UNIVERSITY OF CALIFORNIA SAN DIEGO

Permanent Magnet-Based Localization for Growing Robots in Medical Applications

A thesis submitted in partial satisfaction of the
requirements for the degree of Master of Science

in

Engineering Sciences (Mechanical Engineering)

by

Connor Watson

Committee in charge:

Professor Tania Morimoto, Chair
Professor Michael Tolley
Professor Michael Yip

2019

The Thesis of Connor Watson is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California San Diego

2019

DEDICATION

For my mom and dad.

TABLE OF CONTENTS

Signature Page iii

Dedication iv

Table of Contents v

List of Figures vi

List of Tables vii

Acknowledgements viii

Abstract of the Thesis ix

Chapter 1 Introduction 1
1.1 Soft Robots 3
1.2 Continuum Robots 5
1.3 Growing Robots 7
1.4 Sensing for Continuum Robots 9
1.5 Magnetic Tracking 9
1.6 Contributions 11

Chapter 2 Methods 12
2.1 Dipole Model Based Localization 13
2.2 LSTM Based Localization 15
2.3 A Hybrid Approach to Localization 17

Chapter 3 Experiments and Results 19
3.1 Experimental Setup 19
3.2 Model Comparison 21
3.2.1 The Effect of the Quantity of Data on Model Performance 22
3.2.2 The Effect of Velocity on Model Performance 24
3.3 Localization at Different Scales 25
3.4 Tip Localization of a Growing Robot Deployed in Constrained Environments . 27

Chapter 4 Conclusions and Future Work 29

Bibliography 32

LIST OF FIGURES

Figure 1.1. An application of a soft robot gripper design 4

Figure 1.2. An example of a tendon driven continuum manipulaotor 6

Figure 1.3. A diagram illustrating a concept of a growing robot 8

Figure 2.1. Magnet position during robot growth 12

Figure 2.2. Figure of important localization parameters 14

Figure 2.3. Flowchart of localization approaches 16

Figure 3.1. Figure of test setup 20

Figure 3.2. Model performance with increasing amounts of data 23

Figure 3.3. Model performance tracking a target moving at different speeds 25

Figure 3.4. Workspace area for magnets of different sizes 26

Figure 3.5. Growing robot deployed in multiple constrained environments 28

Figure 3.6. Measured versus actual localiton of the growing robot tip 28

LIST OF TABLES

Table 3.1. Physical properties of the different magnets used in the experiments 19

ACKNOWLEDGEMENTS

First, I would like to thank Professor Morimoto for her invaluable support and guidance throughout the research process. She has not only provided me with the mentor-ship that has helped me to develop this work to its full potential, but also been a consistent role model for me as I develop professionally and learn what it means to be an academic researcher. I am very grateful to have had the opportunity to study under her.

Additionally I would like thank the other members of my committee. In particular, I would like to thank Professor Tolley for having me as a reader for a course he taught on robotics which both helped to support my education as well as further my understanding of the subject. I would also like to thank Professor Yip, whose reinforcement learning course has inspired me to consider new research possibilities for the future.

Further, I would like to thank Dr. Cedric Girerd for his help and advice with regards to this work, as well as his perspective on research in general. Finally, I would also like to thank the other members of the Moromoto Lab for their encouragement and support along the way.

This thesis, in part, has been submitted for publication of the material as it may appear in IEEE Robotics and Automation Letters, 2020, Watson, Connor; Morimoto, Tania K., IEEE, 2020. The dissertation author was the primary investigator and author of this paper.

ABSTRACT OF THE THESIS

Permanent Magnet-Based Localization for Growing Robots in Medical Applications

by

Connor Watson

Master of Science in Engineering Sciences (Mechanical Engineering)

University of California San Diego, 2019

Professor Tania Morimoto, Chair

Growing robots that achieve locomotion by extending from their tip, are inherently compliant and can safely navigate through constrained environments that prove challenging for traditional robots. However, the same compliance and tip-extension mechanism that enables this ability, also leads directly to challenges in their localization and control. In this thesis, we present a low-cost, wireless, permanent magnet-based method for localizing the tip of these robots in 5 degrees of freedom. A permanent magnet is placed at the robot tip, and an array of magneto-inductive sensors is used to measure the change in magnetic field as the robot moves through the workspace. We develop an approach to localization that combines analytical and machine learning techniques and show that it outperforms existing methods, particularly for localizing

fast moving magnets. We also measure the position error over a 50 cm $\times$ 50 cm workspace with different magnet sizes to show that this approach can accommodate growing robots of different scales. Finally, we demonstrate our tracking in real time by growing a 12 mm diameter robot through two different, constrained environments. On average for these experiments, our method of localization achieves a position error of 3.0 ± 1.1 mm and an orientation error of $6.5 \pm 5.4^{\circ}$.

Chapter 1

Introduction

The field of robotics has seen significant advancements over recent years, due to technological contributions spanning multiple disciplines. These developments have dramatically increased the capabilities of robots and have shifted the paradigm of robotics research from one concerned with the study of robots in structured, static environments–such as the factory floor–to one concerned with their interactions in dynamic and sometimes uncertain surroundings–such as the surgical theater. This new research frontier presents significant challenges regarding all aspects of robotics–from design to control. For instance, while robots have been successfully built and programmed to reliably accomplish assembly tasks in the automotive industry, when faced with the slightly more general task of sorting random warehouse goods into individual containers, robots are unable to approach human performance (see the results of the 2015 Amazon Picking Challenge [7]). The reason for this performance disparity is the environmental uncertainty in the later case. Because the environment is not entirely known beforehand, it is not immediately clear what kind of robot is best suited to the task.

Similar difficulties emerge in the study of surgical robotics as well. For example, anatomical variability between patients can make it challenging to design instruments or manipulators that work effectively for a broad range of body types and procedures. Real-time imaging, obtained through methods like ultrasound, can be imprecise and corrupted by noise [53]. The anisotropic stiffness of human tissue and its hard-to-model surface properties make it difficult

to accurately predict robot-tissue interactions [33]. Further, organs are not fixed in place and may move during an operation, changing the environment in real time. In addition to these challenges, surgical robots face even more demanding constraints posed by the medical environment itself [55]. Oftentimes these robots are targeted for use in minimally invasive procedures, meaning that the robot must be able to operate in a highly confined and fragile environment that is out of the direct line of sight of the clinician. Because it is challenging to ensure safety in such a scenario, most of the robotic surgery cases to date have been for procedures that could be accomplished with traditional laparoscopic tools and require the robot to be under the direct control of a highly trained surgeon at all times [46]. Despite these limitations, robotic surgery has already proved useful and sees increasing adoption by medical professionals as the technology develops. It allows for safer execution of existing techniques (see [39] for example) and in some cases even enables entirely new procedures [40]. Robots have great potential to continue to improve and extend the capabilities of modern medicine, but in order to do so, they must be reliably safe, even when faced with uncertainties.

Consequently, researchers have begun exploring design alternatives to traditional rigid-link robots that may be more suited to delicate tasks, such as surgery. In particular, soft robots, which are primarily fabricated out of compliant materials, offer a promising alternative. These robots are able to adapt to their surroundings by physically deforming their own bodies in ways that impart smaller forces on the surrounding environment as compared to rigid robots. These smaller forces make them inherently safer than rigid robots and ideal candidates for use in medical applications [27]. Another design alternative that has been explored is that of the continuum robot. Continuum robots are "snake-like" robots typically consisting of a long, flexible backbone that can be dexterously manipulated to curve and conform around obstacles in their paths [45] . The infinite degrees of freedom they posses in the continuously varying curvature of their backbones allow them to navigate highly restricted environments, as in the human body, in ways that would be impossible for rigid robots. Finally, growing robots, also

known as everting or tip-extending robots, are an example of a subclass of both soft robots as well as continuum robots. These robots consist of a thin-walled, hollow tube that is inverted inside itself, such that when pressurized, the material everts from the tip of the robot, effectively causing the robot to grow [16]. Like continuum robots, growing robots are able to traverse difficult paths by conforming around obstacles in their paths and like soft robots, growing robots are made of compliant materials, meaning they tend to impart smaller forces on their surroundings than rigid robots. This combination of properties makes them especially promising candidates for surgical procedures, which can involve working in conditions that are both highly constrained and fragile, for example in the vasculature of the human brain.

With all of these designs proposed to improve safety however, there comes a trade-off. In general, much of the theory, design methodology, and hardware developed for traditional robots cannot be applied to these new robots. As such, it is necessary to develop methods for fabrication, modelling, and control that accommodate the new physics of compliant robots. Furthermore, in order for these robots to be used in a medical setting, the newly developed methods must also adhere to the constraints posed by the surgical environment. Many researchers have already devoted significant effort to the study of continuum and soft robots for surgical applications, but efforts with regards to growing robots are still in their preliminary stages. This thesis focuses on solving some of the fundamental challenges in robotics for growing robots–specifically the incorporation of sensing into growing robots as well as localization of the robot's tip. The approach detailed below is wireless, low-cost, and suitable for use in a surgical procedure.

1.1 Soft Robots

Traditionally, robots have been fabricated with rigid materials connected at discrete joints because this allows for straightforward design, modelling, and control. Often, however, biological systems outperform rigid robots in tasks like manipulation and locomotion. Consequently, some researchers have begun designing robots with compliant structures, inspired by

nature, in the attempt to improve robotic performance in these challenging areas (see Fig. 1.1 for an example). Though there is not a canonical definition of "soft" in the context of robotics, Rus et. al capture the intent of the word by defining a soft robot to be one that is primarily composed of material with Young's moduli in the range of soft biological materials (10^4–10^9 Pa) [49].

