Domesticating Soft Robotics Research and Development with Accessible Biomaterials

Kiyn Nakeita Agatha Chin

CMU-RI-TR-24-65

September 2024

The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, Pennsylvania

Thesis Committee

Carmel Majidi	Carnegie Mellon University (Chair)
Zeynep Temel	Carnegie Mellon University
Victoria Webster-Wood	Carnegie Mellon University
Josh Bongard	University of Vermont

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Robotics.

"Read to succeed." — Samantha Nakeita Daniels

Abstract

Current robotics design and engineering trends typically focus on high-value applications where high performance, precision, and robustness precede cost, accessibility, and environmental impact. In this paradigm, the capability landscape of robotics is largely shaped by access to capital and the promise of economic return. This thesis explores an alternative paradigm in which the broader utility of robots is achieved via expanding the design space available for building robots using low-cost, accessible, and environmentally sustainable materials. Specifically, I focus special attention on *soft robotic* systems in which robot capabilities are achieved using rubbers, gels, and other compliant materials. This approach aims to lower the cost-of-entry to creating robots such that the technology is *accessible* to those outside highly resourced academic and corporate organizations. In particular, I explore how creating robots with soft and biologically-derived materials achieves these aims.

In the first part of the thesis, I apply well-established sensing and control methods to soft robotic platforms constructed from silicone rubber and other commonly used synthetic soft materials. We adapt these methods to account for the robots' soft material morphologies, sensors, and actuators and characterize the resulting challenges to performance and long-term use.

In the second part, I design and implement a modular robotics research platform that enables data collection across multiple soft material systems. In particular, I demonstrate the ability to collect longer-term data to address the challenges of using soft materials beyond short-term experimentation.

In the third part of the thesis, I examine the mechanics and functional capabilities of a class of *gelatin-based biomaterials* with potential applications in robotics. The purpose of this is to explore the reduction of costs and barriers associated with creating robots via biomaterial composites with the potential to function as various components of an embodied robotic system. Specifically, I formulate composites composed of readily available food-grade materials, agricultural waste, and overabundant by-products, which display physical properties that can be tuned via composition and processing.

Finally, I use this class of biomaterials to create a conductive gelatin-based composite with high electrical conductivity and robust mechanical properties. I fabricate this composite with ingredients and tools that are accessible to a broad population and demonstrate its application in soft and wearable electronics.

Acknowledgements

To Professor Carmel Majidi, my mentor, advisor and greatest advocate throughout my PhD, I cannot thank you enough. You supported me through this winding journey, pushing me forward when I needed encouragement and giving me the space to rest when I needed it. I have been able to learn about so many strange and wonderful things because of the freedom you gave me to explore, and for that I will forever be grateful.

I want to thank my committee, Professors Zeynep Temel, Victoria Webster-Wood and Josh Bongard. Your interest in my work has been a source of encouragement, and your wisdom, questions, and advice have shaped it into something I am proud of. I have been deeply inspired by each of your bodies of work, and I am grateful to have your support.

Thank you to Professor Chris Atkeson for encouraging me to pursue more interesting paths. Our conversations challenged and refined my understanding of not only my own ideas, but also the big picture of robotics. You taught me to see my work as art.

I want to thank all of the collaborators who made this work possible. Thank you to Abhinav Gupta and Vikash Kumar, with whom I interned at FAIR, and without whom I would not have developed the modular manipulation research platform. Thank you to my undergraduate mentees, Zimo Ge and Danielle Brennan, for your contributions and for giving me the opportunity to guide your research. To Eldy S. Lazaro-Vazquez for showing me the wonders of coagulated gelatin and the diversity of biomaterials. Thank you to my labmates for being the best research community I could ask for, through all the lab relocations and changes. Thank you to Yunsik Ohm and Michael Ford for answering my endless questions about conductive materials. Thank you to Michael Vinciguerra, Peter Roberts Olcay, and Dinesh K. Patel, for working alongside me to make conductive coagulated gelatin a reality. Michael, thank you also for endless cheer, delicious baking and for everything you've done to give me the space to write. Thank you to Zach Patterson, Drew Sabelhaus and Xiaonan Huang for trusting me with your precious robots. Zach and Drew, even when things kept breaking, our wonderfully stimulating conversations and camaraderie made working together with you a delight. Thank you to Anthony and Richard for long talks about research, life plans, and making it all work. Thank you Terri for your friendship and support in so many things inside lab and out, throughout all these years. To Tess Hellebrekers, your mentorship, support and friendship through those early years of my PhD were transformative for me. I was thrilled to be able to work with you again at FAIR. This thesis would not have been possible without you.

I want to thank everyone who helped me learn what I needed to on this path. Thank you to all of my friends at the Robotics Institute who made this mountain a bit easier to climb. I would like to especially thank to Ada, Kevin, Thomas, Cherie, Kate, Leo, Aditya, Shaurya, Tabitha and Victoria for office hangouts, giant watch parties, DND, long chats about cheese-bot and all the little moments that made me love my time here. To Himi, thank you for sharing all of your pearls of wisdom, and getting me out of my shell. Thank you to Ananda, for being my first robotics partner. Thank you to Scott Rippetoe and FRC Team 1477: Texas Torque, for showing me what a community can build together. To Kerry, Liz and Jacob, for being my second family during that time. Thank you to Team Guyana for giving me the opportunity to pass on my love for robotics.

Thank you to Dawn and the rest of the UT Dallas computer science crew for helping me grow into a half-decent programmer and infecting me with strong opinions about programming languages. Thank you to Becca, Z, Lexi, Olive, Jon, and Lena for DND,

physics mind-benders, cooking fiascos and road trips. Thank you to Bailey and Tiffany for endless conversations and support throughout the years. Thank you to Turga for your friendship and for sharing vision for an equitable and sustainable future, which changed the path of my research. To my therapist, for keeping me grounded when it all seemed impossible. To Kestrel, Nyx, Icarus, Rose, Casey and the rest of my queer little community in Pittsburgh, I cannot thank you enough. Through all of the challenges and setbacks, you have been there when I needed you. We have shared each other's joys, and we keep each other safe. I am so glad I found yinz.

Thank you to all of my family in Guyana, in Barbados, in Canada and in the US. To my grandma Veronica, who gives the best hugs. To my Papa Sheriff, I am grateful I get to be your granddaughter. To my late great-grandma Agatha, for whom I named myself. To my late grandma June, whose faith and strength taught me to never give up. To my late grandfather Wally, whose craftsmanship taught me to take pride in the work of my hands.

To my parents, David and Heather, and to my siblings Jessica, Kenny, Azalia, and Austen. I would not be the woman I am today without all of your faith, love and support. Mom and Dad, thank you for making sure I knew my strength. I love y'all.

To my late mom, Samantha Nakeita Daniels. I wish I could have shared with you the books that inspired so much of this work (and I thank Butler, Le Guin, Starhawk and Chambers for writing many of them). Mom, you taught me to chase my curiosity and I know you would be so proud to see where it has led me.

Finally, thank you to everyone not named here, without whom I would not have made it this far. To all my teachers, mentors, colleagues, friends, and family. My gratitude is beyond words.

The work presented in this thesis was made possible by the generosity of the National GEM Consortium through the GEM Associate Fellowship, The Robotics Institute at Carnegie Mellon University through the Uber Presidential Fellowship, and the National Science Foundation (NSF) through the National Robotics Initiative (Grant No. CMMI 1830362), and the Graduate Research Fellowship Program (Grant No. DGE 1745016).

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The public imagination of robots is astonishingly varied, with robots being depicted as versatile tools to improve human life, high-tech toys, scientific instruments, replacements for human labor, and even artificial life. The material impacts of real robotic systems on our shared society and ecology depend on their intrinsic characteristics and how we choose to use them [1]. Existing systems shape our vision and conversations around future technology and hold the potential to shape that technology, in turn [2–4]. I hope that the work presented in this thesis enables more people to participate in creating robotic systems and the prefiguration of a more equitable and sustainable future.

Consider the creation of an industrial robotic arm, a process that relies on the precision engineering of robust, rigid materials and the assembly of many parts into articulated kinematic chains that exhibit rigid-body dynamics. The design paradigm allows the use of well-developed modeling, sensing, and control techniques optimized for rigid-body robots. The result is a high-performance device that can operate precisely for extended periods. From the perspective of robots as a form of physical capital, these properties appeal to those with the resources to purchase or develop these systems. The broader impact of these technologies is complex. Their effectiveness at replacing human labor decreases wages and employment in the industries that embrace them [5].

Additionally, the most prolific application for industrial robotics is in automotive manufacturing [6], which currently has a fraught relationship with sustainability [7–9]. However, there is evidence that industrial robotics has the potential to contribute to "green innovation" [10]. The future direction industrial robotics will take has not yet been decided. However, due to the high barrier to entry for developing those systems, it won't be easy to ensure its development reflects the needs and values of the public [1].

As a counterpoint, the rise of additive manufacturing and freely available educational material online has led to engineering knowledge and manufacturing capabilities becoming more distributed among the public [11, 12]. Modern-day cottage industries have sprung up for many goods and services, making production less centralized and more responsive to people's diverse needs and values [13, 14]. Outside of the market model, maker spaces have arisen as community-scale centers of DIY problem-solving, and many libraries have used the increased accessibility of manufacturing technologies to provide free education and services to their communities [15]. Inexpensive robotic technology can be used in education from middle school to graduate levels, and the low cost allows that benefit to spread outside institutions with online courses [16]. By following the pattern of additive manufacturing, focusing on inexpensive robotic systems, robotics could become as accessible and replicate the current boom in open hardware, maker communities and education.

Soft robotics are a promising substrate for exploring this direction of robotic systems development. Soft robotic systems are primarily composed of mechanically compliant and deformable materials, and their governing mechanics are typically dominated by material deformation. Their compliance promises safe integration into domestic environments and contact with humans [17]. Some can be inexpensive to fabricate, constructed with few parts, and from various materials [18]. A particularly striking example of this is Arthrobots, which are made from drinking straws and balloons [19]. The educational applications of inexpensive robots are enhanced by soft robotics. Soft robots can be easier for a broader population to approach than rigid robots due to avoiding established stereotypes among students about the kind of person that can do robotics, increasing diversity in, for example gender [20]. It is also often fairly accessible for educators to bring soft robots into educational environments due to easier safety assurance, some intrinsically inexpensive classes of soft robot [21, 22].

The orientation of this thesis towards the question of sustainability is not in the application of robotics to environmental applications, which is a distinct and better-established articulation of this relationship [23]. Instead, it is the ecological impacts of robotic systems development itself that I aim to address. In the case of soft robotics, there are many potential ways to improve their sustainability [24]. The intimate connection of their function to the materials they are composed provides a compelling justification for starting with materials.

A note about definitions: The term *biomaterial* is established in the field of biomedical engineering to refer to materials designed to interface with biological systems, especially materials designed to be implanted into living bodies or to have living cells grown on them [25]. These materials could be either synthetic or biologically-derived. In this thesis, I instead use the term *biomaterial* in opposition to fossil-fuel-derived materials to refer to materials derived from biological sources built with biologically derived molecules and polymers. The orthogonal factors of material sourcing vs. material biodegradability are shown in Figure 1.2.

In Figure 1.1, we see various stages of life at which to evaluate the sustainability of materials. Biomaterials have sustainability benefits at the extraction stage and often at the disposal stage [26]. Gelatin-based biomaterials are a great example of this and serve as the basis for the materials work presented in this thesis.

1.1 Overview

Challenges in Soft Robot Autonomy

In Chapter 2 of this thesis, I identify aspects of the robotics stack that are impacted by using soft materials in robotic construction. Work on controlling inexpensive systems demonstrates the need for different control methods when we move away from the assumptions of the dominant trends in robotic system design. While rigid-body systems can rely on closed-form solutions to the robot's fundamental kinematics, in the soft robots field, this is often reliant on data-driven empirical methods [28]. I examine diverse soft

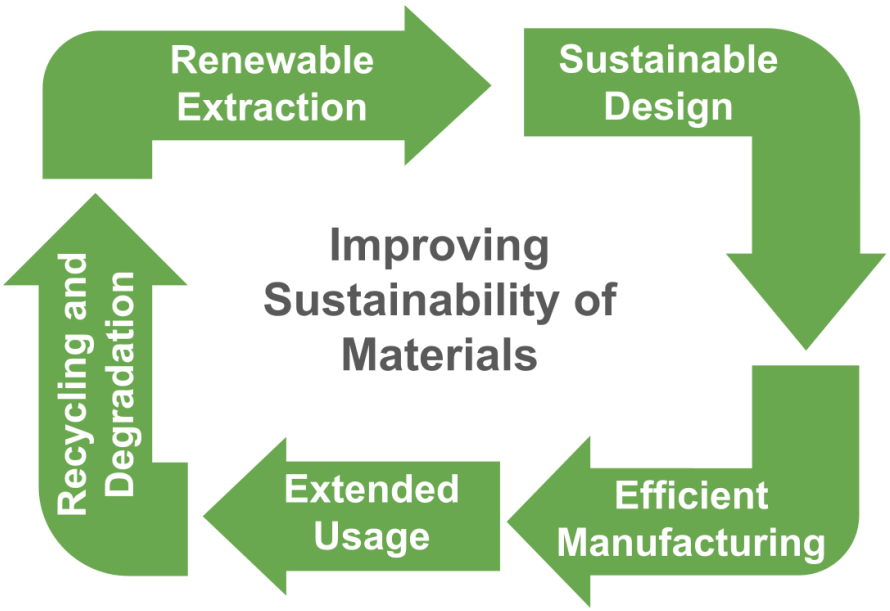

FIGURE 1.1: The sustainability of materials throughout their life-cycle, adapted from Olivetti *et al.* [27].

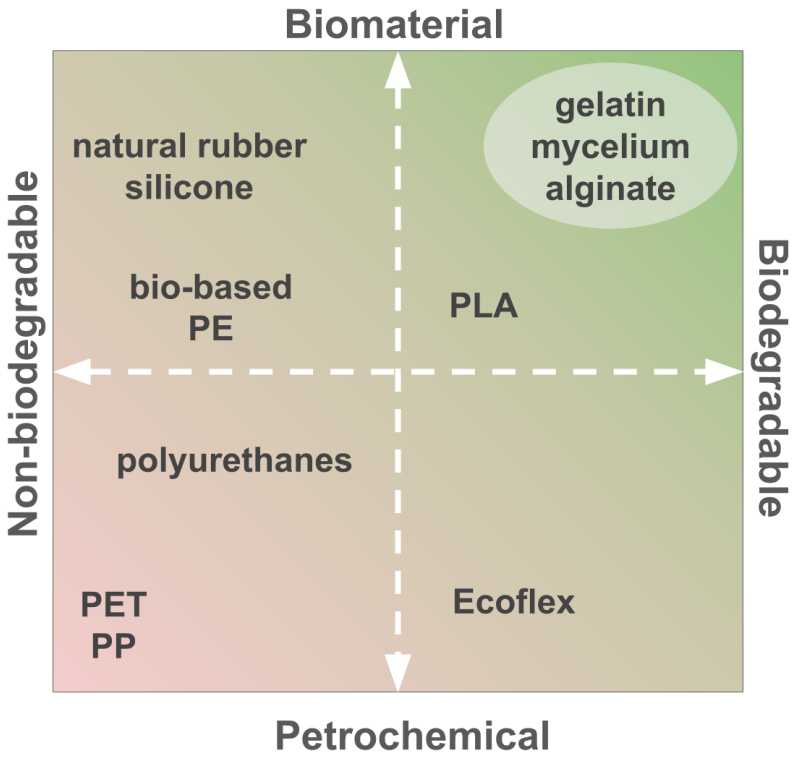

FIGURE 1.2: Orthogonal factors of biodegradability and materials sourcing for potential soft material polymers.

robotics platforms under various task contexts to establish specific challenges in soft robot autonomy.

> **Research Question:** How do the multi-physics dynamics of soft materials challenge effective soft robotic autonomy?
>
> **Research Question:** To what degree are traditional dynamics modeling methods capable of enabling soft robotic autonomy?
>
> **Research Question:** What are the bottlenecks for advancing the capabilities of soft robotic autonomy?

Enabling Soft Robotic Hardware Datasets

I aim to enable the refinement of accessible robotic technologies by developing a platform that allows for long-term experiments, compatible with the kinds of materials developed in Chapters 4 and 5, and suitable for evaluating solutions to the challenges identified in Chapter 2. Additionally, the costs of this platform itself are kept low, hoping it might allow for the kind of collaborative scientific project embodied by endeavors like the Robotic Research Cloud [29].

> **Research Question:** How effective will a research platform constructed from off-the-shelf electric motors and soft material mechanism be at long-term collection of soft robotic hardware data?
>
> **Research Question:** Do the differences in material properties between soft materials result in measurable changes to the dynamics and behavior of a soft parallel manipulator?

Gelatin-based Biomaterials

In Chapter 4, I exploit the coagulation of gelatin with either ethanol or isopropyl alcohol, previously reported in fiber-spinning literature [30, 31] and apply it for the first time to the fabrication of a bulk biomaterial. I use glucose syrup and glycerol as co-solvents and dehydrate the hydrogel to achieve good mechanical properties.

> **Research Question:** How can readily available ingredients derived from biological sources be used to create mechanically?
>
> **Research Question:** How can domestically available tools and techniques (e.g., ovens, food dehydrators) enable synthesis of novel gelatin-based biomaterials?
>
> **Research Question:** How can the fabrication processes and material properties of gelatin be modified to improve suitability of gelatin-based biomaterials for soft robotics?

Accessible Soft Electronics with Conductive Gelatin Composite

In Chapter 5, I add electrical conductivity to the gelatin-based biomaterial formulated in Chapter 4. The fabrication process used is designed to be achievable with consumer-grade equipment. I demonstrate extremely high electrical conductivity, good stretchability, and functionality as a soft wire and electromyography electrode. I further show that by controlled assembly of silver flakes.

> **Research Question:** Can a stretchable, conductive biomaterial suitable for use in soft electronics be produced with accessible fabrication methods?
>
> **Research Question:** Can dynamic material properties of gelatin hydrogel be leveraged to enable controlled assembly of conductive filler particles?
>
> **Research Question:** How do the mechanical properties of highly stretchable coagulated gelatin hydrogel respond to the addition of conductive filler?
>
> **Research Question:** How do mechanical and electrical properties of conductive gelatin hydrogel compare to the state of the art?

Contributions:

The overall flow of this thesis is shown in Figure 1.3. My core contributions are as follows:

- The application of well-established methods of enabling robot autonomy to soft robotic systems and the identification of trends and challenges in soft robotic autonomy deriving from intrinsic soft matter physics (Chapter 2, [28, 32, 33]).
- The design and implementation of a research platform capable of carrying out long-run data collection and machine learning experiments of soft materials in a robotic parallel manipulator (Chapter 3).
- The introduction of a highly stretchable hydrogel leveraging inexpensive biomaterials derived from waste and surplus material stock and an accessible process for its fabrication (Chapter 4, [34]).
- The introduction of a highly stretchable and electrically conductive composite from gelatin hydrogel and silver flakes, and an accessible process for its fabrication (Chapter 5, [34]).

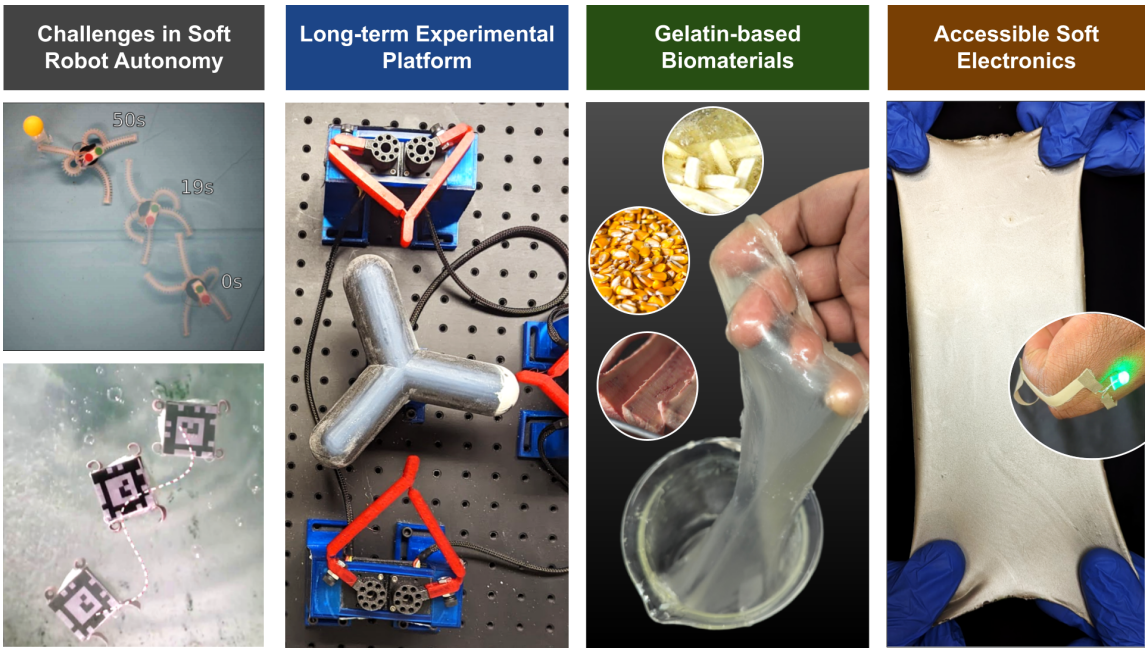

FIGURE 1.3: Overview of the sections of this thesis.

Chapter 2

Challenges in Soft Robot Autonomy

In this chapter, I study the autonomy of soft robot systems constructed from traditional non-degradable soft materials. These systems are of interest in their own right due to their intrinsic potential for operating safely alongside humans derived from their generally lighter construction and ability to conform to the surfaces they contact instead of rigidly colliding [35] [18]. If the development of robotic systems is to be made more accessible, reducing the hardware complexity needed to ensure safe interactions makes soft material systems an attractive option on account of their inherent material robustness and mechanical compliance. In addition to the impacts of their external interactions, the implications for the autonomy of these soft materials might provide a fertile experimental ground for understanding how we can operate soft robots without the massive modeling, sensing, and control simplifications offered by high-precision rigid-body engineering.

When soft materials are integrated into robotic systems, their impact varies depending on which components of the robot architecture are made soft. Soft material sensors can be integrated into existing rigid-body systems with minimal extra operational complexity while providing novel state-sensing capabilities (Figure 3.1A) [28, 36]. On the other hand, when a robot's actuators and moving body are primarily constructed from soft materials, there is no longer an ability to rely on rigid body kinematics. In these systems, where their dynamics are dominated by soft material deformation, state representation, sensing, and control are open problems. 3.1

The work of Deisenroth *et al.* on data-efficient learning on inexpensive robotic hardware shows us some of the challenges an inexpensive rigid-body robotic arm presents for control [37]. For most robotic arms, high-precision joint-space state estimation is ensured via sensors embedded in each joint, usually encoders on the joint motors. But for a minimally-expensive arm like this one, onboard sensors of this kind are not guaranteed. Additionally, joint slop and motor backlash contribute to a partial decoupling between actuator outputs and robot configuration, i.e., it doesn't go quite where it's told, and external forces can move it from the configuration prescribed by the actuators. We can see then the deeper connections between soft and inexpensive systems, where the inexpensive system recapitulates the partially observable and partially controllable features of soft systems, the difficulty of sensor integration into soft systems, and the intrinsic decoupling between actuator state and robot configuration mediated by soft material deformation. As we intend to enable the development of soft and inexpensive systems, further exploration of these features of soft materials is a promising line of research.

2.1 Closed-loop Underwater Crawling of Brittle star

*Z.J. Patterson, A.P. Sabelhaus, **K. Chin**, T. Hellebrekers, and C. Majidi. "An untethered brittle star-inspired soft robot for closed-loop underwater locomotion." In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 8758-8764. IEEE, 2020.* [32]

Author contribution: I formulated the state representation, implemented the vision-based state feedback, and designed and implemented the algorithm for planning over primitives. Z.J. Patterson led the design and fabrication of the soft robot testbed.

2.1.1 Brittle Star Robot

The experimental platform for this work is PATRICK, a brittle star (*ophiuroid)* inspired soft robot 2.1A. PATRICK is the first untethered underwater crawling robot reported in the literature. Like the brittle star, PATRICK moves by multidirectional crawling on submerged surfaces. This robot has a central electronics core, with five radially symmetric legs that can bend about two axes (no torsion) via two pairs of antagonistic actuators embedded in each leg. The legs are composed of silicone (Smooth-On DragonSkin 10 NV silicone elastomer), and their shape is achieved via 3D-printed molds. They are powered by shape memory alloy (0.008in Dynalloy 90ºC Flexinol® with a 0.054in helical diameter) (Figure 2.2).

We evaluate the locomotion performance of PATRICK under tasks requiring closed-loop control. Since PATRICK is not equipped with local proprioceptive sensors, computer vision from an overhead camera was used to track colored markers attached to PATRICK's dorsal surface. The location of a floating goal object (yellow ball) is also tracked this way. (Figure 2.1B) and communicated to the robot via wireless communication (Figure 2.3). The task is then to follow the goal object.

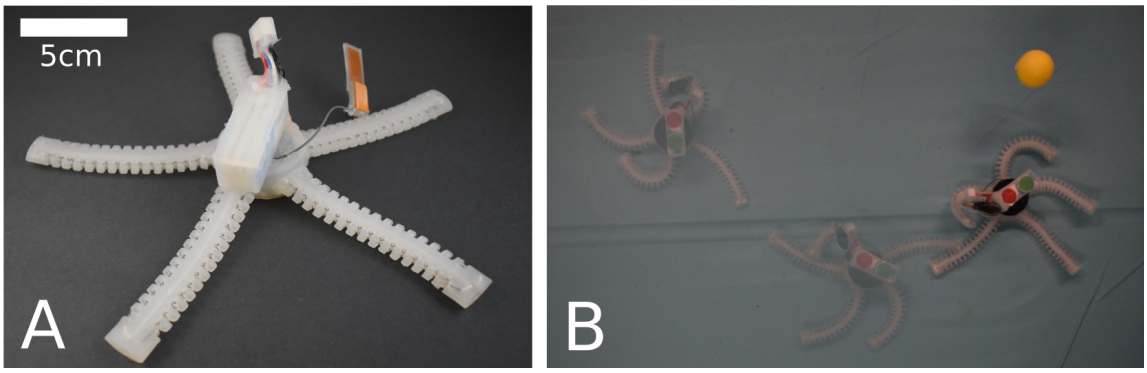

FIGURE 2.1: A) The brittle star-inspired robot, PATRICK. B) Autonomous underwater crawling of PATRICK towards the goal object.

2.1.2 Reduced State and Action Space Representations

Due to the deformable nature of the structure of PATRICK's body and actuators, its motion has many inherent nonlinearities, which is further compounded by the multi-physics contact and drag dynamics of locomotion on an underwater surface. The challenge becomes creating a controllable system representation despite the unmodeled true dynamics. Without a high-fidelity a priori model of the effects of control effort, we need to be

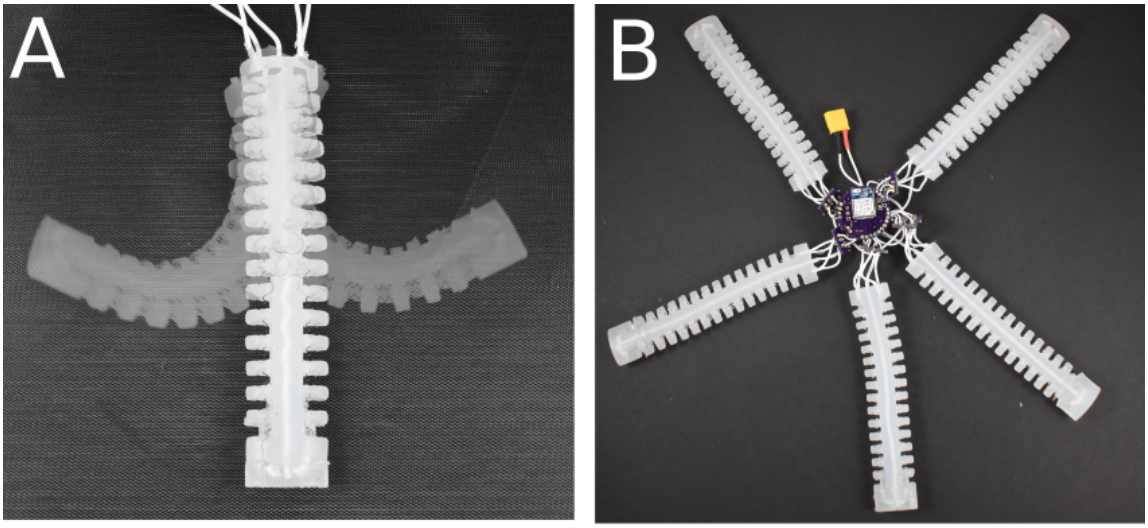

FIGURE 2.2: A) One of the two bending axes is achievable with the activation of shape memory alloy spring actuators. B) PATRICK without silicone core and battery to show electronics.

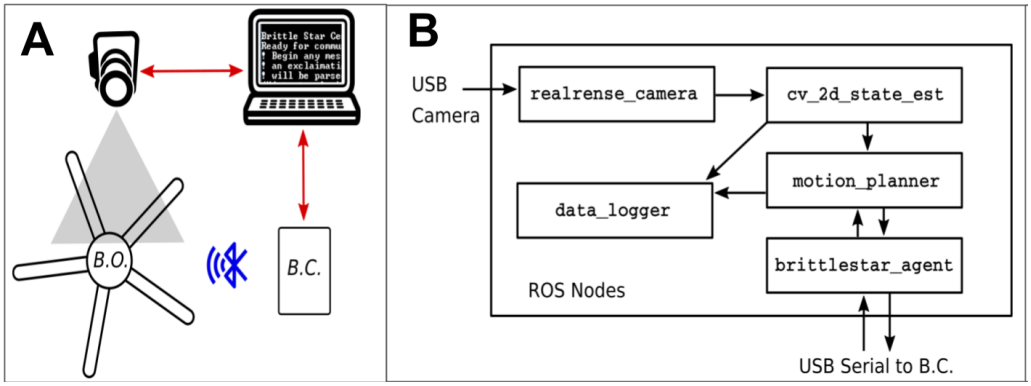

FIGURE 2.3: (A) System architecture for PATRICK and its testbed. A camera pointed at the robot tracks various markers via OpenCV running on a connected computer. Embedded software on the robot communicates over Bluetooth Low Energy to similar software running on a separate board for communication over USB (red) to various ROS nodes on the computer. (B) Architecture for the various nodes used in the Robotic Operating System (ROS) package for the PATRICK robot. The Intel RealSense camera package supplies frames, then four separate nodes perform state estimation from markers on the robot, motion planning, data logging, and communication with the robot itself. Commands are sent to and from the hardware robot via the microcontroller attached over a USB serial port.

able to perform system identification on the hardware in the space of the chosen representation. Given the engineering complexity of the research-grade system and the likelihood of extended operation leading to repairs, we seek as few parameters as possible to limit the amount of hardware data required.

State Representation

The intrinsic state of PATRICK is much more complex than a comparable rigid system. If we consider the robot as a body with five limbs attached, a rigid PATRICK would have five limbs, each with two axes of rotation, resulting in 10 actuated degrees of freedom corresponding to biaxial leg rotation and six unactuated degrees of freedom corresponding to the global pose of the body in 3D euclidean space (SE(3)). PATRICK has the same ten actuated degrees of freedom corresponding to biaxial leg bending and six Euclidean degrees of freedom. However, it also has uncountable additional unactuated degrees of freedom due to soft material deformation. Every silicone polymer chain has many degrees of internal freedom and interacts with other polymer chains, the locomotion surface, and the surrounding water. Thus, we resort to higher-level abstraction in order to write down any representation. We could consider a planar model of the state, incorporating the global 2-dimensional pose of the robot $\{x, y, \theta\}$, and the state of the limbs relative to that frame. For the limbs, we can model curvature without torsion, stretching, or changes to cross-sectional area, resulting in the curvature of the ith limb being expressed as a continuous function along the limb $\mathbf{k}_i(l)$. In this case, the overall state of the robot would be

$$\mathbf{x}_{\text{intrinsic},t} = \{x_t, y_t, \theta_t, \mathbf{k}_{i,t}(l)\}, \quad i \in \{0 \dots 4\}, \quad 0 \leq l \leq \text{limb length L}]$$

Unfortunately, the curvature function would require some additional parameterization to be practically useful. For contact-rich systems, that is somewhat of an open problem (though we will pursue this later in this chapter). Special cases, like the piece-wise constant curvature assumption, have been used to simplify the modeling of many soft limbs [38, 39]. In that case, the state of each limb would be defined by two curvature parameters, which would put our dimensionality at 13, 3 (global pose) + 10 (limb curvatures). However, this does not capture the fact that contact with the ground can change the curvature of the limb, which immediately makes it infeasible as a state representation for any valuable model for motion planning of the system. Finite element methods can capture these contact constraints to a fairly high-fidelity degree but are computationally exceedingly slow, especially if we want to understand the motion of the limb throughout its gait and the resulting effects. Rather than pursue this avenue for PATRICK, we choose to drop explicit modeling of the limb configuration and focus only on the quasistatic pose of the central body, leaving our state as

$$\mathbf{x}_t = \{x_t, y_t, \theta_t\}$$

This state representation then informs the necessary action representation, enabling us to build more compact action representations on these state assumptions.

Action Representation

The shape memory alloy (SMA) used in PATRICK's limbs is a metallic alloy that undergoes a phase change between crystalline forms under deformation and/or heating. This behavior can be exploited to create a temperature-modulated shape change on

the macro-scale. The properties of this thermal phase-change system make it difficult to model as an actuator [40]. Critically, the phase percentage depends on whether the previous state was hotter or cooler than the target state [41]. This temperature-phase hysteresis directly results in temperature-motion hysteresis, as the motion of SMA is directly related to the ratio of material phases in the material. Because of this complexity, there is no precise model of the relationship between applied voltage and the transient deformation of a limb driven by shape-memory alloy. What models have been made are particularly system specific and are not generalized across SMA actuated mechanisms [42].

Despite this analytical complexity, SMA actuation was chosen for PATRICK due to the relative ease of mechanical and electrical integration into a compact untethered form 2.2. When the shape-memory alloy coils embedded in the limbs are Joule heated, they contract, causing bending in that coil's direction. Due to the difficulty of scalar voltage control in this form factor, Joule heating is achieved via temporal control of the time spent in binary on/off voltage states. Therefore, the intrinsic actuation space at a specific moment in time is

$$\mathbf{u}_{t,\text{intrinsic}} \in \{0,1\}^{20},$$

as there are four coils in each of the five limbs, with two voltage levels per coil. As the motion of the SMA coil depends on the thermally-induced phase change caused by the accumulation of heat over time, many candidate actions at this time scale will not produce macro-scale motion. Therefore, we consider the trajectories of these instantaneous inputs,

$$\mathbf{a}_t = \{\mathbf{u}_t, \mathbf{u}_{t+1}, \ldots, \mathbf{u}_{t+T}\}.$$

The dimensionality of this resulting space is $20 \times$ T, where T is the number of time steps in a trajectory. This exact number depends on the trajectories' clock-time length and the voltage control's temporal resolution. However, T could easily be several dozen, making this size of this space dozens to hundreds of dimensions. We thus reduce this dimensionality further by reparameterizing in terms of heating onset and stop time for each SMA coil, producing a 40-dimensional action space. This space is a subset of all possible actions but can represent coordinated motion between the SMA actuators of a limb and across limbs. Even this multiply reduced, relatively compact representation is quite large in absolute terms, especially considering the cost and speed of computing. Let us look at our task's context to find solutions to this. Due to our state representation, actions operate on the planar pose of PATRICK and start and end with zero velocity to enforce the quasistatic assumption. Therefore, the transition model from one state $\mathbf{x}_t$ to the next state$\mathbf{x}_{t+1}$ is:

$$\mathbf{x}_{t+1} = \mathbf{x}_t + F(\mathbf{x}_t, \mathbf{a}_t), \qquad F(\mathbf{x}, \mathbf{a}^i) = [\Delta x^i, \Delta y^i, 0]$$

where $\Delta x^i, \Delta y^i$ are the planar displacement caused by the action $\mathbf{a}_t$. Informed by the "leading-limb" model of real-world brittle star gait, where the headless animal moves in the direction of a chosen limb without turning [43], actions are assumed not to change the robot's orientation. The unknown function $F(\mathbf{x}_t, \mathbf{a}_t)$ captures the result of all interactions between the voltage input sequence, the robot limb dynamics, surface contact, and any hydrodynamic and inertial forces.

To enable autonomous motion from this point, finding a policy that takes in the state of the robot and the goal position and produces actions that move the robot towards the goal is necessary. The transition function $F(\mathbf{x}_t, \mathbf{a}_t)$ is unknown, so an analytical model

is not available to us to calculate the results of actions without hardware testing. Empirical hardware evaluation of the general form of F is not feasible, given that our most compact full action space is 40-dimensional, meaning hardware testing would require a lot of data to sample the space. Therefore, we empirically evaluate F for a small number of actions. This set of actions is found via trial-and-error hand-tuning, guided by expert knowledge of the characteristics of shape memory alloy and spatial intuition about the effect of various limb motions. This library of known actions in an otherwise unknown transition space are called primitives, and this method has been used in rigid robots with complex or high-dimensional dynamics to enable tractable computation of graph based search methods of policy generation [44]. The resulting actions visually resemble the motion of real-world brittle stars, something like a breast-stroke on the bottom of the pool [45].

The displacements these primitives generate are found by reparameterization into polar coordinates,

$$[\Delta x, \Delta y] = [r^i \cos(\theta_t + \phi^i),\ r^i \sin(\theta_t + \phi^i)],$$

Each primitive moves the robot towards five radially symmetric limbs. We set the angle of translation ϕ of the first primitive a^0 to 90°, corresponding to translation along the y-axis, resulting in parameters

$$\begin{bmatrix} r^i \\ \phi^i \end{bmatrix} = \begin{bmatrix} r^0 & r^1 & r^2 & r^3 & r^4 \\ 90^\circ & 18^\circ & 306^\circ & 234^\circ & 162^\circ \end{bmatrix}$$

The true value of $r^0 \dots r^4$ will require characterization on the hardware, so we use an arbitrary value of $r = 5$ cm for all transitions in the primitive library as an initial transition model 2.4A. The next step is composing these primitives into a policy based on this transition model to solve the locomotion planning task.

2.1.3 Closed Loop Underwater Crawling

With our primitives chosen and a rough predictive model of their resulting transitions, we can build a policy by the method shown in Algorithm 1. There is a list of five possible actions $\mathbf{a}_t \in \mathcal{A} = \{\mathbf{a}^0, \dots \mathbf{a}^4\}$ corresponding to the primitive library we created. At each time step, based on visual estimation of the robot state and goal location from the overhead camera, an action $\mathbf{a}^i$ is chosen from the library. The choice of this action is found via the prediction of that action's resulting displacement $\Delta x^i, \Delta y^i$, according to transition model $F(\mathbf{x}, \mathbf{a})$, with a greedy heuristic where states closer to the goal are better than states farther away,

$$c(\mathbf{x}_t) = ||\mathbf{x}_t - \bar{\mathbf{x}}||_2.$$

The planning horizon chosen was a single time-step, as rigorous enough system identification to avoid error accumulation was not performed.

We implement this scheme on the physical robot, resulting in reliable motion towards the goal (Figure 2.4C). From data collected during task completion, we characterize the true displacement magnitudes of the primitives (Figure 2.4B), which average to approximately 2.31 cm, less than half of the 5 cm used in planning. As shown in Figure 2.4D, the robot can consistently decrease its distance to the goal even with this inaccurate model. Notably, since PATRICK's limbs are 10 cm long, the limbs are touching the goal while the center of mass is 10 cm away.

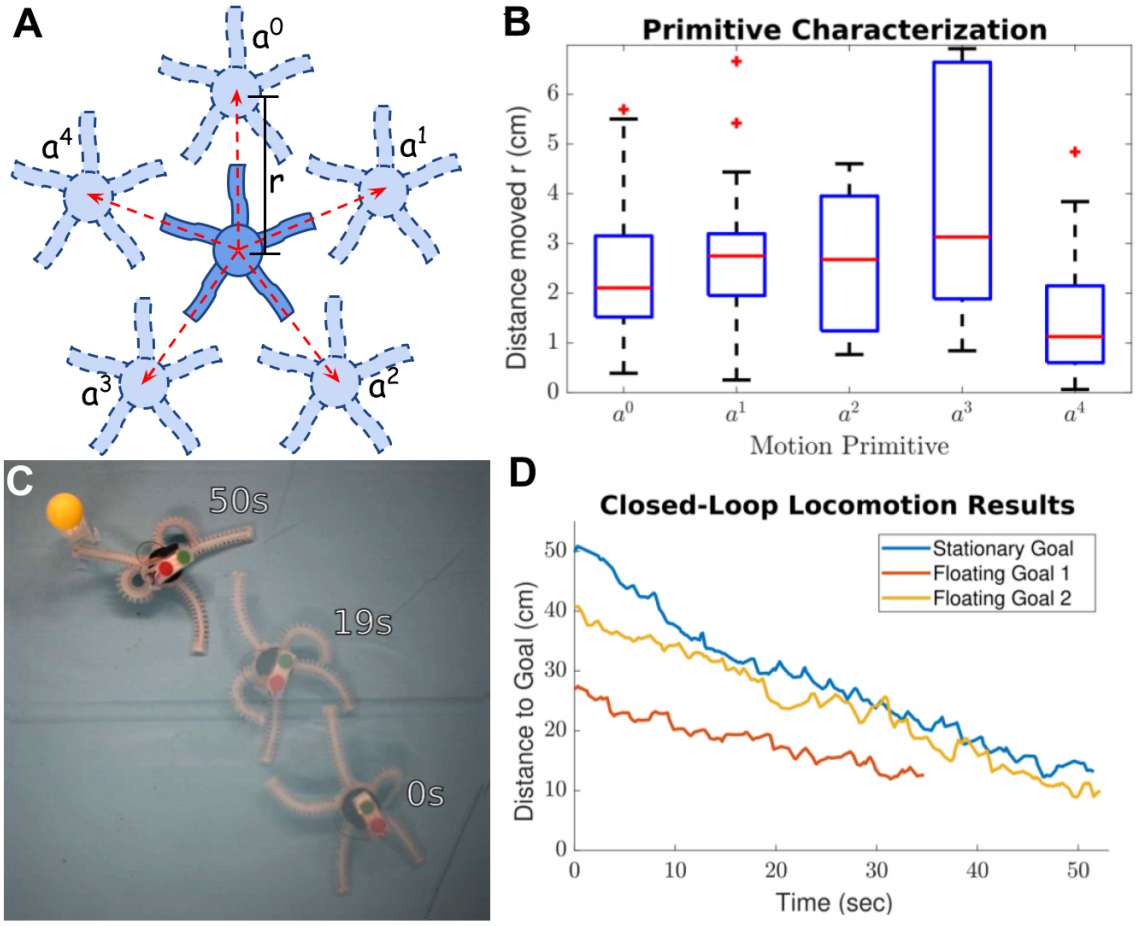

FIGURE 2.4: A) Schematic representation of the nominal state changes induced by primitives, based on assuming some constant displacement and no change to orientation. B) The empirical characterization of the noisy results of the primitives on the real robot over many cycles. C) Visual path of PATRICK moving to a stationary goal via greedy motion planning. D) Numerical distance over time results of goal-seeking performance, measured from center-of-mass of the robot to goal.

Algorithm 1: A greedy model-based policy

Input : Robot state $\mathbf{x}_t$, goal position $\bar{\mathbf{x}}$, set of primitives $\mathbf{a} \in \mathcal{A} = \{\mathbf{a}^0, \mathbf{a}^1, \mathbf{a}^2, \mathbf{a}^3, \mathbf{a}^4\}$, transition function $\mathbf{x}_{t+1} = \mathbf{x}_t + F(\mathbf{x}_t, \mathbf{a}_t)$, tolerance d for distance to goal

Output: Closed loop trajectory of action primitives $\mathbf{a}_t$ to get to the goal

while $||\mathbf{x}_t - \bar{\mathbf{x}}||_2 > d$ **do**
  for $\mathbf{a}^i \in \mathcal{A}$ **do**
    $\hat{\mathbf{x}}_{t+1} = \mathbf{x}_t + F(\mathbf{x}_t, \mathbf{a}^i)$
    $c(\hat{\mathbf{x}}_{t+1}) = ||\hat{\mathbf{x}}_{t+1} - \bar{\mathbf{x}}||_2$
  end
  $\mathbf{a}^t = \arg\min_{\mathbf{a}^i} c(\hat{\mathbf{x}}_{t+1})$
  $\mathbf{x}_{t+1} \leftarrow$ Robot.execute$(\mathbf{a}^t)$
end

2.1.4 Discussion

PATRICK was able to localize to the goal reliably, sequencing action primitives to reduce its distance to the goal. This success shows that low-fidelity transition models can still be useful for closed-loop locomotion in an open environment. The displacement magnitude of each primitive was less than half of the magnitude used in planning, and the primitives induced some rotation, but the qualitative directional accuracy enabled task completion. We could consider finding the correct pose transitions as a system identification task, but this approach has clear challenges. As shown in Figure 2.4B, there is significant variation between primitives and across different execution instances of the same primitive. Many possible sources for this variation emanate from the complex multi-physics interactions of actuators, soft material deformation, and contact with the environment. Treating these variations as noise and following the planning path under uncertainty is possible. However, to address this problem more effectively, let us consider some of the relevant dynamics that we know about. After an action execution, the robot's configuration is not reset back to the same state. The exact limb curvature profile and position of the limbs relative to its center, the thermal and phase state of its shape memory actuators, any internal material stresses, and plastic deformations all change between actions. These changes to the system are not captured in the reduced state used for PATRICK, meaning that even a well-identified system parameterization that characterized expected action outcomes and uncertainty would cease to be accurate over time.

The models and representations that are tractable with soft robotic hardware of this form are constrained by the engineering and logistical costs of collecting hardware data. I locate some of the challenges of soft robot autonomy in this gap between the tractable and the true dynamics of those systems.

2.2 Simulation-aided Motion Planning of Frog-inspired Swimmer

Huang, Xiaonan, Zach J. Patterson, Andrew P. Sabelhaus, Weicheng Huang, ***Kiyn Chin****, Zhijian Ren, Mohammad Khalid Jawed, and Carmel Majidi. "Design and Closed-Loop Motion Planning of an Untethered Swimming Soft Robot Using 2D Discrete Elastic Rods Simulations." Advanced Intelligent Systems 4, no. 10 (2022)* [33]

Author contributions: I programmed the motion planning framework for trajectory following. A.P. Sabelhaus and I formulated the primitive lattice and velocity-space randomization approach for guiding simulation targets for use in online motion planning framework. X. Huang and Z.J. Patterson designed and fabricated the soft robot testbed.

2.2.1 Swimming Frog Robot

The experimental platform for this work is an untethered frog-inspired robot designed and built by my collaborators, Xiaonan Huang and Zach J. Patterson. Compared to PATRICK, this robot can have higher actuation frequencies and swims instead of crawling. The leg actuators are composed in layers, with a loop of shape memory alloy embedded in a liquid metal elastomer composite. This composite has microscale droplets of liquid eutectic gallium-indium (EGaIn) alloy embedded in it, allowing for faster heat transfer, meaning faster cooling of SMA coils and higher actuation frequencies [46]. There is a pre-stretched layer of thermally-conductive silicone embedded in the actuator, introducing a spring force for recovery from the extended leg position, allowing the SMA to return to its curled state faster than by passive cooling, further increasing actuation frequency [47]. Four of these actuators are attached to a buoyant central body, enabling forward motion and controlled rotation via antagonistic actuation. The resulting robot is pictured in Figure 2.5. The published work this section is based on covers many details of the robot simulation and modeling. However, I will discuss these at a high level and focus on the sections most relevant to our goal of understanding the intrinsic challenges in soft robot autonomy. I urge interested readers to seek out that paper.

In the following work, I examine the effectiveness of simulations of soft robot behavior in enabling autonomous trajectory following this soft robotic swimmer.

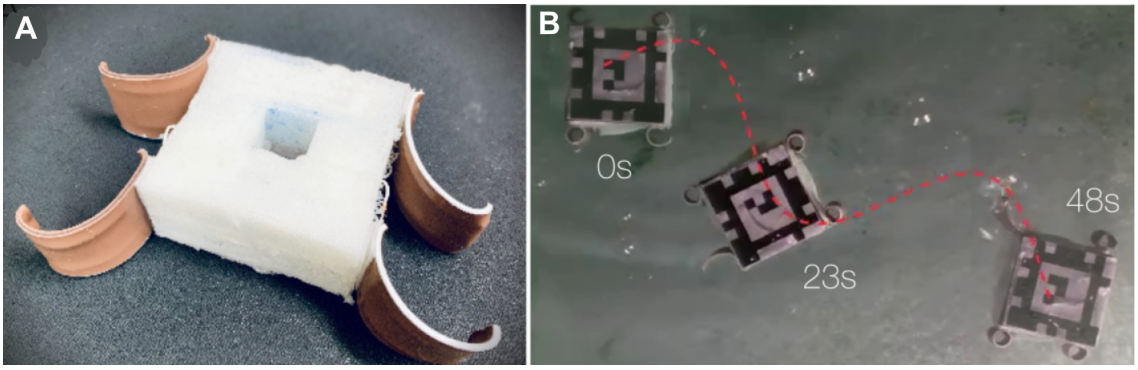

FIGURE 2.5: A) Frog robot with shape memory alloy actuators as designed by my collaborators Huang *et al.* [47]. B) The frog robot swimming along a winding path.

2.2.2 Simulation with Discrete Elastic Rod Model

The robot is simulated via the discrete elastic rod (DER) model of soft material deformation, suitable for structures composed of slender elastic elements [48]. The discretization of the frog robot is shown in Figure 2.6A. In brief, the robot is represented by a series of elastic rods, modeled as masses connected by linear and torsional springs. The energies of those springs are associated with the stretching and bending of each discrete rod, and they produce forces on the masses. External hydrodynamic forces can also be captured as forces on these nodes. These forces are then used to create equations of motion, which are implicitly solved at every timestep via the backward Euler method. There are parameters associated with the stiffnesses of the simulated springs, the mass of the nodes, and drag coefficients for various elements of the limb and robot motion.

Simulation parameters are calibrated on hardware by recording the robot's motion after accelerating to a steady-state speed. The results of this calibration are shown in Figure 2.7. The agreement between simulation and hardware is tested across various actuation frequencies and phase offsets between the front and hind limbs.

