Title

Human Hand Pose Estimation and Artificial Tactile Sensing in Harsh Environments

Author

Peltola, Eric

Publication Date

2022

UNIVERSITY OF CALIFORNIA
Los Angeles

Human Hand Pose Estimation and Artificial Tactile Sensing in Harsh Environments

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Mechanical Engineering

by

Eric Richard Peltola

2022

ABSTRACT OF THE DISSERTATION

Human Hand Pose Estimation and Artificial Tactile Sensing in Harsh Environments

by

Eric Richard Peltola
Doctor of Philosophy in Mechanical Engineering
University of California, Los Angeles, 2022
Professor Veronica J. Santos, Chair

In Study #1, we develop a novel, data-driven method to estimate the kinematic parameters of multi-joint linkages such as the human hand. The method can identify up to two consecutive revolute joint axis orientations between connected rigid bodies. We introduce kinematic constraints into a Generative Topographic Mapping formulation in order to estimate the joint axis parameters. The method was evaluated using simulated motion and via motion capture and a physical 2-DOF mechanism modeled after the metacarpophalangeal joint of the human finger. Our method compares well against state-of-the-art kinematic parameter estimation techniques with regards to reliability and computational efficiency.

In Study #2, we introduce a sensor-embedded soft skin capable of multimodal sensing (contact force and two axes of shear force) in pressurized underwater environments up to 1000kPa. We embed liquid-metal strain gauges within a durable elastomeric skin that is molded around a solid finger core. We demonstrate that the sensor skin is capable of measuring forces up to 220N underwater and while subjected to a range of hydrostatic pressures. We determine that the performance of the sensor skin is unaffected by the submerged, pressurized environment.

In Study #3, we propose a set of design considerations for tactile sensor skins using embedded, microfluidic single-axis strain gauges for the purpose of estimating 3D forces and 1D torque about the skin's surface normal. By displacing shear force taxels such that their principal axes are offset from the point of contact, we are able to more accurately measure torque. We use an experimental testbed to apply force-torque loads to the sensor skin. We develop CNN-based models to evaluate the combined force-torque estimation performance of numerous taxel configurations and provide a detailed discussion of how performance relates to design choices.

In summary, we developed methods that improve sensing in harsh environments such as granular media and underwater. Kinematic and kinetic considerations during hand-object interaction were carefully integrated into the development of novel sensing hardware, data-driven estimation methods, and task-specific sensor design criteria. We advanced the state-of-the-art in tactile sensing using tactile sensor skins and improved the accuracy of hand pose estimation using low-cost motion tracking tools.

The dissertation of Eric Richard Peltola is approved.

Robert Candler

Jacob Rosen

Jonathan Posner

Veronica J. Santos, Committee Chair

University of California, Los Angeles

2022

To my favorite screwdriver

who tuned potentiometers in the thousands

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

To the people without whom this dissertation would have likely still been a dream:

My advisor, Professor Veronica J. Santos, who kept me from straying too far while also pushing me to hunt the tangential questions which would become my first and third articles. Each discussion with her was enlightening and focusing, divergent and convergent. I am grateful to believe I have picked up a fraction of her resolve;

Professor Jonathan D. Posner, who kept me on my toes and taught me to see through the noise;

Professor Jacob Rosen, from whom I have a great deal of memorable and inspiring quotes to live by.

Many thanks to all of my fellow lab members - we collaborated, commiserated, and ate a lot of pizza on the roof. When I think back on the PhD, I'll be thinking mostly of you all.

To my amazing girlfriend Alexandra, whose strength in putting up with the insanity of my last year should be a testament: I love you.

And to my loving parents and family, I have you to thank for literally everything. I am here because of you.

For chapter 1, I'd like to thank Dr. Eunsuk Chong, without whom I would never had learned about the key mathematics that made that paper possible

For chapters 2 and 3, I'd like to thank my long-time UW collaborator Dr. Minyung Song for going above and beyond her call of duty in helping me push through the final project of my PhD.

VITA

2022–Pres.	Principal Engineer, Horizon Surgical Systems, Malibu, California
2017–Pres.	Graduate Student Researcher, UCLA Biomechatronics Lab, Mechanical and Aerospace Engineering Department, UCLA, Los Angeles, California
2017	M.S. Mechanical Engineering, UCLA, Los Angeles, California
2010–2016	Scientist, Ativa Medical, St. Paul, Minnesota
2011	B.A. Physics, Colorado College, Colorado Springs, Colorado

CHAPTER 1

Introduction

1.1 Motivation

A persistent theme in robotics research is the emulation of human capability. While we might take for granted the ability to grasp a sweaty cup on a hot day, this seamless integration of cognition, sensing, and action has remained elusive to generations of roboticists. Perhaps the most elusive is the complexity and elegance of the human hand; 35 actuators, 22 degrees of freedom, and over 17,000 mechanoreceptors act in harmony to carry out countless tasks, often with minimal cognitive burden to the human.

Much as modern technology has its roots in the international space race of the 1960's, the broad pursuit of human-like robotic systems has opened up deep research efforts in a wide variety of disciplines. For instance, the task of recognizing one's place in the environment spurred decades of specialized work in computer science, applied mathematics, electronics, and optics to give us facial recognition and self-driving cars. It is a convenient metaphor to consider the goal of robotics research as bestowing our own senses and behaviors on robotic systems. Now that the robots can see (the sense afforded by computer vision), many researchers are now turning to robot hands and exploring the significantly subtler sense of touch.

1.1.1 Human Hand Motion Tracking

Human motion capture is a robust and mature research field, enabling applications in diverse areas such as sports science, virtual reality, gaming, remote surgery, and rehabilitation [3, 4, 5, 6, 7, 8]. Many of the fundamental aspects of human motion capture are considered a solved problem; optical-based marker methods such as the VICON system, as well as computer vision methods often provide enough accuracy [9]. These systems often come at a high price and are subject to application-specific errors. Chief among the error modes in optical systems is marker occlusion, which can occur often in scenarios where it is desirable to let the subject move freely [10]. If the motion capture system cannot reliably detect marker locations, it will fail to reliably capture motion data.

The occlusion problem is especially difficult when considering the high degree-of-freedom motion of the human hand digging through granular media; resolving the 3D orientation of a rigid body requires placement of marker triads, and these finger-mounted markers are bound to overlap during complex motion. Furthermore, an optical system will fail outright if the subject's hands are buried beneath the surface of granular media.

1.1.2 Artificial Tactile Sensing

The field of tactile sensing in robotics has been exposed by many as the logical "next step" in robotic research, following the successes of robotic perception via vision. In many situations, vision can provide an accurate model of the environment in which a robot works. However, robotic manipulation of objects which is based on vision alone often does not account for the kinetics of a robot-object interaction. Affording a robot with tactile sensing gives it critical information at the point of contact, allowing the measurement of grasp quality, perception of invisible properties of the object, and more intuitive and human-like interaction with the object.

One focus of this dissertation is to answer the question, “what considerations do we need when performing tactile perception and manipulation underwater?” In the human case, much has been reported about the sensitivity of manipulation tasks to the moisture content of the finger pads, predominantly investigated via friction models. Clearly the kinetics of a submerged grasp will change due to buoyancy, but to uncover the full effects of submerged grasp, an experimental plan must be put in place to compare the properties of single and multi-digit grasps in water and out of water. My first research question sets out a series of experiments to address underwater tactile perception by first demonstrating a perception algorithm in a dry environment before performing the same task underwater. Submerged tactile sensing is an unexplored research area, so I intend to develop the experiments in a way that provides transparency to the systemic effects of underwater tactile perception.

In addition to investigating the unique requirements of underwater tactile sensing, I intend to improve the state of tactile sensing for robot-object interaction by developing the capability to detect and quantify in-plane torsion induced between a sensorized gripper and a rotating object or environment. With the ability to detect and quantify rotational forces, robot manipulators will be able to perform more advanced manipulation tasks. As an instructive example, consider the task of pouring liquid out of a cup. Instead of gripping the cup firmly and rotating the entire hand, I believe a more efficient and dexterous method would be to rotate the cup relative to the fingers — this will produce a rotation in-plane, or torsion, on the surface of the sensorized fingertips. I believe that the ability to quantify torsion will enable an artificial manipulator to pour in a way that a human observer would consider more intuitive and human-like.

1.2 Prior Applications

1.2.1 Human Hand Motion Tracking

Capturing and analyzing human motion has been an active research area for many decades, with diverse applications including rehabilitation, sports, gaming, and entertainment. Major success has been found in applications where human movement tracking has contributed to improved results in ambulatory care [11, 12]. The technologies used in capturing human movement are varied and have heavily favored computer vision in the past decade [13]. Reliance upon computer vision limits the use scenarios to those with a well-defined visual setting and low environmental occlusion. Non-visual tracking systems do not suffer this "line-of-sight" problem and are therefore considered more robust in many motion capture scenarios [10].

The number of non-visual human motion capture systems based on Inertial Measurement Units (IMUs) has increased rapidly since feasibility was demonstrated. An IMU utilizes multiple sensors to estimate the inertial state of the body to which it is attached, such as its angular and linear motion, and 3D orientation. Sensor fusion algorithms have been developed to reduce inherent error which would be present in inertial measurements based on the individual sensors' raw data. An algorithm developed by Sebastian Madgewick uses a magnetometer to correct for gyroscopic drift [14]. This results in a 9-DOF IMU which can accurately resolve absolute three-dimensional orientation.

With advancements in production of cheap and accurate IMUs, development of hand-worn sensor gloves emerged as a viable human motion capture method. Liu et al. developed a glove-based system to perform simultaneous hand pose and force sensing in real time [15]. Using 15 IMUs and a network of Velostat piezoresistive fabric, they performed an experiment which quantitatively characterized the accuracy of the kinematics and kinetics of hand-object interaction. Another active research field related to hand-worn IMU sensor gloves is rehabilitation, where dynamic measurements of hand kinematics provides clinicians with a measurement of movement deficiencies in patients with rheumatoid arthritis [16].

My previous work on the design and fabrication of an IMU sensor glove was aimed at alleviating a common and often necessary drawback of IMU gloves, which is the requirement of bulky hardware to ensure proper fastening of sensors. The existing research commonly employs a tight-fitting and robust glove superstructure, which ensures that the motion of the sensors is directly coupled to the hand kinematics. This bulk, however, often impedes the natural motion of the hand and corrupts the analysis of glove-worn hand motion. The team of Kortier et al. approached this issue by using a dorsally-mounted array of sensors and flexible PCBs to leave the palm and finger pads exposed during motion [17]. My palmless IMU glove was designed in a similar manner, using a minimum of attachment hardware as shown in Figure 1. This design requirement allows the wearer to freely explore a tactile environment or to perform complex gestures with as little added constraint as possible.

1.2.2 Artificial Tactile Sensing

Humans can make excellent use of tactile sensing to recognize grasped objects, as well as naturally and adeptly manipulate objects and interact with the environment even in the absence of vision. In order to afford robots and prostheses with these same abilities, research has predominantly focused on building tactile sensor networks that mimic human capability. Among the most mature of these hardware systems are the BioTac [18] and the GelSight [19] sensors.

The BioTac's biomimetic approach is based on a rigid core surrounded by a weakly conductive fluid contained within an elastomeric skin. The deformable nature of the sensor allows for multimodal measurements of the deforming liquid. Electrodes mounted within the conductive fluid perceive changes in voltage as a function of liquid deformation, and pressure and temperature transducers detect changes in the bulk properties of the encapsulated liquid.

The GelSight sensor takes a different approach, borrowing its fundamental sensing capacity from computer vision. Using an elastomeric slab covered with a reflective membrane, the GelSight can achieve high accuracy in estimating contact surface geometry. Bulk movement of the sensor surface can be measured by tracking the optical flow of an array of dots on the inside of the sensor's contact surface.

Current literature surrounding tactile sensing for robotics is dense with examples of utilizing available sensor hardware to demonstrate simple tasks, such as preventing a grasped object from slipping relative to the hand. Su et. al developed an algorithm for reactive control of grip forces when grasping an object [20]. Using a pair of BioTac sensors, finger forces were estimated and fed into their algorithm, shown in Figure 8, to detect and classify slip events. Fig. 9 shows the difference between a grip controller based on position data only (left) compared to their proposed force-based controller (right). This research exposes a key advantage of incorporating tactile information into the control of grip; in addition to the robot's internal kinematic model of finger position, sensory data from the touched surface provides valuable information to ensure an optimal grasp.

The simple model of "if slipping, grip harder" has been employed successfully in many cases where the only goal was to ensure a simplistic form of grasp stability. Wettels et. al presented a method for controlling the grasp of an anthropomorphic prosthetic hand by using the BioTac sensor, Bayesian inference, and simple algorithms for estimation and control [21]. A pair of sensorized fingertips were first employed to estimate the coefficient of friction of the finger/object pair. Using tactile feedback-based force control and the estimated friction model, the prosthetic hand was able to maintain the object within the friction cone to prevent object slip. Song et. al used a similar approach with a more advanced LuGre friction model, which accounts for nonlinearity of the friction coefficient during the "incipient slip" state that occurs shortly before full object slip [22]. Whereas the friction estimation method proposed in [21] was dependent on an initial routine of letting the object slip, the method of Song et al. could perform slip detection without requiring the initial routine.

The above-mentioned methods highlight the importance of an accurate estimation of friction during object manipulation. During underwater manipulation, the friction model is likely to be affected by the moisture content of the finger-object interface. Additionally, there is no guarantee that the modification of an in-air grip control algorithm for underwater use is trivial. It is a major focus of my proposed work to investigate the effect of submersion on tactile perception.

There is a wealth of underwater robots, some employing a form of tactile sensing for submerged manipulation. For instance, a sensorized manipulator with inherent resistance to hydraulic pressure was developed by O'Brian and Lane, whose artificial fingertip with embedded strain gauges and PVDF films could detect grip force and slip via vibration, respectively [23]. The feasibility of underwater use of rigid fingertips with strain gauge-based tactile sensing was evaluated by Sans et. al, whose multi-fingered gripper demonstrated the underwater manipulation of a wooden trunk [24]. Kampmann and Kirchner took this same concept into deep-sea applications, where a complete manipulation system was developed for use at extreme depths up to 6km [25]. This system incorporated a novel optical tactile sensor array using an open-cell foam matrix for exterioceptive tactile input.

Critically, although the available underwater tactile-based manipulators are informed by the submerged application requirements, none of the related work examined the influence of water submersion on the model of automated grasping. Research into the relationship between contact surface wetness and grasp quality is found in the biomechanics community, where fingertip hydration was found to have a dramatic effect on the dynamics of the contact region during the transition between sticking and slipping [26]. The resultant effect of skin hydration was to reduce the transition of stick to slip, regardless of the coefficient of friction. A similar effect was shown by Nacht et. al, who observed a temporary increase of skin friction coefficient after water application [27]. This effect was observed to dissipate after roughly 10 minutes. Presumably, this effect disappeared due to the absorption and evaporation of water from the skin's surface. An interesting extension of the above-mentioned research would be to study the friction properties of a fully submerged hand.

The existing literature surrounding the tactile perception of in-plane torsion is limited. The group of Dr. Danica Kragic has approached this problem through a series of experiments with increasing complexity. First, a simple rubberized two-fingered gripper was used along with a force-torque sensor mounted in the wrist of the gripper [28]. Using the aid of gravity, a grasped object was allowed to rotate relative to the gripper using a controlled-slip method. Critical to this research was an adaptive friction model, which used cues from the force-torque sensor to update the friction coefficient of the rubberized fingertips in real time. See Figure 10 for a visualization of this experiment, where a desired angular trajectory was followed by the real-time controller. Although this work quantified torsion, the method relied on a frictional model of the grasp and did not quantify torsion directly via tactile sensing.

The GelSight sensor, with high-resolution spatiotemporal sensing of a touched object via computer vision techniques, should be readily able to quantify torsion [19]. This can be done by tracking the visual flow of a matrix of dots etched into the elastomeric fingertip (Figure

11). Although a visualization of torsion has been shown, a quantification of rotational motion due to in-plane torsion was not demonstrated. In a similar implementation, Padmanabha et. al created a multi-camera optically-based tactile sensor called OmniTact for multi-axis force-torque estimation. Relying on the high-resolution camera input and a high-curvature finger surface, the sensor was shown to be capable of full 6-axis force-torque estimation.

Building upon the optics-based tactile sensing method, Li et. al developed a finger capable of not only pressure distribution but also 6-axis force-torque measurement. Their finger, called the F-TOUCH, incorporates spring-like mechanisms within the body of the finger to measure applied forces and torques, while the fingertip measures contact distribution via an embedded camera system. Using an experimental apparatus to create repeatable applied loads, the force-torque estimation performance of the sensor was validated. They showed that the additional spring-like elements of their sensor marked a significant improvement in a comparison study against the GelSight sensor.

1.3 Contributions

This dissertation presents contributions to sensing in robotics and biomechanics, inspired by the kinematics and kinetics of the human hand.

Chapter 2 presents a novel data-driven method for the estimation of human hand kinematic parameters based on observations of hand motion. Using the exponential coordinates of rotation, the observations form a manifold that is uniquely determined by the kinematic parameters of the hand. We determine the parameters of this manifold using a novel variant of the Generative Topographic Mapping (GTM) algorithm, which is a nonlinear latent variable model. By modifying the GTM to operate on the Product-of-Exponentials formulation of robot kinematics, we can extract the hand parameters directly from the mapping between observations and the underlying kinematic structure.

Chapter 3 presents a novel skin-like tactile sensor for use in underwater, pressurized environments. Fabricated using microfluidic strain gauges embedded in elastomer, this sensor can exhibit multimodal force sensing (contact normal and two axes of in-plane shear) while standing up to the harsh environment equivalent to 100 meters of ocean depth. We outline the fabrication process, which builds on previous work in tactile sensing via liquid metal strain gauges. We then characterize the environmental sensitivity and force estimation behavior of the sensor skin as a function of the underwater, pressurized environment.

Chapter 4 outlines the process by which skin-like tactile sensor arrays are designed in order to measure in-plane torque. We provide criteria for the geometric placement of individual sensing elements (taxels) within an elastomeric contact medium in order to optimize the sensitivity of the sensor skin during contact. Using the experimental sensor skin of 3, we test the hypothesis that a certain placement of taxels within the medium will enable the sensor skin to accurately measure in-plane torque. We use an experimental testbed to create contact forces and torques at the surface of the sensor skin, and a data acquisition system to capture both taxel readings and ground-truth force-torque readings. To investigate torque estimation capabilities of the sensor, we train a simple CNN model on the data collected with the testbed. By comparing the performance metrics of the CNN model using subsets of the taxel array, we are able to validate our suggested design criteria for torque sensing. In addition to in-plane torque sensing, we show the skin-like sensor is capable of 4-axis force-torque estimation (three axes of linear force and 1 axis of torque).

Chapter 5 provides a summary of the work presented, along with several detailed suggestions for future work on these topics.

CHAPTER 2

Kinematic Parameter Estimation Using Workspace Manifold Mapping

This chapter was based on work co-authored with Dr. Eunsuk Chong and Dr. Xiaoyu Wang.

2.1 Abstract

Objective: Inference of the kinematic parameters of one and two degree of freedom (DOF) joints, such as the metacarpophalangeal and interphalangeal joints of the hand, is achieved by applying a novel data-driven method on observations of the motion of two rigid bodies connected to the joint. *Methods*: Our proposed estimation algorithm relies on a unique geometry that arises within the motion observations. Using the exponential representation of rotations, we describe in detail how we formulate the algorithm to successfully extract kinematic information from the geometry. We verify our method using simulated data from an ideal 2-DOF linkage and real data collected from a motion capture experiment using a 2-DOF biologically inspired mechanical linkage. With the simulated data, we perform model sensitivity analysis with respect to parameters and measurement noise. A quantitative and qualitative comparison study between the proposed algorithm and several benchmark kinematic estimation algorithms is provided with an experiment using a 3D-printed finger and motion capture data. *Results*: The mean estimates of the 2-DOF joint axis orientations deviate from ground truth estimates by 2.46°, with a standard deviation of 3.42°. *Conclusion*: Our method can identify the orientations of a 2-DOF revolute joint axes between two connected rigid bodies, outperforming existing methods in terms of precision, reliability, and computational time. *Significance*: Our proposed method provides improved reliability for the kinematic parameter estimation task compared to existing state-of-the-art methods. This method, demonstrated here in the context of human hand motion, can be extended to any 1- or 2-DOF joint without loss of generality.

2.2 Introduction

The study of multi-joint systems is ubiquitous in fields such as robotics [4], rehabilitation [3, 6], ergonomics [7], animation [8], and sports [5]. In such systems, the underlying kinematic model that represents physical dimensions and joint structure is a major consideration. In practice, errors between a real-world system and its model are inevitable. In mechanical systems, these errors can arise from manufacturing variability and mechanical wear [29]. In biological systems, errors from the non-rigidity and complexity of biological structures, as well as inherent anatomical variability between individuals, render a priori kinematic models inaccurate [30]. Accordingly, there is a need to estimate the kinematic parameters of multi-joint systems, particularly those systems exhibiting many sources of variability, such as the human hand. In this work, **we present a novel data-driven method that uses *workspace manifold mapping* and exponential representations of 3D rotations to estimate the kinematic parameters of multi-joint systems where direct measurement is infeasible**.

2.2.1 Kinematic Parameter Estimation for Rigid Robots

For systems comprised of rigid links, such as traditional robotic arms, specifying the position and orientation of the end-effector requires identification of kinematic parameters, such as link dimensions and joint axis orientations. The field of robot calibration emerged due to a need for implicit parameter measurement [29]. In a typical implementation, a robotic end-effector is calibrated across a series of predefined poses [31]. For each pose, an external reference system, such as a motion capture system, provides a ground truth measurement that is used to refine a kinematic model of the robot. Similar techniques are described in [29, 32, 31].

There are significant challenges to applying the aforementioned robot calibration techniques to biological systems for estimating kinematic parameters. A robot's design specifications provide estimates for link lengths and revolute joint location and orientation. These kinematic parameter estimates can be refined using precise calibration poses and direct measurement of joint angles by encoders. In contrast, design specifications do not exist for biological systems and it is not possible to command precise, repeatable calibration poses.

2.2.2 Kinematic Parameter Estimation for Human Hands

The kinematic parameters of biological systems, such as the human hand, are generally unknown. Benoit et. al performed kinematic parameter estimation for knee kinematics by precisely positioning inertial measurement units (IMUs) according to intra-cortical bone pins inserted into the proximal tibia and distal femur bones [33]. Due to the intricate and comparatively small nature of the human hand, however, it is difficult to precisely position sensors relative to anatomical landmarks non-invasively. In addition, estimation of kinematic parameters can be challenging due to factors such as skin motion artifacts, unexpected sensor movement relative to the skin [33], and soft tissue deformation [34].

A number of methods have been proposed to identify kinematic parameters non-invasively and without precise placement of sensors relative to anatomical landmarks [35]. For example, Taylor et al. identified parameters for a one degree-of-freedom (DOF) revolute joint using a method that exploits the circular motion of two reference frames attached to opposite sides of the joint [30]. This algorithm, called the Symmetrical Axis of Rotation Approach (SARA), is commonly implemented by commercially available motion capture system software, such as Nexus (Vicon, Colorado, USA) [9].

Seel, et al. proposed an alternative algorithm for estimating the kinematics of a 1-DOF revolute flexion/extension axis of the knee joint [36]. Two IMU sensors were attached to opposite sides of the joint. The algorithm relied explicitly on kinematic constraints inherent to the knee structure in order to determine the common rotation axis that linked the motion of the two rotating sensor frames.

Whereas kinematic parameters for 1-DOF joints can be quickly computed via optimization-based methods, the estimation of kinematic parameters for 2-DOF joints can be more challenging. Some estimation methods designed for 1-DOF joints can be applied to 2-DOF joints if motion is constrained such that the range of joint angles about the dominant axis is at least twice as large as the range about the non-dominant axis [37]. These methods, which track optical markers in 3D, rely on the fact that markers will trace out a circle around the dominant axis. When the non-dominant motion becomes too large, the performance of these methods deteriorates.

Muller et al. [2] developed a similar measurement system to the 1-DOF method developed by Seel, et al. [36] to determine the optimal joint axes for a 2-DOF elbow model. Seel, et al.'s

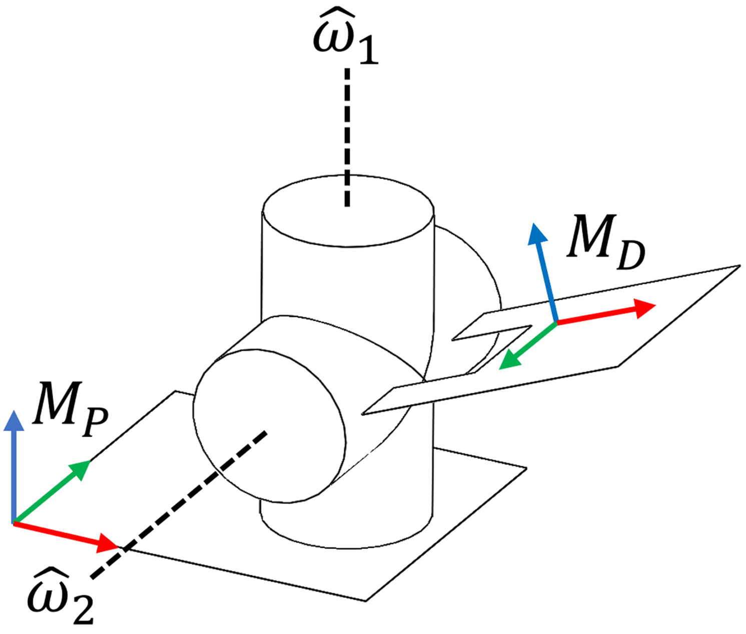

Figure 2.1: Schematic of a 2-DOF kinematic chain. Two reference frames (Distal: $M_D \in SO(3)$, Proximal: $M_P \in SO(3)$) are attached to rigid bodies on either side of two intersecting revolute joints. The joint axis unit vectors $\hat{\omega}_1, \hat{\omega}_2$ are expressed in the world frame.

1-DOF method was combined with numerical optimization methods and gradient descent in order to identify the joint axes in real time using arbitrary motions of the kinematic system. The 2-DOF method was validated in an experiment where IMUs were placed on proximal and distal sides of the elbow joint.

Todorov presented a probabilistic framework for estimating the orientations of 2-DOF joint axes that outputs the most likely kinematic parameters as well as confidence measures for the estimated parameters [1]. Treating the unknown quantities (joint angles and joint axis orientations) as state variables of a dynamical system allows the use of an extended Kalman filter for parameter estimation. One limitation of the approach in [1] is that a sufficiently close a priori estimate of joint axis orientations is required in order to properly initialize the optimization routine and converge to accurate kinematic parameter values. Convergence does not always occur.

As in Seel, et al. [36], Muller, et al. [2], and Todorov [1], we apply a data-driven method to discrete-time observations of 1-DOF and 2-DOF motion in order to estimate joint axis orientations. In contrast to existing algorithms that focus on temporal properties of the motion dataset, we exploit the geometric structure of the motion dataset. Whereas previous methods incrementally update the parameter estimates as new data arrives, our method leverages a batch of discrete-time measurements. We adopt the exponential representation of 3D rotation [38] and characterize a workspace manifold structure that is inherent to the discrete-time kinematic dataset.

We refer to our novel method as "workspace manifold mapping." Critical to our method is a novel implementation of the Generative Topographic Mapping (GTM) approach [39], which is a data-driven method developed for parameterizing the structure of manifolds. In this paper, we show how GTM can be modified to model the exponential representation of the 3D kinematics of a 2-DOF, revolute joint system. We demonstrate that our proposed geometric method offers several advantages over prior approaches for kinematic parameter estimation in terms of accuracy, reliability, and computational time.

The structure of this paper is as follows. In Section 4.4, we introduce the mathematical formulation for solving the kinematic parameter estimation problem using workspace manifold mapping. In Section 2.4, we describe the experimental procedure and performance evaluation using computer simulated data and experimental data from a motion capture system. In Section 2.5, we discuss a number of advantages and limitations of our novel kinematic parameter estimation approach. Section 2.6 concludes with a summary of findings along with directions for future work.

2.3 Framework for Kinematic Parameter Estimation via Workspace Manifold Mapping

The following section describes the process by which the kinematic parameter estimation problem is solved using *workspace manifold mapping*. It begins with an introduction to the requisite techniques in rigid body kinematics and robot modeling. These techniques lead to a key geometric insight in Subsection 2.3.3 that enables the proposed method of parameter estimation.