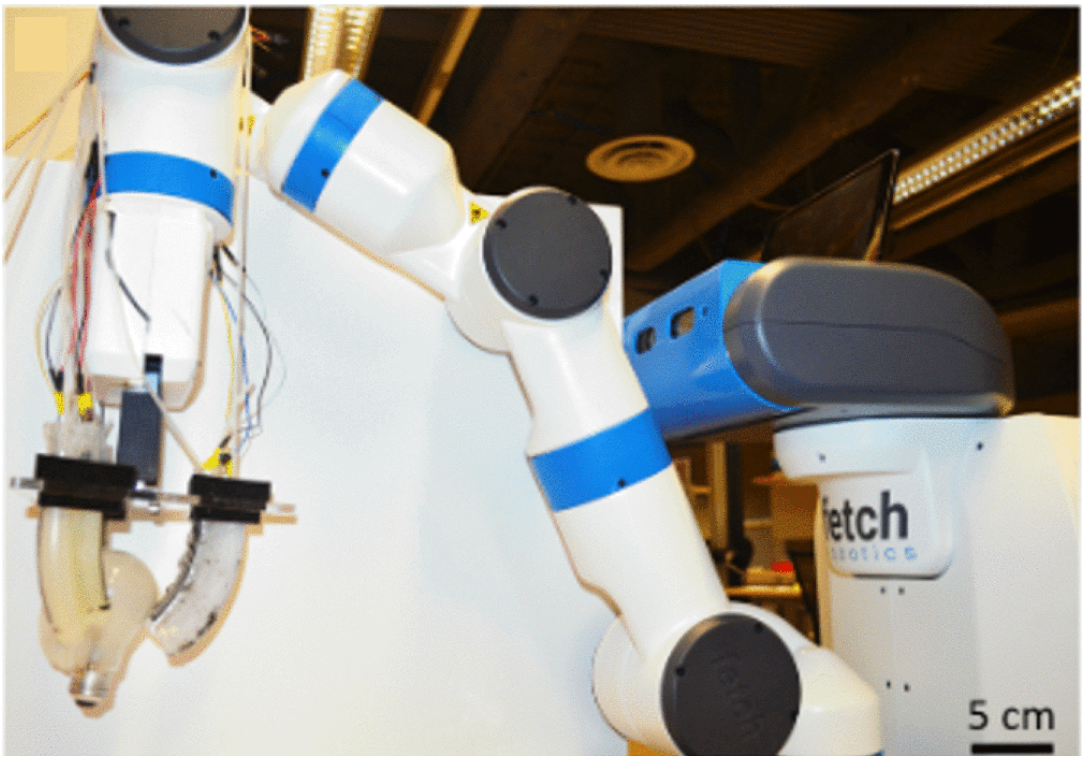

Figure 1.1. This figure, which originally appeared in Shih et al. [54], illustrates the inherent safety of soft robot designs–even if there are inaccuracies in the robot's estimation of the pose of the light bulb or its own end-effector, it is less likely to damage the light bulb because the low stiffness materials out of which its gripper is fabricated can compensate for an overly-aggressive grip.

Examples of soft robot designs explored to date span a range of applications. In [10], Deimel and Brock explored the use of pneumatic actuation to control a soft gripper, while Amend et al. studied granular jamming for a similar mechanism in [2]. Marchese et al., Onal and Rus, and Tolley et al. built robots mimicking the locomotion of fish, snakes, and quadrapeds in [28], [36], and [64], respectively. Soft wearables targeting medical applications have been explored in [38] and [32]. For a more comprehensive review, see [49]. While these applications

demonstrate some of the utility of soft robotics, there are remaining research challenges to address before unlocking their full potential. The soft materials that these robots leverage to enable their novel capabilities are much less studied as compared to rigid materials, meaning that there is an open area of research regarding fundamental problems in soft robotics that already have rigid-based solutions. Such problems include how to best design and place soft sensors and actuators, how to model forward and inverse kinematics and dynamics, and how to implement control laws that take full advantage of the robot's material properties.

1.2 Continuum Robots

The origins of continuum robots are also biologically inspired. The impressive ways in which animals like snakes and cephalopods perform locomotion and manipulation tasks has motivated researchers to build robots capable of similar abilities. The resulting robots have manipulators with a continuous curvature from their proximal to distal ends as compared to traditional robots which have a sharp discontinuity in curvature at each of their joints. Subsequently, this continuity tends to distribute the effects of actuation and environmental interaction along the entire length of the robot [45]. Accordingly, continuum robots are able to compensate for obstacles they contact in their surroundings by bending around them [67].

Early research efforts in this area attempted to directly mimic natural systems– like the underwater eel robot described by McIsaac and Ostrowski in [31] or the "Air-Octor" arm studied by McMahan et al. in [30]. But more recently, as minimally invasive laparoscopic procedures have come to dominate much of modern surgery, continuum robots have emerged as a promising solution for surgeons that need to be able to access hard to reach areas in the body through small openings. Examples of such medically focused continuum robots are the concentric tube robots, originally developed by Dupont et al. [12] and Webster et al. [66], as well as those with tendon-driven mechanisms like the one studied by Camarillo et al. (depicted in Fig. 1.3) in [5]. Some commercial robotic surgery platforms have even begun to incorporate

Figure 1.2. This figure, which originally appeared in Camarillo et al. [5], depicts a tendon driven continuum manipulator. These types of robots are unique in that they possess infinite degrees of freedom in the continuously varying curvature of their flexible backbones.

continuum elements into their designs as well (for example the da Vinci SP by Intuitive [43] and the Flex Robotic System by Medrobotics [60]).

While many groups have made significant progress with regards to the modelling ([21] [47] [65]) and control ([42] [67] [22]) for these robots, the practical application of their methods remains limited because the vast majority of this work does not account for uncertain and distributed forces on the robot from its environment nor does it address how to incorporate physical sensors into the robot body. Both the modelling and sensing questions only become more complex when the continuum robot also happens to be made of soft materials, because although the robot softness makes it safer, soft materials in general cannot be described by

simple elasticity models and they make it difficult to integrate sensors with non-similar stiffness. It is important for researchers to continue to try to address these challenges in order to increase use of these robots in the medical community.

1.3 Growing Robots

Growing robots represent an emerging subclass of both soft and continuum robots that has proved particularly useful in constrained and extremely fragile environments [24] [56]. Growing robots posses a length to diameter aspect ratio that is similar to continuum robots and are likewise flexible along their length–allowing them to naturally conform to obstacles in the environment [15]. Because the materials used to fabricate growing robots are themselves quite compliant, growing robots also exhibit the same small impact on their environment that soft robots do. Further, the nontraditional method by which growing robots achieve locomotion actually makes these robots even less taxing on their surroundings. Because the material of the growing robot does not move relative to its surroundings after it everts, the forces it transmits to its environment are small even compared to other soft robots.

These robots have the potential for use in a wide variety of applications. For example, Luong et al. developed a robot capable of underwater eversion for minimally invasive study of coral reefs [24]. Slade et al. detailed the design and testing of a soft, growing catheter for surgical applications [56]. And Sadeghi et al. developed a robot targeted for applications involving soil penetration [50]. These robots have even been used to deploy reconfigurable antennas [3] and to explore an archeological site [6]. Before the capabilities of these robots can be extended further however, the same research challenges that exist for soft and continuum robots in general must be addressed for growing robots. For example, closed-loop control, and therefore sensor information about the robot's state, is often necessary to overcome the effects of environment interactions, which can change the shape of the robot [41] [26].

Medical applications, in particular, represent scenarios where accurate models and state

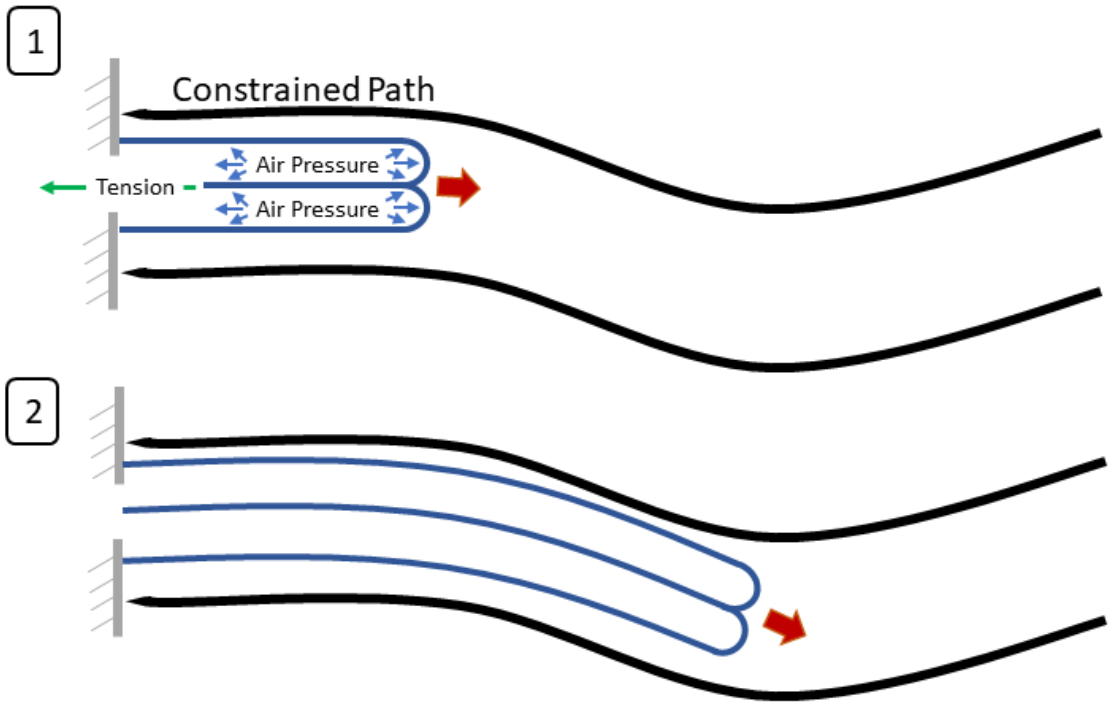

Figure 1.3. This figure depicts a growing robot at different stages of eversion. As the robot is pressurized, material extends from its tip, effectively causing the robot to grow. After the material is everted, it no longer moves relative to its surroundings, resulting in minimal transmission of shear forces to the environment during robot locomotion.

information are critical to ensure safety and efficacy, yet sensing proves even more challenging due to a number of operational constraints. In addition to the relatively small size requirements for many medical robots, maintaining a hollow access channel can be important for delivering surgical tooling or therapy to the target location. It is therefore important that any added sensors do not significantly occlude the center channel by either requiring wires to run through the length of the robot or by requiring an attachment method, such as a cap, that covers the robot tip [6]. This work presents a method for localizing the tip of these growing robots that is accurate, adheres to the constraints described, and is straightforward to implement for applications where the setup of external sensors is feasible.

1.4 Sensing for Continuum Robots

The problem of state estimation for continuum robots typically refers to localizing many points along the robot's backbone relative to an inertial frame—and for some well-studied robots, even estimating additional parameters (e.g. force applied to the end effector [48])—by combining information from different sensors with mathematical models [25]. The choice of sensor and model typically depends on the target application for the robot and the development of both is an active area of research. For scenarios in which there is a direct line of sight to the robot, there are various vision-based algorithms which achieve shape reconstruction of the robot from multiple cameras [9]. In surgical applications, when the robot is occluded from view, solutions involving ultrasound [44], fluoroscopic [18], electromagnetic [58], and fiber Bragg grating [37] technologies have been proposed (see [53] for a review).

For growing robots in particular, the problem of tip localization has yet to be addressed, let alone full shape reconstruction. The most common approaches to deploying these robots in practice require the user to either pre-form the robot into a known shape and deploy it in an open-loop manner [4] or to attach a camera to the tip of the robot which can then be used in combination with flexible actuators for eye-in-hand visual servoing [8]. Imaging the robot under fluoroscopy is feasible [18], but delivers high doses of radiation to patients and operators in the vicinity. Ultrasound methods typically suffer from signal-to-noise ratio problems [53], and while Fiber Bragg grating systems have shown some promising results for other medical robots, in practice, they are still prohibitively expensive for widespread adoption. For these reasons, we focus on magnetic tracking solutions to localization.

