

Improving Robotic Manipulation: Techniques for Object Pose Estimation, Accommodating Positional Uncertainty, and Disassembly Tasks from Examples

by

© *Viral Rasik Galaiya*

A thesis submitted to the
School of Graduate Studies
in partial fulfilment of the
requirements for the degree of
Master of *Science*

Supervisor: Dr. Vinicius Prado da Fonseca

Dr. Xianta Jiang

Department of *Computer Science*

July 2024

Memorial University of Newfoundland

St. John's

Newfoundland

Abstract

To use robots in more unstructured environments, we have to accommodate for more complexities. Robotic systems need more awareness of the environment to adapt to uncertainty and variability. Although cameras have been predominantly used in robotic tasks, the limitations that come with them, such as occlusion, visibility and breadth of information, have diverted some focus to tactile sensing. In this thesis, we explore the use of tactile sensing to determine the pose of the object using the temporal features. We then use reinforcement learning with tactile collisions to reduce the number of attempts required to grasp an object resulting from positional uncertainty from camera estimates. Finally, we use information provided by these tactile sensors to a reinforcement learning agent to determine the trajectory to take to remove an object from a restricted passage while reducing training time by pertaining from human examples.

Acknowledgements

I would never have even started my master's program if it were not for Vinicius' enthusiasm to share the work he was doing after a chance encounter with him while rock climbing. I'd like to thank my supervisors, Dr. Vinicius Prado de Fonseca and Dr. Xianta Jiang, for their invaluable support, guidance, and patience.

I would also like to acknowledge our collaborator, Dr. Thiago Eustaquio Alves de Oliveira, for his ideas and advice on the work I have done, as well as Mohammed Asfour for his input in the second chapter.

This would not have been possible without the encouragement of my family, Rasik Galaiya, Ila Galaiya and Kangana Galaiya. Many thanks to Yanitza Trosel for all her support and advice and her mighty sidekick Suki.

A shoutout to my friends in Computer Science, Marius Seidl, for the engaging conversations, and Elahe Mohammadreza for the camaraderie as well as all my other lab mates and colleagues.

Finally, all the friends I made in Newfoundland who helped make it my new home; Neus Campanya Llovet, Celyn Khoo, John Green, Haris Mubashir, Pranjal Patra, and Ali Ghamartale.

Contents

Abstract	ii
Acknowledgements	iii
List of Tables	vii
List of Figures	viii
1 Introduction	1
1.1 Background	1
1.2 Literature Review	3
1.3 Motivation	5
1.4 Objectives	6
1.5 Contributions	7
1.6 Thesis Outline	8
2 Exploring Tactile Temporal Features for Object Pose Estimation During Robotic Manipulation	9
2.1 Introduction	9

2.2	Literature Review	12
2.3	Materials and Methods	17
2.3.1	Data Collection	17
2.3.2	Data Characteristics	20
2.3.3	Pre-processessing	22
2.3.4	The angle estimation model	25
2.3.4.1	Model Architecture	25
2.3.4.2	Hyperparameters and Window Size Optimization	25
2.4	Results	27
2.4.1	Model Training	28
2.4.2	Window Size	29
2.4.3	Comparing this Temporal Deep Learning Method to Ridge Regression	31
2.5	Discussion	32
2.6	Conclusions	35
3	Grasp Approach Under Positional Uncertainty Using Compliant Tactile Sensing Modules and Reinforcement Learning	37
3.1	Introduction	37
3.2	Related Work	39
3.3	Methodology	42
3.3.1	Compliant Tactile Sensing	42
3.3.2	Reinforcement Learning	44
3.3.3	Object Pose Uncertainty	47

3.4	Experimental Setup	48
3.5	Results	51
3.6	Conclusion and Future Work	52
4	Human Example Pretraining Influences on Reinforcement Learning for Obscured Object Trajectory Guidance	53
4.1	Introduction	53
4.2	Literature Review	55
4.3	Materials and Methods	57
4.3.1	Sensor and Manipulator Setup	58
4.3.2	3D printed Object Set used for Extraction	61
4.3.3	Teleoperation and Data collection for Pretraining	62
4.3.4	Pretraining the Neural Network	62
4.3.5	Object extraction using Reinforcement Learning	63
4.4	Results	64
4.5	Conclusion and Future works	67
5	Conclusion and Future Work	69
	Bibliography	73

List of Tables

2.1	The average frequency of the data obtained from its respective sensors.	21
2.2	The hyperparameters' values of the neural network.	27
2.3	The detailed results of the inspected window size range using MAE and MSE errors in radian, R^2 score, and EXP.	30
2.4	The results of standard regression models using MAE and MSE errors in radian, R^2 score, and EXP.	32

List of Figures

2.1	The underactuated gripper [1] diagram with two fingers, each with two phalanges and their respective sensors.	18
2.2	The sensor with its base attached to the manipulator and close to where the barometer is, and its surface over the MARG sensor is in contact with the object.	18
2.3	The object's two markers to obtain the ground truth angle using computer vision [1].	19
2.4	The angle rotation, the corresponding pressure, and one of the four MARG outputs in one of the trial data collection trials.	21
2.5	The procedure used for sensors alignment. (a) For every angle determined by the camera, the closest corresponding pressure and MARG values were selected. (b) Sensor values were grouped with the corresponding camera frame. (c) Downsampling MARG values to match the frequency of the pressure sensor.	23

2.6	The model architecture and the feed-forward of a single input sample through the model. The architecture consists of 2 LSTM layers with normalization layers with 512 Units and 256 units, respectively, and three fully-connected layers with 128, 64, and 32 neurons.	26
2.7	Average training and validation losses and their standard variation for the different window sizes, highlighting the average epoch of the lowest validation error.	28
2.8	The performance results from varying the window size. (a) Mean Absolute Error (MAE). (b) Mean Squared Error (MSE). (c) Coefficient of Determination (R^2 Score). (d) Explained Variance (EXP).	29
2.9	A comparison of the predicted and ground truth angle for one of the iterations of the model with the best window size of 40 samples. A 40 samples window corresponding to the prediction of angle sample 900 is highlighted. There are persistent deviations on some peaks and troughs such as around sample 1000, 1250 and 2375 which are more likely to occur when the object is static. This may occur because no new information is obtained by the sensor when the object is static and as a result the error in prediction can't be corrected.	31

3.1	The camera provides an estimate of the object’s location. The manipulator attempts to grasp it and return the system’s state, i.e., the success/failure, the force experienced by each sensor and the orientation of the sensor to the reinforcement learning agent. The agent uses that information to update the position to attempt a grasp again if unsuccessful prior.	38
3.2	The blue arrows are the direction of the joint movement of the manipulator. The Environment provides the orientation of the sensors attached to the end effector and the pressure experienced by the sensors. This information is sent to the agent every time a grasp is attempted. The agent sends a new position for trial if unsuccessful.	40
3.3	Tactile sensing module. Components: 1—MARG (magnetic, angular rate, and gravity) system; 2—compliant structure; 3—barometer [2, 3, 4].	42
3.4	The manipulator attempts to grasp the white cuboid by exploring the area within the vicinity of the estimated location provided.	48
3.5	The pressure and the orientation of each sensor on the end effector during each step in the episode.	49
3.6	The trend line is calculated with a 50 moving average. The trend line of the reward increasing with the episodes and the number of steps per episode reducing. As the number of episodes increase the rewards and number of steps reach an asymptote.	51