The objective of the proposed workspace manifold mapping method is to identify the orientations of the axes of rotation that comprise a 2-DOF system. The example shown in this work is based on a 2-DOF joint with intersecting axes, as for the metacarpophalangeal (MCP) joint of the human hand [40]. However, the mathematical formulation applies to a 2-DOF kinematic relationship between any two bodies, regardless of the physical construction of the joint. The revolute joint axes need not be orthogonal, intersecting, or part of the physical system (e.g. "virtual" joints, as in models of the human thumb [41]). For convenience, we will refer to the 2-DOF kinematic relationship as a 2-DOF joint. It should be noted that, while link lengths are required for a complete set of kinematic parameters, link length estimation is beyond the scope of this work.

2.3.1 Problem Statement for a 2-DOF Joint

Kinematic parameter estimation of 1-DOF joints has been demonstrated widely using approaches such as those reviewed in Section 2.2. When the number of revolute joint axes between adjacent links is two or greater, the task of estimating joint axes is non-trivial.

Consider a kinematic system that is comprised of two rigid bodies connected by a series of one or more revolute joints. More specifically, consider a kinematic chain with two intersecting revolute joint axes $[\omega_1, \omega_2]$ in 3D space (Fig. 2.1). Given our interest in the study of human hand motion, we demonstrate our method for the metacarpophalangeal joint, which is typically modeled as having two intersecting and orthogonal axes of rotation [40]. However, our method does not strictly require the joint axes ω_1 and ω_2 to be orthogonal. We affix a distal reference frame M_D and a proximal reference frame M_P to rigid bodies on adjacent sides of the 2-DOF joint (Fig. 2.1). Using observations of relative rotations between reference frames M_D and M_P, we will estimate the revolute joint axes ω_1 and ω_2.

We can express the relative 3D rotation F of the distal reference frame M_D relative to the proximal reference frame M_P as the output of a forward kinematics function **FK** that takes as input the revolute joint axes in unit vector form $\hat{\omega}_1$ and $\hat{\omega}_2$ and joint angles θ_1 and θ_2, about each axis of rotation, respectively.

$$F = \mathbf{FK}([\theta_1, \theta_2]; \hat{\omega}_1, \hat{\omega}_2) \tag{2.1}$$

Using an inverse kinematics function $\mathbf{FK}^{-1}$, joint angles can be calculated using kinematic parameters and a relative 3D rotation F.

$$[\theta_1, \theta_2] = \mathbf{FK}^{-1}(F; \hat{\omega}_1, \hat{\omega}_2) \tag{2.2}$$

For both forward and inverse kinematics, the kinematic parameters are assumed to be known, which allows for the calculation of relative rotation via (2.1) or joint angles via (2.2) (Fig. 2.2).

We assume that the orientations of reference frames M_D and M_P are observable, while

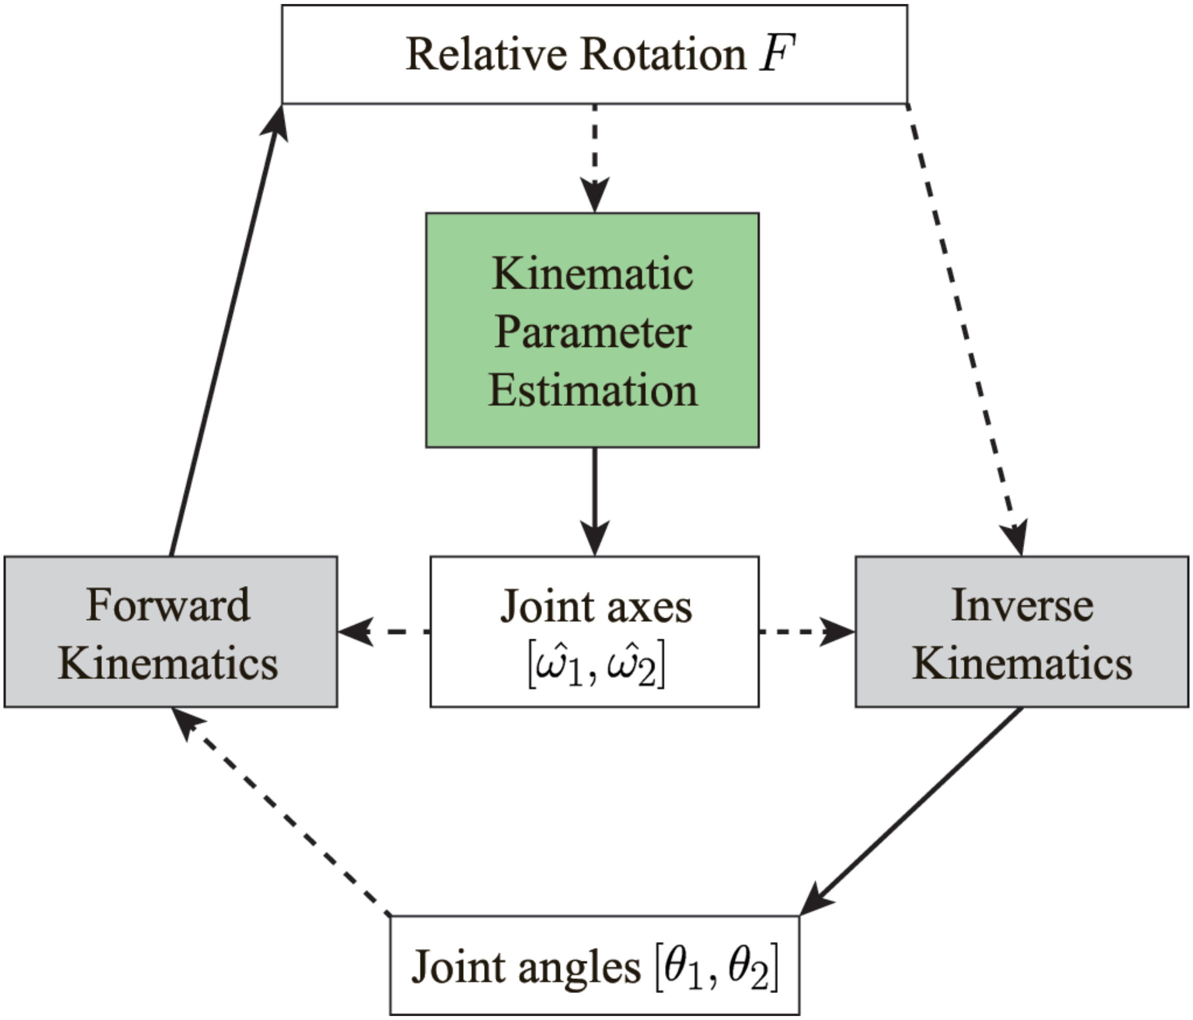

Figure 2.2: Flowchart of quantities and processes related to kinematic chains. White boxes denote the quantities used to represent system parameters. Shaded boxes denote processes that produce an output quantity. Dotted and solid lines represent process inputs and outputs, respectively. Notably, unlike forward and inverse kinematics which can produce a single output for a single input, kinematic parameter estimation requires a collection of end-effector orientation observations.

the constant joint axes and time-varying joint angles are unknown. Under these assumptions, the estimation of the two 3D revolute joint axes in unit vector form $\hat{\omega}_1$ and $\hat{\omega}_2$ poses a six-parameter problem. The estimation of the joint angle trajectories $\{\theta_{1,k}, \theta_{2,k}\}_{k=1}^{N}$ becomes a 2N-parameter problem, where N is the total number of discrete-time observations indexed by k.

Given a set of N observations of relative rotations, $\{F_k\}_{k=1}^N$, our goal is to estimate the kinematic parameters $[\hat{\omega}_1, \hat{\omega}_2]$ and joint angles $\{\theta_{1,k}, \theta_{2,k}\}_{k=1}^N$ that minimize the error $\mathcal{E}$ between the observed rotation F_k and the estimated rotation $\tilde{F}_k$. We define the following objective function $\mathcal{L}$ for the 2-DOF kinematic parameter estimation problem.

$$\underset{\hat{\omega}_1, \hat{\omega}_2, \{\theta_{1,k}, \theta_{2,k}\}}{min} \mathcal{L} = \sum_{k=1}^{N} \mathcal{E}(F_k, \tilde{F}_k) \tag{2.3}$$

The kinematic parameter optimization problem stated in (2.3) is nonlinear, and the solution is not straightforward. In general, closed-form gradients for parameters of kinematic chains can be derived [42], which allow for the use of gradient-based optimization methods. However, the computational cost is high due to operations such as integration and matrix inversion during the gradient evaluation in a high-dimensional space (in our case, a (6 + 2N)-dimensional space). Additionally, solutions to nonlinear optimization problems are sensitive to initial conditions and convergence is not always guaranteed.

Our proposed method for solving the kinematic parameter estimation problem expressed by (2.3) does not rely on computing closed-form gradients. Rather, the workspace manifold mapping approach relies on characterizing a unique geometry that is inherent in the observations $\{F_k\}_{k=1}^N$ when represented using exponential coordinates.

2.3.2 Robotics-based Geometric Representation of 3D Kinematics

Here we provide a brief overview of the exponential representation of 3D rotations and how a product-of-exponentials formulation designed for robotics can be applied generally to kinematic chains. The exponential representation is essential to our geometry-based approach to the kinematic parameter estimation problem. This work builds upon a product-of-exponentials framework for the analysis of robotic systems, which was introduced by Lynch and Park [38].

2.3.2.1 Exponential Representation of 3D Rotation

Underlying the algorithm presented in this paper is the rotation matrix group SO(3), the Lie group of all 3D rotations [38]. A set of observations of relative 3D rotations $\{F_k\}_{k=1}^{N}$ for a kinematic chain, as in Fig. 2.1, can be represented as elements of SO(3), which are 3 × 3 real matrices, $R \in SO(3)$. These matrices have the properties $R^T R = I$ and $det(R) = 1$.

The matrix Lie group $SO(3)$ has an associated Lie algebra, $so(3)$, which is the set of 3 x 3 real skew-symmetric matrices. We can represent an arbitrary 3D rotation vector $\omega = [\omega_x, \omega_y, \omega_z]^T \in R^3$ in skew-symmetric form $[\omega] \in so(3)$.

$$[\omega] = \begin{bmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{bmatrix} \tag{2.4}$$

The relationship between $SO(3)$ and $so(3)$ is the matrix exponential map $R = e^{[\omega]}$. The exponential map, along with its inverse, allows us to transform between $SO(3)$ and $so(3)$.

For any rotation $R \in SO(3)$, the vector $\omega \in \mathbb{R}^3$ represents the 3D orientation of a revolute axis and the magnitude of a rotation about that axis. It is convenient to represent the axis orientation and magnitude separately as $\omega = \hat{\omega}\theta$, where the unit vector $\hat{\omega}$ represents the unit 3D axis of rotation, and θ represents the magnitude (angle) of rotation about the axis. The quantities $\hat{\omega}$ and θ are known as the *exponential coordinates* used for the exponential representation of 3D rotation, and they form the foundation of our proposed kinematic parameter estimation method.

A rotation about $\hat{\omega}$ by $+\pi$ produces the same $SO(3)$ element as a rotation about the same axis by $-\pi$. As such, the set of all 3D rotations in terms of ω can be represented as a solid sphere with radius of π, with antipodal points representing identical rotations. A single instance of ω is shown in Fig. 2.3 as an example.

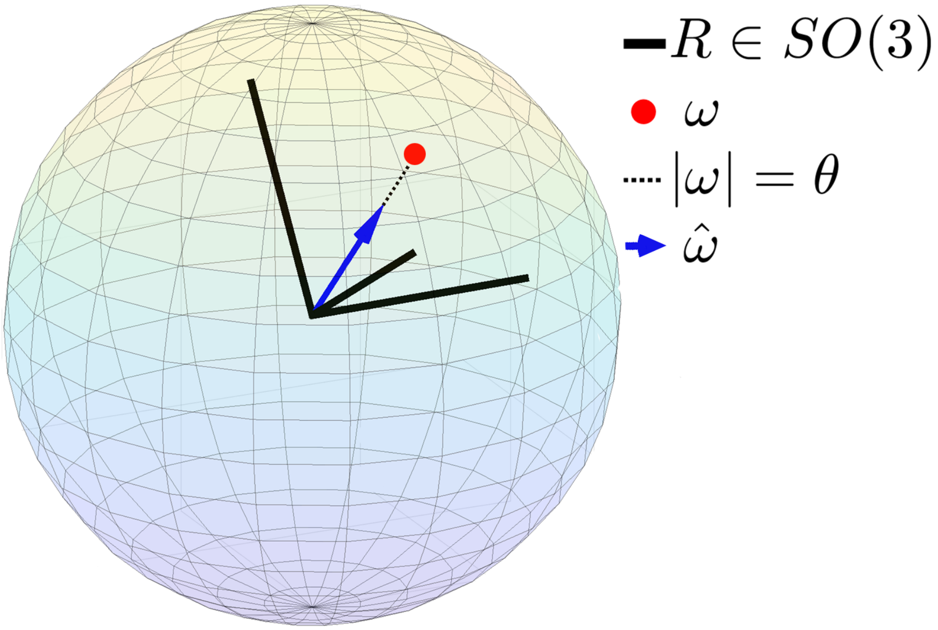

Figure 2.3: Exponential representation $\omega = \hat{\omega}\theta$ of an arbitrary rotation $R \in SO(3)$. The unit vector in the direction of ω represents the unit axis of rotation between the identity matrix and R, and θ represents the rotation magnitude. The sphere has a radius of π radians.

Any 3D rotation $R \in SO(3)$ can be expressed using exponential coordinates and the initial orientation of the rotated frame $R_0 \in SO(3)$.

$$R = e^{\hat{\omega}\theta} R_0 \tag{2.5}$$

That is, a coordinate frame initially given by R_0 is rotated about a unit axis $\hat{\omega}$ by an angle θ to become R. The quantities R_0, R, and $\hat{\omega}$ are represented in the world frame.

2.3.2.2 Product-of-exponentials (PoE) Formulation for Kinematic Chains

The product-of-exponentials (PoE) formulation is derived directly from the exponential representation of 3D rotation. When applied to kinematic chains, the PoE formulation enables the representation of the kinematics of a serial linkage without the need for explicit, a priori expressions of intermediate reference frames, as with the Denavit-Hartenberg parameter approach to kinematic representation [38]. The ability to represent kinematics using only base (proximal) and end-effector (distal) reference frames is particularly beneficial when one's goal is to estimate the location and orientation of multiple joint axes between the frames.

Applying the PoE formulation to the kinematics of the 2-DOF system in Fig. 2.1, we assert the following expression for the rotation of a distal body with respect to the world frame.

$$M_{D,k} = e^{[\omega_P]} e^{[\hat{\omega}_1]\theta_{1,k}} e^{[\hat{\omega}_2]\theta_{2,k}} M_{D,0} \tag{2.6}$$

where $[\hat{\omega}_1]$, $[\hat{\omega}_2] \in so(3)$ represent the intersecting joint axes, $[\theta_{1,k}, \theta_{2,k}]$ represent the angles of rotation about the joint axes for the k^{th} discrete observation, and $M_{D,0}$ represents the initial orientation of the distal reference frame. Additionally, $e^{[\omega_P]}$ represents the rotation of the proximal reference frame M_P with respect to the world frame.

$$e^{[\omega_P]} = M_{P,k} M_{P,0}^{-1} \tag{2.7}$$

The relative 3D rotation $F \in SO(3)$ (or $f \in so(3)$) represents the rotation of the distal reference frame from the perspective of the proximal reference frame. We define the relative rotation $F \in SO(3)$ and its associated Lie algebra $f \in so(3)$ expression as follows.

$$F \triangleq e^{[\hat{\omega}_1]\theta_1} e^{[\hat{\omega}_2]\theta_2} \tag{2.8}$$

$$f = log(F) = log(e^{[\hat{\omega}_1]\theta_1} e^{[\hat{\omega}_2]\theta_2}) \tag{2.9}$$

We can also express the relative 3D rotation F in (2.8) using rotations expressed in the world frame.

$$F = M_{P,0} M_P^{-1} M_D M_{D,0}^{-1} \tag{2.10}$$

Importantly, (2.10) can be evaluated using observations, denoted by F_k and $f_k = log(F_k)$, that can be collected using a motion capture system or IMUs, for example. The kinematic parameters and joint angles in (2.8) and (2.9) are unknown and will be estimated by leveraging experimental data and (2.10).

2.3.3 The Workspace Manifold

In this subsection we describe how the workspace of a kinematic chain can be visualized in 3D using exponential coordinates. We show how 1-DOF and 2-DOF systems of revolute joints result in canonical forms that can be characterized and used to estimate kinematic parameters and joint angles simultaneously.

Recall from Section 2.3.2.1 that any 3D rotation $R \in SO(3)$ can be represented in exponential coordinates by $\hat{\omega}\theta$, where $\hat{\omega}$ is the unit joint axis vector of the rigid body rotation, and θ is the joint angle about that axis.

Consider a kinematic chain with a 1-DOF revolute joint whose joint angle can vary from $\theta = [-\pi, \pi)$. If the system is moved throughout the entirety of its reachable space, the workspace manifold would be represented as a line in $so(3)$ that passes through the center of a 3D sphere (Fig. 2.4a). The linear form of the workspace manifold reflects the fact that there is a single revolute joint axis and the axis has a constant 3D orientation $\hat{\omega}_1$. While the radius of the 3D sphere is π, the length of the line depends on the allowable joint angle range.

Now, consider the extension of this visualization to a kinematic chain with two intersecting revolute joints. As is common for simple kinematic models of MCP joints of the human hand, we have chosen an orthogonal, intersecting pair of axes $\hat{\omega}_1$, $\hat{\omega}_2$ [40]. When the 2-DOF system moves within its workspace, the observations of the end-effector orientation within $so(3)$ take on a unique form. A surface emerges that is a smooth 2D manifold in $so(3)$ and is uniquely determined by the two joint axes (Fig. 2.4b).

The collection of $so(3)$ observations lie on the *workspace manifold* of the kinematic chain that represents the set of all possible relative rotations between the proximal and distal reference frames of the kinematic chain. All points that lie on this workspace manifold surface are explicitly determined by (2.9). The surface is curved, but there exist two straight lines that lie completely within the surface. Those two lines are equivalent to the two unit joint axes $\hat{\omega}_1$ and $\hat{\omega}_2$.

That there exist two straight lines in $so(3)$ that lie completely within a given 2-DOF workspace manifold arises from the fact that the observations represent a combination of two revolute joint motions. If the second joint angle trajectory were fixed at zero, then the system would appear to be 1-DOF, and the workspace manifold would collapse into a single line representing $\hat{\omega}_1\theta_1$, where $\theta_1 = [-\pi, \pi)$. The best fit line to that 1D manifold is the $\hat{\omega}_1$ joint axis. Likewise, if the first joint angle trajectory were fixed at zero, then the workspace manifold would collapse into a single line representing $\hat{\omega}_2\theta_2$. Combining motion for both axes, the manifold is curved due to the non-commutativity of 3D rotations. It is important

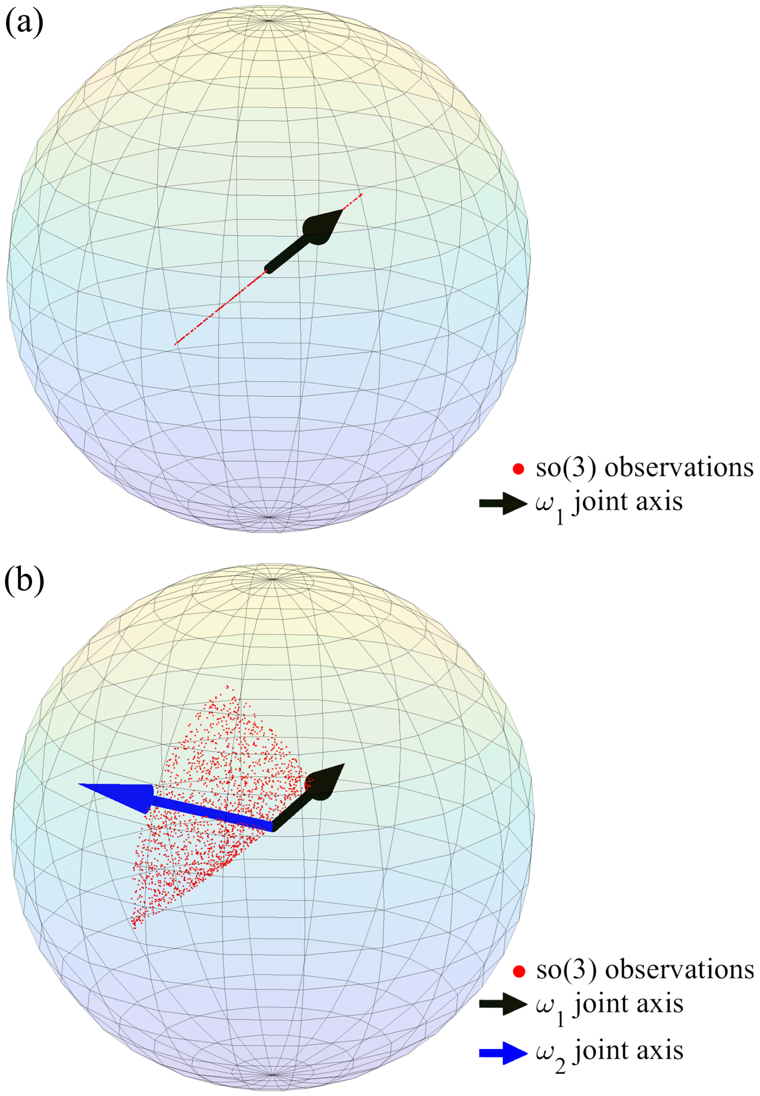

Figure 2.4: (a) 1-DOF workspace manifold comprised of a straight line (red). The red points indicate elements of $so(3)$ corresponding to arbitrary relative rotations of a 1-DOF kinematic system with angles $\theta \in (-180, 180]$ about a common axis $\hat{\omega}_1$. (b) 2-DOF workspace manifold comprised of a curved surface (red). The red points indicate elements of $so(3)$ corresponding to arbitrary relative rotations of the 2-DOF kinematic system represented in (2.9) with two example axes $\hat{\omega}_1$ (black) and $\hat{\omega}_2$ (blue).

to note that the intersection of $\hat{\omega}_1$ and $\hat{\omega}_2$ in $so(3)$ is a result of the $so(3)$ visualization and is entirely independent of whether the two joint axes intersect physically in the kinematic chain. The geometry of the workspace manifold is discussed further in Section 2.5.2.

The task of estimating the joint axes from observed motion of 1-DOF and 2-DOF kinematic chains is, therefore, reduced to the task of finding the lines that exist completely within the observed workspace manifold. The following subsections detail our estimation method that leverages this geometric interpretation of the workspace manifold.

2.3.4 Parameter Estimation of the Workspace Manifold

The workspace manifold is a collection of relative rotations between proximal and distal reference frames in a kinematic chain, which is a 2-DOF system in our example (Fig. 2.1). The PoE formulation provides an analytical approach for expressing a relative 3D rotation in terms of joint axis orientations and joint angles. The problem of estimating joint axis orientations given a single observation of a relative rotation is underdetermined since the joint angles are also unknown.

2.3.4.1 Latent Variable Models

In order to estimate joint axes when joint angle trajectories are unobserved, we consider a *latent variable model*. Latent variable models such as the hidden Markov model [43] are useful when observations of a system can be explained by a function of unobserved variables of the system.

The expectation-maximization (EM) algorithm is a commonly used iterative method for estimating the observed and unobserved (latent) variables [44]. The kinematic parameter estimation problem is well-suited for this approach, as we can treat the joint axis orientations as the model parameters and the joint angles as the model's latent variables. Given an initial estimate for the latent variables, the EM algorithm determines an estimate for the model parameters that maximizes the log-likelihood function of the model given the observed data.

The updated model parameter estimates produce updated estimates for the latent variables, which are then fed back into the likelihood maximization. This process is repeated until the model parameters converge. The following subsubsection explains the process by which we manipulate (2.9) such that maximum likelihood estimates of the joint axes can be found.

2.3.4.2 Generative Topographic Mapping (GTM) Algorithm

The GTM algorithm is a geometrically-inspired, latent variable model that can be used to estimate the parameters of a nonlinear mapping function between a high-dimensional data representation (observations) and a low-dimensional latent representation (latent variables) [39]. Points in a high-dimensional data space $\mathbb{R}^D$ are assumed to have intrinsic dimensionality that is of a lower dimension, called the latent space $\mathbb{R}^L$. A simple example is a curved 2D manifold in $\mathbb{R}^3$ similar to that shown in Fig. 2.4b.

The GTM is often used as a method of nonlinear dimensionality reduction (i.e. $\mathbb{R}^D \rightarrow \mathbb{R}^L$) for visualization of high-dimensional data [39, 45, 46]. One of the GTM's utilities is the estimation of the mapping function $f(\Theta|W), \mathbb{R}^L \rightarrow \mathbb{R}^D$, which maps points from the low-dimensional latent space representation to their corresponding points in the high-dimensional data space $\mathbb{R}^D$. The latent variables are denoted by Θ and the mapping parameters are denoted by W.

The algorithmic structure of the GTM algorithm is based on that of the expectation-maximization algorithm. However, the GTM approach modifies two central components: the definitions of the *prior distribution* and the *likelihood function*. The prior distribution θ

is comprised of a grid of points in latent space $\mathbb{R}^L$. The GTM makes an initial hypothesis for the mapping function parameters W, which results in a hypothetical distribution $\tilde{f}$ in $\mathbb{R}^D$. In general, the distribution f appears to be that same grid of points, transformed to lie on a smooth L-dimensional manifold. The GTM uses a likelihood function that compares the set of observations $\{f\}$ with the hypothetical manifold formed by $\tilde{f}$. Maximizing this likelihood function with respect to the mapping parameters produces the optimal estimates of the latent space representation, which are the joint angle trajectories θ_1, θ_2 in our example.

Bishop et al. proposed a generalized linear regression model for the nonlinear mapping that uses a vector of basis functions $\Phi(\Theta) = [\phi_1(\Theta), \phi_2(\Theta), ...]$ where Θ are the latent variables, and$W$ [39] is the matrix of weight coefficients.

$$f(\Theta; W) = W\phi(\Theta) \tag{2.11}$$

This generalized linear regression model can approximate any mapping from $\mathbb{R}^L \to \mathbb{R}^D$ as a sum of Gaussians. However, the number of basis functions and scaling of the Gaussian distributions must be selected a priori. Improper scaling of the Gaussian distributions can significantly affect the smoothness of the estimated manifold and is commonly avoided by hand-tuning the relative magnitudes of the weights in W [39].

A key contribution of our work is the introduction of a principled alternative to Bishop et al.'s generalized model for the nonlinear mapping [39] that does not require hand-tuning the relative weight magnitudes W in (2.11). Specifically, we modify the generalized linear regression model to reflect the kinematic constraints expressed by (2.9). In the following subsection, we describe our alternative model, which is explicitly defined in terms of the 2-DOF kinematic equations.

2.3.5 GTM Algorithm with Kinematic Constraints (GTM-KC)

Here we present our method of modifying the GTM algorithm for estimating the joint axes in kinematic chains. **As our method introduces kinematic constraints into the GTM formulation, we refer to our algorithm as the *GTM algorithm with kinematic constraints (GTM-KC)*.**

We seek a nonlinear mapping function with the following matrix form:

$$\mathbb{F} = \boldsymbol{\Phi}(\theta)\boldsymbol{W}(\omega) \in \mathbb{R}^{N\times 3} \tag{2.12}$$

where the discrete-time observation matrix $\mathbb{F} = [vec([f_1]), ..., vec([f_k]), ..., vec([f_N])]^T$ for $[f_k] \in so(3)$ comes from the nonlinear mapping (2.9) and contains a collection of observations of the end-effector orientation. The mapping parameters ω represent the joint axes of the kinematic chain, and the joint angles are represented by θ.

The kinematic equation laid out in (2.9) is not immediately compatible for use with the GTM algorithm, which requires the product of two functions as in (2.12): a set of basis functions of the form $\boldsymbol{\phi}(\theta)$ and a parameter matrix $\boldsymbol{W}(\omega)$. At the same time, we recognize there is no closed-form simplification of the right-hand side, $log(e^{[\hat{\omega}_1]\theta_{1,n}}e^{[\hat{\omega}_2]\theta_{2,n}})$ since rotations in $SO(3)$ do not generally commute.