1.5 Magnetic Tracking

Localization based on magnetic field strength has attracted much research attention because it does not require line of sight to the target and because magnetic field strength does not attenuate through human tissue, making it particularly well-suited for medical applications

[20]. The approaches developed to date depend on two primary components—the magnetic source and the magnetic sensor.

In active field approaches, the sensor is collocated with the object to be tracked. An array of magnetic sources then generates a changing magnetic field which can be modeled or precomputed for a finite volume of interest. As the object, and therefore the sensor, move throughout the volume, an online numerical solver can run to find estimates for the position and orientation that minimize the difference between the sensor readings and the generated field [13]. In passive field approaches, the magnetic source—a permanent magnet—is attached to the target, and the magnetic field produced as it moves is measured by an array of sensors. As with the active approach, efficient numerical methods can be used to find position and orientation estimates for the target that minimize the difference between the sensor readings and the modeled field [63].

Both approaches rely on the accuracy of the magnetic sensors, as well as the accuracy of the model for the field generated by the magnetic source. In practice, active sensing approaches have been favored over passive ones because the magnetic field generated from the source can be dynamic, making it easier to sense in the presence of noise, and because multiple targets can be tracked simultaneously. Commercial products implementing this technology, such as the Aurora and trakSTAR (Northern Digital Inc.), are available, and have been used by the surgical robotics community [62] [52], but remain limited by cost and the need to be wired. Because passive field techniques can be implemented in a way that both saves valuable space and is low cost, we adopt this approach to localization. We aim to overcome the limitations of this approach by using low-noise, magneto-inductive sensors and by developing a model that is part analytical and part learned to achieve accurate localization.

1.6 Contributions

The main contributions of this thesis are as follows: (1) We present a method for localizing the tip of a growing robot which is suitable for medical settings. This approach is wireless, low-cost, does not require significant modification of the robot, and allows for localization without line of sight—even through human tissue. (2) We develop a new approach to permanent magnet localization that combines analytical and machine learning techniques and compare it to existing methods. We also characterize its performance at different length scales and across different speeds. This approach is valid beyond tip tracking for growing robots and is suitable for applications involving permanent magnet localization in general. (3) We demonstrate our ability to localize the tip of a growing robot in real-time, while deploying through multiple constrained environments.

Chapter 2

Methods

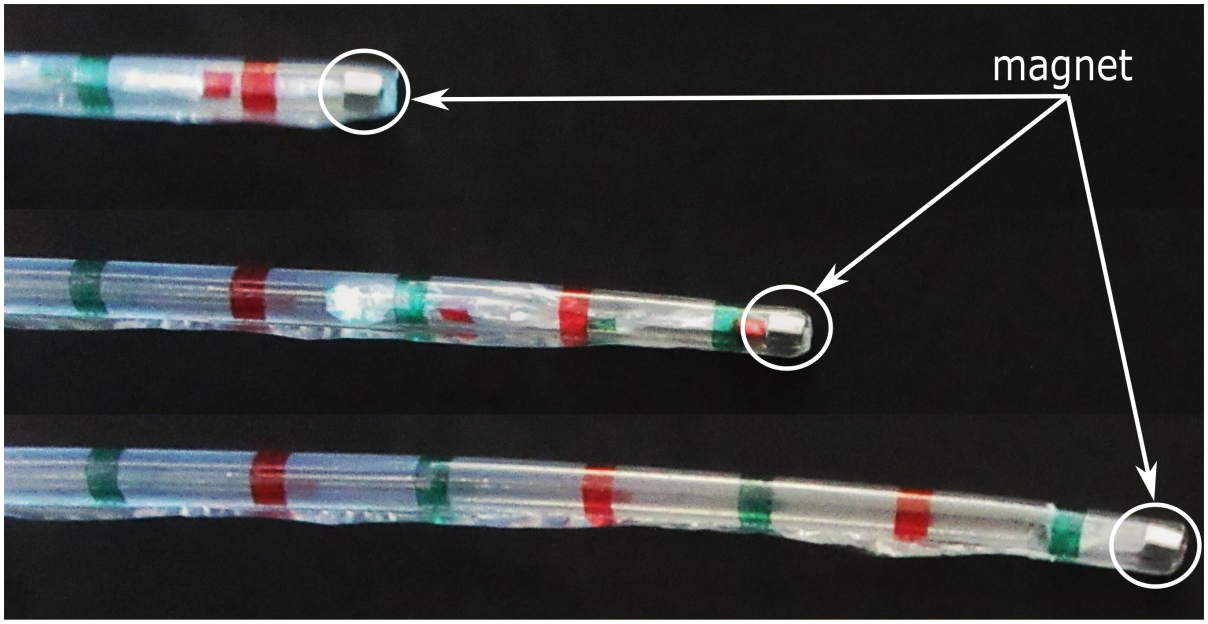

Figure 2.1. A ring magnet, placed at the distal tip of a translucent growing robot, remains at the tip even as the robot everts. This magnet can be used in conjunction with an array of magnetic sensors to localize the tip of the robot to 5 degrees of freedom.

The incorporation of sensing into a growing robot is a challenging problem. The added sensors must not significantly reduce the robots favorable properties of compliance or eversion-based locomotion. Additionally, the sensors must be robust enough to survive the inversion and re-eversion process many times over. Our approach leverages the robot's own unique growing ability to keep a small, rigid magnet at its tip, which we can then use for localization.

A permanent magnet, in the shape of a ring, is sized such that its outer diameter is

slightly smaller than that of the growing robot it will track. By inverting the robot material through the center of the ring magnet and pulling the magnet up to the tip, the friction between the inner layer of the robot and the inside of the ring magnet, together with the applied internal pressure, causes the magnet to stay at the tip of the robot, even as the robot grows (Fig. 2.1). The magnetic field emanating from the magnet can then be sensed with any of a variety of commercially available sensors (Hall effect, magneto-resistive, magneto-inductive, etc.). By combining these sensors in an array that covers the area of interest, it is possible to resolve the position and orientation of the permanent ring magnet, and therefore the tip of the growing robot, to five degrees of freedom—three spatial coordinates and two angular ones. Because the magnetic field from the ring magnet is symmetric about its axis of revolution, it is impossible to determine the third angle that describes the full 6 degree-of-freedom pose of the tip with only a single permanent magnet. We examine three different modeling approaches to permanent magnet localization (Fig. 2.3). The first is the traditional, analytical approach. The second uses a neural network to learn a mapping from sensor data to magnet pose. The third is a hybrid approach that uses both analytical and learning methods for localization.

2.1 Dipole Model Based Localization

In passive magnet tracking, the most common approach to modeling a permanent magnet is to assume that it behaves as a dipole located at the magnet's center [20] [61] [57]. For a magnetic dipole, the magnetic flux density $\mathbf{B}$ at the i^{th} sensor in the array is given by,

$$\begin{aligned} \mathbf{B}_i(\mathbf{r}_i, \hat{\mathbf{m}}) &= \frac{\mu_0 M}{4\pi \|\mathbf{r}_i\|^3} \left(\frac{3(\hat{\mathbf{m}} \cdot \mathbf{r}_i)}{\|\mathbf{r}_i\|^2} \mathbf{r}_i - \hat{\mathbf{m}} \right) \\ \mathbf{r}_i &= \mathbf{p}_i - \mathbf{x}, \end{aligned} \tag{2.1}$$

where $\hat{\mathbf{m}}$ is a unit vector aligned with the magnet's axis of magnetization, $\mathbf{x}$ is the location of the magnet in the global frame, $\mathbf{p}_i$ is the location of the i^{th} sensor in the global frame, and $\mathbf{r}_i$ is the vector from the center of the magnet to the i^{th} sensor (Fig. 2.2). The strength of

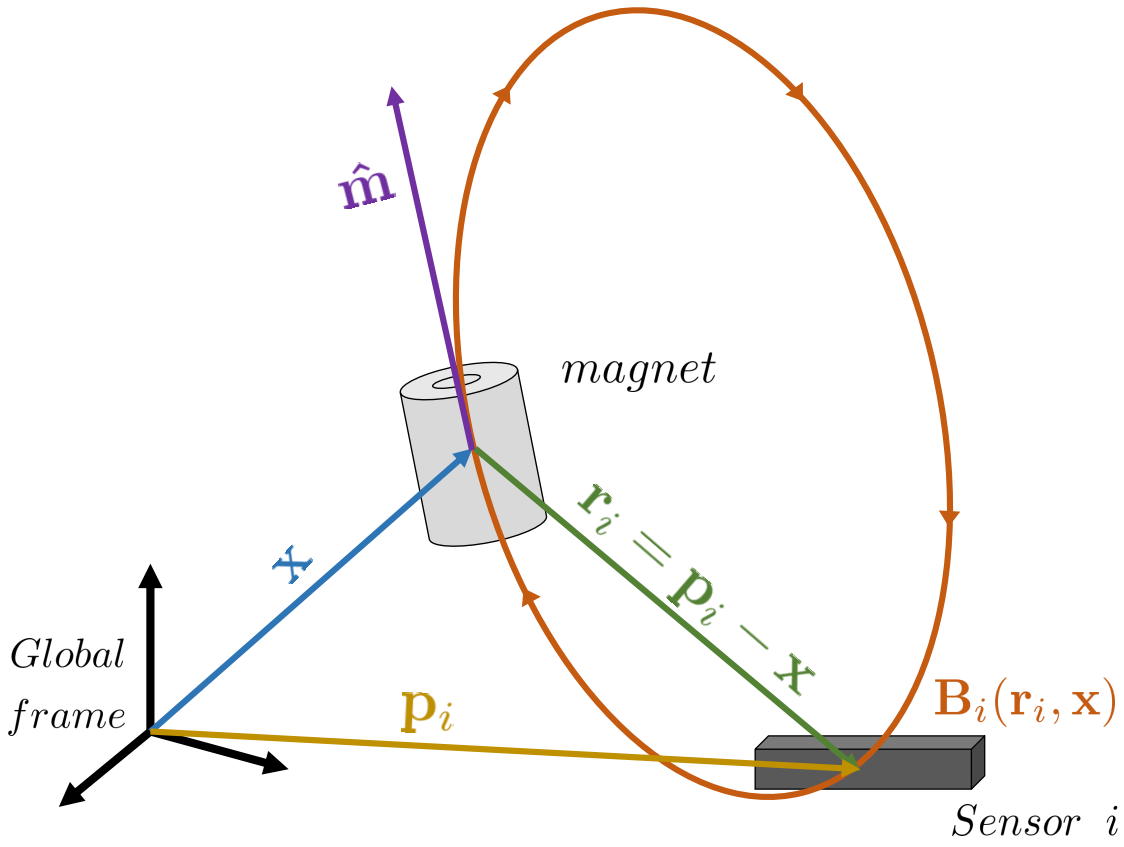

Figure 2.2. Dipole model for the magnetic flux density ($\mathbf{B}_i$) at the i^{th} sensor, where $\hat{\mathbf{m}}$ is the orientation of the magnet, $\mathbf{x}$ is the location of the magnet, $\mathbf{p}_i$ is the location of the sensor, and $\mathbf{r}_i$ is the vector from the magnet to the sensor.

the magnet, M, depends on the material and geometry of the magnet itself, and the constant $\mu_0 = 4\pi \times 10^{-7}$ T $\cdot$ m $\cdot$ A^{-1} is the magnetic permeability of free space [14].

