Don't Over Think It: Mechanically Intelligent Manipulation

by

Gregory Xie

S.B., Electrical Engineering and Computer Science and Mechanical Engineering, Massachusetts Institute of Technology (2022)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2023

Authored by: Gregory Xie
Department of Electrical Engineering and Computer Science
May 19, 2023

Certified by: Daniela Rus
Andrew (1956) and Erna Viterbi Professor
Thesis Supervisor

Accepted by: Katrina LaCurts
Chair, Master of Engineering Thesis Committee

Don't Over Think It: Mechanically Intelligent Manipulation

by

Gregory Xie

Submitted to the Department of Electrical Engineering and Computer Science
on May 19, 2023, in Partial Fulfillment of the
Requirements for the Degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

Developing capable, robust robots that can effectively operate in unstructured environments requires reasoning over how embodiment and cognition impact the capability and complexity. In this thesis, we focus on the design of hands, finding a balance between capability and complexity through the lens of mechanical intelligence: the ability of the body to contribute to functionality. We develop two grippers that serve as examples of this idea, Belt Orienting Phalanges and Flexible Robust Observant Gripper.

Belt Orienting Phalanges (BOP) enables in-hand manipulation through the addition of two belts on each finger of a parallel-jaw gripper, allowing control over the roll, pitch and a translation of a grasped object. We demonstrate how these motion primitives and other aspects of BOP's morphology enable the simple adaption of an existing planning frameworks to perform a complex manipulation tasks.

Flexible Robust Observant Gripper (FROG) eases perception and control through the structure of each finger, allowing for proprioception and robust grasping while being strong and remaining comparable in complexity to other soft grippers. We demonstrate how these features enable FROG to grasp gently and deal with shape and pose uncertainty.

Thesis Supervisor: Daniela Rus
Title: Andrew (1956) and Erna Viterbi Professor

Acknowledgments

First, I would like to thank my advisor Daniela Rus for helping me grow as a researcher, all the way from the summer of freshman year to now, almost 4 years later. I'm truly inspired by her enthusiasm for and wholehearted belief in the future of robotics.

I'm very lucky to have worked with two amazing mentors, Lilly and Rachel. Together, they have been an amazing source of technical knowledge and life advice. Their wisdom and perspective have really helped shape this work and my research direction. In addition, I would like to thank everyone in the lab : Byungchul, Emily, Joseph, Mieke, Steve, Zach, and many more. With them, no tools remain unfound, no coffee remains undrunk, and no free food remains uneaten. A random shoutout to Russ, who somehow knows my face and name even though I've only interacted with him a single-digit number of times.

Since my earliest memories, my parents, Xuebing Xie and Li Feng, and my sister, Audrey, have been supportive and loving. They have continuously encouraged me to follow my dreams and I am deeply grateful for it. I would like to thank all of Putz, specifically Peter, Sofia, Amy, Nha, and Kat for making me feel at home, despite being far from home. Valerie, Megan, and Zach have somehow made me look forward to going to lectures and doing psets. I'd like to thank the entire MakeX crew for sparking my interest in engineering: Jeremy, James, Nathan, and more.

Finally, I would like to thank my girlfriend Grace for supporting me through this period of my life. Her encouragement while I navigated broken projects, a job search, and the end of this chapter of my academic career has been invaluable.

This work was done with the support of National Science Foundation grant number 1830901 and the Gwangju Institute of Science and Technology.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Robotics is the exploration of how the physical embodiment and cognitive abilities of an agent interact and determine the agent's overall capability. We aim to develop capable robots that can effectively manipulate in unstructured environments. In these environments, the surroundings constrain what the robot is able to do or what it is able to perceive and the task at hand constrains how the robot must interact with the world. While solely improving the embodiment or cognition can lead to more capable robots, we also want robots that are minimally complex to reduce points of failure and increase robustness. To design capable robots while minimizing complexity, roboticists must simultaneously consider how a robot interacts with and perceives the world through its hardware as well as how a robot thinks about the world through its decision making and control algorithms. This is particularly important at the hand, the primary point of interaction between a robot and its environment.

However, the interplay between hardware and algorithm design has not been well explored due to the difficulty in bringing together two disparate fields and the challenge of optimizing over a complex design space [18, 61, 38, 85]. As a result, robots typically use hands that completely add the desired functionality in hardware, such as soft or underactuated hands [23, 35], or end effectors that do not meaningfully reduce the burden on the controller, such as parallel-jaw grippers or anthropomorphic hands [57].

We approach this problem from a hardware perspective, designing embodiments

that when combined with simple control or perception algorithms, maximize the increase in the capabilities of the robot while limiting the increase in overall complexity. The unique ability of the body to contribute to capability is known *mechanical intelligence* [81], *intelligence by mechanics* [10], or *morphological computation* [62] in the robotics community. In cognitive science, this concept is known as *embodied cognition* [69, 84].

In this thesis, we focus on developing mechanically intelligent end effectors that enable new and simplify existing perception and control algorithms. We do so by designing grippers with structured fingers, adding actuation or structure to the fingers. This thesis describes two grippers that serve as examples of this idea: Belt Orienting Phalanges and Flexible Robust Observant Gripper.

Belt Orienting Phalanges (BOP) is a belted parallel-jaw gripper that provides a large in-hand grasp manipulation set and is capable of applying significant forces and torques on the grasped object, seen in Fig.1-1-left. Together, the two fingers' belts can impart forces and torques on to a grasped object to control its roll, pitch and translation. These movements form the basis of motion primitives that can be sequenced together for in-hand manipulation of objects. The hardware of BOP provides a simple relationship between control input and object movement and allows for the use of common grasp sampling heuristics, enabling the use of a well-understood motion planning algorithm for planning in-hand manipulation.

Flexible Robust Observant Gripper (FROG) is an underactuated semi-soft gripper that can grasp objects with minimal shape and pose information, grasp objects delicately and forcefully, and estimate the grasp type, defined by which finger segments are contacting the object. FROG is shown in Fig.1-1-right. FROG's finger and actuator design is similar in complexity to other soft grippers but also provides useful observations and strong priors for perception, allowing for proprioception without additional sensors through the measurement of the actuator stiffness.

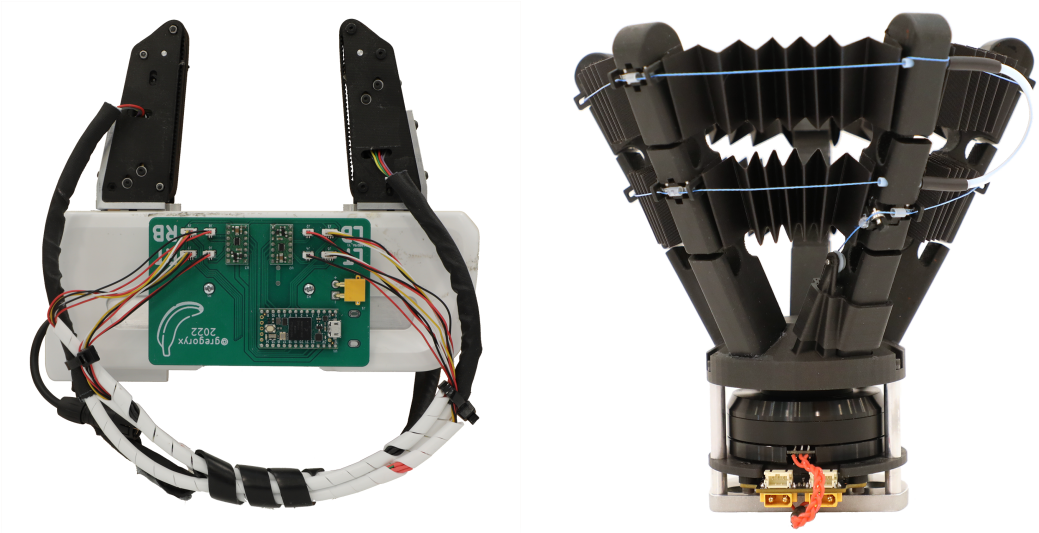

Figure 1-1: (Left) Belt Orienting Phalanges mounted on a Franka Emika Panda hand. (Right) Flexible Robust Observant Gripper

1.1 Contributions

For Belt Orienting Phalanges, we contribute the following:

1. Design and fabrication of Belt Orienting Phalanges (BOP), a belt-driven finger that when paired with a base gripper, allows for in-hand control of object translation, pitch and roll.

2. Characterization of BOP's force and torque limits, as well as reorientation performance under open-loop and closed-loop control.

3. Development of a grasp-to-grasp motion planner and integration into a task and motion planning (TAMP) framework.

4. Demonstration of BOP in accomplishing the real-world manipulation task of screwing a light bulb into a socket as well as executing alternate manipulation skills.

For Flexible Robust Observant Gripper, we contribute the following:

1. Design and fabrication of Flexible Robust Observant Gripper (FROG), a cable-actuated, underactuated semi-soft gripper that can robustly conform to grasped

objects and provides consistent behavior and a strong prior for proprioception and control.

2. Development of two simple feedforward grasping modes to allow FROG to grasp delicately and forcefully.

3. Development of a stiffness-based proprioception algorithm that can estimate the grasp type.

4. Characterization of the strength of FROG's grasps and the performance of FROG's grasp type classifier.

5. Demonstration of FROG's grasping performance on everyday objects and the effectiveness of FROG's soft grasp mode on fragile objects.

1.2 Organization

Chapter 2 gives an overview of end effector design, categorized by joint morphology and actuation methods. We focus on design parameters that are relevant to the design of BOP and FROG.

Chapter 3 introduces Belt Orienting Phalanges (BOP), first describing the hardware, feedback controller, and planner design. We provide an empirical analysis of the reachable grasp set. We then characterize the motion primitives and the forces that BOP can exert. Finally, we demonstrate BOP in a real-world example of screwing a light bulb into a socket and through complex motions such as syringe actuation and fingernail-style lifting.

Chapter 4 introduces Flexible Robust Observant Gripper (FROG), first describing the hardware design, a feedforward grasp force controller, and a stiffness-based grasp classifier. We then characterize the flexures that provide the compliance in our finger design, along with the gripper's holding force and grasp classifier. Finally, we demonstrate manipulation of a variety of objects, including ones that are fragile or soft.

Chapter 5 summarizes the contributions of this thesis and directions for future work.

Chapter 2

Related Work

In this chapter, we give an overview of previous hand designs by exploring two main design parameters: joint/link morphology and actuation. These parameters are the most relevant to the design of BOP and FROG. The classification of end effector design by these design parameters is greatly inspired by [63]. We focus on hands designed for research purposes and do not survey end effectors meant for industrial use.

2.1 Joint and Link Morphology

We first categorize gripper designs by whether the joints and links are rigid or flexible. Joints or links are classified as flexible if the element is intended to comply with contact.

The simplest end effectors to construct have both rigid joints and rigid links - from the perspective of an external force they are essentially immovable. Joint rigidity is typically a result of high-performance position control or non-backdrivable transmissions [7]. This rigidity makes it hard for these end effectors to deal with errors in planning and control - a perfectly position controlled gripper cannot grasp a perfectly rigid object without generating infinite forces. In reality, no controller or object is perfectly stiff and link/object deformation or joint actuator saturation will cause something to comply, intentional or not [57].

The next class of end effectors have flexible joints but rigid links. Joint flexibility can come from mechanically compliant elements (e.g. springs) [23] or from compliant controllers (e.g. stiffness or force control) [21, 7, 67]. In these end effectors, joint deformations allow rigid links to reliably come into contact with external objects. However, the links do not significantly deform after contact, resulting in small contact areas. This morphology allows for a wide range of manipulation capabilities from both simple and complex hands. Anthropomorphic hands with many degrees of freedom (DoFs) have demonstrated impressive in-hand manipulation capabilities, with learned controllers demonstrating finger-gaiting, finger-pivoting, and other dynamic behaviors [2, 30, 17]. Similarly impressive in-hand manipulation results have been accomplished using simple parallel-jaw grippers - using force control to control frictional forces between the gripper and the object while simultaneously using external forces to impart motion [15, 12, 20, 41, 16].

End effectors with rigid joints but flexible links serve as the dual to the previous class of grippers. The compliance of the links serves to insulate the rigid joint from contact, preventing the generation of large forces. Link compliance primarily comes from deformable mechanical elements, such as fingers based off of the finray effect [19]. The links deform significantly on contact with objects, conforming to the surface and providing a large contact area.

Finally, we have a class of end effectors with both flexible joints and flexible links. Because of their low stiffness, these grippers easily adapt to external objects but are limited in the forces they can exert [35]. At the extreme exist continuum soft grippers, where both link and joint are extremely compliant and often cannot be distinguished from each other [37, 65]. There also exist end effectors with both flexible joints and links where there is a clear differentiation between the two [52]. Despite the lack of structure and control when compared the other classes of end effectors, the ability of these grippers to effectively control contact can lead to fine manipulation skills [8, 1].

2.2 Actuation Methods

Just as end effectors have a wide variety in joint morphology, they also have incredible variety in actuation methods.

Like the manipulators they are mounted on, end effectors are commonly actuated using electric motors. Due to the large force requirements on end effector joints, large gear reductions are typically used. The motor is then coupled to the joint through transmission elements: belts [67, 77] and screws [24] are commonly used in end effectors with rigid links while cables [81, 52] are used with both rigid and flexible links. Because gear reductions magnify the apparent motor inertia, elastic (SEA) or variable stiffness (VSA) elements are typically included between the actuator and the joint to reduce impact forces [64, 29, 11, 58, 66]. Quasi-direct drive and direct drive actuators reduce impact forces by reducing the reflected inertia of the motor and also allow for current-based torque control [7].

Soft hands with both flexible links and joints often require more complex actuation methods. Balloon based fluidic actuators, using both positive and negative pressure, are commonly used because their freeform nature allows for easy integration into or around the soft structure [37, 54, 53, 22, 65, 47]. While effort (pressure) control for these actuators is readily available, displacement (volume) measurement and control is not as common [75, 60, 34, 71]. This requires the use of additional sensors if the estimation of deformation or contact is desired. Soft roboticists have also explored using shape memory alloys (SMAs) and dielectric elastomers actuators (DEAs), as they are flexible and easily integrated into flexible bodies. However, these actuators have asymmetric actuation profiles or require highly specialized drivers, limiting their use.

Certain specialized end effectors can also "actuate" the surfaces of their links, typically through modulating their frictional properties [73, 72, 50, 70] or through movement of the surface itself (active surfaces) [77, 40, 28, 87, 86]. This additional actuation is typically used to simplify in-hand manipulation, as control over friction or surface movement can lead to control over object movement.

2.3 Relation to Work

BOP builds off of previous active surface designs to enable simpler in-hand manipulation. Specifically, BOP uses a pair of belt-driven surfaces in each finger, imparting pitch and roll control in addition to translational movement. Compared to the state of the art in in-hand manipulation using high-DOF anthropomorphic hands (flexible joint and rigid link), the use of active surfaces greatly decreases planning complexity by allowing BOP to change grasps while maintaining contact with the object surface [51]. Compared to other active surface grippers, the morphology of BOP enables the use of simple feedback controllers and a well understood sampling-based motion planner instead of handcrafted strategies for specific shapes or learned controllers [87, 86].

