig. C.4. Using these time-series data, we used geometry to compute the position and angle of the object over time.




Figure C.4: Object motion tracking was performed on a grasping test where a marginal grasp occurs. Markers of contact points were estimated by hand from the video.

C.4 Hand Placement Tests

In our exploration of grasping performance as a function of object diameter and hand placement, we used a set of 11 cylindrical objects, as detailed in Table C.1. To give a sense of the actual values in the parameter space that we tested, we have plotted the grasp success and failure regions with the actual points tested. This is shown in Figure C.5

Table C.1: Objects used for grasping characterization. Objects are part of the YCB object set unless marked with a *

Object	Diameter (mm)
#2-56 Bolt*	2.2
1/4-20 Bolt*	6.4
Small Marker	9.8
Syringe*	16.0
Tube Grommet*	25.0
Cup 1 (orange)	46.9
Cup 3 (green)	55.8
Cup 6 (red)	69.4
Cup 8 (blue)	78.8
Cup 10 (yellow)	87.3
Pitcher	116

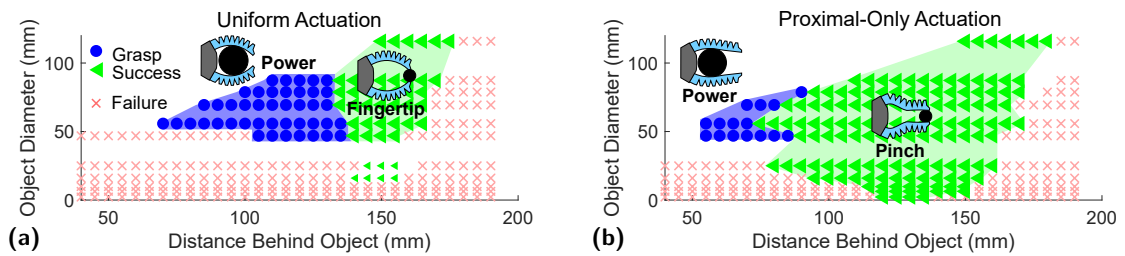


Figure C.5: Hand position and object diameter determine the type of grasp produced upon actuation. The resulting grasp success of each combination of these parameter values that were evaluated are shown. All grasps were performed using fingers with equal-length segments.