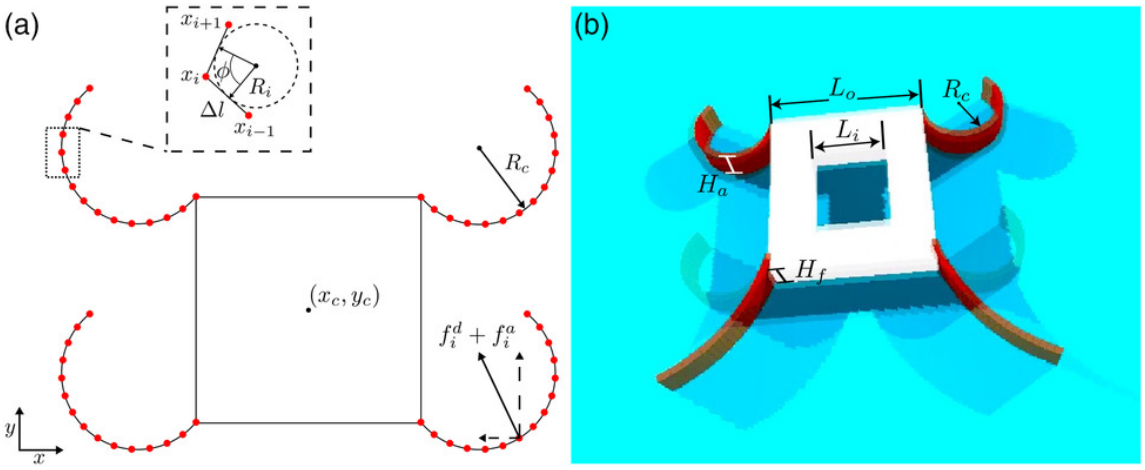

FIGURE 2.6: A) The model used in the simulation of frog robot with discrete elastic rod model. R_c is the radius of curvature of the limb, (x_c, y_c) is the robot's center of mass position, and f_j^d and f_j^a are the drag and virtual mass forces on the jth node respectively. In the inset are properties of the discretized rod geometry, notably the turning angle, ϕ, and the rod length, Δl, which are later used to calculate bending energy. x_{i-1}, x_i, and x_{i+1} are the labels of successive nodes B) Simulated render of frog robot geometry, $L_o = 90mm$, $L_i = 35mm$, $H_f = 35mm$, $H_a = 24mm$.

2.2.3 Motion Planning with Action Primitives

The discrete elastic rod simulation runs faster than real-time, but not fast enough to be used in an online motion planning framework. Instead, we leverage the simulator to collect data, which we use to form a forward model of the robot. We will then use this forward model to enable search-based planning of robot actions based on which actions have the lowest cost as determined by saved state transitions and the task definition.

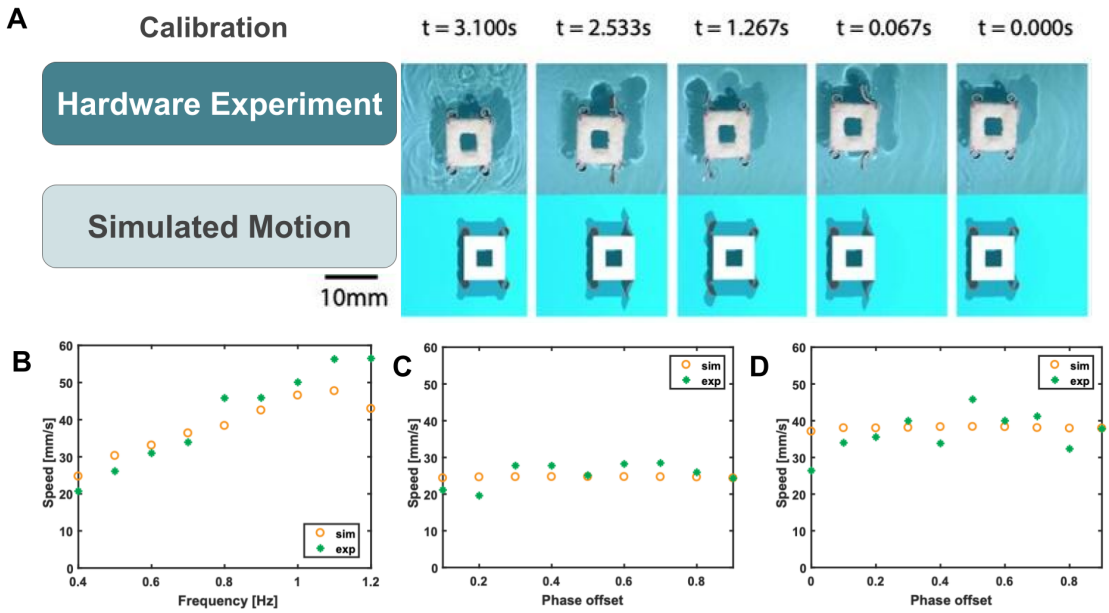

FIGURE 2.7: A) Visual comparison of forward locomotion of frog robot hardware vs. calibrated simulation B) Numerical comparison of agreement between simulation and hardware for swimming speed vs. frequency. C) Numerical comparison of agreement between simulation and hardware for swimming speed vs. phase offset between forelimb and hind limb actuation at 0.4 Hz and D) 0.8 Hz

Similarly to the previous work with PATRICK, the full action space induced by the actuators is very large initially. While simulation experiments are less costly than hardware experiments as there is no risk of needing to perform repairs or replace parts, there is still a need to reduce this dimensionality. Firstly, in a search-based planning framework, the number of actions considered at each time step sets the branching factor for the search tree. Therefore, the action space must be discretized to begin with. In rigid robots with electric motor actuation, a discretization that spans the full configuration space of the robot could be as simple as increasing or decreasing the joint angle by some small constant angular resolution, making the action space dimensionality equal to twice the number of joints (e.g., 12 for an industrial robot arm). However, for shape memory alloy, no such simple generator exists that allows the spanning of the whole space due to the highly temporally coupled nature of the actuator dynamics. Therefore, we again turn to action primitives to lower the dimensionality of our action space.

Action primitives are shape-memory alloy actuation sequences that are hand-chosen to correspond to behaviors like going straight ahead, rotating on the spot, curving left, and curving right, which could be composed to enable path following. Nine of these behaviors are chosen to guide simulation. The basic idea is to start with a stationary simulated robot as the root node in a tree structure. Each primitive is executed starting from that initial state, producing nine child nodes. This context is not quasistatic; the execution of the primitives leaves the robot with non-zero velocity. To achieve better coverage of velocity space, we sample over initial linear and angular velocity distributions, producing over 8.5k initial velocity states from which the nine primitives are executed, resulting in approximately 76k state transitions. These state transitions are stored in tabular form, and this transition table is used as the forward model.

2.2.4 Closed-Loop Trajectory Following

We leverage this simulation-derived forward model in a trajectory-following task with localization. The form of the policy is a receding horizon search-based planner. I define trajectories in the world frame as parametric curves of the form $[x(p), y(p)]$. Waypoints are generated from this curve via dense sampling of parameter p. To enable tracking, the robot's state is captured via overhead cameras localizing to an AprilTag [49, 50]. The trajectory is transformed into the body frame of the robot using this detected pose to allow comparison between points along the trajectory and the stored transitions from the tabular forward model. To choose actions, the nearest neighbor in the velocity space to the robot's current state is found within the transition table. The results of executing all primitives from that state are already stored in the table, so the successor states for the robot's current state are found by indexing. These successor states are transformed back into the world frame to get the potential next states of the robot should it execute any primitive. This process can be repeated recursively from the successor state up to a defined search depth. To select the next action, the action with the lowest cost (for the given planning horizon) is chosen.

The cost function has three components, designed to prioritize staying close to the path (distance to the nearest waypoint), remaining aligned with the path (angular distance from the tangent of the path at the nearest waypoint), and making forward progress along the path (the size of the associated parameter p of the closest waypoint), cost $=$

$w_a * \text{dist} + w_b * \text{ang} + w_c * \text{prog}$. The w's are weights allowing tuning of the relative importance of these features.

This motion planning scheme was tested on multiple linear, sinusoidal, and ellipsoid paths. These results are shown in Figure 2.8. We can see that the robot can follow paths of different levels of complexity fairly effectively, though the tracking precision is not especially high. In practice, the best-performing planner has a search depth of one-time step.

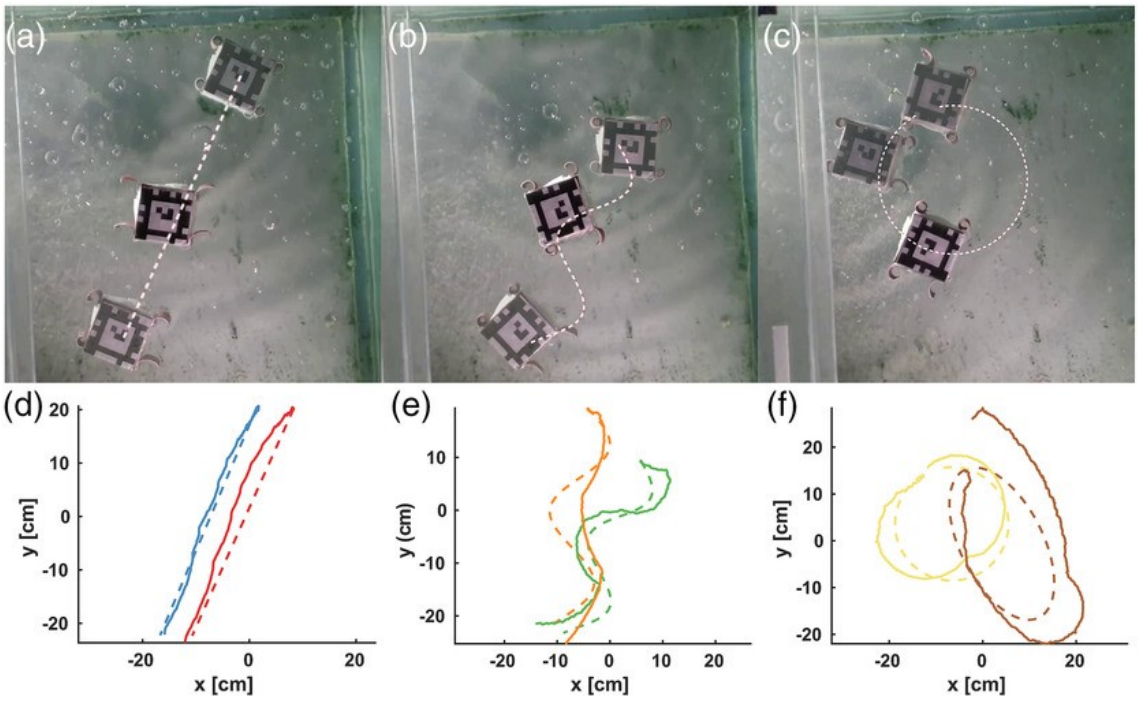

FIGURE 2.8: Camera frames of the robot along A) linear path, B) sinusoidal path, and C) ellipsoid path. Comparison of actual robot motion (solid lines) vs. prescribed trajectories (dotted lines) for two trials of D) linear path, E) sinusoidal path, and F) ellipsoid path

2.2.5 Discussion

The forward model established in discrete elastic rod simulation was sufficient for a successful closed-loop trajectory following across multiple paths. The calibration of the simulation allowed moderate extrapolation to the effects of primitives despite the limited scope of hardware training data. The simulation was calibrated around steady-state forward motion. However, there was no hardware data for the effects of the primitives, states where the system was accelerating, or different velocity values. Within the simulation itself, the data collected was guided by velocity-space sampling to generalize better, but this generalization appears weak. A planning search depth of one outperformed longer-horizon methods, indicating that the tabular forward model was not sufficiently accurate to withstand the compounding error of sequential prediction. It is possible that this could be addressed by a denser sampling of the simulator, perhaps by rolling out the primitive execution recursively in simulation to a deeper search depth as is done by Likhachev *et al.*

[44]. However, intuitively, this would likely lead to the same challenges of moving away from well-validated areas of the state space on the simulation side. This fragility to moving away from validated data is further supported by the fact that interpolation between simulated states for queries performed worse than a nearest-neighbor approach. Another axis of unexplored data is that the DER simulation was constrained to its 2D form, and 3D data might improve this. However, this increases computational costs and the amount of data required to account for the extra parameters.

There are accessibility and resource efficiency benefits to an approach like this, as the speed of tabular querying is better than online solving of the underlying equations of motion. This planning approach of precomputing a library of trajectories based around action primitives has been used in many areas of robotics [51], and has been optimized for more complex tasks like high speed collision avoidance [52]. For rigid systems, this process is more effective, as the simulators available are generally very fast and fairly accurate, especially once system identification has been performed. Discrete elastic rod simulations are significantly faster for compatible systems than the dense, full-order finite element methods for soft material deformation. However, some success has been achieved recently with speeding up FEM simulations for soft robots by using coarse meshes and hardware calibration [53] or reduced order decompositions with dense meshes [54]. DER has also been extended to handle some contact-rich scenarios effectively [55]. However, this has not been tested for hardware, and the method still requires the system morphology to be feasibly decomposed to a network of slender elastic rods, which is compatible with many existing soft robot designs but does constrain the design space of future soft robotic systems. More importantly, performing system identification is significantly more challenging due to the larger number of parameters in elastic rods than in rigid links.

The discrepancy between the model extracted from a locally calibrated simulation and the actual hardware is known as the Sim2Real gap, and it appears to be especially pronounced for soft systems. One method to address this domain randomization, where the parameters of the simulation are sampled from a distribution rather than only using the parameters resulting from limited calibration [56] and has been used to solve systems as complicated as solving a Rubik's cube one-handed with a robotic hand [57]. However, this method requires two things that are difficult to find for soft systems in a resource-constrained context. First, the computation required increases massively as a large parameter sweep is needed to cover enough parameter space to generalize. Secondly, it is still unclear what parameters are necessary to model robots with soft materials. Discrete elastic rod models or other reduced order soft material models potentially provide some parameters, but additional parameters are associated with the physical context, e.g., hydrodynamic parameters in this swimming task and thermal and phase-change actuator parameters when using shape-memory alloy coils. The relevant latent space parameters common to all soft systems, especially low-cost or biomaterial soft systems, are still unknown. For DER, and indeed for other existing soft robot simulation strategies like coarse FEM, the difficulty of identifying relevant parameters is an ongoing challenge.

2.3 An Attempt at Open Loop Walking

2.3.1 Shape-memory Alloy Walker

In this section, I work with Horton, a soft walking robot powered by SMA actuators (Figure 2.9). My collaborator, Andrew P. Sabelhaus, oversaw Horton's design. Horton is intended as a platform for studying walking control using soft robots.

Horton is a planar robot constructed by soft segments molded from Smooth-on Smooth-Sil, a fairly stiff silicone material with a Shore hardness of 50 A. Each segment has two antagonistically arranged shape memory alloy coils. The segments are arranged with three legs with rounded feet, connected by two spine segments. When one SMA is powered, the coil contracts, inducing a bending moment in the segment and causing the segment to bend towards that coil. SMA heating is accomplished via an off-board power supply controlled with a dedicated microcontroller and pulse-width modulation (PWM).

Horton is equipped with two modalities of onboard sensing: five bend sensors, one per limb, and ten thermocouples, one per SMA. Due to the planar arrangement, monocular camera-based sensing is compatible with this platform. These sensors are also connected to the off-board microcontroller.

Detached Horton limbs have been stabilized to desired angles using feedback control with an external camera [58]. Recurrent neural network models have been trained to accurately predict open-loop evolution of Horton limb state on timescales of 10 minutes at a time with little drift, using limb thermocouple data and PWM signal as network input [59]. For the full Horton robot, safe balancing control has been demonstrated [60], yet effective walking trajectories for the full Horton robot remain elusive. While not intended to produce numerically rigorous solutions, the following brief experiment intends to shed some light on the challenges leading to this difficulty.

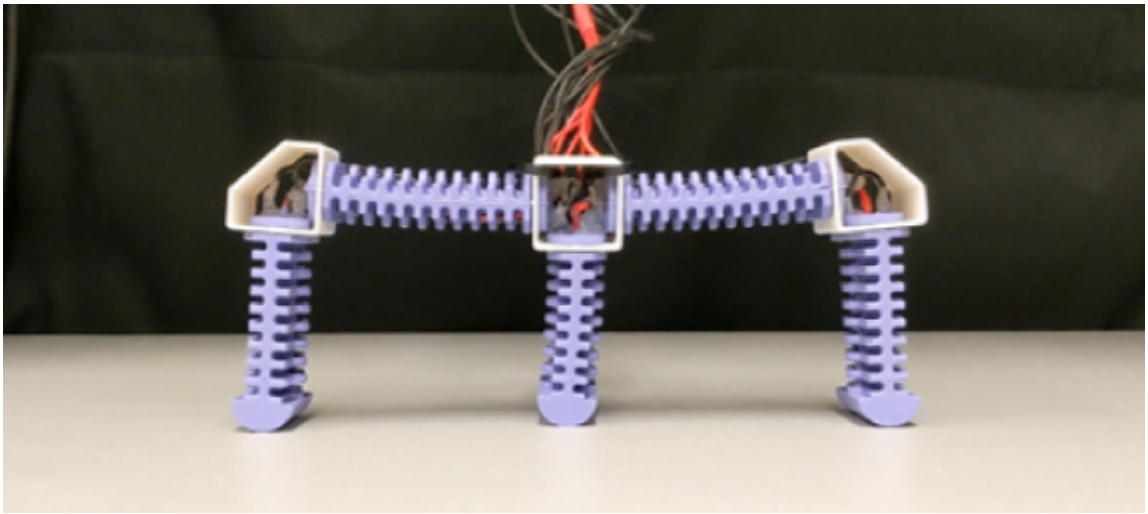

FIGURE 2.9: Horton, a shape memory alloy robot with three legs and two spine segments.

2.3.2 Thermal Dynamics of Open-loop Walking

As a shape-memory actuated robot, the action space of Horton is abstracted to trajectories of pulse-width modulated actuator input. These trajectories are of the form

$$\mathrm{PWM}_{i,t}, \quad i \in [0, 10), \quad t \in [0, T]$$

where i represents the SMA index, and t represents the time-step in period T. Each SMA coil receives a commanded duty cycle at every time step according to the schedule defined by the trajectory. I attempted to determine if there were qualitative differences in how a given trajectory manifested in Horton's motions given repeated actuation. Since actuation occurs in Joule heating, the lingering effects of actuation history should be visible in the robot's thermal signature. To judge this, an experiment was set up as shown in Figure 2.10, using a FLIR thermal camera to record both the temperature and gross motion of Horton.

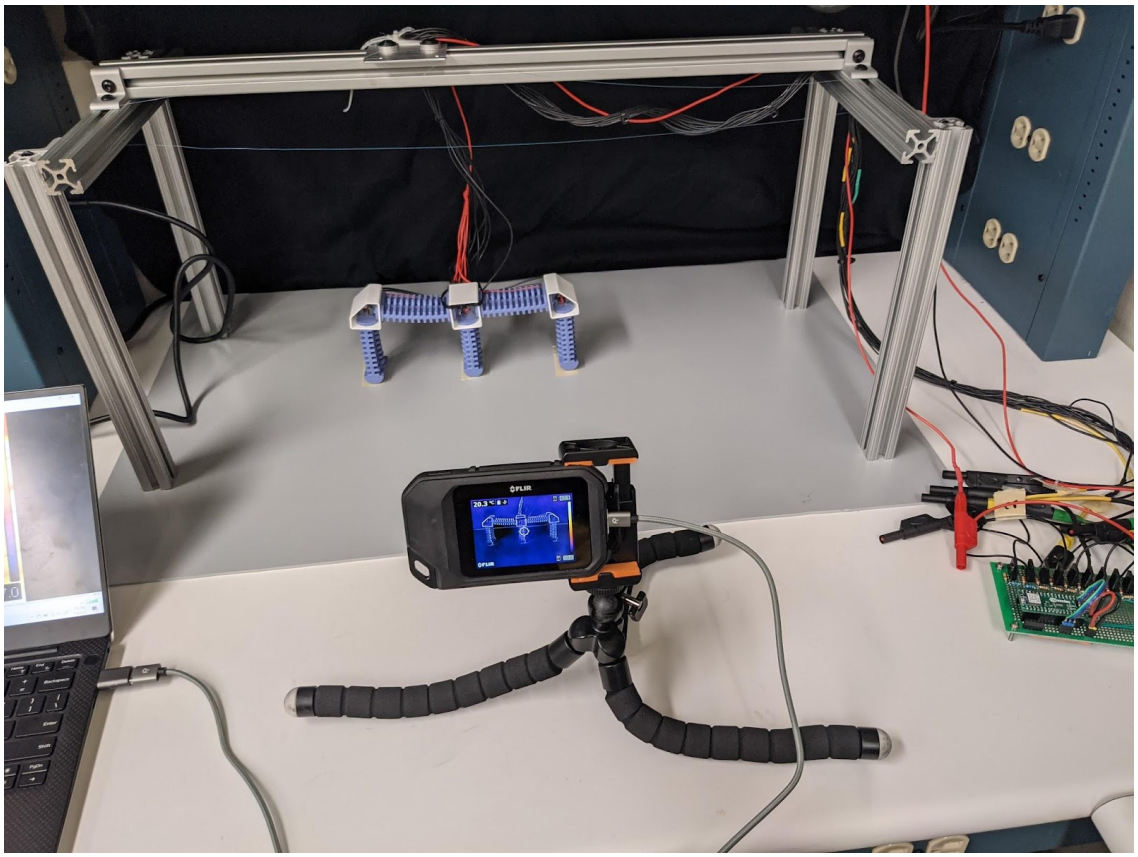

FIGURE 2.10: The experimental setup for mapping thermal state to motion.

The same actuation trajectory was executed 20 times, manually resetting Horton to the initial position and then running again without pause. As shown in Figure 2.11, this ongoing operation of the shape memory alloy actuators leads to an accumulation of heat in the SMA and silicone of the affected legs. While fully cooled, Horton's motion differed from all subsequent motions, which were executed under residual heat. As we see in

Figure 2.12, the differences in motion were not due to noise and could be replicated under similar thermal conditions.

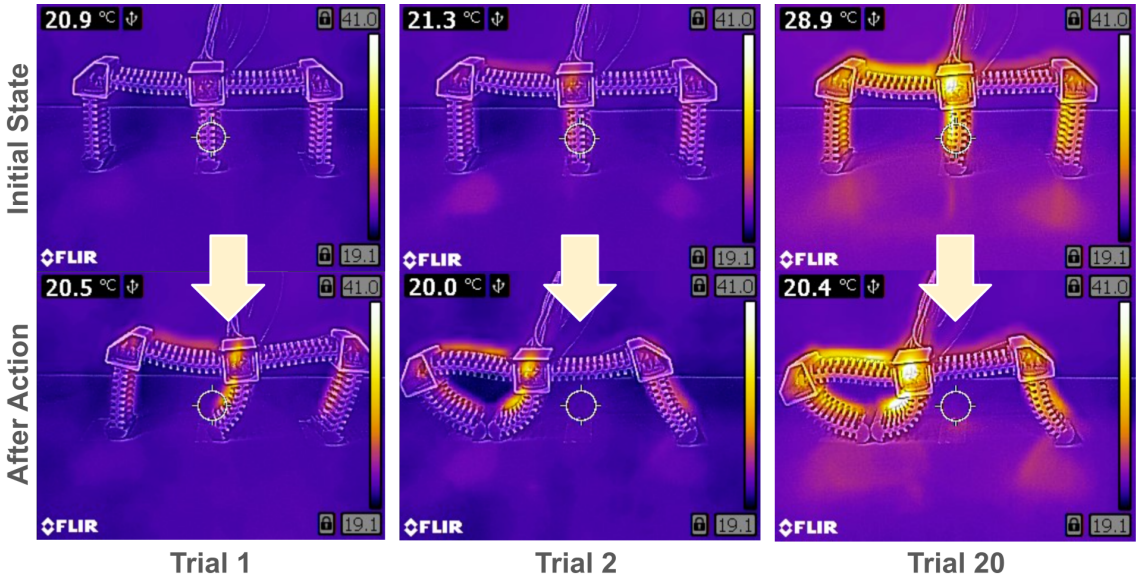

FIGURE 2.11: Motion variance due to accumulation of heat. The same control input is sent to the shape memory alloy actuators over time. The variation is at the level of wholly different contact modes.

Discussion

While I expected minor numerical differences in Horton's trajectory, the observed changes were qualitatively distinct, resulting in entirely different contact modes. Shape memory alloy functions by changing shape based on the volume fraction of different crystalline phases of the alloy. The alloy's temperature determines the volume fraction at a given moment. When this shape change occurs, it is constrained by the loads experienced by the alloy. Due to these two factors, a loaded SMA coil that appears in a given shape can have many internal states that are mostly unobservable but relate to its temperature. The effect of executing a given actuation trajectory on an SMA limb depends on these internal states. This kind of change can cause the system to move across the different contact modes induced by the system's geometric instability, which provides insight into the difficulty of achieving walking. In a contact-rich context, small changes on the microscale of the soft materials and soft actuators can cascade into qualitatively distinct macroscale behavior. The fact that the motion can be qualitatively replicated using similar thermal conditions presents two potential solutions to this fragility of control. The first solution requires the utilization of increased sensing capabilities. If enough of the thermal and force state is captured, extending the state space used for decision-making would be possible. However, this poses the problem of increasing monetary and computational costs, which runs counter to the goal of making these systems function with minimal cost. The second solution requires learning to adapt to the underlying system dynamics as it changes on time, changes over time, through the conditioning of actuation either on estimated internal states or via directly updating the actions themselves. This latter approach is more

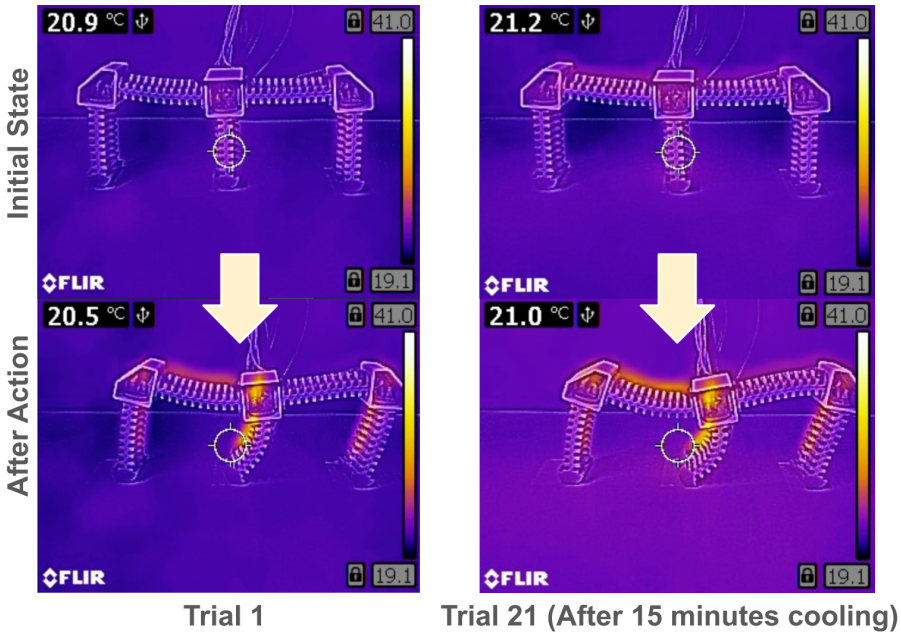

FIGURE 2.12: The effect of thermal buildup is validated by demonstrating that the original effect of the actuator sequence is recovered after sufficient cooldown time.

promising due to the possibility that the latent dynamics governing these changes are not wholly dependent on the observable temperature features or known actuation history. There is reason to suspect this due to the intrinsic time-dependent dynamics of soft materials like hysteresis and creep and the fact that while thermal sensing was able to achieve good performance in single limb control [59], no such success has been demonstrated in the kind of contact-rich contexts that are a significant draw for soft material systems.

Up to this point, the soft robot systems I have shown share some key properties. They are all locomotion platforms actuated by shape memory alloy. Thermal state appears to be a key variable in shape-memory alloy actuated systems, and there is ongoing progress in making approaches that generalize across shape-memory alloy systems using, e.g., Gaussian process models [61]. However, these systems are not the most promising for enabling robotic system development in a domestic or low-resource environment. The challenges to suitability are partly because the specialized nature of expertise required to form strong intuition upon which to build simplifying abstractions like action primitives, which have been key to lowering the amount of data necessary to achieve the operation of these soft systems. In addition, the actuator dynamics of these systems dominate the dynamics of the soft materials themselves, and a robust understanding of these materials is critical to enable the accessible development of soft robotic systems from various context-relevant materials. To generalize my identification of the challenges of soft robots beyond this narrow class and to focus more explicitly on the accessible hardware of soft systems, I study a low-cost tendon-driven soft finger.

2.4 An Attempt at Learning Soft Finger Proprioception

State estimation for soft robots often rely on soft material sensors for external sensing and proprioception [62]. I examine a soft finger designed to explore state estimation and proprioceptive sensing in soft manipulators using a novel foam-like magnetic sensor, based on the work on tactile sensing with a soft magnetic sensor by Hellebrekers *et al.* [63], shown in Figure 2.13. I aim to establish a mapping between changes in the magnetic field measured within the foam and the 2-DOF position of the bending finger.

2.4.1 Soft Tendon-Driven Finger

The hardware is composed of a cylindrical foam finger, with a section of magnetic microparticle and foam composite in the finger's base. The finger is mounted to a 3D-printed stand with a magnetometer in the base (Figure 2.14A, B). The finger can bend along two perpendicular axes, driven by tendons routed through the fabric covering and the printed base (Figure 2.14C).

This system integrates a fabrication strategy based on the foam hand designed by King *et al.* [64]. The finger is molded from Smooth-On FlexFoam-iT! X. This molding is done in three parts: first, the bulk finger is molded, then the magnetic foam is created, and finally, the two are joined via over-molding. Uncured urethane foam is mixed with magnetizable microparticles in a low (<20%) mass fraction to create the magnetic foam. This uncured foam is placed in a small mold mounted rigidly to a neodymium permanent magnet to align and magnetize all the suspended particles during curing. The cured magnetic foam is placed within a larger mold with the bulk finger material. These two foam bodies are then joined by overmolding with additional foam.

A felt sheet is cut into a flattened geometric net of the finger shape and adhered to the surface of the finger with upholstery adhesive, leaving the base free. The edges of the felt form are then sewn securely together. The base is assembled to a platform containing a PCB with a magnetometer in the center, as shown in Figure 2.14. Tendons made of 65 lb test Teflon-coated fishing line are anchored to the top of the finger at four evenly spaced points, then run down opposing sides of the finger via a running stitch in the felt and routed through the holes in the support platform.

2.4.2 Magnetic Proprioception

The magnetometer embedded in the base senses a magnetic vector formed from the integral of the magnetic moments of all the magnetized particles that are embedded within the foam. These moments are initially determined by how the particles are aligned during the curing process. They are further modified by the deformation of the foam as the finger bends. Due to these factors, it is impossible to determine *a priori* what the magnetic field should be given a particular finger configuration. Instead, we create a dataset by sampling servo positions and the resulting image-based ground truth finger orientations and magnetometer readings. The finger position is driven by two servos, one per bending axis, as shown in Figure 2.15. Each servo carries two cable spools in opposing orientations, driving the antagonistic finger tendons. The positions are recorded via an overhead camera, simplifying bending along each axis to displacement in the x and y axes. A 3-DOF magnetic signal was recorded from the magnetometer for a given tuple of servo positions. When plotted across the Cartesian product of each servo in increments of

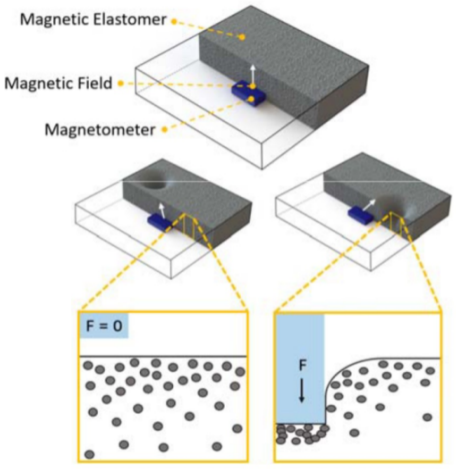

FIGURE 2.13: Function of magnetic elastomer. This work uses an elastomeric foam instead of a solid elastomer. Elastomer foam is embedded in the finger's base, and a magnetometer is mounted to the base stand. Sensor design and this figure derived from Hellebrekers *et al.* [63].

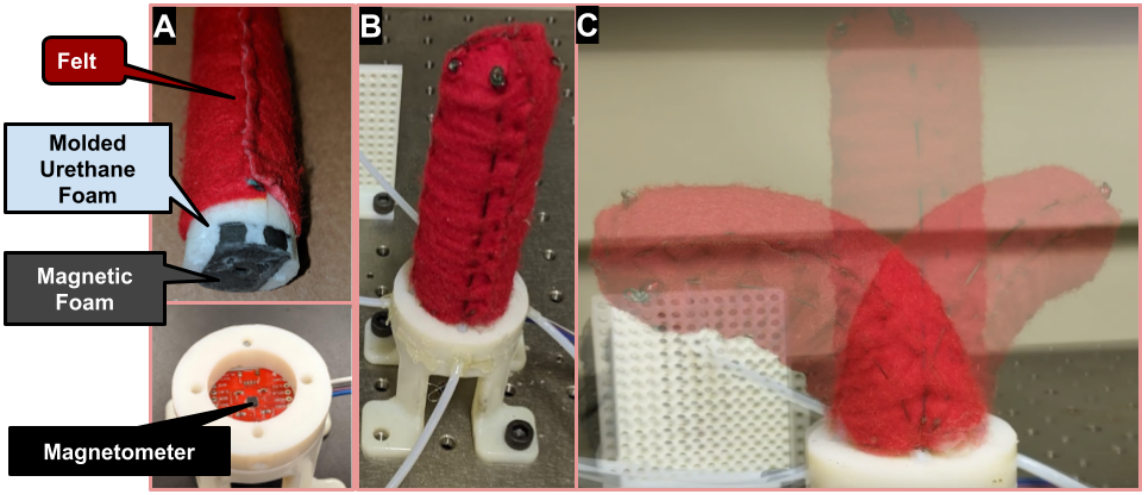

FIGURE 2.14: Design of tendon-driven foam finger with magnetic elastomer sensor in base. **A)** Components of finger and magnetic sensor. textbfB) Assembled finger in base. **C)** Bending motion of finger along one of two available axes.

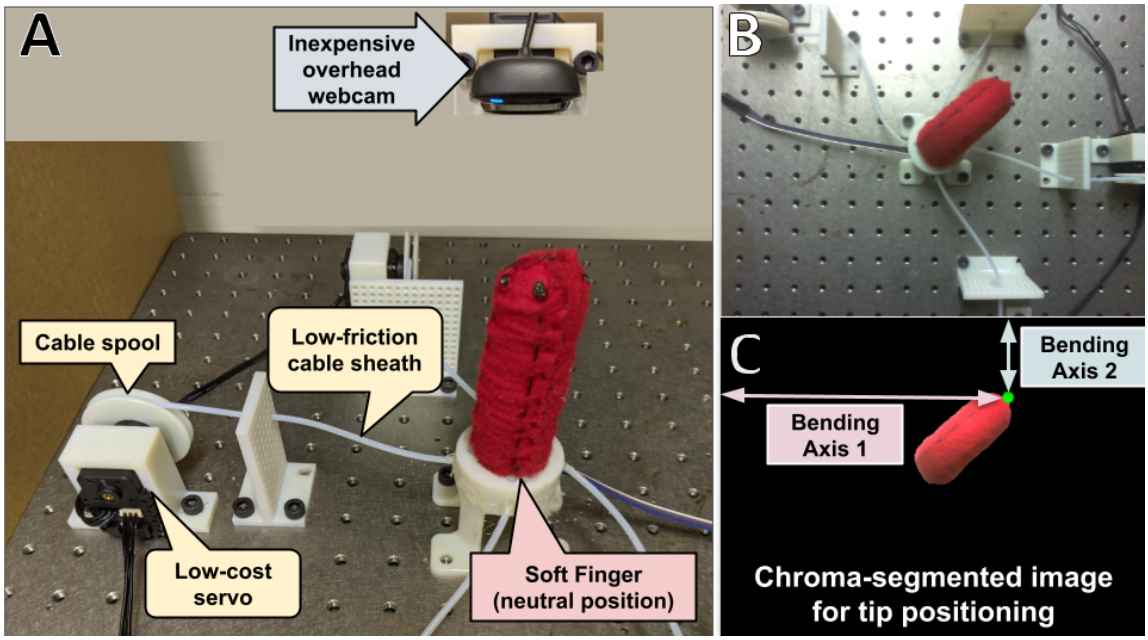

FIGURE 2.15: Experimental setup for determining proprioceptive mapping between finger position and magnetic sensor readings.

1/9 of their range (0 ticks to 1024 ticks), the resulting point cloud is illustrated in Figure 2.16.

This data was then used as the training set of a fully connected feed-forward neural network, which learned a mapping from magnetic-field signatures in the base of the finger to predicted servo position. This mapping is an intermediate step in establishing the desired mapping to the ground-truth camera data. Datasets were collected with the servo ranges discretized to varying resolution, and neural networks were trained with each set. The different networks' predictions were tested against the commanded positions of the test set.

2.4.3 Discussion

Despite the ostensibly simple intermediate task of learning a bijective mapping between servo position and magnetometer reading, a bijective mapping only existed over short time scales. There were multiple potential causes of this, stemming from the fact that the system is an inexpensively fabricated soft robotic system. These include magnetometer drift, fabric detachment, and tendon repair. Magnetometer drift was demonstrated to happen at a timescale of ten days, as would be expected as a result of the drifting magnetic field of the Earth. However, this effect was quantitatively negligible. More significantly, I observe that fabric detaching occurs when the felt skin of the finger no longer adheres to the foam core. Fabric detachment causes changes to the force transmission of the tendon, as motion is converted into deforming the skin more and more, eventually leading to the failure of the finger. Major cases of this were considered fabrication errors and are rectified by remaking the finger, but minor detachment could occur without being detected and would be expected over longer timescales.

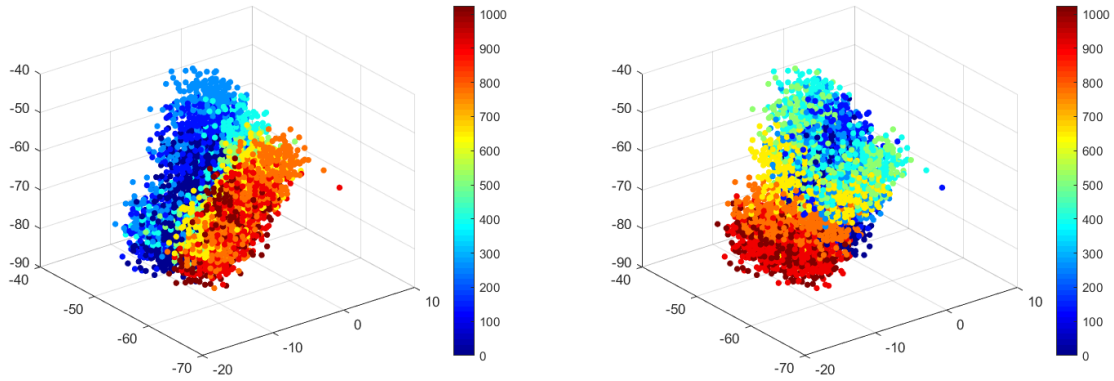

Servo 1 groups Servo 2 groups

FIGURE 2.16: 3D Magnetometer points grouped and colored by biaxial servo positions.

The most relevant source of the system dynamics change I encountered is the effect of tendon replacement. After a tendon breaks, the replacement always has a different tension. Ideally, there would be no loading on the magnetic foam in the zero position, but noticeable tendon force is required to prevent slack in the tendons. Slack induces directional history dependence to the finger's motion, breaking the basic assumption of a bijective proprioceptive mapping. Without slack, as shown in Figure 2.17, the tension changes lead to changes in the magnetometer signal for a given servo configuration. This data was collected while attempting to augment the training set from a previous day after a tendon broke during testing and needed replacement.

The changes to these properties over time are a form of nonstationary. Because of this nonstationary behavior, one-time training of networks for encoding combined system-sensor behavior is infeasible. For the identified sources of nonstationarity, it is conceivable that this problem might be alleviated via supplementary sensing. For example, one might be able to calibrate the magnetometer drift via external magnetometers [63]. If the deformation state of the magnetic foam in the finger base is entirely dependent on the tension of the tendons, then adding force sensors in series with the tendons could allow for correction of that as well. However, there are no simple methods for sensing the adhesion of the fabric to the finger. Critically, these solutions increase the cost and complexity of the platform. In addition, there is no guarantee that these sources are the only causes of nonstationarity, and without identifying the underlying features that change over time, heavy investment in sensing might miss the relevant factors. Some of these factors might be intrinsically hidden. The dynamical effects of aging of urethane in the foam core of the finger, for example, are not well characterized. Our ideal case would be adapting to these nonstationary mappings without increased sensing hardware.

2.5 Identifying Challenges and Promising Techniques

This chapter presents algorithmic approaches for enabling autonomy in soft systems. In Section 2.1, I show that closed-loop behavior in an untethered SMA-powered underwater

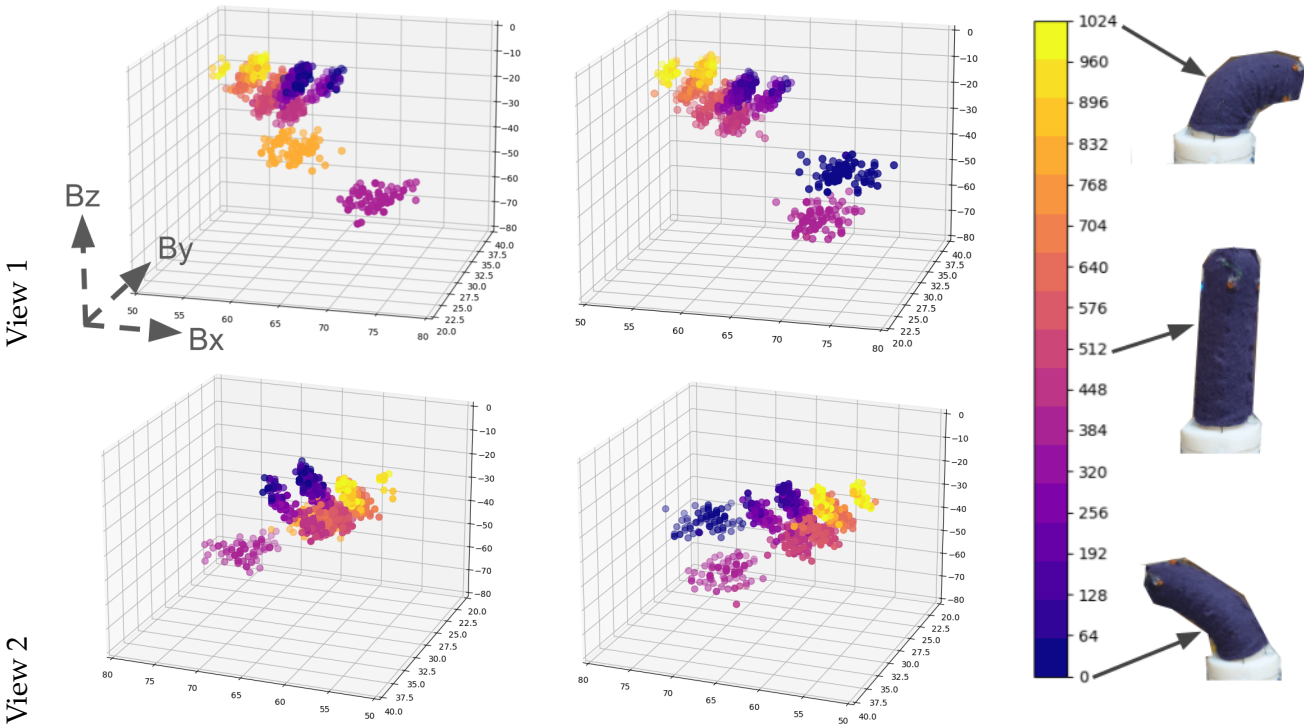

FIGURE 2.17: 3D magnetometer data, where axes correspond to x, y and z components of the magnetic field. The distribution of points for each finger position changes after finger tendon is replaced.

crawling robot can be enabled by external vision-based sensing combined with reduced representations of that robot's state space and action space. Due to the high dimensional action space, discrete action primitives make things more tractable. I show that practical action primitives can be created for even complex shape-memory actuator dynamics. I demonstrate that these enabling techniques rely on expert intuition and are tightly bound to the analysis of the soft actuator dynamics. I show that rough, imprecise forward models of soft robot dynamics can be useful if the task space is forgiving enough, e.g., able to be reduced to a quasistatic state, with feedback and minimal need for accurate prediction of system evolution multiple time steps in the future.

In Section 2.2, I demonstrate that trajectory-following can be achieved in an inertial SMA-powered swimmer via a pre-computed library of transitions and external visual feedback. This strategy is only compatible with systems that can be simulated with high fidelity. There is a trade-off between fidelity and speed. However, for some specific cases, like for systems that can be represented as discrete elastic rods, simulation can be made faster than for general soft systems, though still not nearly as fast as what is available for rigid systems (remaining within an order of magnitude of real-time). Finally, effective calibration of these simulations faces several challenges. First, their fidelity drops significantly as they try to extrapolate to unseen states or even interpolate between known states. The difficulty of generalization provides a solid incentive to collect more hardware data, but this isn't very easy because this kind of fully-soft integrated system is not optimal for long-term experimentation due to the need for frequent repairs. In addition, the dynamics of one shape-memory coil do not intrinsically transfer to another, so as these repairs are required, the data from old actuators is not trivially integrated into this scheme. Secondly, determining which physical phenomena to model is non-trivial, as adding extra physical phenomena increases the state space of the simulation, further increasing the computational cost. In addition, soft materials are susceptible to the effects of many more physical phenomena than rigid material systems, and it has not been well established which phenomena these are or how those phenomena might interact with a robot's behavior as it moves through state space or evolves. In short, it is not clear what the relevant latent states are.

In Section 2.3, I show that small changes to thermal state cascade into macroscale behavior for shape-memory alloy actuated robots in a contact-rich walking task. While this problem is solvable for low-contact, single-finger systems via temperature estimation [58] or additional of thermal sensing [59], these solutions rely on the behavior of the system being dominated by thermal and crystal-phase dynamics of SMA, not generalizing to contact-rich systems or robots with different actuation schemes.

In Section 2.4, I show that magnetic proprioception in an inexpensive soft sensor is feasible at very short timescales. However, the system is prone to changing the relevant dynamics over time. Some of these changes are predictable *a priori* due to known dynamics of soft materials, like hysteresis, polymer aging, creep, and other inelasticities. Other changes arise from the robotic system design and the interactions between components, like the slack induced by tendon stretching or the changes to pressure distributions if repairs are made to the finger. These changes are quite difficult to sense directly if known, and inferring them explicitly requires models that are not available for novel systems.

I identify the following techniques that show promise in enabling autonomy in soft systems to some extent.

1. Simplifying the operational requirements by slowing down time-scales and creating reduced state and action representations.
2. Adding additional sensing so that relevant state can be included in the model.
3. Performing more pre-computation across distributions of simulation parameters to create more robust Sim-2-real behavior.
4. Identifying and characterizing more context-specific phenomena in the operation of soft robotic systems.
5. Using online machine learning or adaptive control methods to enable robots to adapt to changing parameters.

Let us evaluate these approaches in the context of my broader goal of enabling the accessible development of intelligent robotic systems.

Simplifying the operational requirements of the robot has been an influential and practically necessary tool for the systems I examined in this chapter, especially when using soft material actuation like in PATRICK, the frog-inspired swimmer, or Horton. As long as computational resources are limited, reducing the action space of actuators like SMAs is likely necessary. The intrinsically suboptimal results produced by only exploring the space induced by action primitives are the kind of compromise that could be more acceptable in a low-cost development environment. However, the effects of this suboptimal behavior might bleed into the safety features that make soft robots so appealing for domestic environments, or due to, e.g., power consumption inefficiencies, the sustainability concerns that arise when we advocate for the more widespread development of robotic systems. Therefore, relying on this as the primary tool is not appealing and would ideally be kept to the minimum necessary.

Adding additional sensing is effective but intrinsically increases hardware costs, development time, and engineering complexity. Similarly, relying on more computation across varying parameters would be helpful, but the costs associated with computing must be considered, especially considering the cost of high-performance computational hardware. For both of these approaches, it is critical to achieve a better understanding of the underlying dynamics of the systems used. This kind of understanding is deeply tied to the construction of the system and its operational context, meaning that there is likely no general prescription for which sensing modalities need to be added or which phenomena need to be added to a simulation or model's parameter space. However, once relevant phenomena are better understood, incremental incorporation of these strategies can become feasible depending on the exact level of resources available in the research and development environment.

The relevant phenomena we need to characterize will depend on our choices of materials used in the construction of the system and the choice of actuator. In my view, focus should be placed on understanding the dynamics of materials that are accessible and able to be made in increasing amounts without creating a new sustainability issue. The practicalities of working with accessible and sustainable materials will be the focus of Chapters 4 and 5. I focus on biomaterials as an avenue for accessible and sustainable materials. Due to the environmentally sensitive, dynamic, and nonstationary behavior of these systems, it is likely that data collection longer than the tens of minutes seen in existing soft robotic platforms is needed.

To enable machine learning on robotic hardware, it seems clear that platforms that can collect a significantly greater amount of hardware data are needed, as learning in simulation does not seem feasible given the current performance of those systems. Such platforms would also enable longer-term data to be collected on soft material dynamics to inform choices in state estimation and representations. It could also feed back into improving the effectiveness of simulations by providing more data upon which to calibrate and validate [65]. In Chapter 3, my development of a platform designed to enable these avenues of research is shown.

Chapter 3

Enabling Soft Robotic Hardware Datasets

In this chapter, I motivate the design of a robotic platform with several properties that enable soft robotic systems to be developed more accessibly. I then describe the five-bar pendulum, a prototype system I designed and implemented for this goal and the lessons learned from that system. Finally, I describe the modular parallel manipulator, a simple, robust platform that is able to collect data for robotic operation for days at a time. This platform allows for machine learning directly on soft robotic hardware and demonstrates the ability to characterize differences in behavior linked to material properties. I demonstrate the ability to explore the control of soft robotic systems over time and collect data correlated with latent features of soft materials. This data collection ability provides a potential path to creating soft robotic systems from accessible, unmodeled materials and to building controllers for those systems that can handle the long-term dynamics challenges of soft materials.

3.1 Related Work

To clarify the design requirements of a platform for long-term soft material experiments, I first discuss machine learning approaches for enabling soft robot autonomy that rely on the ability to collect hardware data reliably (Figure 3.1). I then discuss approaches to the critical issue of nonstationary dynamics, which govern system dynamics that vary with time.

3.1.1 Machine Learning for Soft Robotics

Content in this section replicated and adapted from ***K. Chin**, T. Hellebrekers, and C. Majidi. "Machine Learning for Soft Robotic Sensing and Control." **Advanced Intelligent Systems**, vol. 2, no. 6, pp. 1900171, John Wiley & Sons, Ltd, 2020.* [28]

Author Contribution: I was the primary author and wrote content reviewing machine learning for the control of soft robots. Tess Hellebrekers wrote content reviewing machine learning for soft sensing.