4.1	The process is broken down into three steps. We first record the actions and the environment using teleoperation for object extraction in the peg-in-hole task. We use the recorded data to train a neural network. We use the trained neural network as a foundation for the Q-learning reinforcement learning setup for the object extraction.	58
4.2	This is the manipulator and sensor setup we are using for this experiment. The four degrees of freedom manipulator with two tactile sensors, one attached at the end of each phalanx. The red arrows represent the joint motions for the four degrees of freedom manipulator. The manipulator is controlled with respect to the end effector frame \mathbb{E} in the directions shown with the black, green, blue and white arrows. The positions with respect to the base frames \mathbb{O} are converted to the positions with respect to the end-effector frame \mathbb{E} . Resulting in motions on the x, y and z axis as well as rotation about the y and z axis with respect to \mathbb{E} . We also obtained the position of the tactile with respect to \mathbb{O} sensor based on the data it provided.	59
4.3	These are pegs and the respective holes used in this peg-in-hole setup. Top right: a slanted, curved and vertical peg. Left: the respective holes for each of the pegs. Bottom right: the cross-section of respective holes.	61
4.4	This figure shows the reward that corresponds to the steps for each episode for some trials until the last episode, where the attempt was successful. Object one pretrained on object one and took more than 60 episodes.	65

4.5	This figure shows a comparison between the number of episodes required to complete an extraction on all objects with pretraining or no pretraining.	66
4.6	This figure shows a comparison between box plots of the number of episodes taken for a successful extraction with any form of pretraining and no pretraining for each object.	66
4.7	This figure shows box plots for the number of episodes required to successfully extract the object based on the pretraining on other objects. This explores the generalizability of the pretraining.	67

Chapter 1

Introduction

1.1 Background

Outsourcing human labor to a being or device has been conceptualized since as early as the Greek-Hellenistic age [5], while there are even records of a musical automaton in the 1100s [6]. Although the idea of humanoid robots has been around in various ages and cultures ranging from Roman mythology to ancient India, the widespread robotics application in the 1950s was machine tools combined with manipulators to perform simple, repeatable activities without information about the environment. With time, these manipulators started becoming more sophisticated, incorporating advances in machine learning, sensor technology [7], and mechanical and material advancements to come full circle to humanoids. These first implementations were very rudimentary compared to today's standards because of the lack of information on the environment and were operated in highly structured environments. With the development of sensors, computing, and application breadth, robotic applications

keep moving closer to unstructured environments. The ability to accurately infer the environment and adapt to this variable environment plays a significant role in robots' wider range of uses. Zamalloa *et al.* [7] describes the trend as moving from mass manufacturing to mass customization.

Since the development of robotics, the breadth and depth have increased tremendously. Despite manipulators being the first at mass adoption [7] and becoming just a subset of the whole field that now includes drones and quadrupeds, it is still a highly researched area due to the range challenges still far from being solved like gripper design [8], control [9] and sensing [10].

Humans have an extensive range of sensing capabilities that enable us to perform very complex tasks seemingly efficiently. We rely highly on visual perception systems but also have other very powerful sensory and information input systems. These systems include sight, touch, sound, and smell. These systems can be used very well together but are also used well in isolation. Helbig *et al.* [11] suggests that since no information-processing system has the ability to provide sufficient information under all conditions, it is imperative to combine multiple systems. This line of thought has encouraged researchers to explore forms of perception systems other than vision directly, ranging from inertial measurement units [12] to rangefinders [13]. In particular, vision has limited capabilities for several grasping and manipulation tasks, specifically for near contact and contact activities. This is because we can't generally obtain information on weight, force, and friction with an unmodified vision system. Albeit, creative methods of using cameras have been invented to obtain tactile information [14].

Despite the wide range of sensor data as well as the quantity of data, there is

intrinsic uncertainty from the sensor’s accuracy, geometric uncertainty from sensor limitations like placement or types, physical uncertainty from properties such as material composition or weight as well as environmental disturbances from the environment, especially in unstructured environments. Reinforcement learning with tactile sensing allows for adapting to these uncertainties and changing environments [15].

1.2 Literature Review

For any manipulation task, it is vital to know the pose of the object being manipulated. It is also important to keep track of changes that are made when manipulating the object so the trajectory of the manipulator can be efficiently calculated. Amongst various other fields, object pose estimation has benefited highly from the advent of deep learning techniques. Du *et al.* [16] explores work done in 2D planner grasps six-degree grasp planning including constraint grasps, object pose estimation using template matching-based methods, voting-based methods, end-to-end grasp estimation using both RGB-D images. They concluded that there needs to be a close similarity of the 3D object model the network is trained on to have accurate grasps. Moreover, the lack of geometry information from the direction not facing the camera affects the accuracy of the shape completion and suggests the use of multi-source data. Therefore, other data sources have been explored, especially close to or in relation to the contact points, such as using tactile sensing or proprioception [17]. These methods generally complement vision-based sensing modalities to address blind spots, increase accuracy or both. Various ways of leveraging the information of tactile sensing for pose estimation have been used, including collisions [18, 19], simulation to real life

[20, 21] and reinforcement learning [22].

Grasping in unstructured environments must overcome significant uncertainties from things like occlusion and noise and accommodate many object properties like geometry and mass. Under-actuation and compliance are great low-cost rectifiers. This is because fewer motors are required to grasp an object, and compliance increases adaptability, reducing the effects of uncertainty. Although these techniques can result in position error because of forces dislodging the object, they can be reconciled using information obtained from the end-effector, such as tactile sensors [23]. Various compliant tactile sensors [24, 25, 26, 2, 27, 4] have been developed to improve grasp force control, prevent slip, improve pose estimation and accommodate positional uncertainty. Kuppuswamy *et al.* [28] use a combined visio-tactile approach to update visual estimates of the object pose such that it converges with less than 15 touches. In contrast, Gasparetto *et al.* [5] uses tactile information for global localization of an object using a convergence based on local geometry. Jiang and team [29] used a visio-tactile based approach as well; the camera providing a region to poke transparent objects and obtain the local information that can be used to increase the likelihood of a successful grasp.

Once the object has been grasped, the object can be used for assembly tasks such as peg-in-hole, slide-in-groove, bolt screwing, pick and place, and pipe connection [30]. Although peg-in-hole [31, 15, 32, 33, 34] is a commonly explored topic, most of the research focuses on collision with the edges and aligning the object with the hole. Dong *et al.* [15] compared the performance of insertion policies using tactile RGB and tactile flow. They also compare the performance of a tactile sensor and a torque/force sensor with a vision-based tactile sensor. Finally they compare supervised learning

and reinforcement learning. Testing with 4 different objects, using the tactile sensor with a tactile flow presentation and reinforcement learning provided the best results. Various methods have been investigated for insertion, such as using perturbations [34], uncertainty distribution exploration [35], reinforcement learning, and supervised learning [36].

1.3 Motivation

In-hand object position is important knowledge for various complicated tasks such as repositioning the object, collaborating with other robots or humans and interacting with the environment. However, these tasks result in changes in the object position with these interactions, and it can be beneficial to keep track of these changes or identify the changes after the interaction to adapt to them and perform tasks based on accurate and updated information. Moreover, when these interactions occur, the proximity to the object is small, the likelihood of the source of the interaction is high, and the obscuring cameras are more likely to occur in these situations. Hence tactile sensors may be used to predict and update the position. Moreover, since these motions are sequential, taking into account the temporal aspect of the signals from the tactile sensors can assist in increasing accuracy.

This occlusion can also occur by the manipulator during approach when grasping. The accuracy of the object pose obtained from the camera can be affected by varying light conditions and object characteristics like colour and environmental characteristics like the wind. The use of tactile sensors and adaptive algorithms can compensate for these factors by updating the internal state of the error from the interactions,

which are both successful and unsuccessful.

These problems associated with occlusion are also prevalent in disassembly tasks because parts of the objects that need to be disassembled are generally occluded, and those parts influence the trajectory required, such as the peg-in-hole. These limitations can be accommodated by using adaptive algorithms like reinforcement learning. In particular, the disassembly trajectory is updated based on the interactions with the environment. Moreover, the system will take a shorter time to improve its performance if it is being implemented over a previously trained system, for instance, from human demonstrations.