To ensure compatibility between the GTM algorithmic structure and the kinematics representation in (2.9), we employ a numerical approximation to the product of matrix exponentials, called the Baker-Campbell-Hausdorff (BCH) formula [47]:

$$\begin{aligned} BCH(A,B) &= log(e^{A}e^{B}) \\ &= A + B + \frac{1}{2}\langle A, B\rangle \\ &+ \frac{1}{12}(\langle A, \langle A, B\rangle - \langle B, \langle B, A\rangle) + ... \end{aligned} \tag{2.13}$$

where angle brackets denote the commutator, $\langle A, B\rangle = AB - BA$. In the case where A and B commute, then $\langle A, B\rangle = 0$ and we are left with the trivial solution, $log(e^{A}e^{B}) = A + B$ (note: this does not apply to $SO(3)$, in general). The BCH formula is an effective and appropriate numerical approximation of the log of the product of two matrix exponentials, as it represents the solution as a series of nested commutator operations, which are easily accessible in the case of the group $SO(3)$.

For clarity, (2.14) shows the BCH formula explicitly applied to the nonlinear mapping (2.9):

$$\begin{aligned} f_k = log(F_k) &= log(e^{[\hat{\omega}_1]\theta_{1,n}}e^{[\hat{\omega}_2]\theta_{2,k}}) \\ &= \hat{\omega}_1\theta_1 + \hat{\omega}_2\theta_2 + \frac{1}{2}\Big\langle \hat{\omega}_1\theta_1, \hat{\omega}_2\theta_2\Big\rangle \\ &+ \frac{1}{12}\Bigg(\Big\langle \hat{\omega}_1\theta_1, \Big\langle \hat{\omega}_1\theta_1, \hat{\omega}_2\theta_2\Big\rangle\Big\rangle \\ &- \Big\langle \hat{\omega}_2\theta_2, \Big\langle \hat{\omega}_2\theta_2, \hat{\omega}_1\theta_1\Big\rangle\Big\rangle\Bigg) + ... \end{aligned} \tag{2.14}$$

where we have modified the commutator operator for vectors, $\Big\langle \omega_1\theta_1, \omega_2\theta_2\Big\rangle \triangleq vec\Big(\Big\langle [\omega_1\theta_1], [\omega_2\theta_2]\Big\rangle\Big)$.

We can express the BCH series in (2.14) as a matrix formula by executing the commutators and recognizing that joint angles $\theta_{1,n}, \theta_{2,n}$ are scalars that can be pulled out of the commutator terms. This allows us to separate (2.14) into the product of two matrix functions as desired in (2.12):

$$\begin{gathered}\boldsymbol{\Phi}(\theta_{1,n}, \theta_{2,n}) = \\ \begin{bmatrix} \theta_{1,n} & \theta_{2,n} & \frac{1}{2}\theta_{1,n}\theta_{2,n} & \frac{1}{12}\theta_{1,n}^2\theta_{2,n} & \frac{-1}{12}\theta_{1,n}\theta_{2,n}^2 \end{bmatrix}\end{gathered} \tag{2.15}$$

$$\boldsymbol{W}(\omega_1, \omega_2) = \begin{bmatrix} \omega_1 \\ \omega_2 \\ [\omega_1, \omega_2] \\ \left[\omega_1, [\omega_1, \omega_2]\right] \\ \left[\omega_2, [\omega_2, \omega_1]\right] \end{bmatrix} \tag{2.16}$$

Separating (2.14) into a product of basis functions that operate on the latent variables (2.15), and the kinematic parameter matrix (2.16) allows the nonlinear mapping of the GTM to be replaced by the 2-DOF kinematics representation of (2.9).

2.3.5.1 PCA Initialization

Iterative optimization methods, of which the GTM-KC method can be considered a member, require an initial estimate where the search within the solution space begins. For the optimization method to be robust against convergence to a suboptimal minimum, it is critical that the initial condition be carefully chosen.

The GTM-KC method determines its own initial condition by solving a low-order approximation of the kinematics equation (2.14). Instead of fitting a workspace manifold to the observations in $so(3)$, the initial iteration of the GTM-KC algorithm fits a plane. This plane is found by performing PCA on the observations. The first two principal components of the data are used as the initial conditions for ω_1 and ω_2.

2.3.5.2 Inverse Kinematics via the GTM-KC Method

Reconstructing the motion of a kinematic chain requires the constant joint axes and time-varying joint angles. With the constant parameters identified via the GTM-KC method, the task becomes estimating the joint angle trajectories via inverse kinematics. The matrix form of our kinematics equations (2.12) can provide insight into the joint angle trajectories that created the end-effector orientation trajectories F. Solving for $\Phi(\theta)$ in a least-squared sense leads to:

$$\tilde{\boldsymbol{\Phi}}(\theta) = \mathbb{F}\boldsymbol{W}(\omega)^{+} \tag{2.17}$$

where we take the Moore-Penrose pseudoinverse of the parameter matrix W. This provides joint angle trajectory estimates, as the first two columns of $\boldsymbol{\Phi}$ are θ_1 and θ_2, respectively. The system of equations being solved in (2.17) is highly nonlinear, so solutions are not unique. In Section 2.4.2, we show that while (2.17) is not sufficient for inverse kinematics alone, the joint angle trajectory estimate is accurate enough to act as the initial condition for typical numerical methods of inverse kinematics, such as the Newton-Raphson method [48].

2.4 Experimental Procedure and Evaluation

With the method now developed, we present a series of case studies that test the performance of the GTM-KC kinematic parameter estimation algorithm and compare its performance against four benchmark methods. Performance metrics for the GTM-KC algorithm were obtained from simulations of an idealized 2-DOF system (Fig. 2.1). With the simulated data we justify a key algorithm design choice and evaluate the robustness of the algorithm against sensor noise.

Additionally, we present the results of an experiment in which a motion capture system was used to measure the movement of a 3D-printed kinematic chain similar to an MCP joint. Using the motion capture data, we provide performance metrics for the GTM-KC algorithm alongside four alternative methods [30],[36],[2],[1].

2.4.1 Performance Metrics

Here, we present the metrics used to assess the performance of our GTM-KC method and benchmark kinematic parameter estimation techniques. We assess the performance of the methods on simulated kinematic data and data from motion capture experiments.

The application of a parameter estimation algorithm on collections of observations of 2-DOF motion will produce estimates for the joint axis orientations $[\omega_1, \omega_2]$. A batch of observations, when fed into the parameter estimation method, produces joint axis estimates. The error of an estimated joint axis from a single batch is defined as the angle between the estimated axis unit vector and the "ground-truth" axis unit vector ω_{GT}. In order to further assess estimation performance, we collect multiple batches for a given kinematic chain. These batches have a common pair of ground-truth axis vectors such that the mean estimated axis $\bar{\omega}$ and standard deviation of the estimation algorithm from batch to batch can be evaluated. For a given number of batches B, the metrics of bias and standard deviation will be used:

$$bias = \angle\big(\bar{\omega}, \omega_{GT}\big)$$

$$std\ dev = \sqrt{\frac{1}{B-1}\sum_{b=1}^{B}\angle\big(\omega_b, \omega_{GT}\big)^2}$$

where ω_b represents the estimated joint axis for batch index b, and $\angle\big(\omega_i, \omega_j\big)$ denotes an operator that calculates the angular difference between two axes ω_i and ω_j.

2.4.2 Performance Assessment: Simulated Model

The first case study of the GTM-KC algorithm was to evaluate estimation performance using simulated data from an idealized 2-DOF kinematic model. Additionally, estimation of 1-DOF motion was evaluated by constraining one joint from the idealized model.

We chose biomechanically relevant joint parameters $[\omega_1, \omega_2, \theta_1, \theta_2]$ using an idealized metacarpophalangeal (MCP) joint of a human index finger as a model. The joint axes were chosen to be orthogonal and intersecting.

First, we generated a pair of orthogonal joint axes, which we created by using the first two columns of a randomly generated $SO(3)$ matrix:

$$\begin{bmatrix} | & | & | \\ \omega_1 & \omega_2 & * \\ | & | & | \end{bmatrix} = R_x(\theta_x)R_y(\theta_y)R_z(\theta_z) \qquad (2.18)$$

$$: \big[\theta_x, \theta_y, \theta_z\big] \sim unif(0, 2\pi)$$

where $R_x(\theta_x)$ represents a rotation about the x axis by the amount given by θ_x, and similarly for y and z. The distribution $unif(a, b)$ notes that $\theta_x, \theta_y, \theta_z$ were drawn from a uniform distribution between a and b.

This method of selecting joint axes enforced orthogonality, since any two columns of an $SO(3)$ matrix are necessarily orthogonal. While we assume orthogonality of ω_1 and ω_2 to represent the MCP joint as in [40], there is no assumption of orthogonality when solving for ω_1 and ω_2 using the GTM-KC method. Assigning θ_1 as an adduction/abduction joint angle and θ_2 as a flexion/extension joint angle, we imposed bioinspired constraints l_1, l_2 on a random sampling of the joint angle space:

$$\begin{bmatrix}\theta_{1,k} & \theta_{2,k}\end{bmatrix} \sim \begin{bmatrix}unif(-l_1, l_1) & unif(0, l_2)\end{bmatrix}, \tag{2.19}$$

where $l_1 = 30°, l_2 = 135°$ based on [49], $k = 1 : N$, and the number of simulated data points N was set to 2000.

One could choose a more biomimetic approach by generating joint angle trajectory functions that more realistically represent the motion of a human finger (e.g. sinusoids or other continuous functions). Given that the performance of the algorithm is independent of the shape and smoothness of the joint angle trajectories, we simply sampled joint angles from uniform distributions as in (2.19). In order to create the timeseries of $so(3)$ observations s, we expressed the forward kinematics in terms of our generated joint parameters as per the process outlined in (2.5) through (2.8):

$$M_k = e^{[\omega_1]\theta_{1,k}} e^{[\omega_2]\theta_{2,k}} M_0 \tag{2.20}$$

A matrix logarithm was then used to create a timeseries in $so(3)$ according to (2.9).

To simulate 1-DOF input data, we simply set $\theta_{2,k} = 0, k = 1 : N$ and use the same generative method as in the 2-DOF case.

We then applied the GTM-KC algorithm to the simulated kinematic data. We generated 100 datasets for both 1-DOF and 2-DOF motion according to the procedure described in (2.19) - (2.20). Table 2.1 and Fig. 3.5 show the GTM-KC algorithm's estimation results for idealized 1-DOF and 2-DOF motion.

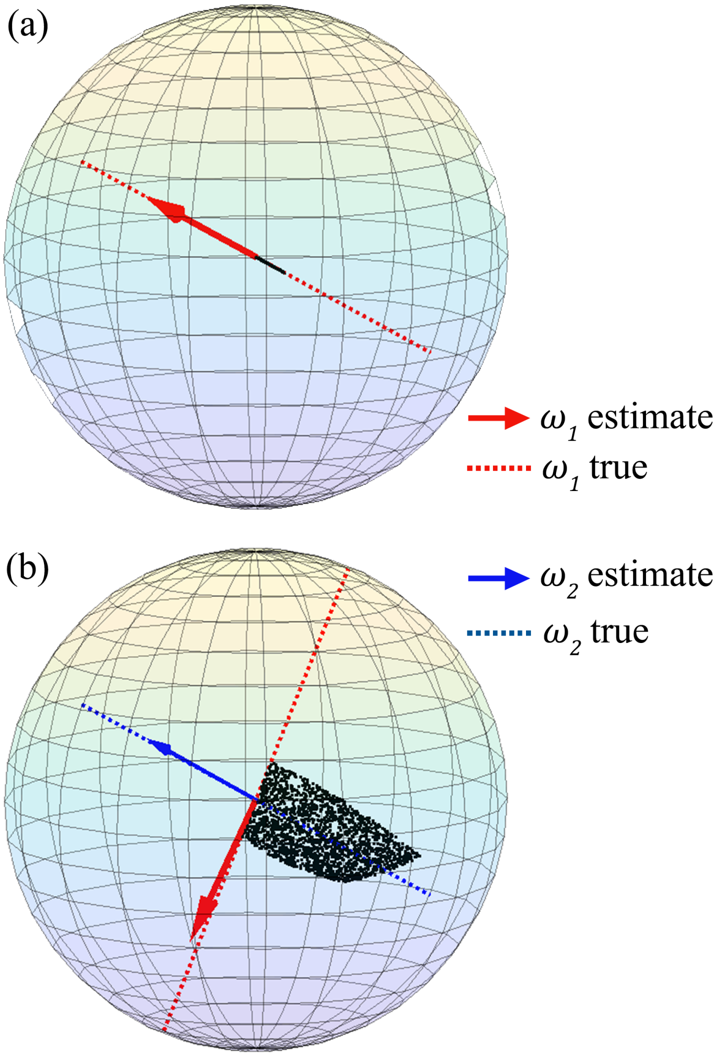

Figure 2.5: (a) 1-DOF and (b) 2-DOF simulation results. Figures show simulated observations (black points), estimated ω_1 and ω_2 axes (red and blue arrows, respectively), and ground truth ω_1 and ω_2 axes (red and blue dotted lines, respectively).

2.4.2.1 Sensitivity to Noise Magnitude

In practice, body-worn sensors suffer from sensor slip and skin motion artifacts. In addition to sensor measurement noise, these artifacts prevent the observed data from lying perfectly within the workspace manifold. In simulation, we can investigate the robustness of the algorithm to noise that is purposely added to each of the simulated end-effector observations s_k.

$$s_k = log(e^{[\omega_1]\theta_{1,k}} e^{[\omega_2]\theta_{2,k}}) + \epsilon_k \tag{2.21}$$

where $\epsilon_k \in \mathbb{R}^3$ represents a stochastic error vector that is added to each simulated observation of M_D. The magnitude of the error vector was drawn from a uniform distribution having a lower bound of 0° and an upper bound that ranged from 3° to 30° in 3° increments. The direction of the error vector was randomized using the first column of a randomly generated $SO(3)$ matrix as in (2.18). The data collection procedure was comprised of 100 trials per maximum noise magnitude.

Fig. 3.6 shows the error in the estimated rotation axes as a function of the maximum noise magnitude $\|\epsilon\|_{max}$. When applied to kinematic data with no added noise, the GTM-KC algorithm yielded axis estimates with an average bias of 0.57°. The bias and standard deviation of the estimated axes increased as the maximum noise magnitude increased. The degradation of performance of the GTM-KC algorithm at increased noise amplitude indicates it should only be used in cases with an acceptable amount of noise. In order to determine this noise threshold, pairwise t-tests were performed between the zero-noise axis estimates and estimates with increasing noise. The null hypothesis (that the axis estimates from a dataset with elevated noise come from the same distribution as the zero-noise dataset) failed to be rejected above 8° of noise. Thus, we can say that the GTM-KC is robust to noise up to 8°.

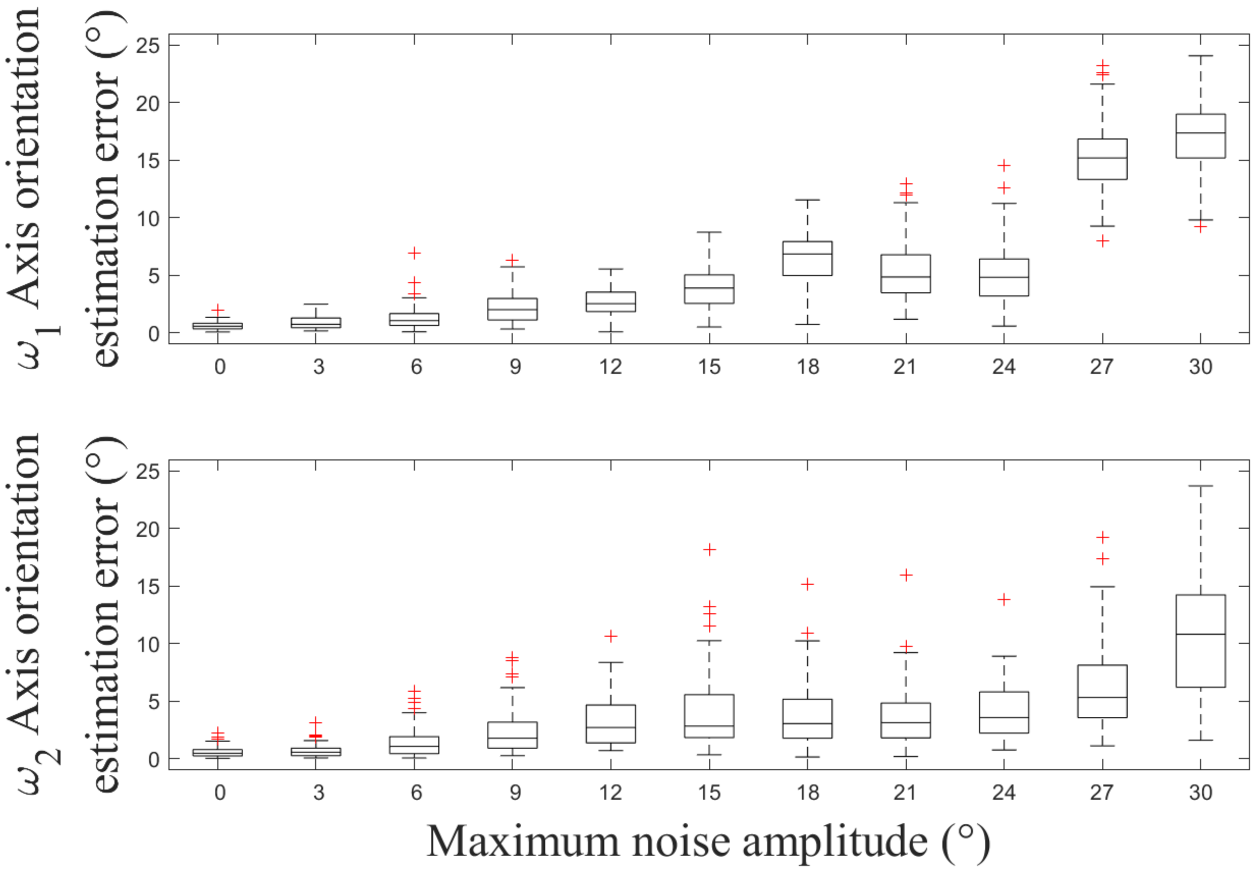

Figure 2.6: Sensitivity of the GTM-KC algorithm to simulated noise for a 2-DOF kinematic chain is shown using box and whisker plots. Each boxplot indicates the 25th, 50th, and 75th percentiles. The whiskers extend to the most extreme data points that are not considered outliers ("+") having values of more than 1.5 times the interquartile range from the top or bottom of the box. The error shown is the pooled average across all axis estimates.

Table 2.1: Kinematic Parameter Estimation, Simulated Results

	1-DOF Simulation Results		2-DOF Simulation Results	
Joint Axis	Bias	Std. Dev	Bias	Std. Dev
ω_1	0.04°	0.13°	0.17°	0.87°
ω_2	-	-	0.03°	0.92°

Table 2.2: Sensitivity of joint axis estimates to order of the BCH approximation function

Order of Approx.	Bias	Std. Dev
1	11.6°	1.62
2	0.91°	1.71
3	0.87°	1.42

2.4.3 Sensitivity to Order of the BCH Approximation Function

The number of terms used in the BCH approximation function (2.14) is a design choice. With higher-order approximations comes higher accuracy, but at the cost of increased computational cost. We assess here the performance of the GTM-KC algorithm as a function of the number of BCH terms used in the mapping.

The nonlinear mapping $F = \mathbf{\Phi}(\theta)\boldsymbol{W}(\omega)$ in (2.12) is the matrix form of the BCH formula, which is an infinite series. The higher order terms contribute to the warped nature of the workspace manifold, with lower-order approximations of (2.14) more closely resembling a flat plane.

In Table 2.2, we show the relationship between the order of the BCH approximation function and estimation error for the simulated 2-DOF kinematic chain data. For an order-1 approximation, only the first two terms of the BCH formula were used. These two terms are simply the linear combination of θ_1 and θ_2. The order-2 and order-3 approximations additionally use nonlinear combinations of θ_1 and θ_2, as in (2.15) and (2.16). The inclusion of the order-2 term greatly improved the performance of the estimation (Table 2.2). Approximations beyond order-3 were tested, but the additional nonlinear terms had a diminishing effect on the performance of the GTM-KC method.

Based on the results from Table II, we recommend using an order-3 approximation to the kinematic equation of (2.14), as it yields accurate estimates at a negligible increased computational cost. In the motion capture study of **Section IV.B**, we used the order-3 approximation, which uses the first five elements as shown in (2.15) and (2.16).

2.4.4 Performance Assessment on Motion Capture Dataset

In order to experimentally validate the GTM-KC algorithm, we collected motion capture data from physical 1-DOF and 2-DOF systems. Critical to this experiment was a well-defined ground truth measurement of joint axis parameters. We designed a 3D-printed, passive mechanical linkage inspired by the dimensions and degrees of freedom of the human finger (Fig. 2.7). In particular, we modeled the MCP joint as two independent, orthogonal and intersecting degrees of freedom [40]. Observations of the orientations of bodies on opposite sides of the 2-DOF joint (M_D, M_P, as in Fig. 2.1) were used to estimate the kinematic parameters of the joint axes.

Evaluating the performance metrics described in Section 2.4.1, required an accurate ground truth measurement must be made. The mechanical setup shown in Fig. 2.7 makes clear the ground truth values for the kinematic parameters $\hat{\omega}_1, \hat{\omega}_2$. It is pictured in the initial pose M_0, with θ_1 and θ_2 set to zero. The initial orientations of the ground truth axes $\hat{\omega}_1, \hat{\omega}_2$ are assumed to be equivalent to the z and y axis of the proximal frame.

A motion capture system (six MX T40 cameras, Vicon, UK) was used to collect observations of motion capture markers affixed to the physical model at a sampling rate of 100 Hz. During data collection, the mechanical linkage was moved randomly through its workspace by an experimenter. The linkage's motion was kept within the bioinspired workspace limits described in Section IV.B. Two datasets were captured: The 2-DOF dataset, in which both axes were allowed to rotate, and the 1-DOF dataset, in which the joint axis ω_1 was fixed to only allow flexion/extension of the distal frame. The total number of observations for each

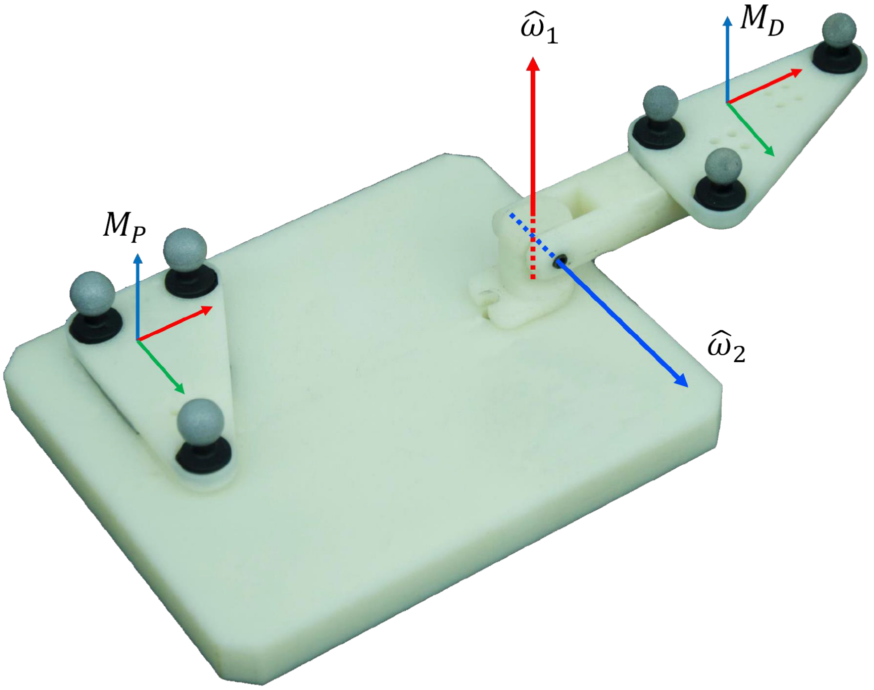

Figure 2.7: Mechanical linkage with motion capture markers representing reference frames that are distal (M_D) and proximal (M_P) to a 2-DOF joint, $[\omega_1, \omega_2]$.

dataset was 60,000. The $so(3)$ observations of a representative subset of 2-DOF motion from the motion capture study are shown in Fig. 2.8.

In order to evaluate the performance of the GTM-KC algorithm and benchmark methods, each method was applied to a number of subsets of the dataset. Each dataset was partitioned into 20 equal batches, such that each batch contained 3,000 sequential observations and no batches overlapped. The ground truth kinematic parameters were held constant within each dataset, so estimates from batch to batch could be evaluated according to the performance metrics laid out in Section 2.4.1.

Fig. 2.9 and Table 2.3 show the performance of the GTM-KC method and benchmark methods applied to the motion capture datasets. On the 1-DOF dataset, the GTM-KC was compared to the symmetrical axis of rotation (SARA) method [30], Todorov's method [1],

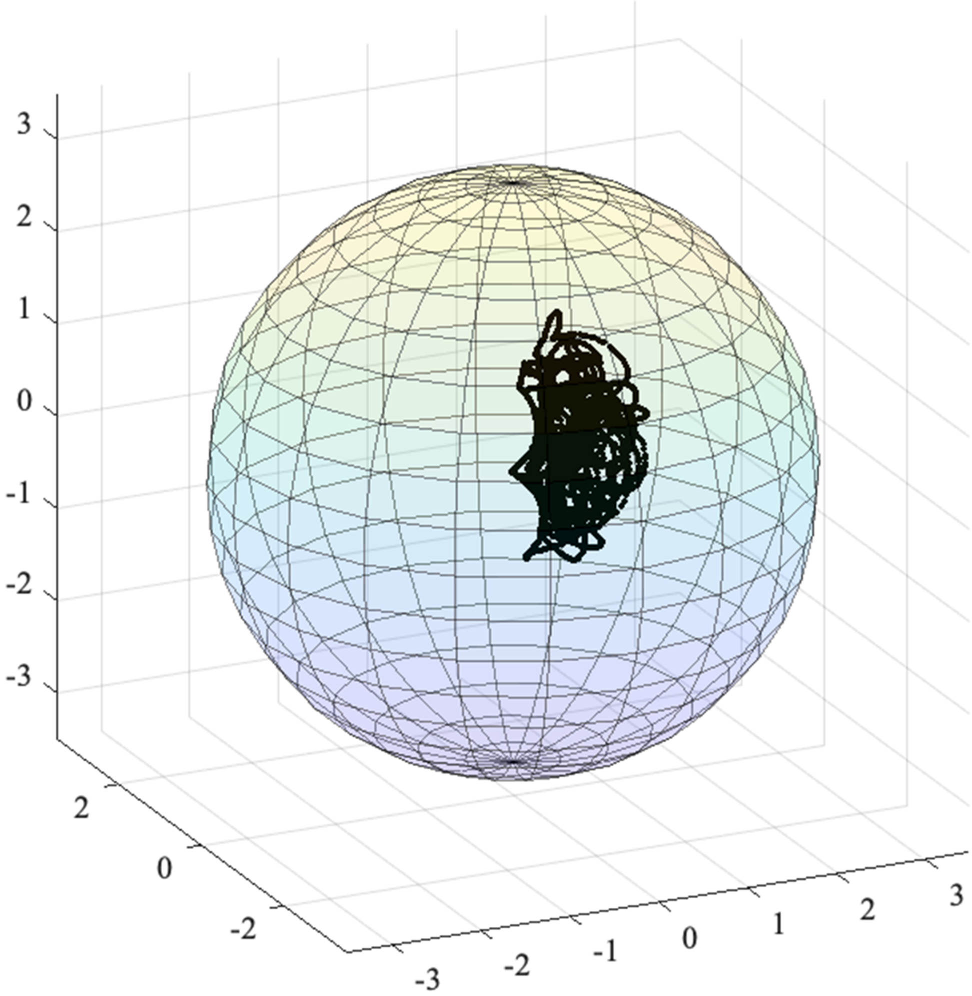

Figure 2.8: A continuous subset of observations of 2-DOF motion taken from the motion capture study.

and the method of Seel et al. [36]. For the 2-DOF dataset, the GTM-KC was compared to the methods of Todorov [1] and Muller et al [2].