Before describing the learning and control testbed, I first review data-driven machine learning approaches to sensing and control systems that aim to manage the dynamics inherent in soft robotics. Soft robotic systems are often operated using open-loop control methods that leverage the inherent compliance and conformability of soft materials as they press against contacting surfaces. Popular examples of this are compliant grippers

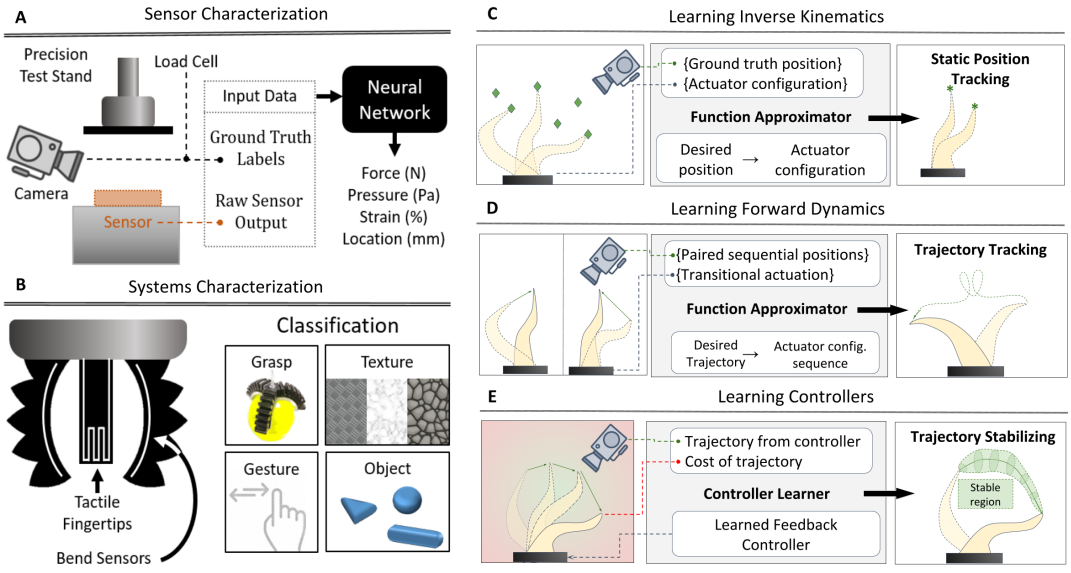

FIGURE 3.1: Overview of machine learning strategies common in the field, reproduced from my paper "Machine Learning for Soft Robotic Sensing and Control" [28]**(A)** Sensor characterization involves collecting raw sensor data and ground truth observations in a controlled and monitored environment (e.g., lab setting). This data is typically provided as input to a neural network to predict further values. **(B)** In contrast, system characterization collects sensor data in a less controlled environment that mimics the entire system's use in the field. In this case, ground truth measurements such as force are more difficult to obtain. Users often circumvent this by mapping to higher-level classifications, such as grasp success and texture recognition. **(C)** Learning inverse kinematics requires training data composed of matched pairs of robot position and actuator configuration. A function approximator (neural network, etc.) encodes the relationship between sampled data to generalize to arbitrary desired positions. **(D)** In learning forward dynamics, sets of two sequential states and the actuator action that caused the transition between them are observed. The dynamic behavior is encoded into a neural net, allowing for open-loop following of arbitrary trajectories. **(E)** Learning controllers requires iteratively evaluating the performance of the resultant trajectory for a controller with respect to some cost function and updating parameters to increase that performance via optimization

that can conform to a wide range of objects with limited sensing [64] [66] and soft robots that can passively shape change as they pass through confined spaces [67]. While material compliance enables intrinsic reactivity in the form of material deformation upon contact, there is still a need for the ability to control the positioning of soft robotic systems in free space or to specify contact beyond simply modulating force. Computing the deformation of soft materials in isolation can be time-intensive. Creating *a priori* models of soft robotic systems, whose behavior might depend on interactions between soft materials with specific, potentially novel compositions or geometries, is usually ineffective. Therefore, to create models or controllers for soft robotic systems, data-driven methods are quite appealing.

Recent efforts to enable the control of soft robots through machine learning have addressed different parts of the control pipeline. The two most common tactics are learning the inverse kinematics/statics of soft manipulators [68] and learning the forward dynamics model to enable predictive control [69]. A less common strategy is the direct learning of controllers for specific behaviors [70].

Supervised learning is the basis of much of the existing work in learning the kinematics of soft robots. If the system is treated as quasistatic, an open-loop mapping between actuator configuration and task-space position is achievable. For redundantly actuated systems, a one-to-one mapping can be enforced by optimizing an additional metric, e.g., actuation effort. Such a bijective mapping is ideal for supervised learning from labeled pairs of training data. To collect training data for learning the kinematics function, a random walk through actuation space on the physical hardware (motor babbling) is a straightforward method and has been used for cable-driven continuum manipulators [71], pneumatic continuum manipulators like the Bionic Handling Assistant [68] and honey-comb pneumatic network manipulator [72], and simulated manipulators [73, 74]. Data can also be collected from human demonstration by manually replicating important configurations and recording the associated actuator parameters [70]. Data-driven methods have been used to achieve locomotion with soft systems as well [53]. However, straightforward supervised learning tends not to scale to the dimensionality and dynamic sensitivity of mobile systems.

The forward kinematics of soft robots are usually nonlinear. For example, in the case of a pneumatic continuum arm, the vector of pressures in segment chambers, i.e., the configuration space, maps to the pose of the manipulator tip, i.e., the task space, through a force equilibrium function that depends on local material properties and design geometry. This mapping between actuator configuration space and task space can also be contingent on the environment. For example, the curvature response of a segment to a given pressure can vary significantly in the presence of external forces like gravity or surface forces. A global representation of the associated Jacobian is difficult to convey analytically, even without concerns of non-stationarity and stochasticity. Inverting the Jacobian once found is also quite difficult. Because of this, a common technique in the field is approximating the global Jacobian with a function approximator, which is most commonly a multi-layer neural network [71–74], but can be of other forms, like linear function approximators [75], constrained extreme learning machines [68], or Gaussian mixture models (GMMs) [76].

It should be noted that this mapping can be found for a local configuration without learning any model by performing gradient descent over the actuator configuration with the desired position as the target [77]. The downside of this method is that it induces

undefined motion iteration in the robot between each commanded position, which is potentially unsafe and results in much slower execution. The global inverse Jacobian can also be approximated by aggregating an ensemble of learned local Jacobians sampled from across the full state space [74, 76]. Explicit global representations are compact, while ensemble methods allow prioritization of performance in the regions of most interest.

By virtue of their mechanical compliance, the equilibrium state of soft systems is dependent on external forces. Replacing the purely kinematic mapping between the robot's configuration and task spaces with a function that incorporates the force equilibrium is effective for position tracking of continuum manipulators [68], [73]. Even without explicitly learning the force balance, the combination of end effector feedback with this compliance can be used for disturbance rejection in simulation [71]. Due to this decoupling of kinematic shape and external force balance, joint position-stiffness controllers can be learned for soft manipulators [75, 78].

Quasi-static approximations of robotic systems are useful for systems with slow time scales or tasks where timing is not critical. However, for many applications, dynamic control is more appropriate. The forward dynamics of soft systems are even more difficult to model analytically than their kinematics, so machine learning techniques have been applied to these problems in various ways. The combination of dynamics modeling with position control has been shown to enable time sensitive trajectory tracking in a hybrid soft-rigid arm [78], and a one-DOF pneumatic finger [79]. Dynamics models can be learned that are adapted to varying dynamic parameters of the system, e.g., with stiffness tuning soft manipulators [80]. Exploiting system dynamics is a benefit of learning the dynamic system behavior and has been shown to increase manipulator workspace for a continuum arm [69]. The work of Thuruthel *et al.* [81] on using learned dynamic models exemplifies a significant step in the evolution of soft manipulators enabled by machine learning methods. Earlier examples of soft continuum manipulators, such as the work of Rolf *et al.* [68], were more mechanically sophisticated than the system used here. However, the full mechanical ability was not leveraged due to the lack of a useful dynamics model. Thuruthel *et al.* demonstrate that an approximate dynamics model can be learned and is sufficient for achieving performance beyond what is possible with quasistatic methods.

Learning kinematic or dynamic models for a soft robot means that while part of the control pipeline relies on learned models, the controller is still engineered. Reinforcement learning is a machine learning strategy that allows controllers to be created through learning from data collected using sequential environmental interactions [82] . Policy gradient-based reinforcement learning converges to a locally optimal controller without an analytical model of the robot dynamics but requires more time and data for training than supervised learning. The need for more time and data is due to the requirement of evaluating the trajectory produced by following a controller from a specific state before updating the model at a given optimization step. A common robotics solution is learning in simulation for many trials, as for the complex tendon-driven humanoid hand in [70]. The controller learned in simulation can be further refined on the physical robot [83]. Trajectory optimization is a common numerical optimization method for generating local controllers, which can be implemented using machine learning. Machine learning-based trajectory optimization has been used to generate a library of trajectories for a mobile tensegrity robot [84] and a forward walking trajectory for a cable-driven soft

quadruped [53]. Zhang *et al.* [84] shows that a global controller can be learned by using local trajectory-stabilizing controllers' outputs as supervised learning labels. Besides leveraging local solutions, other ways off reducing the amount of data required for learning have been used. The humanoid soft hand used by Schlagenhauf *et al.* [70] is much more mechanically expressive than the previous works using continuum soft manipulators. A global dynamic or kinematic model for this system would require an unfeasible amount of data for hardware collection. By creating a basis out of a few finite poses, the manipulator can perform dexterous motion without a full model of system behavior. Leveraging the ability to create useful behavior without globally valid solutions could enable an evolution in the complexity of potential soft robotic systems and the performance of existing systems.

The strengths of empirical, data-driven knowledge and understanding drawn from first principles can reinforce each other. Hybrid approaches allow leveraging existing knowledge so only the most intractable system components need to be learned. Learning the parameters of an analytical dynamics model, similar to traditional adaptive control methods, is fast and effective if such a model can be constructed with enough fidelity [78]. It is also possible to decompose control of multi-actuator systems into analytic kinematic targets, where each actuator achieves the final shape through a system-level controller [68] or individual actuator-level controllers [72]. Model predictive control strategies can be used on systems with uncertain models by using neural networks to encode approximate system dynamics models. [79].

Machine learning has the potential to overcome the challenges of understanding soft material behavior from first principles. However, to enable the operation of soft material-based robotic systems over longer periods of time, as would be required of a useful tool, the way that soft materials change over time is another reason to move towards the ability to collect larger amounts of hardware data.

3.1.2 Control of Systems with Nonstationary Dynamics

While soft robots can be easily built, effective control strategies are less clearly available. Soft robot dynamics are difficult to accurately model analytically due to a multi-physics coupling between shape, forces, physical state (e.g., temperature), and history of motion. To enable the operation of soft robots across longer time scales, it is important to be able to collect hardware data that capture these phenomena, especially those that vary with time. As shown in Figure 3.2, the deformation of elastomeric materials has multiple forms of time-dependent phenomena. There is hysteresis, which means there is no one-to-one mapping from actuation force to deformation – instead, the result of any applied force or deformation depends on the system's recent history. There is also nonstationarity, which is a distribution shift in the underlying dynamics of a system over time and can be caused by wear, external temperature fluctuations, internal strain build-up, or many other sources. Hysteresis and nonstationarity are inherent side effects of using elastomers and other soft materials in constructing soft robot systems. Popular methods often ignore these effects or treat them as unmodeled noise [85] [68]. Hysteresis can theoretically be addressed by incorporating system state history into the input of models. This explicit time dependence can be encoded in structures like recurrent neural networks [69] or by simply concatenating multiple time steps of state data as the input to the model. Nonstationary behavior necessitates models that can adapt to changing dynamics [68, 86].

These inherent dynamics present an essential problem to the autonomy of soft robotic systems (though deliberate changes to dynamics are an emerging feature of soft robotics [87]) and may also function as a reasonable proxy for the general problem of imprecise dynamics. Errors in modeling and the drift in models over time exist in rigid-body robotic systems. As described in a report on the problems faced by teams in the DARPA Robotics Challenge (DRC), on the scale of full humanoid systems, modeling errors due to "wrong kinematic and inertial parameters, cogging and other magnetic effects in electric motors, actuator dynamics and hysteresis, variable viscous or nonlinear dynamic friction depending on the state of hydraulic oil leakage, dust, and dirt, thermal (both external weather and actuator heating) effects, humidity effects, constant delay, variable delay due to computer communications. Joint/actuator/transmission stiction and other static friction effects, foot slip-stick on touchdown and during stance, six-dimensional bearing play, structural link deformation, and material aging" [88] all come into play. More generally, most classes of the robotic systems can encounter problems with dynamics that do not match the model, either due to misspecification or nonstationary dynamics. This challenge is especially relevant for soft robots and systems composed of mechanically compliant materials.

In traditional control theory, the nonstationarity problem is among those problems addressed by the subfield of adaptive control. Traditional adaptive control methods often rely on high-fidelity equations of motion with a small number of uncertain parameters [89]. There has been work integrating these techniques with neural network models, an approach called concurrent-learning adaptive control (CLAC) [90]. Other approaches have leveraged state-estimation techniques to update analytical parameter sets (Kalman) [91].

In the reinforcement learning literature, there have been a few works attempting to solve the problem of learning control under nonstationary dynamics via multiple partial models. These generally train several neural networks for different regions of the drifting dynamics. Multiple Model-based Reinforcement Learning (MMRL) [92] creates a static library of neural models that correspond to a known set of dynamics and are trained as the system moves through those different modes. Reinforcement Learning with Context Detection (RL-CD) [93] is trained similarly, but rather than a static set of models, there is a context detection module which allows for new models to be created as the system encounters different modes. Hierarchical Reinforcement Learning with Context Detection (HRL-CD) [94] combines the technique of hierarchical reinforcement learning for accelerating learning convergence with the RL-CD framework. All of these techniques assume deterministic and discrete dynamics.

There has been work in enabling faster transfer learning for problems in related domains using neural attention and transformer networks [95]. Similarly, automatic domain randomization (ADR) has been shown to allow for the online formation of reactive controllers that can adapt to a wide range of environmental parameters [54, 56, 57]. These methods rely on extremely high-volume data collection for training, requiring high-fidelity simulation and computing time on a scale that can be inaccessible to resource-constrained development environments.

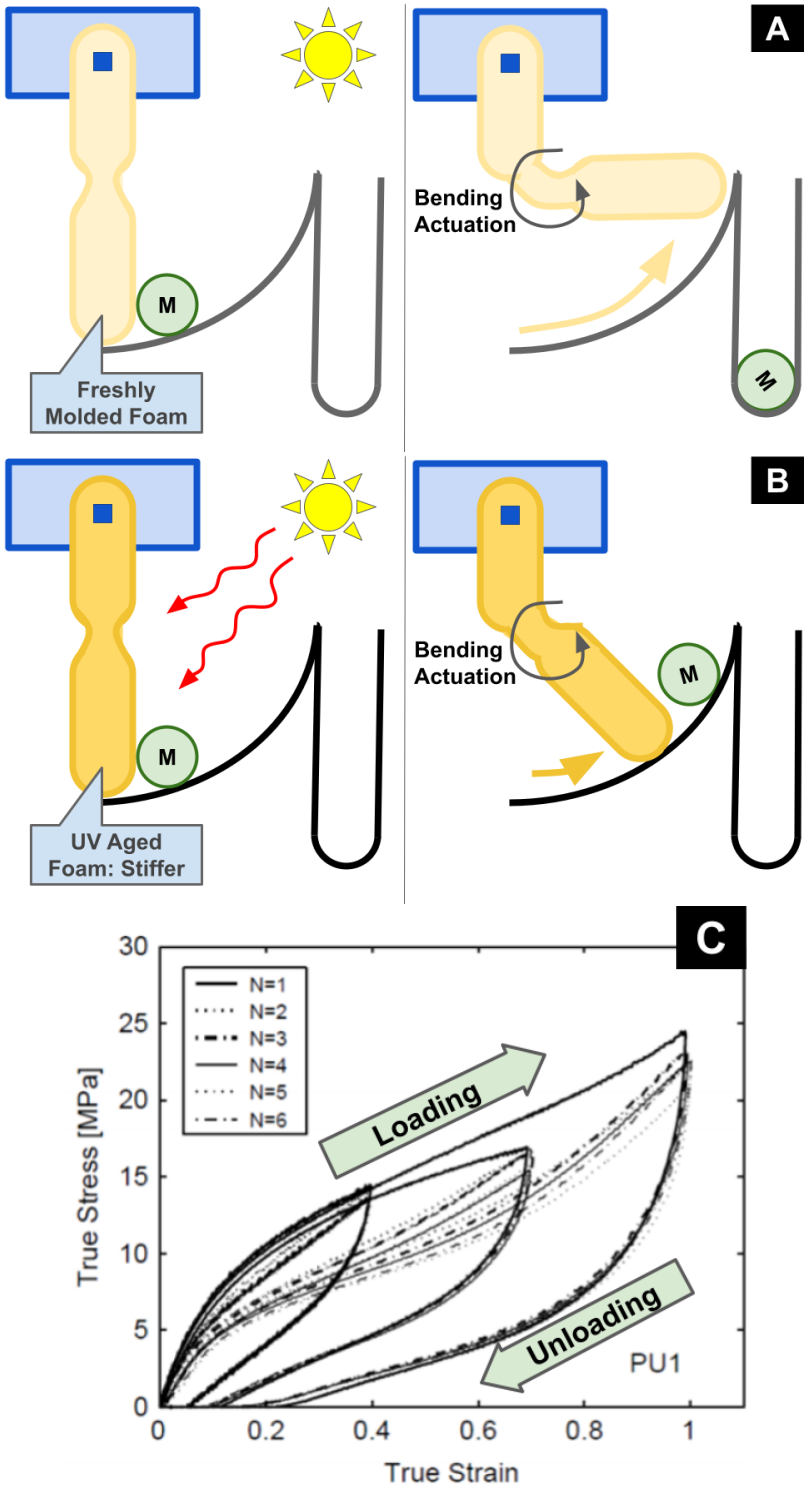

FIGURE 3.2: **A)** Motion of a 1-DOF soft foam manipulator designed to perform a simple open-loop pushing task. **B)** Eventual task failure due to material aging from sun exposure leading to dynamics drift. **C)** Intrinsic time-dependent deformation response of thermoplastic urethane to cyclical loading.

Mathematical Formulation

Learning for nonstationary dynamics can be seen as an extension of the canonical robot model learning problem. For a system with stationary dynamics, there is the assumption that all experience tuples (x_t, u_t, x_{t+1}) are drawn from the same distribution subject to the constraint of the true system dynamics $x_{t+1} = F(x_t, u_t)$, where x_t represents the state of the system at time t and u_t is the control effort at the same time. Since all data provides meaningful, if noisy, information about the true dynamics of the system, all tuples can be used to refine the fidelity of the agent's dynamics estimate. In the nonstationary case, this assumption is violated.

For the nonstationarity case, the ground truth system dynamics vary as a function of time. While this variance can be in the form of the dynamics as well as in the values of dynamics parameters, we assume that there exists some parameterization of the variation such that we can write the dynamics as $x_{t+1} = F(x_t, u_t, \psi_t)$, where ψ_t is the parameter vector determining the dynamics at time t.

The behavior of the vector ψ_t is generally unconstrained and therefore requires domain-specific knowledge to create a reasonable set of assumptions.

Three cases for how the dynamics vector ψ_t might evolve:

- Trends in the dynamics, e.g., material degradation, polymer creep, or thermal buildup. These are gradual, non-periodic changes.
- Random or event-driven changes to dynamics, e.g., mechanism damage, system repair, parts replacement, or power cycling. These changes can be unknown or known and are discrete jumps in the system dynamics.
- Cyclical or oscillating dynamics, e.g., due to environmental changes corresponding with day-night cycles. These are finite periodic changes

It is important that the system designed in this chapter enables the exploration of as many of these cases as possible in the context of soft materials. There are also practical considerations of ensuring the system as a whole is able to operate for long periods of time, and ideally, the system's behavior should be dominated by the dynamics of soft materials.

3.2 Motor-driven Soft Parallel Mechanism

I use reliable electric servomotors as the actuators in my design to create a system that can operate for long periods during data collection for machine learning or other long-term experiments. Many of the moving and interacting parts of the robotic system should be made of soft materials so that the properties of those soft materials can be encoded in data of robot motion. Therefore, I employ the electric motors to actuate a soft matter parallel mechanism, a relatively under-explored strategy that has been shown to enable robust, versatile experimental platforms [66]. The usage of soft materials was focused on the mobile mechanism rather than the whole system because the system components that are not the focus of the study should provide as little disturbance to the system's behavior over time. While soft actuators are diverse, enable many unique properties for robotic systems,

and provide a path to entirely soft systems, they have relatively complex dynamics. Additionally, the kinds of materials used in their manufacture are not as diverse as for passive mechanisms, and many classes of soft actuators require specialized development and operational environments to maintain [96]. Electric motors are especially robust forms of actuation with minimally impactful actuator dynamics. Reliability is further increased when these motors are equipped with internal closed-loop controllers. I base my designs on a planar parallel mechanism, the compliant five-bar.

The choice of a five-bar mechanism as the soft structure the robot uses for interaction is helpful partly because it can be represented in many ways. While I do not explore this path, it is feasible that the geometry of the compliant five-bar could be modeled by the discrete elastic rod simulation environment used on the swimming robot frog discussed in Chapter 2. This simulation strategy is one of the fastest found for soft materials and provides one feature space that could be useful for characterizing our system. Additionally, the approximation of the compliant five-bar via the kinematics of a rigid-body five-bar 3.3A provides an avenue for comparing the effects of the soft materials against a close geometric analog. The difference in fidelity can be intuitively understood by examining the fact that different joint thicknesses produce different behavior of a parallel structure due to changes to spring stiffness (Fig 3.3). However, all joint thicknesses are represented with the same rigid approximation. Parameter estimation of the associated rigid approximation might provide useful data for adaptation. I explore a toy problem of simulated nonstationarity based on changes to the geometry of a rigid-body five-bar in Appendix A.

To explore long-term robotic operation, the five-bar mechanism was used to create an experimental setup that can theoretically perform a task autonomously. I present two such experimental setups.

The first candidate platform I studied for long-term hardware experiments with soft material robotic systems was a five-bar pendulum. I selected this design due to its simplicity of fabrication and its limited sources of uncontrolled nonstationarity. The servos (Dynamixel XC430-W150-T, $120) can control closed-loop position. The angular position of the servos can be directly read from the servos' integrated controllers. Control is run on a desktop computer, and communication is handled by a USB communication controller (Dynamixel U2D2). There is a linear actuator that allows for prescribed changes to the geometry of the system. The five-bar and the pendulum link are made of 3D-printed thermoplastic urethane (TPU), a relatively stiff synthetic elastomer. They are joined by a bolt through the inner diameter of a ball bearing. The state of the five-bar pendulum can be described by the following variables: $\{\theta_1, \theta_2, \theta_3, \dot{\theta_3}\}$. The value of θ_3 can be calculated from $\text{atan2}(y_2 - y_1, x_2 - x_1)$. The points $(x_1, y_1), (x_2, y_2)$ can be extracted directly from video using computer vision as shown in Figure 3.4B.

Pendulum Swing-up

Pendulum swing-up with limited control authority is a classic control problem that cannot be solved via simple linear control. Traditionally, it is composed of a single link pendulum, which is driven by either a motor at its base or the linear motion to a cart on which it is mounted. The task can be broken down into two phases. The first phase is to use limited control authority to add energy to the pendulum to move it from the stable $\theta_3 = \pi$

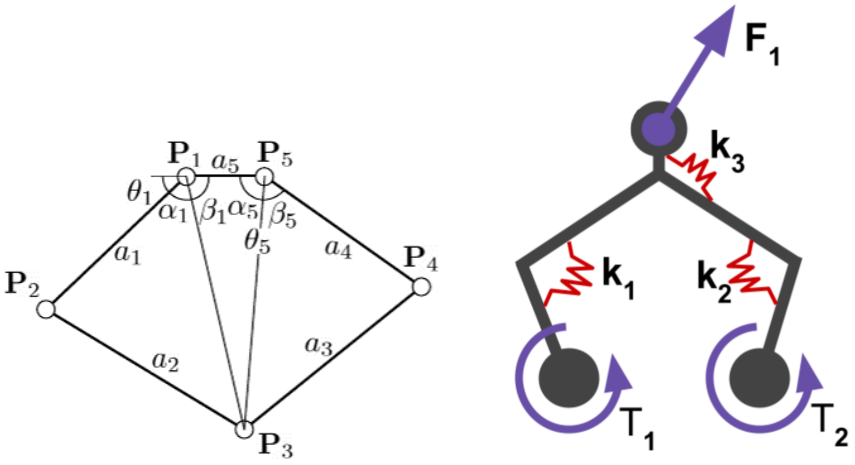

FIGURE 3.3: Left) Rigid body kinematic approximation of five-bar. Right) Approximate forces of soft material five-bar driven by two servos. T1 and T2 are the actuation space of the system, the torques applied by the servos. Springs k_1, k_2, and k_3 are the spring constants of living hinge joints of the five-bar, dependent on hinge material and geometry. The force F_1 is the reaction force that occurs when the tip of the five-bar contacts another object. The spring constants of the five-bar links are not shown but do influence behavior.

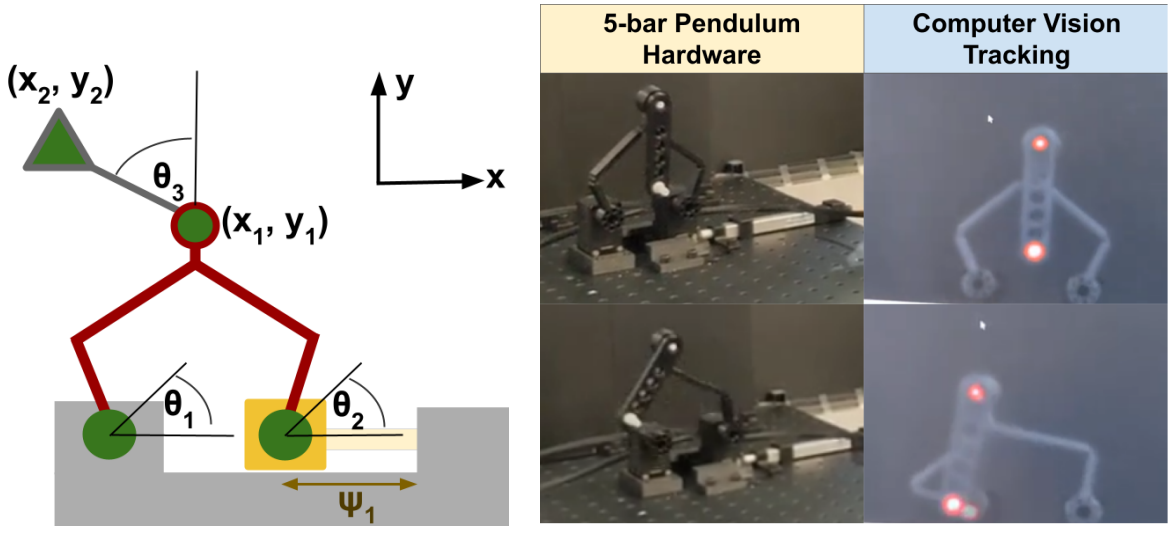

FIGURE 3.4: Left) Kinematics of five-bar parallel mechanism with pendulum. Right) Five-bar pendulum hardware with computer vision tracking of (x_1, y_1) and (x_2, y_2). Video of operation available at [97].

position to the unstable inverted $\theta_3 = 0$ position. The second phase is to exert control authority to keep the pendulum in that position.

Suppose we treat θ_3 as the variable of interest in a pendulum swing-up problem. In that case, our five-bar pendulum can be seen as directly analogous to other versions of the pendulum swing-up problem. One main distinguishing factor is that our actuation is the 2-dimension torque vector of the five-bar servos, which is higher dimensional than other versions of the problem. In addition, actuation is constrained by the interior dynamics of the five-bar itself.

The goal would not be to derive a single static policy for solving pendulum swing-up but rather to derive a method for repeatedly producing such policies as the dynamics of the system drift. The intrinsic nonstationarity inherent to deformable materials is not directly observable, so I introduce two other sources of nonstationarity, ψ_1, ψ_2, which I control and which should dominate the uncontrolled nonstationarity both in magnitude and time scale.

The first source of controlled nonstationarity is the motion of the linear actuator ψ_1, which one of the five-bar servos are mounted to, as shown in Figure 3.4A. This actuator changes the link length of the ground link of the five-bar, which directly affects the global motion of the five-bar. The benefit of this source is that it can be controlled programmatically during training, meaning that different trajectories through the space of the dynamics can be explored automatically.

The second source is replacing the five-bar with a new structure, which could have different link lengths and living hinge thicknesses. The effect of these changes is a modification to both the kinematics [98] and the force equilibrium of the five-bar, as shown in Figure 3.3. The 3D-printed fabrication of the five-bar pendulum lends itself to easily changing the pendulum design. The link-length and joint thickness changes can be encoded in an implementation design vector ψ_2. This vector is a candidate for encoding domain knowledge and making it available to the model so we can address the case where we know what changes to the system have happened but not necessarily their effect on the dynamics.

This system can operate at relatively fast actuation speeds for a soft robot, and the computer vision-based sensing captures its state quite robustly [97]. These initial experiments were not pursued further, however, as it became clear that while this system was an interesting platform for studying nonstationary dynamics, it had some limitations in enabling the control of systems that rely on accessible soft materials.

3.2.1 Limitations

The dominant dynamics of this system are the inertial dynamics of the pendulum. While this non-quasistatic system would enable exploration of solutions of behaviors on shorter time-scales, the impact of deformation of the relatively stiff TPU is not especially pronounced. Furthermore, the changing link length as a source of controlled nonstationarity, while helpful for automating experimentation, does not represent the kinds of dynamics changes expected for soft material systems in general. While this platform possesses some utility to study policy generation under nonstationary dynamics, it is missing the crucial ability to provide insight into the long term behavior of soft materials in robotic systems in general.

3.3 Long-Term Learning with Modular Parallel Manipulator

***K. Chin**, C. Majidi and A. Gupta. "1 Modular Parallel Manipulator for Long-Term Soft Robotic Data Collection." arXiv preprint arXiv:2409.03614, 2024.* [99]

Author contributions: I designed and built parallel manipulator modules, programmed reinforcement learning environment and performed long-term learning experiments.

To better understand the dynamics of soft materials in a long-term experimental environment, I developed a robotic experimental platform that allows automatic performance of a class of simple manipulation tasks using mechanisms made from soft materials. The platform is relatively inexpensive to fabricate and can operate for long periods. A notable improvement from the five-bar pendulum is that the behavior is highly coupled to the material properties of the soft materials used in its construction, allowing insights into those properties. I demonstrate the ability to train reinforcement learning policies using only hardware data collected with this system.

3.3.1 Compliant Five-bar Modules

The modules are actuated by two servomotors with aligned axes of rotation mounted next to each other in 3D printed TPU housing, assembled with friction fit 3.5. The module design is compatible with lower cost servomotor (Dynamixel XC430-W150-T, $120), and higher performance servomotor (Dynamixel XM430-W210, $270). The bases of the modules are attached to a mechanical breadboard with bolts. These mechanical design choices improve the ability of the system to operate for more extended periods without requiring repair, as well as simplifying repair 3.6. The primary moving parts of the modules are compliant five-bar mechanisms that mount to the servomotors. The angle of the driving servos nominally determines the state of an individual module (though the decoupling of nominal state and actual state due to soft material compliance remains). There are no autonomously controllable features for inducing nonstationary dynamics in this system. Instead, the soft components to are the primary source of trends over time in the dynamics. The dynamics can be controlled by swapping out five-bars, which slot onto 3D-printed quick-swap servo horns.

3.3.2 Manipulation Task: Knob Turning

This task is adapted from the ROBEL robotics learning benchmark designed by Ahn *et al.* [100], replacing the rigid-body robot they use with my soft five-bar modules. A knob-like object is in the center of the workspace, mounted to a servomotor. The servomotor enables encoder-based state estimation of the knob's pose and allows the autonomous resetting of the system to a nominal state. The ability to automatically reset the environment is a critical feature for running machine learning experiments autonomously. This need is also a significant motivator of the simple planar geometry chosen. Soft five-bar modules are arranged surrounding the knob and are tasked with turning the knob to a desired pose. This general task outline can be modified by changing the location of modules, the geometry of five-bars, or the pose desired, either static goal poses or time-series pose trajectory. The use of a servo to mount the knob also allows adjustment of the knob's stiffness (as long as the servo used has a torque-control mode).

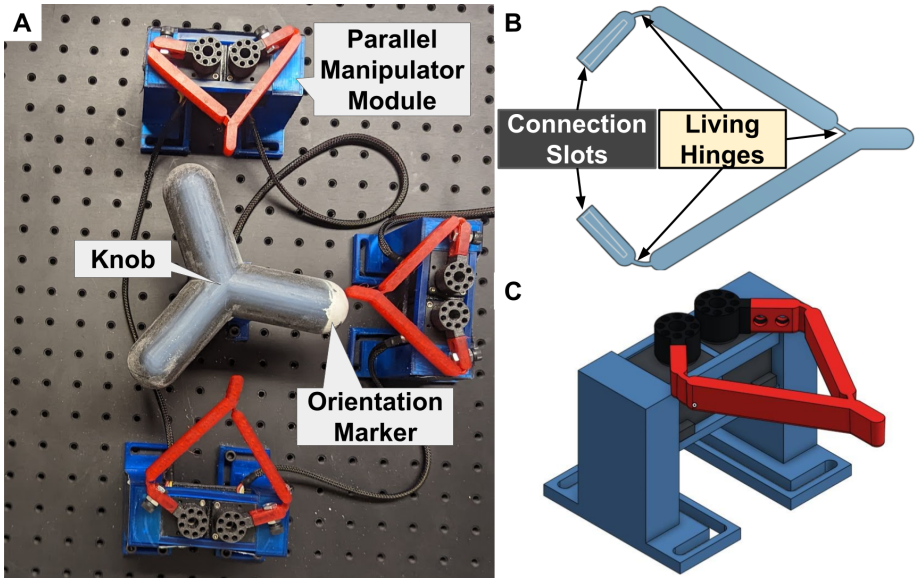

FIGURE 3.5: A) Full experimental manipulator platform B) Soft five-bar mechanism design with quick-swap connection slots. C) Individual module design.

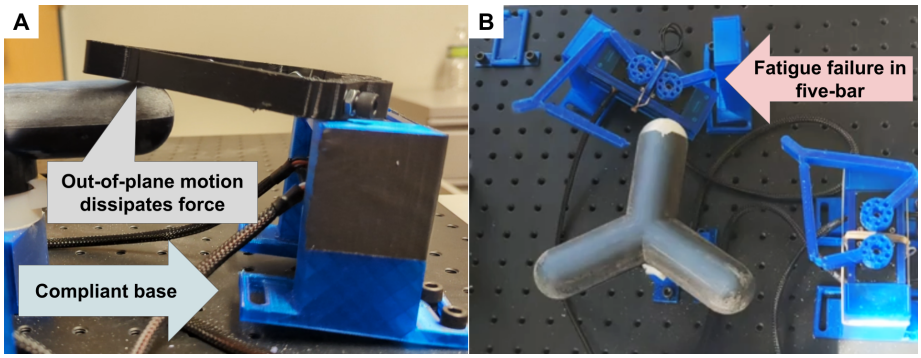

FIGURE 3.6: A) By mounting the compliant base of the module to mechanical breadboard using a linear arrangement of bolts, the module is able to rock back, pushing the finger out-of-plane in response to overly high forces. This relief motion reduces the chance of breaking the system during extended unmonitored data collection. B) The most likely failure modes are module disassembly to friction fit module components and soft five-bar failure, especially at living hinges. Even if both occur undetected and the system keeps trying to operate, the compliant materials used and construction lower the chance of long-term damage to the experimental setup.

I again leverage dimensionality reduction via action space discretization to maximize the learning speed on hardware. Unlike the shape-memory alloy actuators used in Chapter 2, there is a relatively natural choice of discrete actions for electric motors in a quasistatic context: increasing or decreasing the angle of the motor by some static amount. These primitives span the space of available states for a motor. The generated space is a superset of the valid states of a module, but by applying geometry-dependent limits to servo position, these primitives can explore the full nominal state of each module. The full manipulator array would then have an action space of

$$a_t \in \{mag * \Delta\theta(\text{servo}_i)\}, \quad \text{where} \quad i \in \{0, \ldots, 6\} \quad \text{and} \quad mag \in \{0, \pm\text{angular resolution}\}$$

However, this low level of abstraction poses some challenges to the speed of training. For studying the ability to bootstrap autonomy, low-level action spaces are often chosen to minimize learned bias and reliance on human design decisions. However, to make robotic systems development more accessible, there are different sets of considerations. While this system is designed to allow more hardware data to be collected than other soft robots explored earlier in this thesis, there is still a time cost of collecting that data. While the easily replaceable modules and resilient mechanical design minimize damage to the most expensive hardware components, I still aim to minimize the amount of time necessary to run hardware to collect data. Therefore, the action space is chosen to maximize the data efficiency of the learning process, incorporating as much human intuition as possible. First, an intermediate action space that allows the tip position of each module to be controlled in a "extend", "retract", "left", "right" scheme is developed. Then, I build the actual primitives upon this as trajectories of intermediate positions that produce a distinct sweeping motion over several time steps. Specifically, the motion is a sequence of "move left" $\rightarrow$ "extend" $\rightarrow$ "move right" $\rightarrow$ "retract", shown in Figure 3.7. The resulting sweeping motion and its mirror can be performed by each module, resulting in an action space of:

$$a_t \in \text{sweep}(\text{module}_i, \text{dir}), \quad \text{where} \quad i \in \{0, 1, 2\}, \quad \text{and} \quad \text{dir} \in \{left, right\}$$

These primitives are hierarchical, built first upon the lowest-level motor delta primitives and intermediate tip-control primitives. The hierarchical architecture leaves space for learning the dynamics at a lower level, in case more optimal and/or less biased behavior is desired later on. These sweeping primitives are very likely to move the knob if the moving module is aligned with a lobe of the knob geometry. This intrinsic effectiveness means that the learning problem is in determining how to sequence the primitives to complete the task, requiring the learning of some latent representation of the relative geometries of the modules and the knob and the way the modules interact when making contact with the knob, but not needing to learn coherent motion.

For the following results, the task is learning to turn the knob to a defined goal position from a defined start position.

3.3.3 Reinforcement Learning on Hardware

Reinforcement learning was used to learn this task from empirical hardware data. I used policy gradient methods from the Stable Baselines3 repository of standard reinforcement

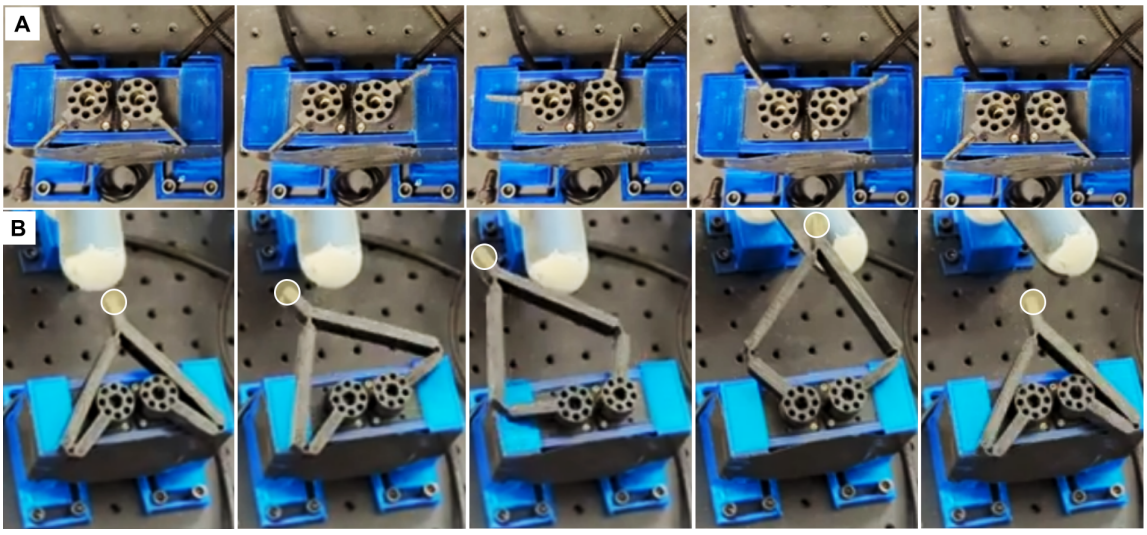

FIGURE 3.7: A) Motion primitive execution without attached finger, showing servo horn finger connectors moving. B) Execution of that primitive with a urethane finger mounted to module.

learning algorithms [101], specifically Proximal Policy Optimization (PPO) with a batch size of 64 and a learning rate of $1e^{-3}$. I wrapped the hardware controller for the modules in a Gym environment. The repository for the reinforcement learning environment can be found at [102].

The action space of the reinforcement learning problem was the sweeping primitives derived in the previous section. The state space is the servo angles of each module, normalized between -1 and 1. Appended to this is the knob pose, represented as a two-element vector of the cosine of the knob angle and the sine of the knob angle, for better numerical performance in the policy network. The reward used for reinforcement learning was

$$\text{Reward} = -5 \times |\Delta\theta_{\text{knob}}| + 10 \times \mathbb{1}\left\{|\Delta\theta_{\text{knob}}| \geq \frac{0.25}{\pi}\right\} + 50 \times \mathbb{1}\left\{|\Delta\theta_{\text{knob}}| \geq \frac{0.10}{\pi}\right\},$$

corresponds to a smoothly increasing cost of 5 times the angle of the knob in radians, penalizing being farther away from the goal position of 0 radians. There is a big reward jump of 10 if within $\frac{0.25}{\pi}$ radians and an even bigger reward of 50 if within $\frac{0.10}{\pi}$ radians of the goal position.

The reinforcement learning experiments were run over time, with experiment length following a powers of two progression. The first experiment was $2^{14} = 16,284$ time-steps and progressed up to $2^{16} = 65,536$ time steps. The next step up, $2^{17} = 131,072$ time-steps, was also run, but it was partway into this that the system finally experienced a hardware failure in the form of one soft finger failing at a hinge and pushing itself out of its housing, like the failure mode shown in Figure 3.6B. The repair process for such a failure was rather quick, requiring only the attachment of a new finger with bolts to the quick-swap servo horn and the fastener-free reassembly of the module.

The length of a time step is the amount of time it takes to execute a primitive, which

is about 5 seconds. Therefore, the longest completed continuous batch of data collection, the 65k batch, ran for 3.76 days straight. In all, counting from the time the system was activated, all the completed batches, as well as the 10,714 time-steps completed in the 131k batch before a hardware failure, the system was active for $2^{14} + 2^{15} + 2^{16} + 10,714 = 125,402$ time-steps, or 7.25 days. There were automated cooldown periods between batches, totaling two 2 hours of downtime, meaning a total of 7.34 days from the beginning of the data collection. Therefore, the system operated for over a week at 98.7% uptime before requiring repair.

The system improved task performance over time, as shown in Figure 3.8. Based on the reward function, achieving a mean reward of 10 was significant as it requires quickly moving the system close to the goal without spending many time steps. The episode length of 4 is significant as that was almost the fastest possible execution (shown in Figure 3.9. Episode lengths of 3 were observed, but they involved hard-to-replicate interactions, where due to the angle of contact, elastic energy was stored in the five-bar before being released and pushing the knob further than the quasistatic sweeping motion.

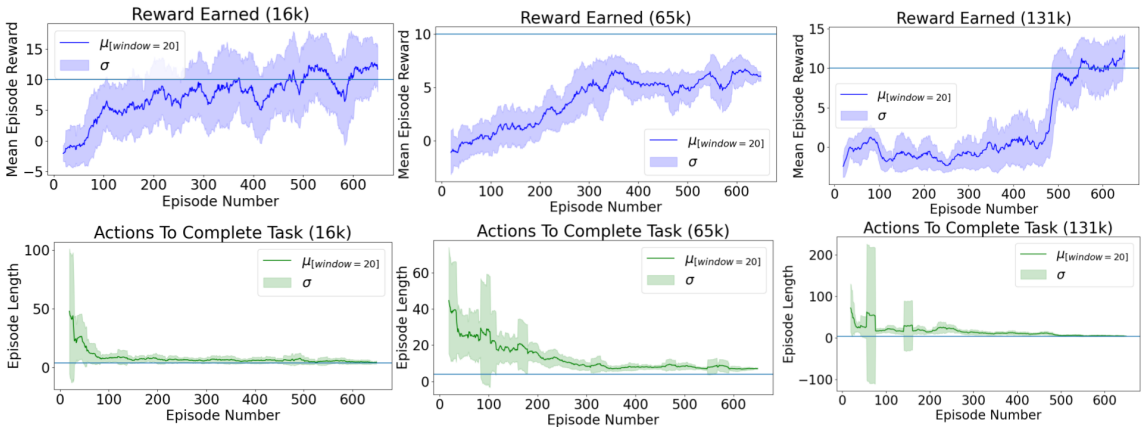

FIGURE 3.8: Policy improvement for different training sessions. These traces are extracted from the first 650 episodes of each training session. The reward value of 10 and episode length of 4 are marked off as they roughly correspond to the best performance found throughout experimentation. The full traces are shown in Figure B.2

Once the ability to collect enough hardware data to perform reinforcement learning was demonstrated, I evaluated the platform's ability to provide insight into the effects of different soft materials on the dynamics.

3.3.4 Policy Transfer Between Soft Materials

Previous results were obtained using 3D-printed thermoplastic urethane five-bars. I created new compliant five-bars from injection-molded silicone (Ecoflex 00-30, Smooth-On), a much softer material. These were manufactured with a desktop injection molding setup powered by compressed air. I trained the system to perform the knob-turning task using the above reinforcement learning system. In Figure 3.10, we can see the execution of this policy on hardware.

I then characterized the performance of the policy learned on the silicone five-bars when executed on silicone five-bars and TPU five-bars. I also characterized the policy

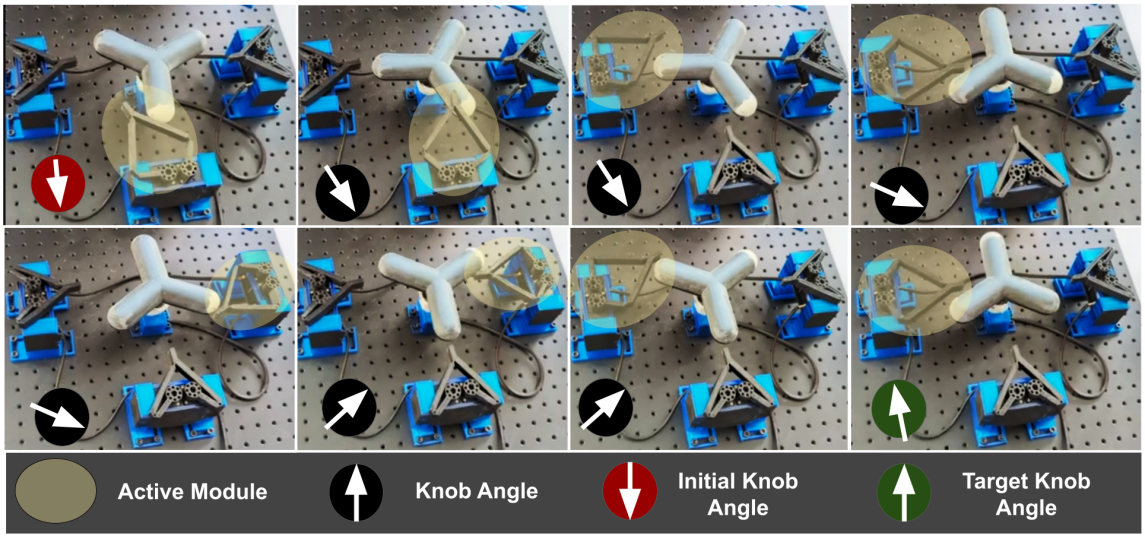

FIGURE 3.9: Successful policy execution leading to completion knob turning task using urethane five-bars. Policy learned via policy gradient reinforcement learning over action primitives.

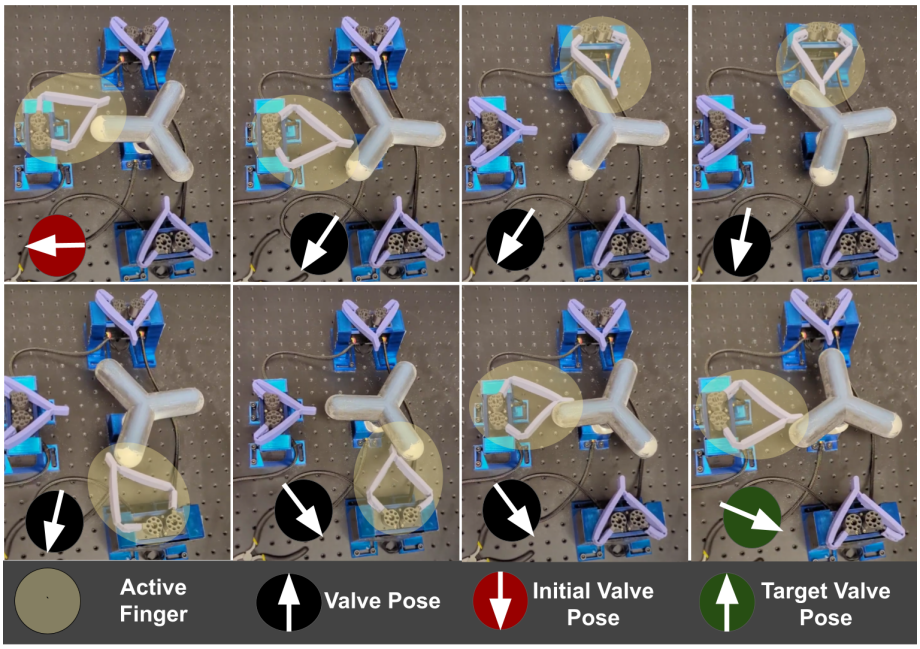

FIGURE 3.10: Successful policy transfer of knob turning task from urethane five-bars to silicone five-bars.

learned on the TPU five-bars to both systems. The results of these experiments are shown in Figure 3.11A.

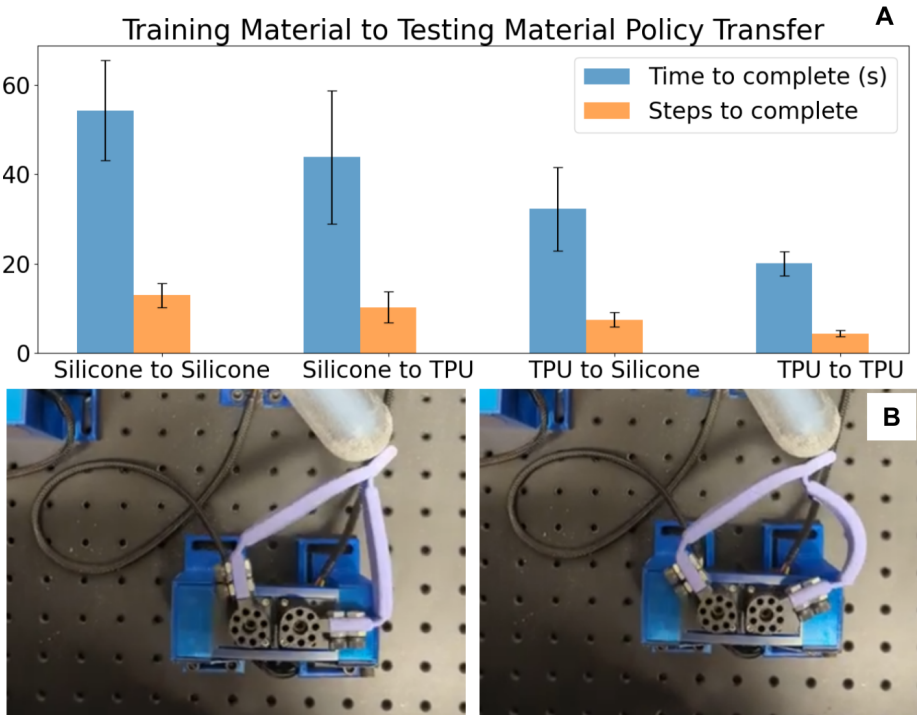

FIGURE 3.11: A) Performance comparison of transferred policies, trained on different soft materials. Each case was tested ten times. B) Significant bending of links in softer silicone indicates greater deviation from nominal kinematics.