FROG finds a middle ground in the class of end effectors that have flexible joints and links, maintaining softness and adaptability while increasing the force that can be applied during a grasp and also allowing for actuator-based proprioception. Compared to the state of the art in strength focused soft grippers (flexible joints and links), FROG's cable actuation method allows for comparable or higher holding forces [47, 27, 76, 46]. The proprioception provided by FROG is of lower quality than proprioception by state of the art sensorized soft grippers, but has the advantage of requiring no additional sensors. [80, 9, 39, 79]. FROG's soft links allow it to create larger contact patches, distributing the grasp contact force and making it easier to grasp fragile objects.

Chapter 3

Belted Orienting Phalanges

Many real-world manipulation tasks require task-specific grasps, where the robot may be required to exert a force through the grasped object or may have geometric constraints on how the object is positioned [48, 31, 49, 42]. For example, to screw a light bulb into a socket, the robot must grasp the bulb such that the bulb's threads are not occluded and such that the robot can exert the torque required to screw the bulb in.

However, in unstructured environments, the robot may not be able to initially select a grasp that fulfills these requirements. In those cases, the robot must transition from an initially reachable grasp to a task-suitable grasp. Regrasping the object by placing it in an intermediate position is one viable strategy to achieve this transition [78], but it is usually significantly slower than performing in-hand object manipulation, especially for dexterous multi-step tasks [57].

Simple and anthropomorphic grippers have demonstrated substantial in-hand object re-orientation capabilities [2, 17, 15, 41], but often require complex algorithms and/or hardware to do so. Meanwhile, it is unclear if grippers specialized for in-hand manipulation are capable of providing the forces needed for dexterous multi-step tasks. We believe a systematic approach that strikes a balance between hardware and algorithmic complexity will be most successful in accomplishing these tasks. We approach this problem from a hardware perspective, designing a gripper that when combined with simple planning algorithm, allows for in-hand manipulation while lim-

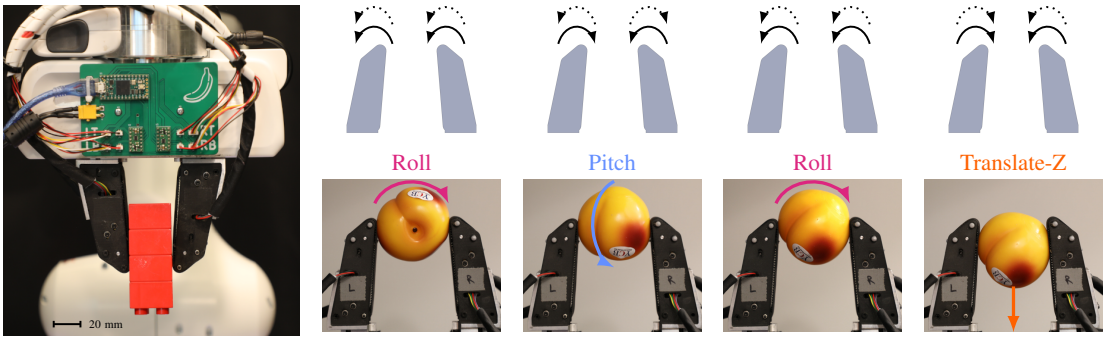

Figure 3-1: The Belt Orienting Phalanges (BOP), mounted on a Franka Emika Panda hand. Each finger has two belts that can move independently from one another. By controlling the direction of the four finger belts, we can translate the object within the hand and control its roll and pitch.

iting the increase in overall complexity.

In this chapter, we present an integrated approach to in-hand manipulation through Belt Orienting Phalanges (BOP), maximizing gripper capability by considering both hardware and algorithm design. BOP is a parallel-jaw gripper where each finger has two sets of belts. As seen in Fig.3-1, through combinations of belt movements, we are able to control the roll, pitch, and translation of the grasped object. We develop simple, proof-of-concept closed-loop controllers for each motion and use them as motion primitives in a geometric planner to generate grasp-to-grasp motion plans. Finally, we demonstrate our gripper and grasp-to-grasp planner in solving a real-world manipulation task by integrating them into a state-of-the-art task and motion planning (TAMP) framework [25].

A video demonstrating the main points of this chapter: the motion primitives, the lightbulb demonstration, and the additional manipulation capabilities can be found at: https://tinyurl.com/bop-rss-main-video

A video showing some of the main limitations of BOP: the necessity of manipulating objects at the fingertips and how the reachable grasp set is limited by geometry object geometry or our planner can be found at: https://tinyurl.com/bop-rss-limits

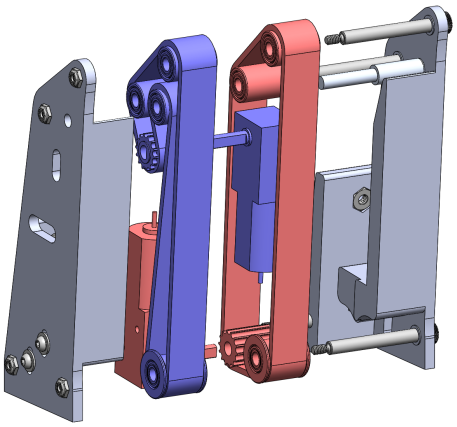

Figure 3-2: Exploded model view of one finger. The blue and red colorations highlight each belt system with associated motor. The drive system for each belt is contained in the finger.

3.1 System Design

To present BOP, we first establish our design goals. We then describe the hardware design, a proof-of-concept closed-loop control strategy and a RRT-based planner that is used to plan grasp-to-grasp motions.

3.1.1 Design Goals

Our design goals for this gripper were to create an end effector that (1) could reach a large set of grasps through in-hand movements (2) had a simple mapping from control input to movement of the grasped object, (3) struck a balance between the amount of mechanical and algorithmic complexity, and (4) was morphologically similar to a traditional parallel-jaw gripper. The first two goals are critical for in-hand manipulation and the third relates to hardware and software robustness. The last goal allows us to leverage off-the-shelf grippers to exert significant grasp forces and existing methods for grasp candidate generation.

Briefly ignoring object geometry, we target the set of grasps that reorient the object in $SO(3)$. We define a frame G which is rigidly attached to the gripper at the center of the line between the grasp contact points and a frame O which is rigidly

attached to the manipulated object (Fig.3-3). Given a starting grasp ${}^{G}X^{O}$ we want to be able to reach all ${}^{G}X^{O'}$ where $\{{}^{G}X^{O'} = R{}^{G}X^{O} | R \in SO(3)\}$. Achieving this requires a minimum of two rotations around non-redundant axes. We choose to target rotations around the Y and Z axes of G, which allows reach of any $SO(3)$ rotation through an extrinsic Z-Y-Z or Y-Z-Y rotation sequence, which follows two popular extrinsic Euler angle conventions. An example of this rotation sequence is shown in Fig. 1, where the gripper performs a roll-pitch-roll (Z-Y-Z) sequence to reorient a sphere to an arbitrary rotation.

With an additional translational degree of freedom, any transform in $SE(3)$ can be reached through a rotation-translation-rotation sequence. We emphasize that this analysis only motivates the controllable degrees of freedom of our gripper and that the actual set of reachable grasps is dependent on the object geometry. The object geometry could reduce the set of reachable grasps if the surface of the object is not sufficiently uniform or if object geometry causes grasps to be in collision. We qualitatively evaluate the size of the reachable grasp set in Sec. 3.2.

As gripper with some degree of force control are common on many manipulator arms [67, 24], we assume access to a base gripper with interchangeable fingers and choose to design active replacement fingers rather than designing an entire end effector assembly.

3.1.2 Hardware Design

Our finger design, as shown in Fig.3-2, has two belts on each finger, allowing for three directions of object movement:

Z-Translation. By driving all belts in a common direction, we are able to translate the grasped object along the Z-axis of G (Fig.3-1-Translate-Z).

Pitch. By driving the belts on one finger in the opposite direction of the belts on the other, we are able to roll the grasped object around an axis parallel to the X-axis of G (Fig.3-1-Roll).

Roll. By driving the belts on each finger differentially, we are able to pitch the object around the Y-axis of G (Fig.3-1-Pitch).

To prevent excessive scrubbing of the object on the belt during rotations, the fingers are slightly tilted inward so that the tips of the fingers are closer together than the base. This allows us to assume that during normal operation, only the tips of the fingers contact the grasped object. If an object contacts the tilted sides and not the fingertips, loss of belt contact may cause loss of control over the pitch rotation. In scenarios where in-hand manipulation is not needed and a stronger grasp is necessary, tilting the fingers in this way allows us to achieve a three-point enveloping power grasp when the grasped object contacts the sides of the fingers and the palm.

A stationary surface is located behind the belts so that, when grasping using the middle of the fingers, the grasp force does not have to be completely supported by belt tension. Low friction PTFE tape is used to minimize friction between the belts and the stationary surface.

We size a gearbox and motor combination (1030:1, N20) so that, at a nominal contact normal force of 40N, the belts slip on the object before the motor stalls. We design for a worst-case scenario, with an unrealistic friction coefficient of 1. We use non-backdrivable worm gearboxes so that the fingers do not have to be actively controlled when in-hand movements are not desired.

Encoders on the motors allow for closed loop velocity and position control. A PCB carries the motor drivers (TI DRV8835), a microcontroller (PJRC Teensy 4.0), and an interface to the host computer through a USB-serial connection. The microcontroller can run either a belt velocity or belt position proportional-integral controller. Including the PCB and wiring to the fingers, but excluding mounting adapters to the base gripper and other wiring, a system with two fingers weighs approximately 230 grams. Each finger is approximately 38x86x28 mm.

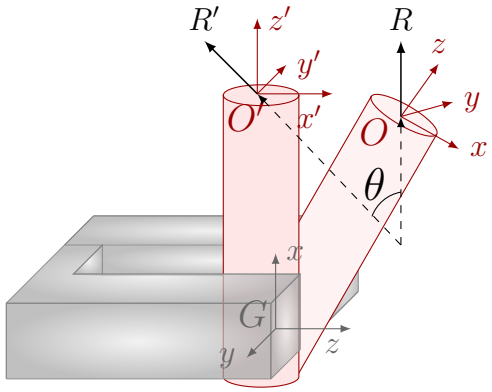

Figure 3-3: The motion primitive controllers use an approximation for the change in object pitch θ. G is the gripper frame, while O is the object frame. As the object moves from start grasp to goal grasp, we estimate θ as the angle between R, the projected ray at the initial grasp onto the XZ plane of G and R', a similarly projected ray from the current grasp onto G. The change in roll is estimated similarly, with projections going onto the XY plane of G. This approximation is only valid when the change in either pitch or roll is small.

3.1.3 Controller Design

We design proof-of-concept closed loop controllers for each of the motion primitives: Z-translation, pitch and roll. The controllers operate on an estimate of the object to grasp transform, ${}^{G}\hat{X}^{O}$. We use a motion capture system (Optitrack) to provide pose feedback. We make no assumptions about where the object frame O is on the object or how it is oriented - only that it is rigidly attached to the object. This allows the controllers to generalize to scenarios where we have an estimate of object pose but no object model.

We use proportional controllers to control roll, pitch and Z-translation of the grasped object through a velocity control interface. The error of each of these controllers is not straightforward to define because directly converting common rotation representations to Euler angles can lead to undesirable behaviors, including discontinuity around singularities. To work around this issue, we make the following approximation to estimate the change in the pitch and roll.

As visualized in Fig.3-3, to estimate the change in pitch, we attach a unit ray R to the object frame, which initially points in the X direction of G, and then project this ray into the XZ plane of G. We estimate the change in pitch (θ) as the angle

between the projected ray at the initial grasp (R) and the projected ray from the current grasp (R'), while tracking full rotations. To estimate the change in roll, we perform the same procedure, but we project into the XY plane instead.

We find this approximation holds when the change in either pitch or roll is small, which, in our case, is generally true because we are interested in controlling motion along one primitive while maintaining zero movement along the others. Further supporting the use of this approximation, we notice that the belt movements corresponding to one of the rotation primitives does not significantly affect the other, and vice versa. This means that during a pitch or roll primitive, we typically do not need to control the other.

To estimate the change in Z-translation, we need to account for how O moves with respect to G. Since O can be located at an arbitrary location on the object, its position relative to G will change depending on how the object has been rotated. Thus, we subtract the effects of roll and pitch before computing the difference between the current object frame and the initial object frame along the Z axis of G.

In our experiments, we notice that Z-translation movements were linked to significant changes in pitch and slight changes in roll. We compensate for this by running multiple controllers simultaneously during a motion primitive. We run all three controllers during the execution of a Z-translation primitive, and run the Z-translation controller with either the roll (or pitch) controller during a roll (or pitch) primitive. We stop the controllers and end the motion primitive if the errors for each of the active controllers are less than user-defined thresholds or the motion takes longer than a user-defined timeout. We evaluate the accuracy of open versus closed loop control in Sec. 3.3.

3.1.4 In-Hand Motion Planning

To plan in-hand motions that move the manipulated object from a start grasp to a goal grasp, we use bidirectional RRT-Connect [45, 44], assuming a known object model. We are able to easily use RRT with few modifications because of the relationship between control input and object movement enabled by our closed-loop motion

Algorithm 1 Sample(obj)

while True **do**
    $p_1, n_1 = \text{SamplePointOnMesh}(obj)$
    $p_2, n_2 = \text{RayCast}(obj, p_1, n_1)$
    $s = \text{ToState}((p_1 + p_2)/2, n_1)$
    if $\text{IsGraspValid}(obj, s, n_1, n_2)$ **then**
        return s

Algorithm 2 Extend(T, s_{near}, s_{target})

candidates = []
for $p \in$ primitives **do**
    for $t \in$ steps **do**
        $s_{candidate} = \text{SimulatePrimitive}(p,\ t,\ s_{near})$
        candidates.Add($s_{candidate}$)
$s_{new} = \text{BestState}(s_{target},$ candidates)
T.AddVertex(s_{new})
T.AddEdge(s_{near}, s_{new})
return s_{new}

primitives.

Our state is the grasp transform represented as a translation and quaternion: $s = (x, y, z, q_x, q_y, q_z, q_w)$. In the standard bidirectional RRT-Connect algorithm, we replace the `Sample()` and `Extend()` functions, given in Algorithm 1 and Algorithm 2, respectively.

To generate random states (Algorithm 1), we rejection-sample grasps using a heuristic from Tedrake [74]. We sample a random point (p_1) on the object's mesh and ray-cast along its mesh normal (n_1) to find the point and its associated normal (p_2, n_2) on the mesh, opposite from the sampled point. We create a grasp transform with the y-axis of the frame aligned with n_1 at the average of p_1 and p_2. We reject grasps if the grasp contact normals are not well aligned, if the grasp width is too large, or the grasp is in collision. Because the space of valid grasps is much smaller than the space of rigid body transforms, we choose to sacrifice probabilistic completeness and only sample grasps using our heuristic.

The `Extend()` (Algorithm 2) uses the motion primitives described in the previous section. For each primitive, we simulate a discrete number of motions in the positive

and negative directions, incrementing the translation or angle by a fixed step size. The results of each simulation are cached to minimize the number of collision checks. To select the best primitive and step, we measure the distance between two states as the weighted L^2 norm between them, with weights chosen so that an error of 10 cm is roughly equivalent to a rotation error of π rad. In the planner, we threshold this distance and consider two states to be equal if the distance is below a user-specified threshold.