On the 1-DOF motion capture dataset, the GTM-KC outperforms the benchmarks with respect to bias, standard deviation, and computation time. Within the $N = 20$ dataset the 1-DOF estimates were robust, in that each algorithm did not include any significant outliers (bias $> 15°$).

Table 2.3: Parameter estimation results for 1-DOF and 2-DOF physical models

	1-DOF Physical Model (N=20)			2-DOF Physical Model (N=20)				
Algorithm	ω_2 Bias(°)	ω_2 Std. Dev (°)	Computation Time (s)	ω_1 Bias(°)	ω_1 Std. Dev (°)	ω_2 Bias(°)	ω_2 Std. Dev (°)	Computation Time (s)
GTM-KC	1.41 (0)	0.20	0.0008	2.46 (0)	3.42	2.41 (0)	2.70	1.72
Seel et al.	2.11 (0)	0.39	0.14	1.52 (0)	4.01	2.29 (4)	9.86	8.54
Todorov	0.85 (2)	0.97	17.01	1.64 (0)	2.68	5.01 (4)	7.41	25.64
Todorov	0.76 (3)	1.10	16.99	4.25 (2)	21.67	2.59 (6)	22.02	25.71
SARA	1.74 (0)	0.22	*					

Parenthetical entries in the $[\omega_1, \omega_2]$ bias columns represent the number of significant outliers for that algorithm (bias > 15°) within the N=20 experiment.

* The SARA method is implemented within the VICON Nexus proprietary development environment, so a comparable computational time estimate could not be taken.

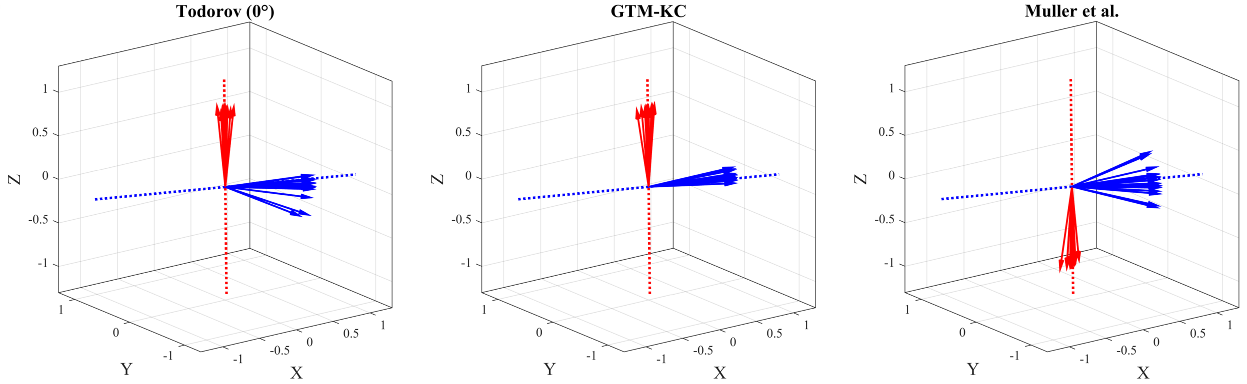

Figure 2.9: 2-DOF joint axis orientation estimation results from the GTM-KC method compared to two existing estimation methods [1, 2]. dotted lines represent ground truth measurements of ω_1 (red) and ω_2 (blue) joint axis orientations.

The behavior of the benchmark algorithms was less repeatable on the 2-DOF set than on the 1-DOF dataset. As such, the performance metrics listed in Table 2.3 do not include significant outliers. In the 2-DOF case, the GTM-KC and the method of Muller et al [2]. both produced biases below 3° for both axes, with zero and four total outliers, respectively. The standard deviation of the estimates of the GTM-KC method were not significantly different from that of Muller et al.'s method for the ω_1 axis ($\sigma = 4.09°$ vs 4.01°, respectively), and outperformed Muller et al.'s method in estimating the ω_1 axis ($\sigma = 3.41°$ vs 9.86°, respectively).

The protocol for applying Todorov's method in [1] to the motion capture dataset was modified in light of the method's sensitivity to initial conditions. An a priori estimate of each joint axis was required to seed their optimization-based algorithm. Using the ground truth values of the joint axes as the initial conditions resulted in joint axis estimates with low bias and low standard deviation (e_1, e_2=(1.64° ,5.01°), σ_1, σ_2=(2.68°,7.41°)). The robustness of Todorov's algorithm also appeared to be better with ideal initial conditions, with zero ω_1 outliers and four ω_2 outliers. To reflect the uncertainty of a real-world scenario, an additional evaluation of Todorov's algorithm was performed, where the a priori joint axis estimates were selected to be 5° offset from the ground truth values. This resulted in increased ω_1 bias (e_1=(4.25°)), increased standard deviations ((σ_1, σ_2)=(21.67°,22.02°)) and an increase in the number of outliers (2 and 6, for ω_1, ω_2, respectively). Without explicit knowledge of Todorov's procedure of selecting these initial conditions, we present these split results in a good-faith attempt at evaluation relative to the other methods.

2.4.5 Assessment of Inverse Kinematics via the GTM-KC

The Product-of-Exponentials (PoE) kinematic model, upon which the GTM-KC performs parameter estimation, can also be used to estimate joint angle trajectories. Using the inverted axis estimation equation (2.17), we solved for joint angle trajectories $[\theta_1, \theta_2]$ on a representative sample of the motion capture dataset (Fig. 2.10).

In addition to the "Inverse-GTM-KC" inverse kinematics results, results from a standard Newton-Raphson iterative root finding method were also provided [48]. For a highly nonlinear joint angle parameter space, fast and robust convergence of the Newton-Raphson method requires a set of initial joint angle estimates that is close to the true joint angle values. As such, the Newton-Raphson results seen in Fig. 2.10 were initialized by the results

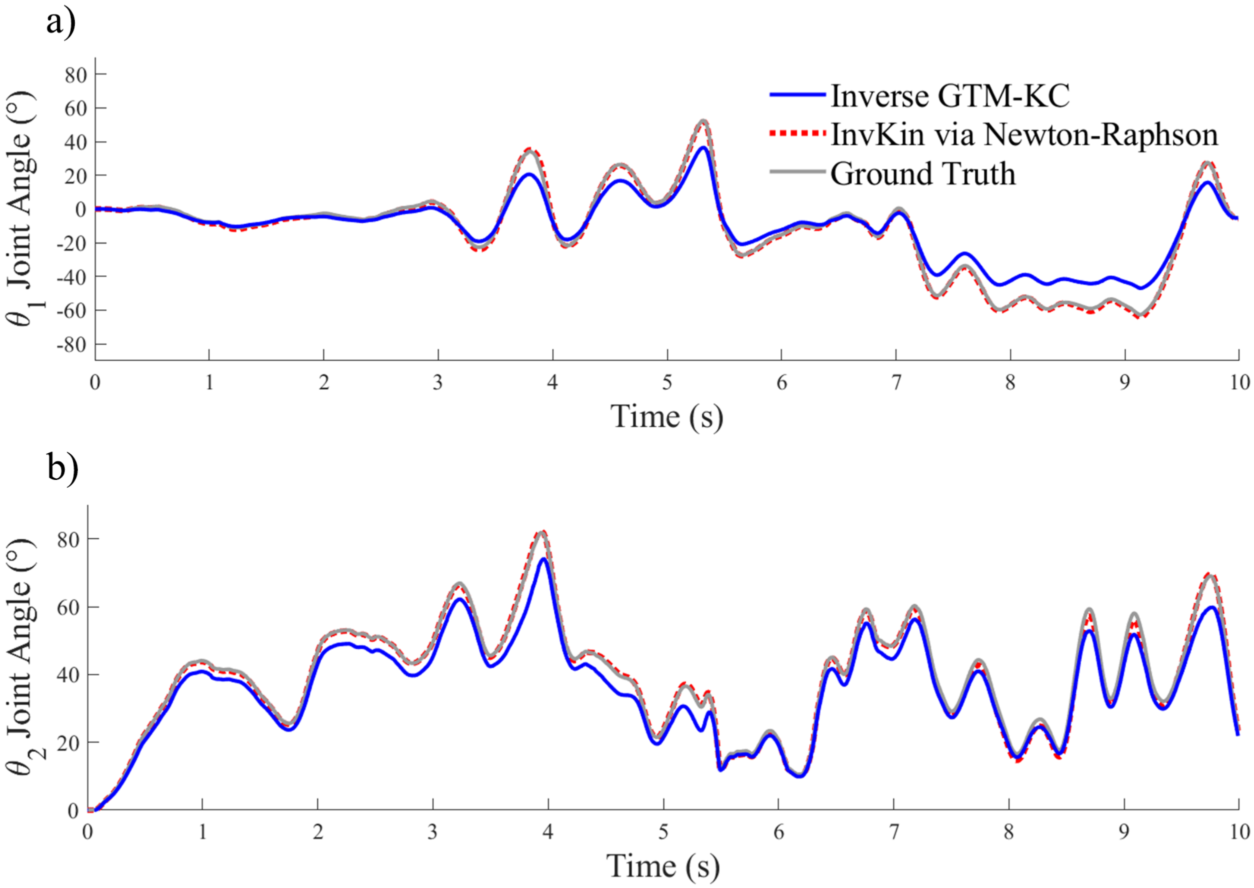

Figure 2.10: The joint angle estimates for a) θ_1 and b) θ_2 are shown for two inverse kinematics approaches: using a Newton-Raphson method (dashed line) and using the inverse of the GTM-KC equation (solid thick line). The true joint angle values are shown with a solid thin line.

from the Inverse-GTM-KC method. While the Inverse-GTM-KC is an intriguing prospect as it provides a single-shot estimate of the inverse kinematics, its performance is inferior to that of a typical iterative Newton-Raphson root-finding method when that method is first seeded by the Inverse-GTM-KC (mean error of 11.2° vs. 1.4°, respectively).

2.5 Discussion

2.5.1 Comparison: GTM-KC and Benchmark Methods

There were two general families of algorithms available for estimating the parameters of a kinematic chain: (i) real-time iterative methods whose estimates converge after a finite number of steps [2],[1] and (ii) a post-process method which relies on the geometry of a sufficiently large set of observations (the GTM-KC). One stated advantage of the iterative methods was the ability to estimate the constant and varying kinematic parameters in real time, with new observations updating the parameter estimates. While the GTM-KC could be modified to compete in this regard, the focus of this paper was to estimate the constant axes representing the MCP joint. With a sufficient number of observations (3000 for the motion capture study), the GTM-KC can produce joint axis estimates accurate to within 3°, which can then be used as a calibrated kinematic model for animation and further motion analysis.

A tricky aspect of the comparison study of Section 2.4.4 was the diversity of methods to address the initial conditions of an optimization. The 2-DOF algorithm of Muller et al. used a hard-coded set of initial conditions for their optimization algorithm which appear to generalize to any set of joint axes [2]. As per Fig. 2.9 and Table 2.3, Todorov's algorithm appears to perform better when the initial conditions are close to the ground truth axes.

The immunity of the GTM-KC to initial condition sensitivity comes from it producing its own initial condition via PCA. This low-dimensional linear model of the kinematics appears to produce estimates which are close enough to the optimal axis orientations such that the nonlinear GTM-KC method does not produce significant outliers.

2.5.2 Geometry of the Workspace Manifold

At the core of this work is the geometry of the observations of 2-DOF kinematics in $so(3)$. The workspace manifold can be understood as the set of all output points of the PoE formula $log(e^{[\hat{\omega}_1]\theta_1}e^{[\hat{\omega}_2]\theta_2})$, with a static pair $\hat{\omega}_1, \hat{\omega}_2$ defining the shape of the manifold, and variables θ_1, θ_2 defining points along that manifold. The PoE formulation and the exponential representation of rotation are often touted for their geometrical interpretation, which our estimation method leverages.

The curved nature of the workspace manifold arises due to the fact that rotations are non-commutative. If the rotations defined by $e^{[\hat{\omega}_1]\theta_1}$ and $e^{[\hat{\omega}_1]\theta_1}$ commute, then the exponents in the PoE formula combine in the trivial fashion. In this case, according to (2.14) the workspace manifold would be defined as $\omega_1\theta_1 + \omega_2\theta_2$, the range of which is a plane in $\mathbb{R}^3$. Since the rotations do not commute, we need the nonzero commutator terms from the BCH formula found in Equation (2.14). As more terms from (2.14) are added, the workspace manifold deviates from this plane.

In reality, the $so(3)$ observations of end-effector orientation are only partially described by the workspace manifold; the real-world observations do not exist precisely on the manifold. The observations deviate from the surface due to approximation errors from the BCH formula in (2.14), sensor noise, and sensor drift. Performing the GTM-KC method can be understood as finding the best-fit 2-DOF workspace manifold that fits the data. The surface normal distance between the best-fit manifold and a given end-effector observation can be interpreted as a component of motion that is not allowed for that specific 2-DOF system. As an example, if the distal sensor in Fig. 2.1 started rotating around an axis orthogonal to both ω_1 and ω_2, this would appear in the sphere of $so(3)$ to be a point moving off of the 2-DOF workspace manifold, along a line which is normal to the manifold. Performing inverse kinematics for a given observation can be understood as projecting the end-effector observation onto the 2-DOF workspace manifold and finding the corresponding parameters $[\theta_1, \theta_2]$, which minimizes the distance to the end-effector observation.

We chose the PoE formulation as a model-based method for describing the two joint axes, and the GTM-KC algorithm as an approach that is built directly on top of the PoE. Alternative model-free methods to estimate the joint axes are also possible, which could identify the two straight lines that best fit within the manifold without explicitly utilizing the PoE formulaic description of the manifold. However, we believe the model-based method has the most merit, since it provides explicit prior knowledge to GTM-KC algorithm.

2.5.3 Limitations

We describe here the limitations of the method introduced in this paper, and the challenges that persist in real-world implementation.

We believe a primary limitation is that the GTM-KC output is sensitive to the quality of motion used to create the input dataset. As an illustrative example, consider a 2-DOF motion dataset where one of the degrees of freedom was restricted to only allow negligible rotations. It is likely that the estimation of this restricted axis would suffer compared to an axis with higher magnitude rotations. While this sensitivity to the characteristics of the input data was addressed briefly in 2.4.2, a more detailed characterization of input data sensitivity is required in order to develop guidelines for proper use of the GTM-KC method with human hand motion.

A critical difference betweeen the GTM-KC method and the benchmark methods is that the GTM-KC is intended to be used off-line. Whereas the 2DOF methods of Muller et. al [2] and Todorov [1] incrementally improve joint axis estimates as new data is fed into the algorithm, the GTM-KC was developed to produce axis estimates only after a batch of data was previously collected. In our study, we collected 3,000 samples over 30 seconds to be used as input for the axis estimation algorithm. The convergence of the iterative Expectation Maximization model underlying the GTM-KC was not directly assessed; the comparison study in 2.4.4 used an $N = 3,000$ dataset, but the performance behavior of the GTM-KC was not assessed at smaller sample sizes.

The case studies presented in 2.4 assumed that the rotation axes of the 2-DOF joint were orthogonal. While we believe that the GTM-KC does not rely on that assumption, the generalization of the GTM-KC algorithm to non-orthogonal axes was not demonstrated.

The primary application of this method, as discussed throughout this paper, has been the study of human hand motion. While we believe that the 3d-printed finger model is a reasonable facsimile to the MCP joint, our work did not evaluate the kinematic parameter estimation method on actual human hand motion.

2.6 Conclusion and Future Work

We introduced the GTM-KC method as a novel and effective approach for estimating the kinematic parameters of 1- and 2-DOF kinematic chains of revolute joints. The GTM-KC method matches or outperforms benchmark algorithms [30, 36, 2, 1] in terms of accuracy and precision, and it outperforms in terms of computational efficiency.

A long-term goal of this work is to apply the GTM-KC method to develop accurate, subject-specific biomechanical models of the human hand. Merging the work discussed in this paper with state-of-the-art techniques in human hand motion capture should be considered as the first step toward this goal. A simple extension of the GTM-KC method should provide the ability to estimate the joint axes for multiple consecutive joints in a kinematic chain, such as a finger. This would involve recursively performing the algorithm; using the terminology introduced in Fig. 2.1, the sensor which is distal to the MCP joint would be used as the proximal sensor M_P while estimating the proximal interphalangeal joint (PIP) parameters.

A complete model of kinematic structures such as the human hand requires the estimation of joint lengths in addition to joint axis orientations. We posit that the extension of the GTM-KC method, founded upon the PoE formulation, provides a potential path towards addressing this challenge. The PoE formulation of open chain kinematics can be extended to include 3D position information about the end-effector. Readers interested in this extension are referred to [38]. This introduces the exciting prospect of building a complete biomechanical model of a multi-joint serial linkage by solving for 3D positions and orientations of joint axes simultaneously.

2.7 Acknowledgement

The authors would like to thank Tingern "Jeremy" Wong for assistance with data collection during the motion capture study, Meng-Hao Li for helpful discussions regarding the implementation of the benchmark algorithms, and Dr. Alexis E. Block for feedback on the manuscript. This work was supported in part by Office of Naval Research Award N00014-16-1-2468.

CHAPTER 3

Multimodal Tactile Sensor Skin for Underwater, Pressurized Environments

This chapter was based on work co-authored by Dr. Minyung Song, Dr. Arif M. Abdullah, Conner J. Smith, Daniel L. Balducci, and Dr. Jonathan D. Posner at the University of Washington.

3.1 Abstract

Challenges in underwater robotic manipulation are addressed by integrating tactile sensing systems that provides the sense of touch in uncertain environments, such as low-light or occluded scenarios where vision-based sensing may fail. We describe the design, fabrication, and characterization of a sensor-embedded soft skin capable of multimodal sensing (contact force and two axes of skin shear) submerged in water and under hydrostatic pressure up to 1000 kPa. Using a multilayer molding process, we embed liquid-metal strain gauges within a durable elastomeric skin that is molded around a solid finger core. We embed strain gauges with two different geometries in strategic locations relative to the surface of the finger to measure both skin shear and contact force. We demonstrate that the sensor skin is capable of measuring forces up to 220 N under water and hydrostatic pressure.

3.2 Introduction

The skin on the human fingertip contains approximately three thousand sensitive nerve endings that detect stimuli such as skin strain, vibration, and temperature. Artificial tactile sensing systems attempt to emulate human fingertips through tactile perception. Tactile sensors have progressively embraced the use of highly deformable materials and their applications in dexterous robotic systems have been increasingly explored because they have the potential to positively transform lives and work practices [50, 51, 52, 53, 54, 55, 56, 57, 58, 59, 60]. Robotic systems can reduce occupational risk for humans when employed in hazardous environments. Robotic systems can handle heavy loads, toxic substances, and repetitive tasks. High-resolution, deformable tactile sensors that use cameras (e.g., GelSight [61], TacTip [62], DIGIT [63], etc.) or multiple transducer signals (e.g. BioTac [64]) have high-dimensional sensor output that can encode complex contact information [65, 66, 67, 68, 69]. These full finger tactile sensors may be problematic to integrate into robotic manipulation systems [51] due to space requirements and sensor geometry. This results in designers having to make compromises in the robotic finger to accommodate the tactile sensing hardware [52].

Flexible skin tactile sensing systems wrap around arbitrary geometries and can conform to surfaces, such as a robotic end effector. A common technique is to produce sensor arrays based on flexible printed circuit boards (PCBs), which can bend uniformly around curved surfaces [70, 71, 72]. PCBs can be carefully designed for complex 3-dimensional geometries; however, they lack the continuous flexibility of elastomer-based tactile sensing systems. Elastomer-embedded soft sensors, in which the sensing elements deform along with the embedding material, allow direct measurement of local strain [56], [73, 74, 75]. Flexible versions of standard electrical components, such as piezoresistors, capacitors, and transistors have been successfully integrated into sensing systems [76, 77, 78, 79]. Soft artificial skin comprised of liquid metal-filled microchannels within a Polydimethylsiloxane (PDMS) substrate are capable of sensing both shear and normal forces [79]. Liquid metal-based tactile sensors achieve high-resolution contact localization and multimodal tactile sensing (i.e. measurement of contact force, skin shear, and vibration) [74, 80].

Top-side controlled vehicle systems (e.g. remotely operated vehicles, unmanned underwater vehicles, autonomous underwater vehicles) with robotic manipulators are being employed to access environments with challenging underwater conditions in applications such as oceanography, deep sea biology, renewable marine energy, drilling for oil and gas, and a wide range of defense applications [81, 82, 83]. Tactile sensors may play an important role in underwater environments where unpredictable disturbances, underwater currents, and poor visual perception in murky waters cause uncertainties in manipulation [81, 84, 85, 86]. To meet the requirements of underwater environments, tactile sensing systems should operate underwater and under significant hydrostatic pressures and should be able to sense large loads [81, 86, 87, 88]. Industrial underwater robots with rigid grippers can exert grip forces of more than 6 kN at depths up to 11 km [89]. Soft robotic grippers have been tested for grip strength up to 150 N in the ocean at depths of 2.4 km [90]. Literature in tactile sensing for underwater manipulators shows a trade-off between maximum operating depth and maximum grip force. As part of the TRIDENT gripper project, a novel optoelectronic tactile sensor was developed that was capable of estimating contact force with grip forces up to 150 N at depths of 5-25 m [91, 92, 93]. A soft fiberoptic tactile sensing system was integrated into the SeeGrip underwater gripper, which performed highly accurate estimation of contact force distribution for grip forces up to 10 N at extreme pressures in a laboratory tank up to 600 bar (6 km ocean depth) [84, 86]. To date, these two examples represent the state of the art in advanced tactile perception as it relates to deep operation and large grip forces.

In this paper, we describe the development and characterization of a multimodal tactile sensor skin—normal strain and two orthogonal directions of shear strain—that operates under large applied loads in an underwater environment at pressures up to 1000 kPa. The tactile skin is constructed using liquid metal-filled microchannel embedded resistive strain gauges encapsulated in multiple layers of three-dimensionally embedded in silicone elastomer. The skin is characterized in dry conditions and underwater at hydrostatic pressure conditions from 0-1000 kPa. We show that the skin force response to loads of up to 220 N is independent of environmental hydrostatic pressure.

3.3 Experimental Methods

This section presents the design and methodologies used to develop and evaluate a multimodal skin tactile sensor array for use deep under water. To balance the performance requirements of tactile sensing with the robustness requirements of underwater operation, we encapsulate an array of liquid metal-based soft strain gauges within a durable, three dimensionally embedded in elastomeric skin, as shown in Figure 1. This multi-layered approach balances the competing skin requirements of robustness and force sensitivity.

We fabricate the sensor through a sequence of molding processes. First, we cast a base layer of elastomer onto a solid finger core. Next, liquid metal filled microchannel sensing elements (taxels) are fixed on the base layer. Finally, an outer layer of elastomer is overmolded to embed the taxels into an all-elastomer tactile sensor. Taxels are embedded in a soft elastomer matrix and experience larger strains in response to a given stress than would be expected if placed directly adjacent to a rigid core [94]. We incorporate circular spiral taxels to measure normal forces and rectangular serpentine channels to measure shear forces [55, 79]. Normal taxels, which measure contact force between the finger and an object, are placed directly in fingerpad areas of anticipated contact. Shear taxels, which measure circumferential skin strain, are placed along the perimeter of the finger at locations of maximum strain and away from locations of direct contact with objects [55].

3.3.1 Liquid metal taxels

The taxels are constructed from microchannels cast in PDMS and filled with liquid metal. The master molds for the microchannels are made of FR4 printed circuit boards with electroless nickel immersion gold surface finish. The taxel geometry was designed with a PCB design tool (EasyEDA, Guangdong, China) and fabricated by a commercial vendor (ChinaPCBOne, Guangdong, China). The measured patterned features were 150 m wide and 60 m tall (3D Optical Profilometer VR-6000, Keyence America). We coat the PCB master molds with Trichloro(1H,1H,2H,2H-perfluorooctyl) silane (Sigma-Aldrich, St. Louis, MO) in a vacuum chamber to reduce the adhesion between the master and the silicone elastomeric layers. The taxels were molded from 10:1 (base: crosslinker) Poly(dimethylsiloxane) (PDMS) (SYLGARD™ 184, Dow corning) mixed using a planetary mixer (THINKY mixer ARV-310, THINKY corporation) for 2 min at 2000 rpm.

We construct the microfluidic channels using two layers of PDMS (one planar and the other with the microfluidic features) that are bonded together (see Supporting Information, Figure S1). The molded channel layer is constructed by spin-coating (400 rpm, 1 min 30 secs) PDMS onto the microfluidic master and a silicon wafer twice to obtain a 500 m thick layer. The first PDMS layer was cured at 80 ºC for 12 min before applying the subsequent coating. Reservoirs are punched in the channel ends using a 15 AWG blunt tip needle to produce holes near the inlet and outlet of each taxel. The planar layer is fabricated using the same process as the channel layer on a silicon wafer. The partially cured layers are brought into contact and cured together at 80 ºC for 30 min to develop a complete microfluidic channel [95, 96].

We fill the microchannels with liquid eutectic gallium-indium alloy (EGaIn) (Indalloy 60, Indium Corporation), using a vacuum filling method [97]. In brief, we place a drop of EGaIn on the inlet and outlet channel reservoirs and place the taxel in a vacuum chamber that is pumped down to 0.08 MPa and then returned to atmospheric pressure. After filling the microchannels, we inserted 29 AWG wire (9511T513, Calmont Inc.) into the EGaIn-filled reservoirs and sealed the inlets and outlets using a drop of the PDMS that is cured at 80°C for 30 min. (An image of the completed taxel is shown in Supporting Information, Figure S2).

3.3.2 Fabrication of taxel-embedded artificial skin

The three-dimensional artificial skin is fabricated using two elastomer molds around a rigid finger core and embedded taxels, as shown in Figure 1. The mold shell components were designed in CAD software (SolidWorks, Waltham, MA), 3D printed using a FDM printer (Ultimaker, Ultimaker S3), and assembled via through bolts (Supporting Information, Figure S2).

The first mold forms a 2 mm thick base layer of skin around the core, as shown in Figure 1a (see Supporting Information, Figure S2). We pour silicone-based elastomer (DS-30, Smooth-On, Macungie, PA) mixed with 10% by weight silicone thinner (Silicone ThinnerTM, Smooth-On, Macungie, PA) (10% by weight of the DS-30 mixture) into the mold. The elastomer is placed in a vacuum chamber (–90 kPa) for 20 min to remove any gas bubbles, cured at room temperature and atmospheric pressure for 2 h and then removed from the mold. Next, we fix normal and shear force taxels onto the 2 mm skin (Figure 1a-b) using silicone adhesive (Sil-poxy, Smooth-On). We then over mold a 2 mm thick elastomer layer to embed the taxels and lead wires in elastomer (Supporting Information, Figure S2). The completed artificial tactile skin with two normal force sensors (one on each of the proximal and distal faces) and three shear force sensors (on distal face) is shown in Figure 1c-e. The shear force taxels on the distal face are arranged to detect shear in two in-plane directions of the distal face.