The best-performing case was the TPU-trained policy that was implemented on TPU. It might be expected that the dominant factor in transfer performance is the similarity between system dynamics during training and system dynamics during execution. However, the next best-performing case is a TPU-trained policy executed on silicone, suggesting that the material properties of TPU are more efficient for learning. The softness of silicone means that the deformation of the five-bar links outside of the living hinges is more significant, as shown in Figure 3.11B. This deformation provides more possible physical configurations for a given motor position, effectively increasing the noisiness of state transitions and lengthening the learning process. The worst performing case is a silicone-trained policy on a silicone five-bar. The fact that a policy trained on silicone performs better on TPU indicates that the stiffness of TPU is a boon for executing actions reliably.

Simple experiments like this allow us to extract insights about the suitability of different materials for this task, supporting the utility of this kind of research platform for studying the effect of soft material dynamics on robotic system operation.

3.4 Discussion

In this chapter, I designed and implemented a platform that enables long-term data collection with parallel mechanisms composed of various soft materials. I show that reinforcement learning can generate control policies from purely the hardware data generated by this platform. By executing policies trained on one material to mechanisms made of different materials, latent information about material properties can be implicitly encoded via differences in performance.

Characterizing the effectiveness of policy transfer between different five-bars presents the opportunity to establish a similarity metric in task execution. By establishing the level of similarity between different five-bars, if the nominal geometry is kept constant, this platform allows for establishing a holistic similarity metric between materials.

I also attempted to transfer policies to five-bars with different geometries, but was not successful. The lowered effectiveness of transfer to different geometries validates the benefit of shared nominal geometry. Establishing similarity could be achieved by executing policies trained on known, well-characterized materials with five-bars constructed from novel materials or materials with unmodeled or unsensed nonstationary material dynamics. Establishing a network of multiple materials and their similarities to each other could provide a path to learning the latent representation of material properties.

In the rest of this thesis, I delve deeper into identifying and creating materials that are accessible to fabricate. Such material systems could benefit from the implicit material characterization enabled by this modular parallel manipulator platform. In Chapter 4, I introduce biomaterial-based hydrogels sourced from accessible ingredients. These materials exhibit changes to morphology and compliance over time, which would induce nonstationary behavior in systems fabricated with them. Additionally, the fabrication of these materials is a versatile process. They can be constructed from ingredients from different sources and processed in several straightforward ways to modify their properties. The option for general fabrication processes to be followed without high-precision measurements and exact ingredient sourcing would improve the ability of people in different contexts to access these materials. Additionally, it would be preferable for people to be able to deliberately modify the materials to fit their needs. These factors would generate many different, related materials depending on the composition and process used to create them. Empirical data that encodes properties of different materials over time would be beneficial to enable the use of the resulting matrix of materials. The parallel manipulator platform presented in this chapter enables the collection of such long-term data. It exemplifies a path for robotic systems built using these materials to be modeled and controlled.

Chapter 4

Gelatin-based Biomaterials

Broader access to creating robots would lead to the creation more robots. This proliferation of robotics technology is desirable from a socioeconomic perspective, developing a world where the labor-saving benefits of robotics are democratized rather than concentrated and hoarded. However, the potential ecological impacts of such proliferation must be considered. In the face of growing climate instability, pursuing more sustainable material choices is a significant concern [24, 27, 103]. Natural materials derived from renewable biological sources have shown the potential to replace materials extracted from petroleum byproducts or other non-renewable sources for use in next-generation soft machines and electronics [104–107]. Thus, I am motivated to start working with these materials now and to prefigure the ecologically sustainable, equitable world after this transition.

In this chapter, I synthesize biomaterial composites based on extremely accessible, inexpensive ingredients derived from agricultural waste and food by-products (gelatin, glucose, glycerol). These ingredients are inexpensive due to their ready availability as shown in Figure 4.1, either derived from animal residues in agricultural waste (gelatin), extracted from plant or animal fats (glycerol), or refined from a wide variety of crops (glucose syrup). Not only are their sources renewable, these ingredients are byproducts and the bulk of their production is driven by co-products with much larger markets: meat or other animal products for gelatin, biofuels for glycerol and subsidized corn and sugar for glucose syrup. Biodiesel production produces a marked excess of glycerol, posing environmental challenges for disposal [108], and the excess of high-fructose corn syrup (glucose syrup) resulting from subsidies has been identified as a driver of public health decline in the United State [109]. By introducing value-added uses for these materials and providing an outlet for existing surplus, the overall sustainability calculus of their sources and co-products can be improved. I focus on the creation of hydrogels, elastic composites of a gelatin polymer matrix swollen with water and biodegradable additives. I apply protein coagulation to produce longer gelatin polymer chains, providing an accessible process for modifying material properties. I demonstrate various phenomena of hydrogel behavior relevant to their use in soft robotics.

4.1 Related Work

A versatile approach to creating material composites from biomaterials is by using biologically derived polymers (biopolymers) in a hydrogel, a composite material composed of a polymer matrix swollen with an interstitial aqueous phase, primarily water. Biopolymer hydrogels are an active and relatively mature area of research outside the field of robotics. In particular, fields where the biocompatibility of a material is of high importance have

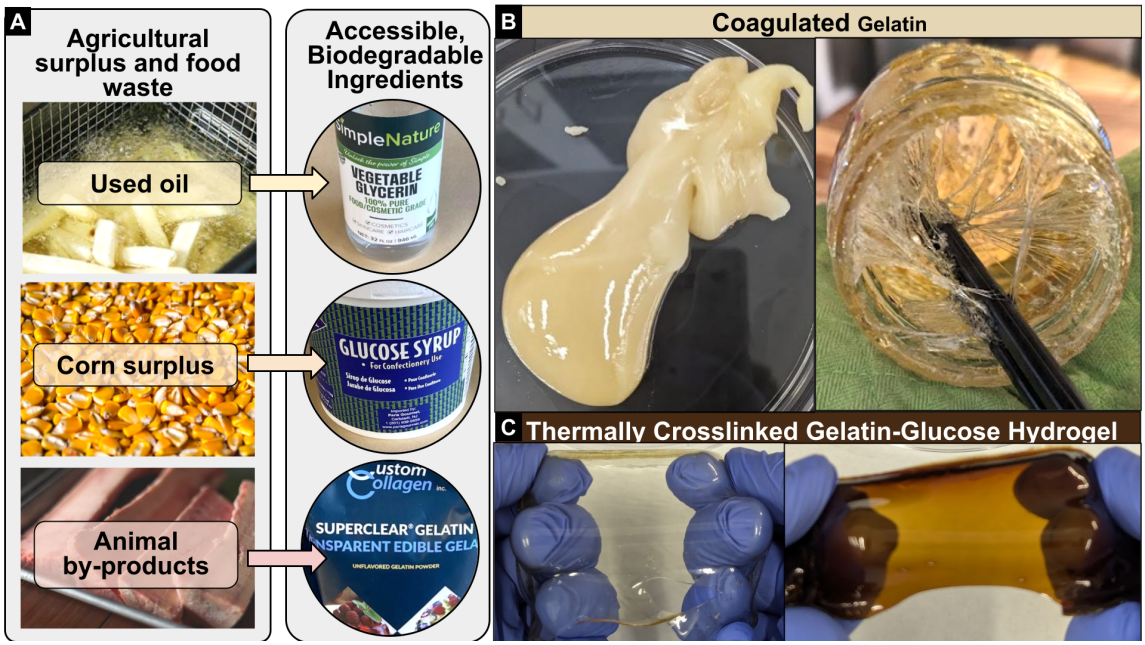

FIGURE 4.1: A) Demonstration of factors enabling the accessibility of biomaterial ingredients for gelatin hydrogel. Images adapted with permission (by Ron Lach, José Roberto Oliveira, and Los Muertos Crew respectively, via Pexels) B) Gelatin coagulation to increase polymer chain length is an accessible way of modifying the polymer characteristics of gelatin. C) Thermal crosslinking with sugar improves the hydrolytic stability of gelatin.

been quite focused on these materials. Food science and biomedical engineering research provides a wealth of inspiration for biomaterial options. For use in accessibly fabricated robotic systems, my materials need to be compatible with being handled in low-resource environments, with domestic environments being the most constraining. This is a relaxation of difficulty compared to the challenges of human biocompatibility faced by other fields. This relaxation allows the exploration of techniques and compositions that differ from existing literature.

Many biomaterials have been used in the creation of hydrogels, including cellulose [110], alginate [111, 112], lignin [113], agar [114] and gelatin [115]. Gelatin is a promising option since it can be designed to exhibit mechanical properties similar to synthetically derived soft elastomers and gels while also being sourced from agricultural waste. It has been used in tissue scaffolds, anatomical models and phantoms, and as drug delivery mechanism, along with being used in cosmetics [116], as food packaging and food themselves [117]. Given these potential benefits, work is underway to incorporate gelatin into emerging technological applications [118]. Gelatin-based hydrogels have been of interest for creating soft structures structures and devices [119].

The accessibility, biodegradability, and mechanical tunability of gelatin have been established by Baumgartner *et al.*, with an additional demonstration of integration into soft robotics [115]. 3D printability of optically clear gelatin for use in waveguides for sensing and as the bulk structure of a soft mechanism was demonstrated by Heiden *et al.* [120]. A functional soft robotic pneumatic actuator using a gelatin-glycerol hydrogel structure has been demonstrated by Shintake *et al.* [121]. Matonis *et al.* fabricated fully edible

origami-inspired actuators from insoluble genipin-crosslinked gelatin [122]. Geckeler *et al.* demonstrated the ability to temporarily deploy sensor packages to tree branches via an origami gripper passively actuated by a pre-stretched, moisture-degraded gelatin hydrogel [123].

Gelatin is extracted by denaturing the triple helices of the collagen protein in animal tissues, especially bones, skin, and connective tissues. The chain of a gelatin polymer is composed of a semi-regular sequence of amino acids, especially glycine, proline, and hydroxyproline, which can form helices, and associate into similar triple helix structures as collagen [124]. The mechanical properties of collagen are evident in its use as one of the main structural components of animal life. Yet, the gelatin derived from this collagen is rarely mechanically impressive when used in a hydrogel in its unmodified, non-composite form, which can be seen in Figure 4.2D. The mechanical performance of hydrogels, in general, has been an issue for their adoption into soft robotics [125]. Improving the mechanical properties of hydrogels can be done in a few ways. One method is to use a double network gel, combining a brittle and densely crosslinked network with an interpenetrating a ductile network with few crosslinks. When stretched, the bonds of the dense network break sacrificially, dissipating energy and improving toughness. This stiffness and toughness comes at the cost of elasticity and repeatability, as those bonds generally do not recover. Gelatin has been used as the energy dissipating network in double network gels with polyacrylamide, and due to thermoreversible sol-gel transition, these hydrogels can be heat treated to initiate self-healing and partial recovery of those bonds [126]. Sun *et al.* showed that if the crosslinks in the dense network are reversible, e.g., via ionic bonding of alginate and calcium, a repeatably elastic and tough hydrogel can be formed, which does not require manual treatment to recover [111]. There is evidence that the hydrogen bonding of gelatin can serve as reversible sacrificial bonds in a network with polyacrylamide and PEDOT:PSS polymers [127]. The work of Baumgartner *et al.* shows that mechanical properties can be improved by the addition of other biomaterials to gelatin hydrogel, namely glucose as a cosolvent and glycerol as a cosolvent and plasticizer [115]. The work of Shimizu *et al.* on gelatin gels with sugar and glycerol proposes exclusion of cosolvents from the surface of biopolymer chains as mechanism for the mechanical stabilization of gel networks by addition of these molecules. Such exclusion would induce gelatin rich regions that promote helix formation and enable high density of reversible hydrogen bonds between these helices [114]. In addition to double network approaches, mechanical properties have been improved by creating "tanglemers", networks where entanglements in the polymer chains greatly outnumber crosslinks between chains . When a highly-entangled polymer or gel is stretched, entanglements allow forces to be transmitted between chains that are not crosslinked, and the location of the entanglements can slip, adding stiffness without making the polymer brittle. When a chain breaks, it dissipates energy stored in that chain, and the many chains that were stretched by being tangled with it, greatly improving the toughness of the gel, while maintaining elasticity [128]. Norioka *et al.* showed that controlling the polymerization environment is a powerful and reliable method of achieving such networks within hydrogels. Specifically, high concentrations of monomer and low concentrations of crosslinker produce longer chains between crosslinks, allowing entanglements to outnumber crosslinks [129]. There have been polyacrylamide tanglemers created with degradable disulfide bonds [130]. However, to my knowledge, there has not been work in creating a tanglemer from only biomaterial ingredients.

The critical step in creating a highly-entangled hydrogels is the acquisition of polymer chains with high molecular weights / long lengths. The fiber spinning literature provides insight into how this can be achieved for gelatin. The molecular mass of gelatin has been increased via coagulation in dimethyl sulfoxide (DMSO), a common solvent for synthetic soft material processing, and tetrafluoroethylene (TFE), the monomer used to make the non-stick coating PTFE [131]. These ingredients are effective coagulants but inaccessible due to toxicity and carcinogenic potential [132, 133]. More promisingly, chain length extension via accessible, safe coagulation in both ethanol and isopropyl alcohol has been demonstrated by Stoessel *et al.* [31, 134]. While most applications of coagulated gelatin are for fiber production, a few other form factors have been demonstrated. Wires that self-wind like plant tendrils have been synthesized via a combination of coagulated gelatin with amyloid fibrils [135]. A related process, complex coacervation, is used to generate a slurry of gelatin microparticles from solution, which is used in a support bath for 3D printing of other biomaterial hydrogels [136]. However, use of coagulated gelatin as the matrix material in a highly entangled hydrogel has not been previously explored.

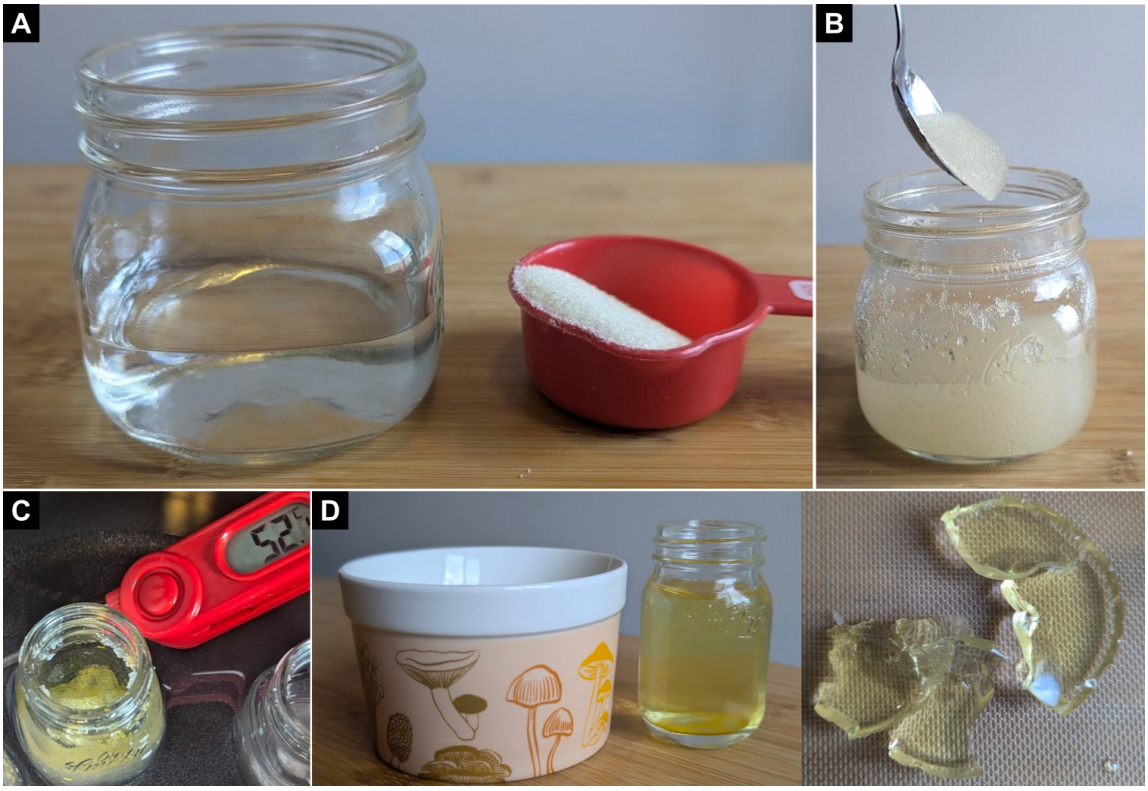

FIGURE 4.2: A) Tap water and pristine gelatin powder. B) Gelatin powder hydrated by slow addition to water (blooming, resulting in a paste). C) Melting gelatin in a stove-top water bath with a commercial probe thermometer for temperature monitoring. D) Ramekin and molten gelatin sol-phase for casting. E) Fresh, swollen, pristine hydrogel without additives

4.2 Coagulated Gelatin for Long-Chain, Highly-Entangled Networks

To create long gelatin chains, I follow the process outlined in the fiber spinning literature by Stoessel *et al.* [30], adapted for kitchen equipment. Briefly, a ternary solution of gelatin powder (10 w.t.%), alcohol (50 w.t.%), and water (40 w.t.%) is placed in water bath at 50°C, with agitation every 5 minutes for 30 minutes until a coagulated phase is formed (Figure 4.3, video available at [137]). This process relies on non-equilibrium dynamics of local phase separation due to the pair-wise interaction energies between polymer, solvent and nonsolvent [138]. Alcohol is a non-solvent of gelatin, so gelatin polymer tends to be excluded from alcohol. Alcohol and water are miscible so they can intermix relatively easily. Water is a solvent of gelatin, so there is energetically favorable interaction between them. While Flory-Huggins solution theory provides some insight into the result of various configurations of pairwise interaction energies [138], the values of relevant parameters are not trivially found. However, even without establishing these parameters, the result of this coagulation process is the slowing of the reaction rate of gelatin polymerization. This produces longer chains, since gelatin polymerizes via condensation polymerization, a process which produces higher molecular weight chains under longer reaction times [139].

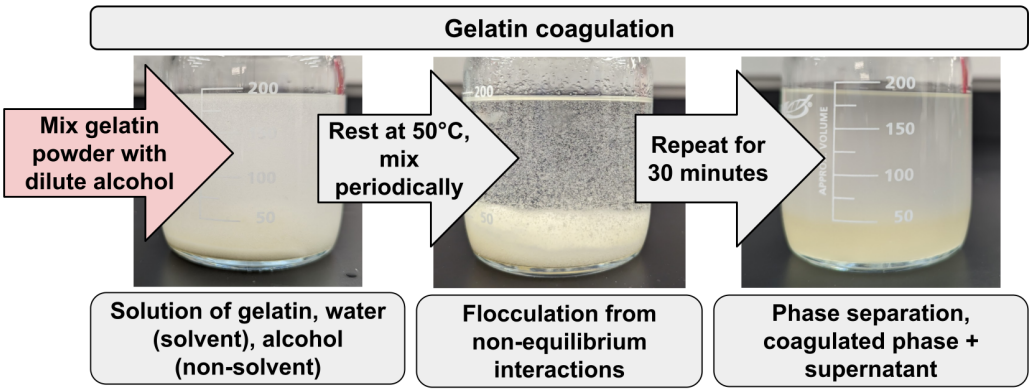

FIGURE 4.3: Process of creating coagulated gelatin based on process in [30, 31] A ternary solution of gelatin powder (10 w.t.%), alcohol (50 w.t.%), and water (40 w.t.%) is placed in water bath at 50°, with agitation every 5 minutes for 30 minutes until a coagulated phase is formed.

The presence of long gelatin chains after the coagulation process becomes evident on the macroscale as the behavior of gelatin changes. The solution begins to contain visible filaments which can be pulled into fibers by hand (Figure 4.4A-B). This coagulation process works best when gelatin and water are not allowed to interact in the absence of alcohol. When dry gelatin powder is exposed to water, it hydrates as shown in Figure 4.2A-B. When this hydrated gelatin is heated above the melting point of gelatin, a solution of polymerized gelatin chains is formed. If alcohol is subsequently added to this solution, long entangled chains are not formed. Instead microparticles of standard length gelatin chains precipitate out, producing a result closer to coacervation [112]. A hydrogel formed from this solution will not be highly entangled, being just as brittle as pristine gelatin as shown in Figure 4.5B, if not more so due to the chains being induced to form distinct

particles during precipitation. If alcohol is instead added to hydrated gelatin prior to melting, the resulting hydrogel appears slightly stronger than if it is added after melting, presumably due to a partial slowing of the rate of polymerization, but the macroscale indications of the presence of long gelatin chains are not observed (Figure 4.5)C. Insight into the differences in the molecular structure of pristine and coagulated gelatin can be found in Figure 4.4C using Fourier transform infrared (FTIR) spectroscopy. There are four major regions of interest, with wavenumbers at 3500 - 2300 cm^{-1} (Amide A), 1656 - 1644 cm^{-1} (Amide I), 1560 - 1335 cm^{-1} (Amide II) and 1240 - 670 cm^{-1} (Amide III). Amide-A represents N-H stretching coupled with hydrogen bonding and free O-H, while Amide I contains C=O stretching vibration with the contribution of C-N bond stretching vibration 1600 and 1700 cm^{-1}. Amide II arises from N-H bending vibration and C-N stretching vibration. Amide III represents vibration in the plane of C-N and N-H groups of bound amide or vibration of CH_2 group. There is a distinct, strong peak around 1550 cm^{-1} in the coagulated gelatin samples, and there are minor peaks from 1170 - 1500 cm^{-1}. Some of the peaks in the coagulated gel are consistent with the presence of retained isopropyl alcohol (IPA). However, the peak at around 1550 cm^{-1} is not present in IPA and is much stronger than in pristine gelatin. Other peaks in pristine gelatin are relatively stronger in coagulated gelatin, providing evidence for increased formation of bonds with energy absorption at those peaks.

The ternary solution used for coagulation can can be created with isopropyl alcohol or ethanol. As I show in figure 4.6, the process works with off-the shelf commercial liquor, as the 75.5% ethanol content is high enough that the non-equilibrium solution is formed when gelatin powder is added. Additionally, I show that extremely pure or technical grade gelatin is not required, as I was able to use gelatin obtained from a local farming community to perform the process. The price of coagulated gelatin hydrogel compares favorably to the synthetic soft materials purchased for use in the compliant five-bars I created in Chapter 3. One pound of thermoplastic urethane filament cost $15.84 [140], while a pound of Ecoflex 00-30 cost $28.69 [141]. Depending on supplier of gelatin and alcohol, the cost of one pound (dry weight) of coagulated gelatin hydrogel with a 1:1:1 ratio of gelatin:glucose:glycerol is between $7.40 and $17.27. The low price and broad availability of ingredients for this gelatin-based biomaterial is encouraging, and supports the vision of being able to create useful materials without relying on long supply chains and centralized suppliers.

4.3 Fabrication Complexities of Gelatin Hydrogels

The first gelatin hydrogels I created beyond the basic recipes used in home kitchens (Figure 4.2) followed the procedure outlined in the work of Baumgartner *et al.* [115]. They characterize many different compositions, I based my initial experiments on G2430, their most stretchable variant (356 $\pm$ 11 % strain-at-break). Briefly, glucose syrup (7g) was heated to 65°C to reduce viscosity, then mixed with glycerol (9g) , citric acid (1g) to prevent microbial growth, and water (4g). This precursor solution was allowed to cool to room temperature, then gelatin powder (6.65g) was slowly added and allowed to bloom for 1 hour, with the particles absorbing the solution and forming a paste like the one shown in (Figure 4.2B). This paste is then melted at 85°C to form a sol phase, which can

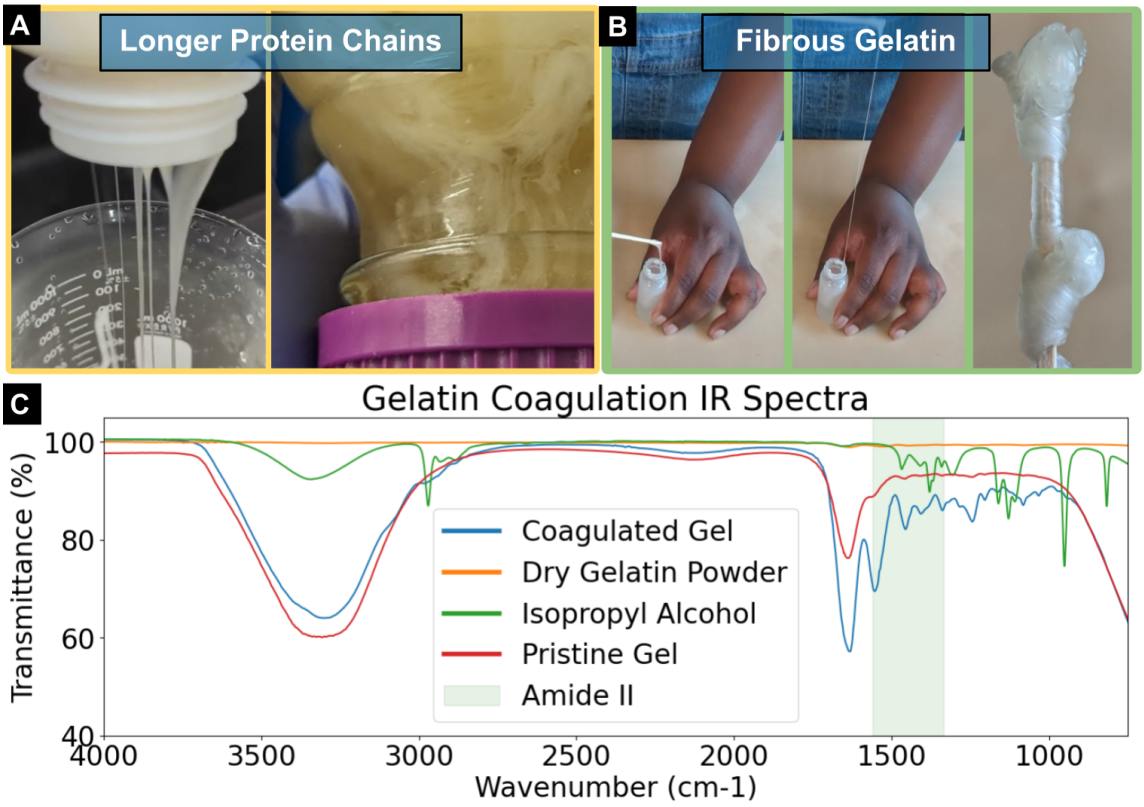

FIGURE 4.4: A) Coagulated gelatin forms longer chains of gelatin protein, leading to higher molecular weight. B) Macroscale behavior of long-chain gelatin polymer shows easy fiber formation, validating long-chain effects. C) Fourier-transform infrared spectroscopy data of different pristine gelatin gel, pristine gelatin powder, coagulated gelatin gel, and isopropyl alcohol.

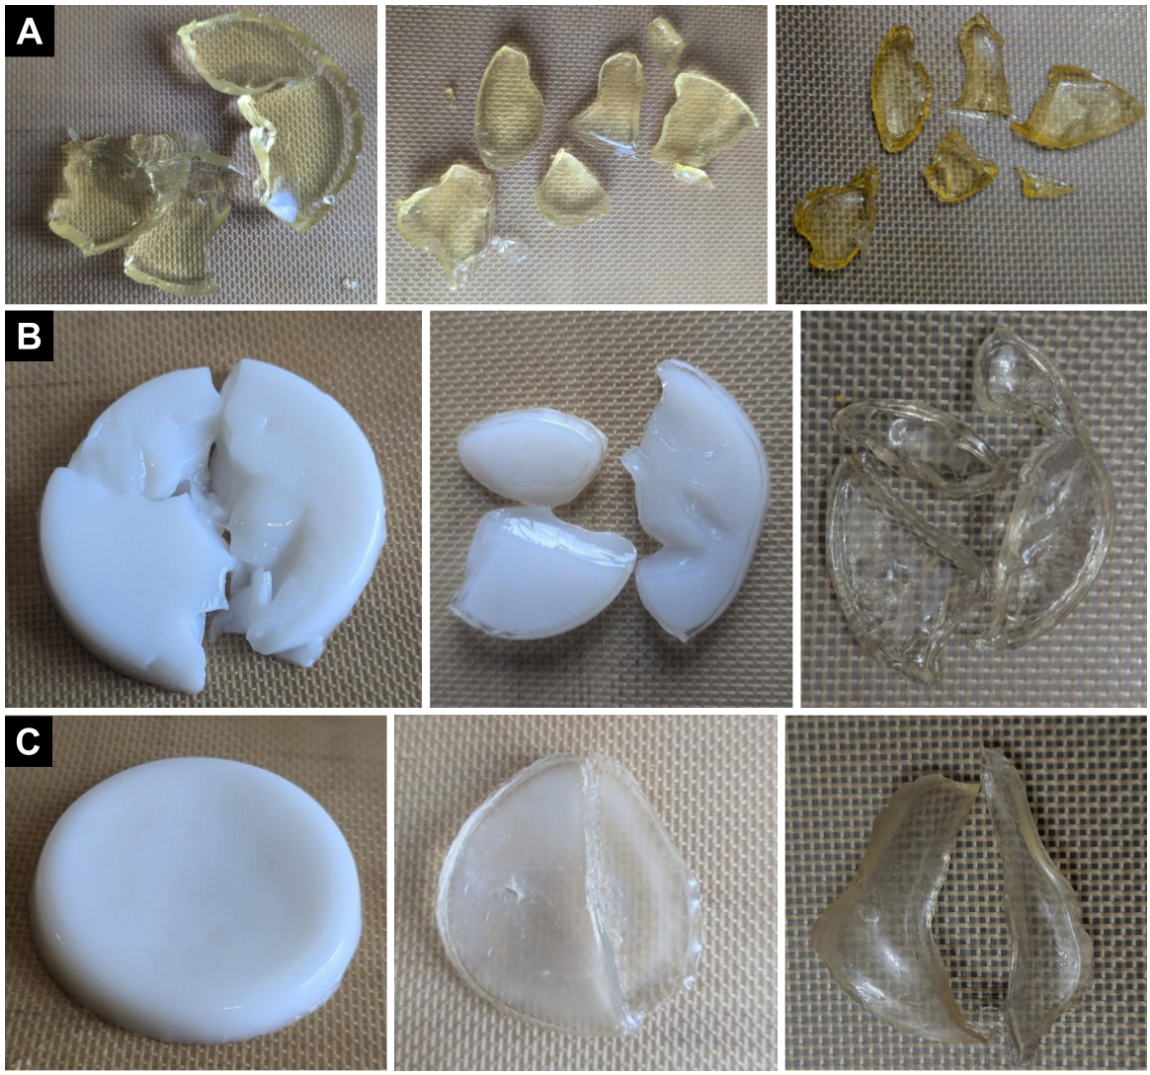

FIGURE 4.5: Drying of different gelatin hydrogels. A) Hydrogel made from pristine gelatin. B) Hydrogel made with gelatin where alcohol was added to a molten solution of fully dispersed and polymerized gelatin. C) Hydrogel made from gelatin where alcohol was added to room temperature hydrated gelatin, then heated to form solution.

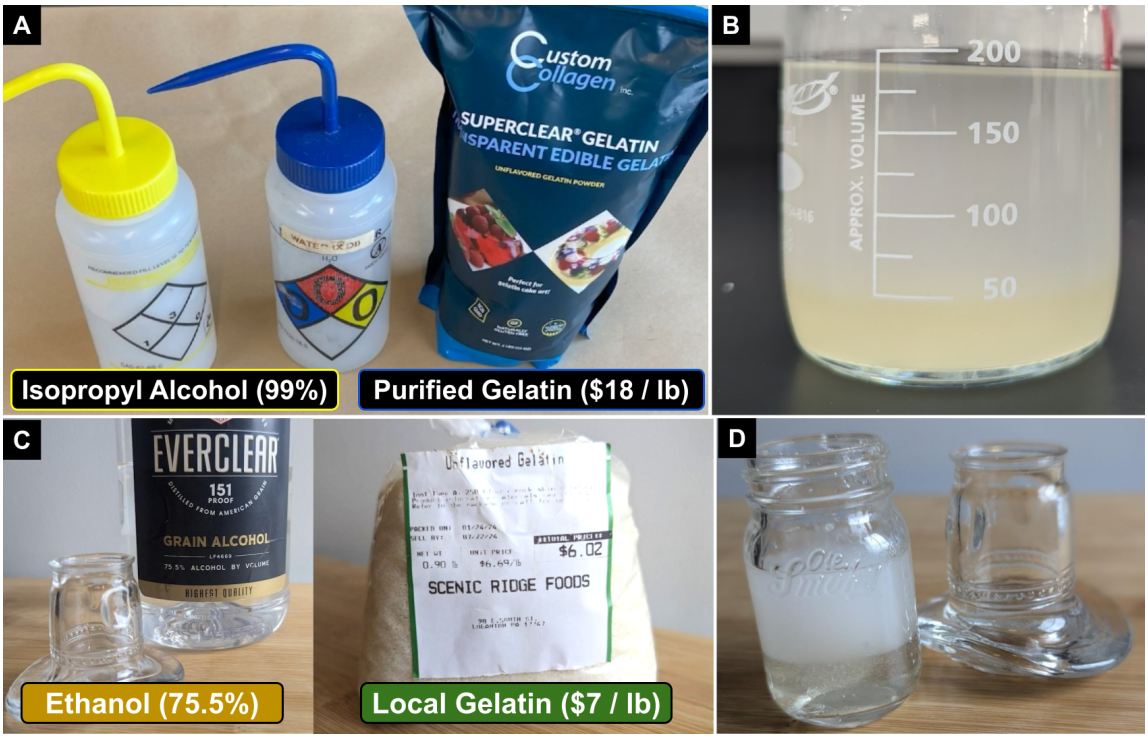

FIGURE 4.6: A)Technical grade, 99% pure isopropyl alcohol, and commercially available purified Type A porcine gelatin powder designed for optical clarity and high bloom strength. B) A two-phase solution of supernatant and coagulated phase from higher-end ingredients. C) Commercially available solution of), with 75.5% ethanol and water (Everclear), and inexpensive Type A porcine gelatin powder locally sourced from Scenic Ridge Foods, an Amish community store in Loganton, PA. D) A two-phase solution of supernatant and coagulated phase from accessible ingredients.

be fabricated into components and mechanisms using processes compatible with thermoplastic polymers. Of these, casting is the most straightforward, where liquid hydrogel is poured into a mold and takes its shape as it cools and gels.

The next phase of this work was adapting this process for use with coagulated gelatin. After the coagulation process, the solution will undergo phase separation if allowed to rest. The majority of the alcohol and water can be removed by pouring off the upper supernatant phase. However, the remaining gelatin-rich phase retains significant amounts of water and alcohol trapped within the gelatin (Figure 4.7)A. If allowed to rest for a while in a container, the gelatin chains compact to form a hydrogel, which is fairly elastic and feels comparable to a rubber. However, in order to add additional ingredients, this form is not particularly convenient. Adding glycerol is especially desired so that the gel retains softness and elasticity as water component evaporates.

One challenge for this process is in controlling the quantity of coagulated gelatin used in making gel. To fairly compare the properties of coagulated gelatin to pristine gelatin hydrogels, the same amount of gelatin polymer should be present in both samples. However, determining the amount of polymer present in a sample of freshly-coagulated gelatin is non-trivial, as the exact proportion of retained alcohol and water is unknown. An initial proportion is calculable via weighing, and immediately after coagulation, there is ≈ 55 weight % of gelatin in the material. However, this proportion changes over time as volatile components evaporate. A simple solution is to use the whole amount of coagulated material, as the amount of gelatin used initially is known. However, this strategy requires going through the whole coagulation process every time a new sample is needed, and it makes parallel fabrication of many samples more difficult.

I initially tried to get accurate gelatin content measurements by drying samples of coagulated gelatin before before using them. In Figure 4.7C, I show what happens if drying is performed rapidly and at high temperatures. The volatiles trapped in the gelatin vaporize, forming many bubbles and causing expansion of sample. This process produces an uneven, rigid foam that shows discoloration, indicating some chemical changes. When this foam is removed from heat, the air within quickly cools and contracts, buckling the surface of the material, and apply large forces to the container it is in, which is sufficient to crack a glass petri dish.

I was able to effectively dry coagulated gelatin by use of a commercial dehydrator, as shown in Figure 4.8. These machines are used for drying food, and apply low heat and airflow over the course of many hours. Coagulated gelatin can be dried at 35°C over the course of 48 hours. It needs to rotated and moved periodically to expose wet surfaces. This process produces a rigid, lightweight material. Any bubbles are too small to see, and there is no indication of chemical changes due to color. However, the material is very stiff and hard to obtain samples from by hand. It would be possible to grind or pulverize it to get powder, however, pulverizing long chain polymers is enough to significantly reduce the average chain length [142], which is the primary property desired from the coagulated gelatin. I therefore process dry coagulated gelatin into a stock material which is able to be easily cut, and in which the proportion of gelatin is known.

In Figure 4.9, I show the process of creating a coagulated gelatin sheet stock. This process begins with the dissolution of known mass dry coagulated gelatin. Dissolving dry coagulated gelatin is slow due to the need to rehydrate and resolubilize gelatin chains. Dissolution can be accelerated by large amounts of water and/or high temperatures. For

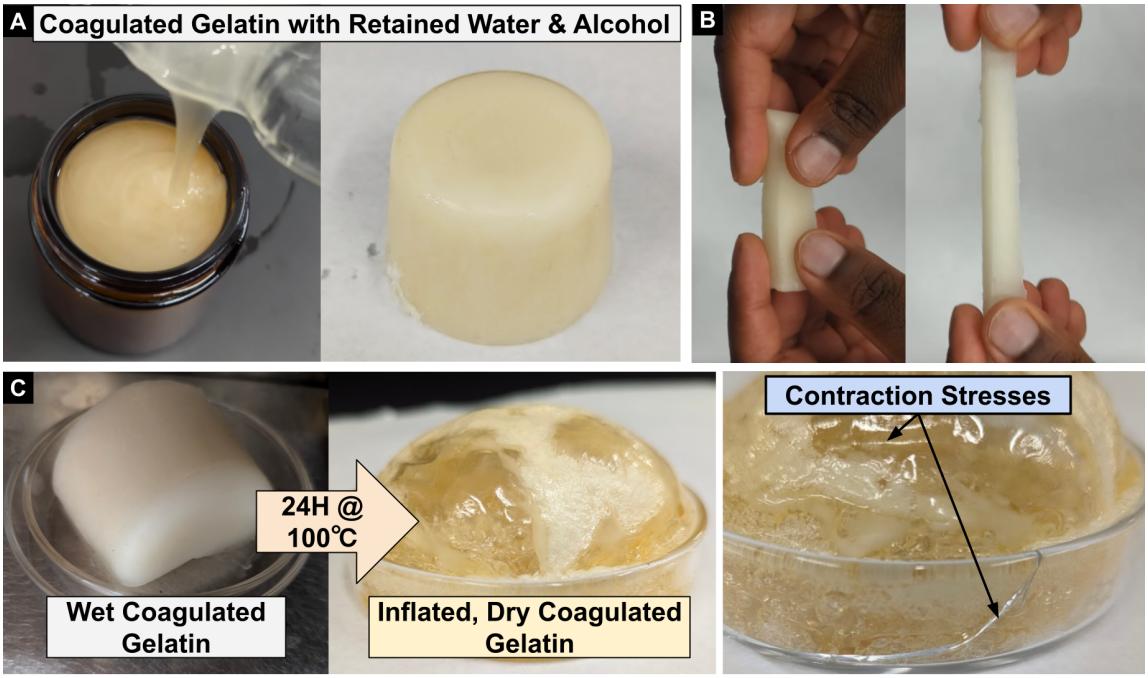

FIGURE 4.7: A) Coagulated gelatin was poured into the container and allowed to sit to form a bulk gel. The mass of bulk is larger than that of gelatin powder due to retained water and alcohol. B) This form of coagulated gelatin is stretchable and soft. C) Heating wet coagulated gelatin at high temperatures produces internal bubbles due to solvent/non-solvent vaporization, resulting in a rigid foam-like structure that contracts when cooled.

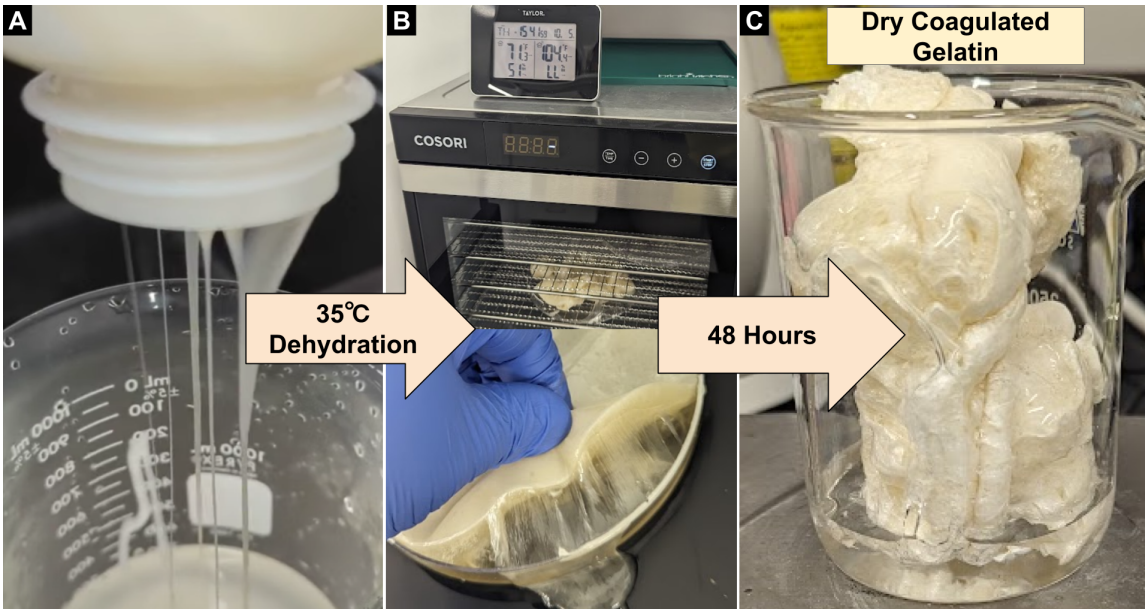

FIGURE 4.8: Drying process for coagulated gelatin. A) Freshly coagulated gelatin is poured into a container. B) Drying in a commercial dehydrator at 35°C for 48 hours. The process can be accelerated by reorienting material and cutting to expose more surface area. C) Dry coagulated gelatin is a rigid, lightweight foam.

1:1 ratio of dry coagulated gelatin to water, full dissolution requires over 2 hours at temperatures above 150°C. Once dissolved, a small mass fraction of glycerol can be added to the solution, and a sheet of stock can be cast. This sheet can be dehydrated to drive off excess water, but retains flexibility due to glycerol. Strips of material can be cut out of this stock for use in hydrogels.

Drying coagulated gelatin and working with dry coagulated gelatin are both fairly labor and time intensive processes, so I sought to remove that step from the process. The creation of a stock of material with glycerol provided a path to doing this. Freshly coagulated gelatin can be dissolved at much lower temperatures and with much less water than dry coagulated gelatin. A batch of gelatin of known input weight can coagulated. This coagulated gelatin can then be immediately dissolved in a 1:1 ratio with water at 85°C. Glycerol can be added as before, and a sheet can be cast. The sheet can then be dehydrated, driving off both water and alcohol. This process of making stock is used for making the sample characterized later in this chapter.

The management of water content is a major theme in the fabrication processes of gelatin hydrogels. This is expected, as the interaction of water and biomaterials is a major determining factor of the properties living biological materials [143]. Additionally, being swollen with water is the default for much of the hydrogel literature, as in biomedical contexts, hydrogels are always used in an aqueous environment. However, when fully swollen with water, hydrogels are less mechanically robust than when they have less water, and indeed the maximum stiffness of gels is achieved when there is minimum water. To achieve material properties comparable to synthetic soft elastomers from a hydrogel form, the water content should fall on the low end of this spectrum. However, it is not easy to work with gels with low water content. If there is insufficient water, liquid gels

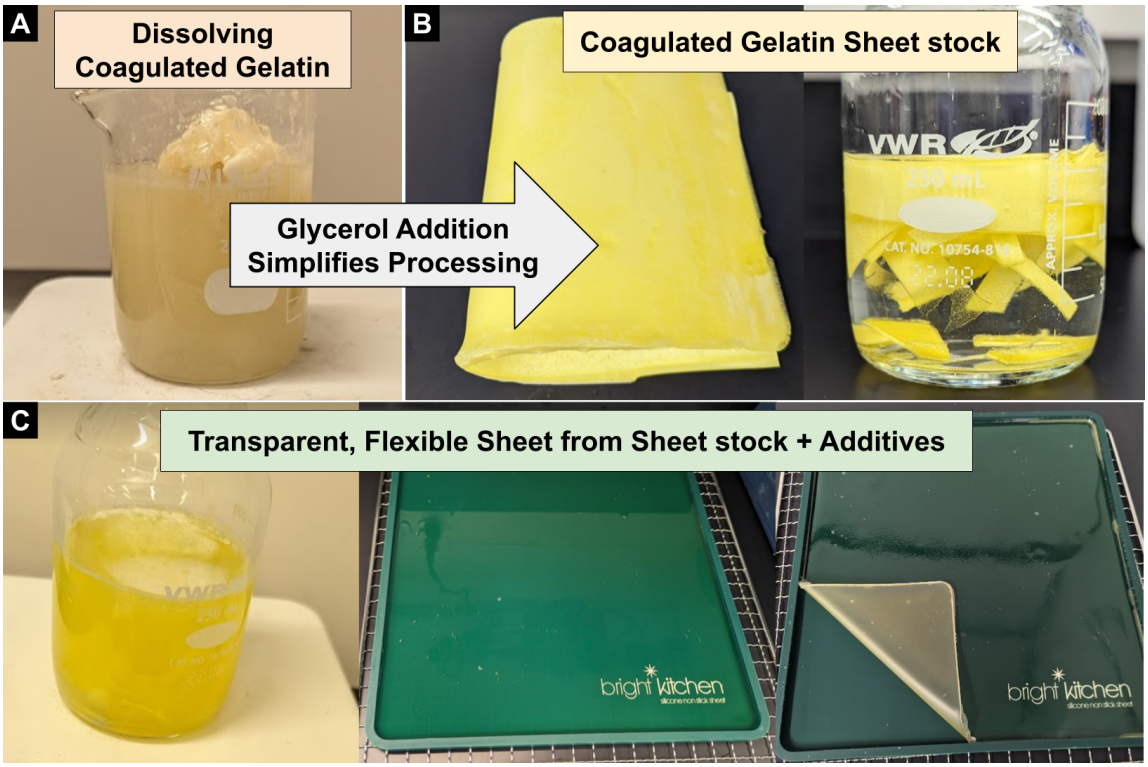

FIGURE 4.9: A) Dissolving coagulated gelatin. B) Adding glycerol to coagulated gelatin solution allows the casting of flexible sheet stock, which can be stored and easily cut for use in making samples. C) Casting a sheet of coagulated gelatin composite from sheet stock prepared with dry gelatin.

quickly set, making it very difficult to even pour a sheet (Figure 4.10A). I therefore take advantage of the fact that these hydrogels do not need to operate in an aqueous environment. I modify gel composition to use much more water than I want in the final composition for processing. After the material is poured into the desired form, it can be dehydrated to the target composition.

This process of dehydration causes dimensional changes, shrinking the hydrogel after it is cast into a particular form. These dimensional changes mean that if the hydrogel is cast into the final form it is intended to be used in, it will shrink and possibly warp, as can be seen in Figure 4.5. Luckily, this can be designed around. First, there is already some benefit by the addition of glycerol. When glycerol is present, the hydrogel does not dehydrate as quickly or completely, as glycerol is incredibly hygroscopic, and large glycerol ratios can prevent much water loss at all [115]. However, large glycerol ratios produce very soft gels which are not rubber-like and do not support their own weight well. This excessive softness would make them unsuitable for use as mechanical components of soft robots like the platform shown in Chapter 3. Therefore, I focus on producing sheets of material, which I dehydrate to produce a stable, equilibrium water content.

From these equilibriated dry sheets, I can cut samples which retain their dimensions. To create specific shapes, like dogbones coupons for tensile testing, I use 3D printed stencils and cut with an exacto knife. This method is inexpensive and simple, however it does require a greater degree of manual dexterity than other parts of the fabrication pipeline. I briefly explore laser-cutting as an alternative process which is more precise and requires less dexterity. In Figure 4.10, I show that hydrogels containing glucose are not compatible with laser-cutting due to the production of carbonized sugar in the air which can coat the laser optics. I was able to laser-cut gelatin hydrogel by removing glucose, so that path is an option. However, I stick with the manual method for the rest of this work.

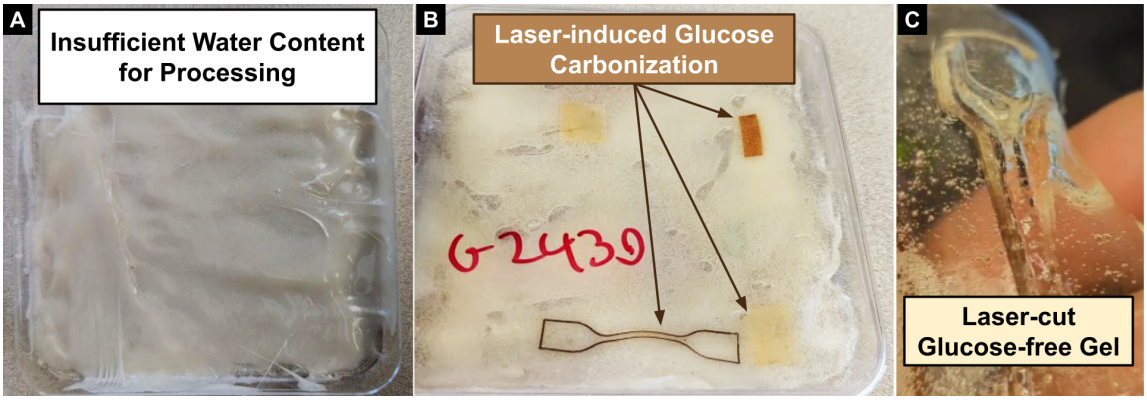

FIGURE 4.10: A) Hydrogel at water content levels optimized for final mechanical properties flows poorly, complicating processing. B) Presence of glucose improves mechanical properties [115], but prevents laser-cutting hydrogel due to rapid deposition of carbonized glucose on laser optical components. C) Glucose-free hydrogel can be laser-cut effectively.

4.4 Mechanical Properties of Dried Gelatin Hydrogels

To characterize the mechanical properties of my gelatin hydrogels, I perform tensile testing to failure and cyclical tensile testing. Experimental parameters for mechanical characterization can be found in Appendix C.