3.2 Empirical Reachability Analysis

We expand on what in-hand manipulations are feasible for BOP to perform with real-world objects, focusing on geometric limitations to complement to the force/torque limitations discussed later in Sec. 3.3.1.

BOP can perform a translation and two extrinsic rotations around perpendicular axes (roll and pitch). Without consideration for object geometry or mass properties, any grasp should be reachable from any other grasp. However, the object's geometry may cause object-gripper collisions or cause the grasp set to be separated into multiple modes, both of which reduce the reachable grasp set. These limitations are not unique to our gripper and are shared with any active surface gripper which performs in-hand manipulation without making or breaking contact. In addition, the grasp sampling heuristic (Algorithm 1) and grasp validity criteria (`IsGraspValid()`) can further restrict the grasp set reachable using our grasp-to-grasp planner.

3.2.1 Collisions between Object and BOP

Collisions between the object and gripper can prevent arbitrary rotations or translations. In Fig.3-4, the gripper initially has control over the Z-translation and pitch of the object. Fig.3-4-left shows loss of control over Z-translation because the width of the object suddenly exceeds the width of BOP's grasp on the object, preventing further translation. Fig.3-4-right shows loss of control over pitch because the top of the object collides with the fingers, preventing further rotation.

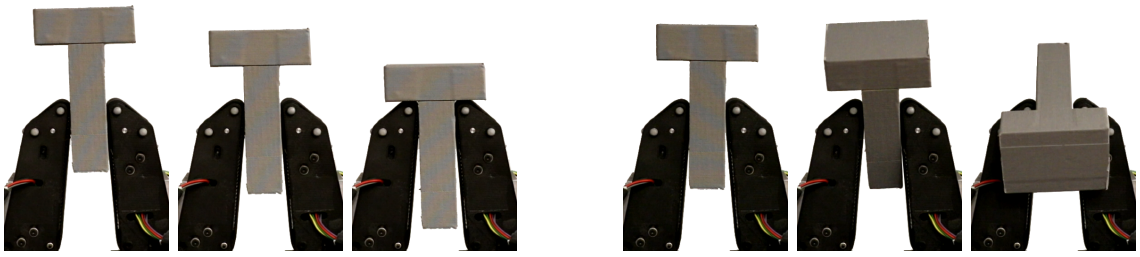

Figure 3-4: Demonstration of how object geometry prevents arbitrary manipulation. The T-shape's horizontal beam prevents arbitrary Z-translation (left) and arbitrary pitch rotation (right) as it collides with the fingers. This motion was run without vision feedback.

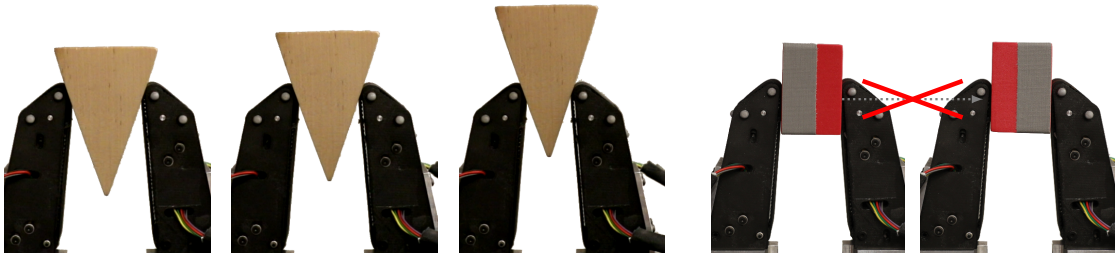

Figure 3-5: (Left) Z-translation of a pyramid. By our grasp validity criteria, no grasp on the pyramid is "valid". However, BOP is still able to manipulate the pyramid. This motion was run without vision feedback. (Right) Display of two valid grasps that can not be manipulated between. Since the grasps belong to different modes, BOP can not go from one grasp to the other.

3.2.2 Overly Conservative Checks on Grasp Validity

The size of the grasp set reachable through our grasp-to-grasp planner is artificially limited by our grasp validity criteria (`IsGraspValid()`). Our criteria rejects grasps if the grasp contact normals are not well aligned, if the grasp width is too large, or the grasp is in collision. For example, depending on the value used for the grasp normal tolerance, the entire grasp set on a square pyramid may be rejected. This makes it impossible to use the grasp-to-grasp planner to find a motion plan, as there are no valid grasps on the object at all. However, as shown in Fig.3-5-left, BOP is still able to successfully grasp and manipulate the pyramid.

This issue could be resolved with more sophisticated grasp validity criteria that would more accurately separate valid grasps from invalid grasps.

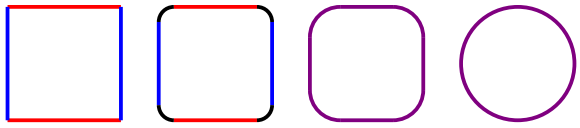

Figure 3-6: Modes on an object, shown in red and blue, are separated by areas of high curvature, shown in black. BOP is unable to manipulate between grasps from different modes (ex. between a top/bottom grasp and a side grasp). This is because the gripper is unlikely to maintain contact with the object while traversing over areas of high curvature. It is ambiguous which mode these high curvature areas belong to, if any. As the curvature is decreased, the two modes combine into one, shown in purple. This allows BOP to manipulate between a top/bottom grasp and a side grasp. The critical curvature depends on the frictional properties between the object and the finger.

3.2.3 Objects with Multi-Modal Grasps

Borrowing the terminology from multimodal motion planning [32], one limitation of our gripper is that we can only accomplish in-hand manipulations within a single mode. We define a mode as the set of object surface-gripper finger contacts. Surfaces are separated from each other along areas of high curvature. In order to perform a mode switch, BOP would need to traverse an area of high curvature. However, we do not allow for this behavior because the fingers are unlikely to maintain contact with the object. Fig.3-6 shows how curvature impacts the modes. As the curvature is reduced, the initially separate modes combine into one. The precise curvature that the fingers can travel over and at which the modes combine will depend on the frictional properties between the object and finger.

As an example, Fig.3-5-right shows two valid grasps on the object that constitute different modes. Since our gripper cannot switch between the two modes, we cannot manipulate the object from the grasp on the left-hand side to the right-hand side.

In addition to mode separation due to object geometry, our grasp-to-grasp planner may artificially split a single mode due to how we sample random grasps. An example of this is shown in Fig.3-7, where the indentation on the side of the cylindrical object causes our sampling method to split the grasps on the side of the cylinder into two modes. However, Fig.3-7-right shows a Z-translation primitive successfully manipulating the object between the two artificially separated modes, showing that

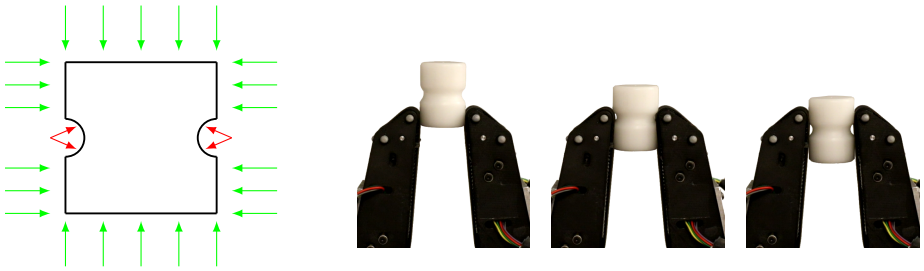

Figure 3-7: Artificial Mode Separation. The side mode is artificially separated into two modes at the indentation due to our sampling method, rendering it hard or impossible for our grasp-to-grasp planner to find a in-hand motion that traverses over the indentation (Left). However, successful Z-translation over the indentation shows that the side mode is actually one mode (Right). This movement was run without vision feedback.

they are actually belong to the same mode.

This issue could be resolved by improving our grasp validity criteria, as well as the addition of uniform sampling to sample grasps missed by our heuristic. Additionally this would restore the probabilistic completeness guarantee of our planner.

3.3 Performance Characterization

We evaluate BOP's performance with two experiments: (1) quantifying the force and torque that can be exerted by the motion primitives and (2) quantifying the accuracy and repeatability of the motion primitives under closed-loop control. These experiments help define the set of objects that can be manipulated in-hand and confirm that several primitives can be accurately sequenced together for grasp-to-grasp planning.

For both experiments, we mount the fingers to a backdrivable test platform driven by a torque-controlled BLDC (mjbots mj5208, mjbots moteus r4.5). In our experiments, we define the grasp force to be two times the contact force at one of the gripper fingers.

3.3.1 Force and Torque

To measure the forces and torques from BOP, a force/torque sensor (ATI gamma SI-32-2.5) is attached to a Franka Emika Panda manipulator. We command zero stiffness

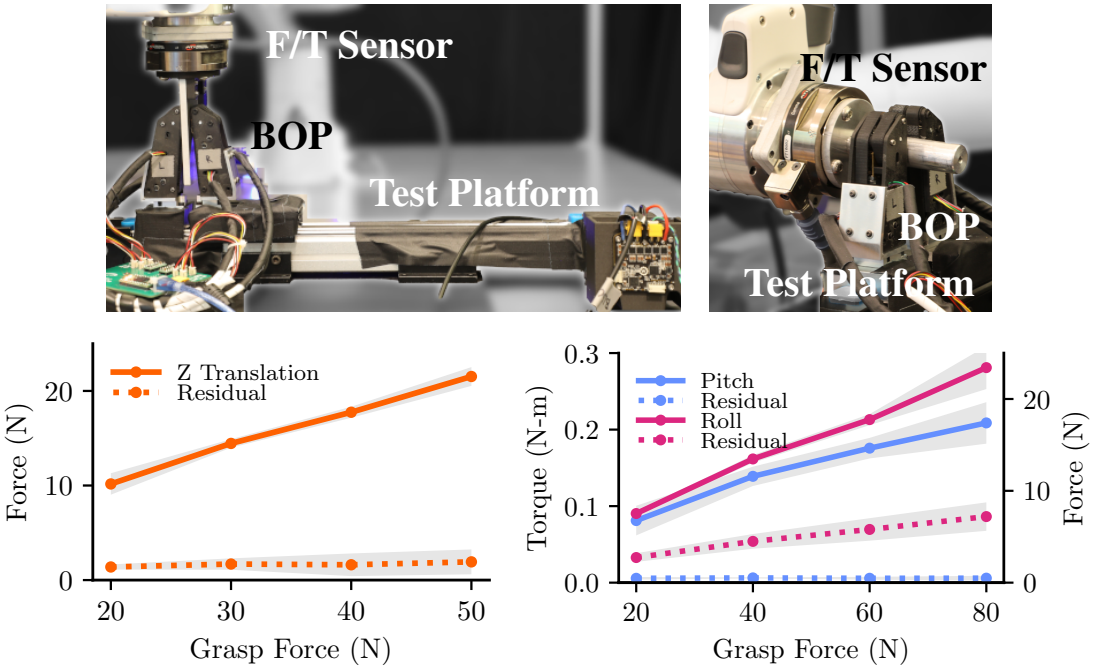

Figure 3-8: Force and Torque Characterization for the three motion primitives. (Far left) The setup used to characterize the force exerted by the Z-translation primitive and the torque exerted by the pitch primitive. (Middle Left) The setup used to characterize the torque exerted by the roll primitive. (Right) Both graphs show the result of 20 trials for each motion primitive swept over a range of grasp forces. Each point is found by averaging across 10 trials taken in the positive direction and 10 trials taken in the negative direction to eliminate forces or torques which result from closing the gripper. We also report a grasp force residual, which is the difference in force that each finger is applying due to non-ideal robot stiffness control. The standard deviation is shown in grey.

in all of the Cartesian axes except for the axis corresponding to the motion primitive being evaluated. To test the force and torque exerted by the Z-translation and pitch primitives, we mount a 6.35 mm flat plate at the XZ plane of the force/torque sensor reference frame (Fig.3-8-far left). To test the torque exerted by the roll primitive, we mount a 19.05 mm cylinder along the Z axis of the force/torque sensor reference frame (Fig.3-8-middle left).

For each trial, we close the test platform on the test plate (or cylinder) with the specified grasp force and command the finger belt velocity to the corresponding primitive. We report the resulting force or torque across 20 trials, 10 in the positive

direction and 10 in the negative direction. Although the robot is commanded to exhibit zero stiffness, joint friction and other effects cause an unequal force to be exerted by each finger, which we report as a grasp force residual.

As shown in Fig.3-8-right, all primitives have a roughly linear relationship with grasp force. We see that the Z-translation primitive can apply a maximum force of approximately 18 N at a grasp force of 40 N. Assuming a conservative friction coefficient (0.23, about half of what was observed in this experiment) and that the object geometry is such that a valid grasp exists, the gripper should be able to manipulate all but four of the objects in the Yale-CMU-Berkeley (YCB) dataset [14] using this primitive.

At the same grasp force, the pitch primitive can apply a maximum torque of approximately 0.14 Nm. Assuming the belts contact the object at a point, we estimate an effective contact-to-contact distance of 7 mm. With a center of mass to grasp distance of 5 cm, the same conservative friction coefficient, and a grasp force of 40 N, the gripper should be able to use its pitch primitive to manipulate objects that weigh less than 230 grams in any orientation. The large difference in allowable object weight between the Z-translation and pitch primitives is due to the small contact-to-contact distance. This effect is shared with any gripper with a small contact patch.

At a grasp force of 40 N, the roll primitive is able to exert a maximum torque of 0.16 Nm on the test cylinder. The maximum torque that the roll primitive can exert should be linearly related to grasp width, which we do not characterize.

Overall, this experiment shows that BOP is able to exert significant forces and torques on the grasped object, enabling its use in forceful manipulation tasks [33].

3.3.2 Open vs. Closed-Loop Primitives

We test the accuracy and consistency of the closed-loop control of our motion primitives on 3 foam objects: a cube (50x50x50 mm), a cylinder (50x50 mm), and a sphere (50 mm). We test our Z-translation and pitch primitives on the cube, all three primitives on the cylinder, and both rotation primitives on the sphere. Because the foam is soft, we use a grasp force of 20 N. As seen in Fig.3-9-left, we place the object

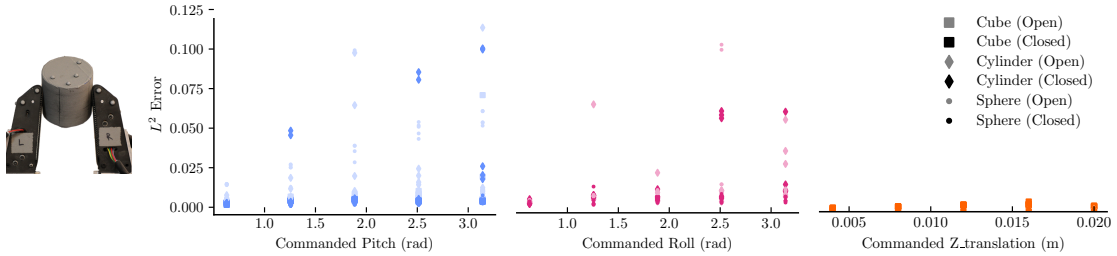

Figure 3-9: Open-Loop Control vs Control-Loop Control for the three motion primitives. For each experiment, a test object is placed in the gripper (Far Left) and a motion primitive is executed. We report the error between the expected and end pose of the object across five trials for each object, using the distance metric described at the end of Sec. 3.1.4. The error for open-loop controller is given in the lighter color and the error for the closed-loop control in the darker color. For the pitch primitive we test a cube, a cylinder and a sphere. For the roll primitive we test a cylinder and a sphere. For Z-translation primitive we test a cube and a cylinder.

in the gripper at a nominal position and command an open or closed-loop motion primitive. The rotation primitives are commanded to a maximum of π radians and the Z-translation primitive is commanded to a maximum of 20 mm. In Fig.3-9 we report the distance between the expected pose and the measured pose for five trials per object, using the distance metric described at the end of Sec. 3.1.4.