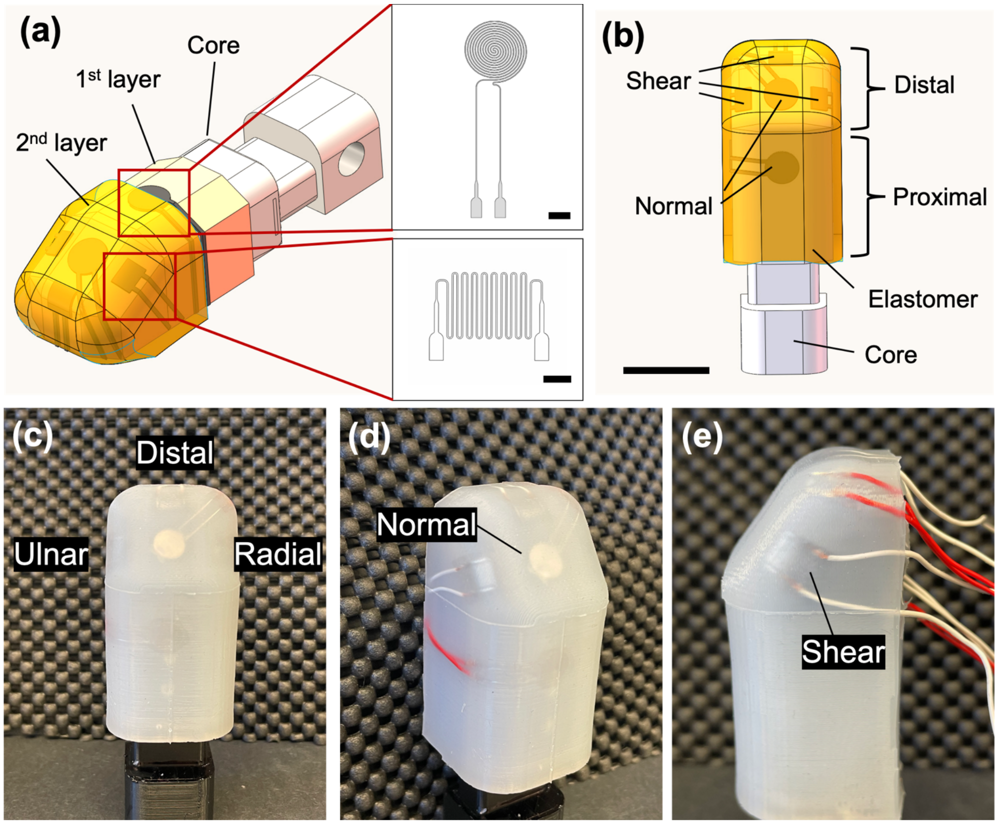

Figure 3.1: Artificial tactile skin with embedded liquid metal microchannel taxels. Schematic of the taxel embedded skin (a) showing 1st and 2nd layer of elastomer and the design of the normal and shear taxels. Scale bar: 2 mm. (b) CAD drawing of the skin with one normal taxel and three shear taxels on the distal face and one normal taxel on the proximal face. Scale bar: 1 cm. (c-e) Image of the taxel embedded skin prototype at different angles.

3.3.3 Dry, static force characterization

The dry, static taxel response at atmospheric conditions is characterized on an optical breadboard stage (PBH11105, Thorlabs Inc., NJ, USA), where the normal and shear forces were applied through the linear motion of three single-axis translation stages (433 translation stage, Newport, CA, USA) as shown in Figure 2. The sensorized end effector and a load cell 6-axis force/torque (F/T) transducer (Nano-25, ATI, Apex, NC) were attached to linear stages so that linear motions of the stages could impart shear force to the skin. The forces were applied via a 3D-printed flat-profile probe that made contact with the entire skin surface. A current source (Keithley Model 6430, OH, USA) supply a 10 mA DC current to each taxels and a data acquisition card (NI USB 6210, National Instruments Corporation, TX, USA) was used for data collection at a frequency of 10 Hz.

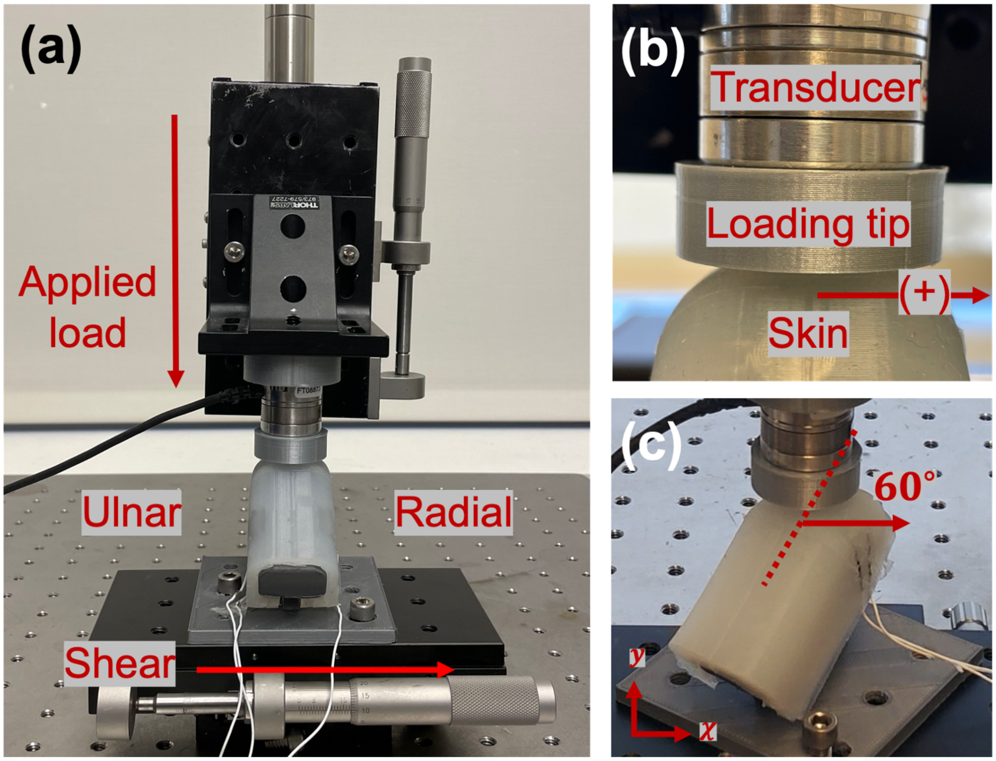

Figure 3.2: (a) Optical table setup for characterizing the shear taxel embedded skin. The 3D printed force loading tip (grey) is connected to the F/T transducer on the top. (b) An intermediate state of the F/T transducer and skin contact during the application of a shear force, indicated by the red arrow. (c) Image of the F/T transducer and skin contact for the two directional shear force application.

3.3.4 Wet testing under hydrostatic pressure

We demonstrate underwater tactile sensing using a modified benchtop watch tester, as shown in Figure 3. We simulate performance of the sensor skin at depths up to 100 m by pressurizing the watch tester with a commercial air compressor to a pressure of 1000 kPa. The sensor skin was submerged below the waterline within the chamber of the watch tester. A single-axis linear actuator (23HS30-3004S, STEPPERONLINE, NY, USA) moved a shaft to transmit loads to the surface of the submerged sensor skin. A single-axis, vented load cell (LC201Z0010-75, Omega Engineering, CT, USA) was fixed to the end of the shaft in order to measure the applied load. The load was applied to the sensor skin through a 3D printed probe tip. The probe tips were designed specifically for each experiment and are detailed in the following section. The hydrostatic pressure within the chamber was collected via a digital pressure gauge (PX309-200G5V, Omega Engineering, CT, USA), and recorded simultaneously along with taxel resistance using a data acquisition system (NI USB-6225, National Instruments Corporation, TX, USA).

We evaluate the normal taxel response of the proximal face underwater at pressures up to 1000 kPa as shown in Figure 3a. We also measure the taxel response to loads that are a combination of normal and two-dimensions of shear force. The finger is oriented so that the load is applied 35° from the surface normal axis and 60° from the radioulnar axis, as shown in Figure 3b.

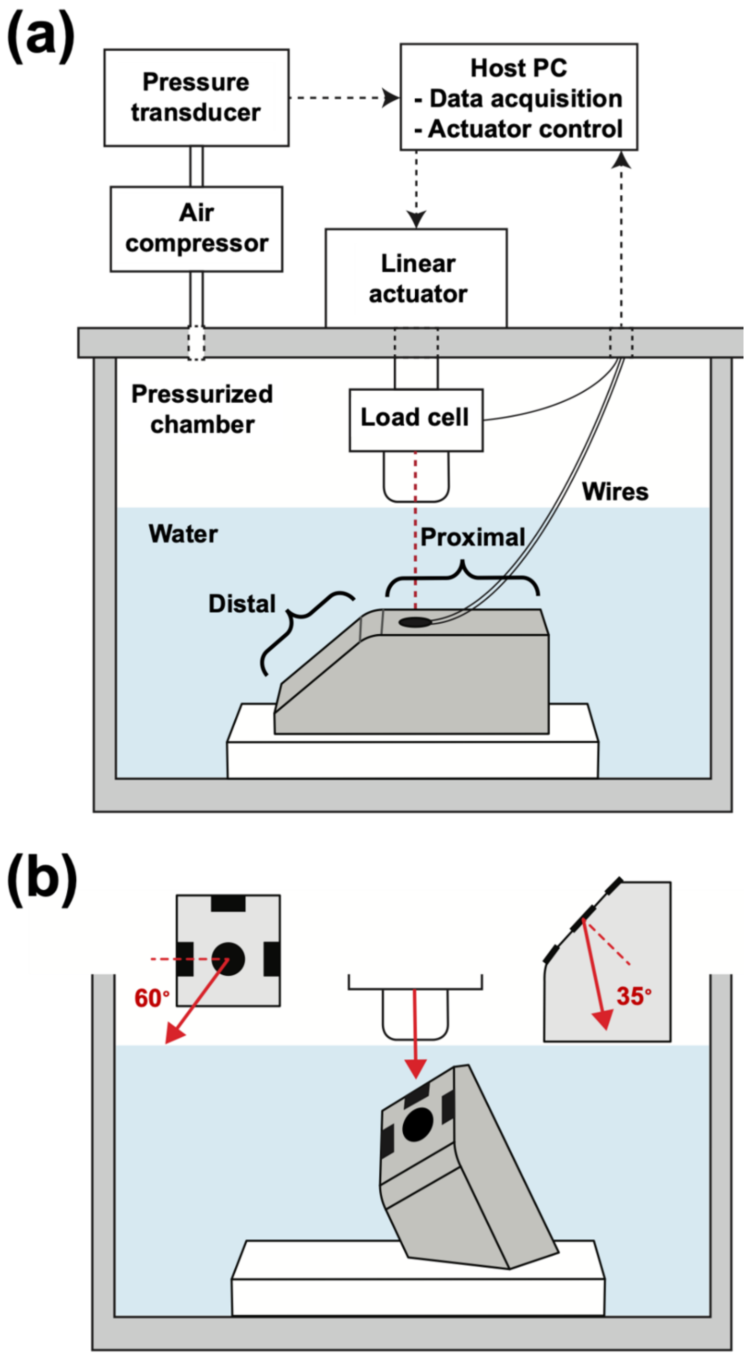

Figure 3.3: Experimental setup for characterization and evaluation of the sensor skin behavior underwater, under hydrostatic pressure. A linear actuation system, equipped with a vented load cell, applies a load to the sensor skin within the pressurized chamber. (a) The sensor skin is mounted to apply a normal force to the proximal face. (b) The sensor skin is mounted at an angle such that the actuation system produces a three-dimensional force relative to the distal face of the sensor skin.

3.4 Results

We present the results of three characterization experiments that evaluate the sensor's behavior in an both and dry and underwater, under hydrostatic pressure environments. We establish the characteristics of the sensor skin in a dry, ambient environment by two dimensions of shear force. We perform underwater characterization of the sensor by evaluating the unloaded taxel response as a function of hydrostatic pressure. Finally, we demonstrate three-dimensional load response of the sensor skin in and underwater, pressurized environment.

3.4.1 Shear taxel characterization in a dry, ambient environment

We evaluate the taxel response to static forces in a dry lab at atmospheric pressure. The shear force was generated by stepwise horizontal motions in increments of 0.1 N. We evaluate not only the sensor responses to the linear shear force (one-dimensional) (Figure 2b), but also to the two-dimensional shear force (Figure 2c). The normalized taxel response at ambient pressure $\widetilde{\Delta R}$ was defined as $\widetilde{\Delta R} = \frac{R-R_0}{R_0}$, where R is the measured resistance at a given load and R_0 is the resistance at zero load.

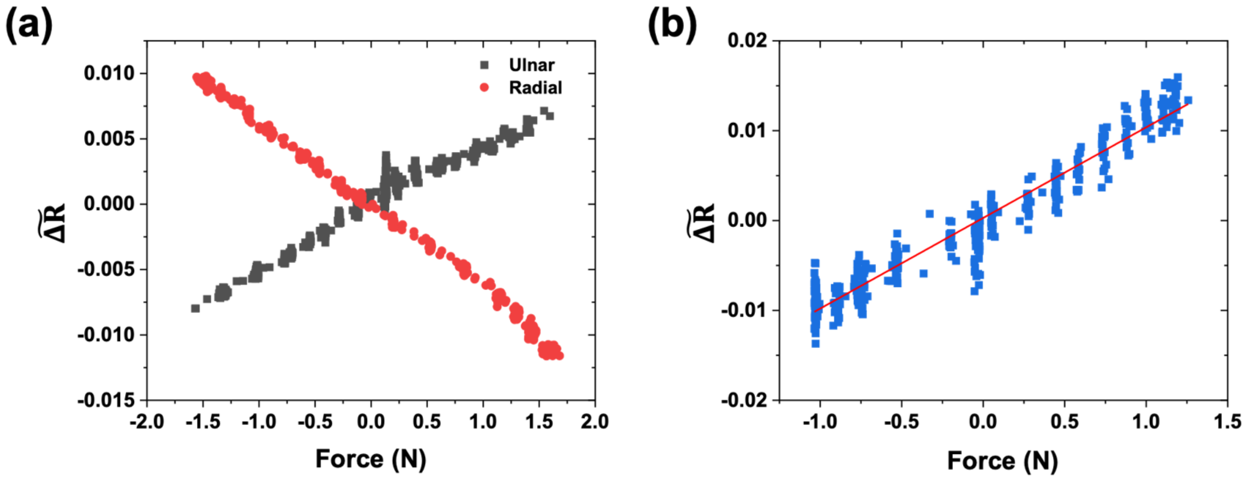

Figure 3.4: Normalized shear taxel response as a function of (a) one dimensional shear force and (b) two dimensional with 5 N normal force applied.

Figure 4 shows normalized resistance due to applied shear forces for an applied 5 N normal force. As shown in Figure 4a, the shear taxel response is linear with applied shear force over the examined range. In this experiment, the end effector is translated in the positive x-direction (radial direction) which imparts a shear force on the skin in the negative x-direction (ulnar direction). The left taxel is compressed and the right taxel is stretched, which results in a decrease and increase in normalized resistance, respectively. Additionally, we measured shear taxel response to two-dimensional shear force (f_x, f_y) as shown in Figure 4b. This characterization shows that the artificial tactile skin with embedded liquid metal taxels exhibit shear forces up to 1.2 N in two dimensions.

3.4.2 Taxel characterization in an underwater pressurized environment

The effect of hydrostatic pressure on taxel behavior is shown in Figure 5. First, we evaluate the sensor in a submerged, pressurized environment without applied load. We submerse the skin in water and cycle the hydrostatic pressure from 103 kPa to 830 kPa in steps between 35 and 105 kPa. Equilibration of the taxel resistance measurement after a stepwise change in pressure varied in duration between 10 sec and 40 min, with higher pressures exhibiting longer equilibration times. The equilibrated taxel responses for three pressure cycles are plotted as a function of hydrostatic pressure in Figure 5a. This data demonstrates that the taxel baseline resistance does not substantially depend on the hydrostatic pressure. This is an expected result because the sensor is constructed of nearly incompressible materials (silicone elastomer, EGaIn).

Figure 5b shows a normal taxel response to applied load while submerged underwater and under hydrostatic pressure. We apply loads up to 220 N orthogonally to the submerged taxel while the watch tester chamber was held at pressures of 0, 345, 690, and 1000 kPa. The normalized taxel response in the underwater shows that the measured normal taxel response, while submerged underwater, is independent of hydrostatic pressure up to 1000 kPa.

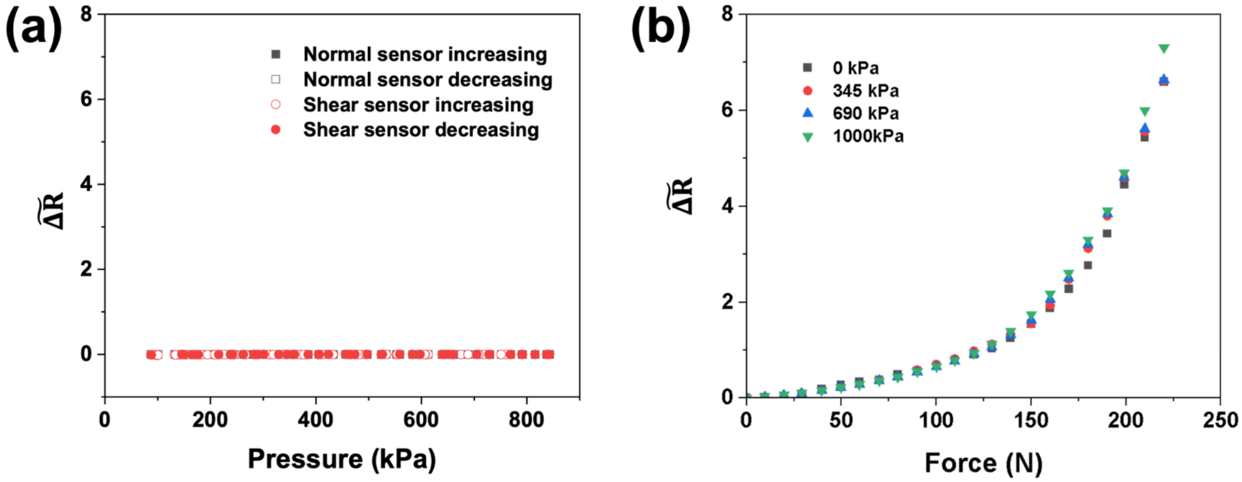

Figure 3.5: (a) Zero-load hydrostatic response of representative taxels of each modality. Data from three pressure cycles are presented. (b) Normal taxel response curves with respect to the loads between 0 and 220N. Each curve was collected while the taxel was subjected to a constant hydrostatic pressure of 0, 345, 690, and 1000 kPa.

3.4.3 Multimodal transient response to applied 3D load at pressure

To demonstrate multimodal sensing underwater at hydrostatic pressure, the sensor skin was submerged 10 mm below the waterline at 690 kPa and oriented such that the linear actuator simultaneously stimulated all four taxels on the finger's distal face. A schematic of the taxel layout relative to the direction of the applied load is shown in Figure 2b. The loading profile consisted of a 5 sec ramp followed by a 7 sec constant hold. All four taxels were measured simultaneously while 3D load was applied, and their normalized taxel responses as a function of time are shown in Figure 6.

The transient 3D load produces tension in the distal and ulnar taxels, and compression in the radial taxel of the sensor skin. The compressed and stretched shear taxels decrease and increase respectively in electrical resistance as shown in Figure 6. The design of the shear taxel enables the measurement of asymmetric strains across the pair of gauges and the shear force applied on the sensor skin. This taxel responses demonstrates the capability of the skin for multimodal, dynamic three-dimensional force sensing underwater and under pressure.

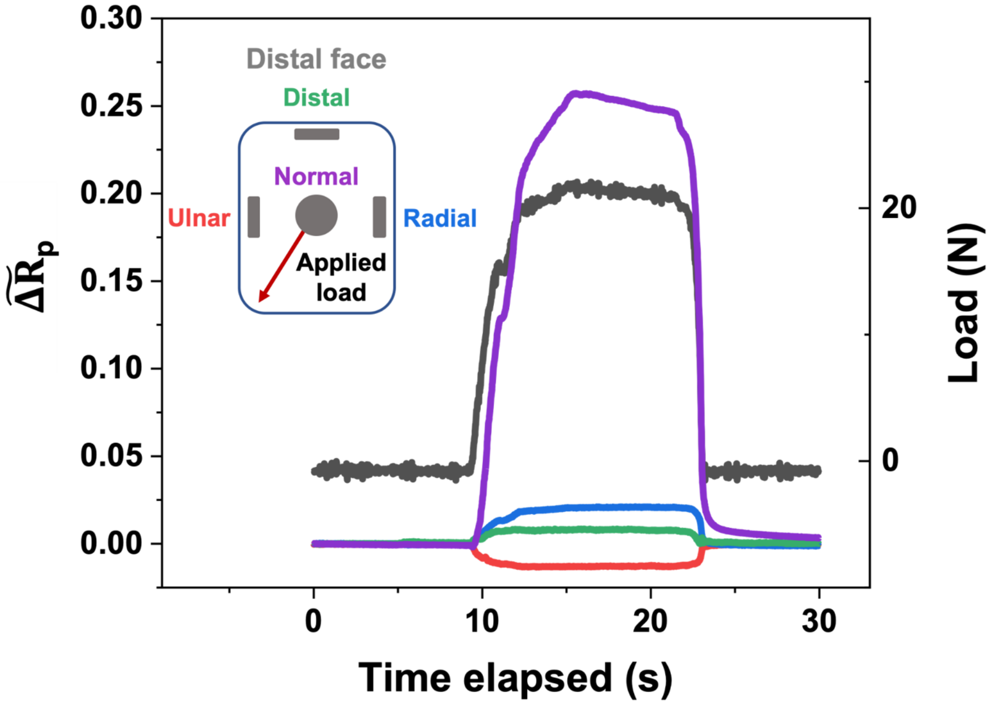

Figure 3.6: Transient response of taxels embedded in the distal face of the skin during the application of a 3-dimensional load at 345 kPa. Inset (legend) shows the schematic of taxel layout relative to the direction of the applied load and the load direction. The true load direction is rotated 35° relative to the surface of the distal face (a more detailed experimental setup is shown in Figure 3b).

3.5 Summary

This paper demonstrates a taxel-embedded artificial skin capable of sensing normal force and two dimensions of shear force under water and a wide range of hydrostatic pressures. The embedded taxels response remained constant under hydrostatic pressures of up to 1000 kPa. The zero-load taxel behavior was shown to be insensitive to hydrostatic pressure, which enables this sensor skin to be used without pressure-dependent calibration. Additionally, the sensor skin is capable of detecting large magnitudes of normal force and two-dimensional shear force, enabling 3-dimensional force sensing in underwater robotic operations. This result provides a guideline for fabricating soft tactile sensors for various potential robotic applications in extreme environments, such as deep in the ocean.

3.6 Acknowledgement

The authors are grateful for support from RE2 Robotics. This work was supported in part by Office of Naval Research Awards N00014-19-C-2008 and N00014-18-S-B003.

CHAPTER 4

Torque Estimation with a Skin-Like Tactile Sensor

4.1 Abstract

Objective: A set of considerations for the design of strain-gauge based multimodal tactile sensor skins is presented, with special focus on the estimation of torque normal to the sensor skin surface. *Methods*: A sensor skin is developed with the purpose of validating these design considerations, and is experimentally investigated by applying 3D linear forces and 1D torques using a custom motion platform. We present single-taxel characterization plots, and also perform multi-axis estimation using a CNN-based model in order to validate the design considerations and demonstrate force-torque sensing. *Results*: The design considerations and analytical sensor skin model provide verifiable assistance in the design and layout of sparsely populated sensor skins. The sensitivity of individual taxels to applied forces and torques is demonstrated to be nonlinear and not entirely predictable. *Conclusion*: While normal force taxels and shear force taxels are intended to measure normal and shear force, respectively, the sensor skin moves monolithically as a result of applied loads and torques. The results demonstrate significant correlation between taxels. Regardless of this correlation, the design consideration that was postulated to enable torque sensing was verified through both single-axis torque response plots as well as combined force-torque estimation results. *Significance*: The proposed design criteria and subsequent evaluation yield the first demonstration of

torque measurement using a tactile sensor skin. With sparsely populated tactile sensor skins, judicious placement of taxels is critical in order to extract as much information as possible from the touch interaction.

4.2 Introduction

With the proliferation of advanced sensing techniques in robotic systems, tactile sensing has emerged as a method for filling in the gaps of other senses, such as vision when a direct line of sight is unavailable. Tactile sensing provides important details during manipulation tasks, such as information about grasp stability and object properties [52]. Tactile sensing is also used as a means to monitor and enable safe human-robot interaction in collaborative human-robot tasks [98]. The past few decades of tactile sensor literature have gradually transitioned from a focus on the development of hardware to the development of perception algorithms and the integration of tactile sensors into comprehensive manipulation pipelines [99], with a focus on motion planning, feedback control, and high-level learning. Integral to the aim of semi-autonomous manipulation is the real-time construction of a world model; the robot must be able to adapt in unstructured environments by integrating multimodal sensory data.

While implementing tactile sensing systems into semi-autonomous robotic systems remains a tricky task, there are a number of hardware designs which have successfully demonstrated advanced capabilities such as force-torque sensing, vibration, temperature, and local geometry estimation through tactile sensing alone [18, 100, 63, 101, 102]. In order to meet the demands of high accuracy and spatial resolution during robotic manipulation, much focus has turned to developing sensorized fingertips. In these devices, the tactile sensing elements as well as the accompanying electronics and mechanical hardware are fit into a finger-like form factor.

The objective of this work is to provide a verifiable set of considerations for the design of tactile sensor skins capable of measuring applied forces and torques. We validate these design considerations using a tactile sensor skin developed in Chapter 3 and demonstrate the measurement of three axes of force and one axis of torque.

The contributions of our work are:

1. A set of design considerations for arranging planar strain gauges for improved multi-modal tactile sensing
2. Rationale for the novel placement of planar strain gauges that enables measurement of torque about the normal to the sensor skin
3. The first demonstration of force-torque estimation using a wrappable skin-like sensor
4. A framework for the design, characterization, and evaluation of skin-like tactile sensors

4.3 Related Work

4.3.1 Tactile Sensor Skins

The sensorized fingertips mentioned in Section 4.2 consist of complex sensor hardware that is integrated within the volume of the finger, and cannot be wrapped around existing robotic fingertips. In contrast to sensorized fingertips, tactile sensor skins allow the tactile sensor design and mechanical gripper design to be approached independently [103]. Tactile sensor skins can be fabricated to conform to and wrap around the surface that needs to be outfitted with tactile sensors. “Electronic skin” devices are being developed for a multitude of purposes, for example as a user input device for consumer electronics [104, 105]. Park et al. produced individual tactile sensing elements (“taxels”) made of microfluidic strain gauge sensors that demonstrated multiple modalities of force sensing [106].

Aside from the physical differences between wrappable and non-wrappable tactile sensors, there are often differences in spatiotemporal resolution. For example, resistive, microfluidic tactile sensor skins can have bandwidths up to 800 Hz [74] while the 30 fps sampling rate for camera-based tactile sensors limits their ability to measure vibration. However, camera-based tactile sensors have an advantageous spatial resolution limited only by the pixel resolution of the camera. In contrast, microfluidic tactile sensors produced using soft lithography techniques typically incorporate fewer taxels and so their data are comparatively sparse. Therefore, care must be taken to place individual tactile sensing elements (taxels) throughout the skin in a way that ensures sensitivity to applied loads of interest, which can include 3D force as well as 3D torque.

4.3.2 Measurement of Torque at Robot Fingertips

Measurement of only linear forces at the fingertip often suffices in manipulation tasks such as pick-and-place and parallel static grasps, but complex and dynamic applications are enhanced by also measuring torque. As an instructive example, consider the task of picking up the end of a block using a precision pinch. Since the block’s center of mass is offset from the grip axis, the block will rotate about the grip axis due to gravity. The rotational slip will occur within the plane of contact about an axis perpendicular to the plane of contact.

Preventing unwanted rotational slip, or controlling rotational slip during in-hand manipulations, is an active research area in tactile sensing for robotics [107, 108, 109].

Existing tactile sensing systems have been either modified or built from the ground up to exhibit not only in-plane torque sensing but also tilt torque. The GelSight sensor and similar optical sensor approaches are capable of estimation of in-plane torque through tracking the optical flow of markers embedded in an elastomeric fingertip material [19]. In a demonstration of dynamic tactile sensing and control, Wang et al. used finger-mounted GelSight sensors to learn the physical features of a grasped object through a set of exploratory motions [109]. The F-TOUCH sensor enhanced the capabilities of the GelSight by combining spring-like elements into the construction of a camera-based tactile sensor. This provides not only optical flow tracking, but also full force-torque tracking of the entire fingertip [101]. Additionally, the OmniTact sensor exploits a high-curvature fingertip surface and multiple cameras to enable multidimensional torque estimation using only optical flow [102].

Notably, none of the state-of-the-art tactile sensing systems that enable in-plane torque sensing can be wrapped around existing robot fingertips. In this work, **we investigate design parameters, such as local taxel design and global taxel placement, that can be exploited to provide in-plane torque sensing capabilities to tactile sensor skins**.

4.4 Framework for Tactile Sensor Skin Design with Planar, Microfluidic Taxels

There are a number of design parameters that can be considered when developing multi-modal tactile sensor skins. In this work, we consider microfluidic tactile sensor skins whose manufacturing process is described in Chapter 3. As such, we assume that the designer can select the (i) local pattern of the microfluidic channel for each taxel as well as the (ii) global placement of taxels throughout the skin as a whole.