I characterized three compositions, one made with pristine gelatin, and two made with coagulated gelatin with different glycerol ratios. (Figure 4.11) . All samples used the same input mass of gelatin and sugar. They were fabricated with a 1:12 ratio of gelatin to water, then dehydrated to equilibrium. the gelatin coagulation process increased the maximum amount that gelatin hydrogels can be stretched, improving strain-at-break from around 85% to around 185%. Comparing coagulated gel with 1:1 gelatin to glycerol to sample with 1:2 gelatin to glycerol ratio, the strain-at-break rises from 185% to 470%. The compliance of the hydrogels follows a similar trend. The pristine gelatin composition is the stiffest at 69.9 ± 17.8 MPa, while the coagulation process lowers that stiffness 12.1 ± 3.6 MPa. The doubling of glycerol content further lowers it to 0.083 ± 0.068 MPa. Let us consider some comparable results from the Baumgartner paper from which the initial recipe for non-coagulated gelatin hydrogels was adapted [115]. Their G1644 composition has a ratio of gelatin to glycerol of 1:2.75 and a strain-at-break of $263 \pm 27\%$. Compared to the double glycerol coagulated gelatin gel which has a ratio 1:2 and a strain-at-break of $471 \pm 75\%$. There is a general trend of glycerol increasing the amount the hydrogel can be stretched before breaking, but this comparison shows that the coagulation process increases this significantly. Their G2430 composition has a 1:1.05:1.35 gelatin to glucose to glycerol ratio compared to our pristine composition which has a 1:1:1 ratio of these. The stiffness of G2430 is $0.26 \pm .006$ MPa compared to the stiffness of the pristine composition at 69.9 ± 17.8 MPa, more than an order of magnitude stiffer. While the pristine has 30% less glycerol than G2430, this difference is stiffness is likely not sufficient to explain the full difference. Instead, the drying process during sheet casting intially drives off moisture that would be lost over time in the G2430 formulation. They show that their gels can become 4 times stiffer over the course of 20 days, and the water content has not stabilized, so the ultimate stiffness would be higher [115].

Cyclical loading of samples was done to characterize the hysteresis features of the different compositions (Figure 4.12). Coagulated compositions demonstrate some strain softening, but loading and unloading curves follow much more similar trajectories than for the pristine. This qualitative trend is validated by the resilience calculation, which compares the energy stored during loading to the energy recovered during unloading. Coagulated gelatin samples have resilience of $73.3 \pm 0.94\%$ which increases to $86 \pm 0.81\%$ for the double glycerol composition. These are both higher and more consistent than the pristine composition, at $53 \pm 4.1\%$ resilience. This difference indicates that the within the 10% strain regime, more of the energy stored in the coagulated gelatin samples is recoverable, which is consistent with long-chain stretching rather than bond breaking. This difference also provides some insight into why the coagulated gelatin appears more compliant than the pristine gelatin during tensile failure testing. Stretching the pristine composition causes bond breakage, which dissipates energy, effectively stiffening the sample. As we can see in the hysteresis data though, this stiffness will not be present on subsequent stretching cycles as those lost bonds are not recovered.

For the purposes of creating soft robotic systems that can manipulate the world around them, support their own weight and perform useful tasks while in close contact with

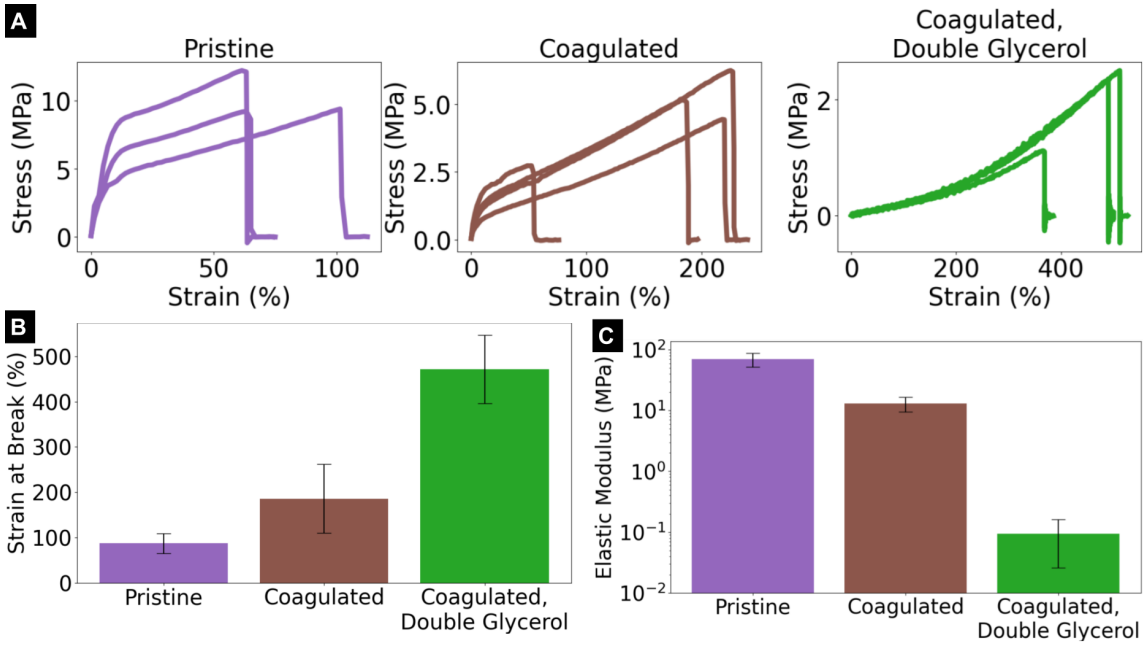

FIGURE 4.11: Tensile failure comparison across hydrogel composites: 1. Hydrogel with 1:1:1 weight ratio of pristine gelatin, glucose, and glycerol. 2. Hydrogel with a 1:1:1 weight ratio of coagulated gelatin, glucose, and glycerol; 3. Hydrogel with a 1:1:2 weight ratio of coagulated gelatin, glucose, and glycerol. A) Raw stress-strain behavior of tensile testing till failure. B) Strain-at-break of pristine

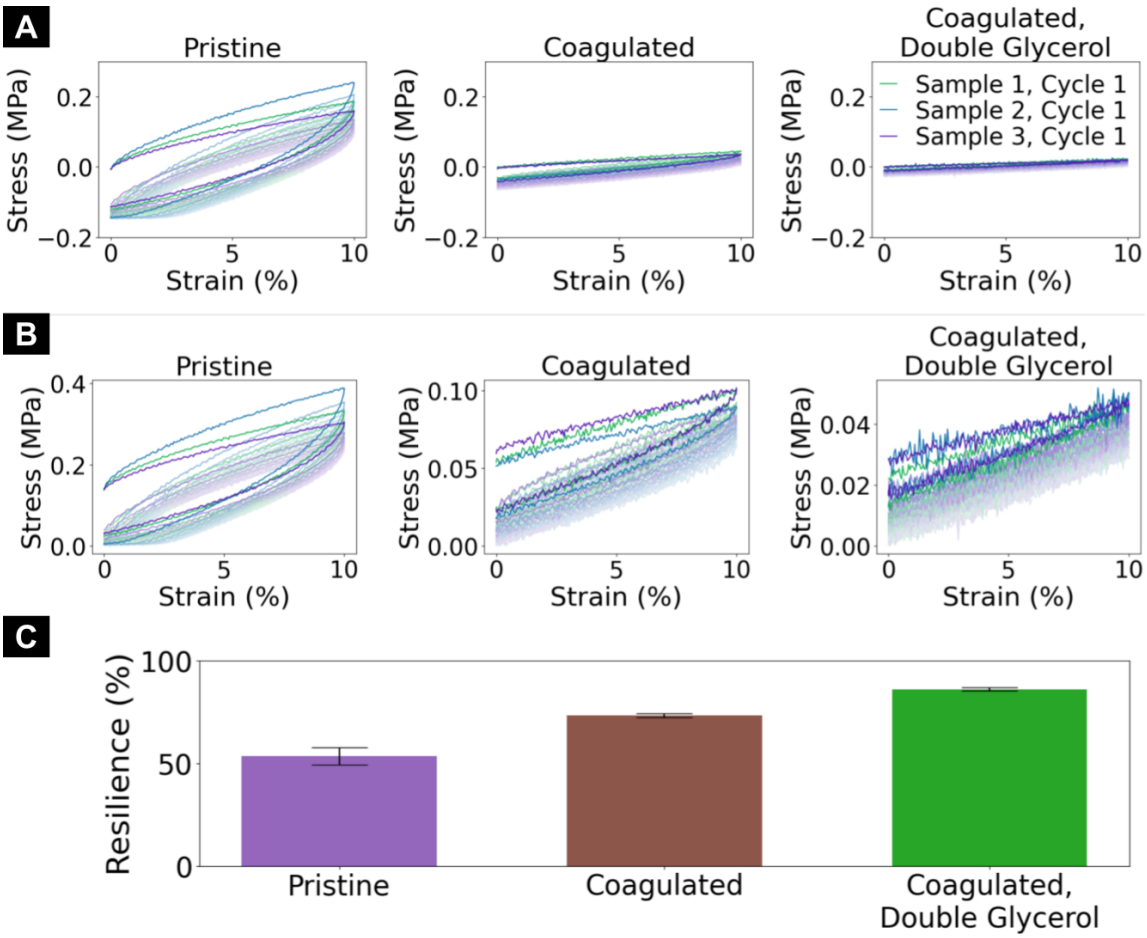

FIGURE 4.12: A) Hysteresis curves of different compositions of gelatin composite. Color saturation decreases, and value increases for later cycles. Axes scaled for quantitative comparison. B) The same data, scaled for qualitative comparison. C) Mechanical resilience was derived from these compositions.

humans, we want tough, stretchable hydrogels with good stiffness. The exact stiffness required will depend on application, but some important benchmarks are the stiffness of silicone (0.1-10 MPa), skin (0.06-0.85 MPa) , natural rubber (around 30 MPa), skeletal muscle (0.005-0.17 MPa) and bone (10,000 - 20,000 MPa) [144]. The most compliant hydrogel I characterized has a stiffness of $\approx$ 0.08 MPa, around the stiffness of human skin and the most compliant silicones. The stiffest hydrogel had a stiffness of $\approx$ 69.9 MPa, stiffer than the stiffest silicones, rubber and muscle, but far less stiff than bone. This range is promising for enabling robots that can interact with the human body with varying degrees of compliance while retaining intrinsic safety benefits from softness.

4.5 Reducing Water Solubility

In this section, I explore the hydrolytic stability of gelatin-glucose hydrogels, and reduce their solubility in water via thermal crosslinking. The work in this section was performed in collaboration with my undergraduate mentee Zimo Ge.

A major factor in the applications of gelatin hydrogels is their ready solubility in water. In traditional tissue scaffolding applications, a water soluble hydrogel degrades quickly as it is constantly exposed to the aqueous environment of the body. To address this, gelatin hydrogels are often crosslinked and otherwise chemically modified to make them insoluble. The standard in that field is the production of gelatin methacryloyl (GelMA), by reaction of gelatin with methacrylic anhydride in a ratio of approximately 0.628:1 to 1.859:1 [145]. Methacrylic anhydride is a fairly expensive reagent, costing around $58 per 100 mL, 94% [146]. Genipin, a compound extracted from *Genipa americana* or *Gardenia jasminoides* fruit has been used to crosslink pristine gelatin hydrogels [122] at a ratio of 100 mol% and coagulated gelatin fibers at a ration of 0.16 weight%. [134]. This compound is expensive as well, costing $4260 per gram from SigmaAldrich [147], though there are less established suppliers claiming $14 per gram [148]. There are less expensive options for crosslinking, formaldehyde (37%) costs $19.95 for 16 oz [149], and has been used to crosslink coagulated gelatin fibers [30], though this option is inaccessible for safety reasons. Glutaraldehyde is a widely used, inexpensive and effective crosslinker for gelatin, which has been used in clinical settings. However, there are concerns of it being released during internal degradation due to the cytotoxity of unbound glutaraldehyde [150]. I pursue the route of crosslinking with sugar due to easy access, safety and low cost, as shown in Figure 4.13.

Gelatin is a biopolymer made of amino acids. The bonding of amino acids to sugars is not an esoteric reaction, in fact, the Maillard reaction, which is responsible for the delicious browning of food during cooking, is a reaction between amino acids and sugars. Glutaraldehyde can be produced via the Maillard reaction, Conveniently, my gels already contain glucose, which can undergo this reaction. Glucose is very inexpensive ($25 for 2.2 lb on Amazon). I compare the prices of crosslinking a pound of gelatin with different crosslinkers and at different degrees of crosslinking in Figure 4.13. Cross-linking between gelatin and sugar can be initiated by ultraviolet irradiation [151], acid-base neutralization [152], and elevated temperatures [153, 154]. I focus on thermal curing to induce crosslinking between gelatin and glucose due to the ability to achieve the required temperatures with domestic equipment and the independence fro additional reagents.

In the first solubility experiment (Figure 4.14), I use samples with a 2:1:2 ratio of pristine gelatin, glucose and glycerol for testing. Samples are baked for 24 hours at different

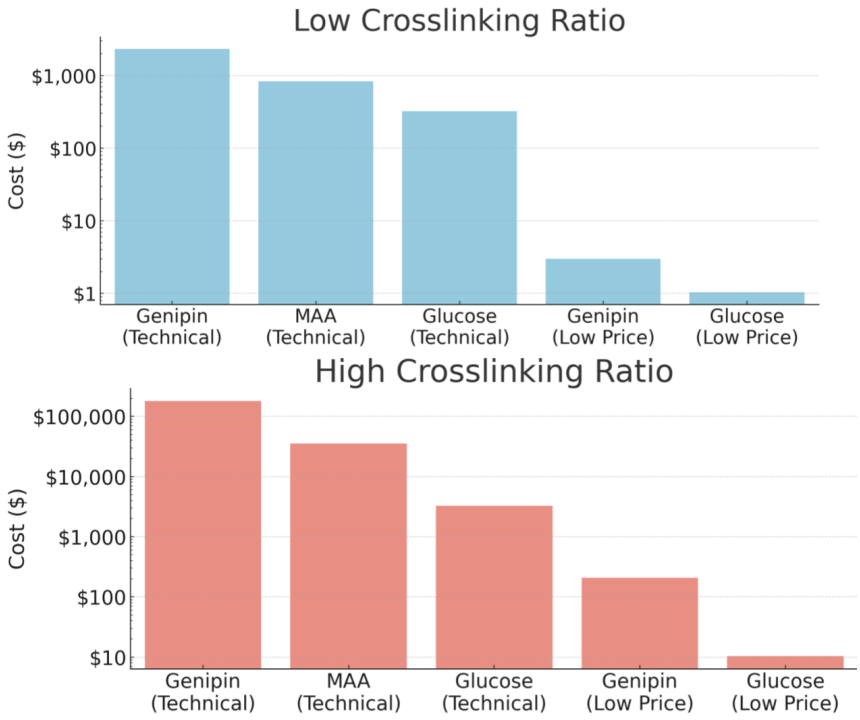

FIGURE 4.13: Cost comparison of crosslinking one pound of gelatin. Technical grade crosslinkers are sourced from SigmaAldrich. Low price crosslinkers are sourced from online retailers. The crosslinking ratio is based on range of ratios found in selected literature. Genipin: 0.544 g/lb to 42.14 g/lb [122, 134], Methacrylic anhydride (MAA): 14.38 mL/lb to 42.57 mL/lb [145], Glucose: 1.6 oz/lb to 16 oz/lb [34, 115].

temperatures, 40°C which is around the melting point of fully hydrated gelatin, 100°C and 120°C. The baked samples are submerged in water and heated for 1 hour at 50°C and agitated periodically to induce hydrolytic stress, dissolving soluble material (Figure 4.14A). The undissolved material is filtered out with filter paper of known mass, and weighed (Figure 4.14B). The paper and undissolved gel are dehydrated for 48 hours at 35°C, then weighed again (Figure 4.14C). The dry mass of the undissolved gel is calculated and the percentage of insoluble material to original mass is calculated (Figure 4.14D). Baking at 40°C does not reduce solubility, but baking above 100°C does. However, at 120°C, there is less mass remaining, indicating some thermal degradation effect. I explore this next.

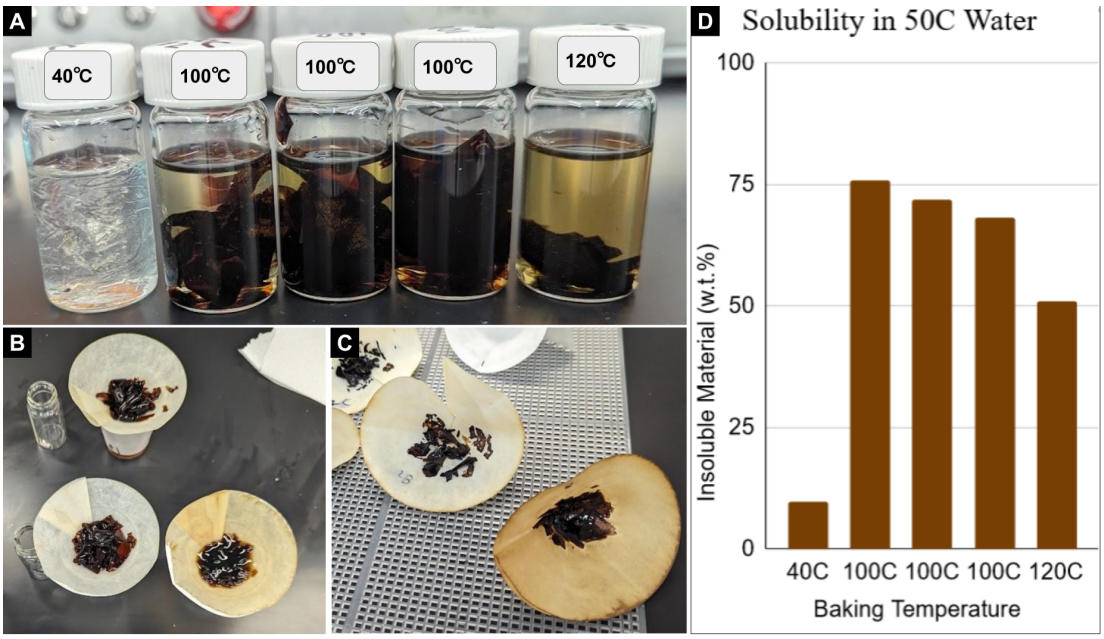

FIGURE 4.14: A) Baked gelatin-sugar hydrogels submerged in water and heated for 1 hour at 50°C and agitated periodically. Brown color indicates chemical reaction B) Undissolved material drained in filter paper. C) Drained material dehydrated and weighed. D) Weight percent of insoluble material to original mass.

The Maillard reaction provides insight into the potential mechanism of thermal degradation, as the Maillard reaction is inhibited by the presence of moisture. To test this, I compared the results of baking for 24H at 100°C in a moisture rich environment, by sealing the lid of the container vs. a low-moisture environment, where the container is open and water is able to evaporate away during baking. As shown in Figure 4.15B, the samples baked in a moisture rich environment lose the ability to gel, and degrade into sticky paste. It has been shown that reducing the molecular mass of gelatin sufficiently prevents it from gelling [124], which might indicates that the high temperature, moisture-rich environment is causing hydrolysis of gelatin chains to smaller polypeptides. If this is the case, the loss of the absorption peaks associated with the broken bonds should be visible in the IR spectra of the material. As we can see in Figure 4.15A, the samples of coagulated and pristine gelatin that were baked with the lid on follow virtually identical spectral curves, with no differences in noticeable peaks. This change is one indication that bonds

are being broken, as the peaks associated with extra bonds formed during coagulation (around 1550 cm^{-1} and between 1170 - 1500 cm^{-1} in the Amide II region) are gone. In samples baked without trapped moisture, which should be thermally crosslinked, there is a very strong peak around 1040 cm^{-1} and two small peaks at 925 cm^{-1} and 850 cm^{-1} in the Amide III region. These crosslinked samples also show another peak around 2920 cm^{-1} in the Amide A region. The 1040 cm^{-1} peak is present in the thermally degraded samples baked in the presence of moisture, but is much reduced.

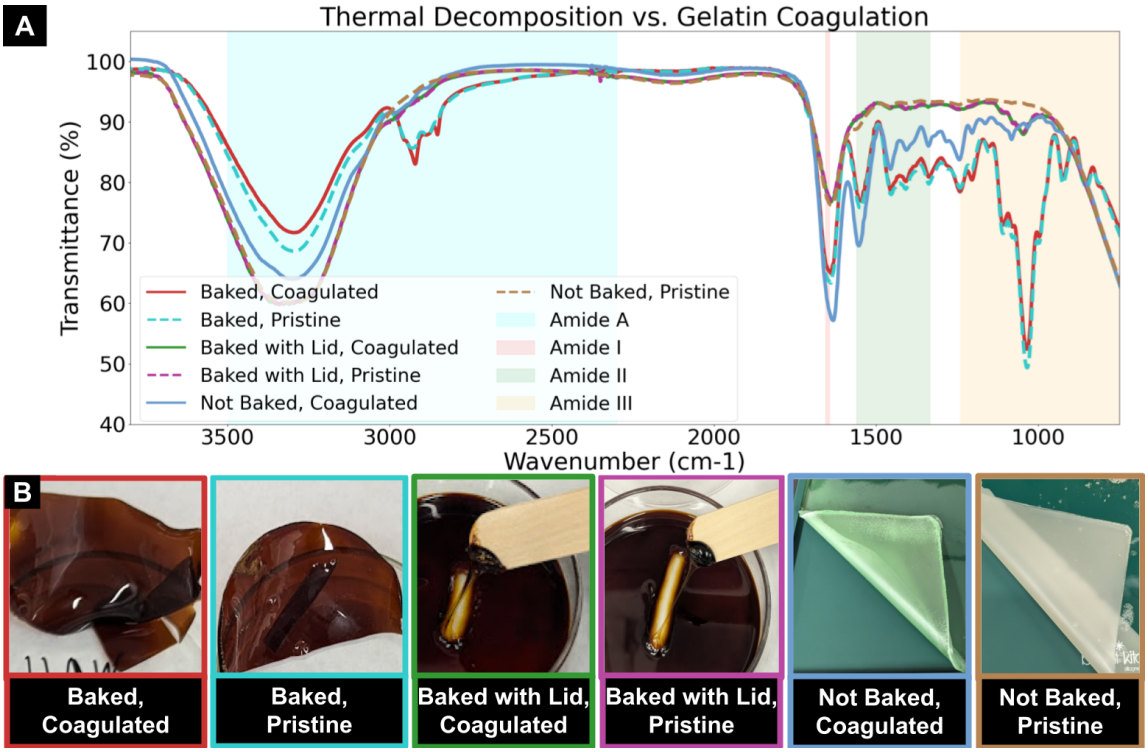

FIGURE 4.15: A) Fourier-transform infrared spectroscopy of gelatin-sugar hydrogels. B) Samples characterized above. Baked samples were baked at 100°C in a glass petri dish for 24 hours, either with an open lid or with the lid on to trap moisture.

I next characterize the solubility of thermally crosslinked samples of different compositions, and whether drying samples before baking improves this (Figure 4.16). These samples were all baked without lids for 24 hours at 100°C. We can see that the solubility of baked samples does not correlate strongly with whether or not the gelatin has been coagulated. We can also see that drying ahead of time does not have a big impact on solubility, indicating that moisture is driven off sufficiently quickly in the thermal curing environment. Looking at the IR spectra 4.16)A, there are again strong peaks at 1040 cm^{-1}, 1550 cm^{-1}, and 2920 cm^{-1} for the baked samples, with minor at 925 cm^{-1}, 850 cm^{-1} and between 1170 - 1500 cm^{-1}. Some of these features are shared between uncured coagulated gelatin hydrogel and the thermally crosslinked samples. The peak around 1550 cm^{-1}, the peak around 1040 cm^{-1}, and the minor peaks between 1170 - 1500 cm^{-1} are shared between coagulated gelatin gel, and the crosslinked pristine and coagulated gels. The peak at 1550 is slightly stronger for coagulated gelatin, while most of the others are slightly stronger for thermally crosslinked gels, with the exception of the 1040 cm^{-1} peak, which

is much stronger for thermally crosslinked gels and is the strongest peak for those samples. Based on these details, the 1040 cm^{-1} peak (primary), 2920 cm^{-1} peak (secondary), 925 cm^{-1} and 850 cm^{-1} peaks (tertiary) are the indicators of thermal crosslinking between gelatin and sugar. Evidence that thermal crosslinking of gelatin with glucose through the Maillard reaction is found by comparing the strong peaks of crosslinking to the IR signature of glutaraldehyde, a known Maillard product which possesses strong peaks at approximately 1040 cm^{-1} and 2920 cm^{-1} [155].

Thermal crosslinking improves solubility significantly, from around 15% insoluble material left after hydrated heating at 50°C to around 75% insoluble under these same conditions. Since I aim for these materials to be used as components of soft robotic systems rather than in aqueous environments, this improvement to hydrolytic stability is very promising. For most contexts, these materials would be less damaged by the presence of water than many of the electronic components of a robotic system.

4.6 Discussion

In this chapter, I fabricate and characterize biomaterial hydrogels using accessible, biodegradable ingredients like glucose, glycerol and gelatin. I leveraged lab equipment for characterization, but put focus on being able to adapt the fabrication process to widely available tools, so that someone trying to work with these materials could likely fabricate them with the existing tools available in their particular work space. This constraint is important for participatory robotic systems development.

I demonstrate the ability to modify mechanical properties of hydrogels via use of accessibly coagulated gelatin. Coagulation of gelatin enhances mechanical performance of hydrogels via increasing the length of polymer chains and can be achieved without laboratory equipment or ingredients, by simply heating and periodically agitating a solution of gelatin, water and alcohol. I show that glycerol can additionally be used to modulate stiffness and elasticity. Freeze-thaw cycling is another accessible way of improving mechanical properties of hydrogels which could be explored [156]. Long chain polymers have been shown to have improved resistance to wear and lower friction [128]. Since gelatin coagulation process produces long protein chains, coagulated gelatin might be useful in applications where wear is the primary failure mode and replacements are frequent. Since coagulated gelatin can be very optically clear, there is likely good synergy with tactile characterization sensors relying on optics, like GelSight [157], where a clear polymer pad is pressed into the surfaces to be characterized and which needs to be replaced very frequently.

I also explore the challenges and opportunities in fabrication of these hydrogels. Throughout this work, managing water content became of primary concern. My results suggest that maintaining a high water content during fabrication, followed by controlled dehydration (mediated by glycerol content) allows a decoupling of the final material strength and the processability. This dehydration induces dimensional changes and even shape change. These effects have been leveraged for hydrogel-based soft actuators, and the rate at which water diffuses can be increased via controlled creation of pores [158]. However, if a stable geometry is desired, gels can be processed into sheets. Due to high aspect ratio, dehydration happens primarily from the exposed top surface, so that shrinkage occurs largely through the thickness dimension of the sheet. This directional shrinkage leads to thin sheets with minimal warping. Parts can be cut from the resulting flat sheets, and

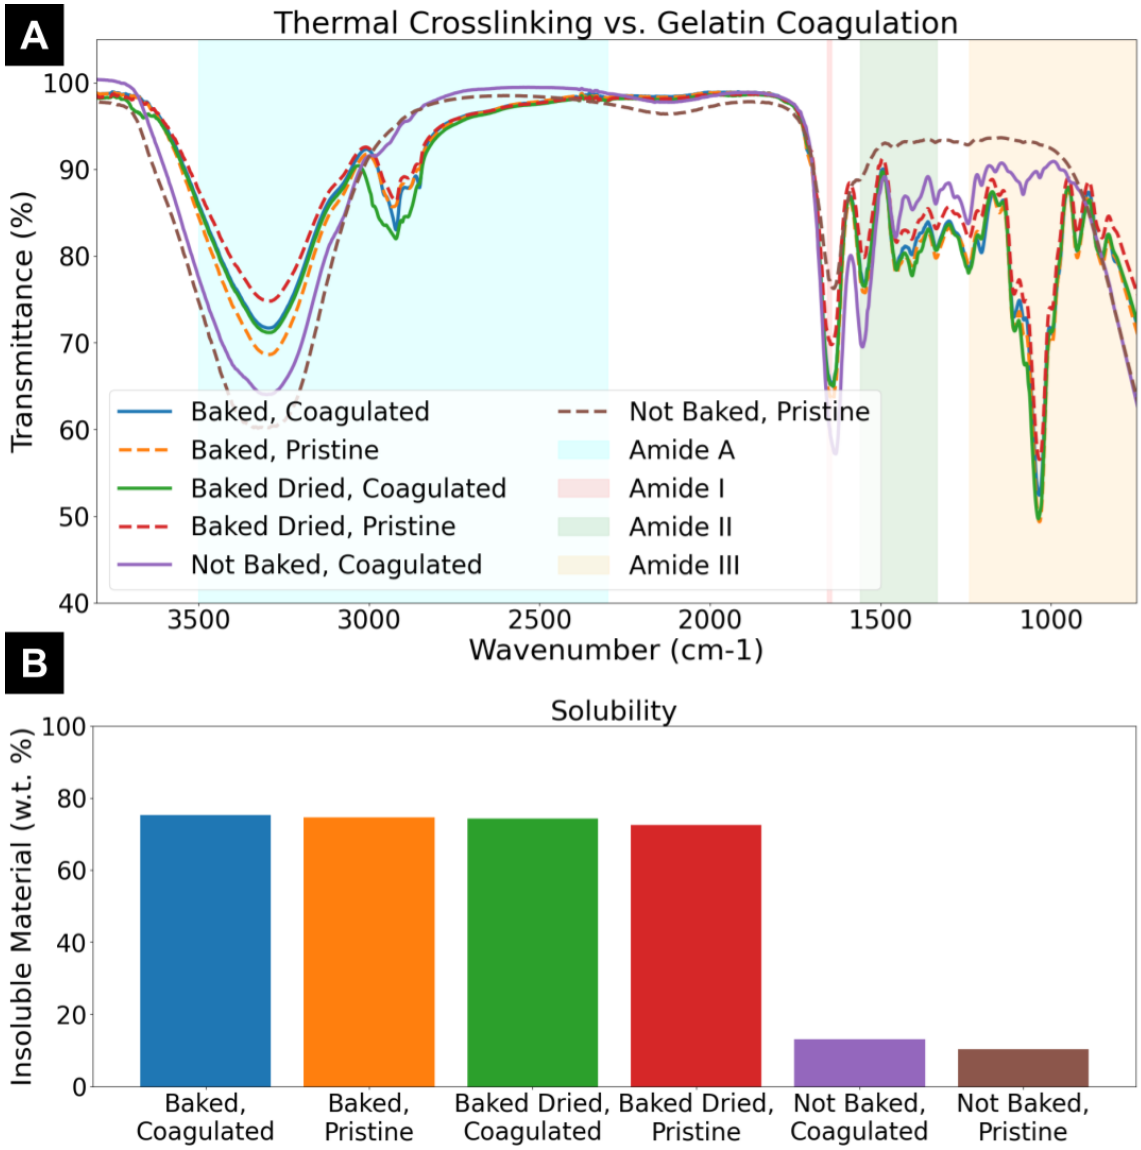

FIGURE 4.16: A) FTIR Spectroscopy of gelatin hydrogels, comparing coagulated vs. pristine and different baking conditions. B) Solubility of gelatin hydrogels.

existing work on origami mechanisms and laminate structures for robotics [159] indicates that a flat form factor after synthesis would still allow a large design space for producing robotic mechanisms. As an alternative to the high moisture fabrication regime, there is work showing that low-moisture hydrogels can be processed by treating them like dough, folding them and kneading them to produce sheets. This process is especially appealing for long-chain, highly-entangled hydrogels, as it produces mechanically robust materials without breaking the long chains [142, 160]. Recent techniques allow for additive manufacturing of synthetic hydrogels with long-chain, entangled polymers via vat photopolymerization aided by a slower redox reaction [161]. 3D printing of gelatin hydrogels has also been demonstrated, [120], and in aqueous environments or with sufficiently high glycerol concentrations and tolerances on dimension, this is a possible path forward. Additionally, additive manufacturing techniques can create structurally complex soft structures, allowing fabrication of artificial tissues and organs [162]. FRESH printing of alginate hydrogels can produce structures with shape complexity on the order of a human heart [163] and a functional robotic arm [164]. Biomaterial hydrogels are not intrinsically limited to simple shapes, and complex structures are increasing feasible.

I demonstrate improvements to solubility of gelatin-based hydrogels by thermal crosslinking between gelatin and sugar via Maillard reaction. I show that if this same thermal treatment is performed in the presence of moisture, hydrogels degrade and coagulation is reversed, which could have implications for controlled degradation of these materials. I briefly explore other ways of modifying gelatin to reduce vulnerability to moisture. First I combine alginate and gelatin hydrogels (Figure D.3) and crosslink with calcium lactate. Alginate is an ionically crosslinked polymer, while gelatin gels via cooling, so careful sequencing of fabrication steps would be required. It might be helpful when thermal treatment is not possible, or if reversing the solubility change is desired . Alginate crosslinks can be removed by chelating agents, some of which are accessible and degradable like sodium citrate which can be made by mixing lemon juice and baking soda. I also experiment with emulsifying hydrophilic gelatin hydrogel with hydrophobic oleogel to form a bigel (Figure D.2). Modeling indicates that depending on temperature and composition, these gels can achieve percolation in neither, one, or both of the phases[165], which implies the possibility of tuning transport of various species, e.g., water via controlling network connectivity. Pure oleogels are of primary interest in controlling the texture and structure of lipids in the food industry [166]. In robotics, they have been used to create soft robotic actuators using shape memory [167] and magnetic actuation [168]. Bigels have been explored for food science [169], drug delivery and cosmetics [170], but these have low polymer fractions and cannot transmit mechanical loads. Hydrogels with phase-separated hydrophobic and hydrophilic networks are possible with synthetic polymers, and demonstrate high strength and low hysteresis [171]. These dual network hydrogels rely on a water-poor hydrophobic hydrogel phase instead of an olegoel phase. Mechanically strong bigels of biomaterial hydrogel and olegel networks have not been demonstrated, but my initial experiments along these lines show that load-bearing biomaterial bigels are possible (Figure D.2). Based on my own work and the implications of the literature, I believe there are diverse possible developmental directions for biomaterial gels suitable for soft robotics.

In the next chapter, I explore the extent to which gelatin-based hydrogels can be used for other elements of robotic systems besides passive mechanical elements, namely as part

of the electronics. I pursue this via creation of conductive gelatin hydrogel, which leverages lessons learned in the importance of managing water content to perform controlled assembly of conductive filler.

Chapter 5

Accessible Soft Electronics with Conductive Gelatin Composite

***K. Chin**, M. Vinciguerra, D.K. Patel, P.R. Olcay, E.S. Lazaro Vasquez, and C. Majidi. "Accessible soft electronics with silver-gelatin conductive hydrogel composite." In Advanced Materials Technologies (In Review), 2024.* [34]

Author Contributions: I designed the conductive gelatin material and the accessible fabrication process. I designed the electromechanical characterization process. I implemented the soft circuitry and EMG electrode demos. M. Vinciguerra developed electrical characterization process. E.S. Lazaro Vasquez provided initial coagulation protocol. D.K. Patel and I performed FTIR experiments and filmed accessible fabrication process. P.R Olcay and I performed EMG impedance characterization.

Electrically conductive hydrogels are a promising class of materials for soft electronics and robotics that mimic the mechanics of natural biological tissue. However, these materials are typically not biodegradable or recyclable and their production typically involves hazardous solvents and monomers that limit accessibility and environmental compatibility. This study introduces a biomaterial hydrogel composite in which a percolating network of silver microflakes is suspended in a natural, gelatin-based matrix. The composite is primarily composed of inexpensive, food-safe ingredients and the fabrication process is designed to be compatible with standard kitchen equipment and can be implemented in a well-ventilated domestic environment (Figure 5.1A). Together, the choice of inexpensive and consumer-grade materials and equipment make this a relatively accessible approach to engineering conductive hydrogel compared to other methods that require access to laboratory-grade chemicals and facilities. The resulting material system is mechanically soft, stretchable up to 470% strain, and highly conductive up to 3.1×10^3 S cm^{-1}, with properties that can be tailored based on material composition and processing conditions. In addition to experimental characterization of its material properties, this conductive gelatin composite is shown to be applicable for a variety of uses cases in soft matter circuitry and bioelectronics.

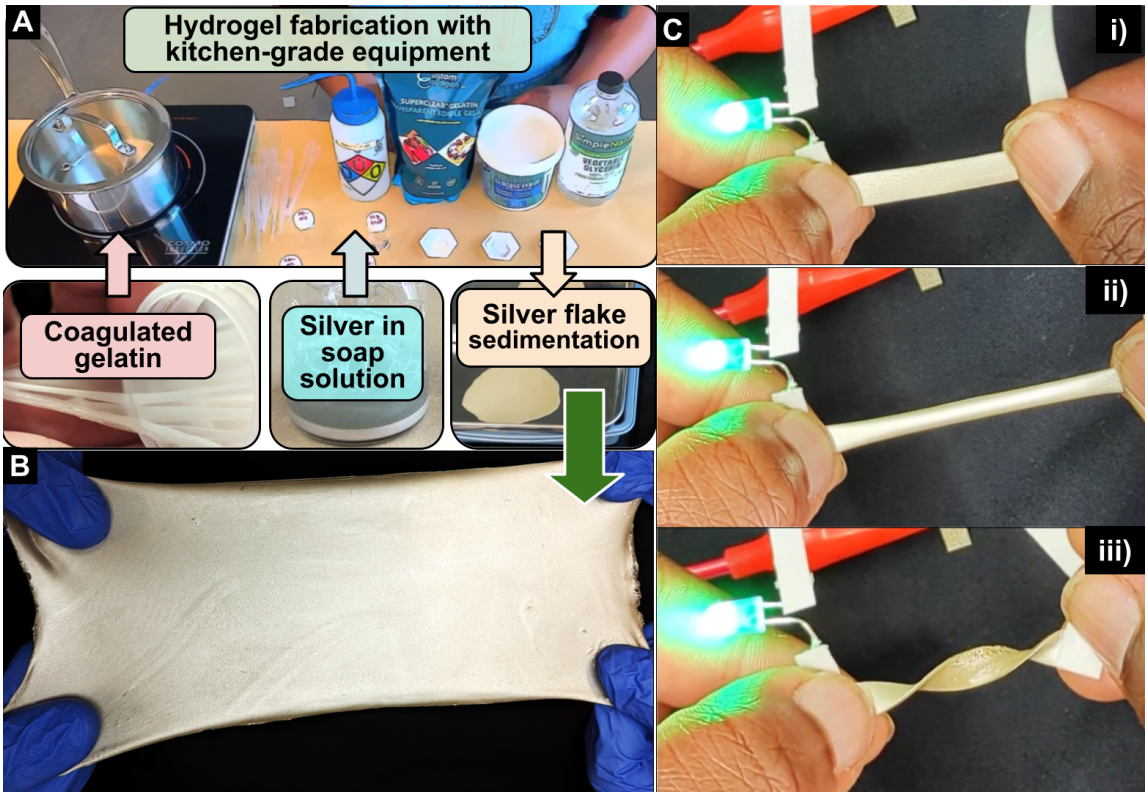

FIGURE 5.1: A) Accessible equipment and ingredients for the fabrication of conductive hydrogel composite. B) Stretched sheet of conductive gelatin sample. C) Strip of GG-C sample functioning as a wire in a circuit under i) normal conditions, ii) stretching, and ii) twisting.

5.1 Overview

There has been growing interest in incorporating gelatin biomaterials, such as the class of materials covered in Chapter 4, into soft electronics and devices [172]. However, there exists an intrinsic tension between mechanical and electrical properties for deformable conductors [173] and the ability to balance these properties is of great importance. Hardman *et al.* showed ionically conductive gelatin-glycerol hydrogel for strain sensing [174]. Jing *et al.* created a conductive filled composite from thiolated gelatin with silver nanowires(GE-SH-AgNWs), designed for wearable strain sensing [175]. Liu *et al.* showed a gelatin hydrogel soaked in ammonium sulfate forming an ionically conductive hydrogel (Gel/AS) [176]. Park *et al.* used reduced graphene oxide (rGO) as a conductive filler in methacrylated (GelMA) gelatin to form conductive hydrogel (r(GO/GelMA)) for applications as a nerve conduit [177]. Han *et al.* combined dopamine grafted gelatin (GelDA) with graphene oxide as conductive filler (GelDA/GO) to form a self-healing, adhesive hydrogel [178].

There have been several simple fabrication processes leveraged in the formation of conductive hydrogels that are promising for accessible fabrication processes. The "salt-out"

effect was leveraged using baths of sodium sulfate or ammonium sulfate after the synthesis of the base hydrogel [175, 176, 179]. Ohm *et al.* showed that partial dehydration of a silver–polyacrylamide–alginate hydrogel can increase conductivity by reducing the space between high aspect ratio conductive particles [180]. Zhao *et al.* showed extremely high conductivities via dry annealing and freeze-thaw processing in a PVA-Borax hydrogel with silver microflakes and gallium-indium alloy liquid metal droplets [156]. Ford *et al.* observed that sedimentation of liquid metal particles in rubber-hydrogel composite could affect conductivity [181]. Gan *et al.* demonstrated sedimentation of liquid metal droplets in gelatin gel and subsequent sintering via gel dehydration to produce thin-films with high stretchability and conductivity [182]. Most of these processes produce either highly conductive hydrogels or are created from sustainable ingredient sources, but not both. The work of Gan *et al.* achieves these goals, at the cost of producing only a thin-film form factor and relying on expensive and low-availability liquid metal conductive filler. An important space remains for a hydrogel produced from sustainable, inexpensive ingredient sources that achieves conductivities over the 10^2 S cm^{-1} threshold for use in a variety of soft matter electronics.

In this chapter, I introduce an electrically conductive hydrogel composed of silver microflakes embedded within a network of food grade biopolymer (gelatin and glucose syrup) and plasticizer (glycerol) (Figure 4.1A). The matrix for this hydrogel is fabricated using the processes and ingredients discussed in Chapter 4. Briefly, I tailor the mechanical properties of the conductive hydrogel composite through the use of glucose and glycerol as biodegradable additives, and coagulated gelatin to form an entangled polymer network. Electrical conductivity is achieved through the use of silver microflakes, which are commonly used for conductive elastomers and gels. An additional benefit to relying on chain entanglement for enhancing mechanical properties of a conductive composite is that it allows the material to maintain thermoplastic properties, preserving the possibility of easily retrieving conductive filler. For characterizing the properties of this material system, I focus on the influence of gelatin composition on mechanical properties (elastic modulus, elastic resilience, strain at break) and electromechanical coupling. To reduce the amount of silver required to achieve percolation and electrical conductivity, I exploit a combination of gravity-driven sedimentation and partial dehydration of the hydrogel matrix [180]. I also examine the influence of silver sedimentation time, which controls the local density of silver microflakes within the gelatin, on electrical conductivity of the composite. To demonstrate the use of the conductive silver-gelatin composite in practical applications, I show that the conductor can be used as a stretchable, finger-mounted electrode that maintains an electrical connection with finger motion and a forearm-mounted bioelectronic sticker that tracks hand motion through electromyography. Collectively, these measurements and demonstrations suggest that silver-gelatin composites could be promising for accessible soft electronics that can be produced with readily-available materials and resources.

5.2 Results

The silver-gelatin composite can be produced in a domestic kitchen environment by mixing silver flakes (47MR-10F, Inframat Advanced Materials) with store-bought, food-grade materials – gelatin, glucose syrup, glycerol and water. Details of the complete fabrication process are presented in Figure 5.2A-B, and are also described in the Experimental Section (video available at [183]). Briefly, as shown in the Figure 5.2A, the glycerol, glucose syrup, and gelatin (pristine powder or long coagulated chains) are mixed with water and melted to form a solution. A separate solution of silver microflakes and dish-soap is prepared under good ventilation (video available at [184]). The two solutions can be mixed, and the resulting mixture is then poured into a mold or dropcast onto a substrate. The gel is kept molten with temperatures over $> 85°C$ in order to maintain a liquid state and allow the silver flakes to settle and form a sedimentation layer (Figure 5.2B). After a sedimentation time of 60 minutes, the material is allowed to cool and gel. Lastly, the silver-gelatin composite is dehydrated in a commercial dehydrator $> 35°C$ for 48 hours, and can then be removed from the mold or substrate. Depending on the silver flake concentration and settling time, the sedimented side of the composite can exhibit a volumetric electrical conductivity up to 3.1×10^3 S cm^{-1}. The following discussion describes how alterations in gelatin composition and silver sedimentation time can influence mechanical, electrical, and electromechanical properties of the composite.

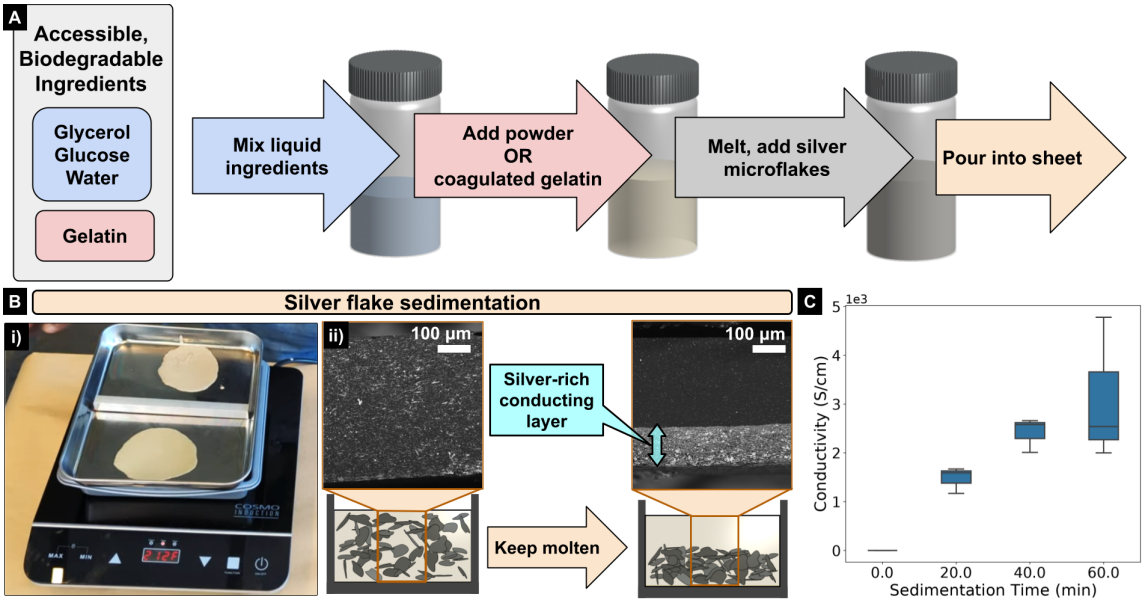

FIGURE 5.2: A) Fabrication process of conductive gelatin hydrogel bulk material. B) Process for gravitational sedimentation of silver microflakes to produce conductive layer. This process is achievable using the double boiler method, using an oven or an induction cooktop. C) Electrical conductivity of CG-C conductive gelatin increasing as a function of sedimentation time.

5.2.1 Material Structure and Composition

The hydrogel matrix is composed of gelatin (coagulated or unmodified), glucose syrup, glycerol and water. The mechanical and electrical properties of the final composite depend on the particular ratios of ingredients and the gelatin type. Rather than a full exploration of this parameter space, we start with a baseline composition of coagulated gelatin (CG), glucose syrup, glycerol and water in a 1:1:1:12 weight ratio. From there, there are two variations in matrix composition. Pristine gelatin (PG) samples use a composition with unmodified gelatin powder in place of the coagulated gelatin, to explore the influence of the longer polymer chains induced by coagulation. Glycerol-rich gelatin (GG) samples use twice the amount of glycerol in order to examine its influence as a plasticizer and viscosity modifier for the molten state of the hydrogel. For all variations in the matrix composition, versions with silver flakes as conductive filler are also characterized separately (e.g. non-conductive pristine gelatin (PG-NC) vs conductive pristine gelatin (PG-C)). All of PG-C, CG-C, and GG-C achieved high electrical conductivity of over 10^2 S cm^{-1}.

In addition to the hydrogel composition, we also examine the influence of silver flake settling and sedimentation. During the silver flake sedimentation step (Figure 5.3), the microstructure of the material transitions from uniformly dispersed and disconnected silver flakes to a bilayer structure of gelatin-rich and silver-rich layers. Within the sedimented silver-rich layer, the silver flakes form a dense, percolating network that is electrically conductive (Figure 5.3A). The thickness of this layer grows throughout this process (Figure 5.3B-C), resulting in a lower electrical resistance on the sedimented side of the composite (Figure 5.3D), and ultimate, increasing conductivity (Figure 5.2C). There appears to be a point in settling where the transition from the conductive layer to insulating layer is no longer a gradient of decreasing silver flake concentration and is instead a more discrete transition. This stress concentration can cause delamination or fracture as shown in Figure D.6. This mechanical failure mode can be avoided by controlling sedimentation time, and should have some dependency on molten matrix viscosity.

5.2.2 Materials Characterization

We examine the influence of material composition and processing conditions on the properties of the composite. These include mechanical characteristics (elastic modulus, elastic resilience, strain at break), electrical conductivity, and electromechanical coupling between elastic strain and electrical resistivity.

Mechanical Properties

Figure 5.4A presents stress-strain curves under tensile loading for each hydrogel composition with and without silver filler. As can be seen in the strain at break data extracted from these curves (Figure 5.4B) the conductive compositions generally exhibit lower elongation at failure (approximately 30-50% for CG-C and PG-C compared to approximately 85-185% for CG-NC and PG-NC). This reduced stretchability is likely due to the conductive filler introducing microscale defects in the hydrogel matrix that initiate failure. We also observe that the GG samples generally exhibit the highest elongation at break. The excess glycerol

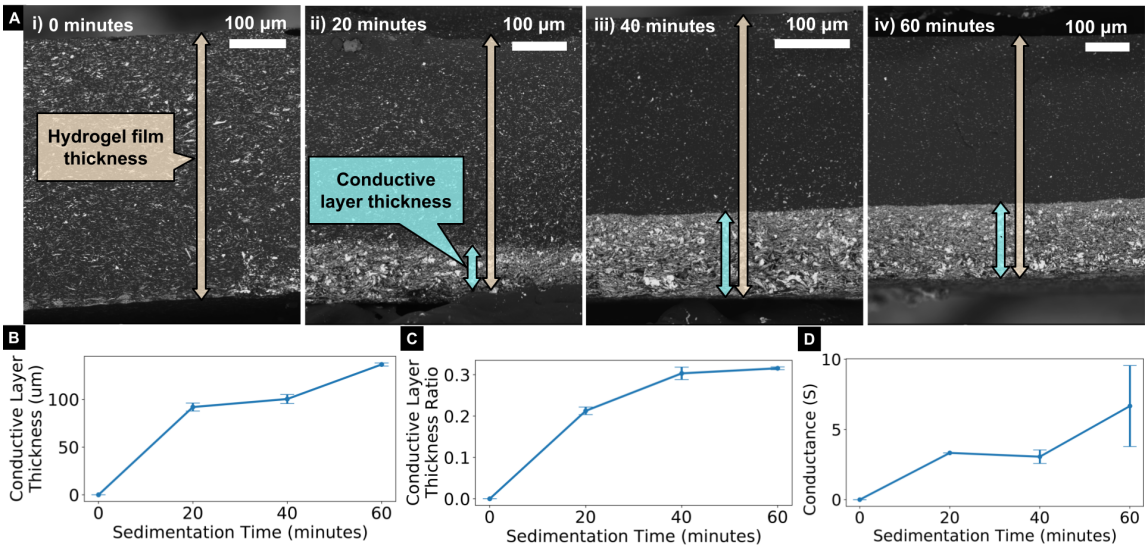

FIGURE 5.3: A) Scanning electron microscopy (SEM) of conductive gelatin composite with a sedimentation time of i) 0 minutes, ii) 20 minutes, iii) 40 minutes and iv) 60 minutes. B) Growth in thickness of the conductive layer as sedimentation time increases. C) Growth in relative thickness of conductive layer as sedimentation time increases. D) Increase in conductance of material samples as sedimentation time increases.

present in these materials acts as an effective plasticizer and creates a much higher elongation before specimen failure, likely due to enhanced mobility of gelatin chains [185]. This effect appears to overwhelm the effect of adding conductive filler, since GG-C and GG-NC have greatly overlapping error bars and very similar mean strains at break (approximately 502% vs. 470%). Finally, it appears that the coagulation process enables larger elongation as well. Comparing CG-NC and PG-NC samples that have no conductive filler, they differ only in that CG-NC uses coagulated gelatin. There is little overlap in their error-bars, with CG-NC having a mean strain at break of 186%, compared to PG-NC with 87%. Coagulation creates longer polymer chains, providing the opportunity for more entanglements that can dissipate energy before bond-breaking [128]. Similar correlations exist when looking at the elastic modulus of the different compositions. As we can see in Figure 5.4C, the presence of conductive filler increases the stiffness of the gelatin hydrogel. This trade-off between high electrical conductivity and compliant mechanical properties appears to be alleviated both by using coagulated gelatin as the matrix polymer and especially by increasing the amount of glycerol plasticizer. Comparing the results of my conductive coagulated gelatin hydrogels to state-of-the-art synthetic hydrogels [180, 186] and gelatin-based hydrogels [175, 176, 178], the GG-C composite is more stretchable (up to 502% strain) than previously reported specimens from either category (Figure 5.5).

