ABSTRACT

Title of Thesis: DEVELOPMENT AND DEPLOYMENT OF A SOFT ROBOTIC GRIPPER: HARNESSING PRESSURE FOR OBJECT RECOGNITION THROUGH DEEP LEARNING

Zachary Coogan
Master of Science, 2025

Thesis Directed by: Professor Eleonora Tubaldi
Department of Mechanical Engineering

Soft robotic manipulators are uniquely suited to handle objects with diverse textures, shapes, and stiffnesses, thanks to their highly deformable structure; however, this same characteristic complicates the collection and interpretation of data for exteroceptive sensing. This challenge is often overcome by using machine learning on a wide array of strain sensors, but this work introduces a methodology that leverages the time-series pressure response of a novel, buckling-enabled soft robotic gripper to enable passive sensing capabilities. The behaviors of this pneumatically controlled, thin-shelled gripper are investigated with numerical modeling and validated experimentally. We detail a robust manufacturing strategy designed to minimize defects and demonstrate several representative use cases of the gripper. To develop built-in sensing capabilities, a dataset of the time series pressure responses during a variety of grasping events is constructed and analyzed. The dataset is initially explored for qualitative insights into exteroceptive sensing. It is then used to train two machine learning classifiers to distinguish between the geometry and size of a

given object at over 85% accuracy across multiple classes. Finally, the ability to recalibrate the model on a new gripper using transfer learning is demonstrated via two grippers with intentionally produced dimensional modifications. By eliminating the need for embedded electronics or structural changes, this strategy enables rapid, low-cost haptic feedback from a single interaction, with broad implications for constrained environments such as surgery and industrial automation.

DEVELOPMENT AND DEPLOYMENT OF A SOFT ROBOTIC GRIPPER: HARNESSING PRESSURE FOR HAPTIC OBJECT IDENTIFICATION THROUGH DEEP-LEARNING

by

Zachary Coogan

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Master of Science
2025

Advisory Committee:
Professor Eleonora Tubaldi, Chair/Advisor
Professor Davis McGregor
Professor Nikhil Chopra

Acknowledgments

First and foremost, I would like to express my deepest gratitude to my advisor, Professor Eleonora Tubaldi, for her unwavering guidance, mentorship, and encouragement throughout both my undergraduate and graduate studies. From my early days in the laboratory to the culmination of this thesis, Professor Tubaldi has continually challenged me to grow as an engineer and a researcher.

I would also like to extend my heartfelt thanks to all my labmates in the Tubaldi Laboratory, whose insights, collaboration, and camaraderie have enriched my research experience. In particular, I am deeply appreciative of Kieran Bervenik, whom I had the pleasure of working with from the onset of my research experience to the culmination in this thesis. I would also like to extend a special thanks to Joseph Messou, who has been a great help in not only the writing of this thesis but in personal development.

To my parents, I owe an immense debt of gratitude. Their unwavering support, encouragement, and belief in my abilities have been the foundation upon which all of my accomplishments rest. I would also like to acknowledge my second family: my (fellow mechanical engineering) roommates, past and present, who have been a defining part of my college experience. Mark Rosser, Evan, Demos, and Ryan Dunning, who I can say wholeheartedly were always a pleasure not only to study with but to make a home with.

Finally, to all who have supported me in ways both big and small throughout this endeavor,

I extend my sincere thanks. This thesis represents the culmination of many shared efforts, and I am deeply appreciative of everyone who has been part of this journey.

UNIVERSITY OF MARYLAND

2181 Glenn L. Martin Hall
College Park, MD 20742-3035
TEL: 301-405-2410
http://www.enme.umd.edu

4/15/2025

Re: Previously Published Materials appearing in Thesis or Dissertation

Dean of the Graduate School,

Zachary Cogan (UID: 117 242 515) has:

☐ NOT INCLUDED any previously published works within their thesis or dissertation.

☒ INCLUDED one or more previously published works within their thesis or dissertation. This letter certifies that the examining committee for the student has determined that the student made a substantial contribution to the previously published work. The inclusion of the previously published work has the approval of the thesis or dissertation advisor and the Graduate Director.

Sincerely,

[signature]

Peter Sandborn

Peter Sandborn
Director of Graduate Studies
Department of Mechanical Engineering
University of Maryland

Sincerely,

Eleonora Tubaldi

Eleonora Tubaldi
Advisor for Zachary Coogan

Table of Contents

Chapter 1: Introduction

1.1 The Role of Soft Robotics

Traditionally, robots have been constructed from rigid metallic structures actuated at specific joints with electrical signals, but advances in soft robotics have begun to expand the realm of robotics to include flexible structures. Soft robotics utilize their inherent compliance to create structures with infinite degrees of freedom, which vastly expands the realm of possibilities for robotic manipulators, making them better suited for tasks that require greater flexibility or a gentle touch. These include challenges such as minimally invasive surgery [1] and usage in confined environments [2], as well as human-robot interaction [3] or interaction with otherwise delicate and fragile objects [4].

These robots have also allowed researchers to look towards biology for novel and creative solutions to problems where conventional robotics fall short [5, 6]. Taking inspiration from biology, a variety of creative and novel actuation can be achieved. One such usage is for energy-efficient locomotion, [7, 8], which can also be inspired by marine invertebrates [9]. This bio-inspiration also helps in the pursuit of attributes hard or impossible to obtain in conventional robotics, such as self-growth, self-responsiveness, and damage resilient systems [10]. Self-growth was originally inspired by plants' abilities to start from a central rooted position and expand almost limitlessly and is usually emulated in soft robotics through the usage of stored

tubing [11]. Soft robotics are also uniquely resistant to damage from blunt force impact or excessive loads since they can often experience large amounts of elastic deformation without affecting their performance. Similarly, soft robotics' high elasticity also reduces the maximum impulse experienced in collisions. Self-responsive soft robots refer to their ability to conform to the environment in beneficial ways [11, 12], such as an octopus' ability to squeeze through small orifices or securely grasp almost any object with its tentacles. In contrast, a hard robotic gripper is usually designed for a specific task and is much less adaptable than compliant systems when attempting to grasp objects it was not specifically designed for. Soft robotic grippers are also particularly well suited for applications requiring delicate touch or for objects difficult to grasp with traditional robotics. As an example, Figure 1.1 showcases a soft gripper's ability to successfully capture a jellyfish with minimal force to prevent harm [4].

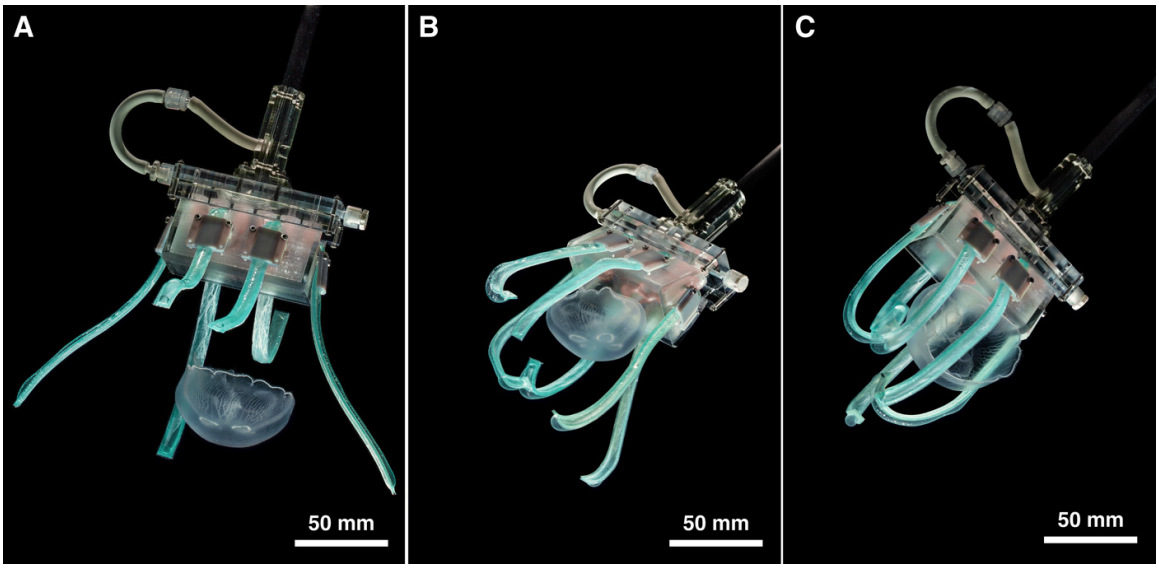

Figure 1.1: An example of an ultra-gentle gripper grasping a jellyfish [4].

1.2 Challenges with soft robotics

Historically, one of the primary challenges in soft robotics has been their inherently limited actuation speed and force output [13, 14]. These limitations arise directly from the compliant and elastic nature of their construction materials, which, while advantageous for safe interaction and adaptability, often compromise performance in terms of actuation speed. One particularly promising approach involves the deliberate incorporation of mechanical instabilities. By leveraging mechanisms such as snap-through and buckling, soft devices can be designed to undergo rapid structural transformations [15, 16].

This strategy of harnessing mechanical instabilities has led to significant advancements across a variety of soft robotic platforms. For example, snap-through and buckling phenomena have enabled the creation of high-speed soft locomotion systems capable of performing tasks such as crawling [17] or jumping [18, 19]. These locomotors harness rapid energy release mechanisms, in which structures initially held in high-energy states overcome an energy barrier and transition to a more stable, lower-energy configuration. These releases achieve rapid and motion profiles that would otherwise be unattainable with conventional soft actuator designs.

Moreover, the versatility of instability-based designs extends to soft robotic grippers. These actuators have demonstrated great promise in grippers operating on a micro scale [20], where conventional robotics remain unfeasible to implement. Instability-enabled soft structures are also being applied to extraterrestrial environments for possibilities such as space structure assembly and space debris disposal [21]. In this application, the instability allows for fast actuation while the compliant nature helps to dissipate energy of potentially fast moving debris. These applications illustrate the broad potential of soft robotics—especially those with mechanically

programmed instabilities—not only for enhancing the speed and strength of actuation but also to enable entirely new functionalities within soft robotic systems.

Classically, sensing in robotics has been essential for the control of autonomous systems, especially those that interact with unknown or changing environments [22]. While soft robots' inherent compliance offers great versatility, this also results in challenges regarding on-the-fly modeling, control, and sensing [23]. Such difficulties can be attributed to their inherently flexible structures yielding less predictable deformations from complex effects such as hysteresis, the viscoelasticity, large strains, and nonlinearities [24]. Sensing is further complicated beyond the capabilities of classical *hard* robotic sensing techniques due to soft robotics' continuum deformation behaviors [23] especially when output data can contain multiple overlapping or conflicting signals and can be highly nonlinear [25]. Overcoming these challenges will be especially important if the preexisting theory of robotic control is to be transferred to the soft robotic space. For instance, soft robotics current sensing limitations prevent classical implementations of motion and force control, which require precise knowledge of kinematic constraints if stability is to be maintained, especially during contact events [22].

1.3 Modeling and Sensing in Soft Robotics

Despite the challenges associated with sensing in soft robotics, many implementations have been successfully achieved using a wide variety of embedded sensors, usually relying on several strategically placed strain sensing tactics for open-loop control of soft robotic movement [26], obtaining force/torque applied [27, 28], and curvature/position [29, 30, 31]. These techniques facilitate proprioceptive sensing, allowing for the reconstruction of behavior and opening the door

to creating closed-loop controllers for soft robotics. One such example of this is the reconstruction of a soft continuum's shape with the use of *strain vector aided sensorization of soft structures* (SVAS^3), which is a technique for determining optimal locations for strain sensors to characterize the deformation of the structure [32]. These strain sensors are then able to characterize the shape of the soft continuum body as it undergoes twisting, pushing, and bending. A similar approach embeds sensors in flexible robotics to determine the weights of the basis functions during deformation [33]. This approach allows for a reduction in the number of sensors required to reconstruct the deformations of the soft robot and can be implemented for the control of continuum manipulators.

Identification of the outside environment (exteroceptive sensing) has been implemented with similar approaches. One approach involves the soft robot conforming to an object and implementing a method of proprioception as mentioned above. Alternatively, arrays of pressure sensors can be used to create a sense of touch similar to the human sense of touch [34], as shown by Figure 1.2. In either application, the large quantity of sensors makes the mapping of raw sensor outputs to usable information difficult, so machine learning (ML) techniques are often employed.

Tackling challenges in sensing has helped push soft robotics into use in rehabilitation devices [35] and minimally invasive surgical devices [1], where flexibility is required alongside precise control. However, despite these advances, many challenges to sensing in soft robotics remain unaddressed and in high demand, such as implementing multimodal sensors, the need for stretchable electrodes, wireless electronic interfaces, and data interpretation methods [24].

For some more advanced tasks requiring object identification or sorting, the above-mentioned strategies require sensory feedback for closed-loop control in soft robots [36]. One commonly

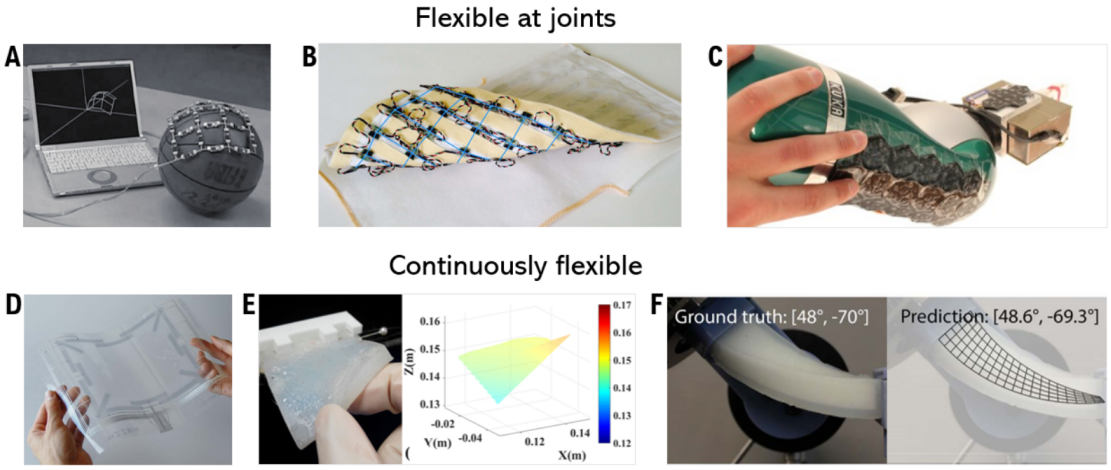

Figure 1.2: Examples of soft robotic skins used for exteroceptive sensing. **(A-C)** arrays of rigid Accelerometers and/or magnetometers measure relative rotation to reconstruct shape. **(D-F)** Continuously flexible devices capture surface deformation to estimate shape, with data-driven methods employed for continuous e-skin shape reconstruction [34].

explored approach towards this goal is the use of arrays of sensors to create a smart skin capable of reconstructing its shape, oftentimes through machine learning techniques [34, 37, 38]. Deep learning has proved to be especially useful for the reconstruction of environmental information from arrays of sensor inputs. As an example, a cicada wing inspired nanostructure with an integrated array of pressure sensors was used to create a highly sensitive haptic sensor [39]. Another bio-inspired soft robotic sensing strategy aims to simulate a human hand for tactile object recognition [40]. Researchers utilized a smaller array of pneumatic tactile sensors on a hand-inspired soft robot to train a ROSE-net deep learning model to classify different grasped objects. Although they demonstrated very high accuracies of over 99%, they required multiple grasps to accomplish this task, and they classified between only a small number of distinctly different everyday objects. Another ML approach involves the use of redundant and arbitrarily placed embedded pressure sensors into a soft robot [41]. Researchers embedded the sensors into a soft robotic finger and store the pressure sensor values alongside video tracking data to be used

as ground truth labels. They then employed a recurrent neural network (RNN) to train the model to determine the shape of the structure as well as the pressure applied to an external object. A similar approach is taken to achieve proprioceptive sensing in a more complex octopus-inspired robotic arm [42]. These researchers again implement a large array of sensors (resistive bend sensors) and raw outputs to train an RNN model. However, in this case, an RNN model was implemented alongside a convolutional encoder to reduce the dimensionality of image data. ML has also been used on time series data to improve the sensing capabilities of soft robotics [27]. This study aims to better model the force and torque applied by a soft robotic actuator being manipulated by Endovascular surgical robots (ESRs) with the goal of enhancing the safety of these ESR-assisted surgeries. A long short-term memory (LSTM) neural network is trained to determine the forces applied to an actuator through time series monitoring of the positional data.