We observe that while closed-loop control is generally more accurate than open-loop control for the pitch and roll primitives, both primitives and control paradigms are sensitive to object geometry and material. For example, with the pitch primitive on the cylinder, variance in the initial object placement can cause only one of the belts on a finger to contact the object. This leads to loss of control over pitch, causing the cylinder to instead roll until the object is jammed in between the fingers and cannot move or translate until the gripper loses contact with the object. With the roll primitive, we found our foam objects are sensitive to perturbations in yaw, resulting in similar failure modes where the object is jammed in between the fingers. In Sec. 5 we discuss how the closed-loop control could be improved to address these issues.

For the Z-translation primitive, the error, as seen in Fig.3-9-right, is significantly lower than the error for the two rotation primitives. Examining the data more closely,

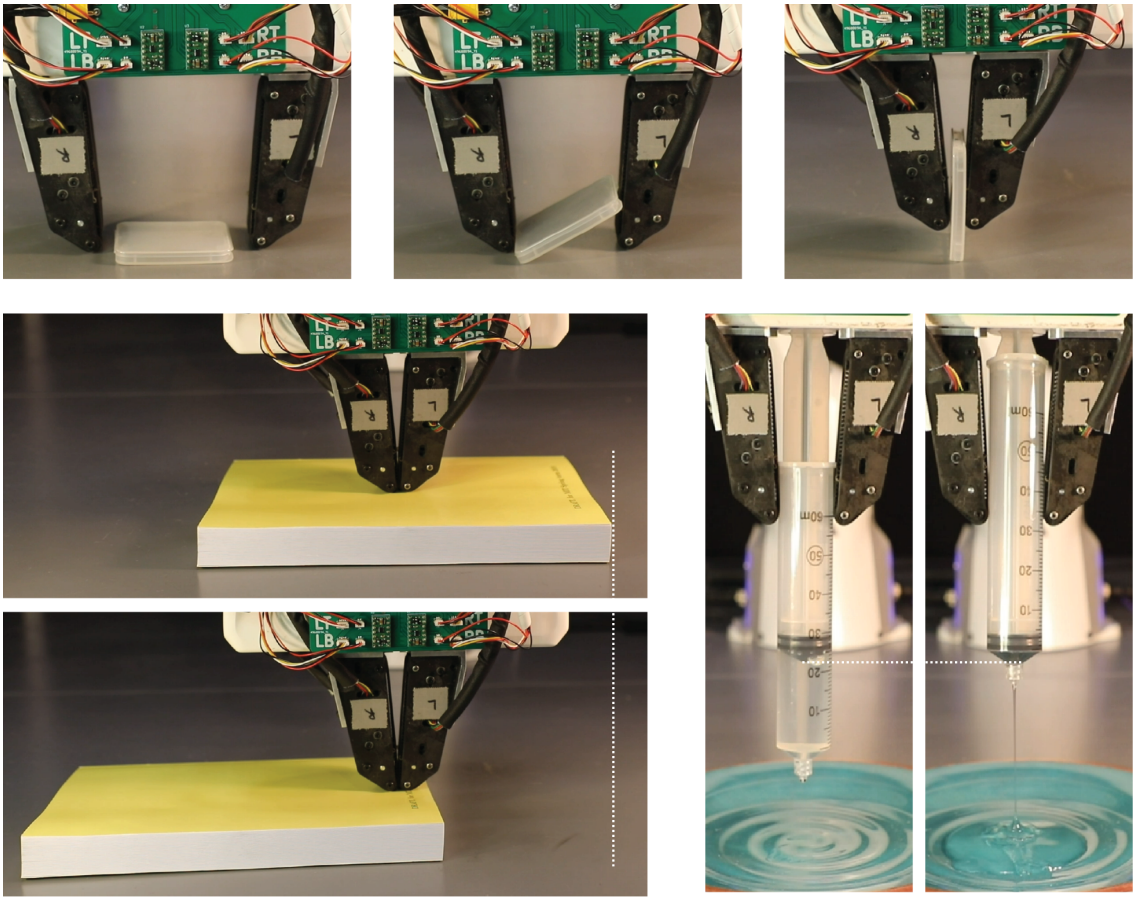

Figure 3-10: BOP's design allow it to achieve unique motions such as (top) fingernail-style picking of a thin object off of a table, (bottom-left) scooting a book along a table, and (bottom-right) dispensing liquid from a syringe. The dotted grey line is used to emphasize the object movement.

the open-loop control performs equal to or better than the closed-loop control with the two experimental objects. This is because the closed-loop control terminates at a user-defined threshold, limiting its accuracy. Qualitatively, we observe that the closed-loop control is more robust to external disturbances and suspect it would also outperform the open-loop control for objects with less idealized geometry.

Overall, experimental results support that, with closed-loop control, our motion primitives are accurate enough to be reliably sequenced into a multi-step plan by our grasp-to-grasp motion planner.

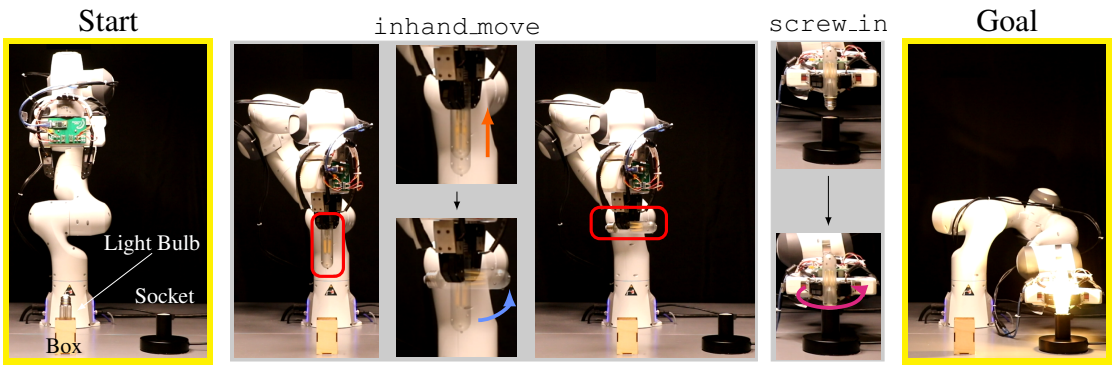

Figure 3-11: We use a TAMP framework, augmented with actions enabled by BOP screw a light bulb into a socket. The box the light bulb starts in constrains the set of available grasps, forcing the robot to manipulate the light bulb in-hand. To complete the task, the robot executes a plan of `move`, `pick`, `inhand_move` (specifically, a z-translation, then a pitch rotation), `move_holding`, `screw_in`. Here we highlight the `inhand_move` and `screw_in` actions.

3.4 Demonstrations

We show how a robot arm with BOP can complete a variety of dexterous manipulation tasks. We first demonstrate how the active surfaces of BOP allow for various other manipulation capabilities beyond our three motion primitives. We script three such example capabilities. We also augment an existing task and motion planning (TAMP) framework to leverage BOP's strengths of in-hand manipulation and force exertion to complete a real-world, multi-step manipulation task: screwing a light bulb into a socket. Please see the supplemental material for a video of the additional grasping capabilities and the light bulb demo.

3.4.1 Additional Manipulation Capabilities

The first manipulation capability, shown in Fig.3-10-top, is inspired by Watanabe et al. [82] and uses a nail-lifting strategy to grasp a thin object. The robot first grasps the object with a weak pinch grasp. The gripper next closes while the left finger belts are run in the direction of the robot's palm, lifting that side of the object up. Once the object is near upright, the hand moves upward while closing the gripper, allowing the robot to more securely grasp the object.

Inspired by the manipulation task taxonomy from Bullock et al. [13], we demonstrate nonprehensile, surface translation and in-hand, syringe actuation in Fig.3-10-bottom left and Fig.3-10-bottom right, respectively.

Surface translation begins with a closed gripper and the robot applying a downward normal force to the surface of the book. Executing the roll primitive, which moves the left finger belts towards the palm and the right finger belts away from the palm, translates the book to the right. Similar to pushing, this nonprensile manipulation could be used to manipulate objects that are too thin or large to be grasped.

Actuating a syringe begins with the syringe grasped in-hand, with the plunger of the syringe in-contact with the palm. Executing the Z-translation primitive moves the body of the syringe towards the palm, thus ejecting the material. This capability differs from the Z-translation primitive introduced in Sec. 3.1, since there is palm-object contact.

These examples serve as additional capabilities that could be sequenced in a multi-step manipulation task, in addition to the in-hand manipulation described in previous sections. Many more manipulation capabilities can be imagined, including the use of the Z-translation primitive to bring objects from a pinch grasp to the three-point enveloping power grasp described in Sec. 3.1.2 or the use of the roll primitive to impart a screwing motion on objects.

3.4.2 Demonstration in a Multi-Step Manipulation Task

We demonstrate this gripper in context of a manipulation task, specifically screwing a light bulb into a socket. As shown in Fig.3-11, BOP is mounted on a Franka Emika Panda and on a table there is a light bulb in a box and a light bulb socket. Since the light bulb is in the box with the threads facing upward, the only reachable grasps occlude the threads, preventing the light bulb from being screwed into the socket. Thus the robot will need to change its grasp on the bulb before inserting the bulb. Additionally, screwing in the light bulb requires many turns, which BOP can do by exerting force while applying a continuous rotation.

To plan the sequence of actions needed to install the light bulb, we cast this as a

task and motion planning (TAMP) problem. Solving a TAMP problem corresponds to selecting the sequence of actions for the robot to execute and selecting the discrete and continuous values, such as grasps, paths, objects, etc., that parameterize those actions [26]. We specifically use the TAMP framework PDDLStream, which reduces this type of hybrid discrete/continuous planning problem to a sequence of discrete planning problems via focused sampling of the continuous and discrete parameters [25]. We assume that we are given geometric models of the robot, the objects and the environment along with the poses of each object. We use a motion capture system to track the pose of the manipulated object, the light bulb, for the closed-loop control of BOP's motion primitives.

Aside from a set of generic manipulator actions (`move`, `move_holding`, `pick`, `place`) we define two custom actions: `inhand_move` and `screw_in`. The `inhand_move` action uses the motion planner defined in Sec. 3.1.4 to plan in-hand grasp-to-grasp motion for a held object. The `screw_in` action first uses a guarded move [83] to insert the bulb into the socket and then rotates the bulb by a fixed amount using the roll primitive. Please see the supplemental material for the full domain specification.

Fig.3-11 shows a plan of: `move`, `pick`, `inhand_move`, `move_holding`, `screw_in`, where the `inhand_move` and `screw_in` actions are emphasized. In this plan, the in-hand manipulation involves the Z-translation primitive followed by the pitch primitive. We believe this task demonstrates BOP's ability to perform in-hand manipulations while simultaneously applying significant forces on the object.

Chapter 4

Flexible Robust Observant Gripper

Many real-world manipulation tasks require the robot to operate in unstructured environments, under uncertain conditions where perfect knowledge of the environment is not obtainable. In these environments, the robot must gracefully deal with uncertainty in motion and estimation while exerting significant forces through the grasped object. [57, 48].

The high stiffness of traditional rigid end effectors allow them to exert the forces required in these tasks, but the same stiffness necessitates the use of controllers to generate the compliance needed to safely deal with unexpected contact [57]. On the other hand, soft grippers are able to easily deal with uncertainty but are unable to exert large grasp forces due to their low stiffness [35].

We approach this problem through the design of Flexible Robust Observant Gripper (FROG), a semi-soft underactuated gripper that is both compliant and strong, allowing it to consistently grasp objects with minimal shape and pose information. We design hardware that is of comparable complexity to other soft grippers, but additionally allows for over the grasp force and proprioception of the grasp type through the addition of simple control and sensing algorithms. We develop a feedforward grasp force controller and a classifier to distinguish between grasp types, defined by which finger segments are contacting the object, using the stiffness of the cable. We characterize the dynamic properties of the joints used in each finger and holding force that the gripper can exert. Finally, we demonstrate our gripper by manipulating a

Figure 4-1: Flexible Robust Observant Gripper (FROG) mounted on a Universal Robots UR5 (left). FROG is simultaneously strong and compliant, shown lifting a 10 lb (4.5 kg) dumbbell (middle) and being deformed by a plastic bat (right).

wide variety of items, including fragile and soft objects.

4.1 System Design

To present FROG, we describe the hardware design, a feedforward grasp force controller, and a grasp type classifier. We then characterize the flexures used in FROG's fingers.

4.1.1 Hardware Design

Our gripper design has 5 semi-soft fingers, driven by a single tension-controlled cable. As seen in Fig.4-2-left, each finger has two flexure joints and three links, and is 3D printed in one piece (Carbon EPU 40). The entire finger is relatively soft, with a Shore hardness of 68A. Because the flexure joint stiffness is low compared to the link stiffness, deformations are mostly localized at the joints. Backstops on the flexures allow for significant force application in grasps where contact occurs on the distal links by increasing the stiffness of the finger, as illustrated in Fig.4-2-right. The flexures additionally allow for two secondary motions, side-to-side (adduction/abduction) and

twisting, enabling the fingers to conform to surfaces that are not normal to the nominal link surface. The fingers are glued into a 3D printed base (Markforged Onyx) with two-part epoxy.

As shown in Fig.4-2-left, the closure of the fingers is actuated by a single cable, routed around the back of the fingers through roller cable guides. Two rollers (682ZZ) on dowel pins and a teflon guide tube are glued into a 3D printed housing (Carbon UMA 90) using cyanoacrylate glue. Cyanoacrylate glue is also used to glue the cable guide assembly into the fingers. Features on the housing help constrain the cable and prevent the cable from falling off the roller if it becomes slack.

The cable is routed through the proximal flexure's center of rotation on its path to the spool, located in the palm of the gripper, to minimize additional torque on the joint. Minimal friction in the cable drive is desired as the return of the fingers to their neutral positions is driven by the flexures. The only sliding friction in the system is the teflon bowden tube connecting the top and bottom cable runs. To further minimize friction a mechanism similar to the one introduced in [43] could be used. The cable guides were designed for a cable tension of 10 N. Polyester bellows between the fingers prevent grasped objects from directly contacting the cables. The bellows are attached to the fingers with double sided tape.