The function relating these quantities to the flux density at the sensor $\mathbf{B}_i(\mathbf{r}_i, \hat{\mathbf{m}})$ is not solvable analytically. Further, it is clear from the dimensions of the equation, which maps five position and orientation magnet parameters to three sensor readings, that information from one sensor is not enough to properly localize the magnet because there is oftentimes more than one magnet pose that will produce the same reading for a given sensor. Instead of approximating an inverse to this function, an optimization problem is posed to minimize the sum of the squared 2-norm of the difference between the modeled and measured values of *all n* sensors over $\mathbf{x}$ and

$\hat{\mathbf{m}}$. This way, information from multiple sensors is combined to better resolve any ambiguities in the magnet pose. The cost function is,

$$\min_{\mathbf{x},\hat{\mathbf{m}}} \sum_{i=1}^{n} \|\mathbf{B}_i - \mathbf{B}_{meas}\|^2 . \tag{2.2}$$

The minimization of Eq. (2.2) can be done in real time with efficient implementations of established algorithms like that developed by Levenberg and Marquardt [23] [29]. The Levenberg Marquardt algorithm alternates between the Gauss-Newton and steepest descent methods as it iterates, making it more robust than Gauss-Newton alone. It can be warm-started at each time step by using the converged solution from the previous time step as the initial guess to significantly decrease the number of iterations until convergence for the current time step. Additionally, because the cost function (Eq. (2.2)) is differentiable, its analytical Jacobian can be supplied to these solvers to improve speed and performance. We use MATLAB's (The MathWorks Inc.) nonlinear least squares solver implementing the Levenberg Marquadt algorithm for the optimization.

2.2 LSTM Based Localization

The classical approach to permanent magnet localization is limited in a number of ways. The dipole model assumed for the permanent magnet, for instance, is only a valid approximation sufficiently far from the magnet, and the algorithms used in the online optimization step are only guaranteed to converge to a local minimum [35]. Additionally, magnetic sensors typically display time dependent nonlinearities such as hysteresis [11]. Each of these phenomena is either difficult to account for analytically or challenging to implement in a way that is computationally tractable for real-time estimation.

Neural networks, however, have been shown to be adept at capturing such nonlinearities with minimal *a priori* knowledge of the underlying system. Given a sufficiently large training set of input-output data tuples and a sufficient number of parameters, a neural network can

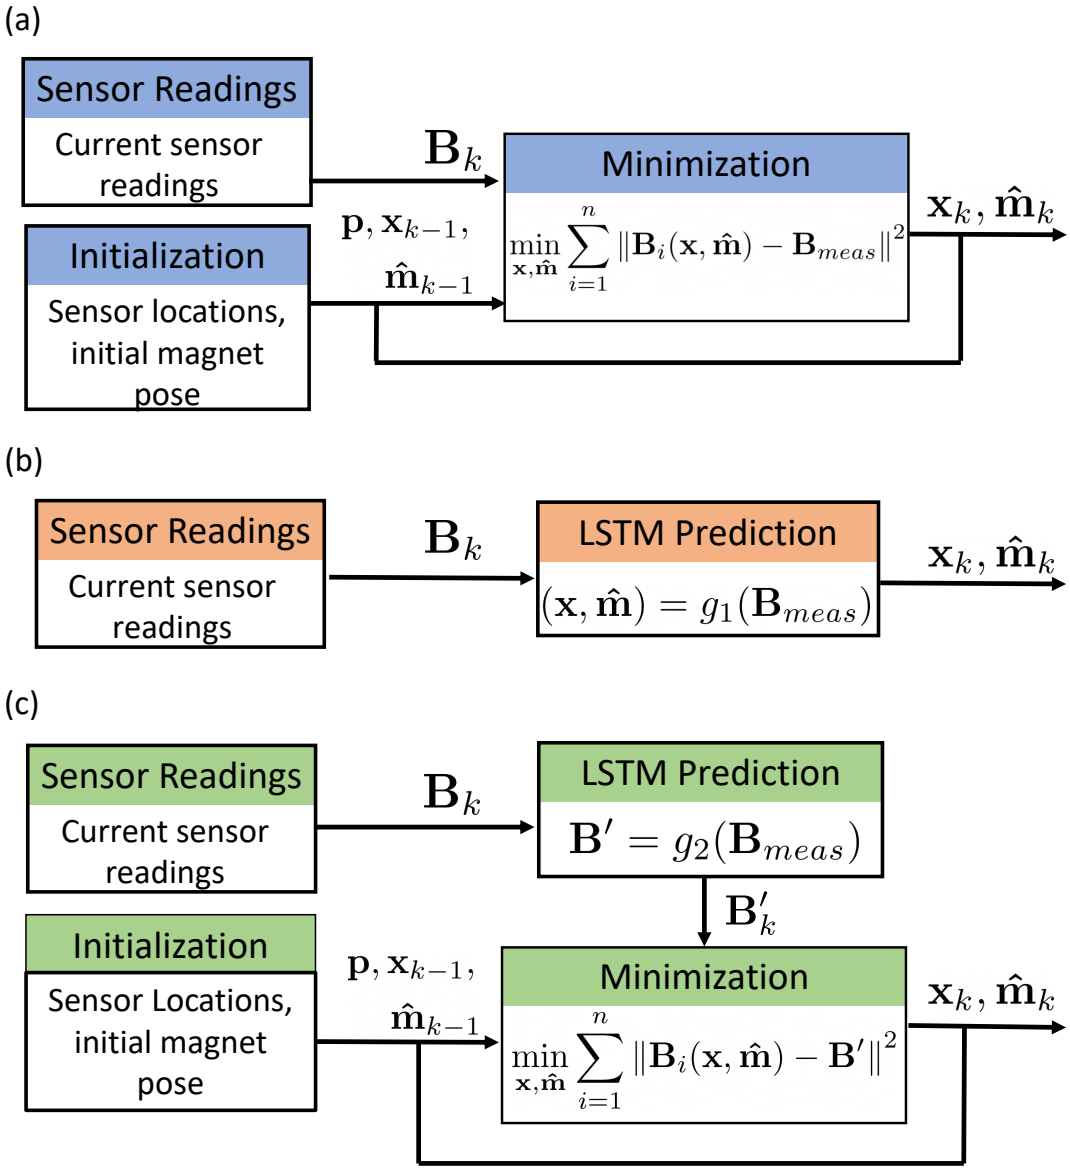

Figure 2.3. Flowchart of three permanent magnet localization schemes, where $\mathbf{B}_k$ is the magnetic flux density read by all of the sensors in the array at the current time step, $[\mathbf{x},\hat{\mathbf{m}}]$ are the magnet position and orientation, and $\mathbf{p}$ is a vector containing the location of the sensors in the sensor array. The localization methods include (a) the traditional approach, which relies on minimizing a difference between modeled and sensed values, (b) a neural network to map directly from sensor data to a localization estimate, and (c) a neural network to map the readings from each sensor to a modeled value, which is then used in an optimization step similar to (a).

approximate any smooth function [19]. This useful property, together with improved computing methods and hardware for training them, have made neural networks useful for prediction in a variety of tasks in recent years. There has even been some preliminary studies applying neural nets to permanent magnet localization [51].

Long-short term memory (LSTM) networks in particular, are a class of recurrent neural network that have seen widespread adoption for predicting on time series data because of their ability to capture trends in data with time lags [17]. We use the MATLAB Deep Learning Toolbox to build a network to learn a function mapping (g_1) directly from input sensor data ($\mathbf{B}_{meas}$) to position and orientation of the magnet ($[\mathbf{x}, \hat{\mathbf{m}}]$) (Fig. 2.3(b)). The network consists of an LSTM layer with size 100, followed by a dropout layer with dropout rate 0.25 to improve performance with noisy training data, and finally a fully connected layer. The network parameters were optimized with the Adam algorithm and miniBatch size 4, and L2 regularization was used to prevent overfitting.

2.3 A Hybrid Approach to Localization

One of the main downsides of using a purely data-driven approach to modeling is the removal of physical significance from the problem, making it difficult for designers to draw insights into the process they are trying to observe. Additionally, when partial information about a system is known, there is not a straightforward way to incorporate this knowledge into a data-driven model. In the passive magnetic tracking problem, even with a perfect network that could exactly map sensor readings to magnet pose, this network would be unable to make useful predictions if the placement of the sensors is changed. Although this change would be straightforward to compensate for using the dipole model described above, the neural network would need to be retrained for each new configuration of sensors.

To incorporate a learned component into an analytical model, a network can be trained to learn the residual between an analytically modeled output and empirical data [34]. Using the

network and model together can lead to more accurate predictions. Unfortunately, this approach would still be tied to a specific sensor configuration. To solve this issue, we train a network to learn to compensate for the residual of each sensor *individually*, by training a network to learn a function (g_2) mapping the readings of an individual sensor, $\mathbf{B}_{i,meas}$, to the values predicted by our dipole model for that sensor, $\mathbf{B}'_i$ (Fig. 2.3(c)). By restricting our sensor array to contain only one type of sensor and assuming that they all exhibit the same type of nonlinear behavior, we can use this network to augment our analytical model in a way that is agnostic to the sensor configuration. We adopt the same network architecture as in the purely data-driven approach to localization above, in order to capture the time dependent nonlinearities in the sensor behavior.

Chapter 3

Experiments and Results

This chapter details the design and results of experiments meant to validate the proposed approach to growing robot tip localization. Specifically, the experiments considered evaluate the accuracy and data efficiency of each of the permanent magnet localization algorithms under different conditions, the area over which differently sized permanent magnets can be effectively located, and the accuracy of the localization scheme for tracking the tip of a growing robot deployed in multiple constrained environments.