Unfortunately, sensorized skins can be difficult and expensive to implement due to the number of sensors required and the increased stiffness the skins would apply to a pre-existing soft robot. More strategically placed sensors have also been implemented to reconstruct the shape of objects using a soft robotic sensor/actuator that rotates an array of pressure sensors around a target, mimicking the process of a human hand for object identification [43]. Rudimentary haptic sensing has also been accomplished by measuring a soft robot's deformation when interacting with a variety of objects using a K-nearest neighbors tactic [44]. In this classification technique, the input vector's Euclidean distance is compared to previously labeled data points in order to classify the original object [45]. While effective on simpler datasets with less available data, this technique can have difficulties distinguishing more complex and higher dimensional patterns.

More recently, monitoring of pressure has been explored as an alternative, non-invasive, and easy-to-implement form of smart sensing that can be retrofitted to a variety of soft robots to

detect the size or stiffness of an object of interest [46]. A summary of some of Zou et al.'s results [46] can be seen in Figure 1.3 , which highlights how the pressure-volume response of a variety of pre-existing soft robotic grippers can be harnessed for sensing. Throughout the remainder of the work, Zou et al. demonstrate the ability to detect relative object size, stiffness, location, and surface roughness using this pressure monitorization technique.

Here, we explore the possibility of more complex and versatile object identification via time-dependent pressure monitoring with the implementation of ML techniques. While effective in laboratory environments, similar results without the use of ML would struggle with noise, distinctions between similar objects, and require multiple interactions of differing orientations to holistically identify an object.

These challenges are overcome and discussed in Chapter 2 of this thesis through the use of a buckling instability-enabled soft gripper, which is extremely sensitive to the outside environment, as proposed by a paper I am the second author on [47]. In this paper, an analytical and numerical approach for representing a novel soft gripper is proposed and compared against experimental samples before demonstrating its ability for exteroceptive sensing. Chapter 2 details the findings of this paper with a focus on my contributions: the experimental works and a refinement of the manufacturing process. Chapter 3 will then showcase and significantly expand upon the gripper's perceptive abilities shown in the first paper through qualitative analysis of the time series pressure data collected during grasping tasks. The process of collecting a large portion of experimental data used for these qualitative analyses is also summarized. In Chapter 4, we harness this dataset to train ML models and demonstrate its ability to achieve holistic haptic identification of a variety of objects. We also present a strategy for recalibration of this model on other hemispherical grippers (with differently tuned dimensions) and demonstrate the model's feasible deployment

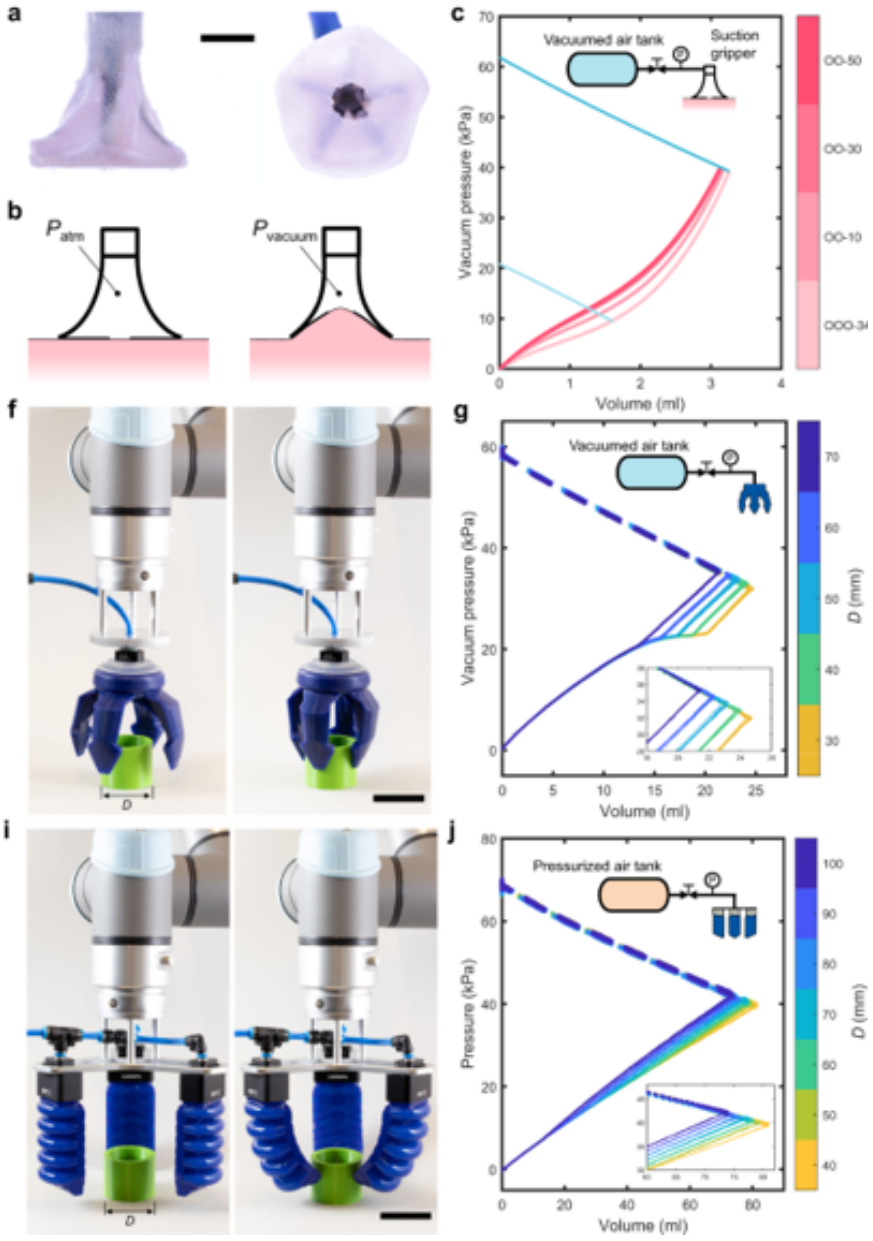

Figure 1.3: A snapshot from [46] which highlights the capability of relative exteroceptive sensing with pressure-volume (PV) curves. (A) an image of a suction enabled gripper, (B) a schematic of this gripper's actuation, (C) PV curves of grasped objects. (F,G) Two *finger* inspired pneumatic grippers, (G,J) and their PV response grabbing different sized objects.

by performing a sorting task.

Chapter 2: Manufacturing and Modeling of the Soft Robotic Gripper

2.1 Overview

The understanding of a soft robot's behavior is critical to predict its interactions with the environment, and especially if we wish to extract tactile information from the response of a soft gripper. Broadly, it is also important for the intelligent design of soft robots, which can be especially challenging from an analytical standpoint given the complexities of soft robotic geometries. This leads us to use computational techniques, but these can be very sensitive to defects or simplifications in numerical models.

This thesis introduces the usage of a novel, soft robotic gripper, which was inspired by a variant of the deep-sea Benthic tunicate, the *Megalodicopia Hians*. This creature developed a hemispherical mouth cavity that collapses around floating organic particles in an energy-efficient feeding mechanism [48]. Our gripper seeks to mimic this motion to develop a new buckling-enabled pneumatic hemispherical gripper shown in Figure 2.1. In its undeformed state, it is a sealed hemispherical shell that is sealed at the equator with a more compliant material, forming a fluidic cavity. Upon a decrease in pressure in the sealed cavity, the shell begins to buckle in a mouth-like fashion, closing around an object.

In this chapter, we detail the design, manufacturing, and modeling of this novel soft robotic gripper used throughout the remainder of this Thesis. We will first present the manufacturing

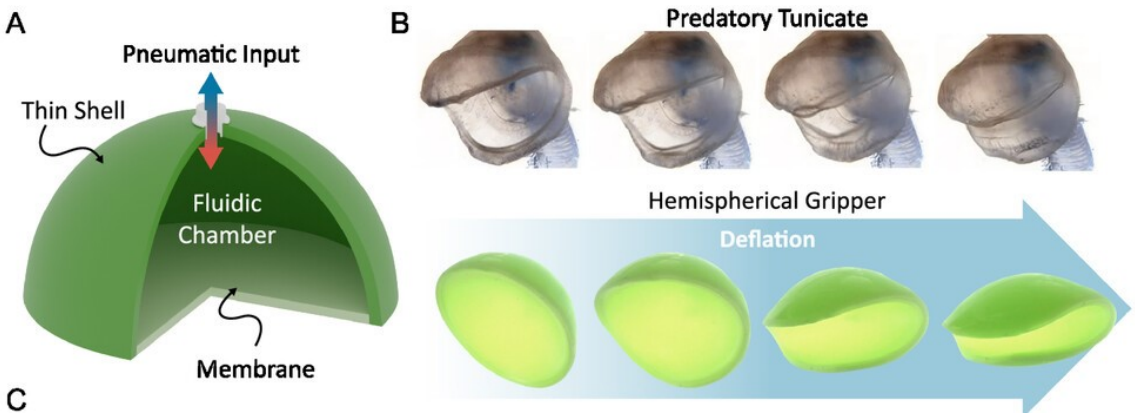

Figure 2.1: A) A schematic of the gripper's structure along with (B) a comparison of the biological inspiration (the *Megalodicopia Hians*) with snapshots of our gripper's actuation.

development process and we will discuss the need for defect control along with a potential strategy for compensating for avoidable defects through the addition of intentional imperfections. We then focus on how we created the functional design space of the gripper using a numerical modeling scheme. Finally, the numerical findings are validated against the experimental trials.

2.2 Design and manufacturing of the gripper

The gripper consists of a thin hemispherical shell of Zhermack Elite Double 32 bonded to a thin layer of ecoflex across the equator of the hemisphere. The resulting soft gripper is a completely sealed two silicone balloon-like structure. A pneumatic valve is then inserted into the top of the hemispherical shell to allow for actuation of the gripper. Although a variety of manufacturing processes were tested, the resulting final process is as follows:

Step 1: The uncured polymer Elite Double 32 is mixed and let sit for approximately 300 seconds. After that, it is poured onto a smooth metal sphere and the elastomer is then allowed to cure. This step can be repeated multiple times to create multiple layers to

increase the thickness of the manufactured shell.

Step 2: Once the final layer is cured, the shell is separated from any excess material that overflowed around the base using a sharp blade.

Step 3: The shell is removed from the metal sphere by gently inverting it, mitigating the risk of material tearing during the process.

Step 4: The shell is placed onto a 3D-printed spherical piece by gently inverting the sample for the second time so that the shell returns to its initial orientation.

Step 5: A sharp blade is used to cut around the equatorial line of the hemisphere, leveraging the inset groove in the 3D-printed sphere to align the cut.

Step 6: The blade is used to create an incision at the top of the hemisphere to insert the pneumatic valve.

Step 7*: The pneumatic valve is inserted into the top of the hemisphere.

Step 8*: A small amount of Sil-Poxy Silicone Rubber Adhesive is applied to the bottom of the pneumatic valve to secure it to the hemispherical shell.

Step 9: A release agent (Ease Release 200 spray, Mann Release Technologies) is applied to a shallow cylindrical well.

Step 10: The uncured polymer Ecoflex 00-30 is mixed, poured into a cylindrical well, and made to cover the entire bottom of the well.

Step 11: The hemispherical shell is placed within the cylindrical well and immersed in Ecoflex 00-30 to ensure a sufficient bonding area to form a robust seal while the polymer cures to form the membrane.

***** denotes steps that vary depending on the application of the hemisphere. If the hemisphere needs to be more durable (as is needed to operate in constrained environments, grasp heavier objects or a very large number of objects), then a more robust pneumatic valve is used. In this application, an additional threaded plastic component is placed on the interior section of the hemisphere and tightly screwed to the valve on the exterior of the hemisphere, tightly sandwiching the shell and enhancing the strength of the bond.

This manufacturing method was inspired by Lee et al.'s method [49] of manufacturing thin spherical shells. They proposed a method for creating thin spherical shells made of silicone materials of varying thicknesses by pouring partially cured solutions of various silicone materials over a spherical shell. They further developed a mathematical model that predicts the thickness of a shell h_f based on the radius of the sphere R, the cure time of the material being poured τ_c, the waiting time from when the material is mixed to when it is poured τ_w, the density of the material ρ, the initial viscosity of the fluid μ_0, gravitational acceleration g, the zenith angle ϕ, and the fitting parameters α and β.

$$h_f \approx \sqrt{\frac{3\mu_0 R}{4\rho g K}}(1 + \frac{\phi^2}{10}) \tag{2.1}$$

where

$$K = (\frac{e^{-\beta\tau_w} - e^{-\beta\tau_c}}{\beta}) + (\frac{\tau_c e^{-\beta\tau_c}}{\alpha - 1}) \tag{2.2}$$

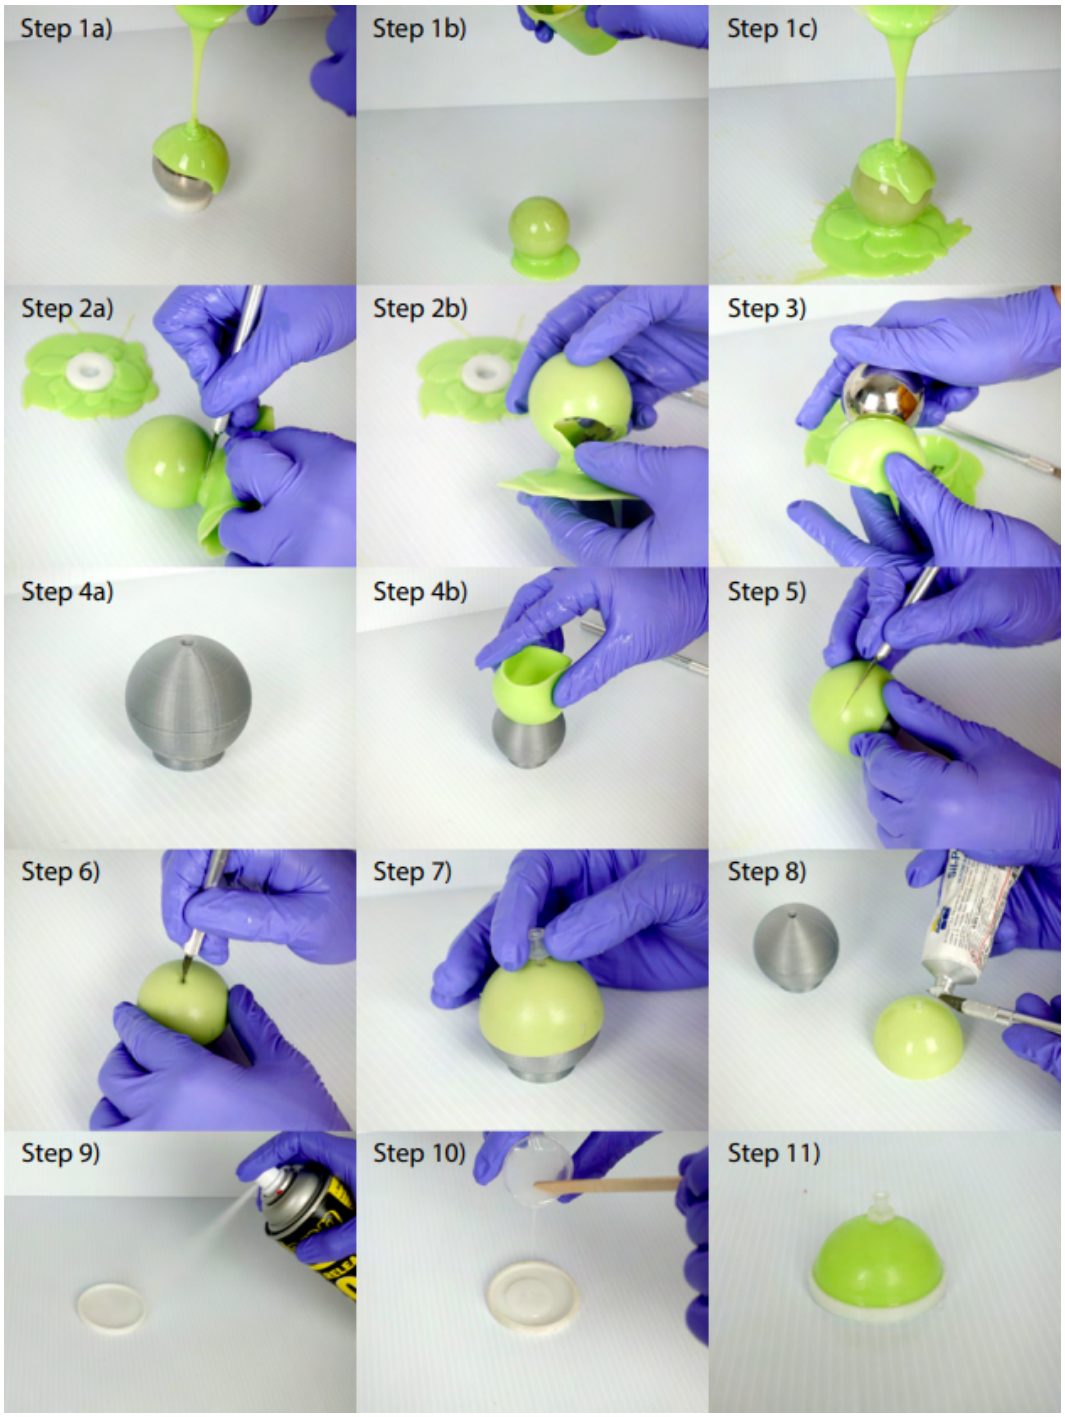

Figure 2.2: Snapshots of the manufacturing process of the hemispherical gripper [47].