1.4 Objectives

Based on this motivation, we defined the following objectives:

- Use tactile sensors to estimate the object pose after grasp. Study deep learning to take advantage of the temporal information obtained from tactile sensors to predict object pose based on external forces.
- Develop methods to estimate object pose by using reinforcement learning to adapt to uncertainties using tactile feedback.
- Evaluate the effect of pre-training on reinforcement learning in extracting objects from a peg-in-hole setup.

The goal of this thesis is to explore the usage of tactile sensing in the various components of robotic manipulation. Specifically, to improve grasping, obtain pose

estimations of objects in hand and guide the extraction process of the peg-in-hole problem.

1.5 Contributions

pretraining In this thesis, we use a long short-term memory network with tactile information to predict the object’s angle. We subsequently use reinforcement learning and a compliant tactile sensor to address pose uncertainty from camera inputs. Finally, we use a pertaining with human examples strategy with reinforcement learning to speed up the learning phase of the agent to extract the peg in the peg-in-hole setup.

The work of this thesis has been published partially in other versions in peer-reviewed conferences and journals. We have the approval of the publishers to use the published findings in the thesis. Published versions of this research Chapter:

- Chapter 2 explores the use of temporal tactile data via a sliding window to predict the object angle under grasp. This paper is published as **Galaiya VR, Asfour M, Alves de Oliveira TE, Jiang X, Prado da Fonseca V. Exploring Tactile Temporal Features for Object Pose Estimation during Robotic Manipulation. Sensors. 2023; 23(9):4535. [37]**
- In Chapter 3, we devise a policy that effectively models object position estimation errors and reduces exploratory sensor contact, improving the grasp execution process. This paper is approved as **Galaiya VR, Alves de Oliveira TE, Jiang X, Prado da Fonseca V. Grasp Approach Under Positional Uncertainty Using Compliant Tactile Sensing Modules and Reinforce-**

1.6 Thesis Outline

In Chapter 2, we explore the effects of incorporating the temporal data of tactile information for pose estimation. We use a long short-term memory network to predict the object's angle and compare it to standard regression models that do not utilize the time series information. Chapter 3 explores the use of reinforcement learning and a compliant tactile sensor to address pose uncertainty from camera inputs. Finally, in Chapter 4, we explore the use of tactile sensors in disassembly tasks using reinforcement learning.

Chapter 2

Exploring Tactile Temporal Features for Object Pose Estimation During Robotic Manipulation

2.1 Introduction

Many areas of human activity, such as mass-production factories, minimally invasive surgeries, and prostheses have adopted robotic manipulation systems. Robotic manipulation is exceptionally reliable when the system has complete information regarding the environment. These systems usually must follow a set of trajectories, interact with objects of known features, and perform repetitive tasks with minimal environmental adaptation, which limits the use of manipulation systems to perform

activities in unstructured settings. Recent advancements in data-driven methods, innovative gripper design, and sensor implementation have reduced the limitations of robotic manipulators in such environments. Nevertheless, there are hurdles to the applications of robotic arms in unstructured environments or dexterous tasks such as complex manipulation of daily objects [38].

One main challenge is estimating the object’s orientation during the after-grasp phase. The object’s orientation can change from an initial visual estimation due to calculation errors, external forces, finger occlusion, and clutter. After a successful grasp, one approach is to use tactile sensors to extract object information, improving the object’s pose estimation.

Robotic hands have immense flexibility despite their use in specific domains, such as prostheses, with limitations regarding the human hand’s size, weight, and shape. By sacrificing initial stability and uncertainty in grasp pose estimation, an under-actuated approach substantially reduces planning time and gripper design complexity [39]. However, it is fundamental for robotic arms to estimate the handled object’s pose to operate optimally in object manipulation applications. For instance, the grasp used by a gripper of a robotic arm or a prosthesis to hold a mug might change if its handle is at a different angle.

Object orientation estimation depends on several aspects, such as the gripper’s configuration, the sensors used, and how the data is analyzed. Different sensors are used in robotic manipulation to categorize the properties of an object, such as its orientation. Pre-grasp poses are commonly obtained using computer vision [40]. However, having visual data only can be insufficient as the gripper approaches the object and the range of occlusion increases. This limitation is particularly pronounced

when the camera’s location is fixed or under unpredictable circumstances, i.e., in unstructured environments. For instance, using a top-view camera to estimate the object pose is not feasible for an arm prosthesis, whereas prosthesis-mounted cameras are susceptible to occlusion. Moreover, once the gripper grasps the object, it will cover at least part of it, making it difficult to estimate its orientation. Furthermore, merely using vision cannot reduce forces and related environmental stimuli, leading to potential errors in the estimation of orientation due to miscalculated geometry, friction, forces, camera occlusion, and clutter [38, 41].

Due to the limitations of visual methods, several applications use tactile sensing [42, 43, 44, 45, 46, 47] while grasping the object, providing more relevant information that is not interrupted [12]. Tactile sensing has shown promise in specific use cases, such as in minimally invasive surgery [48] or cable manipulation [49], and is also being shown to be a good supplement to control system optimization [50, 51]. Sensors such as pressure sensors [47], force sensors [52], and inertial sensors [2] are gradually becoming more prevalent for object pose estimation and object recognition. In addition, tactile sensors provide after-grasp contact information about the object that can be used for control [49] or in-hand manipulation. Nevertheless, there have also been developments of vision-based tactile sensors ranging from using internal reflection [53, 54] to observing the deformation of the surface [55].

Tactile sensing can be a vital link to overcoming computer vision limitations and can result in better performance of robotic manipulation. Previous works have used machine learning models and visual frames of reference to train models that learned the after-grasp object angle, which is later used to estimate the object’s pose [1]. However, previously seen data can affect the current estimation of the object’s state.

So, we hypothesize that estimating the current object pose can be improved by considering temporal data, such as when sliding window sampling. For this reason, in the present work, we study the effect of temporal data based on sliding window sampling to train a deep learning model for object angle estimation.

2.2 Literature Review

Orientation estimation has been a part of pose estimation in robotics research for a long time. Recent studies have made leaps regarding orientation estimation with sufficiently low error due to advancements in sensors technology, most importantly tactile sensors [56, 57, 58, 59].

For instance, Ji *et al.* [60] proposed a novel model-based scheme using a visual-tactile sensor (VTS) [61]. In their study, the sensor comprised a deformable layer that interacted with objects with a depth camera behind said layer to generate a depth map of the deformation caused by the object. They reported orientation errors for three objects under 3° . However, detecting their objects' rotations could have been visually easier compared to more uniform smooth shapes such as cylinders or ellipsoids.

Additionally, Suresh *et al.* [62] formulated the tactile sensing problem as a simultaneous localization and mapping (SLAM) problem, in which the robot end effector made multiple contacts with the object to determine its pose. They reported a rotational root mean square error (RMSE) of 0.09 radians. However, their method assumed the initial pose and scale of the object roughly and neglected factors outside the controlled setup that might change the object's orientation.

Other studies utilized information about the robot arm alongside tactile data to estimate the orientation of objects. Alvarez *et al.* [63] used the kinematic information and a particle filter for pose estimation via tactile contact points, force measurements, and angle information of the gripper’s joints. Their algorithm initiates a pose estimation using visual data, which is refined by a particle filter based on the optical data. After experiments with three objects of different sizes, they report a 0.812° error in their best experiment, which rises to 3.508° in their worst case. Results aside, the method requires a known kinematic model of the robotic arm and a top-view camera for inference, which is infeasible in some applications, like daily activities using prostheses.

To relax the requirement of a detailed kinematic model, recent research has explored underactuated grippers while relying on machine learning methods to build a model of the object pose. For instance, Azulay *et al.* [64] conducted a wide-scope study to investigate objects’ pose estimation and control with under-actuated grippers. They incorporated haptic sensors, joint angles, actuator torques, and a glance at the pose at the start of the gripper’s movement. Using the robotic arm’s kinematic model, they concluded that some combinations of the tested features are better suited for object manipulation than others. They report a root mean square error (RMSE) of $3.0 \pm 0.6^\circ$ for orientation using a neural network with LSTM layers, their best model. Using multiple features alongside the kinematic model can be computationally intensive for processors on devices such as prosthetics.