At the same cable tension, this routing method enables FROG to exert larger grasp forces than traditional cable actuated grippers. This is because the distance between the joint and the cable is significantly larger, resulting in a larger joint torque at the same cable tension. In addition, each time the cable traverses a finger it applies two tension forces. In configurations where the fingers are close to radially symmetric, this results in a inward force of more than 1.5x the cable tension.

The cable is pulled by a BLDC gimbal motor (T-motor GB54-1, mjbots moteus r4.11), chosen for its low cogging torque. From [7], we expect the motor to be able to continuously dissipate 10 W, which results in a continuous torque of 0.23 N-m. The spool has a diameter of 23 mm, and is sized so that the motor can pull the cable to the design tension of 10 N continuously without overheating. The spool is designed to ensure subsequent wraps of the tendon completely overlap with previous

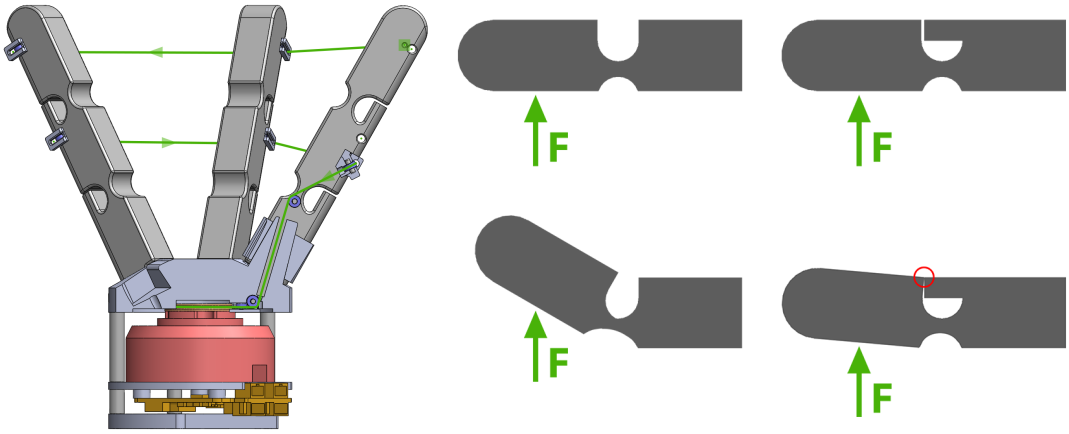

Figure 4-2: (Left) Section view of FROG. The motor is shown in red, motor driver in yellow, spool in orange, and a section of cable in green. Arrows on the cable path show the direction of travel when closing the gripper. A square marks the location at which the cable is attached to the finger. (Right) Diagram demonstrating how the addition of a backstop increases the stiffness of a finger, allowing for more forceful grasps. Contact with the backstop is circled in red.

ones, allowing us to compensate for changes in spool diameter as the cable is wrapped. The motor driver's stock current sense resistors (0.5 mΩ) are replaced with larger ones (5 mΩ) for better current measurements, as the typical current through the motor is less than 2 A.

As seen from the actuator, the effective stiffness of the proximal joint is lower than the distal joint, which results in significant movement of the distal joint only after the proximal link has made contact.

Overall, the gripper is 180 x 170 mm, with a ideal maximum grasp diameter of 130 mm. Because of friction, typically a maximum grasp diameter of only 110 mm is achievable. The bellows and finger thickness limit the minimum grasp diameter to 38 mm.

4.1.2 Grasp Mode Design

We develop controllers for two grasp modes: a soft grasp and a hard grasp. FROG can delicately grasp fragile objects using the soft grasp mode or forcefully grasp rigid objects using the hard grasp mode. We leave the choice of which grasp mode to

use for a particular object to the rest of the robot system. These two grasp modes demonstrate how FROG's mechanical repeatability and actuation method allows for control over grasp behavior, future work could expand on these behaviors. Both grasp modes spool the cable at the same velocity (3.1 rad/s), the controller determines the maximum tension the motor is allowed to exert at each cable position. Because FROG does not have tactile sensors and cannot measure the displacement of the flexures, the controller for each grasp mode is purely feedforward.

Let $f(p)$ be the maximum tension force that the motor is allowed to exert at cable position p. The controller for the hard grasp mode is simple: we command the motor to spool the cable with $f(p) = f_{max} = 10$ N. This exerts the largest grasp forces, resulting in the strongest grasp possible.

For the soft grasp mode, we leverage the physical consistency of our finger design and fine control over gripper actuation effort to do feedforward force control for low forces. We assume that all movement stops after contact, radial symmetry, that only the proximal joint has significant movement, and that the only friction on the cable is from the bowden tube with a friction coefficient of 0.1. Define $f_{nc}(p)$ to be the tension force needed to continue closing the gripper at a set cable position p, and define ϵ to be the increase in tension over $f_{nc}(p)$. For grasps where the object contacts the middle of each proximal link we expect a contact force of 4.7ϵ, for grasps where the object contacts the middle of each distal link we expect a contact force of 1.6ϵ. We do not measure the contact force during the grasp. However, despite the range of possible contact forces, we demonstrate in Sec. 4.3.2 that FROG is still able to grasp a variety of fragile objects using the soft grasp mode.

We measure $f_{nc}(p)$, shown in Fig.4-3-left, by closing the gripper with no object contact. During a soft grasp, we command the motor to spool the cable with $f(p) = f_{nc}(p) + c + \epsilon$, allowing the user to specify a small ϵ. c is an offset to ensure the gripper consistently closes in all orientations, we find that $c = 0.3$ works well.

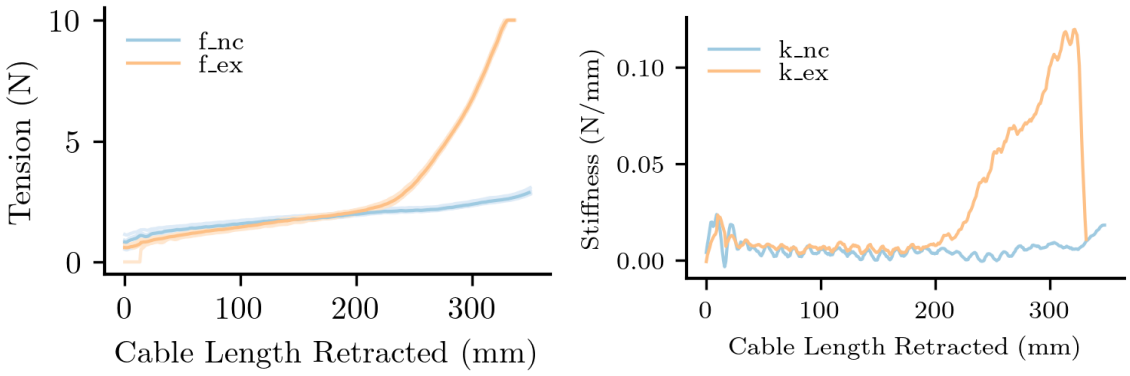

Figure 4-3: Cable tension (left) and stiffness (right) in a grasp with no contact and in an example grasp with both proximal and distal contact, shown in blue and orange respectively. For the tension, 10 trials were taken and each trial is shown in a shaded line. The average is shown in a darker line. The stiffness is obtained by differentiating the average tension to better demonstrate how the stiffness changes during the grasp. As the fingers make contact with the object, the stiffness increases and the force rapidly rises. This allows for estimation of the type of grasp through the stiffness. Because of viscoelastic effects, the fingers continue to move after the maximum tension has been reached, decreasing the stiffness at the end of the grasp.

4.1.3 Grasp Type Proprioception

As the links of an underactuated finger comes into contact with an object, the stiffness seen from its actuator changes [4], seen in Fig.4-3-right. Previous work has focused on fingers with rigid links and have used this effect to estimate the exact location of contact on the links of an underactuated finger [6, 5]. We observe that similar techniques can be used as long as the ratio of joint to link stiffness is low (i.e. the joints are much less stiff than the links).

As our gripper is harder to model than the rigid link fingers explored in previous work, we take a more data-driven approach to contact estimation and use a classifier. We use the stiffnesses observed during a grasp to estimate the grasp type: whether all the fingers are contacting the object with the proximal link (proximal), whether all the fingers are contacting the object with the distal link (distal), or whether all the fingers are contacting the object with both links (both). We require the assumption that all fingers contact the object in the same way because they all share one actuator. This results in a system that is undersensorized - many contact configurations can

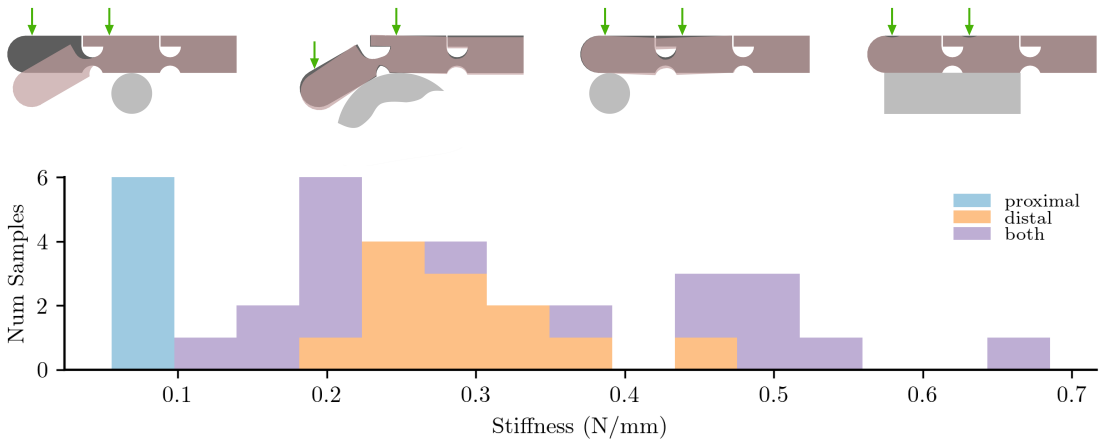

Figure 4-4: The maximum cable stiffness observed during a grasp can be used to identify how FROG is grasping an object. (Top) Diagrams of the distinct types of contact configurations we observe to generate the stiffness groups. Grasps that share the same type of contact configuration tend to result in similar cable stiffnesses. (Bottom) Distribution of maximum cable stiffness for different grasp types.

lead to the same stiffness measurement.

We only do grasp classification in the hard grasp mode because the tension limit in the soft grasp mode prevents us from measuring the stiffness of the fingers after contact. While the gripper is closing in the hard grasp mode, we measure and low pass filter the instantaneous cable stiffness $k(p) = \delta f(p)/\delta p$, recording the maximum instantaneous stiffness. We collect maximum instantaneous stiffness data from different grasps on 8 training objects, listed in Table B.1. The measured stiffnesses are shown in Fig.4-4-bottom. Fig.4-4-top illustrates the types of contact configurations we observe generate each group of stiffnesses, shown for one finger and ordered by increasing stiffness.

After a grasp, we use the maximum instantaneous stiffness to classify the grasp using a k-nearest neighbors classifier ($k = 5$). We evaluate the accuracy of our classifier in Sec. 4.2.2.

4.1.4 Flexure Characterization

To investigate how the time-dependent effects of the flexure stiffness might affect our soft grasp mode, we characterize the dynamic response of the flexures used in the

fingers of FROG. Because the stress in the flexures dissipates over time, the force applied by the soft grasp mode might increase, possibly damaging fragile objects.

As seen in Fig.4-5-left, we mount test flexures to a custom torsion testing system, consisting of a servo (Dynamixel MX-28) and a force/torque sensor (ATI gamma SI-32-2.5). The force/torque sensor measures the reaction torque generated when the servo arm displaces the flexure.

We first test the steady-state torques needed to displace the flexure. We displace the flexure to angles between 0 and 90 degrees and allow the stress in the flexure to relax for 20 minutes before recording the final torque. We see that the steady-state torque is relatively linear with displacement until around 70 degrees, shown in Fig.4-5-middle. This suggests that a linear dynamical model would model the dynamic response of flexure well. We chose to base our model off of the Generalized Maxwell model for viscoelasticity, where the system has an equal number of poles and zeros [3].

We collect displacement and torque data by applying an exponential chirp displacement. The chirp frequency decreases from 0.25 Hz to 0.0005 Hz, and is centered at 45 degrees with an amplitude of 20 degrees. The flexure starts at 0 degrees, and is quickly displaced to 45 degrees to start the chirp. We run 5 trials with 5 identical test flexures, and fit a 5th order model using MATLAB's System Identification Toolbox.

To verify our model, we collect a step response to 45 degrees using the torsion testing system and compare the predicted response to the same displacement input, shown in Fig.4-5-right. Over a time period of 500 seconds, roughly 5x the longest time constant in the model, we find a MSE of $1.2 \cdot 10^{-5}$ N^2-m^2. The fitted model response overestimates the peak of the step response by 7.5%, and slightly underestimates the real response after the first 100 seconds.

We then use the model to estimate how much the force exerted by the soft grasp mode might increase due to stress relaxation in the flexure joints, using the same assumptions as in Sec. 4.1.2. We measure the time needed to reach a proximal flexure angle of 40 degrees, approximately the displacement needed to grasp a 40 mm cylinder, and approximate the closing sequence as a linear ramp up to a set angle,

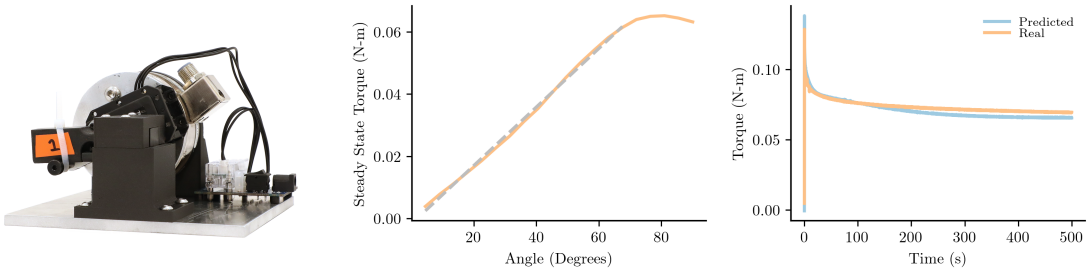

Figure 4-5: Dynamic characterization of the flexures used in the fingers of FROG. (Left) The torsion testing system used to test the flexure. (Middle) Steady state torque of the flexure. The steady state torque is relatively linear with displacement until around 70 degrees. (Right) Step responses predicted by the identified model and measured on the testing system.

after which the angle of the flexures remain constant.

Using the model, we predict that the torque decays to approximately 60% of its value at the end of the ramp after 500 seconds. Using $\epsilon = 0.1$ N and assuming single points of contact, we estimate that the grasp force should increase 120% after 500 seconds of holding an object using the soft grasp mode. Expecting that compensation for this effect is needed, we test the soft grasp mode in Sec. 4.2.1 to see if this increase in grasp force is reflected in the holding force. However, we see that there seems to be no increase in holding force over time and so additional compensation is not needed.

4.2 Performance Characterization

We evaluate FROG's performance with two experiments. We test the strength and gentleness of FROG by measuring its holding force using the soft and hard grasp modes, and evaluating the performance of our grasp classifier.