3.1 Experimental Setup

A non-ferromagnetic test platform of size $L \times L$ (where L = 61 cm) was constructed out of medium density wood fiberboard (MDF) to evaluate the performance of our localization method in a 2D plane while mitigating the effects of background magnetic fields. As shown in Fig. 3.1, the platform was positioned a fixed height (h = 7 cm) above a 2D sensor array consisting of 4 magneto-inductive sensors (RM3100 Breakout Board, PNI), which measure magnetic flux density along 3 perpendicular axes in the range -800 to 800 μT. The sensors

Table 3.1. Physical properties of the different magnets used in the experiments

Magnet	OD (in)	ID (in)	L (in)	Magnet Grade
A	0.5	0.25	0.5	N48
B	0.25	0.125	0.125	N52
C	0.125	0.0625	0.125	N52

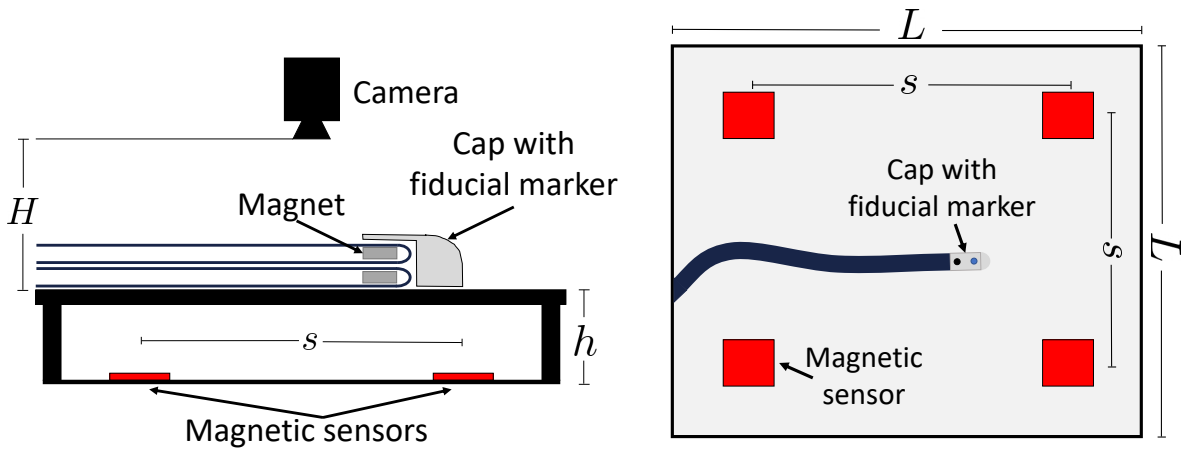

(a) Side view of test setup

(b) Overhead view of test setup

Figure 3.1. The test setup is constructed out of non-ferromagnetic material and positioned above the magnetic sensor array. A cap with a fiducial marker is placed on the tip of the growing robot and tracked by the overhead camera.

communicate over I2C with an Arduino Uno which in turn interfaces with MATLAB via serial USB. During data collection, we measure and subtract out the ambient magnetic field from the sensor readings before the permanent magnet is brought onto the test platform. Three different ring magnets from K&J Magnetics were used in the experiments, with physical properties shown in Table 3.1.

We evaluate the results of our localization against position and orientation determined from tracking fiducial markers using an overhead C920 camera from Logitech with a 30 Hz framerate and 1920 $\times$ 1080 pixel resolution. The camera is positioned a height H = 87 cm above the test platform. We also use the position and orientation data from the camera to train the neural networks, as explained in more detail below. The camera is calibrated using MATLAB's Image Processing Toolbox and the OpenCV Python library is used for tracking. Based on the mean reprojection error of the camera after calibration, the average error is 0.096 mm.

The analysis needed to determine the optimal number of sensors and their placement is left for future work. Here we place the sensors in a symmetric configuration in order to maximize the area over which the magnet can be sensed by at least two sensors simultaneously. The distance between the sensors is calculated based on the effective sensing radius of each

sensor. We take this effective radius to be the distance at which a sensor can sense a given magnet in its worst-case orientation according to the dipole model (i.e. $\hat{\mathbf{m}} \cdot \mathbf{r}_i = 0$) above some small value $\|\mathbf{B}\|_{min}$, where

$$\|\mathbf{B}\|_{min} = \frac{\mu_0 M}{4\pi \|\mathbf{r}_i\|^3}. \tag{3.1}$$

We solve Eq. (3.1) for $\|\mathbf{r}_i\|$ and then choose the distance between the sensors s to be equal to the sensor's effective sensing radius in the test platform located a height h above the sensor array, as stated in Eq. (3.2) and Eq. (3.3).

$$\|\mathbf{r}_i\| = \left(\frac{\mu_0 M}{4\pi \|\mathbf{B}\|_{min}}\right)^{1/3} \tag{3.2}$$

$$s = \sqrt{\|\mathbf{r}_i\|^2 - h^2}. \tag{3.3}$$

Using this heuristic, we can determine a sensor configuration for a given magnet with a dipole moment of magnitude M and a minimum signal value $\|\mathbf{B}\|_{min}$. We take $\|\mathbf{B}\|_{min}$ to be 3 μT, so that it is greater than the background noise read by the sensors.

3.2 Model Comparison

As stated above, we restrict the evaluation of our models to the 2D case by assuming that the component of the magnet's position out of the test plane is zero. We evaluate the models based on the error between their localization estimates and the actual positions and orientations of different validation datasets. Specifically, we compare the two position coordinates and one orientation value from our overhead camera with the corresponding position and orientation estimates from our models at every timestep. We compute the mean absolute position error as,

$$err_{position} = \frac{1}{N}\sum_{k=1}^{N} \|(x_1(k), x_2(k))_{actual} - (x_1(k), x_2(k))_{meas}\|, \tag{3.4}$$

where N is the number of samples in the dataset and x_1, x_2 are the spatial coordinates of the magnet position in the plane of the test platform, and the mean orientation error as,

$$err_{orientation} = \frac{1}{N}\sum_{k=1}^{N} |\theta(k)_{actual} - \theta(k)_{meas}|, \tag{3.5}$$

where θ is the orientation of the magnet in the plane of the test platform.

3.2.1 The Effect of the Quantity of Data on Model Performance

The error in the localization estimates of the learned models depends on the amount of training data available, while the localization error of the analytical model does not. In order to compare these approaches in a way that accounts for different amounts of training data, we first collect a large dataset, consisting of sensor readings and camera-determined magnet poses, by taking Magnet A (Table 3.1) and moving it randomly around the workspace. In total, we collected 40 $\times$ 1-minute videos of data, 20% of which is reserved for evaluating the performance of the models. The neural networks were trained with different amounts of data, drawn at random from the available training data. The performance of all three models was then compared on the validation set in terms of mean absolute position error.

Because the training data is generated randomly, it is not obvious which portions of it will have the most impact on the final learned model. To mitigate the importance of any one part of the overall training set, particularly when training with only a few minutes of data, we repeat this process of randomly selecting data, training the networks, and then evaluating their performance, a total of 5 times. The results in Fig. 3.2 show the mean and standard deviation of the error using each method. The mean position error of the analytical model on the validation set is plotted for reference. It is important to note that although the training data was drawn at random, the 1-minute sequences (referred to as "training sets" in the figure) were kept intact, which is important in order for LSTM networks to properly learn the time dependencies in the data.

The results show that the position error of both of the learned models eventually begins to level off as the amount of data used to train them is increased. Even when trained with the entire training set, the LSTM-based model still has significantly higher error than the analytical and hybrid models, with a mean position error of 28.5 mm, compared to 7.3 mm and 6.6 mm, respectively. The difference in the performance of the LSTM model can be seen as a tradeoff of this approach, which assumes no prior knowledge of the underlying system. The mean position error of the hybrid model matches that of the analytical baseline with just 4 minutes of training data, and begins to outperform the baseline when trained with 8 or more minutes of data. Because both the analytical and hybrid methods achieve substantially lower average position

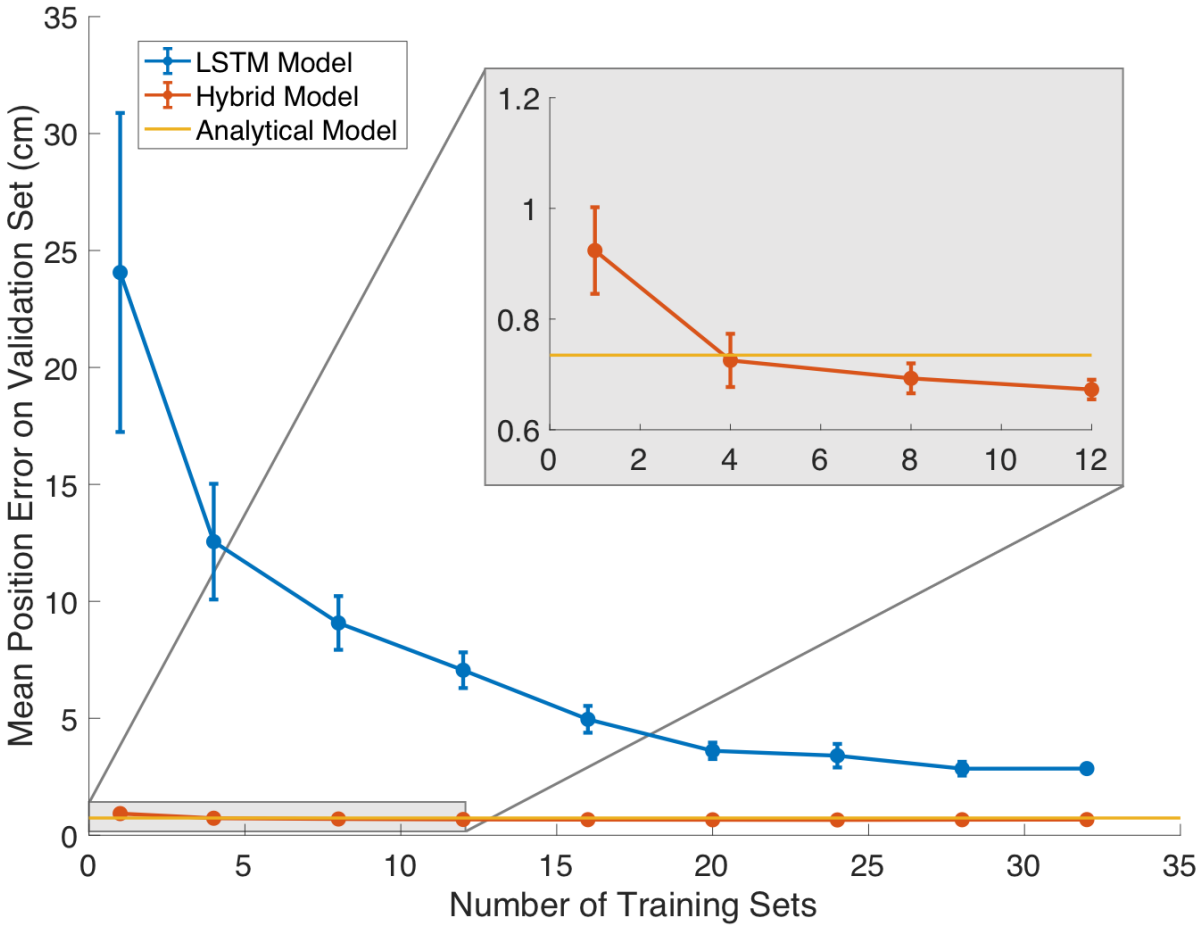

Figure 3.2. Mean position error for each of the three models, as the number of "Training Sets" (1-minute sequences of video data) is increased for the LSTM and Hybrid models. The error bars show the standard deviation of the mean position error of the models evaluated on the validation set after being trained with different amounts of data.

error on the validation set compared to the LSTM model, we disregard the LSTM model for the remainder of our comparison.