Mechanical hysteresis of different compositions under cyclical loading up to 10% strain is shown in Figure 5.6A. Three samples of each are overlaid, to demonstrate consistency within each composition. The first cycle, indicated visually by full saturation, shows a higher stress than subsequent cycles, consistent with stress-softening. Unfilled coagulated gelatin hydrogel with double glycerol GG-NC follows virtually the same path through

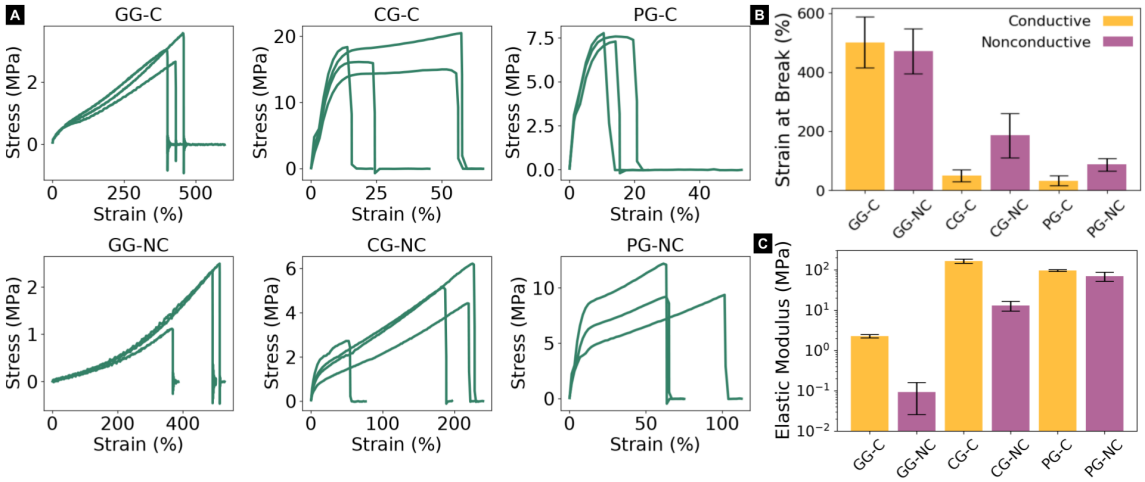

FIGURE 5.4: A) Stress-strain curves of tensile failure testing of different compositions of gelatin composite. B) Strain at break of these compositions. C) Elastic modulus of these compositions.

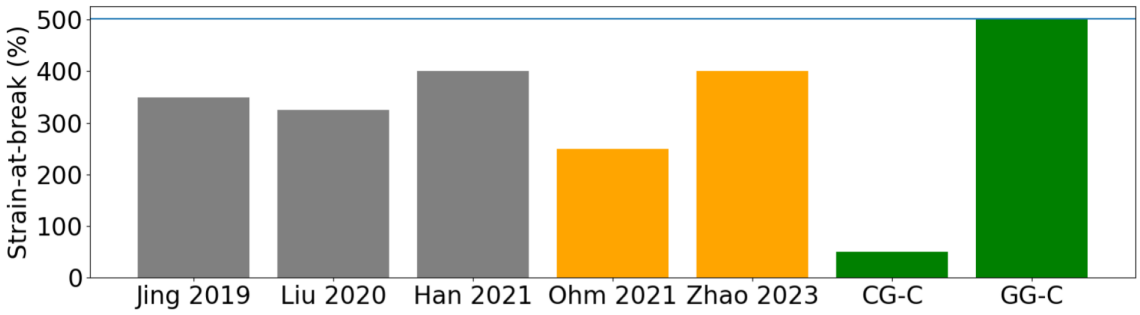

FIGURE 5.5: Comparing stretchability (strain-at-break) of selected compositions from this work (conductive coagulated gelatin (CG-C) and double glycerol conductive coagulated gelatin (GG-C)) to conductive gelatin hydrogels in literature [175, 176, 178] and conductive synthetic hydrogels in literature [180, 186].

stress-strain space after the first cycle, thus displaying minimal hysteresis. Filled coagulated gelatin hydrogel with double glycerol GG-C and unfilled coagulated gelatin hydrogel CG-NC have slightly higher hysteresis, while the filled coagulated gelatin hydrogel CG-C and both filled and unfilled pristine gelatin hydrogels PG-C and PG-NC demonstrate noticeable hysteresis. This is quantified by calculating the average resilience of each composition (Figure 5.6B) across three samples.

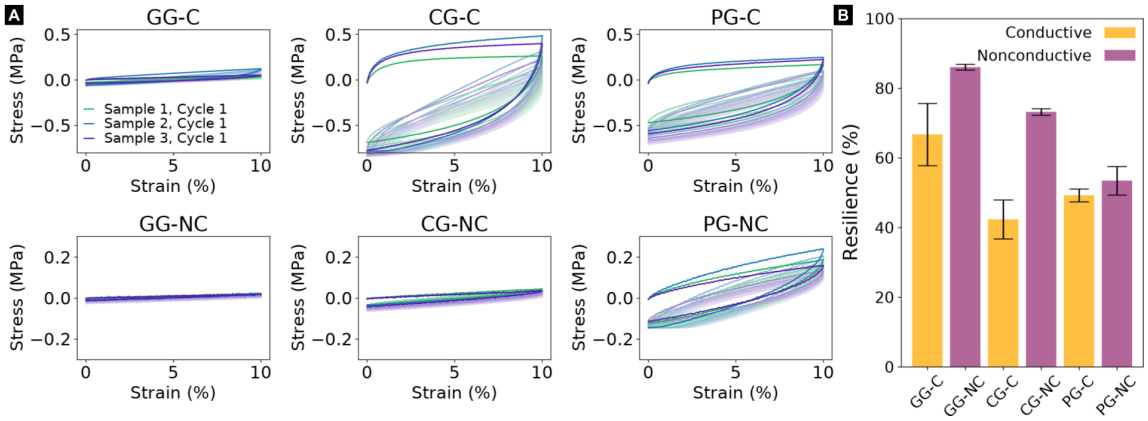

FIGURE 5.6: A) Hysteresis curves of different compositions of gelatin composite. Color saturation decreases and value increases for later cycles. B) Derived mechanical resilience of these compositions.

Elastic resilience for each sample is calculated by first offsetting the values so all stress values are positive, then summing the integrals of the loading stress from 0% to 10% strain for all cycles. This is the energy stored in the material throughout the cyclical process. A sum for the unloading cycles is obtained as well, representing the energy released during unloading. The ratio of the energy released to the energy stored $\frac{U_{r,unloading}}{U_{r,loading}}$ is reported as the resilience of the sample. Nonconductive, unfilled compositions demonstrate higher resilience than their filled counterparts across the board. This is expected due to plasticity induced by rearranging of filler particles during deformation. Unfilled pristine gelatin hydrogel PG-NC demonstrates noticeably less resilience than unfilled coagulated gelatin hydrogel CG-NC. This indicates that the longer polymer chains of coagulated gelatin are better at storing energy elastically than the shorter chains of unmodified gelatin. For the filled versions, CG-C and PG-C, the range of their resilience overlaps, likely because the presence of conductive filler dominates the differences in chain length. The highest resilience filled hydrogel is GG-C, due to the presence of double the amount of plasticizer.

Electrical and Electromechanical Properties

Figure 5.7 shows a comparison of achieved conductivity to literature. Comparing the conductive conductive coagulatd gelatin composites I produced to state-of-the-art gelatin-based conductive hydrogels (0.05 - 0.1 S cm^{-1}) [175, 176, 178], both the GG-C and CG-C composites are at least three orders of magnitude more conductive. Compared to more conductive state-of-the-art synthetic hydrogels, the GG-C composite (381 S cm^{-1}) is comparable in conductivity to the work of Ohm *et al.* (374 S cm^{-1}) [180] and slightly less conductive than the work of Zhao *et al.* (700 S cm^{-1}) [186]. The CG-C composite is more than three times as conductive as either (up to 3110 S cm^{-1}).

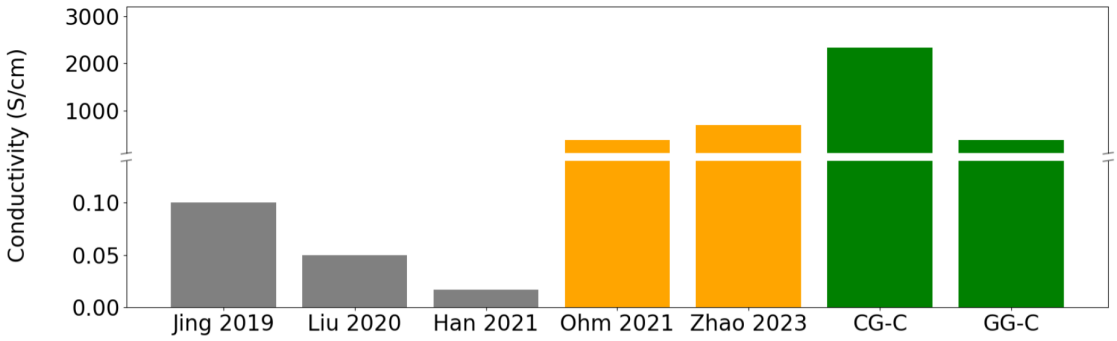

FIGURE 5.7: Comparing conductivity of selected compositions from this work (conductive coagulated gelatin (CG-C) and double glycerol conductive coagulated gelatin (GG-C)) to conductive gelatin hydrogels in literature [175, 176, 178] and conductive synthetic hydrogels in literature [180, 186].

Representative electromechanical response curves for GG-C, CG-C and PG-C hydrogel compositions are shown in Figure 5.8A. These curves were collected from rectangular strips of material using the test rig shown in Figure 5.8B. All compositions demonstrate a coupling between the relative increase in resistance and applied strain. There is a hysteretic element to the increase in resistance as the samples are stretched to higher strains. There is a close alignment of the relaxation trajectory to the stretching trajectory of that cycle. This is an electrical analogue to the stress-softening of filled polymers (Mullins effect [187]), possibly due to changes to the arrangement of filler particles (silver microflakes) when experiencing new maximum strain. This suggests that once stretched to some maximum, the electromechanical coupling would be more consistent within previously experienced strain ranges. Such electromechanical and cyclical behaviors are not unique to this composite and have been previously observed with other soft material systems that use silver as the conductive filler [188].

We note that variations in electromechanical response exist even among samples of the same material composition, as shown in Figure D.9. Such sample-to-sample variation is observed to coincide with differences in the tensile strain at electrical circuit failure. One explanation could be the stochastic nature of material defects within the composites, which manifest themselves in both electromechancial coupling and electrical failure. Aside from these variations, we observe that when comparing between compositions, CG-C and PG-C have similar curves and ultimate strains. GG-C follows a similar pattern, but goes to much higher ultimate strains. This pattern is due to the increased compliance added by having twice as much glycerol as the other compositions.

5.2.3 Demonstrations

The mechanical and electrical properties of this conductive biopolymer composite suggest possible applications in soft and wearable electronics. Here, we present two representative cases in which the material can be placed on the skin and maintain electrical functionality during body motion.

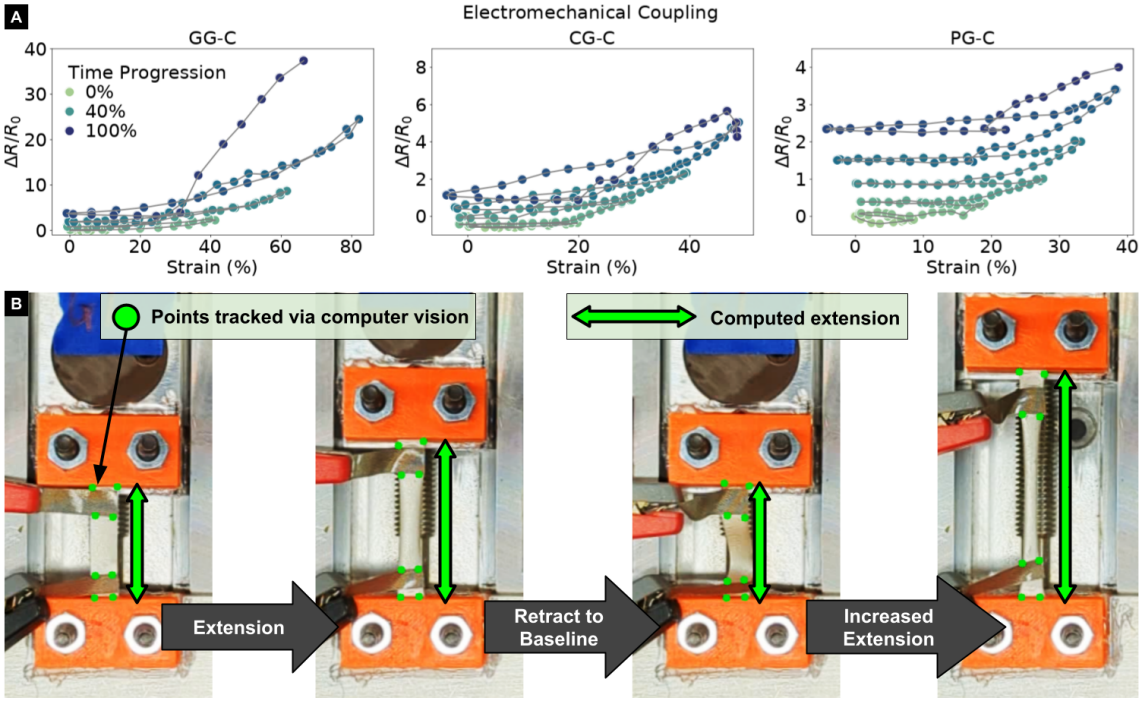

FIGURE 5.8: A) Relative resistance change as strain iteratively increases. B) Experimental setup for electromechanical testing (video available at [189]).

To demonstrate applicability for soft wearable circuits, a rectangular strip of GG-C conductive hydrogel was mounted to the index finger and used for delivering electrical current to a 3-volt LED. Referring to Figure 5.9A, the strip adheres to the finger across the middle and proximal phalanx and remains functional through the full bending range of motion of index finger. Additionally, when detached from the skin, we show that the circuit remains functional under gravitational loading, tensile stretching, and twisting(Figure 5.1C, i-iii).

Furthermore, we demonstrate the ability to use the materials for electromyography (EMG). Generally speaking, electrodes for electrical biosignal recording need to be highly conductive and possesses conformal compliance, features which are present in this conductive gelatin hydrogel. For electromyography, an impedance response curve with sensitivity in the 100Hz to 10kHz range is needed, which is demonstrated by electrodes cut from GG-C composition (Figure 5.9B). This curve was obtained via testing on a pork skin model with a Gamry potentiostat in 2-lead mode (Figure 5.9C). The signals recorded from human muscle activation with attached gelatin electrodes are shown in Figure 5.9D. Electrodes are tested while placed over the flexor digitorum muscle (Figure 5.9D, i-ii) and the extensor digitorum muscle (Figure 5.9D, iii-iv). Recorded voltage traces are visually distinguishable from background, and demonstrate expected features including spikes at both activation and relaxation (video available at [190]), continued activity under movement (Figure 5.9D, iii), and activity extinction while holding position (Figure 5.9D, iv).

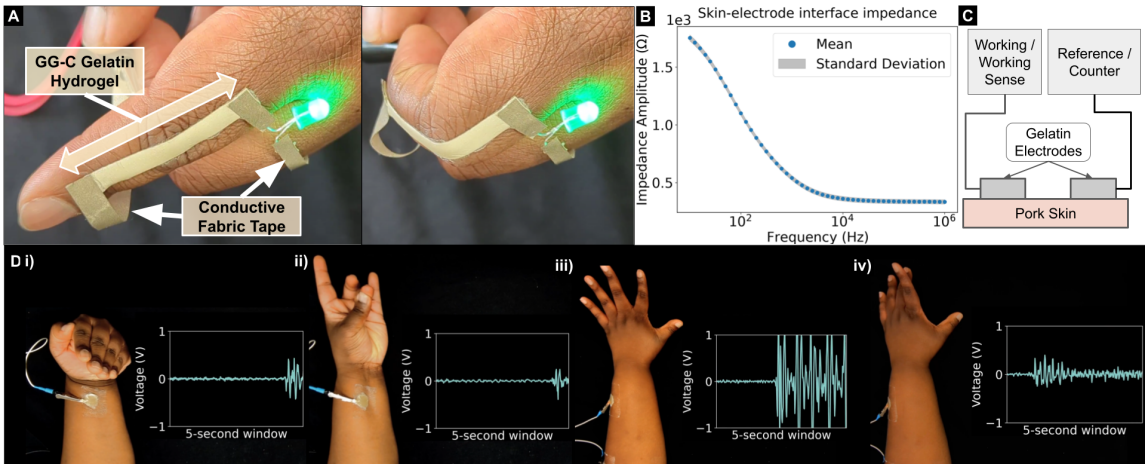

FIGURE 5.9: A) Demonstration of conductive gelatin sample conforming to hand during flexing while functioning as wire. B) Impedance curve of skin-electrode interface in pig-skin characterization experiment (video available at [191]). C) Schematic diagram of impedance analysis setup with Gamry potentiostat. D) Electromyograms of muscle activity during 1) wrist flexion 2) ring-finger flexion 3) wrist and finger extension with rocking 4) finger extension, held till activity extinction. (video available at [190]).

5.3 Discussion & Conclusions

The usage of silver microflakes leaves an opportunity to further increase the sustainability and accessibility of the material system in future work. In a soft electronics context, silver is attractive due to its high conductivity and the relative availability of various form factors that are popular in a multitude of industrial and commercial applications [192, 193]. There are challenges associated with their use as well. Firstly, there are potential respiratory hazards involving breathing particles of this size, and proper ventilation practices during handling and preparation need to be taken (video available at [184]). Research has shown that various morphologies of silver, if released into the environment, pose ecological risks. Silver nanoparticles have potentially useful antimicrobial effects [194], but have been shown to negatively affect the health of fish [195] and aquatic invertebrates [196]. Due to the potential for mechanical or chemical conversion to nanoparticles or silver ions after disposal, life cycle considerations must be made. Carneiro *et al.* demonstrated that soft electronics made with silver flakes and polyurethane composites can be decomposed into polyurethane residue and reusable silver-rich powder via appropriate solvents (IPA) [197]. The gelatin-based composite shown in this work could be dissolved with water as an effective solvent instead of IPA. Reclaiming the silver flakes in this way could allow both recycling of material rather than inciting further resource extraction, as well as preventing the release of silver into aquatic environments. As future work, there is the potential to partially or completely obviate these issues with the use of more biocompatible conductive fillers like graphene oxide [177, 178, 198], PEDOT:PSS [199, 200], zinc [107] or molybdenum [201], and hopefully accessible fabrication methods using these options can be pursued in the future. The structural complexity of components or devices produced from these hydrogels also shows potential for further development, and hydrogel printing techniques [112] potentially provide an avenue to more complex devices fabricated

from these biomaterial hydrogel composites.

Chapter 6

Conclusion

6.1 Contributions

Challenges in Soft Robot Autonomy (Chapter 2): I apply established methods of enabling robot autonomy to soft robotic systems and identify trends and challenges in soft robotic autonomy deriving from intrinsic soft matter physics. I identify collecting hardware data as a critical step and major bottleneck in understanding the effect of soft materials on robotic systems.

Enabling Soft Robotic Hardware Datasets (Chapter 3): I design and implement a research platform capable of extended data collection, sufficient for reinforcement learning. I demonstrate its compatibility with multiple materials and utility in implicitly characterizing materials in the context of a manipulation task.

Gelatin-based Biomaterials (Chapter 4) : I fabricate and characterize a class of hydrogel composites composed of food grade, inexpensive ingredients —gelatin, glucose and glycerol. I demonstrate accessible processes for improving mechanical properties via coagulation, and resistance to water via. thermal crosslinking.I identify water content as an important feature of this biomaterial composite, and manipulate water content of hydrogel to control processability, mechanical properties, susceptibility to thermal degradation.

Accessible Soft Electronics with Conductive Gelatin Composite (Chapter 5): I fabricate conductive composite of gelatin-based hydrogel. I exploit water content, using sedimentation of conductive microparticles in low-viscosity molten hydrogel to achieve high electrical conductivity with moderate conductive filler content. I then improve conductivity and stiffness via controlled dehydration. I demonstrate application as wire in soft electronic circuit and as electrode in electromyography muscle activity recording.

Papers:

*Z.J. Patterson, A.P. Sabelhaus, **K. Chin**, T. Hellebrekers, and C. Majidi. "An untethered brittle star-inspired soft robot for closed-loop underwater locomotion." In 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 8758-8764. IEEE, 2020.* [32]

*Huang, X., Z.J. Patterson, A.P. Sabelhaus, W. Huang, **K. Chin**, Z. Ren, M.K. Jawed, and C. Majidi. "Design and Closed-Loop Motion Planning of an Untethered Swimming Soft Robot Using 2D Discrete Elastic Rods Simulations." In Advanced Intelligent Systems 4, no. 10 (2022)* [33]

***K. Chin**, T. Hellebrekers, and C. Majidi. "Machine Learning for Soft Robotic Sensing and Control." In Advanced Intelligent Systems, vol. 2, no. 6, pp. 1900171, John Wiley & Sons, Ltd, 2020.* [28]

***K. Chin**, C. Majidi and A. Gupta. "1 Modular Parallel Manipulator for Long-Term Soft Robotic Data Collection." arXiv preprint arXiv:2409.03614, 2024.* [99]

***K. Chin**, M. Vinciguerra, D.K. Patel, P.R. Olcay, E.S. Lazaro Vasquez, and C. Majidi. "Accessible soft electronics with silver-gelatin conductive hydrogel composite." In Advanced Materials Technologies, 2024. (Under Review)* [34]

6.2 Outlook

Gelatin hydrogel is attractive as a material for use in soft structures and environmentally sustainable electronics because it can be sourced from agricultural waste. The close proximity of agricultural sources to large population centers is common due to the inherent demand of agricultural products. As policy-makers focus increasingly on climate resilience and short agricultural supply chains fall into greater focus, technology that leverages those sources and provides additional use for waste become increasingly appealing. It is also possible to establish a renewable source of a given biomaterial via the agricultural cultivation of its source organism, so long as an appropriate growing environment is available or can be created. Combined, these factors mean that biomaterial-based systems like the silver-gelatin composites introduced here have the potential to minimize the climate effects of future soft electronic technologies. My hydrogel composites are composed primarily of inexpensive food-grade materials and can be fabricated without expensive equipment, chemically hazardous reagents and restrictive fabrication processes, allowing this material system to be relatively accessible. By focusing on equipment and materials that are purchasable for consumer use cases, this work potentially lowers start-up and ongoing costs, enabling people with less capital to engage with production and study of these materials. By minimizing hazardous material use, people with less training can more easily participate in the creation of these materials. Accessibility to makers, hobbyists, low-resource educational/academic environments and small businesses is a design goal that complements and extends the existing sustainability benefits of biomaterial-based composites. By focusing on accessibility, it becomes possible for this technology to enter into public consciousness directly, without relying on policy change or profitability. Existing efforts in this direction, for example libraries of biomaterial recipes that are open to the public [202], demonstrate demand for this kind of knowledge. Increasing the potential functionality of those materials (e.g. adding conductivity with conductive filler, reducing solubility and improving mechanical robustness) reinforces their utility, contributing to efforts in creating more sustainable electronics [105, 203], and versatile soft biomaterials. Improving access to this technology provides an avenue for broadening education into biomaterial-based technologies, through the process of "citizen science" [204, 205]. This could have downstream improvements to the number of people who study and work

in the field, and in improving popular support for sustainability policy that relies on an understanding of feasible alternatives to fossil-fuel material supply streams. The components of developing soft robotic systems in low resource environments are largely present. There are accessible materials, and options for actuation that allow long-term operation. A new phase of research on integrating these components together is now possible. I hope this thesis demonstrates that this research need not wait for cutting edge methods nor commercial support, and that the power to engage with this technology is available to those that want to build a softer, more sustainable future.

Appendix A

Simulated Nonstationarity

I explore a simulated planning task with very rudimentary dynamics drift and nonstationary adaptation. Next, I explore a more structured simulation environment for studying nonstationarity, with the simple task of model learning and one-step dynamics prediction. Finally, I propose a pipeline for data-efficient, adaptive control of systems with nonstationary dynamics.

Active Sensing Planning with Model Drift

Drift Period	Adapting	Planning Horizon	Execution Horizon	Lifetime Reward
10 steps	No	3 steps	3 steps	183.08
10 steps	Yes	3 steps	3 steps	210.39
10 steps	No	3 steps	1 step	134.76
10 steps	Yes	3 steps	1 step	211.79
10 steps	No	2 steps	2 steps	172.44
10 steps	Yes	2 steps	2 steps	211.00
10 steps	No	2 steps	1 step	135.98
10 steps	Yes	2 steps	1 step	211.34
10 steps	No	1 step	1 step	135.99
10 steps	Yes	1 step	1 step	211.69
20 steps	No	1 step	1 step	109.47
20 steps	Yes	1 step	1 step	187.27
30 steps	No	1 step	1 step	110.87
30 steps	Yes	1 step	1 step	184.92

TABLE A.1: Experimental results for adaptation vs. no adaptation under different planner parameters with continual environment drift.

Overview

Data-driven methods have successfully captured unmodeled behavior of soft systems on short timescales, but begin to fail during lifetime operation due to nonstationarity: changes in the dynamics over time. In this section, I propose a framework for model-based planning for soft robots with nonstationary dynamics and present a sensor coverage task that can be used to evaluate our proposed methods.

Model Similarity	Planning Horizon	Lifetime Reward
100%	5	206.02
100%	3	206.03
100%	1	206.02
99%	1	105.47
95%	1	68.01

TABLE A.2: Effects of constant model error, without continual environment drift.

I evaluate the effect of adaptation under different environmental drift velocities, replanning frequencies, and planning horizon.

Search-based planning relies on having a model of the environment to query internally while forming plans, so it might be sensitive to degradation of this model over time. We posit that interleaving model adaptation with replanning would provide a relatively straightforward approach to dealing with nonstationary dynamics.

Approach

Platform

The platform used for experimentation is a geometric simulation of a real-world compliant planar parallel mechanism, i.e. a compliant 5-bar robot (Figure A.1). An example of a high fidelity simulation of this system would be a discrete elastic rod decomposition [48], relying on mass-spring networks to solve for compliance. However, for the purposes of studying nonstationarity, it is best to start with a small number of simulation parameters. For this purpose, we use the purely kinematic treatment of the system, approximating it as a rigid body 5-bar, parameterized by the five link lengths. However, to incorporate the idea that this is a soft robotic system susceptible to drift, we allow for the link lengths to change. We do a scalar multiplication of the link lengths, keeping the ratios the same, which is an approximation of nonstationary dynamics due to thermal expansion.

Planning Task Specification

To evaluate nonstationarity, we need a task that has a lifetime behavior formulation, in order to allow adaptation to matter. Closed loop adherence to a predifined trajectory would not be a useful metric, nor would a priori planning in the task space allow for the dynamics drift to dominate the performance. For this reason, we use the planning for sensor "coverage" task as defined in Figure A.2 , Figure A.3.

This task is defined by selecting the next action based on a full expansion across a lattice generated by the different primitives of shifting the two servos driving the system. The branching factor is 9.

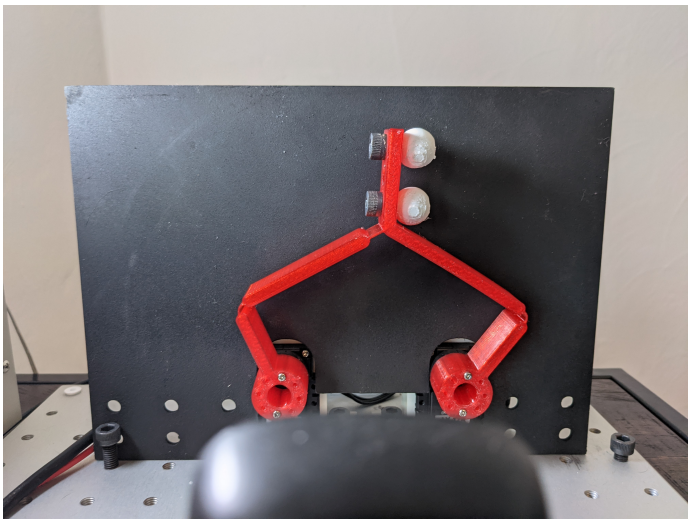

FIGURE A.1: Hardware platform emulated for planning task.

The actual reward behavior in the code is slightly different from the reward spelled out in Figure A.3 however. Instead of returning zero, the reward is attenuated based on how close the state is to all states in the system's memory, as a ratio of the maximum "staleness" of a state. A maximally stale state would be one in which all historical states are within the sensor probe radius of the new state.

Model Adaptation

Adaptation is triggered based on the error between predicted state evolution from the planner's internal dynamics estimate, and the actual state transition shown by the environment. The structure of the system is shown in Figure A.4.

For this project, the model adapter block was simplified to focus on the planning aspect. Adaptation in this case is performed by reducing the distance in parameter space between the planner dynamics and the execution environment dynamics by 50%. This method of adaptation would be impossible for a physical system where the ground truth dynamics are truly unknown. It is used here because of its mathematical properties, i.e., it asymptotically approaches perfect adaptation, wherein the model's parameters are identical to the ground truth. This behavior is an ideal case for us to understand the effect of adaptation while controlling for implementation efficiency.

Experiments

The experiments performed consisted of running the sensor coverage task with variations on the drifting parameters of the execution environment and the adaptation behavior of the planner model. The maximum planning horizon my hardware could handle in a real-time relevant timescale is 3-4 steps. Additionally, the interaction of the planning horizon (how far ahead the planner looked when

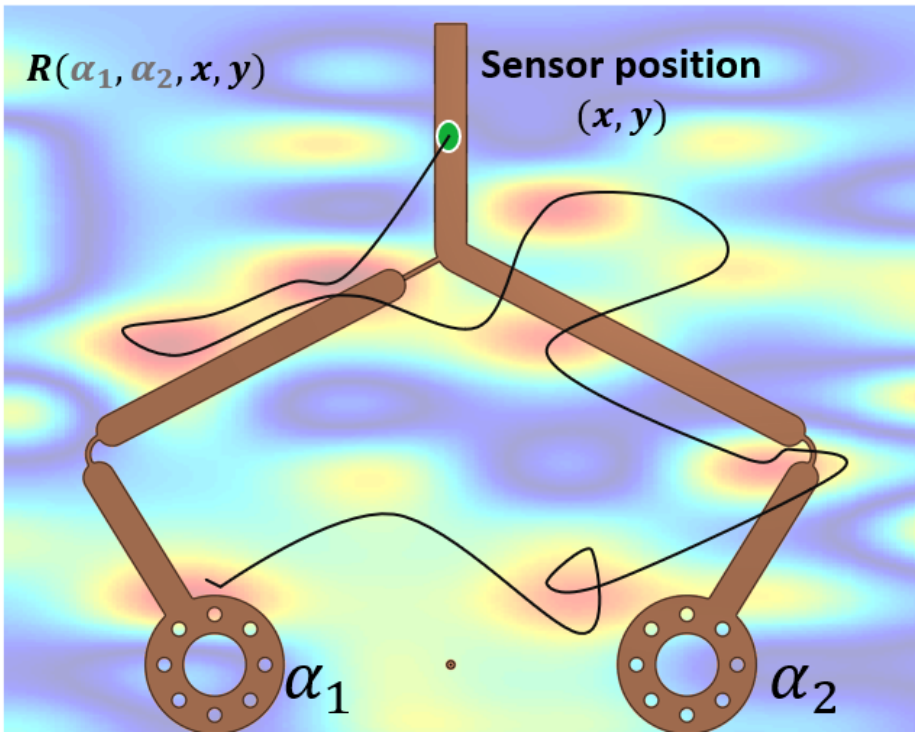

FIGURE A.2: Sensor coverage task description.

State: $s \sim \{ (x, y, \alpha_1, \alpha_2) \}$

Actions: $a \sim \{ (\Delta\alpha_1, \Delta\alpha_2) \}, \quad \Delta\alpha_i \in \{-1, 0, +1\}$

Reward: Pseudo- information gain

$$Reward(s) = \{ R(s), \; s \notin \{s(t-1), s(t-2) \dots s(t_{earliest})\}$$
$$0, \quad otherwise \quad \}$$

FIGURE A.3: Mathematical description of execution environment

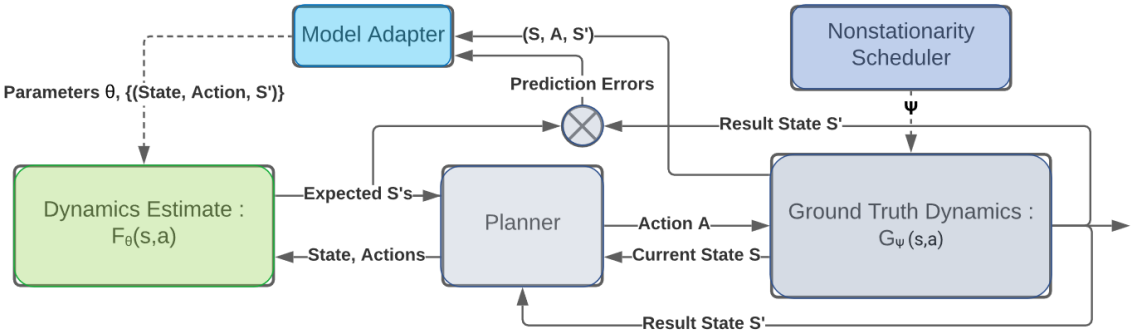

FIGURE A.4: Overview of proposed experimental structure

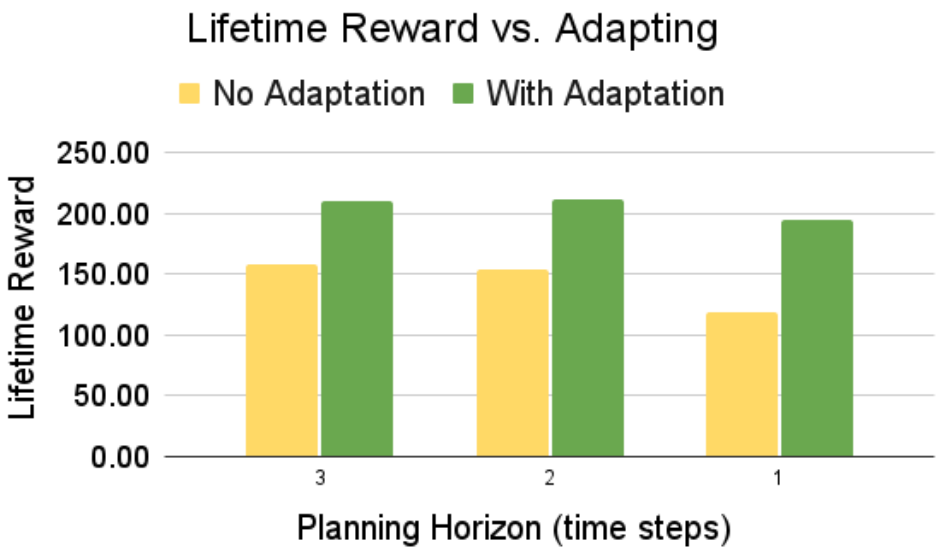

FIGURE A.5: Lifetime rewards for planning with model adaptation vs. without this adaptation. Displayed by the planning horizon used. Columns are averaged from the aggregate data of multiple drift periods, and planning horizons.

choosing an action sequence), the execution horizon (how many actions did the planner execute before replanning) and the model drifting behavior was explored.

Videos

Drifting	Adaptation	Link
None	No	https://youtu.be/R0g6uD0ssCo
Slow	No	https://youtu.be/VREun2EiqJg
Slow	Yes	https://youtu.be/J9L4YRB8ULo
Fast	Yes	https://youtu.be/lAmdXmnMbmA

Conclusions

In all cases, adaptation was helpful in improving the sensor coverage (as measured by Lifetime Reward). This was an expected result, and what is more interesting is examining the interplay between the planner features and the adaptation behavior. Adaptation helps more the more frequently you replan, a bad prediction is worse than moving ahead blindly sometimes. This is reinforced by Table A.2 where we see a steep drop in performance with even slight model errors when the bad model is used over the whole lifetime.

As demonstrated in the videos of execution, this formulation of pseudo-information gain doesn't work very well as the reward map in the context of Dijkstra tree search. The planner ends up stuck patrolling stable loops, even with the penalty

on visiting states in recent memory. This could be alleviated by switching to a non-myopic planner, perhaps using a parameterized target trajectory and attempting to follow it, instead of the forward roll-out method.

We didn't see a strong effect of drift period, which is theoretically an artifact of the unrealistic adaptation behavior which didn't use past experience to adapt. The effect of planning horizon is really only visible in the case without adaptation, and without replanning after every step. I attribute this to partly the fact that our adaptation is too idealized. The adaptation behavior we will implement in the future will be based on parameter estimation + a non-parametric memory based model, using a semi-parametric approach.

Model Learning in Nonstationary Pendulum Environment

Introduction

In this work, I propose a framework for learning the dynamics of systems under nonstationarity, and evaluate its performance on a pendulum environment.

Approach

Pendulum Dynamics

The experimental environment for this work is a simulation of a rigid pendulum with a torque limited motor at the pivot. The simulation is a discrete time state-space system with the following dynamics.

$$\begin{bmatrix}\theta_{k+1}\\ \dot{\theta}_{k+1}\end{bmatrix} = \begin{bmatrix}\theta_k + \dot{\theta}_k \Delta t\\ \dot{\theta}_k + \alpha_k \Delta t\end{bmatrix} \tag{A.1}$$

where θ_k, $\dot{\theta}_k$ are the angular position and velocity of the pendulum at time k. The value of Δt determines the time resolution of the discretization. The angular acceleration at time k α_k is defined with the equation

$$\alpha_k = \frac{3u}{ml^2} - \frac{3g}{2l}\sin(\theta_k + \pi) - b\dot{\theta}_k \tag{A.2}$$

which is parameterized by the pendulum's mass m, length l and damping b. It is also dependent on the force of gravity g, as well as the applied torque at its rotation point, u. For all time steps, the value of u was uniformly sampled from ± 2 N.

Modular Architecture

Because of the multiple cases for nonstationarity, as well as the open question of what components are important to model learning in a nonstationary environment, I decided that my architecture would be constructed out of modules. These modules include the model object, the adaptation scheduler, and the adaptation

mechanism. The setup is very similar to the experimental architecture shown in Figure A.4, with the model serving as the Dynamics Estimate block, and the adaptation scheduler and adaptation mechanism being encapsulated in the Model Adapter block. Rather than having a Planner block, these experiments were performed with an actuation policy which sampled random torques to apply to the pendulum. The nonstationarity scheduler block is preserved unchanged.

Nonstationarity Scheduler

For the pendulum dynamics, the nonstationary behavior is simulated by changing the gravitational force affecting the pendulum over time, i.e. $g = g[k]$. The changes to this force follow a function $\psi[k]$, whose form and frequency is controlled by the nonstationary scheduler. This function is defined of the time horizon of a single trial. The available forms of $\psi[k]$ are the constant, random, oscillating and asymptotic, with an auxiliary flag determining whether non-constant dynamics time-steps were broadcasted for use in supervision. The constant function always returned the same value of 5 $\frac{m}{s^2}$. The random function uses i.i.d. samples from a uniform distribution between 1 and 20 $\frac{m}{s^2}$. The oscillating function oscillated between two values 5 and 15 $\frac{m}{s^2}$. The asymptotic function started at 15 $\frac{m}{s^2}$ and every iteration updated its new result to $0.1g_{max} + 0.9g$, where $g_max = 20\frac{m}{s^2}$.

Model Types

I vary the representation used for the dynamics model. The forms are the black-box model and the parametric physics model.

The black-box model is a neural network function approximator that takes in state-action pairs and outputs subsequent states. This model has no internal representation of the analytic dynamics of the system, and can produce predictions which are incompatible with physical constraints. On the other hand, training can be performed incrementally using constant-time back-propagation calculations. This leads to more flexibility in how training can be set up. Adaptation for this model type can be done by either creating a new neural network and training it with the data relevant to the adaptation, or by feeding the data relevant to adaption into the existing network in order to adjust its weights.

The parametric physics model relies on a first-principles representation of the system dynamics. This means that the performance of this representation relies on the fidelity of the simulation of system dynamics. Training for this representation is defined by solving for the parameters of the system dynamics. If the dynamics are linearizable with relation to these parameters, then training can be performed with linear regression. In general though, this problem is an optimization problem where the objective function is the total prediction error across the relevant dataset. For the pendulum simulation, these system parameters are the mass, damping and length of the pendulum. Adaptation for this representation requires solving this optimization problem over the data relevant to adaptation.

Adaptation Schedule

The adaptation schedule defines when the adaptation step is invoked. The variants of adaptation schedule explored are blind adaptation, supervised adaptation, and detection-based adaptation.

For blind adaptation, adaptation occurs at predefined intervals. This can be every time a new data point is collected, or at given clock-time intervals. In supervised adaptation, adaptation is triggered by an external command. This can be a ground truth knowledge of when the system dynamics has actually drifted, or a human intervention.

Detection-based adaptation is akin to the methods of RL-CD, wherein changes to the dynamics are estimated via some reliability metric evaluated on the existing prediction models. For the pendulum simulation, since the states are continuous, the reliability metric proposed by Da Silva *et al.* [93] is not usable, as that uses finite state calculations. As a replacement, we can use prediction error across the previous N states, where the length of that history and the threshold at which a context shift is triggered are meta-parameters of the learning algorithm.

Adaptation Strategy

The adaptation strategy parameter determines what procedure is performed when an adaptation step is required. It encapsulates the mathematical process by which a new model is created based on new data and, potentially, the existing model. The two strategies I consider are model overwriting and partial models Model overwriting makes the assumption that there is a single correct model which is relevant to the most recent data. This strategy can be trivially implemented if using the neural network model type. In this case, simply training the existing neural model on the small set of new data will cause its internal parameters to shift towards higher prediction performance on that data. After enough new samples, due to the gradient descent update behavior of neural networks, out of date data should theoretically stop affecting the behavior of the current model. For a parametric physics model, the adaptation step should be roughly as slow as training from scratch due to the optimization step jointly solving across all data.

The partial models adaptation strategy involves managing multiple models throughout the experimentation. One of the models is the active model, and only this model's predictions are used. An error metric (the inverse of responsibility) must be tracked for each model. For the pendulum simulation, the error metric is calculated via the rolling average of each models prediction error of the simulations state transitions. The adaptation step happens in two cases. If there is an existing model with an error metric below the error threshold, then adaptation means switching to the existing model with the lowest error. If there is no such model, adaptation means creating a new model and training it on the recent data.

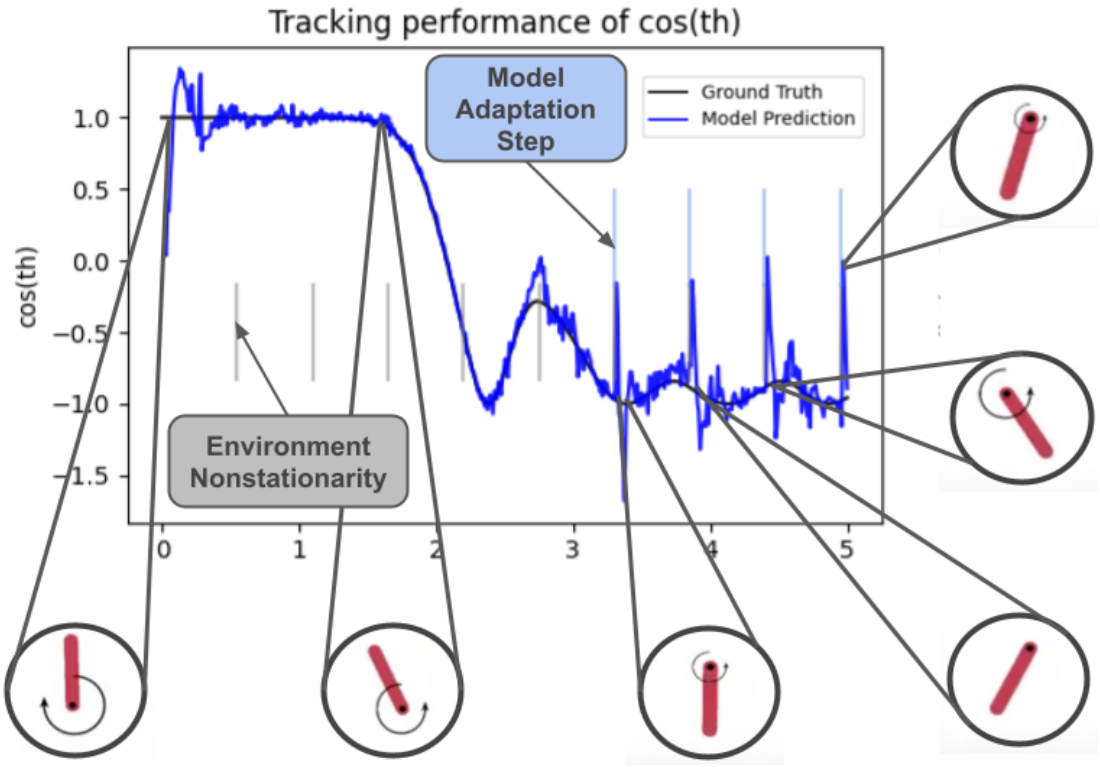

FIGURE A.6: Overview of the structure of a pendulum model learning trial.

Experiments

I performed experiments to gain insight into the interaction of the various parts of the modular architecture, as well as the performance of individual parts.

An experimental trial is structured as follows. The pendulum starts in the inverted position. This can be seen in the first cutout from Figure A.6. The curved arrow in these cutaways represents the direction and magnitude of the torque applied to the pendulum. Every time step, a new torque is uniformly sampled from the -2 N to 2 N. The trial continues for T seconds, swinging down according to the actuated pendulum dynamics. The model predicts the pendulum's next state given the current state and control signal at every time step, where state is broken into $\cos(\theta)$, $\sin(\theta)$, and $\dot{\theta}$, where $\theta = 0$ is the inverted position. Vertical gray lines are shifts to the ground truth dynamics. Vertical blue lines represent adaptation steps. For the parametric physics model (Figure A.7), shifts in ground truth dynamics are shown with vertical blue lines.

Each trial is parameterized by a nonstationarity schedule, a model type, an adaptation schedule, and an adaptation strategy. The ground truth dynamics of Figure A.8 drift according to the random nonstationarity schedule. The dynamics of the parametric model trial, the detection trial (Figure A.9) and the supervised adaptation trial (Figure A.10) drift according to oscillating nonstationarity schedule.

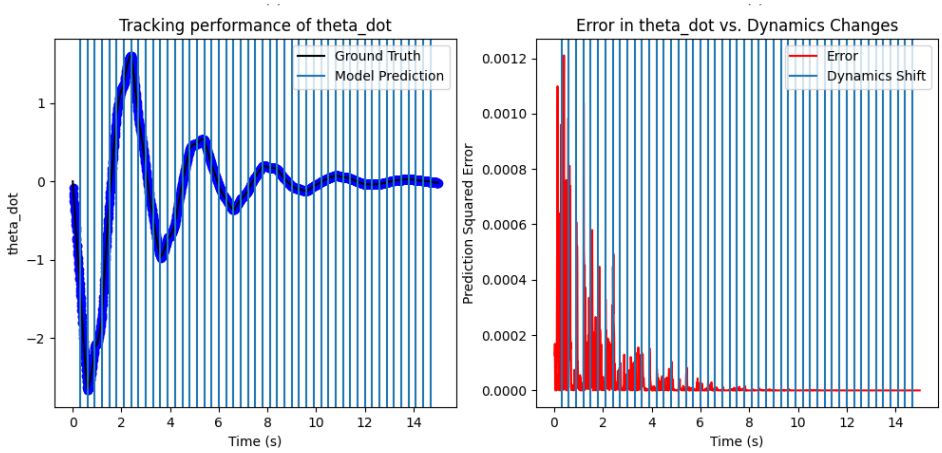

FIGURE A.7: Parametric physics model with oscillating ground truth dynamics.

Results and Discussion

One expectation I had is that those models that have neural network components would be able to achieve lower long-run error for a given parameter set. The parametric physics model, which reflects the knowledge of which system parameters might change, should adapt more quickly relative to the amount of data experienced, but would not necessarily be able to adapt to changes to unmodeled parameters. As shown in Figure A.7, the prediction performance of the parametric physics model even with blind adaptation is nearly perfect. However, this comes at a huge time cost: this trial ran in the order of 400 seconds, as opposed to the 3.5 seconds of all other configurations. This pushed the physics-based adaptation out of the realm of online adaptation, and so it was not explored much further.

For the blind adaptation experiments, e.g. the trial shown in Figure A.8, several configurations collapsed to the same behavior. Firstly, due to the blind nature of the adaptation, there was no difference in performance based on whether drift timing information was available to the algorithm. Additionally, there was no performance gain for experiencing the same dynamics in a periodic fashion, due to the overwriting of any data that ever fell out of relevance to the current dynamics.