However, investigating orientation estimation itself only prior to grasping can limit

the reported results in some situations. For instance, robotic grippers that handle objects can occlude the object partially or fully, affecting visual-based approaches. Furthermore, objects can rotate during handling due to many factors, such as slipping or external forces, thus requiring methods to estimate objects' orientation during the grasp phase.

High-density tactile sensors, akin to the human hand, are another direction that can provide much information. Funaabashi *et al.* [65] used graph convolution neural networks (GCNs) to extract geodesical features from 3-axis tactile sensors across 16 degrees of freedom of a robotic hand providing 1168 measurements at 100Hz. They used eight objects with two different hardness, slipperiness, and heaviness factors. They compared various GCN configurations and a multilayer perceptron, and the GCN model with the most convolution layers was the best performer. The limitation of this method is due to the need for high computational resources, the requirement of a large number of sensors, and the ambiguity of intermediate states, although accounting for different properties, such as hardness and slipperiness, improves possibilities of generalization.

To develop a solution that required minimal finger path planning, relaxed kinematic model requirement, and less needed processing of images, Da Fonseca *et al.* [1] developed an underactuated gripper with four compliant sensing modules on flexible fingers, and investigated the collected sensor data while grasping objects of three distinct sizes. The experiments included a top-view camera to obtain a visual frame of reference for ground truth orientation. The method used tactile sensors' information to represent the object angle, whereas the ground truth angle was obtained from the

camera frame. Finally, the authors compared five regression models trained using tactile data to estimate the object’s angle. The best models reported by the authors were the ridge regression model and linear regression, obtaining a 1.82° average mean square error. The authors used random data sampling for model training in the paper and left possible relationships among the time-series samples as a future research point. Still, given that the tasks were dynamic, we expect that the near samples in the time-series sensor data are correlated with the angle.

Some studies also investigated the fusion of tactile and visual data for orientation estimation during object handling.

Alvarez *et al.* [66] proposed a fusion method of the visual data and tactile data to estimate the object’s pose during grasp. A camera tracked the object during grasp, whereas a particle filter was utilized with the tactile data to reduce the uncertainty of the object’s pose. They reported that their method obtained an orientation error varying from 1° to 9.65° . Their method yielded a high variance of the estimation error, in addition to requiring a 3d model of the handled object for the method to be used.

Dikhale *et al.* [67] proposed sensor fusion of visual and tactile data as well. Their method used neural networks to process the tactile and visual data separately before fusing them to give a final prediction of the object’s pose. They reported an angular error as small as 3° ; however, it reached a high of 24° , showing high variance in the estimation depending on the object.

From the previous studies, we find that different factors affect the orientation estimation performance and eligibility. For instance, computationally-demanding methods, such as ones relying on inverse kinematics or particle filters, are inappropriate for small devices, such as prostheses limited to an onboard processor. Whereas relying on visual data, solely or with sensor fusion, is prone to occlusion during the grasp phase as top-view cameras are not feasible in many applications. Hence, a model must only estimate the object’s angle using only tactile data during grasp without kinematics to reduce computation while providing an acceptable angle error.

We evaluate the use of sliding window sampled tactile data to estimate the yaw angle under the stable grasp of an object while relaxing the kinematic model requirement by using an underactuated gripper, and a compliant bio-inspired sensing module that includes magnetic, angular rate, gravity, and pressure sensing components. We analyze the temporal nature of tactile signals by using a neural network that contains long short-term memory (LSTM) layers to estimate the orientation with the highest precision for objects. The models trained in the present work were based on Da Fonseca *et al.* [1], taking in a window of readings from the sensors mounted on the gripper and then outputting the estimated object’s orientation at the end of this sampling window. As the chapters main topic is the in-hand orientation estimation, our method uses only the initial grasp orientation as a reference and does not require information from the gripper joints, its kinematics model, a multitude of sensors, nor during-grasp visual data.

Our method can be utilized in a multitude of applications from everyday use to factory settings due to its dependency on only a small number of tactile sensors without the need for additional types of sensors. Furthermore, our method does not need

computationally-capable machines as it utilizes only a neural network that can run on a computational device as small as a flash drive, such as Google Coral, due to advancements in computational technology. In addition, the proposed method’s performance is limited to uniform shapes whose orientation change is hard to determine visually, such as rotating cylinders.

2.3 Materials and Methods

Here we describe the data collection and pre-processing methods used for sliding window-sampling tactile data, the models trained for the experiments, and how we organized the sampling strategy for pose estimation.

2.3.1 Data Collection

We used tactile data collected in a previous study [1] from an underactuated gripper with two independently controlled fingers during object-grasping tasks to evaluate the sliding window sampling strategy for pose estimation.

In the gripper developed by Prado da Fonseca *et al.* [1], each phalanx has a fixed tactile sensor developed by Alves de Oliveira *et al.* [2], as shown in Figure 2.1.

Each sensor provides deep pressure information from a barometer in addition to angular velocity, linear acceleration, and magnetic field in all three axes using the 9-degree-of-freedom magnetic, angular rate, and gravity (MARG) system. The barometer as shown in figure 2.2 is encased in a polyurethane structure close to the base, and the MARG sensor is placed closer to the point of contact so it can detect micro-vibrations. The fabrication structure of the sensor enables the pressure to be




Figure 2.1: The underactuated gripper [1] diagram with two fingers, each with two phalanges and their respective sensors.




Figure 2.2: The sensor with its base attached to the manipulator and close to where the barometer is, and its surface over the MARG sensor is in contact with the object.




Figure 2.3: The object’s two markers to obtain the ground truth angle using computer vision [1].

transferred from the contact point to the barometer effectively. The compliant sensor structure allows the contact displacement to be measured by the inertial unit while the deep pressure sensor measures the contact forces. The data is collected using an onboard micro-controller interfacing via I2C with a computer running the Robot Operating System (ROS) framework [68].

Prado da Fonseca *et al.* [1] used the allocentric reference frame from the camera pointed down to calculate the object’s angle. The top view angle of the object is extracted using two colored markers attached to it to identify key points using the OpenCV library as shown in Figure 2.3.

The angle between this the two markers line and the fixed camera frame horizontally in the clockwise direction is established to be the object’s angle, and the object is considered at 90° on the $x - axis$. These points are later compared to a fixed frame of reference at the camera’s center to determine the object position change relative to the specified frame of the gripper. The stable grasp was obtained using

a dual fuzzy controller that obtained micro-vibrations and pressure feedback from the tactile sensor [69]. This procedure was performed with three cylindrical objects with 57 mm, 65 mm, and 80 mm diameters. The objects were rotated manually in CW and CCW directions, simulating external forces causing the object to change its orientation during grasp. Although this motion is at a low speed, the human element of this motion provides inconsistent forces, which the model was able to take into account to provide an accurate prediction. Such movements also simulate the act of parasitic motions, which are undesired motion components which lead to lower manipulation accuracy/quality [70], despite being in stable grasp. Moreover, the three different objects are used to determine the ability of the model to generalize among similar objects. The ground truth angle after rotation is obtained relative to the form of reference from the top view camera as seen in Figure 2.3.

2.3.2 Data Characteristics

The pre-processing methods used in this work depend highly on the time-series details of the data available from Prado da Fonseca *et al.* [1]. For instance, the number of instances in each window sample can be affected by the different frequencies of each sensor. Table 2.1 shows the average sampling frequency of each sensor, in which the slowest sensor is the camera, ranging from 9 to 29.95 Hz . The fastest sensor is the MARG sensor, ranging from 911.33 to 973.50 Hz .

As mentioned in Section 2.3.1, the data collection consisted of a CW and CCW rotation procedure performed by an external operator on three cylindrical objects with 57 mm, 65 mm, and 80 mm diameters. The dataset for each object contains

Table 2.1: The average frequency of the data obtained from its respective sensors.