4.2.1 Holding Force

We evaluate the strength of the grasps generated by FROG by measuring the force needed to pull out grasped objects using a tension testing machine (Instron 5944). We attach FROG to the stationary lower test fixture and the test object to the moving upper test fixture, shown in Fig.4-6-left. The test object is lowered into the gripper

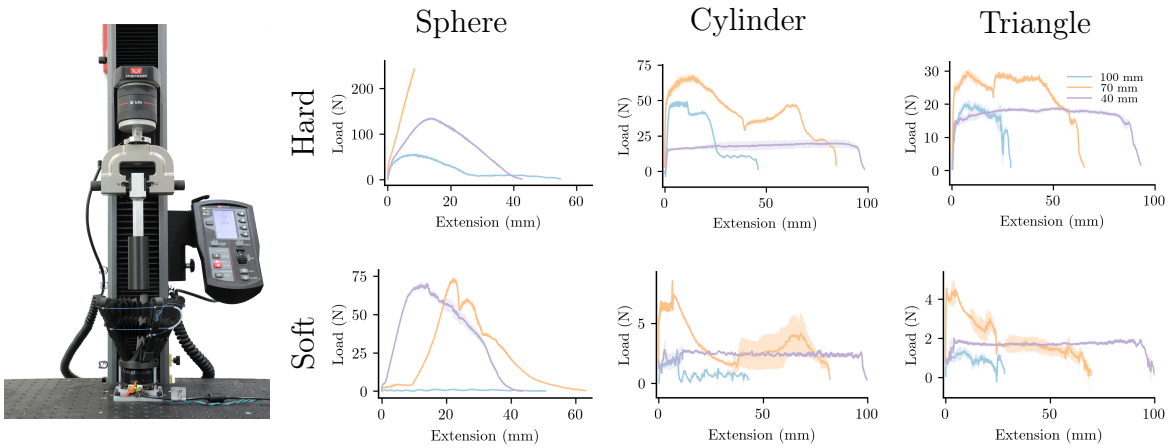

Figure 4-6: Holding force characterization. (Left) The test setup used to measure the holding force. (Right) The graphs show the average holding force for each test object (sphere, cylinder, triangle - 40 mm, 70 mm, 100 mm) with the standard deviation shown by the shaded area. The test is terminated after the object left the gripper or a force limit was reached.

until just before the object contacts any of part of the gripper and before any of the gripper fingers contact the mounting stem during a hard grasp. The same object location is used for the soft and hard grasp tests. The object is grasped using one of the two grasp modes and the object is pulled out of the gripper by moving the upper test fixture at 1 mm/s. The test is terminated after the object leaves the gripper or when a maximum load of 250 N is exceeded.

Because of abnormal wear on the cable prior to this test, a minimum soft grasp ϵ of 0.5 was needed to close the gripper. We use $\epsilon = 0.6$ in the test, as the additional friction adds a constant offset this is equivalent to $\epsilon = 0.1$.

We test three different object geometries in three different sizes - spheres, cylinders, and triangular prisms in 40 mm, 70 mm, and 100 mm. The cylinder tests force closure grasps, the sphere tests both force and form closure grasps, and triangular prism acts as an adversarial object that contacts the cables of FROG. Fig.4-6-right shows the results of this experiment.

Overall, we see that medium sized objects are grasped the strongest, with the maximum holding force reaching $>$ 250 N, 65 N, and 30 N for the 70mm sphere, cylinder, and triangular prism, respectively. The sharp corners of the triangular

prism interfere with the cable and reduce the cable tension while the faces cause the fingers of FROG to twist before they lay flat, reducing the grasp force that FROG is able to exert. Form closure grasps greatly increase the holding force, as the cable tension can now directly resist external forces instead of being transmitted through friction. We see that FROG performs worse at holding objects closer to its minimum or maximum diameters, although it is still able to resist a significant amount of force.

The holding force for the form closure grasps on the various spheres show a distinct peak as the test object is pulled past the fingers. The 70 mm and 100 mm cylinders show distinct regions before and after the proximal links lose contact at around 40 mm and 24 mm of extension, respectively. Although the 70 mm triangle has a similar change in grasp modes, the regions are not clearly distinguishable in the holding force. These regions are less clear in the soft grasp tests. Overall, these results support that FROG is capable of generating strong grasps with high holding forces.

Next, we use holding forces for force closure grasps in the soft grasp mode to confirm that the soft grasp mode greatly decreases the grasp force. In a force closure grasp the holding force is related to the grasp force through the friction coefficient. Assuming a friction coefficient of 1.0 [68], this results in an maximum contact force of 2 N per finger in the soft grasp mode. For grasps where only the distal links contact the test object, the average contact force is 0.4 N per finger. This suggests that the soft grasp mode would be well suited to gentle grasping of fragile objects.

We note that the holding forces for the 40 mm and 70 mm sphere in the soft grasp mode remain relatively large since these are form closure grasps. In addition, the friction in the system may have prevented the motor from correctly regulating the cable tension as the objects push past the fingers, resulting in abnormally high holding forces for these grasps.

Additionally, we test if the holding force in the soft grasp significantly changes due to stress relaxation, as we expect it to based off the analysis in Sec. 4.1.4. We expect the holding force to increase 120% after 500 seconds of grasping. We test the holding force on the 40 mm cylinder 4 times, each test lasting for approximately 120 seconds and maintaining the soft grasp for 300 seconds between tests, seen in Fig.4-7-left. For

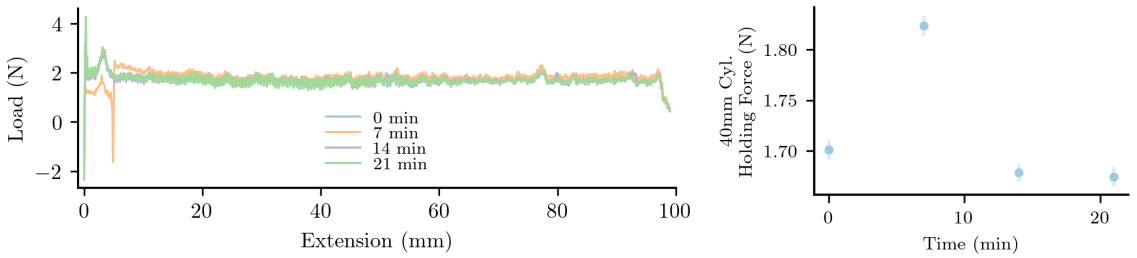

Figure 4-7: Holding force (left) and average holding force (right) on a 40mm cylinder using the soft grasp mode after various grasping times. The error bars show the 95% confidence interval for the average holding force. We expect the holding force to increase as the object is grasped for longer, but do not observe this effect.

each test, we calculate the mean and variance of the holding force after the loading phase but before the unloading phases (between 10 to 85 mm of extension), shown in Fig.4-7-right. We see no discernible relationship between time and holding force, although we expect it to increase over time. This may be due to the cogging torque of the motor or friction along the cable route.

In Sec. 4.3, we qualitatively demonstrate that FROG is able to generate gentle form and force closure grasps on fragile objects.

4.2.2 Grasp Classification

We evaluate the accuracy of FROG's grasp classifier on 21 objects, listed in Table B.1. Objects were held by the experimenter before being grasped by FROG. FROG was able to grasp some test objects with multiple types of grasps, these objects were grasped once per grasp type for a total of 32 trials. We observe that the classifier is able to identify the type of grasp 78% of the time, shown in Table 4.1. All errors were from mixing "distal" and "both" grasps, the classifier was able to identify every proximal grasp. We expect these grasp types to be the most likely to be misclassified as the stiffness distributions for these two grasp types are close and overlap, seen in Fig.4-4-bottom. For distinguishing between "distal" and "both" grasps, we are able to reject the hypothesis that the classifier is worse than random guessing with 99% confidence.

We notice that objects that are very different in size from the training objects

do worse. This is because the object determines the finger configuration which in turn determines the cable stiffness, shifting the stiffness distribution for each grasp type. Grasps on non-rigid objects are not identified well because the compliance of the object results in a lower measured cable stiffness.

Overall, experimental results support that this data-driven approach can be used for proprioception. We are able to use the stiffness seen by the motor to estimate how FROG's fingers are contacting the grasped object.

Grasp Type	Accuracy
Proximal	6/6
Distal	12/17
Both	7/9

Table 4.1: The structure of FROG's fingers allows for effective prioprioception through its actuator, allowing FROG to identify the type of grasp it has on an object.

4.3 Demonstrations

4.3.1 Robust Object Grasping

We demonstrate FROG's ability to handle uncertainty in the shape and pose of the grasped object. First, we demonstrate FROG's ability to grasp objects that differ widely in size and shape. We mount FROG on a robot arm (Universal Robots UR5) and roughly position the target object underneath the gripper, manually controlling the arm to move the gripper up and down. Some objects (toy plane, game controller, brush) are positioned in specific orientations such that they are graspable by FROG.

Table C.1 lists the objects grasped in this experiment, also shown in Fig.4-8. We notice that the bellows between each finger tend to flip to the side of the cables when grasping very irregular objects, which shows that the bellows are not completely effective in separating the grasped object from the cable. Despite this, we did not observe any damage done to the objects or the gripper. The underactuated structure of the gripper allows for the generation of robust grasps without explicitly determining

finger contact locations or sensor feedback, showing FROG's ability to deal with shape uncertainty.

Next, we demonstrate FROG's ability to grasp objects that are not positioned directly underneath the gripper. Table C.2 lists the objects grasped in this experiment. We notice that for small objects, as long as the object starts within the "envelope" of the gripper, FROG tends to cage and center the object before grasping it. This behavior allows the gripper to deal with moderate amounts of pose uncertainty and funnels the object pose into a smaller set of possible locations [56], eliminating uncertainty without sensing. An example of this is shown in Fig.4-9-top. For larger objects, we observe that the object moves less and the fingers of FROG instead move to contact the object, shown in Fig.4-9-bottom. In these cases, the grasp does not not funnel the object as much but still generates a secure grasp. Overall, this demonstration shows FROG's ability to deal with pose uncertainty.

Figure 4-8: FROG's underactuated structure allows it to robustly and forcefully grasp objects without planning contact locations or using sensory feedback, demonstrating the grippers ability to deal with shape uncertainty. Shown in the picture are all the objects that we tested, FROG was able to successfully grasp all of them.

Figure 4-9: Depending on the resistive forces on the grasped object (e.g. friction), FROG's underactuated structure either cages and centers the object (top) or conforms to the object without object movement (bottom), demonstrating the grippers ability to deal with pose uncertainty.

4.3.2 Fragile Object Grasping

Next, we demonstrate the FROG's ability to grasp fragile objects in Fig.4-10. We grasp objects that are fragile due to their geometry, brittleness, or low stiffness using FROG's feedforward soft grasp mode. We set $\epsilon = 0.1$ for all of these tests. All items other than the Play-Doh were grasped off of a stand, the Play-Doh was held by the experimenter before being grasped by FROG.

We first grasp two origami objects, a paper boat and a paper ball, shown in Fig.4-

10-top and Fig.4-10-top middle, respectively. Both objects are crushed when grasped with the hard grasp mode and are undamaged when grasped with the soft grasp mode. The paper ball is grasped in a form closure grasp similar to the one tested in Sec. 4.2.1, demonstrating that the enveloping grasps generated by FROG are gentle. We then demonstrate FROG on potato chips, a delicate and brittle object [36, 55], seen in Fig.4-10-bottom middle. We see a similar result - the hard grasp mode causes the potato chip to fracture while the soft grasp mode is able to lift the chip without damage. Finally, we show FROG grasping a ball of Play-Doh without significant deformation using the soft grasp mode, shown in Fig.4-10-bottom. Once we switch to the hard grasp mode the ball is squished. This demonstrates the effectiveness of our soft grasp mode in grasping fragile or soft objects.

Figure 4-10: FROG's structure and actuator allow it to easily grasp fragile objects. Columns show start (left), result after a soft grasp (middle), and result after a hard grasp (right). Rows show grasping a paper boat (top), paper ball (top middle), potato chip (bottom middle), and ball of Play-Doh (bottom).

Chapter 5

Conclusions

This thesis aims to improve the capabilities of robot hands by developing two new mechanically intelligent end effectors. To limit system complexity a significant amount of computation is done in the hardware, greatly increasing capability when combined with simple perception and control algorithms.

Belt Orienting Phalanges (BOP), introduced in Chapter 3, is a gripper that is capable of a large range of in-hand motions while being able to exert significant forces on manipulated objects and continuously maintain contact with the object. By using belts to create multiple active surfaces within a single finger, BOP is able to avoid excessive hardware complexity while maintaining direct control over three of the grasped object's DoFs. This allows for the use of simple closed-loop controllers and planning algorithms to accomplish real-world tasks.

Flexible Robust Observant Gripper (FROG), introduced in Chapter 4, is a gripper that is capable of generating robust grasps that can be strong or gentle and estimating the type of grasp it has on an object through its actuator. FROG is able to leverage the structure of its fingers to provide strong perception priors, enabling the use of a stiffness-based prioprioception algorithm. In addition, FROG's underactuated fingers allow it to robustly generate grasps on complicated objects without planning contact locations or using sensory feedback.

We hope that these two grippers serve as a starting point for more exploration into this design space and show the advantages of this approach to robot design.

5.1 Lessons Learned

Other than the contributions presented, I believe there are other takeaways from this work that are not purely technical. I hope that someone finds these useful.

1. **If the mapping between effort and progress is bad, consider changing the mapping.** In both engineering and life, you can either try harder or try something else. For a perception system, you can either use more observations or change the physical system to make each observation more valuable. For a control system, you can either use more complicated controllers or change the physical system to make it easier to control. For stressors in life, you can work through the stress or change something to reduce it.

2. **Although nature did not give us telescoping legs, that doesn't mean we can't give robots telescoping legs [10].** While biologically inspired and human-like systems are appealing on both a technical and visceral level, the constraints on living beings are much different than the constraints on robots, and so "optimal" designs may be different.

3. **Sense, think, think some more, and finally act.** Don't rush into doing the first idea that you think of. Spent more time think about what you plan to do and how specific decisions might affect your future self.

5.2 Future Work

To further improve BOP, we are interested in making the overall system more portable and general. We would like to develop more robust closed loop controllers that can reason about belt contact. Another improvement would be to use a wrist mounted camera for 6D pose estimation instead of an external motion capture system. This camera could also improve the generality of our grasp-to-grasp motion planner by estimating the geometry of unknown objects. A neural network could be used to resolve occlusions and the resulting model could be used for planning, replanning as

the object moves and more of the object geometry is known. Additionally, we want to explore how the simple control-to-movement mapping provided by BOP might influence the effectiveness of a learning-based controller.

For FROG, we are interested in further exploring the estimation that can be done with the stiffness information and how the specifics of the flexure design can influence prioprioception quality. The grasp classifier could take into account how the cable stiffness changes over time. As an extension, contact could be detected and the finger pose could be estimated using the timeseries stiffness data. In addition, we noticed that while objects that are grasped further along the fingers are grasped firmly, some translational and rotational degrees of freedom are less stiff due to the low adduction/abduction stiffness of the flexure joints. We believe that the timeseries stiffness information could also be used to estimate the stiffnesses of these degrees-of-freedom, allowing for control over the end-effector stiffness through finger placement.