As will be discussed in Section 4.4.2, the distinction between normal force and shear force taxels is somewhat artificial; the designation is primarily assigned according to the local pattern of the microfluidic channel and the type of load to which the taxel is most sensitive. Normal force or contact pressure is commonly sensed using a spiral microchannel (Fig. 4.1a) [79]. Similar to their rigid strain gage counterparts, microfluidic shear force taxels utilize serpentine microchannels (Fig. 4.1b).

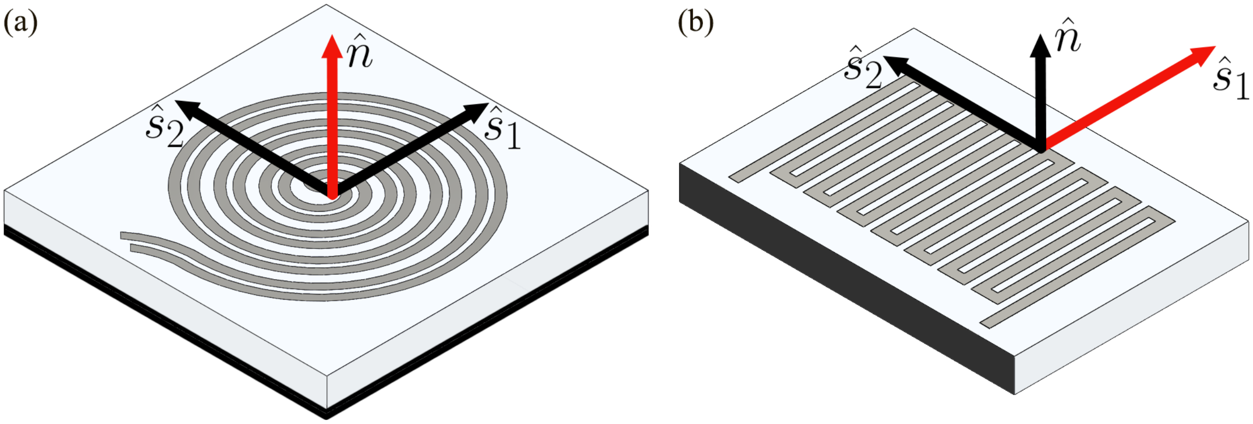

Figure 4.1: Local patterns of microchannels are shown for a single (a) normal force taxel and (b) shear force taxel. Taxels are designed to be most sensitive along their principal axis, denoted by dark red arrows. The black area represents the plane along which the local shear of the elastomer is assumed to be zero.

4.4.1 Individual Taxel: Principal Axis of Sensitivity

At its core, a microfluidic taxel is a highly deformable strain gauge. An elastic medium such as polydimethylsiloxane (PDMS) surrounds a curvilinear channel of embedded liquid metal. The taxel's sensitivity to applied loads is determined by the channel's cross-sectional geometry as well as the channel pattern. As the geometry of the microfluidic channel is mechanically deformed by a load, the voltage drop across the ends of the channel changes.

In this work, all microfluidic channels have a fixed cross-sectional width of $150\mu m$ and height of $60\mu m$, as shown in 4.1 and described in detail in 3.

4.4.1.1 Normal Force Taxels

With its spiral microchannel pattern, the normal force taxel is predominantly sensitive to contact forces along the $\hat{n}$-axis that is orthogonal to the $\hat{s_1} - \hat{s_2}$ plane of the taxel, as shown in Fig. 4.1a. The labels $\hat{n}, \hat{s_1}$, and $\hat{s_2}$ indicate a single normal and two shear basis vectors, respectively. The origin of the local $\hat{s_1} - \hat{s_2} - \hat{n}$ reference frame is located at the center of the spiral pattern.

For the normal force taxel, the $\hat{n}$-axis is the *principal axis of sensitivity*. Compressive contact forces will point in the $(-)\hat{n}$ direction and will result in positive changes in voltage. The microchannels of the normal taxel are designed to have negligible voltage change as a result of shear within the $\hat{s_1} - \hat{s_2}$ plane. However, in practice, planar shear can create a measurable taxel response.

While the contribution of the normal force taxels to multimodal sensing will be discussed in Section 4.6.2, we focus most of our analytical efforts on the contribution of the shear force taxels due to their interesting response characteristics.

4.4.1.2 Shear Force Taxels

With its serpentine microchannel pattern, the shear force taxel is predominantly sensitive to planar shear deformations that are aligned with the long segments of the channel (Fig. 4.1b). For the shear force taxel, the $\hat{s_1}$-axis is the principal axis of sensitivity. The origin of the local $\hat{s_1} - \hat{s_2} - \hat{n}$ reference frame is centered along the edge of the taxel opposite the taxel leads.

An unexpected finding from the work in Chapter 3 was that a shear force taxel of the form shown in Fig. 4.1b could respond to compressive forces along its principal axis if constrained appropriately. In prior work with comparatively thin tactile sensor skins (1 mm thickness), the shear force taxels stretched in response to tensile forces along the principal axis. However, the skin itself would buckle if subjected to compressive forces. As a result, shear force taxels in regions that buckled did not experience deformation of their serpentine microchannels and no voltage changes were observed across the taxel leads [55, 74]. The results presented in 4.6.2 demonstrate that our relatively thick shear taxels are capable of bidirectional shear force sensing.

In the process of enhancing the robustness of the multimodal skin for use underwater and under hydrostatic pressure, as described in Chapter 3, thin taxels were embedded in a 4 mm thick skin. Now embedded and constrained by a thicker skin substrate, shear force taxels subjected to compressive forces along their principal axes experience deformation of their serpentine microchannels, resulting in measurable voltage changes across the taxel leads.

For the framework that follows, we assume that the edge of the shear force taxel having the taxel leads (ends of the microchannel) is fixed in place, and that the microchannel deforms relative to that fixed edge. As such, taxel deformation will be greatest for those regions that are closest to the applied load and will decrease gradually until convergence to an undeformed state along the fixed edge of the taxel (black region noted in Fig. 4.1b). Similarly, we assume that the bottom surface of the normal taxel is supported by a fixed base such that the normal taxel geometry is deformed and not translated as a result of applied force.

Considering shear force taxels constrained within relatively thick skin substrates, the change in voltage due to deformation along the taxel's $\hat{s_1}$ principal axis will be a signed value. Specifically, positive voltages will be measured for tensile strain in the $(+)\hat{s_1}$ direction while negative voltages will be measured for compressive strain in the $(-)\hat{s_1}$ direction (Fig. 4.2(a,b)). A single shear force taxel can provide information about both magnitude and direction of a shear force component along the primary shear $\hat{s_1}$-axis.

Critically, this signed response is only characteristic of the shear force taxel's $\hat{s_1}$ principal axis, or *primary shear axis*. Shear force components along the $\hat{s_2}$-axis, or *secondary shear axis*, will skew the shape of the taxel (Fig. 4.2(c)). However, strain along the $\hat{s_2}$-axis will always yield a positive change in voltage regardless of whether the shear force component points in the $(+)$ or $(-)$ direction along the $\hat{s_2}$-axis (Fig. 4.2(c) and Fig. 4.2(d), respectively). A single shear force taxel can provide information about the magnitude but not the direction of a shear force component along the secondary shear $\hat{s_2}$-axis.

4.4.1.3 3D Deformation Vectors

Consider a tactile sensor skin comprised of n individual planar, microfluidic taxels. For each of $i = 1 \dots n$ taxels, the location of the i^{th} taxel can be expressed as a 3D position vector $\underline{r}^i$ that points from the origin of the skin's global $\hat{x} - \hat{y} - \hat{z}$ reference frame to the origin of the i^{th} taxel's local $\hat{s_1} - \hat{s_2} - \hat{n}$ reference frame (Fig. 4.3).

For simplicity, consider a 3D force-torque load $[\underline{f}, \underline{\tau}]^T$ applied at the origin of the skin's global reference frame. The effect of the load on the deformation of the i^{th} taxel will depend upon (i) the proximity of the taxel to the point of load application and (ii) the orientation of the taxel's microchannels relative to the direction of the load.

The force-based deformation components due to $\underline{f}$ can be expressed for the i^{th} taxel in a

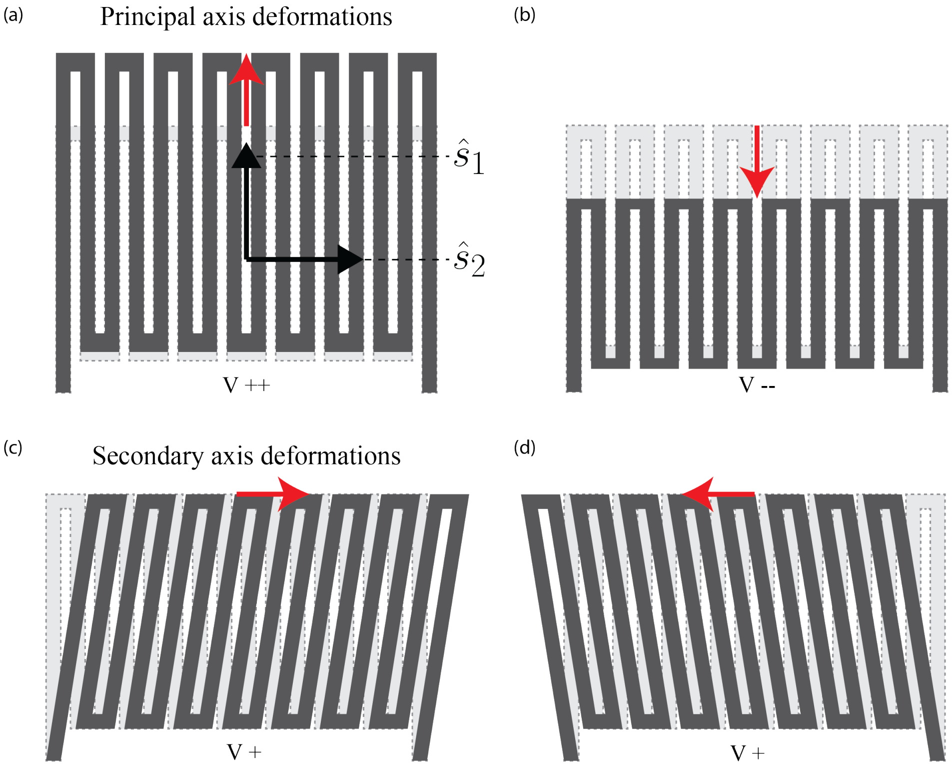

Figure 4.2: Shear force taxel behavior under positive and negative deformations along the principal and secondary axes of the taxel.

normalized manner by considering the projection of the 3D force unit vector $\underline{\hat{f}}$ onto the taxel's local reference frame axes. Although the equations below do not account for the nonlinear behavior of compliant materials and the exact proximity of individual taxels relative to the center of contact, the equations enable the analysis and design of taxel placement for a planar wrappable sensor skin.

$$\delta^i_{f,s1} = \underline{\hat{f}} \cdot \underline{\hat{s}}^i_1 \tag{4.1}$$

$$\delta^i_{f,s2} = \underline{\hat{f}} \cdot \underline{\hat{s}}^i_2 \tag{4.2}$$

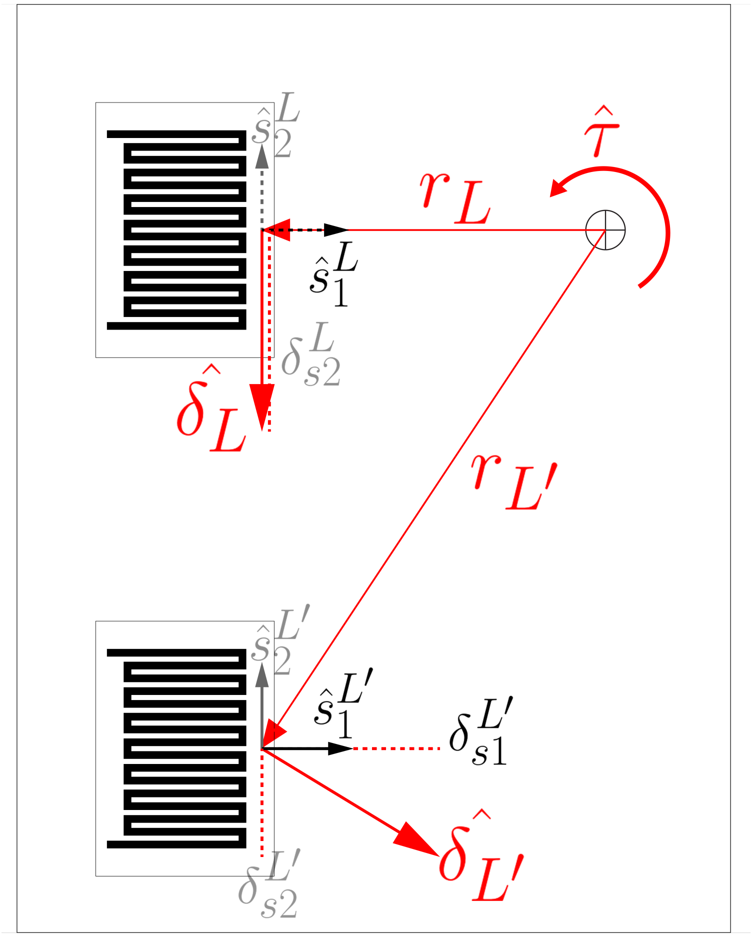

Figure 4.3: Local deformation of two shear taxels as a result of an applied torque, broken into its $\hat{s}_1, \hat{s}_2$ components. The local deformation at the origin of taxel L has no component parallel to the principal axis ($\hat{\delta}^L_{\tau,s1} = 0$).

$$\delta^i_{f,n} = \underline{\hat{f}} \cdot \underline{\hat{n}}^i \tag{4.3}$$

The force-based 3D deformation vector $\underline{\delta}^i_f$ for the i^{th} taxel can be expressed as

$$\underline{\delta}^i_f = [\delta^i_{f,s1}, \delta^i_{f,s2}, \delta^i_{f,n}]^T \tag{4.4}$$

The torque-based deformation components due to $\underline{\tau}$ can be expressed for the i^{th} taxel in a normalized manner by first considering the effects of the torque load on the direction of the torque-based deformation vector $\underline{\hat{\delta}}^i_\tau$

$$\underline{\hat{\delta}}^i_\tau = \underline{\hat{\tau}} \times \underline{\hat{r}}^i \tag{4.5}$$

and then projecting the torque-based deformation unit vector $\underline{\hat{\delta}}^i_\tau$ for the i^{th} taxel onto the taxel's local reference frame axes.

$$\delta^i_{\tau,s1} = \underline{\hat{\delta}}^i_\tau \cdot \underline{\hat{s}}^i_1 \tag{4.6}$$

$$\delta^i_{\tau,s2} = \underline{\hat{\delta}}^i_\tau \cdot \underline{\hat{s}}^i_2 \tag{4.7}$$

$$\delta^i_{\tau,n} = \underline{\hat{\delta}}^i_\tau \cdot \underline{\hat{n}}^i \tag{4.8}$$

The torque-based 3D deformation vector $\underline{\delta}^i_\tau$ for the i^{th} taxel can be expressed as

$$\underline{\delta}^i_\tau = [\delta^i_{\tau,s1}, \delta^i_{\tau,s2}, \delta^i_{\tau,n}]^T \tag{4.9}$$

For normal force taxels, we assume $\delta_n > (\delta_{s1} = \delta_{s2})$. For shear force taxels, we assume $\delta_{s1} > \delta_{s2}$ and $\delta_n = 0$.

4.4.2 Multiple Taxels: Placement Within a Sensor Skin

For most microfluidic tactile sensor skin designs, the soft lithography manufacturing method and relatively large size of the taxels limit the total number of taxels that can be placed on the surface of a wrappable sensor skin. As a result, we must carefully design the type, location, and orientation of each taxel within a sparsely populated sensor skin.

The design recommendations that follow assume that the center of contact on the sensorized fingerpad is generally known and repeatable. In many grasp and manipulation applications, the center of contact is often predetermined by mechanical design (e.g., a robotic gripper with constrained finger motion and convex finger surfaces) and by application (e.g., use of the center of a fingerpad for exploratory movements).

For the purposes of illustrating the design criteria, Fig. 4.4 shows the arrangement of taxels within the sensor skin used in the experimental portion of this paper.

4.4.2.1 Force Sensing

In order to sense a 3D force load, a combination of normal force taxels and shear force taxels can be used to sense deformation normal to the skin's surface and 2D, in-plane deformation, respectively.

Normal force taxels: In this work, the global placement of the normal force taxels within the tactile sensor skin was previously selected based on assumed areas of direct contact between the robot fingertip and grasped objects, as described in Chapter 3. It is recommended that normal force taxels be embedded in close proximity to the expected center of contact, which we show as the origin of the global $\hat{x} - \hat{y} - \hat{z}$ reference frame of the skin (Fig. 4.4) for convenience. Additionally, normal force taxels should be oriented such that their principal

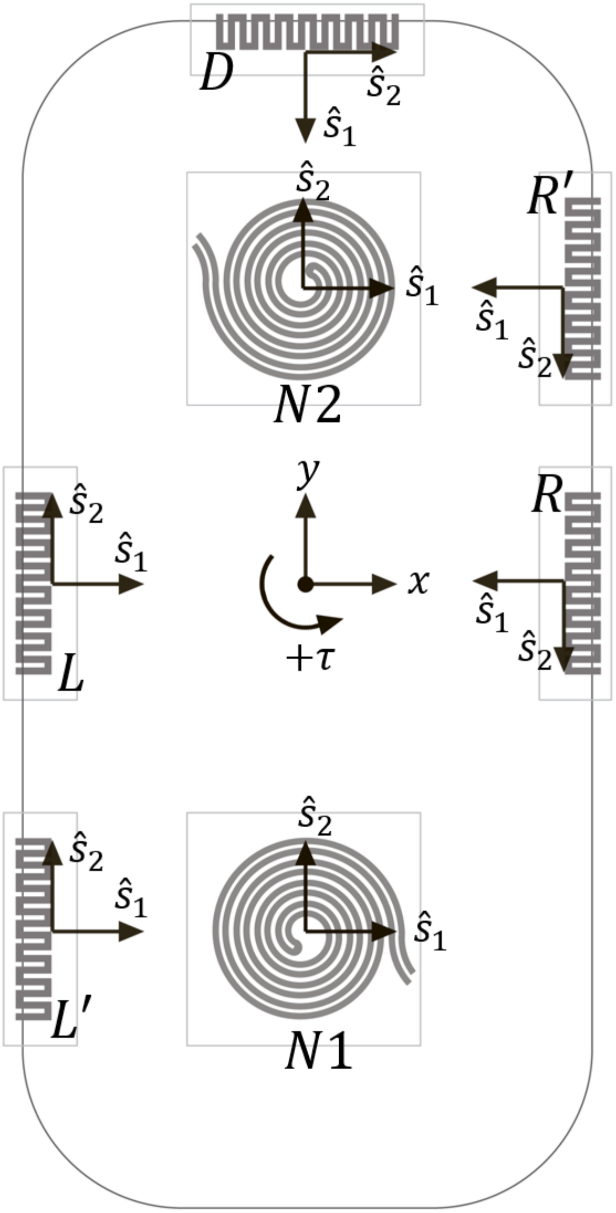

Figure 4.4: Configuration of taxels used in the sensor skin. Off-axis shear force taxels (R', L') are introduced as well as on-axis shear force taxels in order to validate the design consideration for torque sensing.

axis of sensitivity ($\hat{n}$-axis) is aligned with the skin's $\hat{z}$-axis.

Shear force taxels: While designed to sense in-plane deformation, shear force taxels also respond to deformations normal to the skin's surface. This behavior is a consequence of the sensor skin consisting of a single, continuous, compliant body. In order to minimize the covariance of normal and shear force taxel responses due to normal forces, it is recommended to place shear force taxels away from regions of direct contact such that deformation will occur predominantly along the primary ($\hat{s}_1$) and secondary ($\hat{s}_2$) axes. Furthermore, if there is a specific global axis of interest, such as force along the $\hat{x}$-axis, it is recommended to align the primary shear axis $\hat{s}_1$ of the shear force taxel with the global axis of interest. As an example, the 'L' and 'R' shear force taxels shown in Fig. 4.4 were placed such that their $\hat{s}_1$-axes were colinear with the skin's $\hat{x}$-axis.

In prior work with comparatively thin sensor skins [55, 74], pairs of shear force taxels were used to measure shear force along a single axis. The asymmetry in the paired responses was used to estimate shear force. In this work with a thicker 4 mm skin substrate, only a single shear force taxel is necessary to extract magnitude and direction of a shear force component. As an example, the 'D' shear force taxel shown in Fig. 4.4 was placed such that its $\hat{s}_1$-axis were colinear with the skin's $\hat{y}$-axis. Due to constraints such as limited skin area and the need to accommodate gripper kinematics, only a single shear force taxel could be placed along the skin's $\hat{y}$-axis.

For the task of sensing a 2D shear force vector, a minimum of two shear force taxels is needed and their relative orientations must be such that their principal axes of sensitivity are non-parallel. For example, the 'L' and 'D' shear force taxels shown in Fig. 4.4 were placed orthogonally in order to minimize covariance and redundancy of the taxel responses, and for convenient alignment with the skin's $\hat{x}$- and $\hat{y}$-axes, respectively.

4.4.2.2 Torque Sensing

As a first step towards the development of wrappable sensor skins that can sense torque, we focused our efforts on the placement of taxels for the sensing of torque about the skin's $\hat{z}$-axis only. The $\hat{z}$-axis was selected because of its relevance to rotational slip in the robotics literature [110, 109]. The sensing of 3D torque loads is beyond the scope of this work.

To effectively place shear force taxels for the measurement of torque about the $\hat{z}$-axis, consider Fig. 4.3. A shear force taxel, such as L, whose principal axis of sensitivity ($\hat{s}_1$) is aligned with the position vector $\underline{r}^L$, will experience maximal deformation along the $\hat{s}_2$-axis.

Now consider a circle of radius $\underline{r}^L$ emanating from the center of application of a torque load about the $\hat{z}$-axis. Any shear force taxel whose principal axis of sensitivity $\hat{s}_1$-axis intersects with the origin of the global reference frame will not be sensitive to a torque load about the $\hat{z}$-axis. Stated another way, any shear force taxel whose $\hat{s}_2$-axis lies tangent to a circle centered at the origin of the global reference frame will not be sensitive to a torque load about the $\hat{z}$-axis.

A shear force taxel, such as L', is oriented such that it will be sensitive to deformations due to a force load along the $\hat{x}$-axis as well as a torque load about the $\hat{z}$-axis. This is because the principal axis of sensitivity ($\hat{s}_1$) is aligned with neither the radius nor the tangent of the circle centered at the origin of the global reference frame, where the load is assumed to be applied. As a result, the deformation components $\delta^i_{f,s1}$ from eq. 4.4.1.3 and $\delta^i_{\tau,s1}$ from eq. will be nonzero. It is based on this principle, along with the shape of the robot fingertip, that taxels L' and R' were placed for this study (Fig. 4.4).

4.5 Experimental Procedure and Evaluation

In this section, we describe the specific layout of multiple taxels used in this study, the experimental procedure for collecting data from the multi-taxel tactile sensor skin, and the methods for training and testing a CNN model to estimate 3D force and 1D torque loads using the skin.

4.5.1 Multi-taxel Layout of the Tactile Sensor Skin

The layout of normal force and shear force taxels as described in Chapter 3 is extended through the addition of L' and R' shear force taxels on the left and right edges, respectively, of the proximal face of the tactile sensor skin (Fig. 4.4). The multi-taxel layout shown in Fig. 4.4 is a layout that minimizes the number of taxels while leveraging the design criteria described in Section 4.4.2 to achieve 3D force and 1D torque sensing. The multi-taxel layout also considers practical constraints such as limited skin area and the need to accommodate gripper kinematics.

It was hypothesized that normal force taxels N_1 and N_2 will be most sensitive to the $\hat{z}$-component of a 3D force load applied at the origin of the $\hat{x} - \hat{y} - \hat{z}$ global reference frame. The shear force taxels L and R will be most sensitive to the $\hat{x}$-component of the 3D force load. The shear force taxel D, on the distal end of the proximal face of the tactile sensor skin, will be sufficient to measure the $\hat{y}$-component of the 3D force load. Finally, the shear force taxels L' and R' will enable a new sensing capability: the sensing of a torque load about the $\hat{z}$-axis. For the purpose of subsequent discussion, the $L - R$ and $L' - R'$ shear force taxel pairs will be referred to as the "on-axis" and "off-axis" shear force taxel pairs, respectively.

4.5.2 Experimental Setup

We characterized the tactile sensor skin's ability to measure 3D force and 1D torque loads using the experimental testbed shown in Fig. 4.5. The sensorized finger was mounted to the print head of a modified benchtop 3D printer (Prusa, Mk3s). Using an open-source interface (Pronterface), we sent G-code commands to the printer in order to apply loads to the proximal face of the sensorized fingertip.

Three-dimensional force loads were created with the print head by pushing the fingertip against a rigid plate atop a 6DOF force-torque sensor (ATI, Nano17). Torque loads were created by pushing the fingertip against the rigid plate and then simultaneously commanding a rotary stage (PDV, PT-GD201) beneath the 6DOF force-torque sensor. Combining the 3D print head movements with the 1D rotary stage movement (driven by the printer extruder's stepper motor), the testbed enabled 4DOF control of load stimuli.

We acquired voltage signals from the sensor skin taxels and the 6DOF force-torque sensor using two data acquisition (DAQ) boards: National Instruments, NI-DAQ 6363 and NI-DAQ 6225, respectively. Raw taxel resistances were converted to DAQ input voltages using a custom printed circuit board (PCB) having Wheatstone bridge-based signal conditioning circuits. As the nominal resistance of each taxel is on the order of single ohms, each Wheatstone bridge was excited at 0.25 V, and voltage output of the Wheatstone bridge (on the order of microvolts) was sent to an instrumentation amplifier (Analog Devices, AD620). This resulted in analog inputs to the sensor skin DAQ in the [-1.5, 1.5] V range. The sampling rate of the NI-6225 and NI-6363 were synchronized at 500Hz for all data collections.

For each trial, we collect data from the experimental platform while creating loads on the sensor skin. For each load, the platform is actuated and then held constant for two

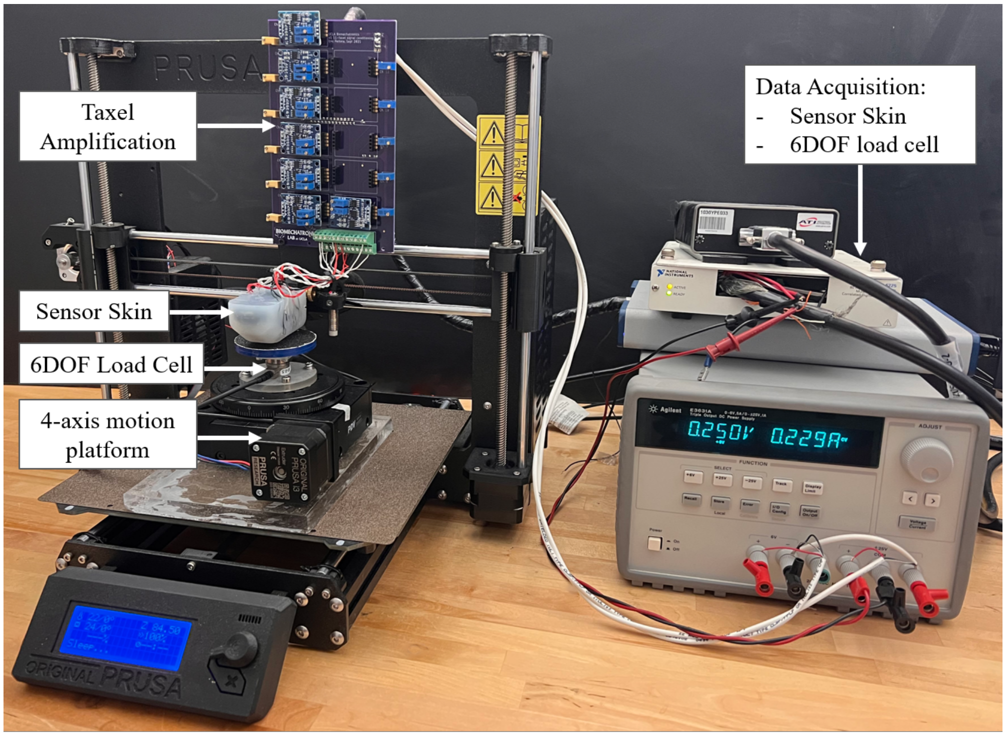

Figure 4.5: Experimental testbed used during characterization of the sensor skin

seconds. Each load contains the following sequence: (1) baseline readings with no applied load, (2) ramp-up of taxel readings while the platform is actuating, (3) equilibrated taxel readings while the platform is held constant, and (4) ramp-down of taxel readings when load is removed.