Camera	Pressure	MARG sensor
29.95Hz	402.19Hz	973.50Hz

sensor readings from five different external rotation operations. Figure 2.4 shows the disturbances of rotation on the pressure, linear acceleration, angular velocity, and magnetic field for one sensor in relation to the angle during external rotations.

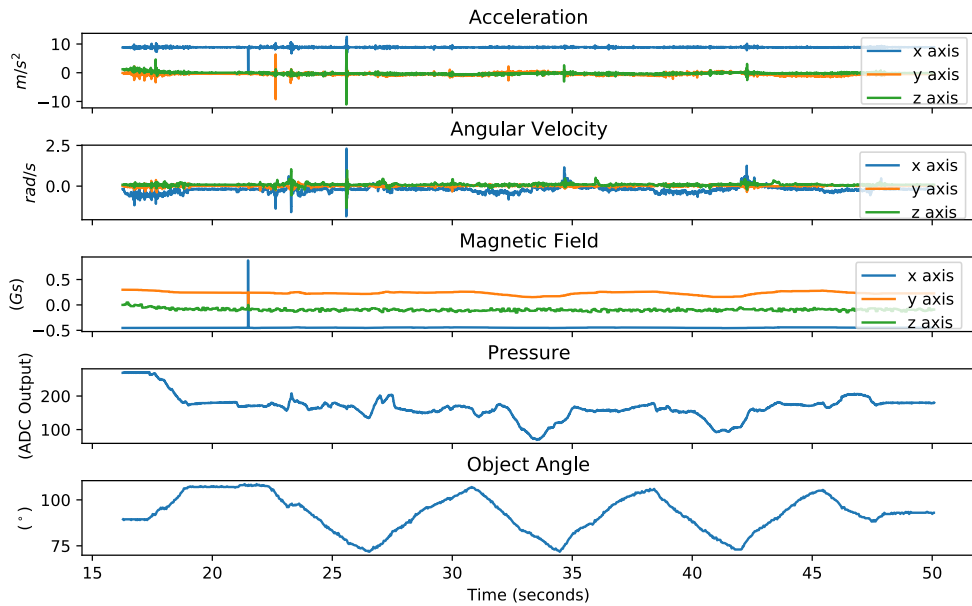


Figure 2.4: The angle rotation, the corresponding pressure, and one of the four MARG outputs in one of the trial data collection trials.

The data characteristics described here are sufficient for our investigation. Further details about the data collection protocol and attributes can be found in the original

data collection study [1].

2.3.3 Pre-processing

Algorithm 1 Pre-processing and experimentation pseudo code

- 1: **for** each barometer reading **do**:
 - 2: keep closest MARG reading
 - 3: discard the rest of the MARG reading
 - 4: **for** For each angle value **do**:
 - 5: take sensor readings of corresponding timestamp
 - 6: take ($WindowSize - 1$) previous sensor readings
 - 7: separate training and test data
 - 8: normalize training and test sensor values using the mean and standard deviation from training data
 - 9: train model using training data
 - 10: obtain performance results using test data
-

The listener ROS node collected data at different time instances since the camera, MARG sensor, and pressure communicated asynchronously. Therefore the signals needed to be aligned for our strategy of window sampling. First, we scaled the data to utilize deep learning methods. Subsequently, to add LSTM layers, we had to reconcile the sampling frequency differences for the various sensors by synchronizing and downsampling their data. Since the lowest frequency was the camera frames, their timestamps acted as a reference for our procedure of sensor alignment presented

in Figure 2.5. Afterward, we reshaped the data to incorporate the previous states for each instance of the ground truth angle.

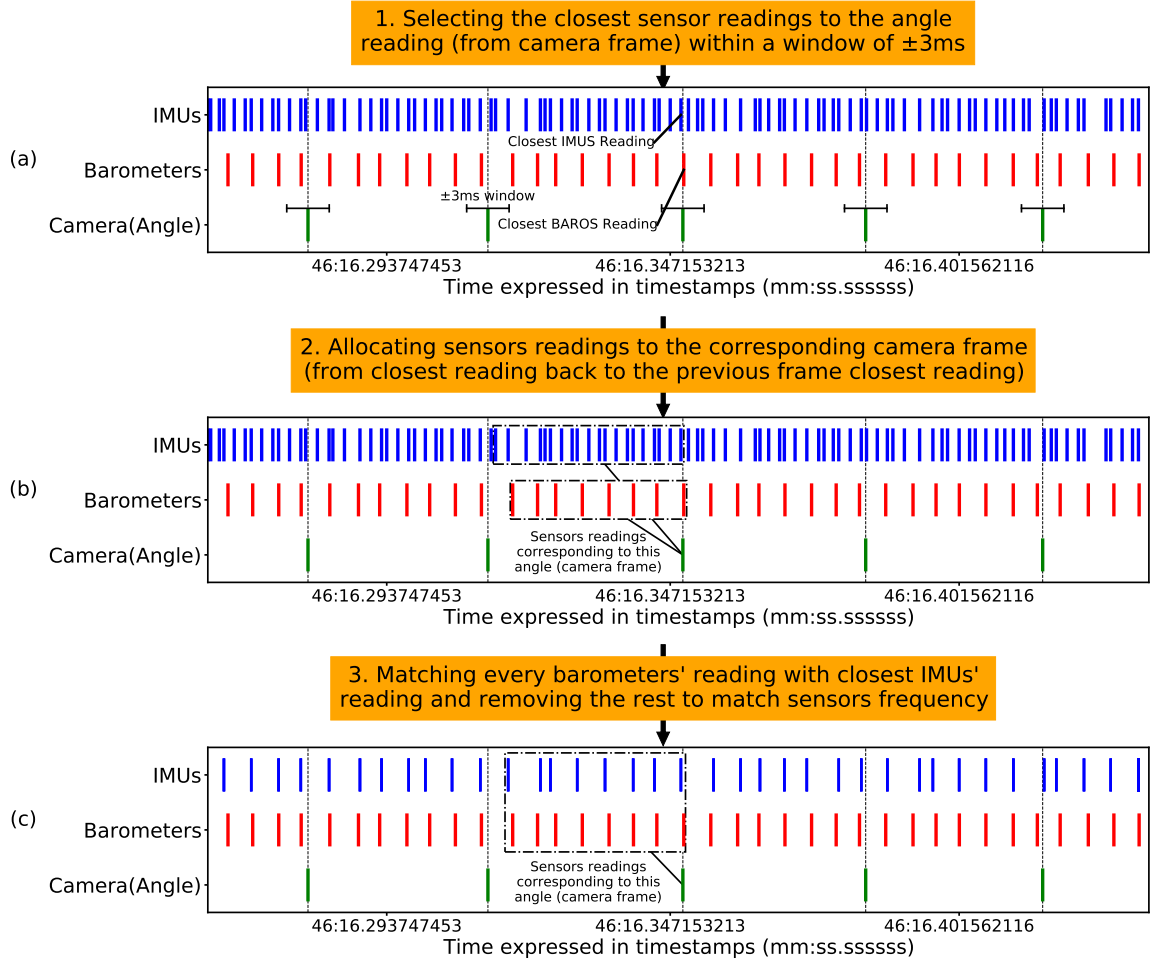


Figure 2.5: The procedure used for sensors alignment. (a) For every angle determined by the camera, the closest corresponding pressure and MARG values were selected. (b) Sensor values were grouped with the corresponding camera frame. (c) Downsampling MARG values to match the frequency of the pressure sensor.

Figure 2.5 (a) shows that the obtained pressure and MARG signals are within

three milliseconds from the angle from the camera frame, on average. Figure 2.5 (b) shows MARG and pressure values collected between two camera frames to correspond to a single frame. Finally, since the pressure is sampled at a lower frequency than the MARG sensor, Figure 2.5 (c) shows the MARG sensor reading closest to the corresponding pressure reading is kept, and the remaining samples in between the selected ones are discarded.