We utilize these models to calculate the appropriate τ_w and number of pours we should use to create a hemisphere of a particular thickness. Since we implement successive pours to construct

our thicker hemispheres, we use the following modified equation to calculate our final hemisphere thicknesses h_f.

$$R_N \approx R_0 + \sum_{i=1}^{N} \sqrt{\frac{3\mu_0 R_i}{4\rho g K}}(1 + \frac{\phi^2}{10}) \tag{2.3}$$

$$h_f = R_N - R_0 \tag{2.4}$$

Here, R_N denotes the thickness after N pours and R_0 denotes the original radius of the mold which is a 0.025 m metal sphere. For modeling the Elite Double 32 material we assume $_0 = 7.1Pa{\cdot}s$, $\alpha = 5.3$, $\beta = 2.06{\cdot}10{-}3$ (as was used by Lee et al. for this material), and $\tau_c = 574s$ and apply equations (2.3) and (2.4) to validate our process as shown below in Figure 2.3. Note that while the relationship may appear linear, the theoretical equatorial thickness is an exponential curve that cannot be neglected for the manufacturing process. The three experimental points shown dictate the three hemispheres that will be the focus of this chapter, each demonstrating a different buckling behavior discussed later.

Although this was the method eventually selected for manufacturing the hemispherical shells, it was not the only method that was attempted. Originally, hemispheres were attempted to be manufactured using a mold-and-cast technique using 3d printed molds. While this technique was also able to moderately successfully create thin shells, as seen in Figure 2.4, it was found to impose defects in several unacceptable ways:

1. Centering of the hemisphere was extremely critical and challenging as an offset of even a tenth of a millimeter would result in wildly inconsistent hemisphere thicknesses.

2. Manufacturing thinner hemispheres of thicknesses around half a millimeter is extremely difficult to achieve both an even coating throughout the mold and removal from the mold

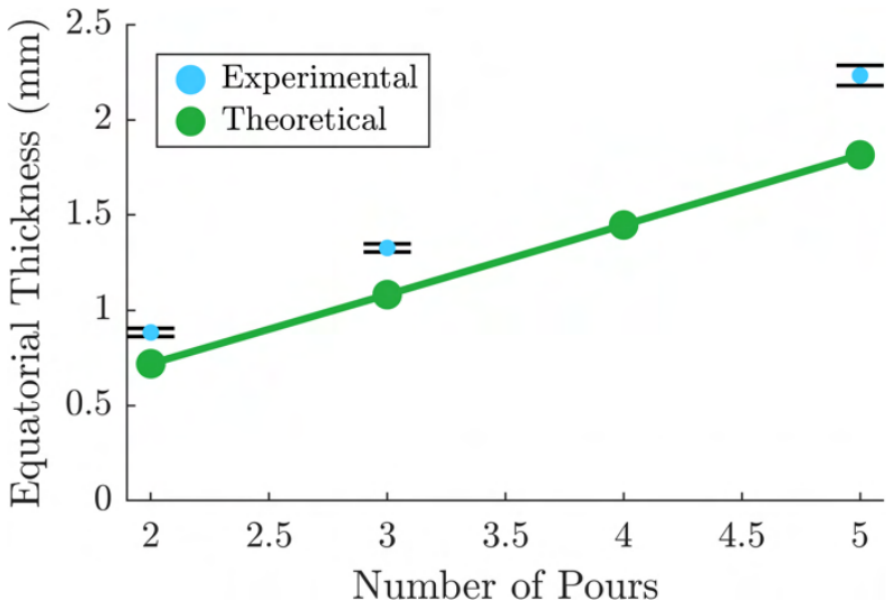

Figure 2.3: Hemisphere thickness plot demonstrating the differences between our manufactured shell thicknesses and those predicted by lee et al. [49].

after curing.

3. Grooves from the FDM printed mold can be seen throughout the hemisphere and can lead to anisotropic material properties, which would compromise the agreement with the analytical and computational models.

4. Spillover from excessive material would lead to defects around the equatorial line, which were difficult to remove

5. Any change in hemisphere thickness would require manufacturing a new mold largely increasing the manufacturing time and reducing flexibility.

Figure 2.4: A hemispherical shell manufactured using a mold-and-cast technique

2.3 Control of Imperfections

It is important to note that these comparisons required extremely precise manufacturing of the hemispheres. Since buckling involves the process of leaving a quasi-stable state when loading is applied, it is an extremely sensitive process to structural imperfections [50]. More specifically, the onset of buckling can occur under lower loads when imperfections are apparent and can also cause higher modes to reveal themselves in the buckling behavior compared to a *more perfect* structure [51]. This discrepancy between critical buckling loads of structures with and without defects is conventionally characterized by a knockdown factor—which refers to the reduction of a critical buckling load from a particular structural defect [52]. These knockdown factors are often very high, so even a small defect can have a very significant effect on the buckling behavior of a structure.

This is particularly evident in the case of buckling-enabled thin-shelled hemispheres, as buckling behavior is highly sensitive to imperfections. Even minor defects in thin shells can have a pronounced impact and are often difficult to eliminate. These challenges significantly jeopardize the validation of analytical and numerical models. Consequently, the control and accurate representation of the imperfections posed a major challenge in this project.

To avoid imperfection, quality control of manufactured hemispheres was essential. After the hemisphere was cut along its equator (Step 5) in the aforementioned manufacturing process, four measurements were taken in 90° intervals across the base of the equator using calipers. If the range of these measurements exceeded a 0.1 mm, the hemisphere was considered too asymmetrical and not suitable for testing.

The perfection of the Ecoflex seal was perhaps more difficult to assess. Indeed, it is also very difficult to achieve an accurate measurement of the Ecoflex film thickness once it is adhered to the hemisphere, so a deconstructive measurement procedure was necessitated. After the curing process, some hemispheres were selected to be scrapped and their Ecoflex layer was removed, cut in half, and then measured with calipers across the diameter. This allowed for accurate measurements across the center of the Ecoflex where it was noticed that a meniscus effect could occur, leading to higher thicknesses along the boundary of the hemisphere. Other unacceptable defects included a slight shift in the Ecoflex, which resulted in one side being slightly thicker than another. Both of these defects were controlled by allowing the Ecoflex to settle at least 10 minutes before placing the hemisphere into the Ecoflex well and ensuring they were cured on a perfectly level work surface.

To measure the thickness of the Ecoflex layer without the need for destructive measurements (measurements using calipers would require folding it over to measure, which yields low-accuracy

results), the weight of the hemisphere was noted before and after the addition of the Ecoflex seal (step 10). The relationship between the weight and the thickness of the Ecoflex is found to be

$$t = \frac{\rho * R^2 * \pi}{\Delta W} \tag{2.5}$$

where t is the thickness of the Ecoflex, ρ the density of Ecoflex (0.934 g / ml), R the radius of the hemisphere, and ΔW the change in weight after the addition of the Ecoflex. This expression was then validated using the measurements conducted from the deconstructive trials to ensure its accuracy.

Defects between the Ecoflex membrane and the Elite Double 32 shell also needed to be accounted for. Physical inspection of the trials would suffice here, but it was often noticed that the hemispherical shell could preemptively buckle if it was too thin, as seen in Figure 2.5. To prevent this, the diameter of the well containing the Ecoflex was tightened to ensure that the well would support the equator of the shell during curing, allowing for perfectly circular hemispherical boundaries.

2.4 Numerical Modeling

Numerical modeling was conducted by Kieran Barvenik in order to construct an effective design space for the gripper and demonstrate our understanding of constitutive modeling and the ability to predict its behaviors. Finite Element (FE) analyses were conducted using the commercial package ABAQUS 2020/Standard. Four-node quadrilateral shell elements (ABAQUS element type: S4R) were used to mesh the hemispherical gripper, while a nearly incompressible neo-Hookean constitutive model [53] with Poisson's ratio $\nu = 0.4998$ and shear modulus $G =$

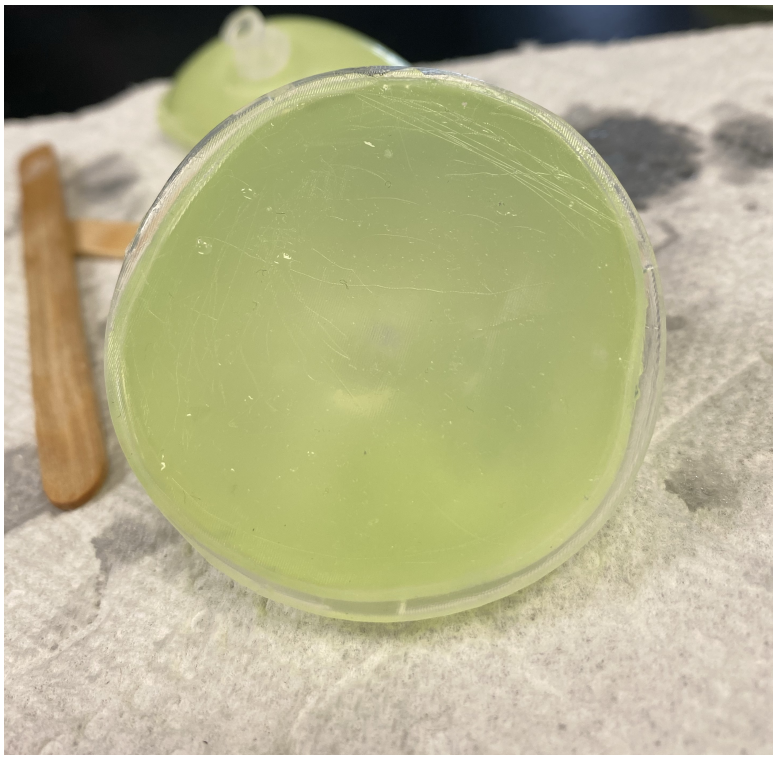

Figure 2.5: A thin hemisphere that was not properly supported in the Ecoflex sealing process.

0.375 MPa was used for the EliteDouble 32 portion of the shell. The Ecoflex membrane was modeled using an incompressible Gent material model [54] with shear modulus $G = 19.43$ kPa and extension limit $J = 37.54$.

The hemisphere is given a single point boundary condition where the pneumatic valve would be inserted in a physical sample and a dynamic implicit solver is implemented to model the behavior of the hemisphere as the volume of the fluid within the shell in decreased. This, in turn, creates a uniform pressure gradient which leads to the buckling of the shell. To ensure a quasi-stable configuration is not found, a mixture of geometric imperfections obtained from a buckling analysis of the shell (similar to the problem described by Equation (3.1)) with both a free and clamped equator was imposed. This acts as a stand-in for physical defects to initially deform the hemisphere just enough to allow for a realistic representation of practical behavior

dictated by uncontrollable defects.

Through preliminary numerical simulations, we observe the hemisphere's response to actuation will take one of three behaviors. A shell with a particularly stiff (thicker) Ecoflex layer shifts the buckling problem towards one characteristic of a thin shell with a fixed equatorial boundary condition. This results in a largely non-useful destructive buckling mode. Alternatively, if the Ecoflex is too compliant (too thin) when compared to the remainder of the shell, its deformation will dominate the response and completely fill in the fluid cavity without allowing the hemispherical shell to buckle. We refer to these behaviors as destructive buckling and film-deformation, respectively, while the optimal behavior in which the hemisphere behaves as intended is denoted as constructive buckling.

2.5 Experimental Validation and Construction of the Design Space

To validate the numerical results, experimental samples to represent each of the behaviors described above were compared to corresponding numerical simulations. These experimental specimens include a deconstructive buckling sample ($\bar{h} = 0.032$, $\bar{t} = 0.076$), a constructive buckling sample ($\bar{h} = 0.058$, $\bar{t} = 0.040$), and a film-deformation sample ($\bar{h} = 0.094$, $\bar{t} = 0.034$). Each gripper was deflated at a constant rate (70 ml/minute) using a syringe pump while an MPX7002 pressure sensor monitored the pressure, and video was captured using a Sony RX100V camera. The syringe pump withdraws 25 ml of air from the shell before pausing and reinflating the hemisphere. A MATLAB script synchronizes these instruments and processes the data into a pressure-volume curve synchronized to the video feed.

We then compare these physical results with a numerical render of the hemisphere's buckling

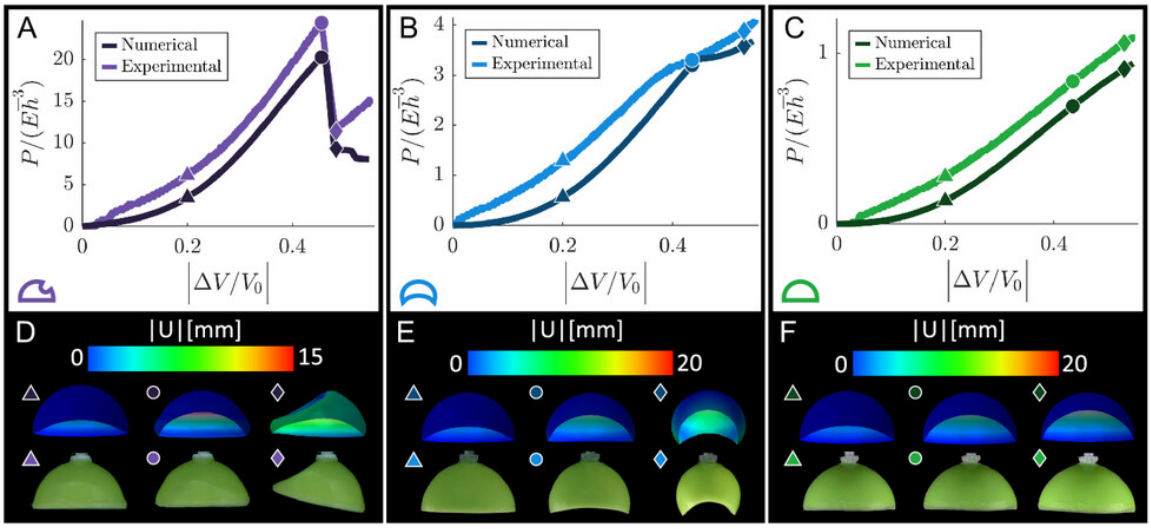

Figure 2.6: (Numerical and experimental results of the PV response and observed deformations for the deconstructive (A), constructive (B), and film-deformation (C) hemispheres.

mode to ensure similar buckling mode shapes and PV curves are observed. The numerical and physical results align closely in the observed buckling modes, the behavior of the PV curves, and the magnitude of the critical buckling pressure [1].

Once the manufacturing technique for the gripper was developed, a wide sweep of differing shell and equatorial membrane thicknesses was produced in an effort to understand what makes an effective gripper. This process was started with the following notion: a thinner gripper would have a lower overall stiffness and would therefore exert less force on objects it interacts with and require less pressure to create this deformation. Alternatively, a thicker gripper would generally have a higher stiffness and would exert more force on the objects it grasps and require more pressure to achieve these deformations. It was generally found that stiffer grippers that exert more force on the objects they grasp were better at grasping tasks; however, certain applications that require a delicate touch may require a thinner actuator. These hemispherical grippers are also

[1] the pressure at which buckling begins as defined by an inflection point in the PV curve. Note this only occurs in the deconstructive and constructive hemispheres.

scale-independent, so the design can be scaled up for the grasping of larger objects or significantly decreased for use on small delicate objects or in minimally invasive surgical procedures.

However, regardless of the desired result, we found that careful tuning of the hemisphere's shell and equatorial thickness was required to obtain the desired grasping buckling behavior. To better understand the behavior, comprehensive numerical sweeps of hemisphere parameters were simulated to create the parameter space seen in Figure 2.7. These numerical results could then be used as a tool for the construction of physical grippers.