3.2.2 The Effect of Velocity on Model Performance

Next, we look to further examine the difference between the hybrid and analytical models. We expect that when there are large changes in the magnetic field, causing high levels of hysteresis in the magnetic sensors, the hybrid model should be better able to compensate. We therefore assess the performance of each of these two models by moving Magnet A along a fixed path that passes near the sensors at different velocities. The faster the magnet traverses this path, the larger signal change it induces in the sensors. We measure speed by a finite difference derivative approximation of the position data captured by the camera followed by a low pass filter to reduce the effect of noise. We then take the mean value of the speed to be representative of the speed at which the magnet traversed the path.

The results in Fig. 3.3 show the mean and standard deviation of the position error for five different mean velocities. The hybrid model achieves lower position error than the analytical model on every trial, and this difference in performance increases for the faster trials. At the fastest, when the magnet is moving an average of 25.5 cm/s around the path, the position error of the analytical model increases to 10.4 ± 2.3 mm, while the position error of the hybrid model is 6.6 ± 1.6 mm. This result indicates that the hybrid model is better at tracking a faster moving magnet. This difference in performance is likely due to the hybrid model's ability to better compensate for phenomena that are not modeled analytically, but become important for large changes in sensor readings—such as when the magnet quickly passes a sensor. Because the hybrid model achieves lower error than the analytical model on both the validation set from the previous test and in tracking the magnet at different speeds, this method is used for all subsequent tests.

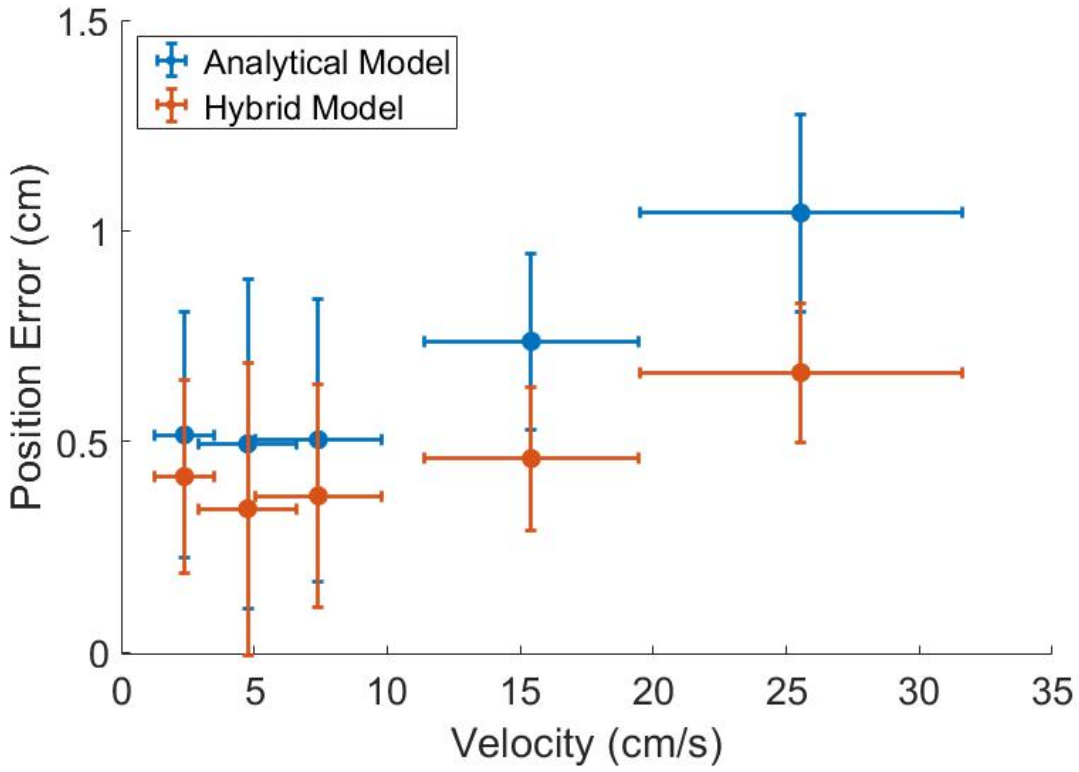

Figure 3.3. Comparison of the hybrid and analytical approaches to permanent magnet localization for a magnet travelling along the same path at different speeds. As the speed of the magnet increases the position error of the analytical model increases significantly more than the hybrid model, resulting in 10.4±2.3 mm and 6.6±1.6 mm at 25.5 cm/s, respectively.

3.3 Localization at Different Scales

Growing robots can easily be fabricated to meet a wide range of size requirements depending on the specific application (e.g. [56], [3], [6]). Our method of localization is able to accommodate these changes in dimensions by varying the size of the magnet used to track the tip of the robot. To assess how magnet size changes the effective workspace over which we are able to accurately localize the robot tip, we test magnets of three different sizes (see Table 3.1).

To maximize the effective workspace for each magnet, we change the distance between the magnetic sensors (s) according to our heuristic above (see Eq. (3.2), Eq. (3.3)). Using this method for sensor placement gives a distance between the sensors of 30.2 cm for Magnet

A, 15.1 cm for Magnet B, and 6.9 cm for Magnet C. We place each magnet at 36 different points spaced evenly across a 50 cm $\times$ 50 cm grid in the center of the test platform, holding the orientation constant, and we compare the result of the magnetic localization to the actual position and orientation. We repeat this for 8 different magnet orientations and compute the average position error at each test point. We linearly interpolate between the average error at each point to visualize the effective workspace for each magnet (Fig. 3.4). We consider any position estimates with an error greater than 15 mm to be points where the magnetic localization is unreliable and plot the corresponding area in black.

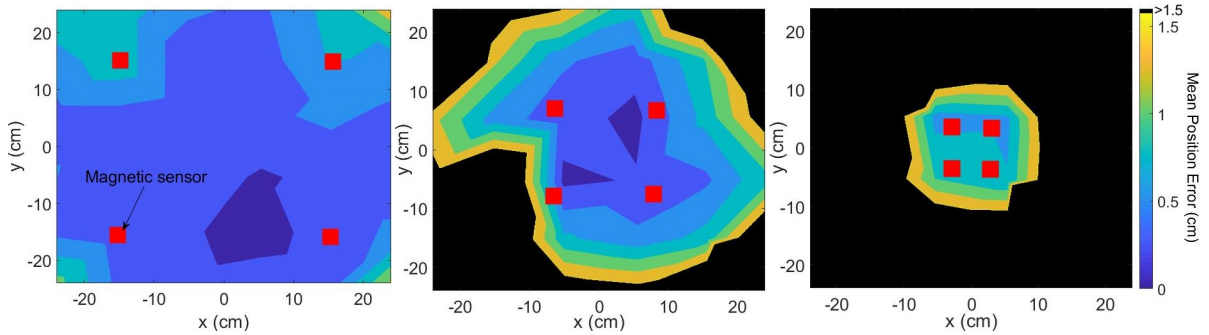

(a) Magnet A **(b)** Magnet B **(c)** Magnet C

Figure 3.4. Magnets of different sizes are positioned at 36 gridpoints spaced across a 50 cm $\times$ 50 cm area in 8 orientations at each point. The average position error of the magnetic localization estimate at each point is plotted to visualize the effective workspace for the magnet used. Points where the magnetic position estimate has $>$15 mm error are considered unreliable and plotted in black.

As shown in Fig. 3.4, for the area circumscribed by the sensors, the magnetic localization achieves similar accuracy across magnet size. The effective area of the workspace however, where position error is $<$15 mm, shrinks with magnet size. This decrease is due to the fact that the magnitude of the signal read by the sensors is proportional to the strength of the magnet, M, and inversely proportional to the cube of the distance between the sensor and the magnet, $\|\mathbf{r}_i\|^3$ (see Eq. (2.1)). As magnet size decreases, so does M, resulting in a weaker signal strength. Additionally, for smaller magnets, we position the sensors closer together according to Eq. (3.2) and Eq. (3.3) in order to achieve good tracking in a region close to all the sensors. While our method of localization successfully tracks magnets of different sizes, it is clear that for smaller

magnets, more sensors are necessary in order to achieve the same size of effective workspace. Further study is needed to determine the optimal number and placement of sensors for a given size magnet and target workspace.

3.4 Tip Localization of a Growing Robot Deployed in Constrained Environments

Finally, we evaluate the performance of our hybrid magnetic localization method on a robot as it grows through two constrained environments. For each deployment, the growing robot is positioned at the entrance to a highly curved path—one path resembling an "S" shape, and the other a spiral—which is centered over the array of magnetic sensors. A 3-D printed cap with fiducial markers is attached to the distal end of the robot (Fig. 3.1) to obtain the actual position and orientation of the robot tip for comparison. The robot is pressurized with air, causing it to grow until it reaches the target marked with a blue "X", as seen in Fig. 3.5. The objective of the experiments is to validate that the permanent magnet based approach to localization successfully estimates the position and orientation of the growing robot in the plane, even as the robot grows and interacts with the environment.

The robot in both scenarios has an outer diameter of 12 mm and an overall length of 645 mm. Magnet B is used for tracking the tip, and the same sensor spacing ($s = 15.1$ cm) from the previous experiment is used. Images from four different time steps for each experiment can be seen in Fig. 3.5, and the results of the measured and actual position can be seen in Fig. 3.6. For the "S" path, the magnetic localization achieved an average position error of 3.7 ± 1.2 mm and an average orientation error of $5.5 \pm 5.5^\circ$. For the spiral path, the average position error was 2.2 ± 0.9 mm and the average orientation error was $7.5 \pm 5.1^\circ$.

These experiments demonstrate that even as the robot interacts with the environment, the magnet reliably stays at the tip, and, using this method of localization, we can accurately track the position and orientation of the tip in the plane.