In this approach, small window sizes would only utilize signals corresponding to the selected camera frames. In contrast, overlapping with signals corresponding to previous frames is used to obtain more data for large window sizes.

After alignment, the final dataset contains five runs for each of the three object sizes. Each run consists of 900 camera frames, which had an average of 8 corresponding samples from the sensors per frame. We used the data for all object sizes to ensure the dataset size was sufficient for model training. Since all the sensors have different magnitudes and distributions, all the data apart from the object angle is scaled. Finally, we implement standardization on the rest of the dataset. We use the following equation to normalize each sensor’s data.

$$N^{(i)} = \frac{X^{(i)} - \mu^{(i)}}{\sigma^{(i)}} \quad (2.1)$$

Where $N^{(i)}$ is the standardized signals of the i^{th} sensor. $X^{(i)}$ are the raw signals of the i^{th} sensor, $\mu^{(i)}$ is the mean signal value of the i^{th} sensor. $\sigma^{(i)}$ is its signal’s standard deviation.

2.3.4 The angle estimation model

Since tactile sensing measurements from objects under grasp manipulation are continuous and sequential, we used time series-based neural networks, specifically long-short-term memory (LSTM) based networks, to analyze windows sampling.

2.3.4.1 Model Architecture

Using a small baseline model initially, we arrived at the final model after adding layers that provided the best marginal improvement in performance for its size without overfitting, as increasing the model’s size overfitted the training data.

Figure 2.6 shows the final model architecture we established consisting of 2 LSTM layers with normalization layers with 512 Units and 256 units, respectively, and three dense fully connected layers with 128, 64, and 32 neurons, respectively. All of the experiments are conducted on Compute Canada, an Advanced Research Computing platform, using the Python programming language and Tensorflow [71] library to preprocess and model training.

We used the mean absolute error (MAE) between the angle’s estimated and actual values as the training loss function. Moreover, we choose MAE as it diminishes in value much slower than a mean square error (MSE) as the model’s estimation gets closer to the actual angle and has a value of less than one.

2.3.4.2 Hyperparameters and Window Size Optimization

Various experiments were performed to provide an understanding of the data and identify the effects of hyperparameters and performances corresponding to their vari-




Figure 2.6: The model architecture and the feed-forward of a single input sample through the model. The architecture consists of 2 LSTM layers with normalization layers with 512 Units and 256 units, respectively, and three fully-connected layers with 128, 64, and 32 neurons.

ations. In particular, we manipulated batch sizes and windows and explored regularization methods. We explored the trade-off between window size and performance

based on the best results to find the best gain in accuracy for a small model size. This trade-off is fundamental in mobile robotics, with less memory and computational time leeway. We performed a grid search to determine the hyperparameters over the number of epochs, learning rate, and batch size. We chose the best configuration of hyperparameters to conduct the study and investigate the window sampling technique. We used cross-validation with four folds, with six iterations for the model per fold on a rolling basis, to ensure the consistency of the model’s performance, report any variance in the metrics scores and prevent data leakage. Table 2.2 shows the neural network hyperparameters.

Table 2.2: The hyperparameters’ values of the neural network.

Hyperparameter	Value
Learning Rate	0.00025
Batch Size	128
Epochs	400
K-Folds	4
Iterations	6

2.4 Results

Here we present the results of our experiment to estimate in-hand objects’ orientation using a sliding window sampling strategy and evaluate with LSTM models. The evaluation metrics used are mean squared error (MSE), mean absolute error (MAE),

Coefficient of determination (R^2), and explained variance score (EXP).

2.4.1 Model Training

Figure 2.7 depicts the training and validation losses during the training phase while highlighting the average epoch of the lowest validation error averaged over folds and model iterations.

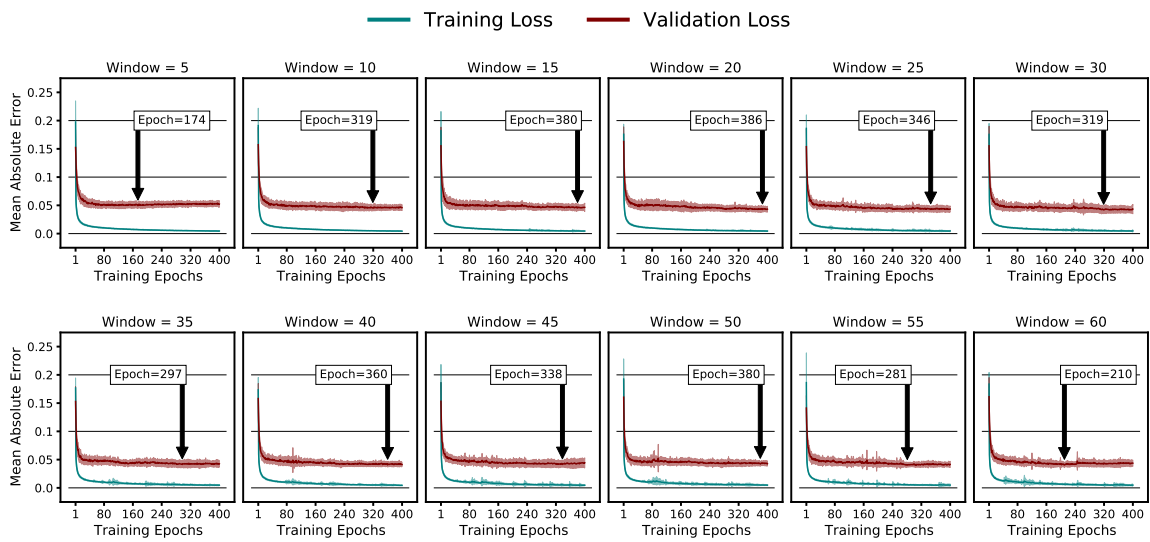


Figure 2.7: Average training and validation losses and their standard variation for the different window sizes, highlighting the average epoch of the lowest validation error.

We prevent overfitting by training the models for 400 epochs and selecting the model weights at the epoch of the lowest validation loss.

2.4.2 Window Size

The primary factor of temporal data explored was the window size. Figure 2.8 shows that a window size of 40 achieved the lowest error. It revealed a performance improvement as the window size expanded; however, the improvement magnitude decreased asymptotically.

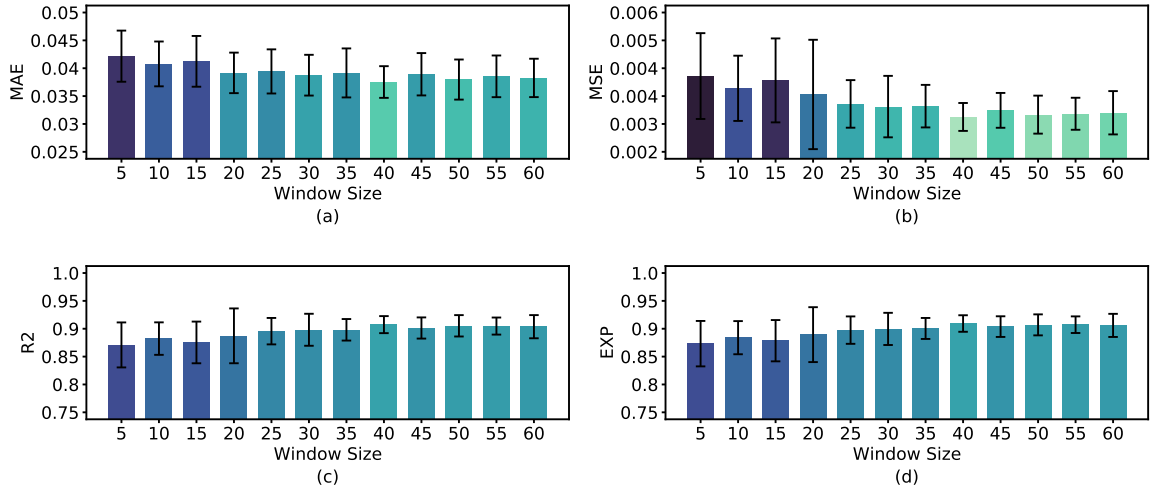


Figure 2.8: The performance results from varying the window size. (a) Mean Absolute Error (MAE). (b) Mean Squared Error (MSE). (c) Coefficient of Determination (R^2 Score). (d) Explained Variance (EXP).