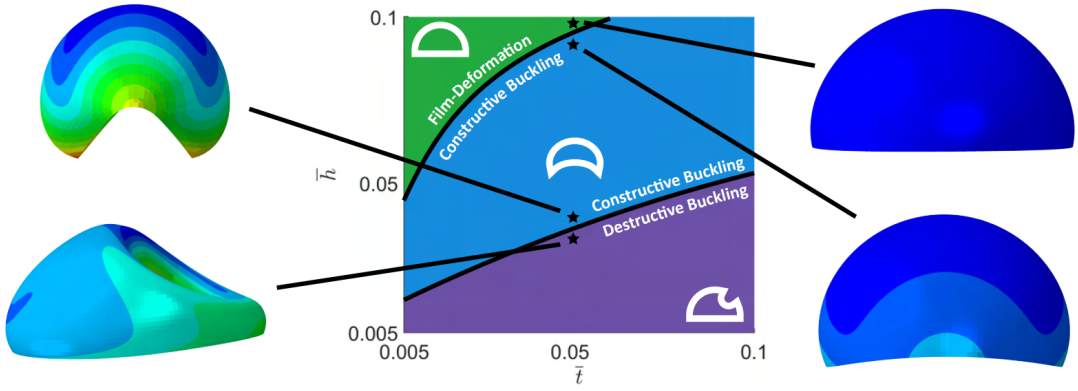

Figure 2.7: Identification of the different buckling regimes as found from the numerical results. The black lines demonstrate the boundaries between the three buckling regimes the hemisphere may envelop [47].

2.6 Gripper Functionality

This gripper is unique from other soft grippers largely in part due to its extremely sensitive nature. This can be attributed to the fact that the Ecoflex (the most compliant component) points toward the object it is grasping; we generally expect the gripper to be more sensitive than alternatives that have their least compliant components pointed toward the object. This gripper is also capable of grasping small and slippery objects, as demonstrated in Figure 2.8 , which

showcases the grasping process as well as its ability to protect and controllably release a payload. Furthermore, this gripper is capable of grabbing a wide variety of objects (Fig. 2.9), many of which are slippery or delicate and would pose significant challenges to other grippers.

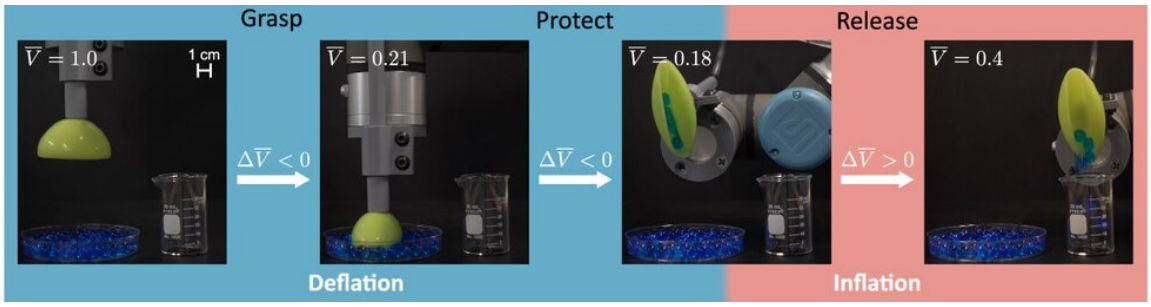

Figure 2.8: The soft gripper grabs small highly slippery *orbeez* and protects them as they are transported and controllably released into a small container [47].

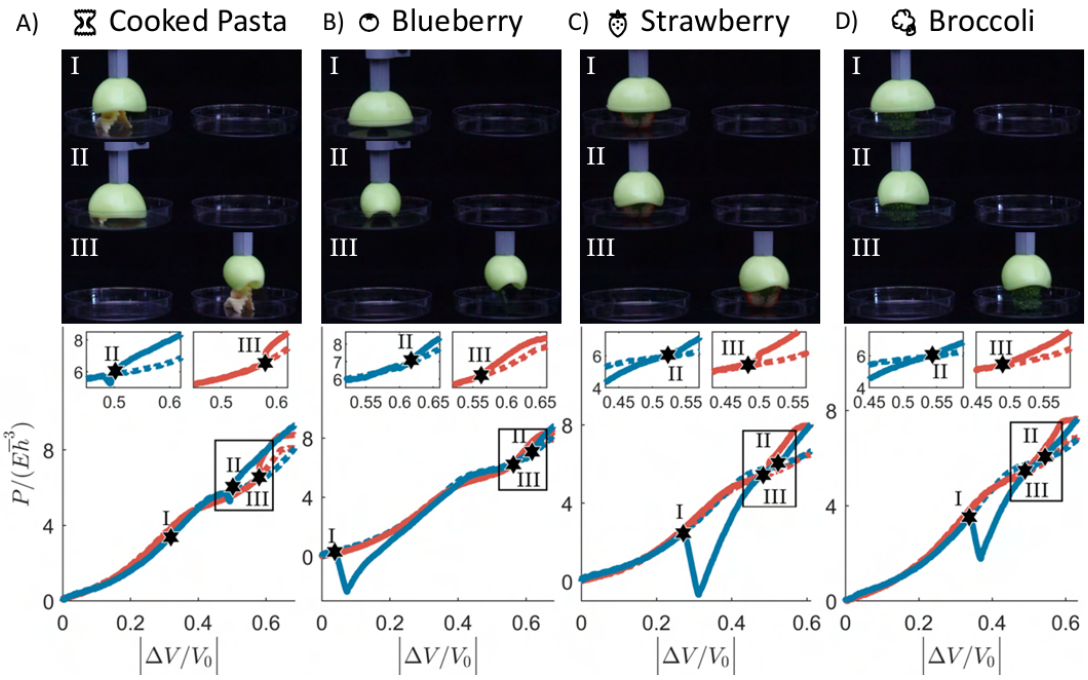

Figure 2.9: Demonstrations of the soft gripper's ability to grasp a variety of slippery and or delicate objects [47].

2.7 Conclusion

Throughout this chapter, we summarize the majority of the results found in Barvenik et al.[47] while highlighting my contributions to the manufacturing process and experimental validation. In the original work, analytical, numerical, and experimental results are compared before we demonstrate the gripper's capability for sensing tasks. These sensing tasks were limited to qualitative analyses, which determined an object's size, time of contract/grasp, as well as the time of release and will be explored and significantly expanded upon in the next chapter. We also expand beyond the original paper in this chapter through a discussion of the importance of controlling imperfections on the gripper as well as a briefly investigated proposed mechanism for harnessing intentional imperfections to increase repeatability and modeling effectiveness. Finally, we showcased some of the unique use cases of the gripper.

Chapter 3: Collection and Analysis of Pressure Data

3.1 Overview

The overarching goal of the remainder of this thesis will be to implement the soft robotic gripper described in Chapter 2 as a tactile sensor via the monitoring of change in pressure over time throughout the grasping trials. This can be done in one of several ways: some qualitative and explainable and others utilize a gray box strategy implementing vaguely recognizable patterns requiring deep learning techniques to exploit. For a trial example of the former, consider that an outside stimulus would deform the hemisphere slightly and increase pressure; therefore, monitoring for pressure increases can be implemented as a form of object detection. This can be used to determine when the hemisphere closes its *mouth* around an object and is thus ready to manipulate it and/or determine when a foreign object makes contact with the hemisphere (allowing us to determine when to actuate the gripper), a task essential for closed-loop control.

While this is perhaps the simplest form of sensing the hemispherical gripper, inspection of more pressure-volume (PV) and pressure-time curves will reveal more useful and harder-to-distinguish patterns. This section will focus first on the data collection process required to assess the sensing capabilities of the gripper, followed by a brief inspection of these data to determine qualitative patterns that could prove valuable in practical applications.

3.2 Intelligent Selection of Objects grasped

Perhaps the most challenging, novel, and useful implementation of this technique is the detection of the geometric shape of an object. While it is nearly impossible to discern patterns between all shapes —especially with the addition of noise— we can identify the effects that certain properties of shapes yield. To exasperate these patterns, objects selected to be grasped were associated with particular vibrational mode shapes in the hemisphere's buckled structure as it grasps each object. Analytical modeling of the buckling modes of the hemisphere are described by

$$u(\theta,\phi) = -A_m tan^m(\frac{1}{2}\theta)sin(\theta)cos(m\phi) \tag{3.1}$$

$$v(\theta,\phi) = -A_m tan^m(\frac{1}{2}\theta)sin(\theta)sin(m\phi) \tag{3.2}$$

$$w(\theta,\phi) = A_M(m + cos(\theta)tan^m(\frac{1}{2}\theta cos(m\phi) \tag{3.3}$$

where u, and w represent the displacements in the azimuthal and polar angles, respectively, θ and ϕ represents the cap and angle and azimuthal angle (which denotes our position on the hemisphere) m is an integer that represents the mode imposed and A_m represents a scaling factor. These buckling modes result in the shapes represented by Figure 3.1, showcasing the different analytical mode shapes of the thin hemispherical shell. In each, as the deformation increases, the hemisphere closes in around its central cavity, making this thin-shelled hemisphere an optimal geometry for a gripper.

Figure 3.2 demonstrates how varying levels of these mode shapes can be observed depending

Figure 3.1: **(A)** The first buckled mode shape of the hemisphere, which is the only naturally apparent buckling mode without boundary stimulus. **(B)** the second (3-lobed) buckled mode shape of the hemisphere. **(C)** the third (four-lobed) buckled mode shape of the hemisphere.

on the shapes that the hemisphere grasps. The four-lobed mode is brought out more while grasping the quadrilateral, and similarly, the three-lobed mode is brought out by grasping a triangular shape. The circle imposes the least change in the mode shape compared to not grasping an object and most highly preserves the first buckling mode. Inspection of the pressure-time curve exhibits how as the hemisphere enters less favorable configurations (higher buckled modes), this affects the pressure-time curves. To further *bridge* the gap between the shapes shown in Figure 3.2 , the remainder of shapes selected *meld* the differences between these shapes: trapezoids bridge the differences between a square and a triangle, and the pentagon bridge the difference between the square and the circle.

Here, it can be seen how varying the shape grasped affects the presence of different buckled mode shapes, which in turn results in changes in the experimental pressure-time curves throughout the grasping process. We notice that as the shape becomes *less circular* and contains less of the first mode shape in the deformation, there is a distinct shift in the buckling behavior around time $t = 13$ seconds.

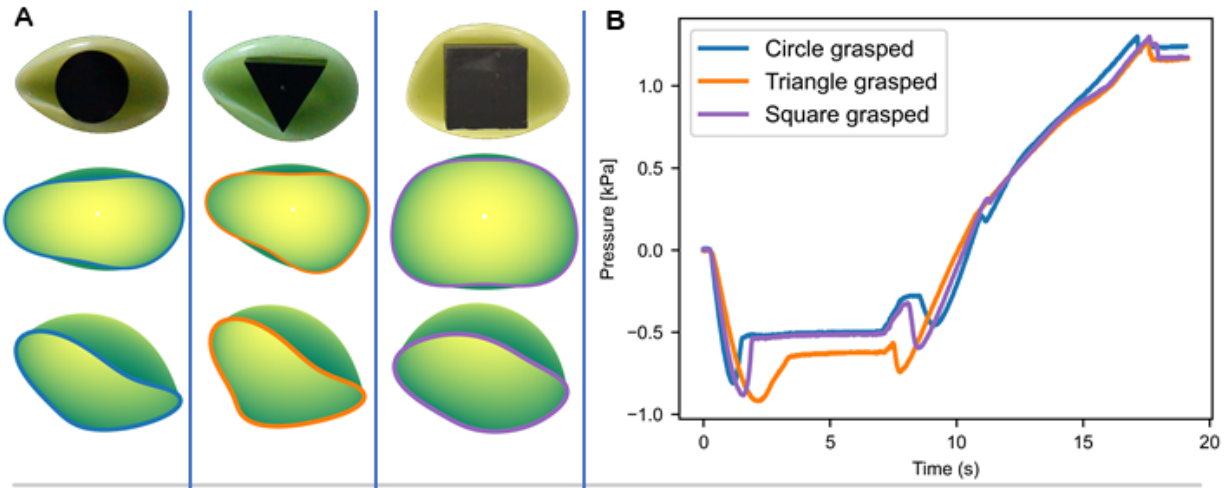

1.7	1.4	0.9	A_m (scaling factor) for mode 1
0.2	0.4	0.15	A_m (scaling factor) for mode 2
0.0	0.0	0.1	A_m (scaling factor) for mode 3

Figure 3.2: A) Comparisons of the experimentally observed grasping of the hemisphere alongside a render of the bottom-up and isometric view of the reconstructed mode shapes using eqn. (3.1). The scaling factors used for each of the mode shapes can also be seen to differ for each shape in the bottom table. B) A portion of the experimental pressure-time curve for each of the shapes shown.

3.3 Experimental Setup

The experimental setup was constructed with the primary goal of collecting extremely repeatable results for thousands of trials spanning a variety of input parameters. The final evolution of the experimental setup described in Figure 3.3 shows the physically required equipment to obtain these trials. In this setup, a single main MATLAB script serves as the control for the remainder of the system. Matlab communicates with the UR3e through a Real Time Data Exchange (RTDE) system, to the pump through serial communication commands, and to an Arduino Uno. The Arduino receives raw data from the pressure sensor and controls the servo motor, which before every trial opens a valve to equalize the pressure in the hemisphere with the

environment, allowing us to zero the system and the sensor.

A trial is defined as the grasping and release process of one object. Grasping is always conducted using volume control—meaning the pressure within the fluidic cavity of the gripper is reduced by imposing a constant decrease in fluidic volume. In each trial, the arm first centers itself over the object of interest before linearly plunging down such that the hemisphere's equatorial plane is 8 mm beneath the top of the object. At this point, the syringe pump begins withdrawing 26 ml of air from the hemisphere (approximately 80% of the fluidic cavity) at a rate of 140 ml/min[1] , which closes the hemisphere's mouth around the object. The robotic arm will then lift the hemisphere 15 mm such that it is completely suspended in the air before the syringe pump re-infuses the air into the hemisphere in the same fashion, thereby releasing the object.

This process can then be repeated an arbitrary number of times to obtain multiple trials. For our final goal of building an ML based model for tactile identification, each object was tested an average of approximately 50 times, resulting in approximately 1500 total trials. To ensure that the data varied some amount and to guarantee a more robust model, noise was added to the robotic arm's position. This accounts for the fact that in a realized application, perfect control of the location of the object may not be realistic and allows for the eventual creation of a more robust ML model. This noise is randomly selected using a uniform distribution of $\pm$ 2.5 mm along orthogonal axes in the planar direction, along with another uniform offset of $\pm$ 0.5 mm in the vertical direction. To ensure proper grasping occurs in each trial, any offset outside a planar radius of 2.5 mm was regenerated.

[1]Note that the actuation speed of the hemisphere is actually much higher than this and is limited by the speed of the pump and to reduce kinematic effects

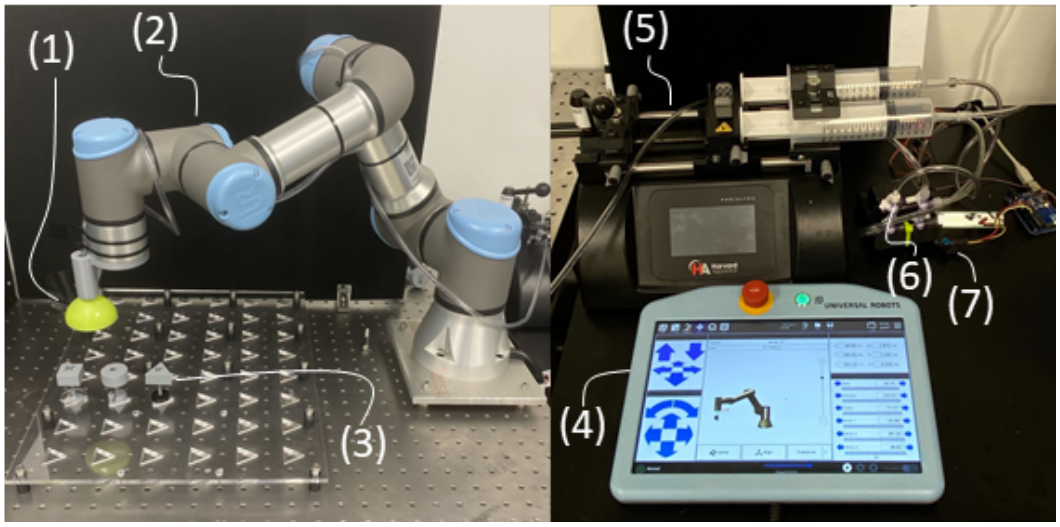

Figure 3.3: The experimental setup: (1) the soft robotic gripper, (2) a UR3e Robtic arm, (3) The shapes to be grasped, (4) the UR3 control pendent, (5) a syringe pump (PHD Ultra 4400 Series Harvard Apparatus), (6) a servo powered automatic zeroing apparatus, (7) a pressure sensor (MPXV7002DP NXP).