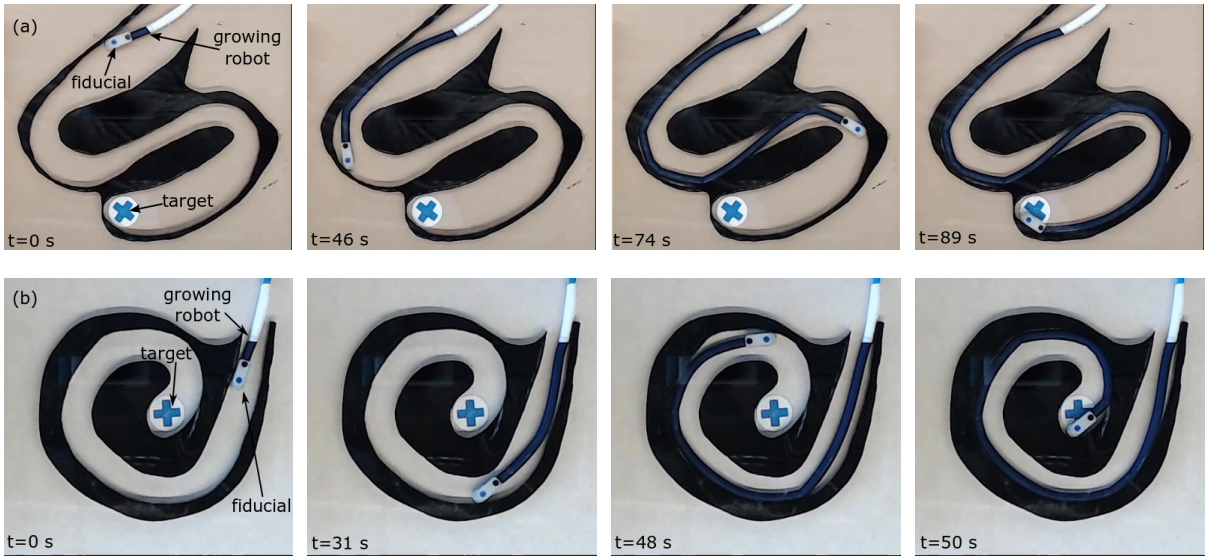

Figure 3.5. Time-series images of a growing robot with a ring magnet at its tip as it extends through constrained paths resembling (a) an "S" and (b) a spiral. In each scenario, the robot must conform to the curvature of the path in order to reach the target at the end, marked with a blue X.

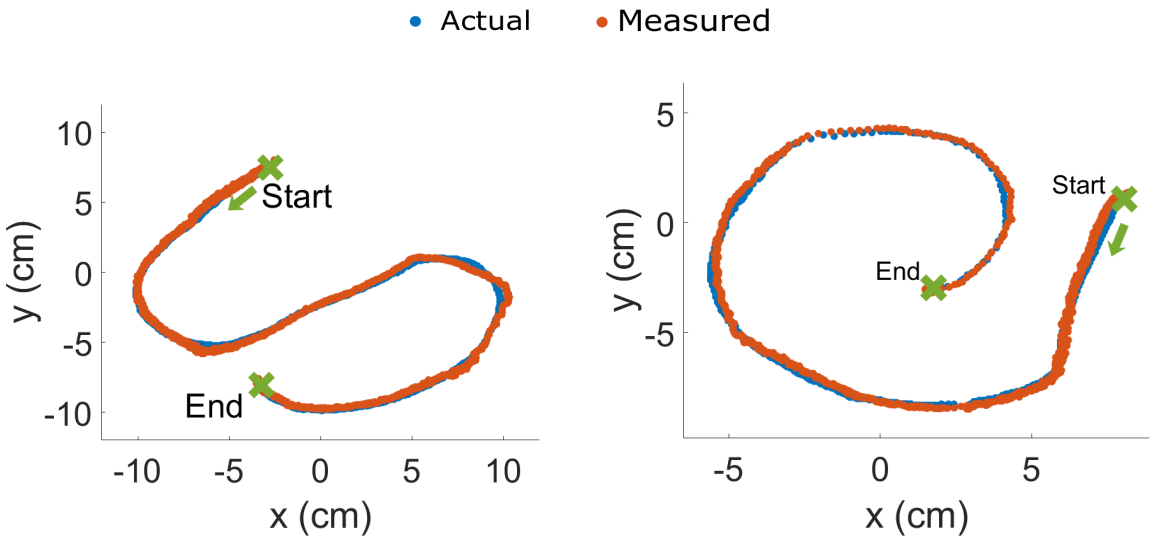

(a) "S" path

(b) Spiral path

Figure 3.6. Hybrid model for permanent magnet based localization tracks the tip of the growing robot as it navigates (a) an "S" shaped path, with average position error 3.7±1.2 mm and average orientation error 5.5±5.5° and (b) a spiral shaped path, with average position error 2.2±0.9 mm and average orientation error 7.5±5.1°.

Chapter 4

Conclusions and Future Work

The future of medical robotics promises to continue to improve existing surgical procedures and to enable new methods of treatment entirely, provided that the robots are guaranteed to operate in a safe manner. Growing robot designs in particular show great potential ability to address these safety concerns, but require researchers to develop new methods for their design, fabrication, and control. This thesis focused on a solution to the problem of tip localization for growing robots that is suitable for use in a surgical setting. We have demonstrated effective tracking of the tip of a growing robot using a permanent ring magnet and an array of magneto-inductive sensors. We developed a modeling approach that leverages both analytical and machine learning techniques to achieve high-accuracy localization for a permanent magnet that is agnostic to the configuration of the magnetic sensors. We validated this model by comparing its performance against existing approaches to permanent magnet localization and tested its ability to track magnets of different sizes, ensuring its versatility. Finally, we deployed a growing robot in two different, constrained environments, in order to validate that our method is capable of reliably resolving the position and orientation of the robot's tip.

It is important to note some of the limitations inherent with the proposed approach. Permanent magnet-based localization is sensitive to noise from other magnetic fields, which can come from nearby electronic devices that emit electromagnetic fields or induced in ferromagnetic materials by the permanent magnet itself. Additionally, it is only possible to track the position

and orientation of one permanent magnet at a time, meaning this method could not be extended to do full shape reconstruction for growing robots. Finally, because permanent magnet localization requires multiple magnetic sensors to be near the magnet at all times, this method for localization is not suitable for use in large unknown environments, such as an archaeological site, because sensors need to be placed in precisely known locations at the start. Despite these limitations, there are clear avenues for the future of this work that can expand its usefulness and usability in real-world settings.

For example, this approach can be extended to a 3D environment. Though the analysis in this work was confined to the 2D case, the 3D extension is very straightforward from an algorithmic perspective–it simply requires that the full 3D localization obtained from the online optimization step be used instead of projecting it onto a 2D plane. Validation of this approach would require a very accurate, ground-truth system that could be used to compare against the permanent magnet-based tracker. Such a system might consist of an array of cameras together with an appropriate set of image processing algorithms to determine the pose of the tip of the growing robot. Likely, the primary challenge with this approach would be developing these image processing algorithms such that they are not overly sensitive to the changes at the tip of the growing robot–which does not maintain a uniform shape as it everts.

Additionally, sensor placement can be fully investigated to ensure optimal tracking of the robot tip for specific tasks. In a medical setting, combining sensor placement algorithms with preoperative imaging and the proposed localization method could ensure that the growing robot is tracked accurately for the entirety of an operation. A hypothetical workflow to investigate this approach might begin by running a growing robot path planner for a given patient image and target to reach. By framing the problem of localizing the robot over a finite time window as tracking a point moving along this planned path with a finite velocity, we can optimize sensor placement to maximize some notion of information gain over the time window [1] [59].

Lastly, by combining localization with actuators that can change the orientation of the

robot's tip, we can develop closed-loop control techniques to enable accurate and autonomous positioning for this class of robot, similar to those that exist for other continuum manipulators. Of particular interest in this line of work would be to develop hybrid methods for control that parallel the techniques in this thesis used for localization. These proposed methods would combine model-based control techniques for the robot with those that can learn from data to compensate for model inaccuracies and environmental disturbances in an online and data-efficient way.

ACKNOWLEDGEMENTS

This thesis, in part, has been submitted for publication of the material as it may appear in IEEE Robotics and Automation Letters, 2020, Watson, Connor; Morimoto, Tania K., IEEE, 2020. The dissertation author was the primary investigator and author of this paper.

Bibliography

[1] Optimal sensor placement and motion coordination for target tracking. *Automatica*, 42(4):661 – 668, 2006.

[2] J. R. Amend, E. Brown, N. Rodenberg, H. M. Jaeger, and H. Lipson. A positive pressure universal gripper based on the jamming of granular material. *IEEE Transactions on Robotics*, 28(2):341–350, April 2012.

[3] L. H. Blumenschein, L. T. Gan, J. A. Fan, A. M. Okamura, and E. W. Hawkes. A tip-extending soft robot enables reconfigurable and deployable antennas. *IEEE Robotics and Automation Letters*, 3(2):949–956, April 2018.

[4] L. H. Blumenschein, N. S. Usevitch, B. H. Do, E. W. Hawkes, and A. M. Okamura. Helical actuation on a soft inflated robot body. In *2018 IEEE International Conference on Soft Robotics (RoboSoft)*, pages 245–252, April 2018.

[5] D. B. Camarillo, C. F. Milne, C. R. Carlson, M. R. Zinn, and J. K. Salisbury. Mechanics modeling of tendon-driven continuum manipulators. *IEEE Transactions on Robotics*, 24(6):1262–1273, Dec 2008.

[6] Margaret M. Coad, Laura H. Blumenschein, Sadie Cutler, Javier A. Reyna Zepeda, Nicholas D. Naclerio, Haitham El-Hussieny, Usman Mehmood, Jee-Hwan Ryu, Elliot W. Hawkes, and Allison M. Okamura. Vine robots: Design, teleoperation, and deployment for navigation and exploration. *CoRR*, abs/1903.00069, 2019.

[7] Nikolaus Correll, Kostas E. Bekris, Dmitry Berenson, Oliver Brock, Albert Causo, Kris Hauser, Kei Okada, Alberto Rodriguez, Joseph M. Romano, and Peter R. Wurman. Lessons from the amazon picking challenge. *CoRR*, abs/1601.05484, 2016.

[8] Joseph D. Greer, Tania K. Morimoto, Allison M. Okamura, and Elliot Hawkes. A soft, steerable continuum robot that grows via tip extension. *Soft Robotics*, 6, 10 2018.

[9] M. M. Dalvand, S. Nahavandi, and R. D. Howe. High speed vision-based 3d reconstruction of continuum robots. In *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 000618–000623, Oct 2016.

[10] Raphael Deimel and Oliver Brock. A novel type of compliant and underactuated robotic hand for dexterous grasping. *The International Journal of Robotics Research*, 35(1-3):161–185, 2016.

[11] Dora Domajnko and Dejan Krizaj. Lagging-domain model for compensation of hysteresis of xmr sensors in positioning applications. *Sensors*, 18:2281, 07 2018.

[12] P. E. Dupont, J. Lock, B. Itkowitz, and E. Butler. Design and control of concentric-tube robots. *IEEE Transactions on Robotics*, 26(2):209–225, April 2010.

[13] Alfred Franz, Tamas Haidegger, Wolfgang Birkfellner, Kevin Cleary, Terry Peters, and Lena Maier-Hein. Electromagnetic tracking in medicine-a review of technology, validation, and applications. *IEEE Transactions on Medical Imaging*, 05 2014.

[14] Edward P Furlani. Permanent magnet and electromechanical devices: Materials, analysis, and applications. 2001.

[15] J. D. Greer, L. H. Blumenschein, A. M. Okamura, and E. W. Hawkes. Obstacle-aided navigation of a soft growing robot. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018.

[16] Elliot W. Hawkes, Laura H. Blumenschein, Joseph D. Greer, and Allison M. Okamura. A soft robot that navigates its environment through growth. *Science Robotics*, 2(8), 2017.