The trial shown in Figure A.9 shows the effect of supervision. The way the supervision was implemented did not force a model switch every time the dynamics drifted. Instead the active model changed only if both the time-step contained a drift event, and the active model performed worse than an existing model in the library. All trials of the supervision type kept the same model until the last 40% of the trial. This could indicate that the error threshold for creating a new partial model was too high, or that the partial models did not become significantly different until they had experienced a significant portion of data. Lowering the partial model creation threshold, however, causes the system to create more partial models than the number of dynamics modes. This behavior can be seen in

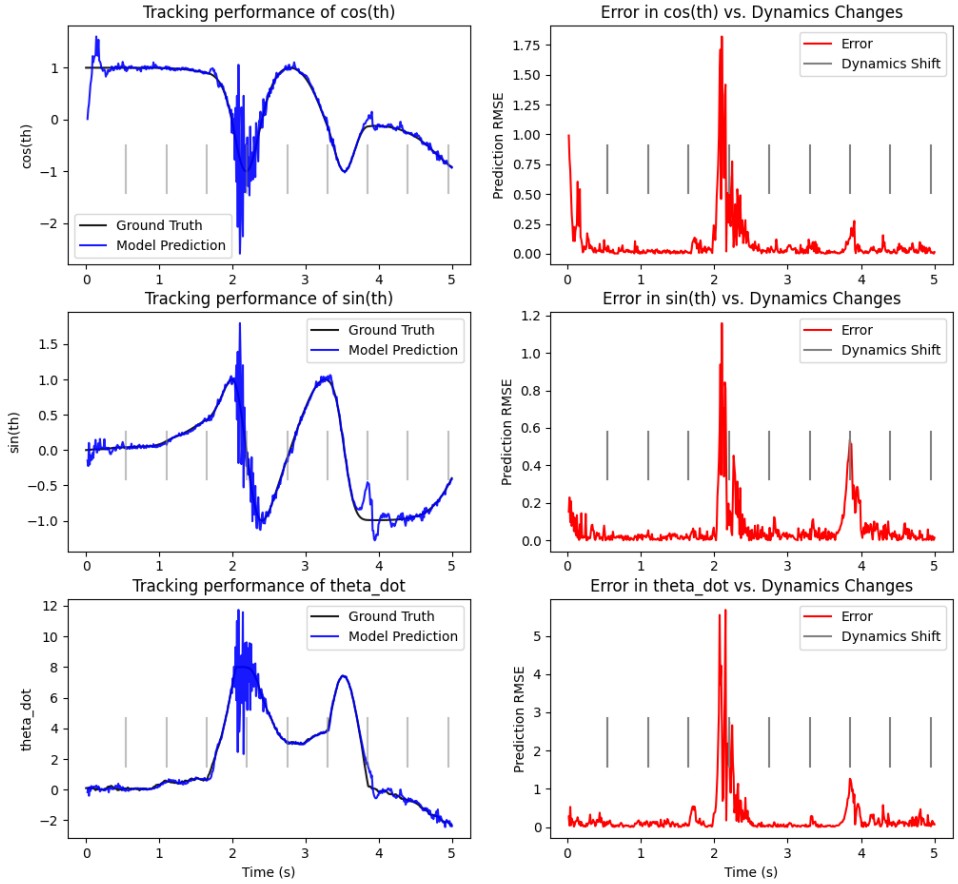

FIGURE A.8: Neural network model with blind adaptation schedule and model overwriting adaptation strategy.

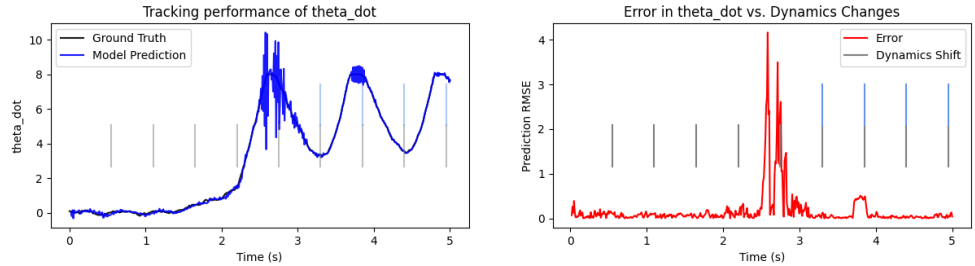

FIGURE A.9: Neural network model with supervised adaptation schedule and partial models adaptation strategy.

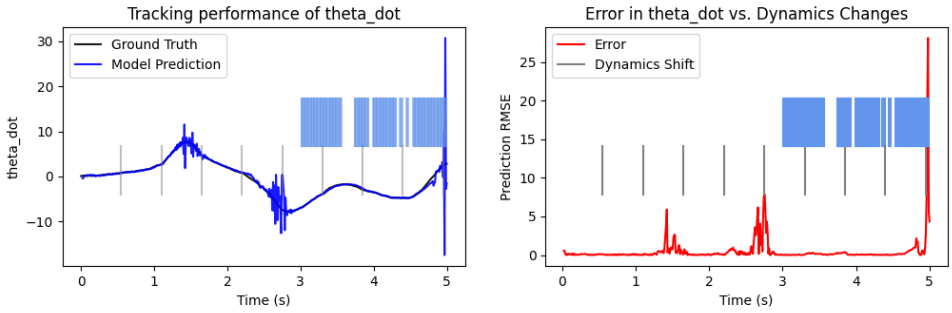

FIGURE A.10: Neural network model with detection-based adaptation schedule and partial models adaptation strategy.

greater intensity in the detection-based adaptation trials, e.g. Figure A.10.

Conclusion and Future Work

The main conclusion reached by this experimentation is that the random policy pendulum is a poor experimental task for studying nonstationarity. One problem is the lack of a task-oriented performance metric. Mean squared prediction error can only function as a comparative metric, and makes it difficult to meaningfully assign success. Without a task with a real fail-state, this hard to interpret metric remains.

The lack of differentiation between configurations and the relatively high performance of the blind, model-overwriting configuration on even oscillating dynamics indicates that this task is also too simple. Though the aforementioned problems with the performance metric make it difficult to be sure, it seems clear that no real techniques for addressing nonstationarity is necessary for solving the nonstationary characteristics of this task, beyond continual learning. Part of this is probably due to the fact that the changes induced are too subtle; changes to 'gravity' are unlikely to affect overall short-term behavior beyond slight numerical adjustments. This idea is supported by the fact that this was the case when damping was used as the nonstationary factor, except in the most drastic of circumstances. Changes to other factors during a single trial are not possible however, due to the non-physical energy injections and the resulting instabilities.

Future work involves finding a task with characteristics that address these issues. In particular, a task that requires the development of a policy would solve the performance interpretation problem. Likewise, a task that can reasonably be assumed to only drift between episodes would reduce the time burden which stymied development of the physics model and prevented more distinctive dynamics changes of the pendulum parameters would allow those hurdles to be conquered.

Approach	Multiple Dynamics	Online Data	Analytical Parameterization	Precomputation
MMRL	Limited	Yes	No	No
RL-CD	Yes	Yes	No	No
HRL-CD	Yes	Yes	No	Yes
ADR	Yes	No	No	Intense
CLAC	Yes	Yes	Required	No
Kalman	Yes	Yes	Required	No

TABLE A.3: Overview of approaches to controlling robotic systems with potentially nonstationary dynamics

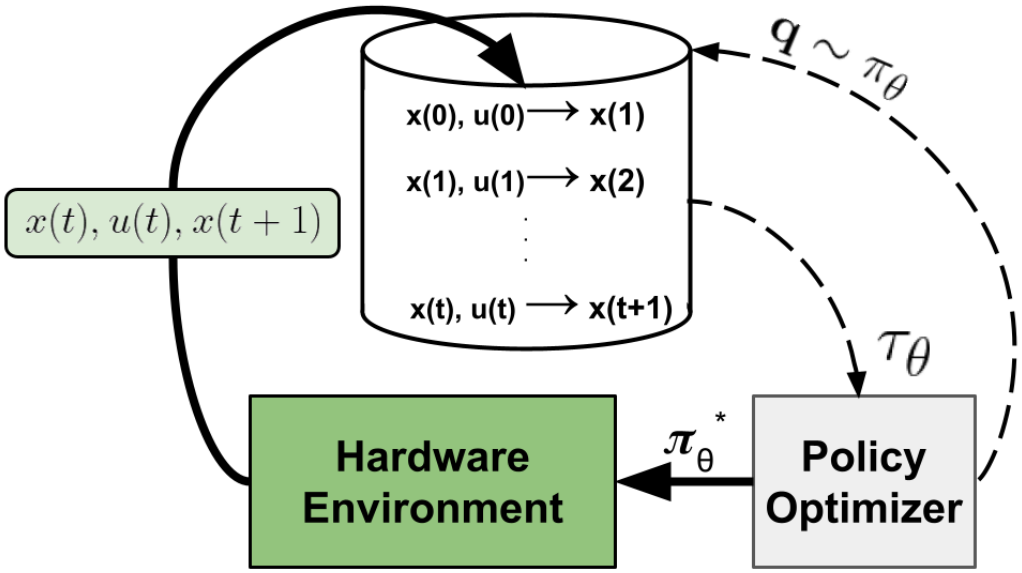

FIGURE A.11: For the five-bar pendulum, $x(t)$ is minimally represented by $\{\theta_1, \theta_2, \theta_3, \dot{\theta}_3\}$. $u(t)$ would be the torques of the servos over one time step,$\{T_1, T2\}$. Queries q are composed of the current x, u being considered by the policy optimizer. τ_θ is the trajectory of outputs to queries put to the model sequentially.

Memory based adaptation

I propose a framework for applying memory-based learning to the nonstationarity problem, shown in Figure A.11. The framework is composed of the memory-based forward model and the policy learner. The memory-based forward model accumulates data from hardware trials, and its active data represents the current model of the hardware. The policy run on the hardware is derived from an inner reinforcement learning loop interacting with the memory-based model as a simulator. This is a model-based reinforcement learning formulation. The two primary components are the memory-based forward model of the system, and the parametric policy learner. Atkeson et. al show that locally weighted models have distinct advantages over neural networks with regard to interference [206]. For

the policy learner, we will select from several existing methods based on compatibility and performance. Memory-based methods are susceptible to high dimensionality of data, due to their inference step. Querying a memory-based model involves calculating a weight for points in the neighborhood of the query. Evaluating which points are in the neighborhood involves either calculating a distance to every point in the memory bank, or solving a k-nearest neighbors problem. Both of these approaches scale poorly with the dimensionality of the problem, so reducing the dimensionality of our action representation via primitives provides an opportunity for performance gain. However, there is a tension here, as adding states to our representation might allow us to provide hints to the adaptation process, e.g. by adding the time the data was collected. Feature scaling in our distance function would allow for us to use the augmented states to filter the relevance of data to adaptation hierarchically with the relevance to location in state space [206]. Resolving the tension between augmenting the state space and dimensionality reduction will depend on balancing computational costs with the costs of collecting more hardware data due to slower adaptation.

There are three observations that motivate extensions to the general memory-based reinforcement learning structure. The first is that data are not collected from noisy dynamics, but rather from temporally coherent nonstationary dynamics. The second observation is that there is no precomputation on the data, nor are we storing the result of any computations performed on the data long-term. Finally, we note that for most robotic systems, we know more about the general structure of their dynamics than simply the representation of their state and actuation spaces. It should be possible to leverage temporal coherence by performing extra computation on our stored data in the form of temporal clustering. If we use hard assignment, then we could add to each data point a flag which indicates the cluster we believe it belongs to. We can use this flag as an augmentation to the state space as described above. We could also divide our memory bank into multiple models using this, effectively lowering the computational cost by only considering data in the same cluster. This would have the downside of requiring a module for switching clusters and triggering re-clustering however. For soft cluster assignment, the analogous approach would be modifying the weights of data based on cluster assignment. For both cluster switching and cluster-based weighting, we would need to calculate cluster statistics and parameters to facilitate insight into a cluster's contents without iterating through it every time. Let's assume we have domain knowledge in the form of a low-fidelity parametric approximation of our system. We could estimate which parameter set best explains the data in a temporal cluster, providing an N-dimensional augmentation to our state. The other use for such a low-fidelity approximation would be using it as the basis for the model prediction at every query, and changing the data to represent a correction to the low-fidelity prediction.

Appendix B

Reinforcement Learning Additional Data

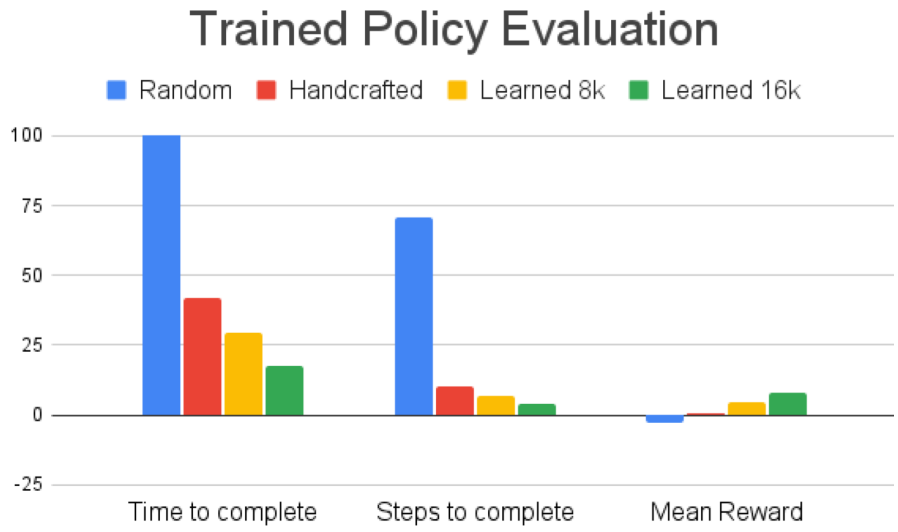

FIGURE B.1: Performance comparison between a random actuation policy, a hand-crafted policy, and policies learned for 8k+ and 16k+ time steps.

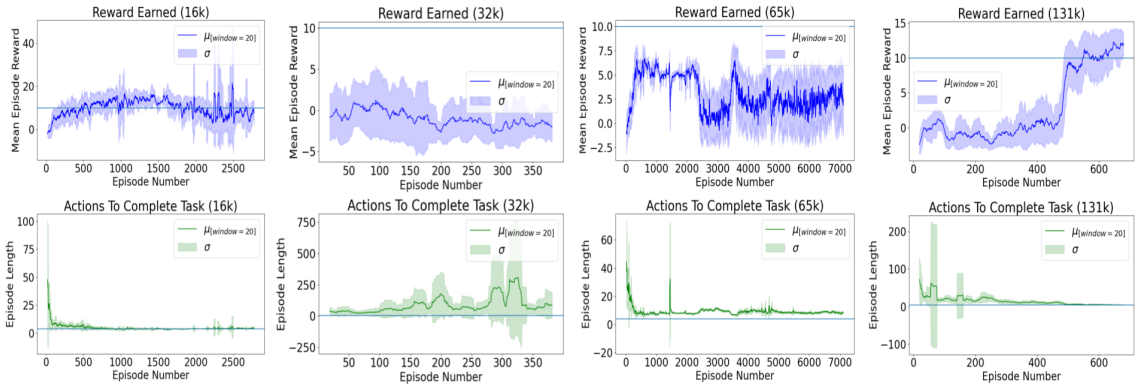

FIGURE B.2: Training results of increasing data collection sizes. Note, 131k session terminated early at 10k time-steps.

Appendix C

Gelatin Biomaterials Experimental Methods

Materials

Hydrogel composite composed of the following: gelatin powder (300 Bloom Type A Porcine, Superclear, Custom Collagen), glucose syrup (DE 45-49, Pastry 1) and glycerol (vegetable glycerin, SimpleNature). For conductive compositions, additionally used silver microflakes (99.9% pure, APS 2-5 microns, 47MR-10F, Inframat Advanced Materials) and concentrated dishwashing liquid (Ultra, Dawn). For coagulation of gelatin in CG and GG variants, isopropyl alcohol (Technical Grade 99%, R-615-1, PTI Process Chemicals) is used.

Hydrogel Fabrication

Gelatin Coagulation

The gelatin coagulation process is based on Stoessel et. al [30] as shown in Supplemental Information Video S3. Distilled water (4 grams / 40 wt%) is mixed with isopropyl alcohol (5 grams / 50 wt%). Gelatin powder (1 gram / 10 wt%) is added to solution. Due to the high solubility of gelatin in water and its low solubility in alcohol, a ternary solution with two phases is formed. Vial with solution is placed in water bath, kept at 50° C using an induction cooktop (COS-YLIC1, Cosmo) and allowed to rest undisturbed for 5 minutes. After resting time, the solution is shaken by hand to agitate for 30 seconds. This cycle of resting and agitation is repeated for 30 minutes, increasing the length of gelatin chains over time. Afterwards, the supernatant phase is decanted, leaving behind a dope phase composed of coagulated gelatin and residual absorbed water/isopropyl for use in hydrogel synthesis.

Silver Flake Preparation

Silver microflakes are prepared under fume hood ventilation as shown in Supplemental Information Video S4. Distilled water (3 grams) is mixed with concentrated dishwashing liquid (0.2 grams) as a surfactant in a scintillation vial to

produce a soapy solution. Silver microflakes (3 grams) are dispersed in soapy solution and shaken to ensure all flakes are wetted to prevent aerosolization during subsequent hydrogel synthesis.

Hydrogel Synthesis

Hydrogel composites are prepared according to process shown in Supplemental Information Video S2. For PG-C composition, in a scintillation vial glucose syrup (1 gram) is mixed with glycerol (1 gram), distilled water (9 grams) and a suspension of silver filler particles (3.2 grams soapy solution, 3 grams silver microflakes). Mixture is heated in a water bath to 50°C to lower the viscosity of the glucose syrup, then agitated by hand to mix. Gelatin powder is slowly added to allow granules to hydrate without clumping, producing a 1:1:1:12 weight ratio of gelatin powder, glycerol, glucose syrup and water. Once all gelatin powder is hydrated, the sealed vial is placed in a water bath at 80°C for 1 hour, until molten. For CG-C composition, coagulated gelatin (1 gram, plus absorbed water and isopropyl alcohol) is added instead of gelatin powder, producing a 1:1:1:12 weight ratio of coagulated gelatin, glycerol, glucose powder and water. For GG-C composition, coagulated gelatin (1 gram, plus absorbed water and isopropyl alcohol) is also used, and an additional 1 gram of glycerol is added, producing a 1:2:1:12 weight ratio of coagulated gelatin, glycerol, glucose syrup and water. PG-NC, CG-NC and GG-NC nonconducting variants are synthesized similarly to their respective conductive counterparts by replacing the silver microflake suspension with distilled water (3 grams), and then skipping the subsequent sedimentation step. Larger sheets of material are produced by multiplying ingredient quantities; sheets used for consistent comparison in this work were fabricated with 16x the reference quantities and poured onto aluminum or silicone sheets to cast (Figure D.7).

Silver Flake Sedimentation

Silver flakes are assembled into a conducting layer via sedimentation in the molten hydrogel matrix. Hydrogel is kept above its composition-dependent melting point of around 85° C via double-boiler method over an induction cooktop (COS-YLIC1, Cosmo) as shown in Supplemental Information Video S2. Due to presence of surfactant, flakes do not remain floating on the top surface, and all microflakes sink over time under force of gravity due to higher filler density than matrix density. This process is illustrated in Figure 5.2C and Figure 5.3A. Hydrogel is kept in this molten state for 1 hour, then allowed to cool, resulting in sheets of conductive hydrogel composite.

Hydrogel Dehydration

After hydrogel samples cool down enough to gel, they are dehydrated for 48 hours in a food dehydrator (CP267-FD, Cosori) at 35° C, kept below 10% humidity as tracked by a humidity monitor (1731, Taylor) inside the dehydration chamber.

Excess water is driven off, with stable amounts retained by the humectant properties of glycerol.

Material Characterization

Electrical Characterization

Conductivity measurements were performed using a digital multimeter (Fluke 177 True-RMS Digital Multimeter, Fluke) in 2-probe mode. Contact resistance between probe and material was reduced by immersion in droplets of gallium-based liquid metal, produced in house from 75% Gallium (99.99%, Rotometals) and 25% Indium (99.99% Rotometals). Rectangular strips of conductive hydrogel material are cut from sheets and dimensions are measured using digital calipers (6 in. 3-Mode Digital Fractional Caliper, Husky), three times for each dimension with median measurement used. Conductivity calculated according to $\sigma = \frac{l}{R*w*h}$, where R is resistance in Ohms across the strip, minus contact resistance, and l, w, h are the length, width and thickness of the sample.

Tensile Failure

Data for strain-at-break and elastic modulus presented in Figure 5.4B, 5.4C was collected via a materials universal testing machine (5969, Instron) with a 50N load cell (2530-50N, Instron). Samples of composite were prepared as dogbone coupons (Die C, 1/2 scale, ASTM D412) and secured with manual wedge action grips (2716-016, Instron). Samples were strained at 3 mm/min rate until failure. Experiments were repeated 3-4 times per composition.

Mechanical Hysteresis

Mechanical hysteresis measurements were taken using materials universal testing machine (5969, Instron) with 50N load cell (2530-50N, Instron). Samples of composite were prepared as dogbone coupons (Die C, 1/2 scale, ASTM D412) and secured with manual wedge action grips (2716-016, Instron). A strain curve was programmed at 10 repetitions, with 5 mm/min strain rate and 0-10% strain range. All samples experienced slack removal via pre-stress of less than 1N. Experiments were repeated 3 times per composition.

Electromechanical Characterization

Electromechanical characterization was performed on the setup as shown in Figure 5.8B. Rectangular coupons of conductive hydrogel composite were cut from a sheet and dimensions were recorded for strain normalization. Coupon is clamped to manual linear stage (Unislide A25 Series, Velmex) using textured clamping blocks printed on a fused deposition 3D printer (MK3S, Prusa) from PLA (Prusament, Prusa). Double sided silver conductive tape (Gray-T-C-25-20, JUFU) was bonded to the coupon at both ends using silver epoxy (8331D, MG Chemicals), and probes were clamped to this tape to connect to a digital multimeter (34401A,

Agilent) for continuous monitoring. During testing (shown in Supplemental Information Video S5), coupon is stretched by hand-cranking the linear stage, then returning it to the initial position. Each cycle, the amount stretched is increased by one revolution for PG-C and CG-C and two revolutions for GG-C (0.1 inches advanced per revolution). Test is monitored optically by video recording using an overhead mounted smartphone camera (Pixel 6, Google). Bounding boxes are selected around visual intersection of the coupon and clamping block, and strain over time is calculated using computer vision box-tracking with OpenCV Python library.

Scanning Electron Microscopy

Scanning electron micrscope images collected from samples after 48 hour-dehydration, using a benchtop scanning electron microscope (PW-100-018, Phenom XL, Phenom). Settings for the microscope were 10kV beam strength with 0.1 Pa vacuum pressure.

Demonstrations

Wearable soft circuit

A soft circuit is constructed with an LED (20mA,3-3.2V, 60F5W-YT-6SE-6SE, CHANZON) as the load. LED is connected to double sided silver conductive tape (Gray-T-C-25-20, JUFU) using silver epoxy (8331D, MG Chemicals), which is then epoxied to conductive hydrogel composite GG-C. Hydrogel composite is placed across index finger joint as shown in Figure 5.9A, or wrist joint as shown in Supplemental Information Video S6, conformally attached using a medical transfer adhesive (1524A, 3M). Joint is flexed repeatedly while current is carried through conductive composite, demonstrating a stable percolating network.

Impedance Characterization

Electrodes are cut from conductive hydrogel sheet (GG-C) using a 10cm diameter stencil. Electrodes are attached conformally to the pork skin substrate with tape (539H, 3M), and connected to a potentiostat (Interface 1000, Gamry) in 2-probe mode as shown in Figure 5.9C. Electrodes placement distance is 1.5x the electrode diameter of 10cm. Impedance curve is measured 10 times to produce the mean and standard deviation curves shown in Figure 5.9B,

Electromyography electrodes

Two conductive hydrogel electrodes are placed across relevant muscle groups, extensor digitorum for extension and flexor digitorum for flexion, with placement derived from Atlas of Muscle Innervation Zones [207]. Electrodes are attached conformally using medical transfer adhesive (1524A, 3M). Disposable self-adhesive clinical grade electrodes (HEX Dual Electrodes #272S, Noraxon) are used

as reference electrodes, placed near the olecranon prominence of the elbow (electrodes shown in Supplemental Information Figure S4). Both the hydrogel electrodes and reference electrodes are connected to an open source biosensing board (Cyton V3, OpenBCI), which communicates to a neighber computer through a wireless USB receiver. During demonstration, activity from finger and wrist motion is recorded through hydrogel electrodes as shown in Figure 5.9D and Supplemental Information Video S7. Recorded data is filtered using a Butterworth bandpass filter from 5Hz-50Hz followed by notch filter at 50Hz and 60Hz.

Appendix D

Gelatin Biomaterials Supplemental Figures

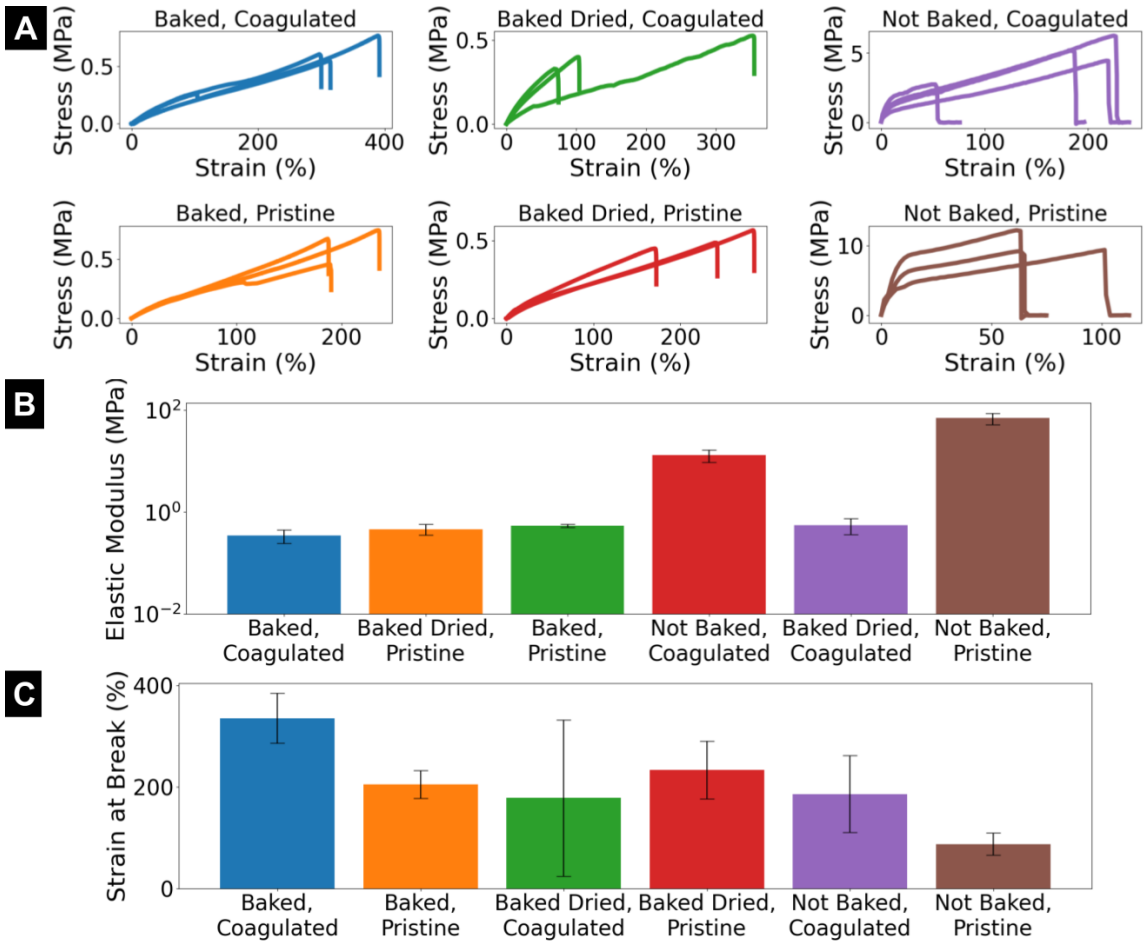

FIGURE D.1: A) Stress-strain curves for tensile failure testing of different baking conditions of gelatin hydrogels. B) Elastic modulus of different baking conditions. C) Strain-at-break of different baking conditions.

FIGURE D.2: A) Oleogel and its ingredients. B) Emulsifier, molten oleogel, and coagulated gelatin hydrogel. C) Combined bigel, flexible and slightly stretchable.

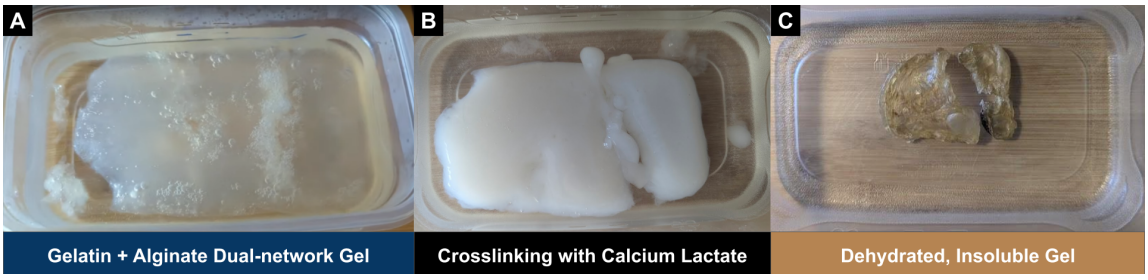

FIGURE D.3: A) Gelatin-alginate composite hydrogel. B) Color change due to crosslinking. C) Significant dimensional and morphological changes under dehydration

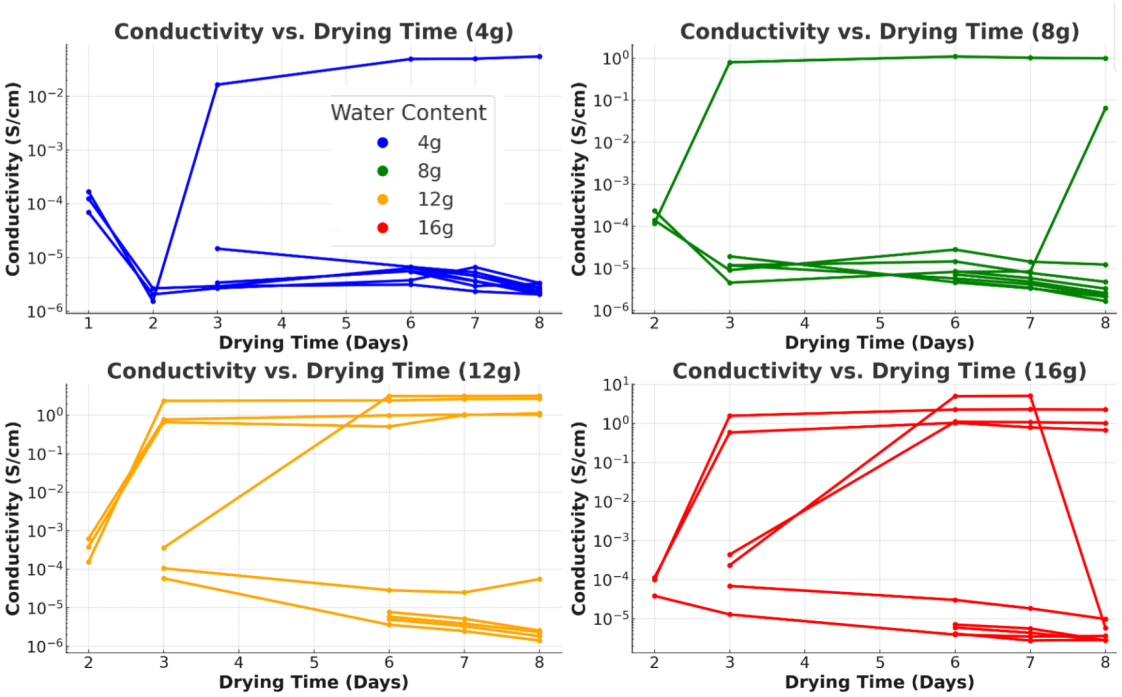

FIGURE D.4: Dehydration-induced conductivity in hydrogel samples with different amounts of initial water. Ionic conductivity decreases as water evaporates by reducing mobility of ions through the gel. Electrical conductivity increases as water evaporates, as conductive filler is brought closer together as gel shrinks. There is a threshhold effect, as this does not happen until a percolating network is able to be formed by the conductive flakes.

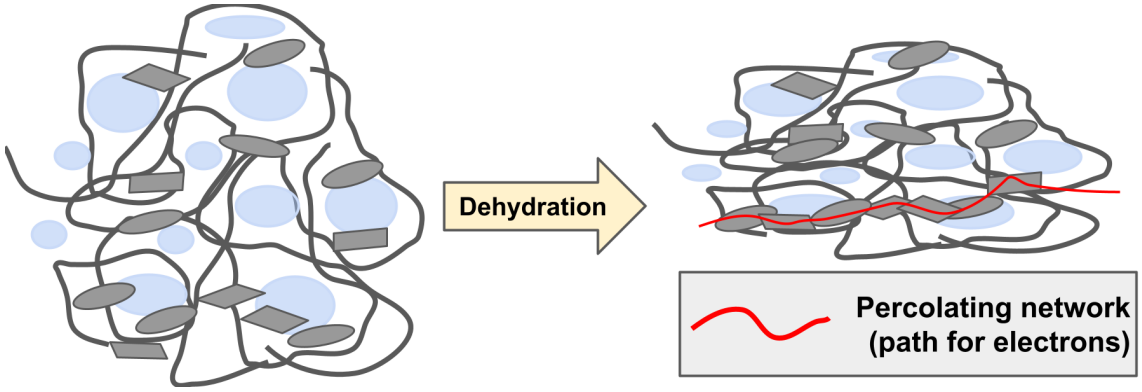

FIGURE D.5: Illustration of percolating network forming due to dehydration-induced shrinkage of hydrogel matrix.

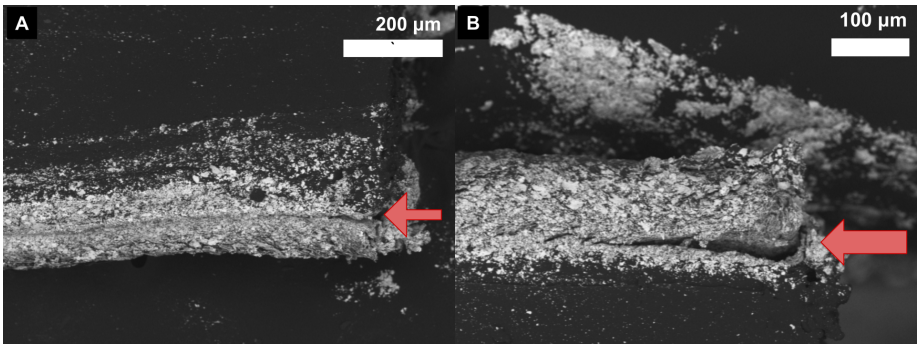

FIGURE D.6: Coagulated gelatin sample displaying fracture between conductive layer and rest of matrix.

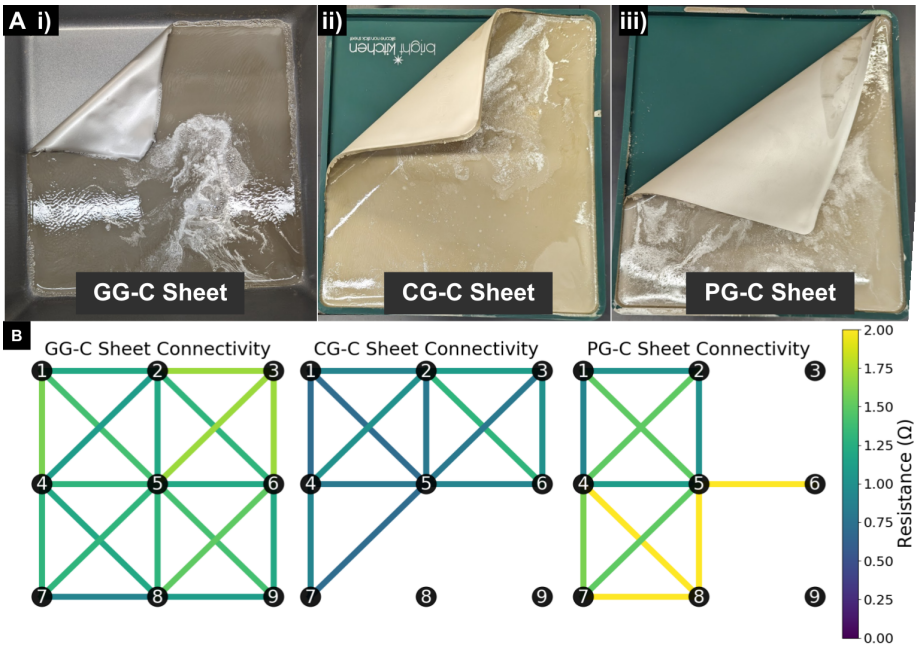

FIGURE D.7: A) Sheets of silver-gelatin conductive composites: i) high glycerol composition, ii) coagulated gelatin composition, iii) pristine gelatin composition. B) Point to point electrical connectivity of conductive sheets, showing process variation. Each sheet has 9 reference points chosen in a grid, spaced by 120 mm / 130 mm (horizontal / vertical) for CG-C and PG-C and by 85 mm in both directions for GG-C.

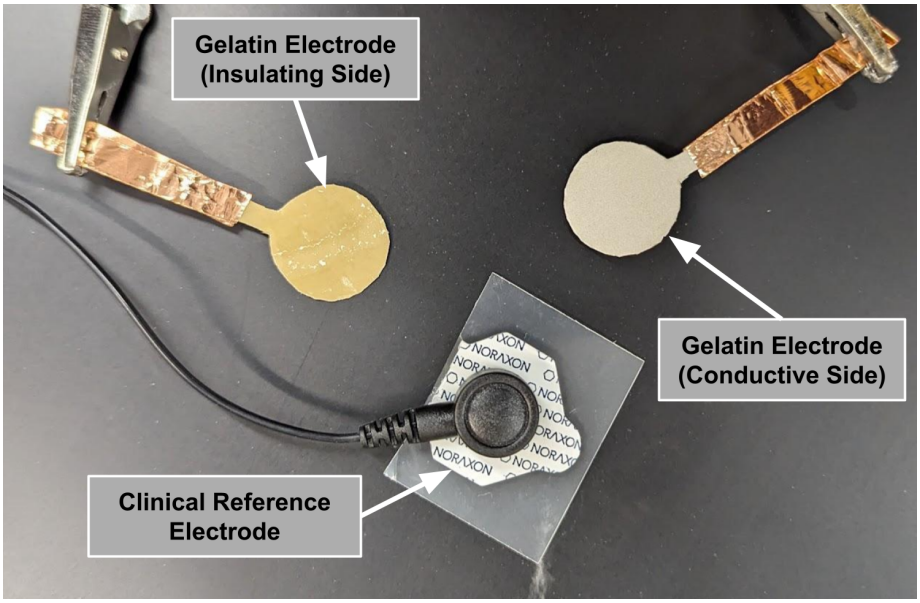

FIGURE D.8: Electrode setup for electromyography. Silver-gelatin conductive electrodes used for recording, clinical-grade electrode used as reference electrode.

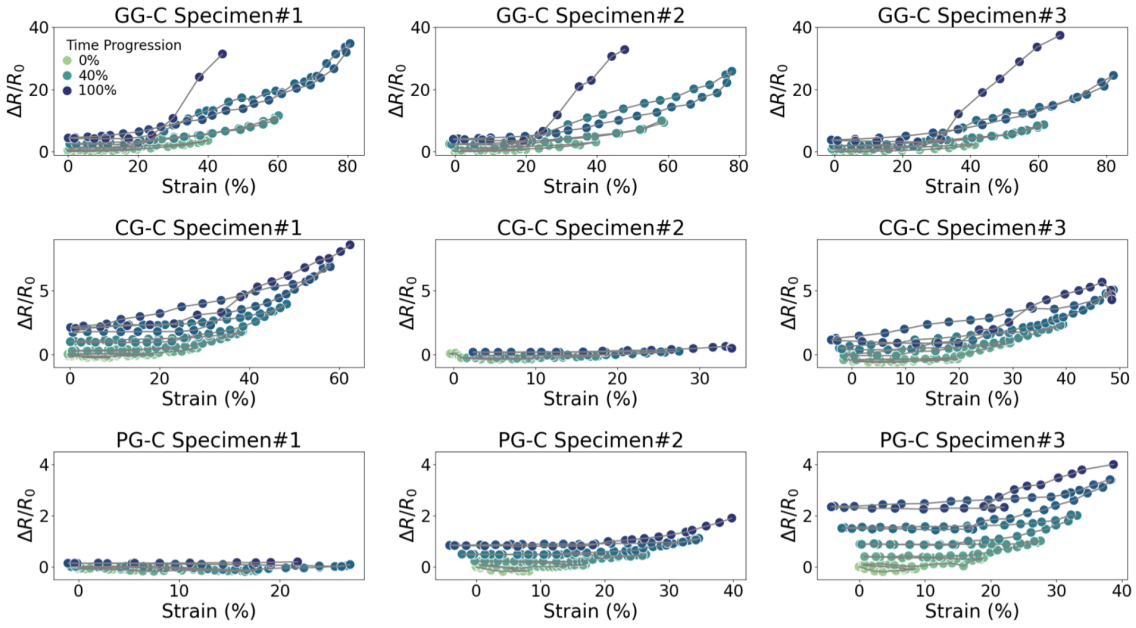

FIGURE D.9: Comparison of all electromechanical experiments for GG-C, CG-C and PG-C specimens. CG-C and PG-C have specimens which demonstrate notable difference from shape. This correlates with lower ultimate strain, potentially indicating early mechanical failure rather than variation in electromechanical properties.

Bibliography

[1] Langdon Winner. "Do Artifacts Have Politics?" In: *Computer Ethics*. Routledge, May 2017, 177–192. ISBN: 9781315259697. DOI: 10.4324/9781315259697-21. URL: http://dx.doi.org/10.4324/9781315259697-21.

[2] Juan David Reina-Rozo. "Art, Energy and Technology: the Solarpunk Movement". In: *International Journal of Engineering, Social Justice, and Peace* 8.1 (Mar. 2021), 55–68. ISSN: 1927-9434. DOI: 10.24908/ijesjp.v8i1.14292. URL: http://dx.doi.org/10.24908/ijesjp.v8i1.14292.

[3] Brian David Johnson. *Science fiction prototyping: Designing the future with science fiction*. Morgan & Claypool Publishers, 2011.

[4] Michael L. Wong and Stuart Bartlett. "Asymptotic burnout and homeostatic awakening: a possible solution to the Fermi paradox?" In: *Journal of The Royal Society Interface* 19.190 (May 2022). ISSN: 1742-5662. DOI: 10.1098/rsif.2022.0029. URL: http://dx.doi.org/10.1098/rsif.2022.0029.

[5] Daron Acemoglu and Pascual Restrepo. "Robots and Jobs: Evidence from US Labor Markets". In: *Journal of Political Economy* 128.6 (June 2020), 2188–2244. ISSN: 1537-534X. DOI: 10.1086/705716. URL: http://dx.doi.org/10.1086/705716.

[6] Michal Bartoš et al. "An overview of robot applications in automotive industry". In: *Transportation Research Procedia* 55 (2021). 14th International scientific conference on sustainable, modern and safe transport, pp. 837–844. ISSN: 2352-1465. DOI: https://doi.org/10.1016/j.trpro.2021.07.052. URL: https://www.sciencedirect.com/science/article/pii/S2352146521004543.

[7] D. R. Peters et al. "Public Health and Climate Benefits and Trade-Offs of U.S. Vehicle Electrification". In: *GeoHealth* 4.10 (Oct. 2020), e2020GH000275. DOI: 10.1029/2020GH000275.

[8] Jamie Morgan. "Electric vehicles: the future we made and the problem of unmaking it". In: *Cambridge Journal of Economics* 44.4 (June 2020), pp. 953–977. ISSN: 0309-166X. DOI: 10.1093/cje/beaa022. eprint: https://academic.oup.com/cje/article-pdf/44/4/953/33469998/beaa022.pdf. URL: https://doi.org/10.1093/cje/beaa022.

[9] Ahmad Mayyas et al. "Design for sustainability in automotive industry: A comprehensive review". In: *Renewable and Sustainable Energy Reviews* 16.4 (2012), pp. 1845–1862. ISSN: 1364-0321. DOI: https://doi.org/10.1016/j.rser.2012.01.012. URL: https://www.sciencedirect.com/science/article/pii/S1364032112000135.

[10] Lin Liang, Liujie Lu, and Ling Su. "The impact of industrial robot adoption on corporate green innovation in China". In: *Scientific Reports* 13.1 (Oct. 2023), p. 18695. ISSN: 2045-2322. DOI: 10.1038/s41598-023-46037-8. URL: https://doi.org/10.1038/s41598-023-46037-8.

[11] Mohsen Attaran. "The rise of 3-D printing: The advantages of additive manufacturing over traditional manufacturing". In: *Business Horizons* 60.5 (Sept. 2017), 677–688. ISSN: 0007-6813. DOI: 10.1016/j.bushor.2017.05.011. URL: http://dx.doi.org/10.1016/j.bushor.2017.05.011.

[12] M. Richardson. "Designer/Maker: The Rise of Additive Manufacturing, Domestic-Scale Production and the Possible Implications for the Automotive Industry". In: *Computer-Aided Design and Applications* PACE (Aug. 2012), 33–48. ISSN: 1686-4360. DOI: 10.3722/cadaps.2012.pace.33-48. URL: http://dx.doi.org/10.3722/cadaps.2012.PACE.33-48.

[13] Thierry Rayna and Ludmila Striukova. "From rapid prototyping to home fabrication: How 3D printing is changing business model innovation". In: *Technological Forecasting and Social Change* 102 (2016), pp. 214–224. ISSN: 0040-1625. DOI: https://doi.org/10.1016/j.techfore.2015.07.023. URL: https://www.sciencedirect.com/science/article/pii/S0040162515002425.

[14] Martin Hannibal and Gary Knight. "Additive manufacturing and the global factory: Disruptive technologies and the location of international business". In: *International Business Review* 27.6 (Dec. 2018), 1116–1127. ISSN: 0969-5931. DOI: 10.1016/j.ibusrev.2018.04.003. URL: http://dx.doi.org/10.1016/j.ibusrev.2018.04.003.

[15] Heather Moorefield-Lang. "Change in the Making: Makerspaces and the Ever-Changing Landscape of Libraries". In: *TechTrends* 59.3 (2015), pp. 107–112. ISSN: 1559-7075. DOI: 10.1007/s11528-015-0860-z. URL: https://doi.org/10.1007/s11528-015-0860-z.

[16] Melisa Orta Martinez et al. "3-D printed haptic devices for educational applications". In: *2016 IEEE Haptics Symposium (HAPTICS)*. 2016, pp. 126–133. DOI: 10.1109/HAPTICS.2016.7463166.

[17] Carmel Majidi. "Soft-Matter Engineering for Soft Robotics". In: *Advanced Materials Technologies* 4.2 (2019), p. 1800477. DOI: https://doi.org/10.1002/admt.201800477. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/admt.201800477. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/admt.201800477.

[18] Jonathan Rossiter and Helmut Hauser. "Soft Robotics - The Next Industrial Revolution? [Industrial Activities]". In: *IEEE Robotics & Automation Magazine* 23.3 (2016), pp. 17–20. DOI: 10.1109/MRA.2016.2588018.

[19] Alex Nemiroski et al. "Arthrobots". In: *Soft Robotics* 4 (3 2017), pp. 183–190. DOI: 10.1089/soro.2016.0043. URL: www.liebertpub.com.

[20] Andrew Jackson, Nathan Mentzer, and Rebecca Kramer-Bottiglio. "Increasing gender diversity in engineering using soft robotics". In: *Journal of Engineering Education* 110.1 (2021), pp. 143–160.

[21] Eric Markvicka et al. "Low-cost wearable human-computer interface with conductive fabric for STEAM education". In: *2018 IEEE Integrated STEM Education Conference (ISEC)*. 2018, pp. 161–166. DOI: 10.1109/ISECon.2018.8340469.

[22] Israel Ulises Cayetano-Jiménez et al. "Experimenting With Soft Robotics in Education: A Systematic Literature Review From 2006 to 2022". In: *IEEE Transactions on Learning Technologies* 17 (2024), pp. 1261–1278. DOI: 10.1109/TLT.2024.3372894.

[23] Aimee van Wynsberghe and Justin Donhauser. "The Dawning of the Ethics of Environmental Robots". In: *Science and Engineering Ethics* 24.6 (Oct. 2017), 1777–1800. ISSN: 1471-5546. DOI: 10.1007/s11948-017-9990-3. URL: http://dx.doi.org/10.1007/s11948-017-9990-3.

[24] Florian Hartmann et al. "Becoming Sustainable, The New Frontier in Soft Robotics". In: *Advanced Materials* 33 (19 May 2021), p. 2004413. ISSN: 1521-4095. DOI: 10.1002/ADMA.202004413. URL: https://onlinelibrary.wiley.com/doi/full/10.1002/adma.202004413https://onlinelibrary.wiley.com/doi/abs/10.1002/adma.202004413https://onlinelibrary.wiley.com/doi/10.1002/adma.202004413.

[25] David F. Williams. "On the nature of biomaterials". In: *Biomaterials* 30.30 (2009), pp. 5897–5909. ISSN: 0142-9612. DOI: https://doi.org/10.1016/j.biomaterials.2009.07.027. URL: https://www.sciencedirect.com/science/article/pii/S0142961209007261.

[26] A. K. Mohanty, M. Misra, and L. T. Drzal. "Sustainable Bio-Composites from Renewable Resources: Opportunities and Challenges in the Green Materials World". In: *Journal of Polymers and the Environment* 10.1 (2002), pp. 19–26. ISSN: 1572-8900. DOI: 10.1023/A:1021013921916. URL: https://doi.org/10.1023/A:1021013921916.

[27] Elsa A. Olivetti and Jonathan M. Cullen. "Toward a sustainable materials system". In: *Science* 360.6396 (2018), pp. 1396–1398. DOI: 10.1126/science.aat6821. eprint: https://www.science.org/doi/pdf/10.1126/science.aat6821. URL: https://www.science.org/doi/abs/10.1126/science.aat6821.

[28] Keene Chin, Tess Hellebrekers, and Carmel Majidi. "Machine Learning for Soft Robotic Sensing and Control". In: *Advanced Intelligent Systems* 2 (6 June 2020), p. 1900171. ISSN: 2640-4567. DOI: 10.1002/AISY.201900171. URL: https://onlinelibrary.wiley.com/doi/full/10.1002/aisy.201900171https://onlinelibrary.wiley.com/doi/abs/10.1002/aisy.201900171https://onlinelibrary.wiley.com/doi/10.1002/aisy.201900171.

[29] Victoria Dean, Yonadav G Shavit, and Abhinav Gupta. "Robots on Demand: A Democratized Robotics Research Cloud". In: *Proceedings of the 5th Conference on Robot Learning*. Ed. by Aleksandra Faust, David Hsu, and Gerhard Neumann. Vol. 164. Proceedings of Machine Learning Research. PMLR, Nov. 2022, pp. 1769–1775. URL: https://proceedings.mlr.press/v164/dean22a.html.

[30] Philipp R. Stoessel et al. "Spinning Angora Rabbit Wool-Like Porous Fibers from a Non-Equilibrated Gelatin/Water/2-Propanol Mixture". In: *Advanced Functional Materials* 24.13 (Nov. 2013), 1831–1839. ISSN: 1616-3028. DOI: 10.1002/adfm.201303321. URL: http://dx.doi.org/10.1002/adfm.201303321.