4.5.3 Taxel Characterization

4.5.3.1 Data Collection

We assess the performance of the sensor skin with respect to two-dimensional force loads. For this assessment, we used the experimental testbed shown in Fig. 4.5 to apply 200 randomly generated linear loads along either the x- or y-axis, as defined in Fig. 4.4. Each trial begins with an initial linear movement along the z-axis in order to generate a static normal force load of 10 N ($\mu = 10.02N, \sigma = 0.04N$), which was held for 0.5 seconds. The subsequent linear motion command (along either the x- or y-axis of the sensor reference frame) was held for 1 second.

We assess the performance of the sensor skin as a result of one-dimensional torque loads. For this assessment, we used the experimental testbed to apply 200 randomly generated torques between -30 N-mm to 30 N-mm to the surface of the sensor skin, located at the skin's global origin, as defined in Fig. 4.4. In order to prevent slip during the experiment, each rotation command was preceded by a linear movement along the z-axis, in order to generate a static normal force load of 10 N ($\mu = 10.02N, \sigma = 0.04N$), which was held for 0.5 seconds. The subsequent rotary command (about the positive z-axis in the counterclockwise direction) was held for 1 second.

4.5.3.2 Data Processing

Each data point in Fig. 4.7 represents the relative change in taxel voltage as a result of the applied torque:

$$\widetilde{\Delta V_k} = \frac{V_k - V_{k,0}}{|V_{k,0}|} \tag{4.10}$$

where V_k is the mean taxel voltage during the k^{th} applied static torque, and $V_{0,k}$ denotes the mean taxel voltage during the k^{th} applied static normal force. Normalizing the torque response data against the static normal load data allows us to visualize taxel response to torque by decoupling the two simultaneous stimuli.

4.5.4 Development of a CNN Model for 3D Force and 1D Torque

Tactile images are used as input to a three-layer CNN model, diagrammed in Fig. ZZ. The architecture of this CNN has been demonstrated to work well for tactile estimation applications, as in [111] and [112] where this architecture was used to estimate 2D shear force and contact force using the BioTac tactile sensor. The shear estimation of [112] assumed that the input was interpolated, whereas the normal force estimation method of [111] was left sparse. We continue in the vein of [ref] without interpolating between inputs and rely on the convolutional layers of the CNN to learn the taxel-to-taxel interactions during training.

For each configuration, we trained the CNN model described in Section 4.5.4 on the full force-torque dataset to create estimates for $[f, \tau]$ in the sensor's global frame. Only the taxels included in the configuration were used to create the tactile images input into the model, with the unused taxel values set to zero.

4.5.4.1 Data Collection

A total of 1000 random deformations were applied to the sensor skin, during which the data acquisition system was sampling Nano17 force-torque and sensor skin readings at 500 Hz. We used a 50 Hz moving average filter to remove the effects of signal noise. The linear forces applied during the experiments ranged in magnitude between $[-5 \ldots 5]$ N in the x-y plane, and $[0 \ldots 10]$ N along the z axis. The torques applied during the experiments ranged in magnitude between $[-30 \ldots 30]$ mN-m about the z axis. After data was collected, the raw taxel values were converted to *tactile images* through a preprocessing step as described in the following section.

4.5.4.2 Data Processing

For each trial and i^{th} taxel, the mean of the baseline (undeformed) sensor data $V^i_{k,0}$ is subtracted from the sensor data during loading $V^i_{k,n}$. After the removal of the baseline measurements, the sensor data $\tilde{V}^i_{k,n}$ are used as inputs for training a 3D force/1D torque estimation model:

$$\tilde{V}^i_{k,n} = V^i_{k,n} - V^i_{k,0} \tag{4.11}$$

where $n \in (1, 40)$ represents 40 samples taken from the equilibrated readings from the i^{th} taxel during phase (3) of the k^{th} trial. All samples from all deformations are consolidated into a single array representing the processed data. Each row of the array consists of all $i = 1 \ldots N$ taxel values for a given sample.

$$\tilde{V}_0 = [\tilde{V}^1_0, \tilde{V}^2_0, ..., \tilde{V}^N_0] \tag{4.12}$$

And the full dataset matrix is defined as

$$\tilde{V} = [\tilde{V}_0, \tilde{V}_1, ..., \tilde{V}_{n\times k}]^T \tag{4.13}$$

4.5.4.3 Tactile Images

The final step to prepare the taxel readings for use with the CNN-based force-torque estimation model is to shape the data into a 2D array which represents the geometric layout of the sensor skin. We chose to represent the data as a sparse array T where non-zero elements are defined as such:

$$T = \begin{bmatrix} 0 & 0 & \tilde{V}_D & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \tilde{V}_{N2} & 0 & \tilde{V}_{R'} \\ 0 & 0 & 0 & 0 & 0 \\ \tilde{V}_L & 0 & 0 & 0 & \tilde{V}_R \\ 0 & 0 & 0 & 0 & 0 \\ \tilde{V}_{L'} & 0 & \tilde{V}_{N1} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \tilde{V}_D & 0 & 0 \end{bmatrix} \tag{4.14}$$

4.5.4.4 Model Training

During training, Nano17 force-torque data is used as ground truth to produce labeled outputs for each tactile image in the training dataset. The entire dataset [X,y] where X is an Nx7x5 array of taxel reading inputs, y is a Nx4 array of force-torque outputs, and N is the length of the dataset (in this case, N=40,000), is split such that the model is trained on 80% of the dataset and tested on 20% of the dataset.

During testing, we performed a study comparing different sensor configurations in order to assess the performance of individual taxels as well as multi-taxel sensor configurations. For each sensor configuration, the data from specified taxels were removed from the dataset (and taxel readings set to zero). Comparison between taxel configurations was achieved by training separate CNN models on each configuration, using the same training dataset.

The sensor skin used in this study is comprised of seven taxels, five shear and two normal, arranged in the pattern of Fig. 4.4. The taxels labeled $L2$ and $R2$ are the taxels which are hypothesized to enable torque measurement via the CNN model. As per the mathematical walkthrough of Section XX, these two taxels are displaced from the intended contact location such that their principal axes are not coincident with the center of contact. This “lever arm” of the taxel allows applied torque to be seen locally as a principal deformation. The taxels labeled $L1$ and $R1$ are placed such that their principal axes are collinear with the center of contact, which should result in poorer torque estimation.

4.6 Results

In this section we present data that corroborates the design considerations and demonstrate force-torque estimation. (1) taxel-specific responses to single-axis forces, and (2) combined force-torque estimation results. The goal of this section is to infer whether the predictions of 4.4 are correct, and to act as a stand-alone empirical model of how multiple taxels embedded in a single elastomer react to applied forces and torques.

4.6.1 Individual Taxel Responses

This section presents individual taxel responses to single-axis loads, in order to motivate and inform the following section on combined force-torque estimation. Development of typical strain-gauge sensors and force-torque sensors includes generating *calibration curves*, which are functions that represent transducer response to applied loads. While the final force-torque estimation method of this paper will use a data-driven method that does not need explicit taxel response curves, we present them here in order to provide more transparency into the analysis of the force-torque estimation results.

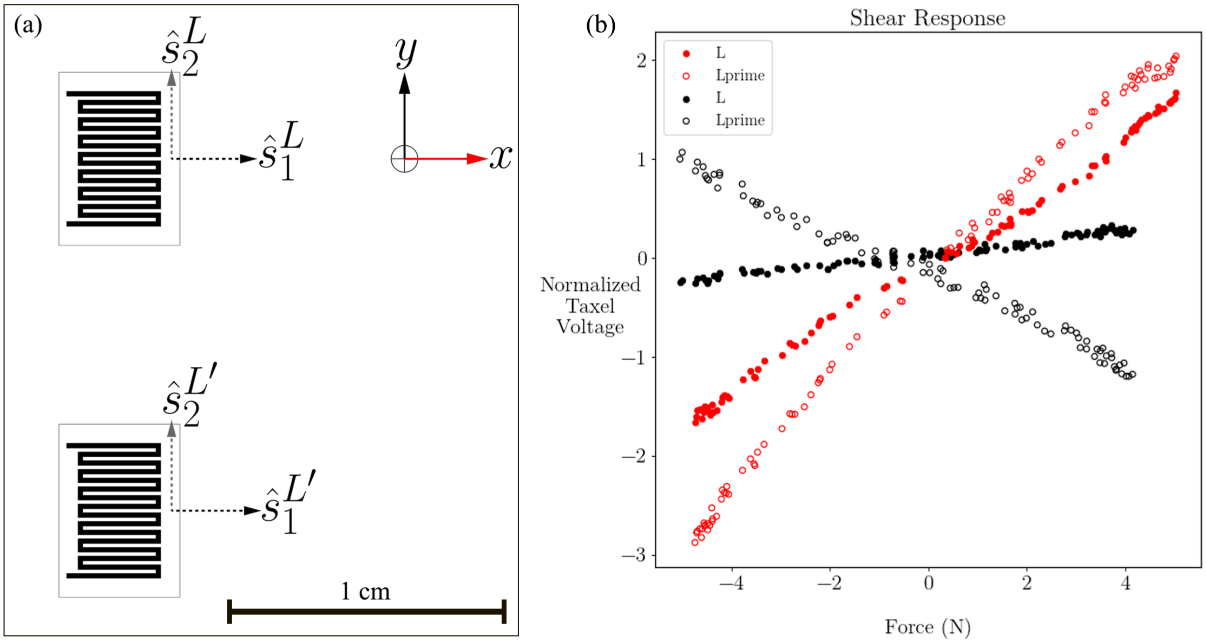

Figure 4.6: Shear force taxel responses to applied linear force within the XY plane. Principal deformations abide by the behavior outlined in the theoretical model, with positive and negative deformations resulting in positive and negative taxel responses.

4.6.1.1 Responses to a Shear force Load

The key assumption that underpins the hypothesis that off-axis shear placement enables torque sensing is investigated in Fig. 4.6. This assumption, that a taxel's response to a local deformation along the secondary axis is unsigned, while true per the model, has not been shown in practice.

This result could indicate several things. During the experiment, care was taken to locate the contact patch to be colinear with the on-axis shear force taxels. Due to either variability during fabrication, or failure to find the optimal contact location, the on-axis shear force taxels were, in actuality, slightly off-axis. Another assumption that could be incorrect is that the secondary axis is assumed to be unfixed, that a right-handed skew transformation is in practice not equivalent to a left-handed transformation. This again could be due to fabrication issues, such as a thicker material surrounding one edge of the taxel vs. the other. Most generally, the assumption of a linear sensor model with point contact might not be applicable, and the model might not capture the mechanics of soft body contact.

4.6.1.2 Responses to 1D Torque Load

We conducted an experiment to investigate taxel response as a function of applied torque in Fig. 4.7. The function of this experiment was to corroborate and supplement the theoretical discussion and modeling of Section 4.4, while also providing an analytical basis for the combined force-torque estimation results in the following section.

It is clear from Fig. 4.7 that the shear taxels labeled $[L', R']$ are the most sensitive to applied torque. According to the geometric argument of Section 4.4.2.2, it is expected that shear taxels whose principal axes makes the greatest angle with their taxel origin vector should exhibit the highest torque sensitivity.

Notably, while normal taxels are not designed in a way which is conducive to torque measurement, the normal taxel labeled $N2$ shows an appreciable response to applied torque. The local deformation at the origin of $N2$ as a result of an applied torque at the global sensor origin should be coincident with the taxel's local s_2 axis. While the taxel response plot of this taxel becomes nonlinear around -20 Nmm, the torque response is nonetheless bidirectional. This behavior could arise from taxel placement during manufacturing, in which the taxel's principal axis $\hat{n}$ creates a nonzero angle with the taxel origin vector. Alternatively, the applied torque could cause a change in the normal taxel's microchannel geometry that is inherently bidirectional, such as furling and unfurling the spiral.

It is worth considering the sign of the slope of the taxel response curve. For a positive torque $+\tau$ applied at the sensor's global origin, the direction of local taxel deformations extend from their origins in the direction tangent to a circle centered at the origin according to the right hand rule. The curves for $[R', L']$ are clearly positive, as their local deformations due to torque are most aligned with their principal axes. However, the slopes of the other taxels are less intuitive. It is possible, in the case of taxel R for example, that the effective taxel origin is displaced slightly along the $-y$ global axis. If this were the case, a positive torque at the sensor origin would result in a compressive force along the taxel's principal axis s_1, creating a negative change in taxel voltage. This is likely the case, as taxel R shows a weakly negative slope – weekly negative because it's principal axis is not aligned with the local deformation, and negative because a positive torque results in a negative voltage change. A similar rhetoric can be used to explain the slopes and signs of the other taxels.

The behavior of the shear taxel D as a function of torque appears to be a piece-wise linear function, with a greater sensitivity to negative torque than to positive torque. Whereas the shear taxels $[R, R', L, L']$ are designed and manufactured identically, D was designed differently. According to requirements inherited from the work in 3, shear taxel D was embedded in the region between the proximal and distal contact faces of the sensor skin (the latter of which was not used in this study). As a result, the orientation of taxel D within the surrounding elastomer is unique. Additionally, the piece-wise behavior could be due to compressible inhomogeneities in the surrounding elastomer, which is a known problem with this taxel. These compressible inhomogeneities could create a situation in which the surrounding elastomer was more stiff in one direction than the other.

Ultimately, while the idealized model would predict that only $[R', L']$ exhibit an appreciable response to applied torque, it is clear that all taxels exhibit some sensitivity. Whether this is due to imperfections in manufacturing or imperfections in the idealized model is

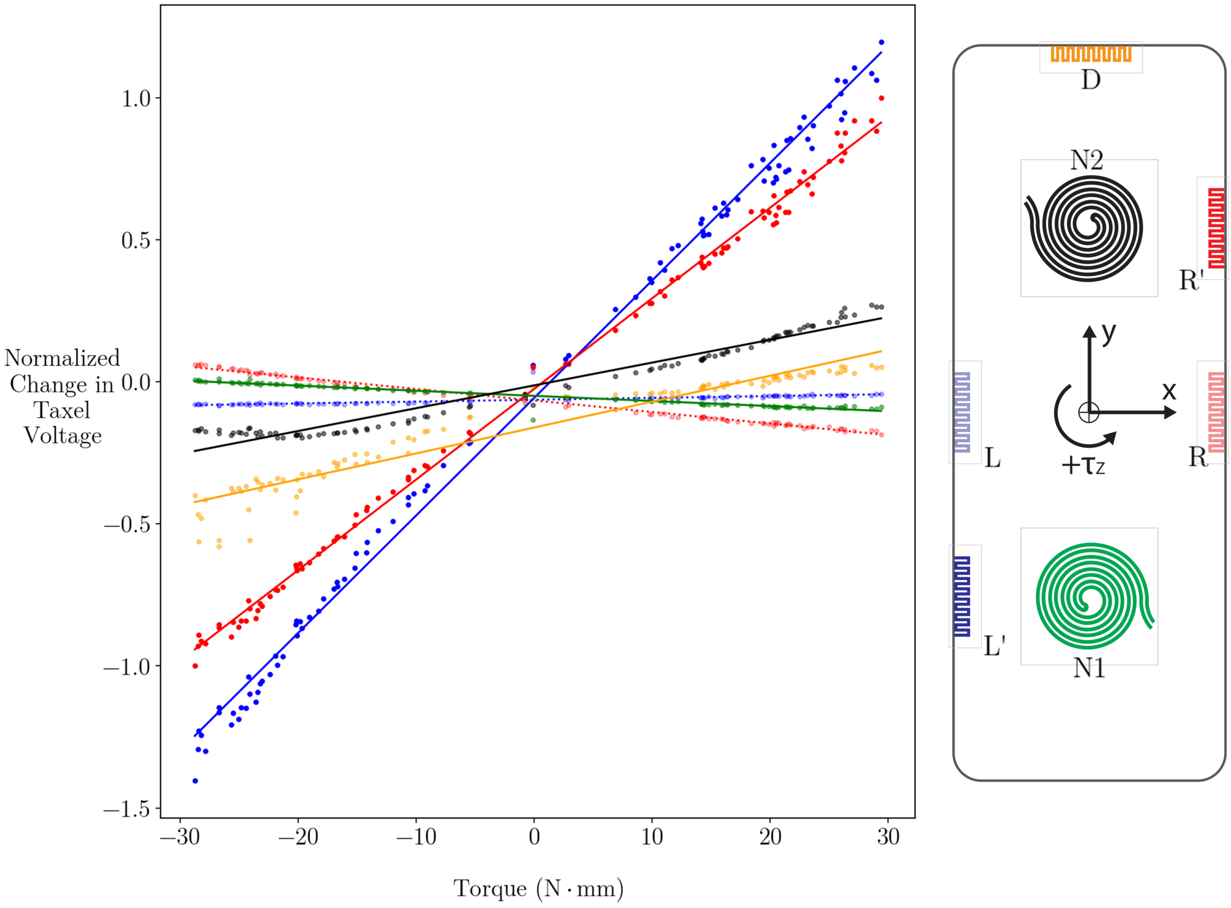

Figure 4.7: Plot of taxel responses to applied torque. Each data point represents an equilibrated taxel response to a static load.

unknown.

4.6.2 Combined Force-Torque Estimation

With the taxel behavior now characterized in terms of individual linear forces and torques, we now turn to the task of simultaneous force-torque estimation using an arrangement of multiple taxels within a single sensor skin. The purpose of this section is two-fold: (1) to demonstrate the multi-axis estimation capabilities of our skin-like tactile sensor, and (2) to

Table 4.1: Force-torque estimation performance of a selection of sensor configurations. Configurations are ordered from left to right according to increasing force-torque estimation performance. The highest R^2 value for each configuration is in bold. Row in blue represents the R^2 averaged across features for a given configuration.

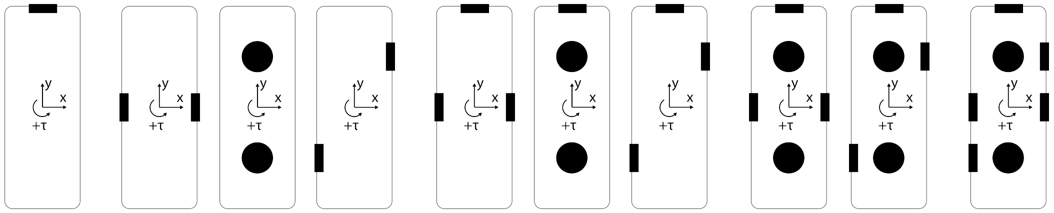

R^2	Distal Only	On-axis Only	Normal Only	Off-axis Only	On-axis + Distal	Normal + Distal	Off-axis + Distal	On-axis + Distal + Normal	Off-axis + Distal + Normal	All
f_x-only	0.023	**0.991**	0.432	**0.875**	**0.994**	0.751	**0.990**	**0.997**	**0.997**	**0.998**
f_y-only	**0.918**	0.081	0.347	0.309	0.978	**0.970**	0.979	0.993	0.995	0.996
f_z-only	0.104	0.711	**0.699**	0.313	0.929	0.887	0.719	0.988	0.987	0.991
τ_z-only	0.102	0.514	0.588	0.694	0.763	0.877	0.923	0.984	0.995	0.995
R^2 mean for $\underline{f}, \tau_z$	*0.287*	*0.574*	*0.516*	*0.548*	*0.916*	*0.871*	*0.903*	*0.991*	*0.993*	***0.995***

empirically support the modeling and characterization efforts of the previous sections.

To this end, we evaluate not only the performance of the complete seven-taxel sensor array shown in 4.4, but also a number of configurations that use a subset of all available taxels. Table 4.6.2 contains the configurations used in the study, which are arranged from left to right in order of increasing number of taxels.

It is useful to consider that individual taxels provide information to the CNN model, inasmuch as a certain taxel should be important for the encoding of a corresponding force-torque feature. For example, the *Distal Only* configuration results show that the distal taxel is important for the estimation of f_y. This corroborates both the microfluidic design and the taxel placement intuition developed earlier, as the taxel's microfluidic features and its principal axis align with the sensor's global y-axis. Moving to the right it's equivalently clear that the *On-axis Only* configuration, with two taxels aligned with the global x-axis, successfully encodes f_x. Interestingly, this pattern does not hold as clearly for the *Normal Only* configuration, instead showing only marginally higher R^2 for f_z estimation than for the other force-torque features. This is in part due to the surprising fact that the *Normal Only* configuration appears to provide a significant amount of information to the CNN model as it relates to the other force-torque features. The torque response experiment (Figure 4.7 also supports this finding, as the normal taxel $N1$ showed a significant sensitivity to applied torque.

Comparing the force-torque estimation performance of each configuration against the performance of the *All* configuration reveals the relative performance of each configuration, denoted as $\Delta R^2(\tau_z - only)$. This allows true side-by-side comparison between configurations, which we can use to corroborate the model and the taxel characterization studies and verify our taxel design considerations of 4.4.2.

Enabling torque estimation by moving the shear force taxels to be off-axis was the primary hypothesis of this paper, and close inspection of 4.6.2 supports this hypothesis. The configurations are grouped according to the number of taxels in the configuration, and sorted in increasing order. Within each grouping, the highest-performing configuration in terms of torque estimation performance was the configuration with off-axis shear force taxels.

Another suggestion was the placement of normal taxels closest to the anticipated point of contact. Surprisingly, the result that most clearly supports this suggestion is seen by comparing the *On-axis Only* and *Normal Only* configurations. Even though the on-axis shear force taxels are not designed to encode normal force, the configuration outperforms the *Normal Only* configuration in estimation of f_z. It is clear that taxels of any modality are capable of encoding force-torque features which are not their primary purpose. In the case of on-axis shear force taxels, this is amplified by the taxels' close proximity to the point of contact; the on-axis shear force taxels are the closest to the sensor's origin, where the load is applied.

4.7 Discussion

4.7.1 Sufficient Torque Estimation Performance from On-axis Shear Force Taxels

The model laid out in Sec XYZ predicts that the on-axis shear force taxels of the sensor skin should perform poorly on torque estimation since the local taxel deformation due to torque is predominantly along the taxel's secondary axis. As shown in Fig A and Fig XY, the experimental results do not entirely corroborate this prediction. While the torque sensitivity of on-axis shear force taxels is less than that of the off-axis taxels, the hypothesis that on-axis shear force taxel response is unsigned doesn't appear in the experimental results. We believe that, in reality, the "on-axis" shear force taxels are slightly off-axis due to inaccuracy in the experimental setup and during fabrication. We believe the reason the on-axis model effectively measures torque is because the CNN is robust enough to extract features from the on-axis taxel readings, as poor and insensitive as they might be.

We believe that the results presented Fig ABA, showing higher torque sensitivity of off-axis taxels, is sufficient to corroborate the hypothesis that in-plane torque estimation is enhanced by placing shear force taxels remotely to the axis of applied torque.

4.7.2 Limitations

4.7.2.1 Generalization to arbitrary contact geometry

The mathematical formulation presented in Sec. 4.4 is an oversimplification of the contact dynamics and nonlinear strain behavior of the sensor skin elastomer. In the experiments of Sec. 4.5, care was taken to ensure that the center of contact and orientation of the fingertip were held constant. Additionally, the orientation and surface properties of the contacted surface were held constant. The resultant dataset showed repeatable taxel response and showed excellent force/torque estimation performance.

While holding variables constant allowed us to demonstrate a hypothesis about force/torque estimation, the full solution to the problem of "tactile state estimation" is of a much higher dimension than our experimental testbed was capable of exploring.

4.7.2.2 Overfitting

The full-sensor force-torque estimation performance in Table 4.6.2 should be taken with a grain of salt, as the dataset it was trained on was generated in a repeatable and heavily constrained environment. This work was focused on validating our proposed design criteria for sparse tactile sensor skins, and not on creating a generalizable force-torque sensing model.

4.8 Conclusion

A set of design considerations has been proposed for tactile sensor skins using embedded single-axis strain gauges with the purpose of enabling the estimation of applied 3D forces and 1D torque about the skin's surface normal. By displacing the shear force taxels such that their principal axes do not intersect the anticipated center of contact, the sensor is able to more accurately measure torque about the sensor skin's surface normal. Force-torque estimation results from a CNN-based estimator confer with both the single-axis load characteristics and the analytical modeling that informed the design considerations.

4.9 Acknowledgement

The authors are grateful for support from RE2 Robotics. This work was supported in part by Office of Naval Research Awards N00014-19-C-2008 and N00014-18-S-B003.

CHAPTER 5

Summary and Conclusion

5.1 Contributions

Kinematic Parameter Estimation of the Human Hand: As was made clear by the palmless glove experiments, hand sensor networks are prone to miscalibration and unwanted sensor movement. To remediate these issues, we developed a data-driven method to identify the underlying hand kinematics using a sensor glove in the presence of sensor error. The method derived in Chapter 2 is used to estimate the parameters of a 2-DOF joint, such as the metacarpophalangeal (MCP) joint of the human hand. This method was verified to estimate joint axis orientations to within a few degrees of the ground truth measurements, and it outperforms existing estimation methods in terms of reliability and computational complexity.

Skin-like Tactile Sensor Array for Underwater, Pressurized Environments: We developed a skin-like tactile sensor array capable of multimodal sensing that is invariant to the effects of an underwater, pressurized environment. Designed as a multi-layer stack of durable elastomer and liquid-metal strain gauge, the sensor was evaluated in both dry and underwater environments. Taxel sensitivity to applied loads was shown to be constant as a function of simulated water depth. We demonstrated the sensor's multimodal force sensing capability by simultaneously applying normal force and two directions of shear force while in an underwater, pressurized environment.

Tactile Sensor Skin Design for Force-Torque Sensing: We provide criteria for the geometric placement of individual sensing elements (taxels) within an elastomeric contact medium in order to optimize the sensitivity of the sensor skin during contact. We test the hypothesis that a certain placement of taxels within the medium will enable the sensor skin to accurately measure in-plane torque. We use an experimental testbed to create contact forces and torques at the surface of the sensor skin, and a data acquisition system to capture both taxel readings and ground-truth force-torque readings. To investigate torque estimation capabilities of the sensor, we train a simple CNN model on the data collected with the testbed. By comparing the performance metrics of the CNN model using subsets of the taxel array, we are able to validate our suggested design criteria for torque sensing. In addition to in-plane torque sensing, we show the skin-like sensor is capable of 4-axis force-torque estimation (three axes of linear force and 1 axis of torque).

5.2 Future Work

5.2.1 Kinematic Parameter Estimation of the Human Hand

The kinematic parameter estimation method presented in Chapter 2 represents the starting point toward creating anatomically correct kinematic models of the hand. The methods developed so far are extensible in two ways. First, while Chapter 2 was focused on the estimation of kinematic parameters of a single joint, the method can be used recursively and in parallel to estimate the multiple serial linkages of human hands. Secondly, the mathematical framework used in the method provides the ability to estimate not only joint axes (the focus of Chapter 2), but also the common normal distance between axes. With the additional estimation of the common normal, the extended estimation method could be used to identify a complete set of kinematic parameters of the hand. This method could become part of a calibration sequence prior to interactive VR applications, such as the user donning a pair of sensor gloves and quickly being able to enter into the virtual setting intimately and accurately.

5.2.2 Force-Torque Estimation in Underwater Environments via Tactile Sensor Skins

Chapters 3 and 4 lay out a framework for advanced tactile sensing in an underwater, pressurized environment. This framework includes hardware, design optimization, and software considerations for a sparse taxel array consisting of embedded liquid metal strain gauges. Other than the hardware validation and brief characterization of Chapter 3, the ultimate goal of underwater force-torque estimation using a skin-like tactile sensor was not comprehensively demonstrated. The technical risk of using the sensor skin at ocean depth was addressed partially, as Chapter 3 demonstrated taxel sensitivity was not appreciably affected by hydrostatic pressure. However, the holistic behavior of the sensor skin during submerged applications could succumb to a variety of issues. A full demonstration of the capabilities laid out in Chapters 3 and 4 would require an underwater robotic testbed, where forces and torques could be applied to the sensor skin in situ. In the same way that Chapter 3 laid out individual taxel response as a function of hydrostatic pressure, a more complete treatment of underwater force-torque estimation would verify the constraints under which the sensor skin can be operated in underwater, pressurized environments.