[17] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9:1735–80, 12 1997.

[18] M. Hoffmann, A. Brost, M. Koch, F. Bourier, A. Maier, K. Kurzidim, N. Strobel, and J. Hornegger. Electrophysiology catheter detection and reconstruction from two views in fluoroscopic images. *IEEE Transactions on Medical Imaging*, 35(2):567–579, Feb 2016.

[19] Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural Networks*, 4(2):251 – 257, 1991.

[20] C. Hu, M. Li, S. Song, W. Yang, R. Zhang, and M. Q. . Meng. A cubic 3-axis magnetic sensor array for wirelessly tracking magnet position and orientation. *IEEE Sensors Journal*, 10(5):903–913, May 2010.

[21] J. Jung, R. S. Penning, N. J. Ferrier, and M. R. Zinn. A modeling approach for continuum robotic manipulators: Effects of nonlinear internal device friction. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5139–5146, Sep. 2011.

[22] Kai Xu and N. Simaan. Actuation compensation for flexible surgical snake-like robots with redundant remote actuation. In *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006.*, pages 4148–4154, May 2006.

[23] KENNETH Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168, 1944.

[24] J. Luong, P. Glick, A. Ong, M. S. deVries, S. Sandin, E. W. Hawkes, and M. T. Tolley. Eversion and retraction of a soft robot towards the exploration of coral reefs. In *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)*, pages 801–807, April 2019.

[25] Arthur Mahoney, Trevor Bruns, Ron Alterovitz, and Robert III. *Design, Sensing, and Planning: Fundamentally Coupled Problems for Continuum Robots*, pages 267–282. Jan 2018.

[26] M. Mahvash and P. E. Dupont. Stiffness control of surgical continuum manipulators. *IEEE Transactions on Robotics*, 27(2):334–345, April 2011.

[27] Carmel Majidi. Soft robotics: A perspective—current trends and prospects for the future. *Soft Robotics*, 1(1):5–11, 2014.

[28] Andrew D. Marchese, Cagdas D. Onal, and Daniela Rus. Autonomous soft robotic fish capable of escape maneuvers using fluidic elastomer actuators. *Soft Robotics*, 1(1):75–87, 2014. PMID: 27625912.

[29] Donald W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.

[30] W. McMahan, B. A. Jones, and I. D. Walker. Design and implementation of a multi-section continuum robot: Air-octor. In *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2578–2585, Aug 2005.

[31] K. A. Melsaac and J. P. Ostrowski. A geometric approach to anguilliform locomotion: modelling of an underwater eel robot. In *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C)*, volume 4, pages 2843–2848 vol.4, May 1999.

[32] Yiğit Mengüç, Yong-Lae Park, Hao Pei, Daniel M. Vogt, Patrick M. Aubin, Ethan Winchell, Lowell Fluke, Leia A. Stirling, Robert J. Wood, and Conor James Walsh. Wearable soft sensing suit for human gait measurement. *I. J. Robotics Res.*, 33:1748–1764, 2014.

[33] Sarthak Misra, K. T. Ramesh, and Allison M. Okamura. Modeling of tool-tissue interactions for computer-based surgical simulation: A literature review. *Presence (Cambridge, Mass.)*, 17(5):463–463, Oct 2008. 20119508[pmid].

[34] N. Mohajerin, M. Mozifian, and S. Waslander. Deep learning a quadrotor dynamic model for multi-step prediction. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2454–2459, May 2018.

[35] Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, New York, NY, USA, second edition, 2006.

[36] Cagdas D Onal and Daniela Rus. Autonomous undulatory serpentine locomotion utilizing body dynamics of a fluidic soft robot. *Bioinspiration & Biomimetics*, 8(2):026003, mar 2013.

[37] Yong-Lae Park, Santhi Analytis, Bruce Daniel, Seok Chang, Mihye Shin, Joan Savall, R Black, B Moslehi, and Mark Cutkosky. Real-time estimation of three-dimensional needle shape and deflection for mri-guided interventions. Jan 2010.

[38] Yong-Lae Park, Bor rong Chen, Néstor O Pérez-Arancibia, Diana Young, Leia Stirling, Robert J Wood, Eugene C Goldfield, and Radhika Nagpal. Design and control of a bio-inspired soft wearable robotic device for ankle–foot rehabilitation. *Bioinspiration & Biomimetics*, 9(1):016007, jan 2014.

[39] Manish N. Patel, Mahendra Bhandari, Mani Menon, and Craig G. Rogers. Robotic-assisted partial nephrectomy: Has it come of age? *Indian journal of urology : IJU : journal of the Urological Society of India*, 25(4):523–528, 2009. 19955680[pmid].

[40] Jessie Osborne Paull, Abdullah I. Alalwan, and Vincent Obias. *Next-Generation Robots for taTME*, pages 465–474. Springer International Publishing, Cham, 2019.

[41] R. S. Penning, J. Jung, J. A. Borgstadt, N. J. Ferrier, and M. R. Zinn. Towards closed loop control of a continuum robotic manipulator for medical applications. In *2011 IEEE International Conference on Robotics and Automation*, pages 4822–4827, May 2011.

[42] Ryan S. Penning and Michael R. Zinn. A combined modal-joint space control approach for continuum manipulators. *Advanced Robotics*, 28(16):1091–1108, 2014.

[43] Giuseppe Maria Prisco, Craig R. Gerbi, Theodore W. Rogers, and John Ryan Steger. United States Patent: 8545515 - Curved cannula surgical system, October 2013.

[44] H. Ren and P. E. Dupont. Tubular structure enhancement for surgical instrument detection in 3d ultrasound. In *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 7203–7206, Aug 2011.

[45] III Robert J. Webster and Bryan A. Jones. Design and kinematic modeling of constant curvature continuum robots: A review. *The International Journal of Robotics Research*, 29(13):1661–1683, 2010.

[46] Hyunsuk Frank Roh, Seung Hyuk Nam, and Jung Mogg Kim. Robot-assisted laparoscopic surgery versus conventional laparoscopic surgery in randomized controlled trials: A systematic review and meta-analysis. *PLOS ONE*, 13(1):1–12, 01 2018.

[47] W. S. Rone and P. Ben-Tzvi. Continuum robot dynamics utilizing the principle of virtual power. *IEEE Transactions on Robotics*, 30(1):275–287, Feb 2014.

[48] D. C. Rucker, B. A. Jones, and R. J. Webster III. A geometrically exact model for externally loaded concentric-tube continuum robots. *IEEE Transactions on Robotics*, 26(5):769–780, Oct 2010.

[49] Daniela Rus and Michael T. Tolley. Design, fabrication and control of soft robots. *Nature*, 521:467 EP –, May 2015.

[50] A. Sadeghi, A. Tonazzini, L. Popova, and B. Mazzolai. Robotic mechanism for soil penetration inspired by plant root. In *2013 IEEE International Conference on Robotics and Automation*, pages 3457–3462, May 2013.

[51] N. Sebkhi, N. Sahadat, S. Hersek, A. Bhavsar, S. Siahpoushan, M. Ghovanloo, and O. Inan. A deep neural network-based permanent magnet localization for tongue tracking. *IEEE Sensors Journal*, pages 1–1, 2019.

[52] J. Sganga and D. Camarillo. Orientation estimation of a continuum manipulator in a phantom lung. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2399–2405, May 2017.

[53] C. Shi, X. Luo, P. Qi, T. Li, S. Song, Z. Najdovski, T. Fukuda, and H. Ren. Shape sensing techniques for continuum robots in minimally invasive surgery: A survey. *IEEE Transactions on Biomedical Engineering*, 64(8):1665–1678, Aug 2017.

[54] B. Shih, D. Drotman, C. Christianson, Z. Huo, R. White, H. I. Christensen, and M. T. Tolley. Custom soft robotic gripper sensor skins for haptic object visualization. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 494–501, Sep. 2017.

[55] Nabil Simaan, Rashid M. Yasin, and Long Wang. Medical technologies and challenges of robot assisted minimally invasive intervention and diagnostics. *CoRR*, abs/1807.03731, 2018.

[56] P. Slade, A. Gruebele, Z. Hammond, M. Raitor, A. M. Okamura, and E. W. Hawkes. Design of a soft catheter for low-force and constrained surgery. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 174–180, Sep. 2017.

[57] S. Song, C. Hu, M. Li, W. Yang, and M. Q. . Meng. Two-magnet-based 6d-localization and orientation for wireless capsule endoscope. In *2009 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pages 2338–2343, Dec 2009.

[58] S. Song, Z. Li, M. Q. . Meng, H. Yu, and H. Ren. Real-time shape estimation for wire-driven flexible robots with multiple bending sections based on quadratic bézier curves. *IEEE Sensors Journal*, 15(11):6326–6334, Nov 2015.

[59] John R. Spletzer and Camillo J. Taylor. Dynamic sensor planning and control for optimally tracking targets. *The International Journal of Robotics Research*, 22(1):7–20, 2003.

[60] Joseph A. Stand, Robert A. Didomenico, Brett Zubiate, Ian J. Darisse, and J. Christopher Flaherty. United States Patent: 10238460 - Highly articulated robotic probes and methods of production and use of such probes, March 2019.

[61] Zhenglong Sun, Shaohui Foong, Luc Maréchal, U-Xuan Tan, Tee Hui Teo, and Asim Shabbir. A non-invasive real-time localization system for enhanced efficacy in nasogastric intubation. *Annals of Biomedical Engineering*, 43(12):2941–2952, Dec 2015.

[62] Philip J. Swaney, Arthur W. Mahoney, Bryan I. Hartley, Andria A. Remirez, Erik Lamers, Richard H. Feins, Ron Alterovitz, and Robert J. Webster. Toward transoral peripheral lung access: Combining continuum robots and steerable needles. *Journal of Medical Robotics Research*, 2(1):1750001, 2017.

[63] T. D. Than, G. Alici, H. Zhou, and W. Li. A review of localization systems for robotic endoscopic capsules. *IEEE Transactions on Biomedical Engineering*, 59(9):2387–2399, Sep. 2012.

[64] Michael T. Tolley, Robert F. Shepherd, Bobak Mosadegh, Kevin C. Galloway, Michael Wehner, Michael Karpelson, Robert J. Wood, and George M. Whitesides. A resilient, untethered soft robot. *Soft Robotics*, 1(3):213–223, 2014.

[65] D. Trivedi, A. Lotfi, and C. D. Rahn. Geometrically exact models for soft robotic manipulators. *IEEE Transactions on Robotics*, 24(4):773–780, Aug 2008.

[66] R. J. Webster, A. M. Okamura, and N. J. Cowan. Toward active cannulas: Miniature snake-like surgical robots. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2857–2863, Oct 2006.

[67] M. C. Yip and D. B. Camarillo. Model-less feedback control of continuum manipulators in constrained environments. *IEEE Transactions on Robotics*, 30(4):880–889, Aug 2014.