In this thesis, I've taken a hardware focused approach. It would be interesting to explore this problem from a software perspective, designing algorithms that when combined with simple hardware changes maximize the increase in the capabilities of the robot while limiting the increase in overall complexity. A critical missing piece in evaluating the effectiveness of any of these designs is a method to properly quantify the complexity of a system. Overall, we are interested in continuing to explore the overall design space, combining the strengths of the grippers that we design along the way to develop more and more capable, robust, and dexterous robot hands.

Appendix A

BOP Domain Description

To plan for multistep manipulation tasks, we use PDDLStream [25], a state-of-the-art task and motion planning (TAMP) framework [26]. PDDLStream searches for a sequence of actions for the robot to execute and samples the discrete and continuous values, such as grasps, paths, objects, etc., that parameterize those actions. The set of actions and the set of fact types which define the state, and the derived facts are specified in the Planning Domain Definition Language (PDDL) [59]. To include continuous domains, PDDLStream extends PDDL to allow for the sampling of continuous parameters subject to constraints. These samplers are specified in PDDL and implemented in Python.

The actions, derived facts and samplers for the light bulb domain are given in Table A.1a, Table A.1b and Table A.1c, respectively. Each action is defined by its preconditions and effects. Preconditions are the set of facts that must be true in order to perform the action. Effects are the changes to set of facts as a result of performing the action. We can define derived facts as facts that are true as the result of other facts. Finally, samplers are conditional generators that output facts that certify that the set of parameters satisfy some set of constraints.

In our domain, we include six actions: `move`, `move_holding`, `pick`, `place`, `inhand_move` and `screw_in` (Table A.1a). The last two actions are specific to BOP's capabilities.

The `inhand_move` action plans in-hand motion that moves an object o from a

starting grasp g_o^1 to a goal grasp g_o^2. The sampler `plan-inhand-motion` (Table A.1c) uses the motion planner defined in Sec. III-D to generate the in-hand trajectory.

The `screw_in` action first uses a guarded move [83] to insert the bulb into the socket and then rotates the bulb by a fixed amount using the roll primitive. These trajectories are generated by sampler `plan-tighten-motion`. The sampler `sample-insertion-pose` samples valid bulb poses that parameterize the guarded move and the sampler `sample-insertion-grasp` samples valid grasps that enable the appropriate roll motion.

In our light bulb domain, we define the goal of the task as `(Tightened` o`)`, which is acheived by the `screw_in` action.

While we demonstrate multistep planning in the context of the light bulb domain, these primitives could be used to complete a wide range of tasks. We could further leverage BOP's dexterity by encoding each of the additional manipulation capabilities described in Sec. IV-A as action within this framework.

Operator	Preconditions	Effects
move_free	(AtConf a q_1) $\wedge$ (HandEmpty a) $\wedge$ (FreeMotion a q_1 q_2 t) $\wedge$ ($\neg$ (TrajUnsafe a t))	($\neg$ (AtConf a q_1)) $\wedge$ (AtConf a q_2)
move_holding	(AtConf a q_1) $\wedge$ (Movable o) $\wedge$ (AtGrasp a o g_o) $\wedge$ (HoldingMotion a o g_o q_1 q_2 t) $\wedge$ ($\neg$ (TrajUnsafe a t))	($\neg$ (AtConf a q_1)) $\wedge$ (AtConf a q_2)
pick	(AtConf a q) $\wedge$ (HandEmpty a) $\wedge$ (Movable o) $\wedge$ (AtPose o p_o) $\wedge$ (Kin a o p_o g_o q t) $\wedge$ ($\neg$ (TrajUnsafe a t))	(AtGrasp a o g_o) $\wedge$ ($\neg$ (AtPose o p_o)) $\wedge$ ($\neg$ (HandEmpty a))
place	(AtConf a q) $\wedge$ (AtGrasp a o g_o) $\wedge$ (Kin a o p_o g_o q t) $\wedge$ ($\neg$ (TrajUnsafe a t))	(AtPose o p_o) $\wedge$ (HandEmpty a) $\wedge$ ($\neg$ (AtGrasp a o g_o))
inhand_move	(AtConf a q) $\wedge$ (AtGrasp a o g_o^1) $\wedge$ (RegraspMotion a o g_o^1 g_o^2 q t)	$\wedge$ (AtGrasp a o g_o^2) $\wedge$ ($\neg$ (AtGrasp a o g_o^1))
screw_in	(AtConf a q) $\wedge$ (InsertionGrasp a o g) $\wedge$ (AtGrasp a o g) $\wedge$ (InsertionPose o p h) (TwistMotion a o p g q t)	(Tightened o)

(a) Actions

Derived Facts	Definition
(Holding o)	$\exists a, g_o$ (AtGrasp a o g_o)
(On o r)	$\exists p_o$ ((Supported o p_o r) $\wedge$ (AtPose o p_o))
(TrajUnsafe a_1 t_1)	($\exists o, p$ ((AtPose o p) $\wedge$ ($\neg$ (ObjCollisionFree a_1 t_1 o p))))

(b) Derived Facts

Sampler	Inputs	Outputs	Certified Facts
sample-pose	o r	p_o	(Pose o p_o) $\wedge$ (Supported o p_o r)
sample-grasp	a o	g_o	(Grasp a o g_o)
inverse-kinematics	a o p_o g_o	q t	(Kin a o p_o g_o q t) $\wedge$ (Conf q) $\wedge$ (Traj t)
plan-free-motion	a q_0 q_1	t	(FreeMotion a q_1 q_2 t) $\wedge$ (Traj t)
plan-holding-motion	a q_1 q_2 o g_o	t	(HoldingMotion a o g_o q_1 q_2 t) $\wedge$ (Traj t)
test-obj-collision	a_1 t_1 o p		(ObjCollisionFree a_1 t_1 o p)
sample-insertion-pose	o h	p_o	(Pose o p_o) $\wedge$ (InsertionPose o p_o h)
sample-insertion-grasp	a o	g_o	(Grasp a o g_o) $\wedge$ (InsertionGrasp a o g)
plan-inhand-motion	a g_o^1 g_o^2 o q	t	(RegraspMotion a o g_o^1 g_o^2 q t) $\wedge$ (Traj t)
plan-tighten-motion	a o p_o g_o	q t	(TwistMotion a o p_o g_o q t) $\wedge$ (Traj t) $\wedge$ (Conf q)

(c) Samplers

Table A.1: The light bulb domain is defined by the set of actions, the derived predicates and the samplers. The final two actions (`inhand_move`, `screw_in`) and the final four samplers (`sample-insertion-pose`, `sample-insertion-grasp`, `plan-inhand-motion`, `plan-tighten-motion`) are unique to BOP's capabilities. Throughout the table we use the symbols: a is a robot arm, o is an object, p_o is a pose of object o, g_o is a grasp on object, o, q is a configuration, r is a region, t is a trajectory.

Appendix B

Items Used in Grasp Classifier

Training Items	Testing Items
Baseball, Plastic Peach, Plastic Pear, Can of Tuna, Plastic Bottle, Metal Tumbler, Spray Can, Aluminum Stand	Tape Measure, Plastic Lemon, Plastic Apple, Large Construction Block, Whiteboard Eraser, Water Bottle, Plastic Box, Mustard Bottle, Lubricant Spray Can, Soup Can, Plastic Tumbler Lid, Single Serving Cereal, Glass Beaker, Small Roll of Tape, Large Hex Nut, Dried Gourd, "Easy" Button, Empty Plastic Soda Bottle, Large Roll of Tape, Computer Mouse

Table B.1: Items used to fit and test the stiffness-based grasp classifier. All items were manually presented to the gripper.

Appendix C

Items Grasped by FROG

Items
Cleaning spray, Toy plane, Large construction block, Game controller, Mug, Whiteboard Eraser, Bag of marbles, Brush, Watering can, Pot lid, Water bottle, Soup can, Plastic chain, Box of sugar, Softball, Baseball, Golf ball, Plastic peach, Plastic lemon, Plastic Apple, Inflatable ball, Jar of multivitamins, Small clamp, Medium clamp, Large clamp, Measuring spoon, Can of coffee, Mustard bottle, Large box of jello, Small box of jello, Can of tuna

Table C.1: Items successfully grasped by FROG, used to test the effect of shape uncertainty. All items with the exception of the game controller and brush were grasped from a flat surface.

Items
Small cardboard box, Small roll of tape, Medium roll of tape, Glass beaker, Lock, Mug, Roll of solder, Plastic bottle, Compressed air can

Table C.2: Items used to test the effect of pose uncertainty on grasping with FROG. All items were grasped from a flat surface.

Appendix D

Design Files

The design files for BOP can be found at:
https://www.dropbox.com/sh/ncvrn6378ods4tg/AAA6-KCB-UBm5BPeDUEWdtfha

The design files for FROG can be found at:
https://www.dropbox.com/sh/ylzpein76tme5xy/AABCkezieWjY3APB7vWcnT0va

These files include CAD files of the grippers and test platforms (Solidworks), PCB schematics and layouts (KiCad), firmware, and code used to control the grippers.

Bibliography

[1] Sylvain Abondance, Clark B Teeple, and Robert J Wood. A dexterous soft robotic hand for delicate in-hand manipulation. *IEEE Robotics and Automation Letters*, 5(4):5502–5509, 2020.

[2] OpenAI: Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Józefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, Jonas Schneider, Szymon Sidor, Josh Tobin, Peter Welinder, Lilian Weng, and Wojciech Zaremba. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research (IJRR)*, 39(1):3–20, January 2020.

[3] Behzad Babaei, Ali Davarian, Kenneth M. Pryse, Elliot L. Elson, and Guy M. Genin. Efficient and optimized identification of generalized maxwell viscoelastic relaxation spectra. *Journal of the Mechanical Behavior of Biomedical Materials*, 55:32–41, 2016. doi: 10.1016/j.jmbbm.2015.10.008.

[4] Bruno Belzile and Lionel Birglen. A compliant self-adaptive gripper with proprioceptive haptic feedback. *Autonomous Robots*, 36(1-2):79–91, 2013. doi: 10.1007/s10514-013-9360-1.

[5] Bruno Belzile and Lionel Birglen. Stiffness analysis of double tendon underactuated fingers. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014. doi: 10.1109/icra.2014.6907845.

[6] Bruno Belzile and Lionel Birglen. Stiffness analysis of underactuated fingers and its application to proprioceptive tactile sensing. *IEEE/ASME Transactions on Mechatronics*, 21(6):2672–2681, 2016. doi: 10.1109/tmech.2016.2589546.

[7] Ankita Bhatia, Aaron M. Johnson, and M. T. Mason. Direct drive hands: Force-motion transparency in gripper design. *Robotics: Science and Systems (RSS)*, 2019. URL `http://www.roboticsproceedings.org/rss15/p53.pdf`.

[8] Aditya Bhatt, Adrian Sieler, Steffen Puhlmann, and Oliver Brock. Surprisingly robust in-hand manipulation: An empirical study. *arXiv preprint arXiv:2201.11503*, 2022.

[9] R. Adam Bilodeau, Edward L. White, and Rebecca K. Kramer. Monolithic fabrication of sensors and actuators in a soft robotic gripper. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2015. doi: 10.1109/iros.2015.7353690.

[10] Reinhard Blickhan, Andre Seyfarth, Hartmut Geyer, Sten Grimmer, Heiko Wagner, and Michael Günther. Intelligence by mechanics. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365 (1850):199–220, 2006. doi: 10.1098/rsta.2006.1911.

[11] L. B. Bridgwater, C. A. Ihrke, M. A. Diftler, M. E. Abdallah, N. A. Radford, J. M. Rogers, S. Yayathi, R. S. Askew, and D. M. Linn. The robonaut 2 hand - designed to do work with tools. In *International Conference on Robotics and Automation (ICRA)*. IEEE, 2012.

[12] David Lawrence Brock. Enhancing the dexterity of a robot hand using controlled slip. In *International Conference on Robotics and Automation (ICRA)*, pages 249–251. IEEE, 1988.

[13] Ian M Bullock, Raymond R Ma, and Aaron M Dollar. A hand-centric classification of human and robot dexterous manipulation. *Transactions on Haptics*, 6 (2):129–144, 2012.

[14] Berk Calli, Arjun Singh, Aaron Walsman, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. The YCB object and model set: Towards common benchmarks for manipulation research. In *International Conference on Advanced Robotics (ICAR)*, pages 510–517. IEEE, 2015.

[15] Nikhil Chavan-Dafle, Alberto Rodriguez, Robert Paolini, Bowei Tang, Siddhartha S Srinivasa, Michael Erdmann, Matthew T Mason, Ivan Lundberg, Harald Staab, and Thomas Fuhlbrigge. Extrinsic dexterity: In-hand manipulation with external forces. In *International Conference on Robotics and Automation (ICRA)*, pages 1578–1585. IEEE, 2014.

[16] Nikhil Chavan-Dafle, Rachel Holladay, and Alberto Rodriguez. Planar in-hand manipulation via motion cones. *The International Journal of Robotics Research (IJRR)*, 39(2-3):163–182, 2020.

[17] Tao Chen, Jie Xu, and Pulkit Agrawal. A system for general in-hand object re-orientation. In *Conference on Robot Learning (CoRL)*, pages 297–307. PMLR, 2022.

[18] Tianjian Chen, Zhanpeng He, and Matei Ciocarlie. Hardware as policy: Mechanical and computational co-optimization using deep reinforcement learning. *arXiv preprint arXiv:2008.04460*, 2020.

[19] Whitney Crooks, Gabrielle Vukasin, Maeve O'Sullivan, William Messner, and Chris Rogers. Fin ray® effect inspired soft robotic gripper: From the RoboSoft grand challenge toward optimization. *Frontiers in Robotics and AI*, 3, 2016. doi: 10.3389/frobt.2016.00070.

[20] Silvia Cruciani and Christian Smith. In-hand manipulation using three-stages open loop pivoting. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 1244–1251. IEEE/RSJ, 2017.

[21] M.R. Cutkosky. On grasp choice, grasp models, and the design of hands for manufacturing tasks. *IEEE Transactions on Robotics and Automation*, 5(3): 269–279, jun 1989. doi: 10.1109/70.34763.

[22] Raphael Deimel and Oliver Brock. A compliant hand based on a novel pneumatic actuator. In *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013. doi: 10.1109/icra.2013.6630851.

[23] Aaron M Dollar and Robert D Howe. The sdm hand: A highly adaptive compliant grasper for unstructured environments. In *Experimental Robotics: The Eleventh International Symposium*, pages 3–11. Springer, 2009.

[24] Franka Emika. Franka emika research 3, 2023. URL `https://www.franka.de/research`.

[25] Caelan Reed Garrett, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. PDDLStream: Integrating symbolic planners and blackbox samplers via optimistic adaptive planning. In *International Conference on Automated Planning and Scheduling (ICAPS)*, volume 30, pages 440–448, 2020.

[26] Caelan Reed Garrett, Rohan Chitnis, Rachel Holladay, Beomjoon Kim, Tom Silver, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Integrated task and motion planning. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(1):265–293, 2021.