[31] Philipp R. Stoessel et al. "Porous, Water-Resistant Multifilament Yarn Spun from Gelatin". In: *Biomacromolecules* 16.7 (June 2015), 1997–2005. ISSN: 1526-4602. DOI: 10.1021/acs.biomac.5b00424. URL: http://dx.doi.org/10.1021/acs.biomac.5b00424.

[32] Zach J Patterson et al. "An Untethered Brittle Star-Inspired Soft Robot for Closed-Loop Underwater Locomotion". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS* (2020), pp. 8758–8764. URL: https://github.com/softmachineslab/brittlestar.

[33] Xiaonan Huang et al. "Design and closed-loop motion planning of an untethered swimming soft robot using 2D discrete elastic rods simulations". In: *Advanced Intelligent Systems* 4.10 (2022), p. 2200163.

[34] Kiyn Chin et al. "Accessible soft electronics with silver-gelatin conductive hydrogel composite". In: *Advanced Materials Technologies* (2024). In Review.

[35] Carmel Majidi. "Soft Robotics: A Perspective—Current Trends and Prospects for the Future". In: *Soft Robotics* 1.1 (2014), pp. 5–11. DOI: 10.1089/soro.2013.0001. eprint: https://doi.org/10.1089/soro.2013.0001. URL: https://doi.org/10.1089/soro.2013.0001.

[36] Tess Hellebrekers et al. "Soft magnetic tactile skin for continuous force and location estimation using neural networks". In: *IEEE Robotics and Automation Letters* 5.3 (2020), pp. 3892–3898.

[37] Marc Peter Deisenroth, Carl Edward Rasmussen, and Dieter Fox. "Learning to Control a Low-Cost Manipulator Using Data-Efficient Reinforcement Learning". In: *Robotics: Science and Systems VII*. The MIT Press, June 2012. ISBN: 9780262305969. DOI: 10.7551/mitpress/9481.003.0013. eprint: https://direct.mit.edu/book/chapter-pdf/2269822/9780262305969_car.pdf. URL: https://doi.org/10.7551/mitpress/9481.003.0013.

[38] III Robert J. Webster and Bryan A. Jones. "Design and Kinematic Modeling of Constant Curvature Continuum Robots: A Review". In: *The International Journal of Robotics Research* 29.13 (2010), pp. 1661–1683. DOI: 10.1177/0278364910368147. eprint: https://doi.org/10.1177/0278364910368147. URL: https://doi.org/10.1177/0278364910368147.

[39] Zach J. Patterson et al. *Design and Control of Modular Soft-Rigid Hybrid Manipulators with Self-Contact*. 2024. arXiv: 2408.09275 [cs.RO]. URL: https://arxiv.org/abs/2408.09275.

[40] Jaronie Mohd Jani et al. "A review of shape memory alloy research, applications and opportunities". In: *Materials & Design (1980-2015)* 56 (2014), pp. 1078–1113. ISSN: 0261-3069. DOI: https://doi.org/10.1016/j.matdes.2013.11.084. URL: https://www.sciencedirect.com/science/article/pii/S0261306913011345.

[41] Andrea Sellitto and Aniello Riccio. "Overview and Future Advanced Engineering Applications for Morphing Surfaces by Shape Memory Alloy Materials". In: *Materials* 12.5 (2019), p. 708. DOI: 10.3390/ma12050708.

[42] Konstantinos Andrianesis et al. "Development and Control of an Ultra-Lightweight Anthropomorphic Modular Finger Actuated by Shape Memory Alloy Wires". In: June 2007.

[43] Henry C Astley. "Getting around when you're round: quantitative analysis of the locomotion of the blunt-spined brittle star, Ophiocoma echinata". In: *Journal of Experimental Biology* 215.11 (2012), pp. 1923–1929.

[44] Maxim Likhachev. *Search-based planning with motion primitives*. 2010.

[45] Zach J Patterson et al. *Brittle Star Robot (IROS 2020)*. URL: https://youtu.be/j18NgpCnn3c.

[46] Michael D. Bartlett et al. "High thermal conductivity in soft elastomers with elongated liquid metal inclusions". In: *Proceedings of the National Academy of Sciences of the United States of America* 114.9 (2017). PMCID: PMC5338550, pp. 2143–2148. DOI: 10.1073/pnas.1616377114. URL: https://doi.org/10.1073/pnas.1616377114.

[47] Xiaonan Huang et al. "Chasing biomimetic locomotion speeds: Creating untethered soft robots with shape memory alloy actuators". In: *Science Robotics* 3.25 (2018). PMID: 33141693, eaau7557. DOI: 10.1126/scirobotics.aau7557. eprint: 33141693. URL: https://doi.org/10.1126/scirobotics.aau7557.

[48] Miklós Bergou et al. "Discrete Elastic Rods". In: *SIGGRAPH* (2008).

[49] Edwin Olson. "AprilTag: A robust and flexible visual fiducial system". In: *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2011, pp. 3400–3407.

[50] John Wang and Edwin Olson. "AprilTag 2: Efficient and robust fiducial detection". In: *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Oct. 2016, pp. 4193–4198. DOI: 10.1109/IROS.2016.7759617.

[51] M. Stolle and C.G. Atkeson. "Policies based on trajectory libraries". In: *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.* May 2006, pp. 3344–3349. DOI: 10.1109/ROBOT.2006.1642212.

[52] Vaibhav K. Viswanathan et al. "Efficient Trajectory Library Filtering for Quadrotor Flight in Unknown Environments". In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Oct. 2020, pp. 2510–2517. DOI: 10.1109/IROS45743.2020.9341273.

[53] James Bern et al. "Trajectory Optimization for Cable-Driven Soft Robot Locomotion". en. In: *Robotics: Science and Systems XV*. Robotics: Science and Systems Foundation, June 2019. ISBN: 978-0-9923747-5-4. DOI: 10.15607/RSS.2019.XV.052. URL: http://www.roboticsproceedings.org/rss15/p52.pdf (visited on 10/09/2021).

[54] Charles Schaff et al. "Sim-to-real transfer of co-optimized soft robot crawlers". In: *Autonomous Robots* 47.8 (2023), pp. 1195–1211. DOI: 10.1007/s10514-023-10130-8. URL: https://doi.org/10.1007/s10514-023-10130-8.

[55] Dezhong Tong et al. "A fully implicit method for robust frictional contact handling in elastic rods". In: *Extreme Mechanics Letters* 58 (2023), p. 101924. ISSN: 2352-4316. DOI: https://doi.org/10.1016/j.eml.2022.101924. URL: https://www.sciencedirect.com/science/article/pii/S2352431622002000.

[56] Xiaoyu Chen et al. "UNDERSTANDING DOMAIN RANDOMIZATION FOR SIM-TO-REAL TRANSFER". English (US). In: Publisher Copyright: © 2022 ICLR 2022 - 10th International Conference on Learning Representationss. All rights reserved.; 10th International Conference on Learning Representations, ICLR 2022 ; Conference date: 25-04-2022 Through 29-04-2022. 2022.

[57] OpenAI et al. "Solving Rubik's Cube with a Robot Hand". In: (Oct. 2019). URL: https://arxiv.org/abs/1910.07113v1.

[58] Anthony Wertz, Andrew P. Sabelhaus, and Carmel Majidi. "Trajectory Optimization for Thermally-Actuated Soft Planar Robot Limbs". In: *2022 IEEE 5th International Conference on Soft Robotics (RoboSoft)*. 2022, pp. 439–446. DOI: 10.1109/RoboSoft54090.2022.9762226.

[59] Andrew P. Sabelhaus et al. "In-Situ Sensing and Dynamics Predictions for Electrothermally-Actuated Soft Robot Limbs". In: *Frontiers in Robotics and AI* 9 (2022). ISSN: 2296-9144. DOI: 10.3389/frobt.2022.888261. URL: https://www.frontiersin.org/journals/robotics-and-ai/articles/10.3389/frobt.2022.888261.

[60] Ran Jing et al. *Safe Balancing Control of a Soft Legged Robot*. 2022. arXiv: 2209.13715 [cs.RO]. URL: https://arxiv.org/abs/2209.13715.

[61] Andrew P Sabelhaus and Carmel Majidi. "Gaussian process dynamics models for soft robots with shape memory actuators". In: *2021 IEEE 4th International Conference on Soft Robotics (RoboSoft)*. IEEE. 2021, pp. 191–198.

[62] Hongbo Wang, Massimo Totaro, and Lucia Beccai. "Toward Perceptive Soft Robots: Progress and Challenges". In: *Advanced Science* 5.9 (2018), p. 1800541. DOI: https://doi.org/10.1002/advs.201800541. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/advs.201800541. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/advs.201800541.

[63] Tess Hellebrekers, Oliver Kroemer, and Carmel Majidi. "Soft Magnetic Skin for Continuous Deformation Sensing". In: *Advanced Intelligent Systems* 1 (4 Aug. 2019), p. 1900025. ISSN: 2640-4567. DOI: 10.1002/AISY.201900025. URL: https://onlinelibrary.wiley.com/doi/full/10.1002/aisy.201900025https://onlinelibrary.wiley.com/doi/abs/10.1002/aisy.201900025https://onlinelibrary.wiley.com/doi/10.1002/aisy.201900025.

[64] Jonathan P. King et al. "Design. Fabrication, and Evaluation of Tendon-Driven Multi-Fingered Foam Hands". In: *IEEE-RAS International Conference on Humanoid Robots* 2018-November (Jan. 2019), pp. 540–545. DOI: 10.1109/HUMANOIDS.2018.8624997.

[65] James M. Bern et al. "Soft Robot Control With a Learned Differentiable Model". In: *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*. 2020, pp. 417–423. DOI: 10.1109/RoboSoft48309.2020.9116011.

[66] Pragna Mannam et al. "A Low-Cost Compliant Gripper Using Cooperative Mini-Delta Robots for Dexterous Manipulation". In: *Robotics Science and Systems* (). URL: https://sites.google.com/view/mini-delta-robots.

[67] Elliot W. Hawkes et al. "A soft robot that navigates its environment through growth". In: *Science Robotics* 2.8 (2017), eaan3028. DOI: 10.1126/scirobotics.aan3028. eprint: https://www.science.org/doi/pdf/10.1126/scirobotics.aan3028. URL: https://www.science.org/doi/abs/10.1126/scirobotics.aan3028.

[68] M. Rolf et al. "A multi-level control architecture for the bionic handling assistant". In: *Advanced Robotics* 29.13 (July 2015), pp. 847–859. ISSN: 15685535. DOI: 10.1080/01691864.2015.1037793.

[69] Thomas George Thuruthel et al. "Learning dynamic models for open loop predictive control of soft robotic manipulators". In: *Bioinspiration and Biomimetics* 12.6 (2017). ISSN: 17483190. DOI: 10.1088/1748-3190/aa839f.

[70] Cornelia Schlagenhauf et al. *Control of Tendon-Driven Soft Foam Robot Hands*. Tech. rep.

[71] Thomas George Thuruthel et al. "Learning closed loop kinematic controllers for continuum manipulators in unstructured environments". In: *Soft Robotics* 4.3 (2017), pp. 285–296. ISSN: 21695180. DOI: 10.1089/soro.2016.0051.

[72] Hao Jiang et al. "A two-level approach for solving the inverse kinematics of an extensible soft arm considering viscoelastic behavior". In: *Proceedings - IEEE International Conference on Robotics and Automation*. Institute of Electrical and Electronics Engineers Inc., July 2017, pp. 6127–6133. ISBN: 9781509046331. DOI: 10.1109/ICRA.2017.7989727.

[73] Thomas George Thuruthel et al. "Learning global inverse statics solution for a redundant soft robot". In: *ICINCO 2016 - 13th International Conference on Informatics in Control, Automation and Robotics, Doctoral Consortium* 2 (2016), pp. 303–310. DOI: 10.5220/0005979403030310.

[74] Kit Hang Lee et al. "Nonparametric Online Learning Control for Soft Continuum Robot: An Enabling Technique for Effective Endoscopic Navigation". In: *Soft Robotics* 4.4 (Dec. 2017), pp. 324–337. ISSN: 21695172. DOI: 10.1089/soro.2016.0065.

[75] Y. Ansari et al. "Multiobjective Optimization for Stiffness and Position Control in a Soft Robot Arm Module". In: *IEEE Robotics and Automation Letters* 3.1 (2018), pp. 108–115. ISSN: 23773766. DOI: 10.1109/LRA.2017.2734247.

[76] Jie Chen and Henry Y.K. Lau. "Learning the inverse kinematics of tendon-driven soft manipulators with K-nearest Neighbors Regression and Gaussian Mixture Regression". In: *Proceedings - 2016 the 2nd International Conference on Control, Automation and Robotics, ICCAR 2016* (2016), pp. 103–107. DOI: 10.1109/ICCAR.2016.7486707.

[77] Robert K. Katzschmann, Andrew D. Marchese, and Daniela Rus. "Autonomous Object Manipulation Using a Soft Planar Grasping Manipulator". In: *Soft Robotics* 2.4 (Dec. 2015), pp. 155–164. ISSN: 2169-5172. DOI: 10.1089/soro.2015.0013. URL: https://www.liebertpub.com/doi/10.1089/soro.2015.0013.

[78] Maja Trumic, Kosta Jovanovic, and Adriano Fagiolini. *Kernel-based Nonlinear Adaptive Control of Stiffness and Position for Soft Robots Actuators*. Tech. rep.

[79] Morgan T. Gillespie et al. "Learning nonlinear dynamic models of soft robots for model predictive control with neural networks". In: *2018 IEEE International Conference on Soft Robotics, RoboSoft 2018* (2018), pp. 39–45. DOI: 10.1109/ROBOSOFT.2018.8404894.

[80] Cosimo Della Santina et al. "Controlling Soft Robots: Balancing Feedback and Feedforward Elements". In: *IEEE Robotics and Automation Magazine* 24.3 (2017), pp. 75–83. ISSN: 10709932. DOI: 10.1109/MRA.2016.2636360.

[81] Thomas George Thuruthel et al. "Stable Open Loop Control of Soft Robotic Manipulators". In: *IEEE ROBOTICS AND AUTOMATION LETTERS* 3.2 (2018). DOI: 10.1109/LRA.2018.2797241. URL: http://www.ieee.org/publications.

[82] Jens Kober, J. Andrew Bagnell, and Jan Peters. "Reinforcement learning in robotics: A survey". In: *The International Journal of Robotics Research* 32.11 (2013), pp. 1238–1274. DOI: 10.1177/0278364913495721. eprint: https://doi.org/10.1177/0278364913495721. URL: https://doi.org/10.1177/0278364913495721.

[83] Haochong Zhang et al. "Toward effective soft robot control via reinforcement learning". In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Vol. 10462 LNAI. Springer Verlag, 2017, pp. 173–184. ISBN: 9783319652887. DOI: 10.1007/978-3-319-65289-4_17.

[84] Marvin Zhang et al. "Deep reinforcement learning for tensegrity robot locomotion". In: *Proceedings - IEEE International Conference on Robotics and Automation* (2017), pp. 634–641. ISSN: 10504729. DOI: 10.1109/ICRA.2017.7989079.

[85] Chiwon Lee et al. "Soft robot review". In: *International Journal of Control, Automation and Systems* 15.1 (2017), pp. 3–15. ISSN: 20054092. DOI: 10.1007/s12555-016-0462-3.

[86] Josh Bongard, Victor Zykov, and Hod Lipson. "Resilient Machines Through Continuous Self-Modeling". In: *Science* 314.5802 (2006), pp. 1118–1121. DOI: 10.1126/science.1133687. eprint: https://www.science.org/doi/pdf/10.1126/science.1133687. URL: https://www.science.org/doi/abs/10.1126/science.1133687.

[87] Q. Tyrell Davis et al. "Subtract to Adapt: Autotomic Robots". In: *2023 IEEE International Conference on Soft Robotics (RoboSoft)*. 2023, pp. 1–6. DOI: 10.1109/RoboSoft55895.2023.10122102.

[88] Christopher Atkeson et al. "What Happened at the DARPA Robotics Challenge Finals". In: Apr. 2018, pp. 667–684. ISBN: 978-3-319-74665-4. DOI: 10.1007/978-3-319-74666-1_17.

[89] Joseph E Gaudio et al. *Connections Between Adaptive Control and Optimization in Machine Learning*. 2019.

[90] Girish V. Chowdhary and Eric N. Johnson. "Theory and flight-Test validation of a concurrent-learning adaptive controller". In: *Journal of Guidance, Control, and Dynamics* 34 (2 May 2011), pp. 592–607. ISSN: 15333884. DOI: 10.2514/1.46866. URL: https://arc.aiaa.org/doi/abs/10.2514/1.46866.

[91] Abulikemu Abuduweili and Changliu Liu. "Robust Online Model Adaptation by Extended Kalman Filter with Exponential Moving Average and Dynamic Multi-Epoch Strategy". In: *arXiv* 120 (Dec. 2019), pp. 1–14. URL: http://arxiv.org/abs/1912.01790.

[92] Doya K et al. "Multiple model-based reinforcement learning". In: *Neural computation* 14 (6 June 2002), pp. 1347–1369. ISSN: 0899-7667. DOI: 10.1162/089976602753712972. URL: https://pubmed.ncbi.nlm.nih.gov/12020450/.

[93] Bruno C Da Silva et al. *Dealing with Non-Stationary Environments using Context Detection*. 2006.

[94] Yiğit E. Yücesoy Yücesoy and M. Borahan Tümer. "Hierarchical Reinforcement Learning with Context Detection (HRL-CD)". In: *International Journal of Machine Learning and Computing* 5 (5 Oct. 2015), pp. 353–358. ISSN: 20103700. DOI: 10.7763/ijmlc.2015.v5.533.

[95] Ashish Vaswani et al. "Attention is all you need". In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. NIPS'17. Red Hook, NY, USA: Curran Associates Inc., 2017, 6000–6010. ISBN: 9781510860964.

[96] Weidong Li, Diangang Hu, and Lei Yang. "Actuation Mechanisms and Applications for Soft Robots: A Comprehensive Review". In: *Applied Sciences* 13.16 (2023). ISSN: 2076-3417. DOI: 10.3390/app13169255. URL: https://www.mdpi.com/2076-3417/13/16/9255.

[97] Kiyn Chin. *Five Bar Pendulum Controlled Nonstationary Dynamics*. URL: https://www.youtube.com/watch?v=laMng1LhdvQ.

[98] Pragna Mannam, Oliver Kroemer, and F Zeynep Temel. "Characterization of Compliant Parallelogram Links for 3D-Printed Delta Manipulators". In: ().

[99] Kiyn Chin, Carmel Majidi, and Abhinav Gupta. *1 Modular Parallel Manipulator for Long-Term Soft Robotic Data Collection*. 2024. arXiv: 2409.03614 [cs.RO]. URL: https://arxiv.org/abs/2409.03614.

[100] Michael Ahn et al. *ROBEL: Robotics Benchmarks for Learning with Low-Cost Robots*. 2019. arXiv: 1909.11639 [cs.RO]. URL: https://arxiv.org/abs/1909.11639.

[101] URL: https://stable-baselines3.readthedocs.io/en/master/.

[102] Kiyn Chin. *Modular Manipulator Reinforcement Learning with Hierarchical Primitives Gym Environment*. URL: https://github.com/kiynchin/module-nonstationarity/blob/primitives/primitives_module_env.py.

[103] Charles Fletcher et al. "Earth at risk: An urgent call to end the age of destruction and forge a just and sustainable future". In: *PNAS Nexus* 3.4 (Apr. 2024), pgae106. ISSN: 2752-6542. DOI: 10.1093/pnasnexus/pgae106. eprint: https://academic.oup.com/pnasnexus/article-pdf/3/4/pgae106/57335971/pgae106.pdf. URL: https://doi.org/10.1093/pnasnexus/pgae106.

[104] Ehsanul Kabir et al. "Prospects of biopolymer technology as an alternative option for non-degradable plastics and sustainable management of plastic wastes". In: *Journal of Cleaner Production* 258 (2020), p. 120536. ISSN: 0959-6526. DOI: https://doi.org/10.1016/j.jclepro.2020.120536. URL: https://www.sciencedirect.com/science/article/pii/S0959652620305837.

[105] Rita G. Fonseca et al. "Photodegradable Non-Drying Hydrogel Substrates for Liquid Metal Based Sustainable Soft-Matter Electronics". In: *Advanced Materials Technologies* (2023), p. 2301007. ISSN: 2365-709X. DOI: 10.1002/ADMT.202301007. URL: https://onlinelibrary.wiley.com/doi/full/10.1002/admt.202301007https://onlinelibrary.wiley.com/doi/abs/10.1002/admt.202301007https://onlinelibrary.wiley.com/doi/10.1002/admt.202301007.

[106] Danli Luo et al. "Autonomous self-burying seed carriers for aerial seeding". In: *Nature* 614.7948 (2023), pp. 463–470.

[107] Subash Dahal et al. "Degradability of Biodegradable Soil Moisture Sensor Components and Their Effect on Maize (Zea mays L.) Growth". In: *Sensors* 20.21 (Oct. 2020), p. 6154. ISSN: 1424-8220. DOI: 10.3390/s20216154. URL: http://dx.doi.org/10.3390/s20216154.

[108] César A G Quispe, Christian J R Coronado, and João A Carvalho. "Glycerol: Production, consumption, prices, characterization and new trends in combustion". In: *Renewable and Sustainable Energy Reviews* (2013). DOI: 10.1016/j.rser.2013.06.017. URL: http://dx.doi.org/10.1016/j.rser.2013.06.017.

[109] Karen R. Siegel et al. "Association of Higher Consumption of Foods Derived From Subsidized Commodities With Adverse Cardiometabolic Risk Among US Adults". In: *JAMA Internal Medicine* 176.8 (Aug. 2016), pp. 1124–1132. ISSN: 2168-6106. DOI: 10.1001/jamainternmed.2016.2410. eprint: https://jamanetwork.com/journals/jamainternalmedicine/articlepdf/2530901/ioi160036.pdf. URL: https://doi.org/10.1001/jamainternmed.2016.2410.

[110] Sarah Husnaini Zainal et al. "Preparation of cellulose-based hydrogel: A review". In: *Journal of Materials Research and Technology* 10 (2021), pp. 935–952.

[111] Jeong-Yun Sun et al. "Highly stretchable and tough hydrogels". In: *Nature* 489.7414 (Sept. 2012), 133–136. ISSN: 1476-4687. DOI: 10.1038/nature11409. URL: http://dx.doi.org/10.1038/nature11409.

[112] Wenhuan Sun et al. "Long-Fiber Embedded Hydrogel 3D Printing for Structural Reinforcement". In: *ACS Biomaterials Science and Engineering* 8 (1 Jan. 2022), pp. 303–313. ISSN: 23739878. DOI: 10.1021/ACSBIOMATERIALS.1C00908/SUPPL_FILE/AB1C00908_SI_003.MP4. URL: https://pubs.acs.org/doi/full/10.1021/acsbiomaterials.1c00908.

[113] Yue Zhang and Gang Cheng. "Lignin-containing hydrogels: Transforming a low-value byproduct to absorbents, wound dressings and strain sensors". In: *Chemical Engineering Journal* 475 (Nov. 2023), p. 146460. ISSN: 1385-8947. DOI: 10.1016/j.cej.2023.146460. URL: http://dx.doi.org/10.1016/j.cej.2023.146460.

[114] Seishi Shimizu and Nobuyuki Matubayasi. "Gelation: The Role of Sugars and Polyols on Gelatin and Agarose". In: *The Journal of Physical Chemistry B* 118.46 (Nov. 2014), 13210–13216. ISSN: 1520-5207. DOI: 10.1021/jp509099h. URL: http://dx.doi.org/10.1021/jp509099h.

[115] Melanie Baumgartner et al. "Resilient yet entirely degradable gelatin-based biogels for soft robots and electronics". In: *Nature Materials 2020 19:10* 19 (10 June 2020), pp. 1102–1109. ISSN: 1476-4660. DOI: 10.1038/s41563-020-0699-3. URL: https://www.nature.com/articles/s41563-020-0699-3.

[116] Stanisław Mitura, Alina Sionkowska, and Amit Jaiswal. "Biopolymers for hydrogels in cosmetics: review". In: *Journal of Materials Science: Materials in Medicine* 31.6 (May 2020). ISSN: 1573-4838. DOI: 10.1007/s10856-020-06390-w. URL: http://dx.doi.org/10.1007/s10856-020-06390-w.

[117] Cesar Vinicius Toniciolli Rigueto et al. "Gelatin films from wastes: A review of production, characterization, and application trends in food preservation and agriculture". In: *Food Research International* 162 (2022), p. 112114. ISSN: 0963-9969. DOI: https://doi.org/10.1016/j.foodres.2022.112114. URL: https://www.sciencedirect.com/science/article/pii/S0963996922011723.

[118] Meiyazhagan Ashokkumar and Pulickel M. Ajayan. "Materials science perspective of multifunctional materials derived from collagen". In: *International Materials Reviews* 66.3 (Apr. 2020), 160–187. ISSN: 1743-2804. DOI: 10.1080/09506608.2020.1750807. URL: http://dx.doi.org/10.1080/09506608.2020.1750807.

[119] Farwa Mushtaq et al. "Preparation, properties, and applications of gelatin-based hydrogels (GHs) in the environmental, technological, and biomedical sectors". In: *International Journal of Biological Macromolecules* 218 (Oct. 2022), pp. 601–633. ISSN: 0141-8130. DOI: 10.1016/J.IJBIOMAC.2022.07.168.

[120] A. Heiden et al. "3D printing of resilient biogels for omnidirectional and exteroceptive soft actuators". In: *Science Robotics* 7 (63 Feb. 2022). ISSN: 24709476. DOI: 10.1126/SCIROBOTICS.ABK2119. URL: https://www.anthropocenemagazine.org/2022/02/squishy-eco-friendly-robots-made-with-gelatin-and-sugar/.

[121] Jun Shintake et al. "Soft pneumatic gelatin actuator for edible robotics". In: *IEEE International Conference on Intelligent Robots and Systems* 2017-September (Dec. 2017), pp. 6221–6226. ISSN: 21530866. DOI: 10.1109/IROS.2017.8206525.

[122] Spencer J. Matonis et al. "Edible Origami Actuators Using Gelatin-Based Bioplastics". In: *ACS Applied Polymer Materials* 5.8 (July 2023), 6288–6295. ISSN: 2637-6105. DOI: 10.1021/acsapm.3c00919. URL: http://dx.doi.org/10.1021/acsapm.3c00919.

[123] Christian Geckeler, Benito Armas Pizzani, and Stefano Mintchev. "Biodegradable Origami Gripper Actuated with Gelatin Hydrogel for Aerial Sensor Attachment to Tree Branches". In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*. 2023, pp. 5324–5330. DOI: 10.1109/ICRA48891.2023.10160316.

[124] Bjorn Vergauwen et al. *Gelatin*. Mar. 2016. DOI: 10.1002/14356007.a12_307.pub2. URL: http://dx.doi.org/10.1002/14356007.a12_307.pub2.

[125] Antonio López-Díaz, Andrés S. Vázquez, and Ester Vázquez. "Hydrogels in Soft Robotics: Past, Present, and Future". In: *ACS Nano* 18.32 (2024). PMID: 39099317, pp. 20817–20826. DOI: 10.1021/acsnano.3c12200. eprint: https://doi.org/10.1021/acsnano.3c12200. URL: https://doi.org/10.1021/acsnano.3c12200.

[126] Xiaoqiang Yan et al. "High strength and self-healable gelatin/polyacrylamide double network hydrogels". In: *J. Mater. Chem. B* 5 (37 2017), pp. 7683–7691. DOI: 10.1039/C7TB01780D. URL: http://dx.doi.org/10.1039/C7TB01780D.

[127] Hongling Sun et al. "Ultra-Stretchable, durable and conductive hydrogel with hybrid double network as high performance strain sensor and stretchable triboelectric nanogenerator". In: *Nano Energy* 76 (Oct. 2020), p. 105035. ISSN: 2211-2855. DOI: 10.1016/J.NANOEN.2020.105035.

[128] Junsoo Kim et al. "Fracture, fatigue, and friction of polymers in which entanglements greatly outnumber cross-links". In: *Science* 374.6564 (2021), pp. 212–216.

[129] Chisa Norioka et al. "A universal method to easily design tough and stretchable hydrogels". In: *NPG Asia Materials* 13.1 (Apr. 2021). ISSN: 1884-4057. DOI: 10.1038/s41427-021-00302-2. URL: http://dx.doi.org/10.1038/s41427-021-00302-2.

[130] Meixuanzi Shi et al. "Highly entangled hydrogels with degradable crosslinks". In: *Extreme Mechanics Letters* 59 (2023), p. 101953. ISSN: 2352-4316. DOI: https://doi.org/10.1016/j.eml.2022.101953. URL: https://www.sciencedirect.com/science/article/pii/S2352431622002292.

[131] Giuseppe Tronci et al. "Wet-spinnability and crosslinked fibre properties of two collagen polypeptides with varied molecular weight". In: *International Journal of Biological Macromolecules* 81 (Nov. 2015), 112–120. ISSN: 0141-8130. DOI: 10.1016/j.ijbiomac.2015.07.053. URL: http://dx.doi.org/10.1016/j.ijbiomac.2015.07.053.

[132] M. Verheijen et al. "DMSO induces drastic changes in human cellular processes and epigenetic landscape in vitro". In: *Scientific Reports* 9.1 (Mar. 2019). ISSN: 2045-2322. DOI: 10.1038/s41598-019-40660-0. URL: http://dx.doi.org/10.1038/s41598-019-40660-0.

[133] Dario Consonni et al. "Cancer Risk Among Tetrafluoroethylene Synthesis and Polymerization Workers". In: *American Journal of Epidemiology* 178.3 (July 2013), pp. 350–358. ISSN: 0002-9262. DOI: 10.1093/aje/kws588. eprint: https://academic.oup.com/aje/article-pdf/178/3/350/17341255/kws588.pdf. URL: https://doi.org/10.1093/aje/kws588.

[134] Eldy S. Lazaro Vasquez et al. "Desktop Biofibers Spinning: An Open-Source Machine for Exploring Biobased Fibers and Their Application Towards Sustainable Smart Textile Design". In: *Proceedings of the CHI Conference on Human Factors in Computing Systems*. CHI '24. New York, NY, USA: Association for Computing Machinery, 2024, pp. 1–18. ISBN: 9798400703300. DOI: 10.1145/3613904.3642387. URL: https://doi.org/10.1145/3613904.3642387.

[135] Viviane Lutz-Bueno et al. "Self-Winding Gelatin–Amyloid Wires for Soft Actuators and Sensors". In: *Advanced Materials* 32.48 (2020), p. 2004941. DOI: https://doi.org/10.1002/adma.202004941. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/adma.202004941. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/adma.202004941.

[136] ARHA Lee et al. "3D bioprinting of collagen to rebuild components of the human heart". In: *Science* 365.6452 (2019), pp. 482–487.

[137] Kiyn Chin and Dinesh K. Patel. *Gelatin Coagulation*. URL: https://youtu.be/LzlcVa6-kvs.

[138] Da-Ming Wang, Antoine Venault, and Juin-Yih Lai. "Chapter 2 - Fundamentals of nonsolvent-induced phase separation". In: *Hollow Fiber Membranes*. Ed. by Tai-Shung Chung and Yingnan Feng. Elsevier, 2021, pp. 13–56. ISBN: 978-0-12-821876-1. DOI: https://doi.org/10.1016/B978-0-12-821876-1.00009-3. URL: https://www.sciencedirect.com/science/article/pii/B9780128218761000093.

[139] J. S. Temenoff and A. G. Mikos. *Biomaterials: The Intersection of Biology and Materials Science*. Pearson, 2009, p. 65. ISBN: 9780130097101. URL: https://www.pearson.com/en-us/subject-catalog/p/biomaterials-the-intersection-of-biology-and-materials-science/P200000003248/9780130097101.

[140] *Thermoplastic Urethane Costs*. URL: https://www.amazon.com/SainSmart-Flexible-Printing-Filament-Dimensional/dp/B00TI3JUW4.

[141] *Ecoflex 00-30 Costs*. URL: https://shop.smooth-on.com/ecoflex-00-30.

[142] Xianyang Bao et al. "Low-intensity mixing process of high molecular weight polymer chains leads to elastomers of long network strands and high fatigue threshold". In: *Soft Matter* 19 (31 2023), pp. 5956–5966. DOI: 10.1039/D3SM00687E. URL: http://dx.doi.org/10.1039/D3SM00687E.

[143] Steven G. Harrellson et al. "Hydration solids". In: *Nature* 619.7970 (2023), pp. 500–505. DOI: 10.1038/s41586-023-06144-y. URL: https://doi.org/10.1038/s41586-023-06144-y.

[144] Daniela Rus and Michael T. Tolley. *Design, fabrication and control of soft robots*. May 2015. DOI: 10.1038/nature14543.

[145] Xin Zhao et al. "Photocrosslinkable Gelatin Hydrogel for Epidermal Tissue Engineering". In: *Advanced Healthcare Materials* 5.1 (2016), pp. 108–118. DOI: https://doi.org/10.1002/adhm.201500005. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/adhm.201500005. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/adhm.201500005.

[146] *Sigma-Aldritch Methacrylic anhydride*. URL: https://www.sigmaaldrich.com/US/en/product/aldrich/276685.

[147] *Sigma-Aldrich Genipin*. URL: https://www.sigmaaldrich.com/US/en/product/sigma/g4796.

[148] *Amazon Genipin*. https://www.amazon.com/Grams-genipin-HPLC-Grade-6902-77-8/dp/B0BNWRQL9C/.

[149] *Formaldehyde Costs*. URL: https://www.sciencecompany.com/Formaldehyde-Formalin-37-16oz-P16385.

[150] A Bigi et al. "Mechanical and thermal properties of gelatin films at different degrees of glutaraldehyde crosslinking". In: *Biomaterials* 22.8 (2001), pp. 763–768. ISSN: 0142-9612. DOI: https://doi.org/10.1016/S0142-9612(00)00236-2. URL: https://www.sciencedirect.com/science/article/pii/S0142961200002362.

[151] Evan M. Masutani et al. "Increasing Thermal Stability of Gelatin by UV-Induced Cross-Linking with Glucose". In: *International Journal of Biomaterials* 2014 (2014), 1–9. ISSN: 1687-8795. DOI: 10.1155/2014/979636. URL: http://dx.doi.org/10.1155/2014/979636.

[152] Salih Birhanu Ahmed, Nurcan Doğan, and D. "A Novel Approach to Crosslink Gelatin Nanofibers Through Neutralization-Induced Maillard Reaction". In: *Food and Bioprocess Technology* 17.2 (2023). ISSN: 1935-5149. DOI: 10.1007/s11947-023-03146-6. URL: http://dx.doi.org/10.1007/s11947-023-03146-6.

[153] Kaido Siimon et al. “Effect of glucose content on thermally cross-linked fibrous gelatin scaffolds for tissue engineering”. In: *Materials Science and Engineering: C* 42 (2014), pp. 538–545. ISSN: 0928-4931. DOI: https://doi.org/10.1016/j.msec.2014.05.075. URL: https://www.sciencedirect.com/science/article/pii/S0928493114003580.

[154] Hyo Won Kwak et al. “Effect of crosslinkable sugar molecules on the physicochemical and antioxidant properties of fish gelatin nanofibers”. In: *Food Hydrocolloids* 111 (Feb. 2021), p. 106259. ISSN: 0268-005X. DOI: 10.1016/j.foodhyd.2020.106259. URL: http://dx.doi.org/10.1016/j.foodhyd.2020.106259.

[155] Pengli Zuo et al. “Spontaneous Formation of Fluorescent Carbon Nanoparticles in Glutaraldehyde Solution and Their Fluorescence Mechanism”. In: *Journal of Fluorescence* 31.2 (2021), pp. 509–516. ISSN: 1573-4994. DOI: 10.1007/s10895-020-02678-w. URL: https://doi.org/10.1007/s10895-020-02678-w.

[156] Yongyi Zhao et al. “A self-healing electrically conductive organogel composite”. In: *Nature Electronics* 6.3 (2023), pp. 206–215.

[157] Wenzhen Yuan, Siyuan Dong, and Edward H. Adelson. “GelSight: High-Resolution Robot Tactile Sensors for Estimating Geometry and Force”. In: *Sensors* 17.12 (2017). ISSN: 1424-8220. DOI: 10.3390/s17122762. URL: https://www.mdpi.com/1424-8220/17/12/2762.

[158] Margarethe Hauck et al. “Overcoming Water Diffusion Limitations in Hydrogels via Microtubular Graphene Networks for Soft Actuators”. In: *Advanced Materials* 35.41 (Aug. 2023). ISSN: 1521-4095. DOI: 10.1002/adma.202302816. URL: http://dx.doi.org/10.1002/adma.202302816.

[159] Daniel M. Aukes and Robert J. Wood. “PopupCAD: a tool for automated design, fabrication, and analysis of laminate devices”. In: *Micro- and Nanotechnology Sensors, Systems, and Applications VII*. Ed. by Thomas George, Achyut K. Dutta, and M. Saif Islam. Vol. 9467. International Society for Optics and Photonics. SPIE, 2015, 94671B. DOI: 10.1117/12.2177576. URL: https://doi.org/10.1117/12.2177576.

[160] Guodong Nian et al. “Making Highly Elastic and Tough Hydrogels from Doughs”. In: *Advanced Materials* 34.50 (Oct. 2022). ISSN: 1521-4095. DOI: 10.1002/adma.202206577. URL: http://dx.doi.org/10.1002/adma.202206577.

[161] Abhishek P. Dhand et al. “Additive manufacturing of highly entangled polymer networks”. In: *Science* 385.6708 (2024), pp. 566–572. DOI: 10.1126/science.adn6925. eprint: https://www.science.org/doi/pdf/10.1126/science.adn6925. URL: https://www.science.org/doi/abs/10.1126/science.adn6925.

[162] Ferry P.W. Melchels et al. “Additive manufacturing of tissues and organs”. In: *Progress in Polymer Science* 37.8 (2012). Topical Issue on Biorelated polymers, pp. 1079–1104. ISSN: 0079-6700. DOI: https://doi.org/10.1016/j.progpolymsci.2011.11.007. URL: https://www.sciencedirect.com/science/article/pii/S0079670011001328.

[163] Eman Mirdamadi et al. “FRESH 3D Bioprinting a Full-Size Model of the Human Heart”. In: *ACS Biomaterials Science & Engineering* 6.11 (2020). PMID: 33449644, pp. 6453–6459. DOI: 10.1021/acsbiomaterials.0c01133. eprint: https://doi.org/10.1021/acsbiomaterials.0c01133. URL: https://doi.org/10.1021/acsbiomaterials.0c01133.

[164] Avery S. Williamson et al. "FRESH-Printing of a Multi-actuator Biodegradable Robot Arm for Articulation and Grasping". In: *Biomimetic and Biohybrid Systems*. Ed. by Fabian Meder et al. Cham: Springer Nature Switzerland, 2023, pp. 130–141. ISBN: 978-3-031-38857-6.

[165] Daniel de las Heras, José Maria Tavares, and Margarida M. Telo da Gama. "Bicontinuous and mixed gels in binary mixtures of patchy colloidal particles". In: *Soft Matter* 8 (6 2012), pp. 1785–1794. DOI: 10.1039/C1SM06948A. URL: http://dx.doi.org/10.1039/C1SM06948A.

[166] Ashok R. Patel and Koen Dewettinck. "Edible oil structuring: an overview and recent updates". In: *Food Funct.* 7 (1 2016), pp. 20–29. DOI: 10.1039/C5FO01006C. URL: http://dx.doi.org/10.1039/C5FO01006C.

[167] Qiukai Qi et al. "Edible, optically modulating, shape memory oleogel composites for sustainable soft robotics". In: *Materials & Design* 235 (2023), p. 112339. ISSN: 0264-1275. DOI: https://doi.org/10.1016/j.matdes.2023.112339. URL: https://www.sciencedirect.com/science/article/pii/S0264127523007542.

[168] Sarah Costrell et al. "A Magnetic Soft Device for Tactile Haptic Actuation of the Fingertip". In: *2023 IEEE World Haptics Conference (WHC)* (2023), pp. 48–55. URL: https://api.semanticscholar.org/CorpusID:261128203.

[169] Ruting Zheng et al. "Microstructure and physical properties of novel bigel-based foamed emulsions". In: *Food Hydrocolloids* 134 (2023), p. 108097. ISSN: 0268-005X. DOI: https://doi.org/10.1016/j.foodhyd.2022.108097. URL: https://www.sciencedirect.com/science/article/pii/S0268005X22006178.

[170] Ahmad Shakeel et al. "Key characteristics and modelling of bigels systems: A review". In: *Materials Science and Engineering: C* 97 (2019), pp. 932–953. ISSN: 0928-4931. DOI: https://doi.org/10.1016/j.msec.2018.12.075. URL: https://www.sciencedirect.com/science/article/pii/S0928493118305964.

[171] Guogao Zhang et al. "Hydrogels of arrested phase separation simultaneously achieve high strength and low hysteresis". In: *Science Advances* 9.26 (2023), eadh7742. DOI: 10.1126/sciadv.adh7742. eprint: https://www.science.org/doi/pdf/10.1126/sciadv.adh7742. URL: https://www.science.org/doi/abs/10.1126/sciadv.adh7742.

[172] Xuechuan Wang et al. "Engineered gelatin-based conductive hydrogels for flexible wearable electronic devices: Fundamentals and recent advances". In: *Journal of Science: Advanced Materials and Devices* (2022). DOI: 10.1016/j.jsamd.2022.100451. URL: https://doi.org/10.1016/j.jsamd.2022.100451.

[173] Jiangxin Wang et al. "Deformable conductors for human–machine interface". In: *Materials Today* 21.5 (June 2018), 508–526. ISSN: 1369-7021. DOI: 10.1016/j.mattod.2017.12.006. URL: http://dx.doi.org/10.1016/j.mattod.2017.12.006.

[174] David Hardman, Thomas George Thuruthel, and Fumiya Iida. "Self-healing ionic gelatin/glycerol hydrogels for strain sensing applications". In: *NPG Asia Materials* (2022). DOI: 10.1038/s41427-022-00357-9. URL: https://doi.org/10.1038/s41427-022-00357-9.

[175] Xin Jing et al. "Stretchable gelatin/silver nanowires composite hydrogels for detecting human motion". In: *Materials Letters* 237 (2019), pp. 53–56. ISSN: 0167-577X. DOI: https://doi.org/10.1016/j.matlet.2018.11.078. URL: https://www.sciencedirect.com/science/article/pii/S0167577X1831841X.

[176] Chunlin Liu et al. "Electrically Conductive Tough Gelatin Hydrogel". In: *Advanced Electronic Materials* 6.4 (2020), p. 2000040. DOI: https://doi.org/10.1002/aelm.202000040. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/aelm.202000040. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/aelm.202000040.

[177] Junggeon Park et al. "Electrically Conductive Hydrogel Nerve Guidance Conduits for Peripheral Nerve Regeneration". In: *Advanced Functional Materials* 30.39 (2020), p. 2003759. DOI: https://doi.org/10.1002/adfm.202003759. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/adfm.202003759. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/adfm.202003759.

[178] Kai Han et al. "Gelatin-based adhesive hydrogel with self-healing, hemostasis, and electrical conductivity". In: *International Journal of Biological Macromolecules* 183 (2021), pp. 2142–2151. DOI: 10.1016/j.ijbiomac.2021.05.147. URL: https://doi.org/10.1016/j.ijbiomac.2021.05.147.

[179] Kai Chen et al. "Highly Stretchable, Tough, and Conductive Ag@Cu Nanocomposite Hydrogels for Flexible Wearable Sensors and Bionic Electronic Skins". In: *Macromolecular Materials and Engineering* 306 (10 Oct. 2021). ISSN: 14392054. DOI: 10.1002/MAME.202100341.

[180] Yunsik Ohm et al. In: *Nature Electronics 2021 4:3* 4 (3 Mar. 2021), pp. 185–192. ISSN: 2520-1131. DOI: 10.1038/s41928-021-00545-5. URL: https://www.nature.com/articles/s41928-021-00545-5.

[181] Michael J. Ford et al. "Controlled Assembly of Liquid Metal Inclusions as a General Approach for Multifunctional Composites". In: *Advanced Materials* 32.46 (2020), p. 2002929. DOI: https://doi.org/10.1002/adma.202002929. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/adma.202002929. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/adma.202002929.

[182] Tiansheng Gan et al. "Conformally Adhesive, Large-Area, Solidlike, yet Transient Liquid Metal Thin Films and Patterns via Gelatin-Regulated Droplet Deposition and Sintering". In: *ACS Applied Materials and Interfaces* 14 (37 Sept. 2022), pp. 42744–42756. ISSN: 19448252. DOI: 10.1021/ACSAMI.2C12880/ASSET/IMAGES/LARGE/AM2C12880_0007.JPEG. URL: https://pubs.acs.org/doi/full/10.1021/acsami.2c12880.

[183] Kiyn Chin and Dinesh K. Patel. *Gelatin Hydrogel Fabrication*. URL: https://www.youtube.com/watch?v=14h__ZCPNJQ.

[184] Kiyn Chin and Dinesh K. Patel. *Conductive Filler Preparation - Wetting Silver Flakes*. URL: https://youtu.be/ArBiZhxyVNM.

[185] Matthew Wei Ming Tan et al. "Toughening Self-Healing Elastomers with Chain Mobility". In: *Advanced Science* (June 2024). ISSN: 2198-3844. DOI: 10.1002/advs.202308154. URL: http://dx.doi.org/10.1002/advs.202308154.

[186] Zhijie Zhao et al. "Natural polymers-enhanced double-network hydrogel as wearable flexible sensor with high mechanical strength and strain sensitivity". In: *Chinese Chemical Letters* 34.6 (June 2023), p. 107892. ISSN: 1001-8417. DOI: 10.1016/j.cclet.2022.107892. URL: http://dx.doi.org/10.1016/j.cclet.2022.107892.

[187] Julie Diani, Bruno Fayolle, and Pierre Gilormini. "A review on the Mullins effect". In: *European Polymer Journal* 45.3 (2009), pp. 601–612. ISSN: 0014-3057. DOI: https://doi.org/10.1016/j.eurpolymj.2008.11.017. URL: https://www.sciencedirect.com/science/article/pii/S0014305708006332.

[188] Mahmoud Tavakoli et al. "EGaIn-Assisted Room-Temperature Sintering of Silver Nanoparticles for Stretchable, Inkjet-Printed, Thin-Film Electronics". In: *Advanced Materials* 30.29 (2018), p. 1801852. DOI: https://doi.org/10.1002/adma.201801852. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/adma.201801852. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/adma.201801852.

[189] Kiyn Chin. *Electromechanical Testing Process*. URL: https://youtu.be/UmY3zIjtm1E.

[190] Kiyn Chin. *Electromyography Electrode Recording*. URL: https://youtu.be/0l3jRjZYjm4.

[191] Kiyn Chin. *Soft Wearable Circuit Conformal Motion*. URL: https://youtu.be/nM3g3vU6lbY.

[192] Sila Temizel-Sekeryan and Andrea L. Hicks. "Global environmental impacts of silver nanoparticle production methods supported by life cycle assessment". In: *Resources, Conservation and Recycling* 156 (May 2020), p. 104676. ISSN: 0921-3449. DOI: 10.1016/J.RESCONREC.2019.104676.

[193] Alexandre Larmagnac et al. "Stretchable electronics based on Ag-PDMS composites". In: *Scientific Reports* 4.1 (2014), p. 7254. DOI: 10.1038/srep07254. URL: https://doi.org/10.1038/srep07254.

[194] Maria Do et al. "Effect of tannic acid as crosslinking agent on fish skin gelatin-silver nanocomposite film". In: *Food Packaging and Shelf Life* (2018). DOI: 10.1016/j.fpsl.2018.11.005. URL: https://doi.org/10.1016/j.fpsl.2018.11.005.

[195] José María Lacave et al. "Waterborne exposure of adult zebrafish to silver nanoparticles and to ionic silver results in differential silver accumulation and effects at cellular and molecular levels". In: *Science of the Total Environment* 642 (Nov. 2018), pp. 1209–1220. ISSN: 18791026. DOI: 10.1016/J.SCITOTENV.2018.06.128. URL: https://www.sciencedaily.com/releases/2018/09/180918110905.htm.

[196] Adriano Magesky and Émilien Pelletier. "Cytotoxicity and Physiological Effects of Silver Nanoparticles on Marine Invertebrates". In: *Cellular and Molecular Toxicology of Nanoparticles*. Cham: Springer International Publishing, 2018, pp. 285–309. ISBN: 978-3-319-72041-8. DOI: 10.1007/978-3-319-72041-8_17. URL: https://doi.org/10.1007/978-3-319-72041-8_17.

[197] Manuel Reis Carneiro et al. "Recyclable Thin-Film Soft Electronics for Smart Packaging and E-Skins". In: *Advanced Science* 10.26 (2023), p. 2301673. DOI: https://doi.org/10.1002/advs.202301673. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/advs.202301673. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/advs.202301673.

[198] Junaid Khan and M Mariatti. "Effect of natural surfactant on the performance of reduced graphene oxide conductive ink". In: *Journal of Cleaner Production* 376 (2022), p. 134254. DOI: 10.1016/j.jclepro.2022.134254. URL: https://doi.org/10.1016/j.jclepro.2022.134254.

[199] Hao He et al. "Biocompatible Conductive Polymers with High Conductivity and High Stretchability". In: *ACS Applied Materials & Interfaces* 11.29 (2019), pp. 26185–26193. DOI: 10.1021/acsami.9b07325. URL: https://doi.org/10.1021/acsami.9b07325.

[200] Baoyang Lu et al. "Pure PEDOT:PSS hydrogels". In: *Nature Communications* (2019). DOI: 10.1038/s41467-019-09003-5. URL: https://doi.org/10.1038/s41467-019-09003-5.

[201] Amal George Kurian et al. "Multifunctional Molybdenum-based nanoclusters engineered Gelatin Methacryloyl as in situ photo-cross-linkable hybrid hydrogel dressings for enhanced wound healing". en. In: *ACS Appl. Mater. Interfaces* (June 2024).

[202] *Materiom*. URL: https://commons.materiom.org/materials-database.

[203] *Jiva Materials*. URL: https://www.jivamaterials.com/.

[204] Rick Bonney et al. "Citizen Science: A Developing Tool for Expanding Science Knowledge and Scientific Literacy". In: *BioScience* 59.11 (Dec. 2009), pp. 977–984. ISSN: 0006-3568. DOI: 10.1525/bio.2009.59.11.9. eprint: https://academic.oup.com/bioscience/article-pdf/59/11/977/19403922/59-11-977.pdf. URL: https://doi.org/10.1525/bio.2009.59.11.9.

[205] Rick Bonney et al. "Can citizen science enhance public understanding of science?" In: *Public Understanding of Science* 25.1 (2016). PMID: 26445860, pp. 2–16. DOI: 10.1177/0963662515607406. eprint: https://doi.org/10.1177/0963662515607406. URL: https://doi.org/10.1177/0963662515607406.

[206] Christopher G Atkeson, Andrew W Moore, and Stefan Schaal. "Locally Weighted Learning". In: *Artificial Intelligence Review* 11 (1997), pp. 11–73. URL: http://www.cc.gatech.edu/fac/Chris.Atkesonhttp://www.cc.gatech.edu/fac/Stefan.Schaalhttp://www.cs.cmu.edu/awm/hp.html.

[207] Marco Barbero, Roberto Merletti, and Alberto Rainoldi. "Atlas of Muscle Innervation Zones". In: *Atlas of Muscle Innervation Zones* (2012). DOI: 10.1007/978-88-470-2463-2.