3.4 Object Detection

Inherent object detection is an important part of closed-loop processes for any task involving object manipulation or avoidance. In this case, we demonstrate that we can determine both when the hemisphere comes into contact with an object and when its lobes close around an object. The first case was briefly discussed previously and is trivial, as any significant disturbance to the pressure reading implies an interaction with an object. For the second case, we employ a reference curve of the hemisphere deforming without environmental disturbances to achieve a baseline of expected pressure when the gripper is in that state and compare this pressure with the experimental trial, as shown in Figure 3.4. This technique demonstrates a rudimentary form of proprioceptive sensing since any hemisphere will demonstrate repeatable deformation under a set pressure differential, assuming no external stimulus is imposed. By analyzing this deformation,

the shape of the hemisphere can be inferred at a given moment, enabling an estimate of the size of an object being grasped, specifically along the axis perpendicular to the hemisphere's buckling plane of symmetry. However, it is important to note that this method yields information about only a single one-dimensional measurement of the object. In Chapter 4, we will discuss how ML techniques can be used to achieve a more comprehensive schema of an object's size in only one grasping event.

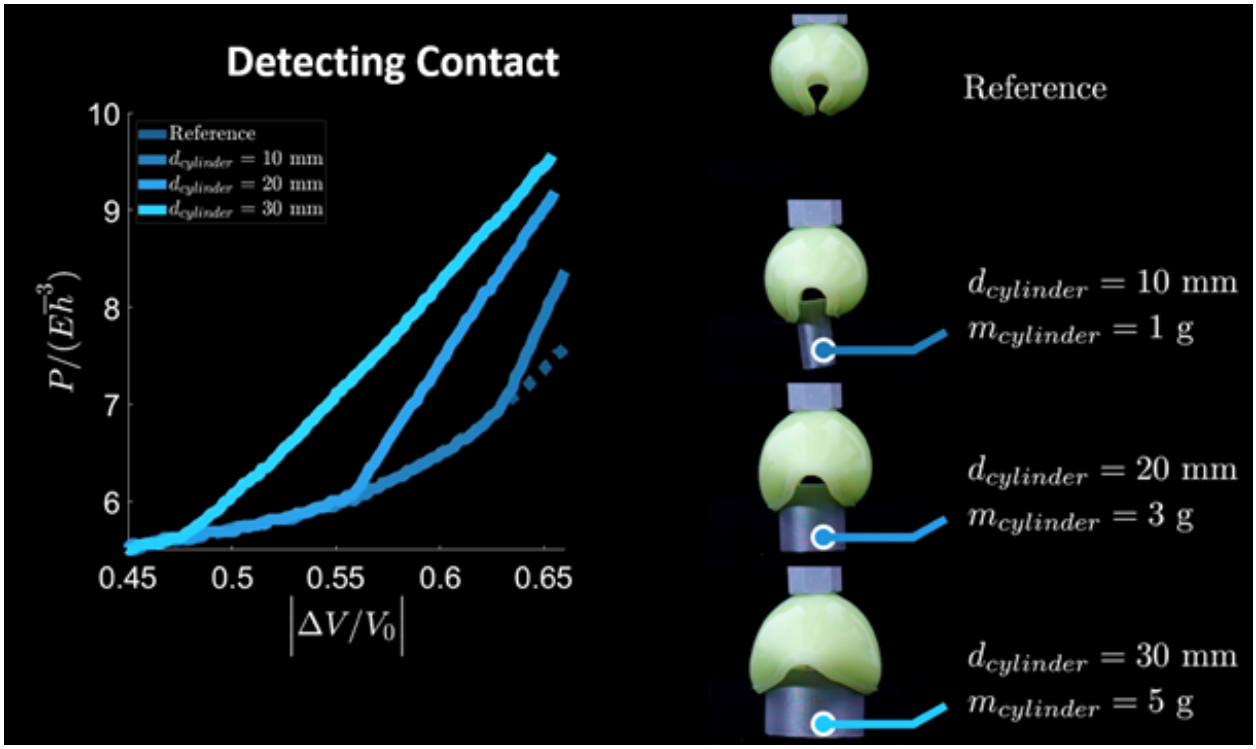

Figure 3.4: A demonstration of how the detection of the grasping of an object occurs and can be used to determine its relative diameter [47].

This pattern can also be used in a similar fashion to determine the release of an object. By monitoring when the pressure-time curve re-converges with a baseline reference trial, we can tell when the hemisphere releases an object. This is essential for closed-loop object manipulation, but we can also utilize this information to determine a grasping event failed by monitoring if the

pressure curve converges before re-inflation begins.

3.5 Cross-Sectional Area

Similarly to how we can identify the one-dimensional size of an object perpendicular to the hemisphere's buckled line of symmetry, we can use another technique to identify the approximate cross-sectional area of a grasped specimen. As the hemisphere contacts the object, there is a distinct pressure response—since the volume is still kept constant—as the object interfaces with the hemisphere, as seen in Figure 3.7. We inspect the rate at which the pressure increases with respect to the plunge depth of the hemisphere to use as a basis for approximating the cross-sectional size of an object; since it is clear that a larger object would attempt to deform the gripper more as it is pressed into it, resulting in higher pressure responses. In these examples, the fully inflated hemisphere is pressed into objects of varying sizes in preparation for the grasping task. As this occurs, we monitor the pressure-time curve in Figure 3.5 and observe a nonlinear relationship between the size of the object and the initial slope of the pressure-time curve [2]. Notice that we cannot use the maximum or steady state pressure of the deformed hemisphere since larger objects can cause a secondary buckling of the hemisphere to occur (Figure 3.6), resulting in the local pressure maxima seen in the 35 mm shape. This necessitates the utilization of the derivative of pressure with respect to time to accomplish this sensing task.

We concede that this metric does not account for an exact relationship to the cross-sectional area due to the nature of the hemisphere's deformation. In Figure 3.7, we notice that the Ecoflex boundary layer deforms more than the cross-sectional area of the shape due to its elastic behavior.

[2]Note that the change in pressure with respect to time is linearly proportional to the change in pressure with respect to the plunge depth since the robotic arm moves at a constant velocity

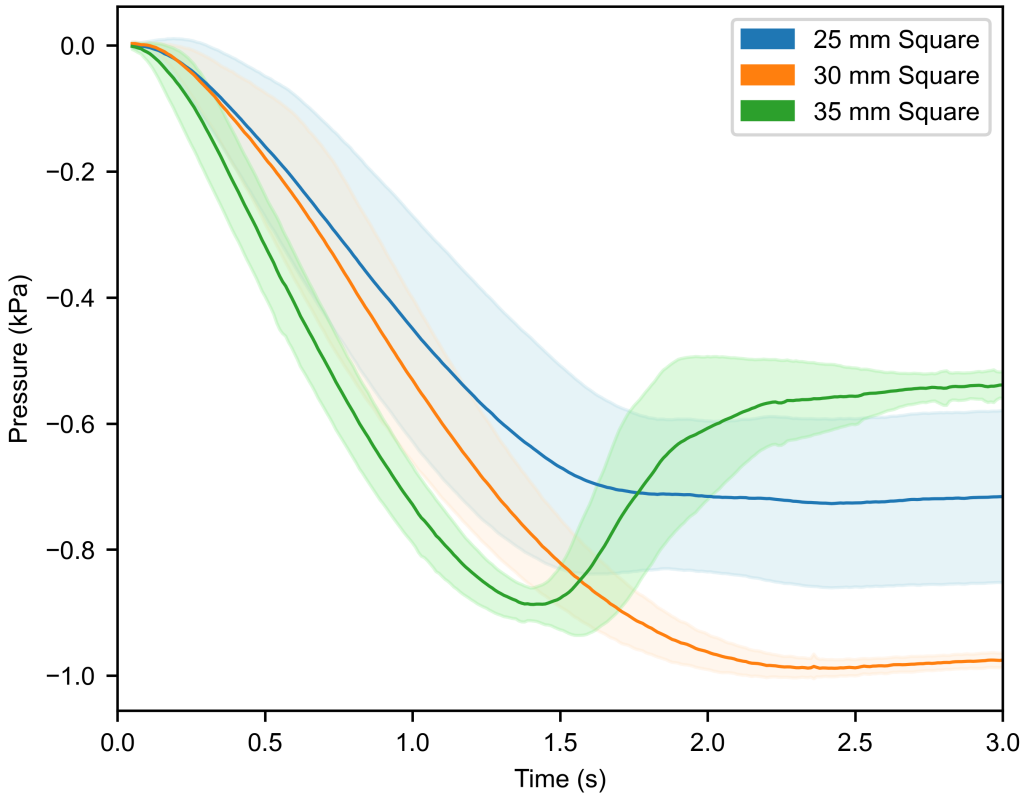

Figure 3.5: The mean and standard deviation of 25, 30, and 35 mm inscribed square shapes as the hemisphere makes initial contact with the object.

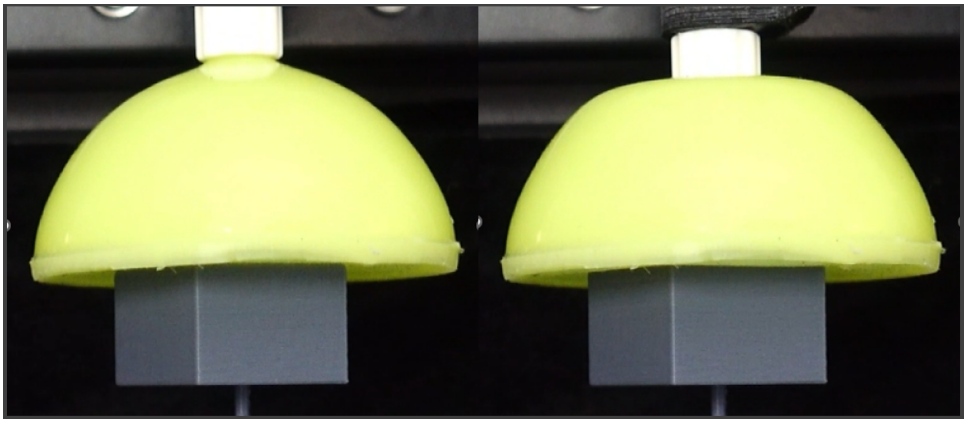

Figure 3.6: As the hemisphere is forced downwards onto this larger shape, the shell buckles, lowering the energy stored in the system and resulting in a reduction of internal pressure.

For this reason, it is perhaps more accurate to propose that this metric is correlated to a slightly higher cross-sectional area of a shape contacted, which would vary non-trivially based on the

geometry of the shape. This phenomenon would have higher effects on shapes with larger circumferences; this results in a small effect on a square and an even smaller effect on a circle. These patterns are difficult to distinguish, but it is not unfeasible to figure that an ML model could harness these patterns in its recognition processes, perhaps even beyond a metric to determine the cross-sectional area alone when correlated with other variables.

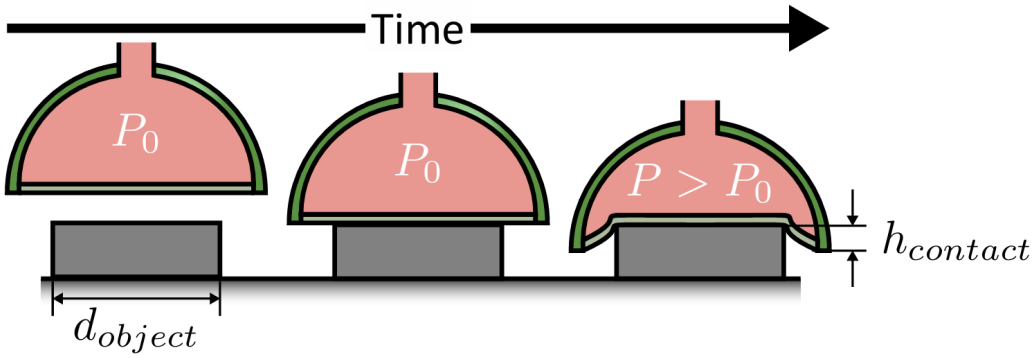

Figure 3.7: A schematic showing a cross section of the experimental process used to create Figure 3.5 [47].

3.6 Effects of Offsets

As the hemisphere is about to grab an object, some offset is inevitable. In this environment, we have a very high degree of control and precision over these offsets, however, this may not always be the case. Here, we investigate the effects that different prescribed offsets have on the pressure-time curves and see which patterns could potentially be extracted.

The effects of a variety of offsets on a triangular shape can be observed in Figures 3.8 and 3.9. It can be seen that, in all cases, there are clear shifts in the pressure curves, especially for z offsets. In the case of the Z offsets we can see the value of the initial slope is approximately equal—which is to be expected since the rate at which the hemisphere is lowered is constant—

yet we notice distinct correlation in the magnitude of the pressure around time $t \approx 5$ to the amount of extra displacement of the hemisphere. Following this event, we see this trend continue throughout the curve until the 13-second mark when the curves collapse. This phenomenon arises from the loss of contact between the hemisphere and the object as the Ecoflex membrane undergoes upward deformation due to suction forces while the mouth begins its closure. This is the case for the majority of grasping events the hemisphere undergoes throughout this study and can be determined comparing with a reference curve as shown in Figure 3.5.

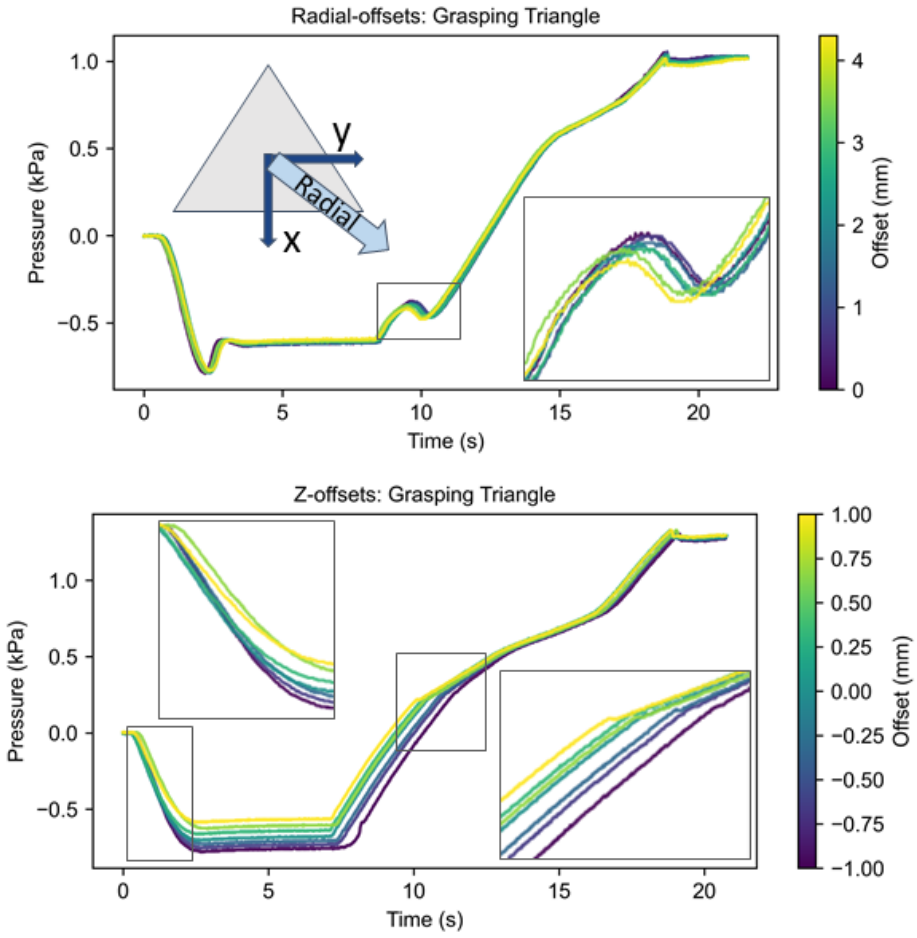

Figure 3.8: Prescribed offsets of varying magnitudes in the z axis, defined as coming out of the page, and an arbitrary radial direction defined as the diagonal vector between the x and y axes.

Investigating radial offsets, we again notice shifts in the pressure response curve, especially at the local maxima and minima. At these points, the hemispherical shell is undergoing buckling as described in Figure 3.6. In the enhanced portion of this plot, this buckling is being undone by the deflation of the hemisphere. It can again be seen that the more off-centered a hemisphere is—leading to more force exerted on the shell portion of the gripper—the earlier it returns to its unbuckled state.