[27] Paul Glick, Srinivasan A. Suresh, Donald Ruffatto, Mark Cutkosky, Michael T. Tolley, and Aaron Parness. A soft robotic gripper with gecko-inspired adhesive. *IEEE Robotics and Automation Letters*, 3(2):903–910, 2018. doi: 10.1109/lra.2018.2792688.

[28] Jesús M Gómez-de Gabriel and Helge A Wurdemann. Adaptive underactuated finger with active rolling surface. *IEEE Robotics and Automation Letters*, 6(4): 8253–8260, 2021.

[29] Markus Grebenstein, Alin Albu-Schaffer, Thomas Bahls, Maxime Chalon, Oliver Eiberger, Werner Friedl, Robin Gruber, Sami Haddadin, Ulrich Hagn, Robert Haslinger, Hannes Hoppner, Stefan Jorg, Mathias Nickl, Alexander Nothhelfer, Florian Petit, Josef Reill, Nikolaus Seitz, Thomas Wimbock, Sebastian Wolf, Tilo Wusthoff, and Gerd Hirzinger. The DLR hand arm system. In *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011. doi: 10.1109/icra.2011.5980371.

[30] Ankur Handa, Arthur Allshire, Viktor Makoviychuk, Aleksei Petrenko, Ritvik Singh, Jingzhou Liu, Denys Makoviichuk, Karl Van Wyk, Alexander Zhurkevich, Balakumar Sundaralingam, Yashraj Narang, Jean-Francois Lafleche, Dieter Fox, and Gavriel State. Dextreme: Transfer of agile in-hand manipulation from simulation to reality. *arXiv preprint arXiv:2210.13702*, 2022.

[31] Robert Haschke, Jochen J Steil, Ingo Steuwer, and Helge Ritter. Task-oriented quality measures for dextrous grasping. In *International Symposium on Computational Intelligence in Robotics and Automation (SCIRA)*, pages 689–694. IEEE, 2005.

[32] Kris Hauser and Victor Ng-Thow-Hing. Randomized multi-modal motion planning for a humanoid robot manipulation task. *The International Journal of Robotics Research*, 30(6):678–698, 2011.

[33] Rachel Holladay, Tomás Lozano-Pérez, and Alberto Rodriguez. Planning for multi-stage forceful manipulation. In *International Conference on Robotics and Automation (ICRA)*, pages 6556–6562. IEEE, 2021.

[34] Donal Padraic Holland, Colette Abah, Marielena Velasco-Enriquez, Maxwell Herman, Gareth J. Bennett, Emir Augusto Vela, and Conor James Walsh. The soft robotics toolkit: Strategies for overcoming obstacles to the wide dissemination of soft-robotic hardware. *IEEE Robotics and Automation Magazine*, 24(1):57–64, 2017. doi: 10.1109/mra.2016.2639067.

[35] Josie Hughes, Utku Culha, Fabio Giardina, Fabian Guenther, Andre Rosendo, and Fumiya Iida. Soft manipulators and grippers: A review. *Frontiers in Robotics and AI*, 3, nov 2016. doi: 10.3389/frobt.2016.00069.

[36] Josie Hughes, Shuguang Li, and Daniela Rus. Sensorization of a continuum body gripper for high force and delicate object grasping. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, may 2020. doi: 10.1109/icra40945.2020.9196603.

[37] Filip Ilievski, Aaron D. Mazzeo, Robert F. Shepherd, Xin Chen, and George M. Whitesides. Soft robotics for chemists. *Angewandte Chemie International Edition*, 50(8):1890–1895, 2011. doi: 10.1002/anie.201006464.

[38] Rebecca H. Jiang, Neel Doshi, Ravi Gondhalekar, and Alberto Rodriguez. Shape and motion optimization of rigid planar effectors for contact trajectory satisfaction. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022. doi: 10.1109/iros47612.2022.9981918.

[39] Tao Jin, Zhongda Sun, Long Li, Quan Zhang, Minglu Zhu, Zixuan Zhang, Guangjie Yuan, Tao Chen, Yingzhong Tian, Xuyan Hou, and Chengkuo Lee. Triboelectric nanogenerator sensors for soft robotics aiming at digital twin applications. *Nature Communications*, 11(1), 2020. doi: 10.1038/s41467-020-19059-3.

[40] Atsushi Kakogawa, Hiroyuki Nishimura, and Shugen Ma. Underactuated modular finger with pull-in mechanism for a robotic gripper. In *International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2016.

[41] Yiannis Karayiannidis, Karl Pauwels, Christian Smith, Danica Kragic, et al. In-hand manipulation using gravity and controlled slip. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 5636–5641. IEEE/RSJ, 2015.

[42] Mia Kokic, Johannes A Stork, Joshua A Haustein, and Danica Kragic. Affordance detection for task-specific grasping using deep learning. In *International Conference on Humanoid Robotics (Humanoids)*, pages 91–98. IEEE, 2017.

[43] Jiun-Yih Kuan, Kenneth A Pasch, and Hugh M Herr. A high-performance cable-drive module for the development of wearable devices. *IEEE/ASME Transactions on mechatronics*, 23(3):1238–1248, 2018.

[44] James J Kuffner and Steven M LaValle. RRT-connect: An efficient approach to single-query path planning. In *International Conference on Robotics and Automation (ICRA)*, volume 2, pages 995–1001. IEEE, 2000.

[45] Steven M LaValle. *Planning Algorithms*. Cambridge university press, 2006.

[46] Haili Li, Jiantao Yao, Pan Zhou, Xinbo Chen, Yundou Xu, and Yongsheng Zhao. High-load soft grippers based on bionic winding effect. *Soft Robotics*, 6(2):276–288, 2019. doi: 10.1089/soro.2018.0024.

[47] Shuguang Li, John J. Stampfli, Helen J. Xu, Elian Malkin, Evelin Villegas Diaz, Daniela Rus, and Robert J. Wood. A vacuum-driven origami "magic-ball" soft gripper. In *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, may 2019. doi: 10.1109/icra.2019.8794068.

[48] Zexiang Li and S Shankar Sastry. Task-oriented optimal grasping by multifingered robot hands. *IEEE Journal on Robotics and Automation*, 4(1):32–44, 1988.

[49] Yun Lin and Yu Sun. Grasp planning to maximize task coverage. *International Journal of Robotics Research (IJRR)*, 34(9):1195–1210, 2015.

[50] Qiujie Lu, Angus B. Clark, Matthew Shen, and Nicolas Rojas. An origami-inspired variable friction surface for increasing the dexterity of robotic grippers. *IEEE Robotics and Automation Letters*, 5(2):2538–2545, 2020. doi: 10.1109/lra.2020.2972833.

[51] Raymond R. Ma and Aaron M. Dollar. In-hand manipulation primitives for a minimal, underactuated gripper with active surfaces. In *Volume 5A: 40th Mechanisms and Robotics Conference*. American Society of Mechanical Engineers, aug 2016. doi: 10.1115/detc2016-60354.

[52] Mariangela Manti, Taimoor Hassan, Giovanni Passetti, Nicolò D'Elia, Cecilia Laschi, and Matteo Cianchetti. A bioinspired soft robotic gripper for adaptable and effective grasping. *Soft Robotics*, 2(3):107–116, 2015. doi: 10.1089/soro.2015.0009.

[53] Andrew D. Marchese and Daniela Rus. Design, kinematics, and control of a soft spatial fluidic elastomer manipulator. *The International Journal of Robotics Research*, 35(7):840–869, 2015. doi: 10.1177/0278364915587925.

[54] Ramses V. Martinez, Jamie L. Branch, Carina R. Fish, Lihua Jin, Robert F. Shepherd, Rui M. D. Nunes, Zhigang Suo, and George M. Whitesides. Robotic tentacles with three-dimensional mobility based on flexible elastomers. *Advanced Materials*, 25(2):205–212, 2012. doi: 10.1002/adma.201203002.

[55] Ryoji Maruyama, Tetsuyou Watanabe, and Masahiro Uchida. Delicate grasping by robotic gripper with incompressible fluid-based deformable fingertips. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013. doi: 10.1109/iros.2013.6697148.

[56] M. Mason. The mechanics of manipulation. In *Proceedings. 1985 IEEE International Conference on Robotics and Automation*. Institute of Electrical and Electronics Engineers. doi: 10.1109/robot.1985.1087242.

[57] Matthew T Mason. Toward robotic manipulation. *Annual Review of Control, Robotics, and Autonomous Systems*, 1(1), 2018.

[58] Yoky Matsuoka, Pedram Afshar, and Michael Oh. On the design of robotic hands for brain–machine interface. *Neurosurgical Focus*, 20(5):1–9, 2006. doi: 10.3171/foc.2006.20.5.4.

[59] Drew McDermott, Malik Ghallab, Adele Howe, Craig Knoblock, Ashwin Ram, Manuela Veloso, Daniel Weld, and David Wilkins. PDDL-the planning domain definition language, 1998.

[60] Jifei Ou, Felix Heibeck, and Hiroshi Ishii. TEI 2016 studio. In *Proceedings of the TEI '16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*. ACM, 2016. doi: 10.1145/2839462.2854119.

[61] Xinlei Pan, Animesh Garg, Animashree Anandkumar, and Yuke Zhu. Emergent hand morphology and control from optimizing robust grasps of diverse objects. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021. doi: 10.1109/icra48506.2021.9562092.

[62] Rolf Pfeifer and Gabriel Gómez. Morphological computation – connecting brain, body, and environment. In *Creating Brain-Like Intelligence*, pages 66–83. Springer Berlin Heidelberg, 2009. doi: 10.1007/978-3-642-00616-6_5.

[63] C. Piazza, G. Grioli, M.G. Catalano, and A. Bicchi. A century of robotic hands. *Annual Review of Control, Robotics, and Autonomous Systems*, 2(1):1–32, 2019. doi: 10.1146/annurev-control-060117-105003.

[64] G.A. Pratt and M.M. Williamson. Series elastic actuators. In *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots*. IEEE Comput. Soc. Press, 1995. doi: 10.1109/iros.1995.525827.

[65] Steffen Puhlmann, Jason Harris, and Oliver Brock. RBO hand 3: A platform for soft dexterous manipulation. *IEEE Transactions on Robotics*, 2022.

[66] Shadow Robot. Dexterous hand series, 2023. URL https://www.shadowrobot.com/dexterous-hand-series/.

[67] Weiss Robotics. Weiss robotics wsg series, 2023. URL https://weiss-robotics.com/servo-electric/wsg-series/product/wsg-series/.

[68] Sho Sato, Takeshi Yamaguchi, Kei Shibata, Toshiaki Nishi, Kenta Moriyasu, Kenichi Harano, and Kazuo Hokkirigawa. Dry sliding friction and wear behavior of thermoplastic polyurethane against abrasive paper. *Biotribology*, 23:100130, 2020. doi: 10.1016/j.biotri.2020.100130.

[69] Lawrence Shapiro and Shannon Spaulding. Embodied Cognition. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2021 edition, 2021.

[70] Jun Shintake, Samuel Rosset, Bryan Schubert, Dario Floreano, and Herbert Shea. Versatile soft grippers with intrinsic electroadhesion based on multifunctional polymer actuators. *Advanced Materials*, 28(2):231–238, 2015. doi: 10.1002/adma.201504264.

[71] Ali Shtarbanov. FlowIO development platform – the pneumatic "raspberry pi" for soft robotics. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. ACM, 2021. doi: 10.1145/3411763.3451513.

[72] Adam J. Spiers, Berk Calli, and Aaron M. Dollar. Variable-friction finger surfaces to enable within-hand manipulation via gripping and sliding. *IEEE Robotics and Automation Letters*, 3(4):4116–4123, 2018. doi: 10.1109/lra.2018.2856398.

[73] Ian H Taylor, Nikhil Chavan-Dafle, Godric Li, Neel Doshi, and Alberto Rodriguez. Pnugrip: An active two-phase gripper for dexterous manipulation. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 9144–9150. IEEE/RSJ, 2020.

[74] Russ Tedrake. *Robotic Manipulation*. 2022. URL http://manipulation.mit.edu.

[75] Clark B. Teeple. Ctrl-p: Control pressure, 2023. URL https://ctrl-p.cbteeple.com/latest/.

[76] Witchuda Thongking, Ardi Wiranata, Ayato Minaminosono, Zebing Mao, Shingo Maeda, and and. Soft robotic gripper based on multi-layers of dielectric elastomer actuators. *Journal of Robotics and Mechatronics*, 33(4):968–974, 2021. doi: 10.20965/jrm.2021.p0968.

[77] Vinicio Tincani, Manuel G Catalano, Edoardo Farnioli, Manolo Garabini, Giorgio Grioli, Gualtiero Fantoni, and Antonio Bicchi. Velvet fingers: A dexterous gripper with active surfaces. In *International Conference on Intelligent Robots and Systems (IROS)*, pages 1257–1263. IEEE/RSJ, 2012.

[78] Pierre Tournassoud, Tomás Lozano-Pérez, and Emmanuel Mazer. Regrasping. In *International Conference on Robotics and Automation (ICRA)*, volume 4, pages 1924–1928. IEEE, 1987.

[79] Ryan L. Truby, Michael Wehner, Abigail K. Grosskopf, Daniel M. Vogt, Sebastien G. M. Uzel, Robert J. Wood, and Jennifer A. Lewis. Soft somatosensitive actuators via embedded 3d printing. *Advanced Materials*, 30(15):1706383, 2018. doi: 10.1002/adma.201706383.

[80] Ryan L. Truby, Lillian Chin, Annan Zhang, and Daniela Rus. Fluidic innervation sensorizes structures from a single build material. *Science Advances*, 8(32), aug 2022. doi: 10.1126/sciadv.abq4385.

[81] Nathan T. Ulrich. *Grasping With Mechanical Intelligence*. PhD thesis, University of Pennsylvania, 1989.

[82] Tetsuyou Watanabe, Kota Morino, Yoshitatsu Asama, Seiji Nishitani, and Ryo Toshima. Variable-grasping-mode gripper with different finger structures for grasping small-sized items. *IEEE Robotics and Automation Letters*, 6(3):5673–5680, 2021.

[83] Peter Will and David Grossman. An experimental system for computer controlled mechanical assembly. *Transactions on Computers*, 100(9):879–888, 1975.

[84] Andrew D. Wilson and Sabrina Golonka. Embodied cognition is not what you think it is. *Frontiers in Psychology*, 4, 2013. doi: 10.3389/fpsyg.2013.00058.

[85] Jie Xu, Tao Chen, Lara Zlokapa, Michael Foshey, Wojciech Matusik, Shinjiro Sueda, and Pulkit Agrawal. An end-to-end differentiable framework for contact-aware robot design. *arXiv preprint arXiv:2107.07501*, 2021.

[86] Shenli Yuan, Austin D. Epps, Jerome B. Nowak, and J. Kenneth Salisbury. Design of a roller-based dexterous hand for object grasping and within-hand manipulation. In *International Conference on Robotics and Automation (ICRA)*. IEEE, 2020.

[87] Shenli Yuan, Lin Shao, Connor L. Yako, Alex Gruebele, and J. Kenneth Salisbury. Design and control of roller grasper V2 for in-hand manipulation. In *International Conference on Intelligent Robots and Systems (IROS)*. IEEE/RSJ, 2020.