The above result indicates that a window of 40 samples effectively captures the necessary amount of tactile information for estimating the object’s orientation, regardless of the metric used. Larger window sizes, beyond 40 samples, did not result in any further improvement in model performance. This finding is further supported by Table 2.3.

The model achieved the best MAE of 0.0375 radians with a window of 40 and

Table 2.3: The detailed results of the inspected window size range using MAE and MSE errors in radian, R^2 score, and EXP.

Window	MAE	MSE	R^2	EXP
5	0.0422 ± 0.0046	0.0042 ± 0.0012	0.8710 ± 0.0404	0.8732 ± 0.0408
10	0.0408 ± 0.0040	0.0038 ± 0.0009	0.8823 ± 0.0292	0.8840 ± 0.0297
15	0.0412 ± 0.0046	0.0041 ± 0.0012	0.8754 ± 0.0374	0.8785 ± 0.0370
20	0.0392 ± 0.0036	0.0036 ± 0.0016	0.8873 ± 0.0492	0.8894 ± 0.0493
25	0.0394 ± 0.0040	0.0034 ± 0.0007	0.8956 ± 0.0237	0.8975 ± 0.0246
30	0.0388 ± 0.0037	0.0033 ± 0.0009	0.8981 ± 0.0287	0.8997 ± 0.0289
35	0.0392 ± 0.0044	0.0033 ± 0.0006	0.8981 ± 0.0193	0.9005 ± 0.0188
40	0.0375 ± 0.0028	0.0030 ± 0.0004	0.9074 ± 0.0153	0.9094 ± 0.0148
45	0.0389 ± 0.0038	0.0032 ± 0.0005	0.9013 ± 0.0190	0.9038 ± 0.0185
50	0.0380 ± 0.0036	0.0031 ± 0.0005	0.9053 ± 0.0192	0.9069 ± 0.0188
55	0.0385 ± 0.0037	0.0031 ± 0.0005	0.9048 ± 0.0153	0.9073 ± 0.0149
60	0.0383 ± 0.0034	0.0031 ± 0.0006	0.9037 ± 0.0208	0.9060 ± 0.0207

an average error of 0.0408 with a window size as small as 10 samples. The model also obtained high R^2 and EXP scores of 0.9074 and 0.9094, respectively for the best window size.

We use one of the iterations of the best model to illustrate its angle prediction compared to ground truth in Figure 2.9. The figure also shows a window of 40 samples of sensors' readings that correspond to a single angle prediction, test point no. 900.

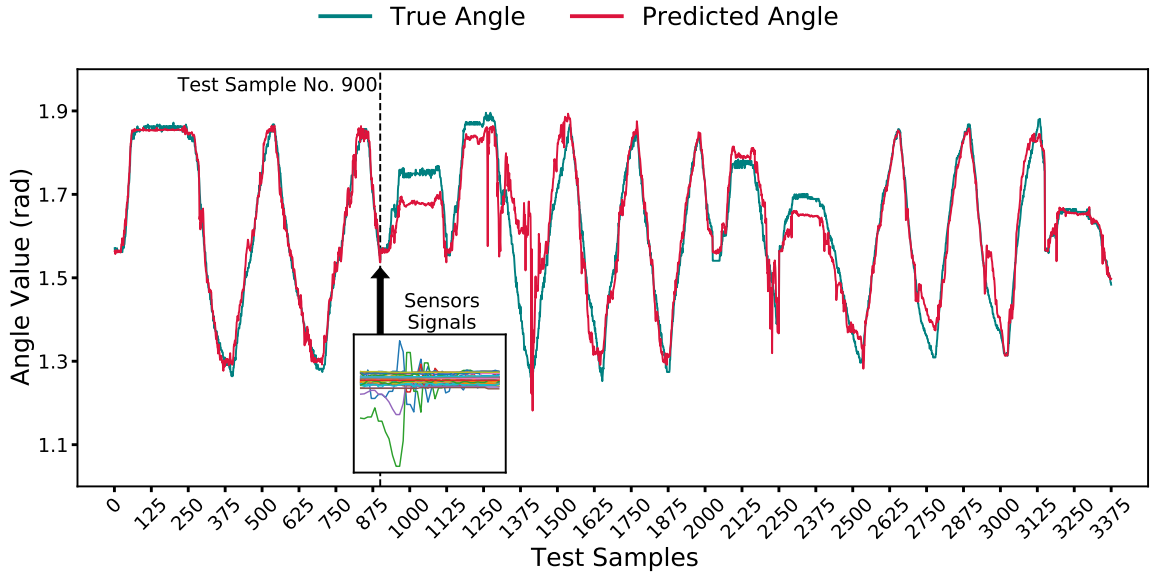


Figure 2.9: A comparison of the predicted and ground truth angle for one of the iterations of the model with the best window size of 40 samples. A 40 samples window corresponding to the prediction of angle sample 900 is highlighted. There are persistent deviations on some peaks and troughs such as around sample 1000, 1250 and 2375 which are more likely to occur when the object is static. This may occur because no new information is obtained by the sensor when the object is static and as a result the error in prediction can't be corrected.

2.4.3 Comparing this Temporal Deep Learning Method to Ridge Regression

Nevertheless, we trained linear and ridge regression models with the same data protocol we applied for the neural networks for comparison. Notably, these two classifiers present the best results in an approach that does not use the time-series relation in the data [1], in which the models were trained per object size and not using all the

sizes at once. Although we can not conclude the advantage of temporal data from the MAE and MSE values due to different normalization and scaling procedures ranges, the R^2 and EXP scores highlight that point in the previous study [1]. The results of these two models are reported in Table 2.4 using our preprocessing procedure for comparison.

Table 2.4: The results of standard regression models using MAE and MSE errors in radian, R^2 score, and EXP.

Model	MAE	MSE	R^2	EXP
Ridge Regressor	0.0677	0.0088	0.6875	0.7033
Linear Regressor	0.0678	0.0089	0.6862	0.7021

2.5 Discussion

This study aimed to determine if pose estimations relate to the time series tactile data captured by the sliding window sampling strategy adopted. We analyzed the performance of using a neural network with LSTM layers in estimating the angle of the handled object by tactile sensing robotic hand, considering different sliding window sizes of input samples. The deep learning model is compared to standard regression models to showcase the improvement due to their temporal tactile data incorporation.

We presented a data processing procedure to align collected data from multiple asynchronous sensors and approximate their reading timestamps to yield multi-sensor

temporal data in Figure 2.5. The data was then used to train and evaluate a deep learning model that we optimized its architecture, as shown in Figure 2.6, and training hyperparameters using grid search.

By testing a range of window sizes between 5 and 60 to investigate the degree of the impact of the temporal relation between tactile data, we demonstrated the importance of such relations between sensor readings in estimating the angle of an object under grasp. We found that incorporating a small window size of 5 inputs gives an acceptable performance of 0.0422 radians, equivalent to 2.417 degrees, and scores above 0.87 for both R^2 and EXP metrics. Compared to the standard classifiers tested in this study, we found that the smallest window can improve about 26% and 24% for the R^2 and EXP scores, respectively, and a reduction of 0.0256 and 0.0047 for MAE and MSE, respectively. Thus, it shows that the temporal relationships of the sensor readings can improve estimating the objects' angle as evident in Tables 2.3 and 2.4.

Furthermore, these results gradually improved by integrating more sensor information from larger window sizes of up to 40 samples per window, after which the performance saturates. Including more past readings beyond 40 samples did not add valuable information to the instantaneous angle value prediction as seen in Figure 2.8. This result shows that despite the importance of temporal relationships in tactile data for estimating the object's angle during manipulation, these relationships diminish asymptotically after a threshold.