Similar effects can be observed in Figure 3.9 , which both demonstrate similar patterns in the case of the Y-offset; however, for z offsets, this pattern is only apparent in the initial onset of the buckling event at $t \approx 3$ but not the removal of the buckling at time $t \approx 10$ as was observed in the other offset studies. This discrepancy could allow for the detection of not only the magnitude of an offset but also a general sense of the radial direction of this offset.

3.7 Fatigue Testing

A verification of the fatigue of the hemisphere was deemed necessary due to the large number of trials collected and the sensitivity of the hemisphere to the environment. The fatigue test consisted of a new hemisphere subjected to 2,126 withdrawal and infusion (actuation) cycles with its pressure-time response displayed in Figure 3.10. We notice a very slight discrepancy in the first 100 trials where the newer—presumably stiffer—gripper reaches marginally higher pressures ($\tilde{1}.3\%$ greater at its peak). Interestingly enough, these effects are only apparent in the withdrawal process.

While these effects are fairly negligible, especially when comparing near-concurrent trials, it is recognized that this could bias a dataset used for training ML models. In order to mitigate

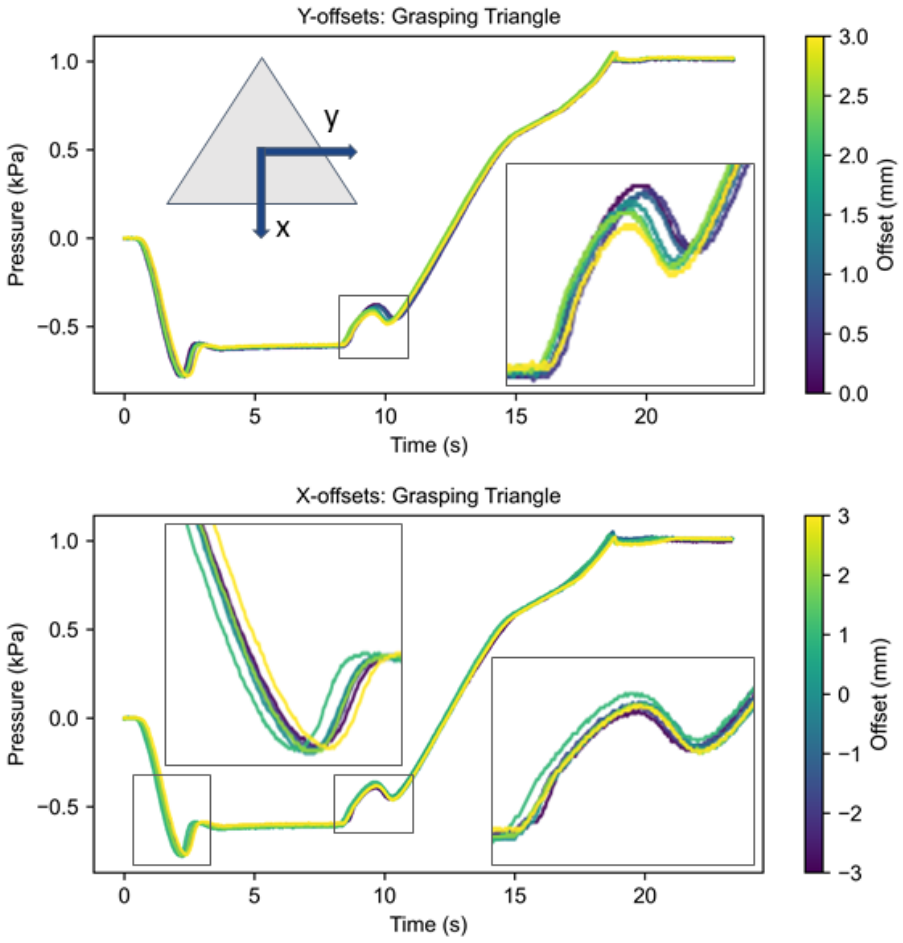

Figure 3.9: Prescribed offsets of varying magnitudes in the isolated x then y directions as defined by the axis in the diagram enclosed by the top figure

the potential effects of fatigue, the gripper used in the creation of the dataset used for ML (as described in Chapter 4) was pre-cycled approximately 200 times before being implemented. Furthermore, the order in which the trials were collected was randomized so as to prevent any biases from hemisphere age or other uncontrollable environmental conditions from tainting the dataset.

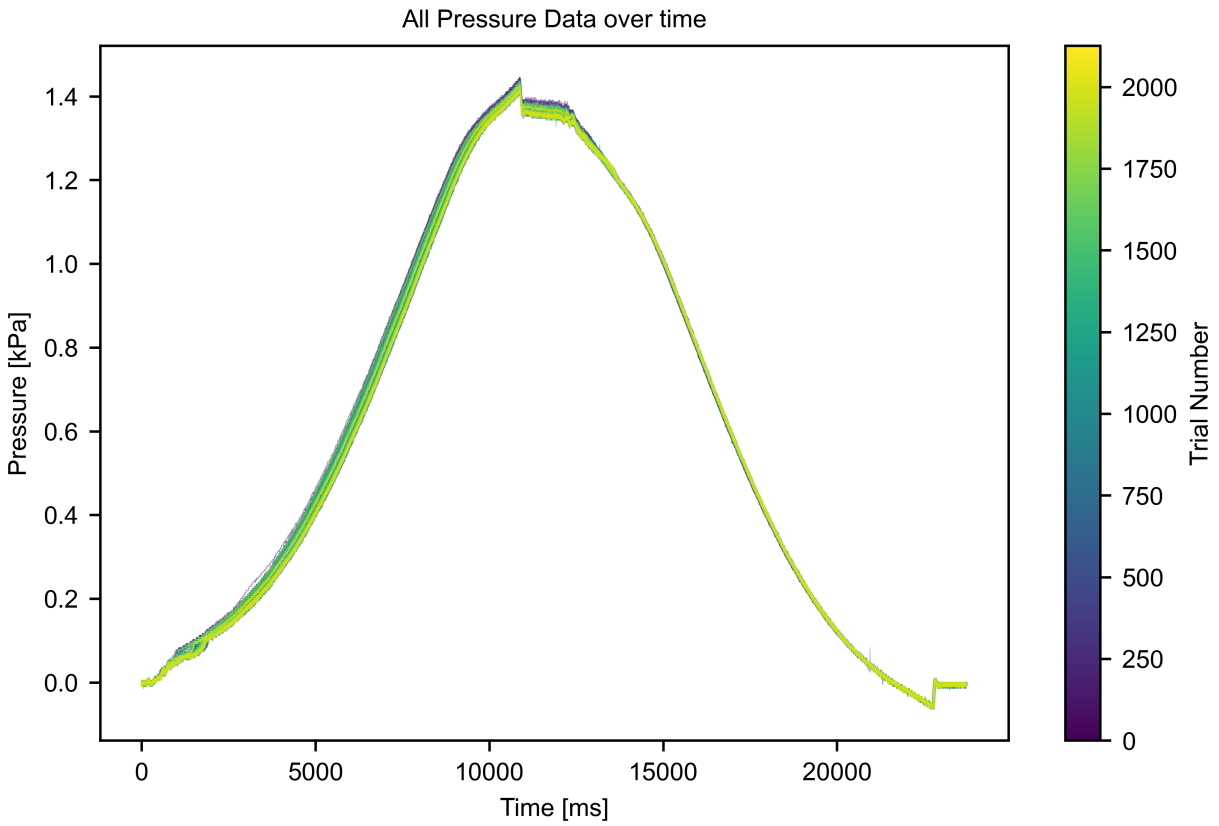

Figure 3.10: The pressure time response of the hemisphere over 2,126 inflation/deflation cycles.

3.8 Conclusion

Throughout this chapter, we have over-viewed our dataset collection process and began qualitative analysis of the collected time series pressure data for the purpose of exteroceptive sensing. We detail the hemisphere's ability to sense an object's diameter, grasping time, release time (which can in turn be used to determine grasping success), and its ability to determine offsets in targeting positions through the use of time series pressure data during actuation. We also investigated the information obtainable from the initial contact with an object and found that the change in pressure with respect to the gripper's location can be correlated with an object's cross-sectional area. This section demonstrates this soft gripper's ability to get a sense of

certain characteristics about its environment without having to modify its pre-existing gripping procedures, many of which are novel to the realm of soft robotics. We will continue this sensing investigation in the following chapter, while implementing ML to further exemplify the gripper's perceptive abilities.

Chapter 4: Harnessing Machine Learning for Object Identification

4.1 Overview

In this section, we utilize the dataset collected in Chapter 3 with the goal of developing haptic identification grasped objects using several ML techniques. We will first describe the data pre-processing procedures before training two sets of classifiers—one for geometric shape identification and the other to determine the size of the shape—which together give a holistic determination of the object that was grasped, as shown in Figure 4.1. A K-nearest neighbors (KNN), support vector machine (SVM), and a multilayer perceptron (MLP) are tested on this classification challenge. However, most of the analyses will focus on the more effective MLP. Finally, we introduce a way of recalibrating the MLP model through the use of transfer learning such that a new hemisphere can be used without the need for large amounts of training data.

In a KNN classifier, the Euclidean distance between a data point to be classified (one out of the test set) is compared to all of the data-points in the training set. The K nearest points are identified, and the point of interest is then classified according to the most frequently occurring label among these K neighbors [45]. The SVM classifier constructs a hyperplane to attempt to separate the training data based on its labels [55]. This separation is constructed by attempting to maximize the margins between classes, which helps classify boundary cases more effectively. A more powerful model is the MLP, which implements layers of perceptrons (neurons) that are

assigned weights to classify data-points, allowing for highly complex nonlinear classification patterns [56]. By exploring ML models of varying complexity, we ensure we do not unnecessarily add complexity to the problem and give a baseline for the more complex models.

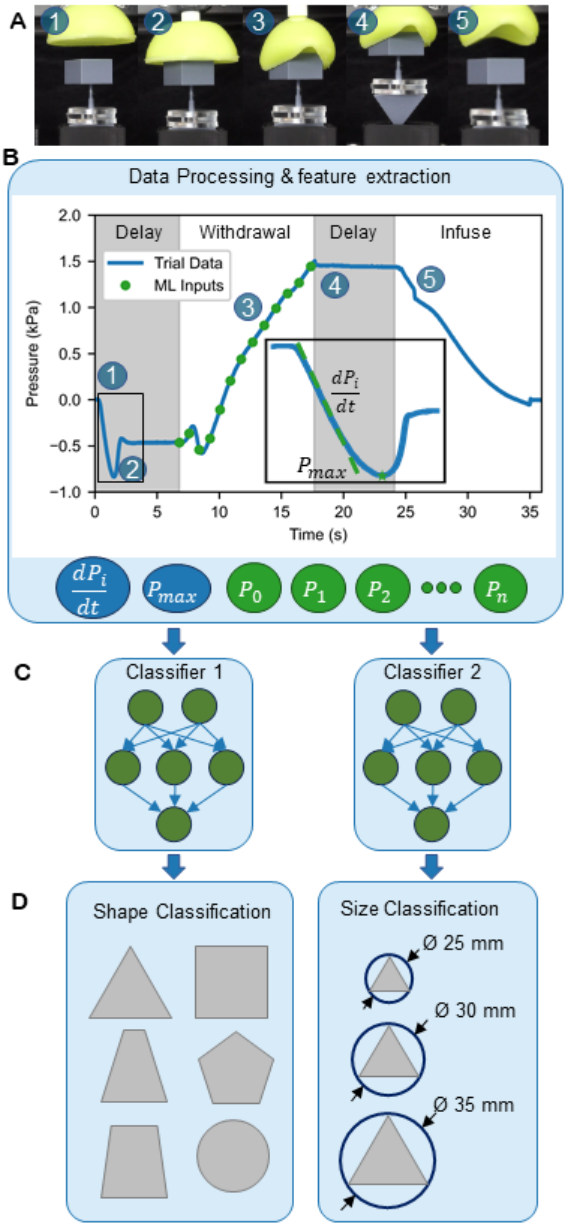

Figure 4.1: The overview of the classification process: A) Still frames of the hemisphere grasping an object. B) The pressure-time curve collected which highlights key extracted data points. C) a representation of the two classifiers used. D) the outputs of each of the classifiers

.

4.2 Data Pre-processing

The portion of the pressure-volume curve that contains the most information regarding an object grasped is the withdraw portion of the pressure curve since it is the section in which the hemisphere's deformation is most affected by the shape of the object grasped. For this reason, the only *raw* portion of the pressure data fed into the models will be this portion of the curve, defined by the times when the pump was withdrawing air from the hemisphere (actuation). This means that this portion of the curve is volume-dependent and can thereby be thought of as a PV curve, making the model easily adaptable to other withdrawal speeds. Once this portion is extracted, we down-sample the data by a factor of five (sampling rate of 14 samples/second), which was selected by qualitative inspection of the PV curve to maximize down-sampling while ensuring minimal loss of any patterns in the PV data.

While the withdraw portion of the curve contains the most information about the shape of the object, the initial portion of the pressure data also contains valuable information, as demonstrated in Figure 3.5. However, we will strategically reduce the size of our input data to the ML models by encompassing this data in only two data points.

We first extract the initial slope of pressure with respect to time, as described in Section 3.6. This process is obtained using a Python script that takes the moving average of the derivative of pressure with respect to time over the entire initial portion of the pressure curve. The max of this array is then extracted and used as the first data point in the data fed into the ML models.

Following this, the maximum pressure in the initial contact portion of the pressure data is also noted and used as the second data point fed into the ML models. Using these two data points in conjunction with the beginning of the withdrawal portion of the pressure curve gives a useful

Classification Method	Hyperparameters Selected
SVM	**Kernel**: polynomial (2^{nd} order) $\boldsymbol{L_2^2}$ **Regularization** = 10
MLP	**Hidden Layers:** (200,200) **Activation**: hyperbolic tangent **Solver:** LBFGS $\boldsymbol{L_2^2}$ **Regularization** = 2

Table 4.1: Selected hyperparameters for SVM and MLP.

representation of the initial contact portion of the dataset with only these two data points.

4.3 Training and Tuning of ML Models

The Machine Learning pipeline implemented is as follows:

1. Data is split into training (n=1,029) and test sets (n=182)
2. The training set is split using 5-fold cross-validation, and a grid search is implemented using a wide breadth of hyperparameters.
3. The model with the best validation accuracy is selected
4. Steps 2 and 3 are repeated for each model tested (KNN, SVM, MLP)
5. The test set is used to evaluate the classification accuracy of each of the models tested as shown in Table 4.2

We notice that while KNN was demonstrated to be a fairly successful classification technique, the results are better for the more complex models. Furthermore, we can see that as the number of neighbors increases in KNN, accuracy decreases. This is especially apparent in the case of shape classification, which saw a decrease of 2.7% as we increased to three neighbors. Furthermore,

ML Technique	Size Classification Accuracy	Shape Classification Accuracy
KNN (1-neighbor)	86.8%	79.1%
KNN (3-neighbor)	86.5%	76.4%
SVM	91.8%	79.7%
MLP	**94.0%**	**87.9%**

Table 4.2: Machine Learning Results for Shape Classification

while slightly more effective, the SVM also struggled to create an effective classifier, hinting that the data is not easily separable. This suggests that nonlinearities or interactions between points may be particularly important for classifying shape.

An MLP's ability to perform its own feature extraction [57] is likely what sets this method apart from the rest since the model is left to analyze the majority of the PV data on its own. One simple potential example of this would include the relationship between the extracted contact slope and the point at which the hemisphere grasps the object (as described in Figure 3.4, which would also require a model capable of its own feature selection to determine). This informs the model of the cross-sectional area and distance between the *jaws* of the gripper, respectively; this ratio would be invaluable for differentiating between, say, a triangle and a circle. Of course, the black box nature of an MLP makes discerning its feature selections very difficult, this is just one example as to what may result in its higher predictive power.

While the MLP was found to be the most effective in predicting both the shape and size of an object, less powerful methods could still potentially be used in practice if an insufficient amount of data to train the deeper model was available. Regardless, Figure 4.2 will focus on the selected MLP model for both classifiers.