5.2.3 Tactile State Estimation via Tactile Sensor Skins

Chapter 5 provides a walk-through of how to use a microfluidic-based tactile sensor array to estimate 3D forces and 1D torques. It lays out a principled framework for the design of such sensors, and shows force-torque estimation results using a CNN-based tactile force-torque estimation method. The contribution of Chapter 5 can be extended both inward, toward the mathematical foundations, and outward, toward additional sensing capabilities and applications. Regarding the mathematical foundations, the characterization of elastomer-embedded strain gauges requires treatment of higher-order effects. As all taxels are embedded in a single elastomer, there are inter-taxel correlations and nonlinear stress-strain behavior of the embedding material that need to be accounted for. A more accurate sensor model could be attained through material mechanics analysis, multiphysics simulation such as COMSOL, or both.

Additionally, the tactile force-torque estimation demonstrated in Chapter 5 was achieved using a repeatable testbed that provided a high-quality dataset. In order to use the proposed method in an unstructured environment, further work is needed to test the model on unstructured data. It is anticipated that refining the pre-trained CNN-based force-torque estimation model of Chapter 5 using data gathered in a less repeatable environment should improve the generalization of tactile force-torque measurements.

REFERENCES

[1] Emanuel Todorov. Probabilistic inference of multijoint movements, skeletal parameters and marker attachments from diverse motion capture data. *IEEE Transactions on Biomedical Engineering*, 54(11):1927–1939, 2007. ISBN: 0018-9294 (Print)\n0018-9294 (Linking).

[2] Philipp Muller, Marc Andre Begin, Thomas Schauer, and Thomas Seel. Alignment-Free, Self-Calibrating Elbow Angles Measurement Using Inertial Sensors. *IEEE Journal of Biomedical and Health Informatics*, 21(2):312–319, 2017.

[3] Jeffrey A. Reinbolt, Jaco F. Schutte, Benjamin J. Fregly, Byung Il Koh, Raphael T. Haftka, Alan D. George, and Kim H. Mitchell. Determination of patient-specific multi-joint kinematic models through two-level optimization. *Journal of Biomechanics*, 38(3):621–626, 2005.

[4] Mitja Veber and Tadej Bajd. Assessment of human hand kinematics. *Proceedings - IEEE International Conference on Robotics and Automation*, 2006(May):2966–2971, 2006. Publisher: IEEE ISBN: 0780395069.

[5] Christina Strohrmann, Holger Harms, Cornelia Kappeler-Setz, and Gerhard Tröster. Monitoring kinematic changes with fatigue in running using body-worn sensors. *IEEE Transactions on Information Technology in Biomedicine*, 16(5):983–990, 2012.

[6] Xi Chen. Human motion analysis with wearable inertial sensors. 2013.

[7] Luis Unzueta, Manuel Peinado, Ronan Boulic, and Ángel Suescun. Full-body performance animation with Sequential Inverse Kinematics. *Graphical Models*, 70(5):87–104, 2008. Publisher: Elsevier Inc.

[8] Eng Jon Ong and Adrian Hilton. Learnt inverse kinematics for animation synthesis. *Graphical Models*, 68(5-6):472–483, 2006.

[9] Vicon Nexus Reference Guide. page 356.

[10] Huiyu Zhou and Huosheng Hu. Human motion tracking for rehabilitation—a survey. *Biomedical Signal Processing and Control*, 3(1):1–18, 1 2008.

[11] Scott C. White and Robert M. Lifeso. Altering asymmetric limb loading after hip arthroplasty using real-time dynamic feedback when walking. *Archives of Physical Medicine and Rehabilitation*, 86(10):1958–1963, 10 2005.

[12] H J Luinge and P H Veltink. Measuring orientation of human body segments using miniature gyroscopes and accelerometers. *Kinematics Med. Biol. Eng. Comput*, 43:273–282, 2005.

[13] Hilton Moeslund and Ger Krü. A survey of advances in vision-based human motion capture and analysis. 2006.

[14] Sebastian Madgwick. An efficient orientation filter for inertial and inertial/magnetic sensor arrays.

[15] Hangxin Liu, Xu Xie, Matt Millar, Mark Edmonds, Feng Gao, Yixin Zhu, Veronica J. Santos, Brandon Rothrock, and Song-Chun Zhu. A glove-based system for studying hand-object manipulation via joint pose and force sensing. Vancouver, Canada, 9 2017. [Online; accessed 2017-09-13].

[16] James Connolly, Joan Condell, Brendan O'Flynn, Javier Torres Sanchez, and Philip Gardiner. Imu sensor-based electronic goniometric glove (iseg-glove) for clinical finger movement analysis. *IEEE Sensors Journal*, pages 1–1, 2017.

[17] Henk G Kortier, Victor I Sluiter, Daniel Roetenberg, and Peter H Veltink. Assessment of hand kinematics using inertial and magnetic sensors. *Journal of NeuroEngineering and Rehabilitation*, 11(1):70, 2014.

[18] Nicholas Wettels, Veronica J. Santos, Roland S. Johansson, and Gerald E. Loeb. Biomimetic tactile sensor array. *Advanced Robotics*, 22:829–849, 8 2008.

[19] Wenzhen Yuan, Rui Li, Mandayam A. Srinivasan, and Edward H. Adelson. Measurement of shear and slip with a gelsight tactile sensor. pages 304–311, Seattle, WA, USA, 5 2015. 2015 IEEE International Conference on Robotics and Automation (ICRA), IEEE. [Online; accessed 2020-02-19].

[20] Zhe Su, Karol Hausman, Yevgen Chebotar, Artem Molchanov, Gerald E. Loeb, Gaurav S. Sukhatme, and Stefan Schaal. Force estimation and slip detection/classification for grip control using a biomimetic tactile sensor. In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, pages 297–303, Seoul, South Korea, November 2015. IEEE.

[21] N. Wettels, A.R. Parnandi, Ji-Hyun Moon, G.E. Loeb, and G.S. Sukhatme. Grip control using biomimetic tactile sensing systems. *Mechatronics, IEEE/ASME Transactions on*, 14(6):718–723, 2009.

[22] Xiaojing Song, Hongbin Liu, Kaspar Althoefer, Thrishantha Nanayakkara, and Lakmal D. Seneviratne. Efficient break-away friction ratio and slip prediction based on haptic surface exploration. *IEEE Transactions on Robotics*, 30(1):203–219, 2 2014.

[23] D. J. O'Brien and D. M. Lane. Force and explicit slip sensing for the amadeus underwater gripper. *International Journal of Systems Science*, 29(5):471–483, 5 1998.

[24] P.J. Sanz, A. Peñalver, J. Sales, D. Fornas, J.J. Fernández, J. Pérez, and J. Bernabé. Grasper: A multisensory based manipulation system for underwater operations. pages 4036–4041. 2013 IEEE International Conference on Systems, Man, and Cybernetics, 10 2013. ISSN: 1062-922X.

[25] Peter Kampmann and Frank Kirchner. A Tactile Sensing System for Underwater Manipulators. page 5.

[26] T. André, V. Lévesque, V. Hayward, P. Lefèvre, and J.-L. Thonnard. Effect of skin hydration on the dynamics of fingertip gripping contact. *Journal of The Royal Society Interface*, 8(64):1574–1583, 11 2011. PMID: 21490002.

[27] Sergio Nacht, Jo-ann Close, and David Yeung. *Synopsis*.

[28] F. E. Barrientos, Y. Karayiannidis, K. Pauwels, C. C. Smith, and D. Kragic. In-hand manipulation using gravity and controlled slip. pages 5636–5641, Hamburg, Germany, 10 2015. Proc IEEE/RSJ Intl Conf on Intelligent Robots and Systems.

[29] Zvi S. Roth, Benjamin W. Mooring, and Bahram Ravani. An Overview of Robot Calibration. *IEEE Journal on Robotics and Automation*, 3(5):377–385, 1987.

[30] W. R. Taylor, E. I. Kornaropoulos, G. N. Duda, S. Kratzenstein, R. M. Ehrig, A. Arampatzis, and M. O. Heller. Repeatability and reproducibility of OSSCA, a functional approach for assessing the kinematics of the lower limb. *Gait and Posture*, 32(2):231–236, 2010. Publisher: Elsevier B.V.

[31] David Daney, Yves Papegay, and Blaise Madeline. Choosing measurement poses for robot calibration with the local convergence method and tabu search. *The International Journal of Robotics Research*, 24(6):501–518, 2005.

[32] Kiyoshi Hirose, Hitoshi Doki, and Akiko Kondo. A measurement method of the 2DOF joint angles and angular velocities using inertial sensors. *SICE Annual Conference (SICE), . . .*, pages 366–371, 2012. Publisher: IEEE.

[33] Daniel L. Benoit, Dan K. Ramsey, Mario Lamontagne, Lanyi Xu, Per Wretenberg, and Per Renström. Effect of skin movement artifact on knee kinematics during gait and cutting motions measured in vivo. *Gait and Posture*, 24(2):152–164, 2006.

[34] Eugene J. Alexander and Thomas P. Andriacchi. Correcting for deformation in skin-based marker systems. *Journal of Biomechanics*, 34(3):355–361, 2001.

[35] Rainald M. Ehrig, William R. Taylor, Georg N. Duda, and Markus O. Heller. A survey of formal methods for determining functional joint axes. *Journal of Biomechanics*, 40(10):2150–2157, 2007.

[36] Thomas Seel, Thomas Schauer, and Jörg Raisch Raisch. Joint Axis and Position Estimation from Inertial Measurement Data by Exploiting Kinematic Constraints. pages 0–4, 2012. ISBN: 9781467345040.

[37] Agneta Gustus, Georg Stillfried, Judith Visser, Henrik Jörntell, and Patrick van der Smagt. Human hand modelling: kinematics, dynamics, applications. *Biological Cybernetics*, 106(11-12):741–755, December 2012.

[38] Kevin M Lynch and Frank C Park. *Modern robotics.* Cambridge University Press, 2017.

[39] Christopher M Bishop and Markus Svens. GTM : The Generative Topographic Mapping. (1982), 1998.

[40] Ian M. Bullock, Julia Borras, and Aaron M. Dollar. Assessing assumptions in kinematic hand models: A review. In *2012 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, pages 139–146, Rome, Italy, June 2012. IEEE.

[41] DJ Giurintano, AM Hollister, WL Buford, DE Thompson, and LM Myers. A virtual five-link model of the thumb. *Medical engineering & physics*, 17(4):297–303, 1995.

[42] Frank C. Park and James E. Bobrow. Geometric optimization algorithims for robot kinematic design. *Journal of Robotic Systems*, 12(6):453–463, 1995.

[43] L. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, February 1989.

[44] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society*, B39:1–38, 1977.

[45] Adriano O. Andrade, Slawomir Nasuto, Peter Kyberd, and Catherine M. Sweeney-Reed. Generative topographic mapping applied to clustering and visualization of motor unit action potentials. *BioSystems*, 82(3):273–284, 2005.

[46] Nikolaos Gianniotis and Peter Tiňo. Visualization of tree-structured data through generative topographic mapping. *IEEE Transactions on Neural Networks*, 19(8):1468–1493, 2008.

[47] Dragomir Ž. Djoković. An elementary proof of the Baker-Campbell-Hausdorff-Dynkin formula. *Mathematische Zeitschrift*, 143(3):209–211, October 1975.

[48] Andrew A. Goldenberg, B. Benhabib, and Robert G. Fenton. A Complete Generalized Solution to the Inverse Kinematics of Robots. *IEEE Journal on Robotics and Automation*, 1(1):14–20, 1985.

[49] John Lin, Ying Wu, and T.S. Huang. Modeling the constraints of human hand motion. In *Proceedings Workshop on Human Motion*, pages 121–126, Los Alamitos, CA, USA, 2000. IEEE Comput. Soc.

[50] R.S. Dahiya, G. Metta, M. Valle, and G. Sandini. Tactile Sensing—From Humans to Humanoids. *IEEE Transactions on Robotics*, 26(1):1–20, February 2010.

[51] Akihiko Yamaguchi and Christopher G. Atkeson. Recent progress in tactile sensing and sensors for robotic manipulation: can we turn tactile sensing into vision? *Advanced Robotics*, 33(14):661–673, July 2019.

[52] Zhanat Kappassov, Juan-Antonio Corrales, and Véronique Perdereau. Tactile sensing in dexterous robot hands — Review. *Robotics and Autonomous Systems*, 74:195–220, December 2015.

[53] Hanna Yousef, Mehdi Boukallel, and Kaspar Althoefer. Tactile sensing for dexterous in-hand manipulation in robotics—A review. *Sensors and Actuators A: Physical*, 167(2):171–187, June 2011.

[54] Wei Chen, Heba Khamis, Ingvars Birznieks, Nathan F Lepora, and Stephen J Redmond. Tactile sensors for friction estimation and incipient slip detection—toward dexterous robotic manipulation: A review. *IEEE Sensors Journal*, 18(22):9049–9064, 2018.

[55] Jianzhu Yin, Veronica J. Santos, and Jonathan D. Posner. Bioinspired flexible microfluidic shear force sensor skin. *Sensors and Actuators, A: Physical*, 264:289–297, 2017. Publisher: Elsevier B.V.

[56] Ruben D. Ponce Wong, Jonathan D. Posner, and Veronica J. Santos. Flexible microfluidic normal force sensor skin for tactile feedback. *Sensors and Actuators, A: Physical*, 179:62–69, 2012. Publisher: Elsevier B.V.

[57] Moaed A Abd, Mostapha Al-Saidi, Maohua Lin, Genevieve Liddle, Kunal Mondal, and Erik D Engeberg. Surface Feature Recognition and Grasped Object Slip Prevention With a Liquid Metal Tactile Sensor for a Prosthetic Hand. page 6.

[58] Gereon H. Büscher, Risto Kõiva, Carsten Schürmann, Robert Haschke, and Helge J. Ritter. Flexible and stretchable fabric-based tactile sensor. *Robotics and Autonomous Systems*, 63, Part 3:244–252, January 2015.

[59] Ho-Hsiu Chou, Amanda Nguyen, Alex Chortos, John W.F. To, Chien Lu, Jianguo Mei, Tadanori Kurosawa, Won-Gyu Bae, Jeffrey B.-H. Tok, and Zhenan Bao. A chameleon-inspired stretchable electronic skin with interactive colour changing controlled by tactile sensing. *Nature Communications*, 6(1):8011, November 2015.

[60] Frank L. Hammond, Rebecca K. Kramer, Qian Wan, Robert D. Howe, and Robert J. Wood. Soft tactile sensor arrays for micromanipulation. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 25–32, October 2012. ISSN: 2153-0866.

[61] Liang Zou, Chang Ge, Z. Wang, Edmond Cretu, and Xiaoou Li. Novel Tactile Sensor Technology and Smart Tactile Sensing Systems: A Review. *Sensors*, 17(11):2653, November 2017.

[62] Benjamin Ward-Cherrier, Nicholas Pestell, Luke Cramphorn, Benjamin Winstone, Maria Elena Giannaccini, Jonathan Rossiter, and Nathan F. Lepora. The TacTip Family: Soft Optical Tactile Sensors with 3D-Printed Biomimetic Morphologies. *Soft Robotics*, 5(2):216–227, April 2018.

[63] Mike Lambeta, Po-Wei Chou, Stephen Tian, Brian Yang, Benjamin Maloon, Victoria Rose Most, Dave Stroud, Raymond Santos, Ahmad Byagowi, Gregg Kammerer, Dinesh Jayaraman, and Roberto Calandra. DIGIT: A Novel Design for a Low-Cost Compact High-Resolution Tactile Sensor With Application to In-Hand Manipulation. *IEEE Robotics and Automation Letters*, 5(3):3838–3845, July 2020. Conference Name: IEEE Robotics and Automation Letters.

[64] Yashraj S. Narang, Balakumar Sundaralingam, Karl Van Wyk, Arsalan Mousavian, and Dieter Fox. Interpreting and predicting tactile signals for the SynTouch BioTac. *The International Journal of Robotics Research*, 40(12-14):1467–1487, December 2021. Publisher: SAGE Publications Ltd STM.

[65] Yanmei Li and Imin Kao. A review of modeling of soft-contact fingers and stiffness control for dextrous manipulation in robotics. In *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No.01CH37164)*, volume 3, pages 3055–3060 vol.3, May 2001. ISSN: 1050-4729.

[66] Matthew T Mason. *Mechanics of robotic manipulation*. MIT Press, 2001.

[67] Richard M. Murray, Zexiang Li, and S. Shankar Sastry. A Mathematical Introduction to Robotic Manipulation, 1994.

[68] Yashraj Narang, Balakumar Sundaralingam, Miles Macklin, Arsalan Mousavian, and Dieter Fox. Sim-to-Real for Robotic Tactile Sensing via Physics-Based Simulation and Learned Latent Projections. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6444–6451, May 2021. ISSN: 2577-087X.

[69] A. Bicchi and V. Kumar. Robotic grasping and contact: a review. In *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*, volume 1, pages 348–353 vol.1, April 2000. ISSN: 1050-4729.

[70] G. Cannata and Marco Maggiali. Design of a Tactile Sensor for Robot Hands. December 2008.

[71] Giuseppe De Maria, Ciro Natale, and Salvatore Pirozzi. Force/tactile sensor for robotic applications. *Sensors and Actuators A: Physical*, 175:60–72, 2012.

[72] Alexander Schmitz, Perla Maiolino, Marco Maggiali, Lorenzo Natale, Giorgio Cannata, and Giorgio Metta. Methods and Technologies for the Implementation of Large-Scale Robot Tactile Sensors. *IEEE Transactions on Robotics*, 27(3):389–400, June 2011. Conference Name: IEEE Transactions on Robotics.

[73] Daniel M. Vogt, Yong-Lae Park, and Robert J. Wood. Design and Characterization of a Soft Multi-Axis Force Sensor Using Embedded Microfluidic Channels. *IEEE Sensors Journal*, 13(10):4056–4064, October 2013.

[74] Jianzhu Yin, Peter Aspinall, Veronica J. Santos, and Jonathan D. Posner. Measuring Dynamic Shear Force and Vibration with a Bioinspired Tactile Sensor Skin. *IEEE Sensors Journal*, 18(9):3544–3553, 2018.

[75] Nguyen Thanh-Vinh, Nguyen Binh-Khiem, Hidetoshi Takahashi, Kiyoshi Matsumoto, and Isao Shimoyama. High-sensitivity triaxial tactile sensor with elastic microstructures pressing on piezoresistive cantilevers. *Sensors and Actuators A: Physical*, 215:167–175, August 2014.

[76] Yuanzhao Wu, Yiwei Liu, Youlin Zhou, Qikui Man, Chao Hu, Waqas Asghar, Fali Li, Zhe Yu, Jie Shang, Gang Liu, Meiyong Liao, and Run-Wei Li. A skin-inspired tactile sensor for smart prosthetics. *SCIENCE ROBOTICS*, page 9, 2018.

[77] Sihong Wang, Jie Xu, Weichen Wang, Ging-Ji Nathan Wang, Reza Rastak, Francisco Molina-Lopez, Jong Won Chung, Simiao Niu, Vivian R. Feig, Jeffery Lopez, Ting Lei, Soon-Ki Kwon, Yeongin Kim, Amir M. Foudeh, Anatol Ehrlich, Andrea Gasperini, Youngjun Yun, Boris Murmann, Jeffery B.-H. Tok, and Zhenan Bao. Skin electronics from scalable fabrication of an intrinsically stretchable transistor array. *Nature*, 555(7694):83–88, March 2018.

[78] Lili Wang, Joshua A. Jackman, Ee-Lin Tan, Jae Hyeon Park, Michael G. Potroz, Ee Taek Hwang, and Nam-Joon Cho. High-performance, flexible electronic skin sensor incorporating natural microcapsule actuators. *Nano Energy*, 36:38–45, June 2017.

[79] Yong-Lae Park, Bor-Rong Chen, and R.J. Wood. Design and Fabrication of Soft Artificial Skin Using Embedded Microchannels and Liquid Conductors. *IEEE Sensors Journal*, 12(8):2711 –2718, August 2012.

[80] Xiaomei Shi, Ching-Hsiang Cheng, Yongping Zheng, and P K A Wai. An EGaIn-based flexible piezoresistive shear and normal force sensor with hysteresis analysis in normal

force direction. *Journal of Micromechanics and Microengineering*, 26(10):105020, October 2016.

[81] Rafsan Al Shafatul Islam Subad, Liam B. Cross, and Kihan Park. Soft Robotic Hands and Tactile Sensors for Underwater Robotics. *Applied Mechanics*, 2(2):356–383, June 2021.

[82] DANA R. YOERGER, ALBERT M. BRADLEY, MICHAEL JAKUBA, CHRISTOPHER R. GERMAN, TIMOTHY SHANK, and MAURICE TIVEY. AUTONOMOUS AND REMOTELY OPERATED VEHICLE TECHNOLOGY FOR HYDROTHERMAL VENT DISCOVERY, EXPLORATION, AND SAMPLING. *Oceanography*, 20(1):152–161, 2007. Publisher: Oceanography Society.

[83] Giacomo Marani, Song K. Choi, and Junku Yuh. Underwater autonomous manipulation for intervention missions AUVs. *Ocean Engineering*, 36(1):15–23, January 2009.

[84] Achint Aggarwal, Peter Kampmann, Johannes Lemburg, and Frank Kirchner. Haptic Object Recognition in Underwater and Deep-sea Environments. *Journal of Field Robotics*, 32(1):167–185, 2015. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/rob.21538.

[85] Dongsik Chang, Scott Chow, Timothy R Player, and Geoffrey A Hollinger. Adaptive and Informative Planning for an Underwater Vehicle-Manipulator System. page 6.

[86] Peter Kampmann and Frank Kirchner. A tactile sensing system for underwater manipulators. 2012.

[87] Nawid Jamali, Petar Kormushev, Arnau C. Viñas, Marc Carreras, and Darwin G. Caldwell. Underwater robot-object contact perception using machine learning on force/torque sensor feedback. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3915–3920, May 2015. ISSN: 1050-4729.

[88] Giovanni Gerardo Muscolo and Giorgio Cannata. A novel tactile sensor for underwater applications: Limits and perspectives. In *OCEANS 2015 - Genova*, pages 1–7, May 2015.

[89] Satja Sivčev, Joseph Coleman, Edin Omerdić, Gerard Dooly, and Daniel Toal. Underwater manipulators: A review. *Ocean Engineering*, 163:431–450, September 2018.

[90] Daniel M. Vogt, Kaitlyn P. Becker, Brennan T. Phillips, Moritz A. Graule, Randi D. Rotjan, Timothy M. Shank, Erik E. Cordes, Robert J. Wood, and David F. Gruber. Shipboard design and fabrication of custom 3D-printed soft robotic manipulators for the investigation of delicate deep-sea organisms. *PLOS ONE*, 13(8):e0200386, August 2018.

[91] J.R. Bemfica, C. Melchiorri, L. Moriello, G. Palli, U. Scarcia, and G. Vassura. Mechatronic Design of a Three-Fingered Gripper for Underwater Applications*. *IFAC Proceedings Volumes*, 46(5):307–312, 2013.

[92] David Ribas, Pere Ridao, Alessio Turetta, Claudio Melchiorri, Gianluca Palli, José Javier Fernández, and Pedro José Sanz. I-AUV Mechatronics Integration for the TRIDENT FP7 Project. *IEEE/ASME Transactions on Mechatronics*, 20(5):2583–2592, October 2015. Conference Name: IEEE/ASME Transactions on Mechatronics.

[93] J. R. Bemfica, C. Melchiorri, L. Moriello, G. Palli, and U. Scarcia. A three-fingered cable-driven gripper for underwater applications. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2469–2474, May 2014. ISSN: 1050-4729.

[94] Xuebo Yuan, Sang Min Won, Mengdi Han, Youshan Wang, John A. Rogers, Yonggang Huang, and Heling Wang. Mechanics of encapsulated three-dimensional structures for simultaneous sensing of pressure and shear stress. *Journal of the Mechanics and Physics of Solids*, 151:104400, June 2021.

[95] Mark A Eddings, Michael A Johnson, and Bruce K Gale. Determining the optimal PDMS–PDMS bonding technique for microfluidic devices. *Journal of Micromechanics and Microengineering*, 18(6):067001, June 2008.

[96] Caffiyar Yousuff, Mohd. Danish, Eric Ho, Ismail Kamal Basha, and Nor Hamid. Study on the Optimum Cutting Parameters of an Aluminum Mold for Effective Bonding Strength of a PDMS Microfluidic Device. *Micromachines*, 8(8):258, August 2017.

[97] Yiliang Lin, Olivia Gordon, M. Rashed Khan, Neyanel Vasquez, Jan Genzer, and Michael D. Dickey. Vacuum filling of complex microchannels with liquid metal. *Lab on a Chip*, 17(18):3043–3050, 2017.

[98] Markus Fritzsche, Norbert Elkmann, and Erik Schulenburg. Tactile sensing: A key technology for safe physical human robot interaction. In *Proceedings of the 6th International Conference on Human-robot Interaction*, pages 139–140.

[99] Ziwei Xia, Zhen Deng, Bin Fang, Yiyong Yang, and Fuchun Sun. A review on sensory perception for dexterous robotic manipulation. *International Journal of Advanced Robotic Systems*, 19(2):17298806221095974, 2022.

[100] Wenzhen Yuan, Siyuan Dong, and Edward H. Adelson. Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors*, 17(12), 2017.

[101] Wanlin Li, Akram Alomainy, Ivan Vitanov, Yohan Noh, Peng Qi, and Kaspar Althoefer. F-touch sensor: Concurrent geometry perception and multi-axis force measurement. *IEEE Sensors Journal*, 21(4):4300–4309, 2020.

[102] Akhil Padmanabha, Frederik Ebert, Stephen Tian, Roberto Calandra, Chelsea Finn, and Sergey Levine. Omnitact: A multi-directional high-resolution touch sensor. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 618–624. IEEE, 2020.

[103] Ravinder S Dahiya, Philipp Mittendorfer, Maurizio Valle, Gordon Cheng, and Vladimir J Lumelsky. Directions toward effective utilization of tactile skin: A review. *IEEE Sensors Journal*, 13(11):4121–4138, 2013.

[104] Ravinder Dahiya. E-skin: from humanoids to humans [point of view]. *Proceedings of the IEEE*, 107(2):247–252, 2019.

[105] Martin Weigel, Tong Lu, Gilles Bailly, Antti Oulasvirta, Carmel Majidi, and Jürgen Steimle. Iskin: flexible, stretchable and visually customizable on-body touch sensors for mobile computing. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, pages 2991–3000, 2015.

[106] Yong-Lae Park, Carmel Majidi, Rebecca Kramer, Phillipe Bérard, and Robert J Wood. Hyperelastic pressure sensing with a liquid-embedded elastomer. *Journal of Micromechanics and Microengineering*, 20(12):125029, nov 2010.

[107] Andrea Cirillo, Pasquale Cirillo, Giuseppe De Maria, Ciro Natale, and Salvatore Pirozzi. Control of linear and rotational slippage based on six-axis force/tactile sensor. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1587–1594. IEEE, 2017.

[108] Kanishka Ganguly, Pavan Mantripragada, Chethan M. Parameshwara, Cornelia Fermüller, Nitin J. Sanket, and Yiannis Aloimonos. Gradtac: Spatio-temporal gradient based tactile sensing. *Frontiers in Robotics and AI*, 9, 2022.

[109] Chen Wang, Shaoxiong Wang, Branden Romero, Filipe Veiga, and Edward Adelson. Swingbot: Learning physical features from in-hand tactile exploration for dynamic swing-up manipulation. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5633–5640, 2020.

[110] Qiang Li, Oliver Kroemer, Zhe Su, Filipe Fernandes Veiga, Mohsen Kaboli, and Helge Joachim Ritter. A review of tactile information: Perception and action through touch. *IEEE Transactions on Robotics*, 36(6):1619–1634, 2020.

[111] Balakumar Sundaralingam, Alexander Sasha Lambert, Ankur Handa, Byron Boots, Tucker Hermans, Stan Birchfield, Nathan Ratliff, and Dieter Fox. Robust learning of tactile force estimation through robot interaction. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 9035–9042. IEEE, 2019.

[112] Kenneth Gutierrez and Veronica J Santos. Perception of tactile directionality via artificial fingerpad deformation and convolutional neural networks. *IEEE Transactions on Haptics*, 13(4):831–839, 2020.