In Figure 4.2 A and D, we typically hope to see classification (denoted with higher numbers)

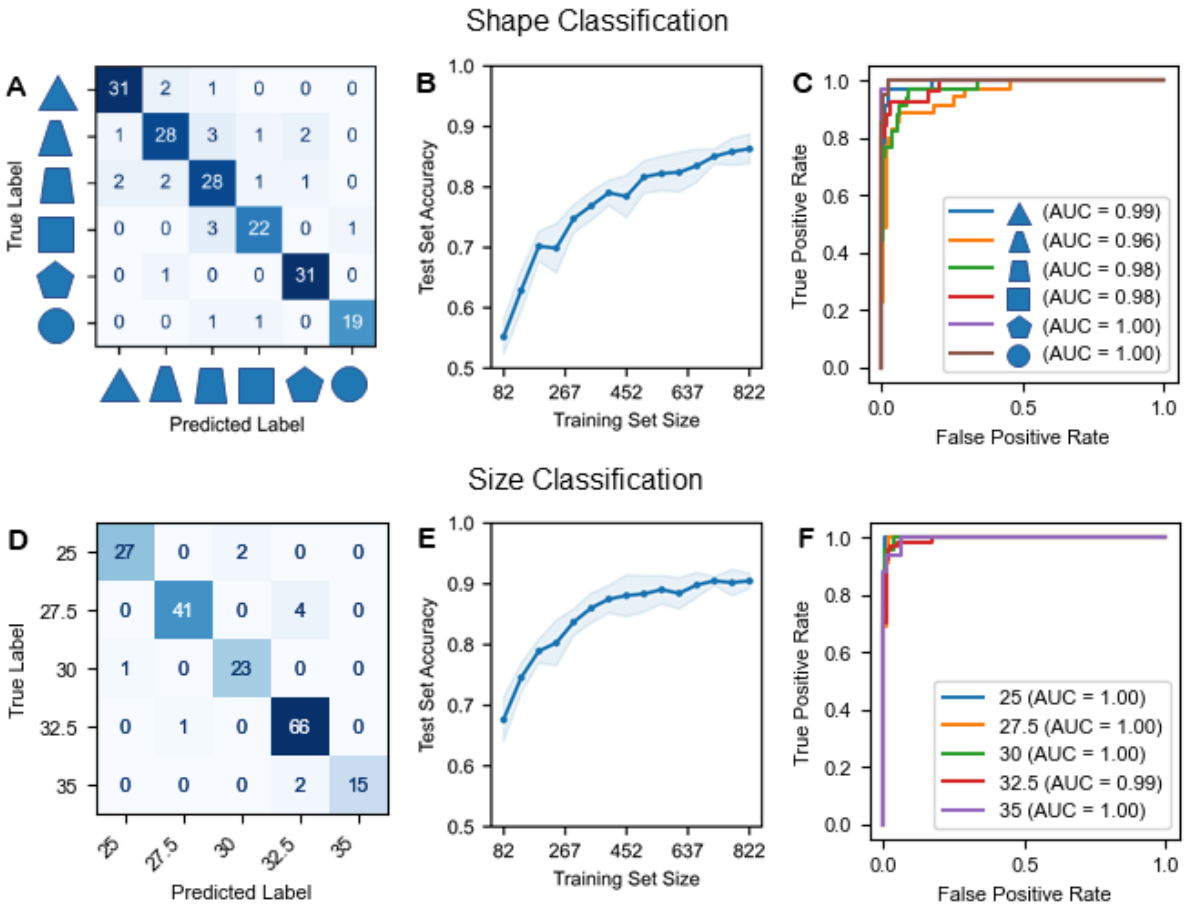

Figure 4.2: (A,B) the confusion matrices for shape and size classification. (B,E) a learning curve for shape and size classification. (C,F) ROC curves for shape and size classification.

on the diagonal of the confusion matrix, as this indicates a successful prediction. However, in the case of classifying shapes (consider that a triangle is much closer to a trapezoid than a circle) and sizes, it is also beneficial to be close to the correct guess, which is shown here by landing close to the diagonal. Here, we see that this is indeed the case for both classifiers, especially the otherwise less accurate shape classifier. We also notice in both this confusion matrix and the Receiver operating characteristic (ROC) curve (Fig. 4.2 (C, F)) that the trapezoids are generally harder to classify. This is expected, as these shapes induce fewer deformations in the hemisphere, leading to data that is harder to differentiate. Similarly, when analyzing the Shape confusion

matrix (Fig.4.2-D), we notice that the model is at most off by 5 mm when predicting the size of the shapes.

Looking at the learning curves shown in Figure 4.2 (B, E) it can be observed that a *mostly sufficient* amount of data has been collected. In the case of size classification, the curve appears to be saturated, indicating that slightly more data would not be overly beneficial to increasing the accuracy of this model. However, the learning curve shown for the shapes has not quite leveled off and is indicative that a larger dataset could further increase the accuracy of the model beyond the 87.9% that was obtained. In any case, an increase in input data is always beneficial to a model and could certainly be investigated in future works, especially in expanding the dataset to include different shapes or sizes. Nonetheless, this serves as a proof of concept for the ability to use deep-learning techniques on passively observed pressure data to extract geometric properties from a grasped object. Future works would likely include the specialization of these classification techniques to employ this *smart* gripper's sense of touch in applications such as minimally invasive surgeries or quality control in manufacturing.

4.4 Recalibration of ML Models

In practice, it would be inconvenient to repeat this training set data collection process every time a new gripper was produced, especially since each gripper's size and thickness could be tuned to fit a specific use case. Furthermore, if a gripper's properties slightly change over time, which prevents the ML model from behaving effectively (either due to fatigue or small damages), it would be beneficial to be able to easily recalibrate the model.

In Figure 4.3, we demonstrate how two newly manufactured hemispheres with intentionally

varied thicknesses could be implemented without the need to collect thousands of grasping trials. These thicknesses could hypothetically be varied either by accidental defects in manufacturing the hemisphere, repairs to an existing hemisphere, or an intentionally modified hemisphere with thicker or thinner shells more specialized for grasping a particular object [1]. Regardless of the application, Figure 4.3-A demonstrates the differences in pressure response as an arbitrarily selected triangular shape is grasped. This significantly challenges current models, yielding accuracies that are only marginally above random chance. To overcome this, a recalibration technique is employed through a process known as homogeneous transfer learning [58].

In this process, the amount of data required for a training set can be vastly reduced if the model already has some knowledge of how to classify the object. Since the input and output of the model remain unchanged, this becomes especially feasible. This is akin to a person's ability to transfer skills in, say, playing a guitar to playing bass. Both are distinct skills that require practice to learn; however, many of the skills learned from one—such as reading music—can help the training process of learning the other. In this case, the introduction of a new gripper slightly changes the classification processes, but we can use the old models' weights to serve as the pre-existing experience.

To perform the transfer learning, we first collect a new dataset using each of the newly constructed hemispheres. This dataset included all shapes but was limited to only sizes 25, 30, and 35 mm in order to speed up the data collection process. Data was collected until shape classification accuracies were found to be at least within 10% of the original set model, resulting in datasets containing 485 and 225 trials, which were then split 80%-20% for training to test. The

[1] thicker hemispheres are general better suited for larger/heavier objects, while thinner hemispheres are optimal for smaller, more delicate, and lighter objects.

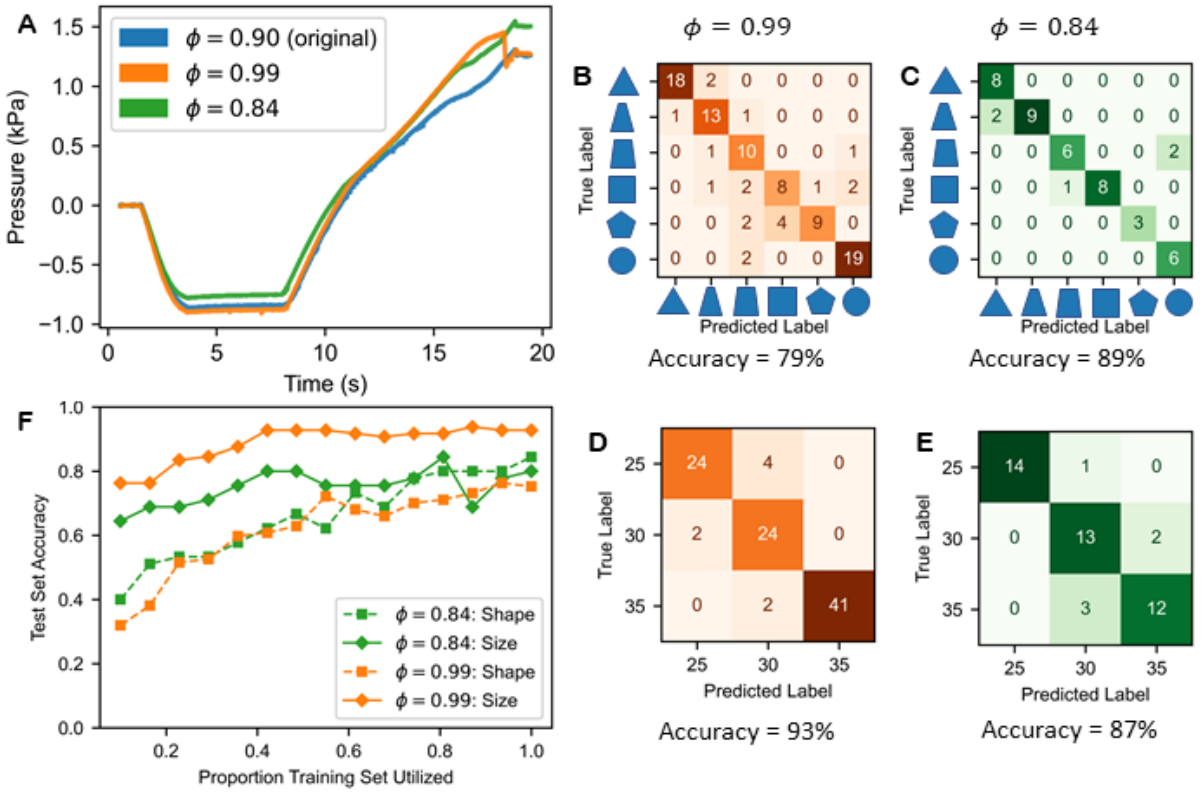

Figure 4.3: Demonstrates the potential for implementing the original model on new hemispheres using smaller recalibration datasets of size 485 and 225 for $\phi = .99$ and $\phi = 0.84$, respectively. (A) Differences in pressure response, (B,C) show the size confusion matrices, and (D,E) shows the shape confusion matrices of the two new hemispheres.

learning curve in Figure 4.3-F indicates that despite the decrease in dataset size, the saturation point is still approached. This training set is then used to train a new MLP classifier with the following exceptions:

1. The initial weights used are that of the previously selected MLP model
2. The learning rate is reduced from 10^{-4} to 10^{-6}
3. Training proceeds for 800 epochs

In this manner, we ensure that the model starts with something close to an ideal classifier

and can then nudge this in the correct direction using the much lower learning rate. This allows the same patterns that were previously recognized to ideally be tuned to fit with the new model. Going back to our guitarist analogy, we can imagine the muscle memory required to fret the bass is only slightly different and would require minimal training to relearn. The confusion matrices of the results can be seen in Figure 4.3 (B,C,D,E), which again showcases the new models' comparable accuracies.

4.5 Sorting Demonstration

Demonstrating high levels of accuracy within the dataset is promising, but to prove the effectiveness of a model deployment is always the final test. To prove the effectiveness of our model in predicting shape—the more challenging classification problem—we showcase its ability to sort several shapes by size in Figure 4.4. In this setup, the gripper is moved to the first shape, grasps it, and is lifted. At this point, a Python script is called to preprocess the data, feed them into the shape classification model, and output the predicted shape. The UR3 then moves the shape to its labeled location before the shape is released, and the process is repeated. This demonstration proves the

4.6 Conclusion

Through the application of ML techniques, we have significantly furthered the ability of the gripper to sense and even identify specific objects based on their shape and size. ML has often been applied to complex sensing problems in soft robotics; however, the monitoring of internal pressure during actuation has had limited investigation and has never been explored with the help

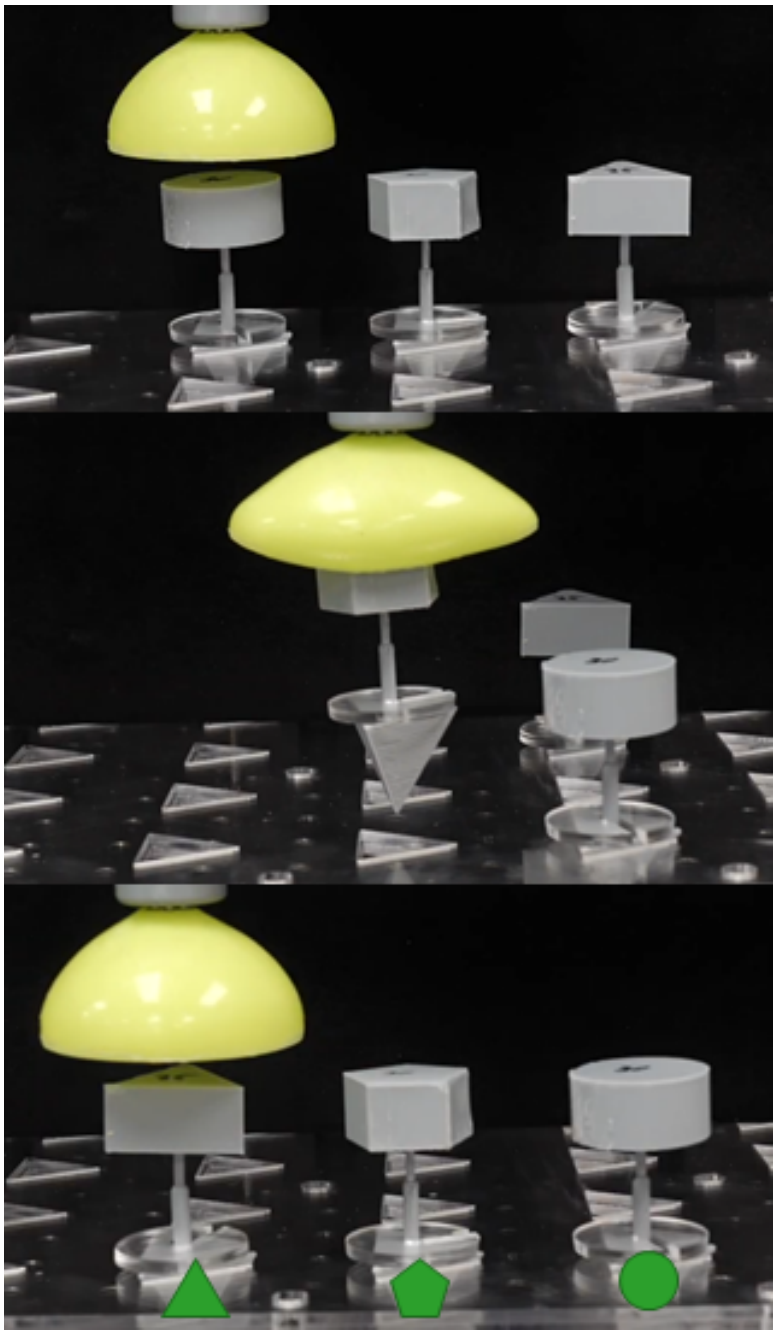

Figure 4.4: Snips from a brief demonstration which sorts shapes from *least to most circular*.

of ML techniques. This advancement could enable real-time object sensing in applications such as surgical procedures or quality control in manufacturing without requiring modifications to the gripper design or its operational sequence.

Chapter 5: Conclusion

Throughout this thesis, we have outlined the usage of a new class of soft robotic grippers capable of extracting information from one-dimensional time series pressure data during normal grasping tasks. We first overview the work "Tactile Sensing and Grasping Through Thin-Shell Buckling" [47], of which I am the second author. In the second chapter, we demonstrate our ability to accurately predict and tune its behavior using numerical modeling, which is then validated with experimental trials.

In Chapter three, we finish discussing the sensing investigated in [47] and begin to expand upon these findings with our much larger dataset of time series pressure data during a wide variety of grasping events. This dataset was collected with the intention of training ML datasets and, therefore, is excessively large and contains intentional noise by reducing the precision of the grasping location (intentionally missing the center of objects) as would be present in practical applications. We then used a qualitative analysis of the collected pressure-time curves to extract information about the environment.

Finally, in chapter four, this dataset is harnessed for the training of ML models. We focus on two ML classifiers that work in parallel to determine a grasped object's geometric shape and size, effectively expanding upon the gripper's capacity for haptic sensing. An MLP model was found to be most effective for both of these tasks, which can also be recalibrated on similar

data sets via transfer learning. This technique was used to recalibrate the models on two newly constructed hemispherical grippers (with intentional variations in dimensions), allowing them to achieve similar perceptive abilities of the original model while requiring only a fraction of the dataset.

This paper fills a void in the literature on soft robotic sensing by employing machine learning techniques on data readily available to pneumatically actuated soft robotics. This avoids the need for expensive, delicate, or difficult to implement additional sensors in a soft robotic actuator and also avoids the challenges associated with flexible electronics. Our approach demonstrates the ability to achieve comparable object characterization without additional rotations or measurements, enabling the completion of the grasping task in a single interaction. Moreover, both shape and size classification exhibited high sensitivity, successfully distinguishing between geometries differing by only a few millimeters.

While this thesis explored only an arbitrarily selected set of shapes of a set material, a more varied dataset could be investigated to expand its breadth of sensing capabilities since any attribute that creates a predictable change in the pressure response could be investigated with these techniques. In this way, this work serves as a proof of concept of the potential for tackling major issues other soft sensors face by mitigating the need for electrodes, multimodal sensors, and wired electronic interfaces. We are also interested in specializing the sensor toward a particular application, such as monitoring the successful grasp of debris in a minimally invasive surgical application or detecting a potentially defective part on a manufacturing line. Furthermore, since it has been demonstrated that many pneumatically controlled soft robotics can capture information about their environment [46], this sensing procedure could be leveraged to retrofit existing soft robotic actuators into passive sensors. Overall, this novel implementation of ml-enabled sensing

requiring only easily obtainable pressure series data of a soft robotic gripper demonstrates the potential for faster and cheaper haptic object identification than existing methods.

Bibliography

[1] Mark Runciman, Ara Darzi, and George P. Mylonas. Soft Robotics in Minimally Invasive Surgery. *Soft Robotics*, 6(4):423–443, August 2019. Publisher: Mary Ann Liebert, Inc., publishers.

[2] Edward Gough, Andrew T. Conn, and Jonathan Rossiter. Planning for a Tight Squeeze: Navigation of Morphing Soft Robots in Congested Environments. *IEEE Robotics and Automation Letters*, 6(3):4752–4757, July 2021. Conference Name: IEEE Robotics and Automation Letters.

[3] Panagiotis Polygerinos, Nikolaus Correll, Stephen A. Morin, Bobak Mosadegh, Cagdas D. Onal, Kirstin Petersen, Matteo Cianchetti, Michael T. Tolley, and Robert F. Shepherd. Soft Robotics: Review of Fluid-Driven Intrinsically Soft Devices; Manufacturing, Sensing, Control, and Applications in Human-Robot Interaction. *Advanced Engineering Materials*, 19(12):1700016, 2017. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/adem.201700016.

[4] Nina R. Sinatra, Clark B. Teeple, Daniel M. Vogt, Kevin Kit Parker, David F. Gruber, and Robert J. Wood. Ultragentle manipulation of delicate structures using a soft robotic gripper. *Science Robotics*, 4(33):eaax5425, August 2019. Publisher: American Association for the Advancement of Science.

[5] Faheem Ahmed, Muhammad Waqas, Bushra Jawed, Afaque Manzoor Soomro, Suresh Kumar, Ashraf Hina, Umair Khan, Kyung Hwan Kim, and Kyung Hyun Choi. Decade of bio-inspired soft robots: a review. *Smart Materials and Structures*, 31(7):073002, June 2022. Publisher: IOP Publishing.

[6] Liang Zhou, Lili Ren, You Chen, Shichao Niu, Zhiwu Han, and Luquan Ren. Bio-Inspired Soft Grippers Based on Impactive Gripping. *Advanced Science*, 8(9):2002017, 2021. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/advs.202002017.

[7] Jeffrey Ackerman and Justin Seipel. Energy Efficiency of Legged Robot Locomotion With Elastically Suspended Loads. *IEEE Transactions on Robotics*, 29(2):321–330, April 2013. Conference Name: IEEE Transactions on Robotics.

[8] Kalina Bonofiglio, Lauryn Whiteside, Maya Angeles, Matthew Haahr, Brandon Simpson, Josh Palmer, Yijia Wu, and Markus P. Nemitz. Soft Fluidic Closed-Loop Controller for Untethered Underwater Gliders. In *2023 IEEE International Conference on Soft Robotics (RoboSoft)*, pages 1–6, April 2023. ISSN: 2769-4534.

[9] Michael Krieg, Isaac Sledge, and Kamran Mohseni. Design considerations for an underwater soft-robot inspired from marine invertebrates. *Bioinspiration & Biomimetics*, 10(6):065004, October 2015. Publisher: IOP Publishing.

[10] Yang Yang, Zhiguo He, Pengcheng Jiao, and Hongliang Ren. Bioinspired Soft Robotics: How Do We Learn From Creatures? *IEEE Reviews in Biomedical Engineering*, 17:153–165, 2024. Conference Name: IEEE Reviews in Biomedical Engineering.

[11] Emanuela Del Dottore, Alessio Mondini, Nick Rowe, and Barbara Mazzolai. A growing soft robot with climbing plant–inspired adaptive behaviors for navigation in unstructured environments. *Science Robotics*, 9(86):eadi5908, January 2024. Publisher: American Association for the Advancement of Science.

[12] Yegor Piskarev, Antoine Devincenti, Vivek Ramachandran, Pierre-Etienne Bourban, Michael D. Dickey, Jun Shintake, and Dario Floreano. A Soft Gripper with Granular Jamming and Electroadhesive Properties. *Advanced Intelligent Systems*, 5(6):2200409, 2023. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/aisy.202200409.

[13] Jun Shintake, Vito Cacucciolo, Dario Floreano, and Herbert Shea. Soft Robotic Grippers. *Advanced Materials*, 30(29):1707035, 2018. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/adma.201707035.

[14] Weiqiang Dou, Guoliang Zhong, Jinglin Cao, Zhun Shi, Bowen Peng, and Liangzhong Jiang. Soft Robotic Manipulators: Designs, Actuation, Stiffness Tuning, and Sensing. *Advanced Materials Technologies*, 6(9):2100018, 2021. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/admt.202100018.

[15] Aniket Pal, Vanessa Restrepo, Debkalpa Goswami, and Ramses V. Martinez. Exploiting Mechanical Instabilities in Soft Robotics: Control, Sensing, and Actuation. *Advanced Materials (Deerfield Beach, Fla.)*, 33(19):e2006939, May 2021.

[16] Yinding Chi, Yanbin Li, Yao Zhao, Yaoye Hong, Yichao Tang, and Jie Yin. Bistable and Multistable Actuators for Soft Robots: Structures, Materials, and Functionalities. *Advanced Materials (Deerfield Beach, Fla.)*, 34(19):e2110384, May 2022.

[17] B. Grossi, H. Palza, J. C. Zagal, C. Falcón, and G. During. Metarpillar: Soft robotic locomotion based on buckling-driven elastomeric metamaterials. *Materials & Design*, 212:110285, December 2021.

[18] Benjamin Gorissen, David Melancon, Nikolaos Vasios, Mehdi Torbati, and Katia Bertoldi. Inflatable soft jumper inspired by shell snapping. *Science Robotics*, 5(42):eabb1967, May 2020. Publisher: American Association for the Advancement of Science.

[19] Yang Yang and Yongquan Wang. Snapping for 4D-Printed Insect-Scale Metal-Jumper. *Advanced Science*, 11(3):2307088, 2024. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/advs.202307088.

[20] Xuan Zhang, Yue Wang, Zhihao Tian, Manar Samri, Karsten Moh, Robert M. McMeeking, René Hensel, and Eduard Arzt. A bioinspired snap-through metastructure for manipulating micro-objects. *Science Advances*, 8(46):eadd4768, November 2022. Publisher: American Association for the Advancement of Science.

[21] Yuhang Liu, Kai Luo, Shuai Wang, Xiaodong Song, Zhijuan Zhang, Qiang Tian, and Haiyan Hu. A Soft and Bistable Gripper with Adjustable Energy Barrier for Fast Capture in Space. *Soft Robotics*, February 2023. Publisher: Mary Ann Liebert, Inc., publishers 140 Huguenot Street, 3rd Floor New Rochelle, NY 10801 USA.

[22] V. Ortenzi, R. Stolkin, J. Kuo, and M. Mistry. Hybrid motion/force control: a review. *Advanced Robotics*, 31(19-20):1102–1113, October 2017. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/01691864.2017.1364168.

[23] Cosimo Della Santina, Christian Duriez, and Daniela Rus. Model-Based Control of Soft Robots: A Survey of the State of the Art and Open Challenges. *IEEE Control Systems Magazine*, 43(3):30–65, June 2023. Conference Name: IEEE Control Systems Magazine.

[24] Hongbo Wang, Massimo Totaro, and Lucia Beccai. Toward Perceptive Soft Robots: Progress and Challenges - Wang - 2018 - Advanced Science - Wiley Online Library, July 2018.

[25] Juntian Qu, Guangming Cui, Zhenkun Li, Shutong Fang, Xianrui Zhang, Ang Liu, Mingyue Han, Houde Liu, Xueqian Wang, and Xiaohao Wang. Advanced Flexible Sensing Technologies for Soft Robots. *Advanced Functional Materials*, 34(29):2401311, 2024. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/adfm.202401311.

[26] Thomas George Thuruthel, Egidio Falotico, Mariangela Manti, and Cecilia Laschi. Stable Open Loop Control of Soft Robotic Manipulators. *IEEE Robotics and Automation Letters*, 3(2):1292–1298, April 2018. Conference Name: IEEE Robotics and Automation Letters.

[27] Chuqiao Lyu, Shuxiang Guo, Yonggan Yan, Yongxin Zhang, Yongwei Zhang, Pengfei Yang, and Jianmin Liu. Deep-Learning-Based Force Sensing Method for a Flexible Endovascular Surgery Robot. *IEEE Transactions on Instrumentation and Measurement*, 73:1–10, 2024. Conference Name: IEEE Transactions on Instrumentation and Measurement.

[28] Wyatt Felt and C. David Remy. Smart braid: Air muscles that measure force and displacement. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2821–2826, September 2014. ISSN: 2153-0866.

[29] John Morrow, Hee-Sup Shin, Calder Phillips-Grafflin, Sung-Hwan Jang, Jacob Torrey, Riley Larkins, Steven Dang, Yong-Lae Park, and Dmitry Berenson. Improving Soft Pneumatic Actuator fingers through integration of soft sensors, position and force control,

and rigid fingernails. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5024–5031, May 2016.

[30] Wyatt Felt, Khai Yi Chin, and C. David Remy. Contraction Sensing With Smart Braid McKibben Muscles. *IEEE/ASME Transactions on Mechatronics*, 21(3):1201–1209, June 2016. Conference Name: IEEE/ASME Transactions on Mechatronics.

[31] Charbel Tawk, Marc in het Panhuis, Geoffrey M. Spinks, and Gursel Alici. Soft Pneumatic Sensing Chambers for Generic and Interactive Human–Machine Interfaces. *Advanced Intelligent Systems*, 1(1):1900002, 2019. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/aisy.201900002.

[32] Utku Culha, Surya G. Nurzaman, Frank Clemens, and Fumiya Iida. SVAS3: Strain Vector Aided Sensorization of Soft Structures. *Sensors*, 14(7):12748–12770, July 2014. Number: 7 Publisher: Multidisciplinary Digital Publishing Institute.

[33] Yoel Shapiro, Gábor Kósa, and Alon Wolf. Shape Tracking of Planar Hyper-Flexible Beams via Embedded PVDF Deflection Sensors. *IEEE/ASME Transactions on Mechatronics*, 19(4):1260–1267, August 2014. Conference Name: IEEE/ASME Transactions on Mechatronics.

[34] Benjamin Shih, Dylan Shah, Jinxing Li, Thomas G. Thuruthel, Yong-Lae Park, Fumiya Iida, Zhenan Bao, Rebecca Kramer-Bottiglio, and Michael T. Tolley. Electronic skins and machine learning for intelligent soft robots. *Science Robotics*, 5(41):eaaz9239, April 2020. Publisher: American Association for the Advancement of Science.

[35] Panagiotis Polygerinos, Zheng Wang, Kevin C. Galloway, Robert J. Wood, and Conor J. Walsh. Soft robotic glove for combined assistance and at-home rehabilitation. *Robotics and Autonomous Systems*, 73:135–143, November 2015.

[36] 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, November 2016. Publisher: Frontiers.

[37] C. Larson, B. Peele, S. Li, S. Robinson, M. Totaro, L. Beccai, B. Mazzolai, and R. Shepherd. Highly stretchable electroluminescent skin for optical signaling and tactile sensing. *Science*, 351(6277):1071–1074, March 2016. Publisher: American Association for the Advancement of Science.

[38] Amjadi. Stretchable, Skin-Mountable, and Wearable Strain Sensors and Their Potential Applications: A Review - Amjadi - 2016 - Advanced Functional Materials - Wiley Online Library, 2016.

[39] Yexi Jin, Hao Shen, Lining Sun, Xingwen Zhou, and Liguo Chen. Multilevel nanostructured pressure sensor for object recognition with Deep-Learning assistance: A strategy for high sensitivity and wide detection range - ScienceDirect, March 2024.

[40] Oliver Shorthose, Luca Scimeca, Alessandro Albini, and Perla Maiolino. Bioinspired Tactile Object Identification Leveraging Deep Learning and Soft Body Compliance. *Advanced Intelligent Systems*, n/a(n/a):2400802. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/aisy.202400802.

[41] Thomas Thuruthel, Benjamin Shih, Cecelia Laschi, and Michael Tolley. Soft robot perception using embedded soft sensors and recurrent neural networks | Science Robotics, January 2019.

[42] Gabor Soter, Andrew Conn, Helmut Hauser, and Jonathan Rossiter. Bodily Aware Soft Robots: Integration of Proprioceptive and Exteroceptive Sensors. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2448–2453, May 2018. ISSN: 2577-087X.

[43] Benjamin Shih, Dylan Drotman, Caleb Christianson, Zhaoyuan Huo, Ruffin White, Henrik I. Christensen, and Michael T. Tolley. Custom soft robotic gripper sensor skins for haptic object visualization. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 494–501, September 2017. ISSN: 2153-0866.

[44] Bianca S. Homberg, Robert K. Katzschmann, Mehmet R. Dogar, and Daniela Rus. Haptic identification of objects using a modular soft robotic gripper. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1698–1705, September 2015.

[45] J. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, volume 5.1, pages 281–298. University of California Press, January 1967.

[46] Shibo Zou, Sergio Picella, Jelle de Vries, Vera G. Kortman, Aimée Sakes, and Johannes T. B. Overvelde. A retrofit sensing strategy for soft fluidic robots. *Nature Communications*, 15(1):539, January 2024. Publisher: Nature Publishing Group.

[47] Kieran Barvenik, Zachary Coogan, Gabriele Librandi, Matteo Pezzulla, and Eleonora Tubaldi. Tactile Sensing and Grasping Through Thin-Shell Buckling. *Advanced Intelligent Systems*, 6(9):2300855, 2024. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/aisy.202300855.

[48] John D. Gage and Paul A. Tyler. *Deep-Sea Biology: A Natural History of Organisms at the Deep-Sea Floor*. Cambridge University Press, April 1991. Google-Books-ID: vqgfH9DXXNAC.

[49] A. Lee, J Marthelot, G Balestra, F Gallaire, and P.M. Reis. Fabrication of slender elastic shells by the coating of curved surfaces | Nature Communications, 2016.

[50] Ali Hasan Nayfeh and Dean T. Mook. *Nonlinear oscillations*. Wiley, New York, 1995.

[51] Stephan Prokof'evitch Timoshenko and James M. Gere. *Theory of elastic stability*. Dover publ, Mineola, 2nd ed edition, 2009.

[52] Anna Lee, Francisco López Jiménez, Joel Marthelot, John W. Hutchinson, and Pedro M. Reis. The Geometric Role of Precisely Engineered Imperfections on the Critical Buckling Load of Spherical Elastic Shells. *Journal of Applied Mechanics*, 83(111005), September 2016.

[53] R. W. Ogden. *Non-linear Elastic Deformations*. Courier Corporation, January 1997. Google-Books-ID: 2u7wCaojfbEC.

[54] A. N. Gent. A New Constitutive Relation for Rubber. *Rubber Chemistry and Technology*, 69(1):59–61, March 1996.

[55] Bernhard E. Boser, Isabelle M. Guyon, and Vladimir N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the fifth annual workshop on Computational learning theory*, COLT '92, pages 144–152, New York, NY, USA, July 1992. Association for Computing Machinery.

[56] Pierre A. Devijver and Josef Kittler. *Pattern recognition: a statistical approach*. Prentice/Hall International, Englewood Cliffs, N.J, 1982.

[57] Fatma Shaheen, Brijesh Verma, and Md. Asafuddoula. Impact of Automatic Feature Extraction in Deep Learning Architecture. In *2016 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–8, November 2016.

[58] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A Comprehensive Survey on Transfer Learning. *Proceedings of the IEEE*, 109(1):43–76, January 2021.