For the best window size of 40 samples, we found that it achieves an acceptable error for many applications with an average of 0.0375 MAE in radians and can explain most of the variance in the distribution, shown in 0.9074 R^2 score and 0.9094

EXP score. This performance is sufficient in multiple applications without a camera reference during the grasp phase, thus supporting the use of temporal tactile data for orientation estimation of in-hand objects in unstructured environments.

Notably, the model can achieve such results after training on data from objects with differing sizes, thus incorporating more variation in the data, making the temporal relation harder to capture. Therefore, improving on the previous results [1] where only a per object angle estimation was performed. Additionally, using different object sizes also generalizes the model performance. This generalization also extends to being applied in an under-actuated system which experiences larger effects of parasitic motions (compared to fully actuated systems). However, this is a limited application that does not account for the other dimensions, and, as a result, future work can include all other axis and provide a complete object pose description and improve the robustness. In addition, we can not directly compare the metrics because of different normalization methods, as they use a normalized degree unit, whereas we use radians.

Future research can use our results as a reference and investigate a tactile dataset with objects of different shapes as well as remaining degrees of freedom to determine the complete change in the object’s pose, not only its yaw orientation. Moreover, feature engineering can be an additional step alongside the temporal tactile data to enhance the model further. Future studies can benefit from the proposed alignment of asynchronous sensors that we illustrated in Figure 2.5.

Finally, collecting a dataset of both arm-mounted and gripper-mounted tactile data for object orientation estimation can further illustrate the benefits of temporal tactile sensing compared to other techniques.

2.6 Conclusions

This paper illustrates the importance of temporal tactile data in estimating the orientation of in-hand objects by proposing a model architecture with LSTM layers that uses signals from tactile sensors on the fingers. We evaluated these experiments' performance using MAE, MSE, R^2 , and the EXP metrics. The results show that including temporal data benefits the orientation estimation of the objects up to an asymptotic threshold, as investigating a range of window sizes concluded that the smallest window studied boosts the performers compared to standard regression models, such as linear and ridge regression. The best window size in the investigated range is 40 input samples, which could predict the object angle with an average MAE of 0.0375 radians. Our model also has an R^2 value of 0.9074 and an EXP value of 0.9074, respectively. By comparison, the ridge regressor yields an average MAE of 0.0677 radians, 0.6875 R^2 score, and 0.7033 EXP value. Therefore, the relationship between the tactile signals object's angle is better explained with time-series models that utilize the temporal relationships of the sensors' readings. These results highlight the benefits of using previous state information, particularly because manipulation tends to be sequential. At the same time, it presents a simple architecture that uses less processing and computational power compared to setups with high-density tactile sensors. Moreover, our tactile data model can work with objects such as symmetric cylinders that may look fixed from the visual sensors' perspective. Finally, it also presents the viability of pose estimation without needing 3d models.

Our proposed model can be included in future research investigating the pose estimation problem using tactile data and the importance of their temporal relations

with different modes of pose change. Future studies can also benefit from our proposed preprocessing procedure to match the timestamps of readings obtained from asynchronous sensors.

Chapter 3

Grasp Approach Under Positional Uncertainty Using Compliant Tactile Sensing Modules and Reinforcement Learning

3.1 Introduction

Grasping objects is a fundamental yet intricate task in robotics applications ranging from industrial automation to assistive technologies. Traditional approaches to robotic grasping have predominantly relied on visual sensory information to discern the environment and guide the manipulation strategies [72]. However, visual perception is susceptible to challenges such as variable lighting conditions, object transparency, occlusion, camera calibration errors, and clutter, despite its rich detail.

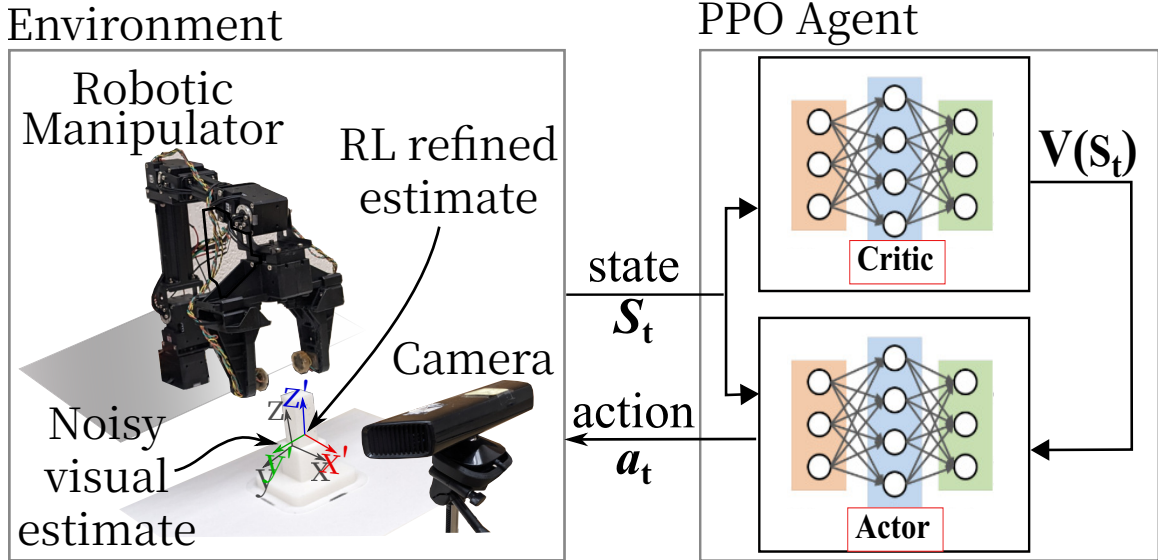


Figure 3.1: The camera provides an estimate of the object’s location. The manipulator attempts to grasp it and return the system’s state, i.e., the success/failure, the force experienced by each sensor and the orientation of the sensor to the reinforcement learning agent. The agent uses that information to update the position to attempt a grasp again if unsuccessful prior.

These factors can significantly impede the robot’s ability to accurately estimate an object’s pose, especially within unstructured environments [73, 74, 75].

Tactile sensing may offer a way to circumvent some of the limitations associated with visual perception. Tactile sensors in robotic tasks have been used in various capacities, such as for pose estimation [37, 3], texture classification [76, 77], and object recognition [78, 79]. Additionally, a wide breadth of features can be collected from them, including contact force, contact location, and contact deformation [80]. Compliant tactile sensors can deform upon contact and provide direct feedback on the robot’s interaction with its environment. This form of feedback is invaluable,

especially when visual information is unreliable or insufficient.

Moreover, the potential of compliant tactile sensors providing nuanced feedback can be amplified when integrated with perception capabilities empowered by techniques like reinforcement learning (RL). RL can enable robots to learn from interactions with their environment, refining their performance through trial and error. This learning paradigm promises to significantly enhance the efficiency and reliability of robotic grasping in the face of positional uncertainties.

Our preliminary findings demonstrate the potential of our method to significantly reduce the number of attempts required to achieve a successful grasp, highlighting the efficacy of integrating tactile feedback with reinforcement learning for adaptive grasping under uncertainty. We do this by using an estimate of the object position from sources such as cameras and use reinforcement learning to explore the vicinity of the estimate to identify and grasp the object based on the reward determined by the tactile information as shown in Fig. 3.1.

3.2 Related Work

Recent research has focused on addressing this challenge by investigating feature sensing and robotic grasping of objects with uncertain information [81]. Significant work has been done to improve tactile robotic grasping using reinforcement learning. Most methods consist only of vision-based sensing methodologies that use supervised learning, reinforcement learning, or a combination of both [82]. Also, methods like domain randomization [83] have reduced the gap from simulation to reality, although learning in simulation did not generally transfer better to reality.

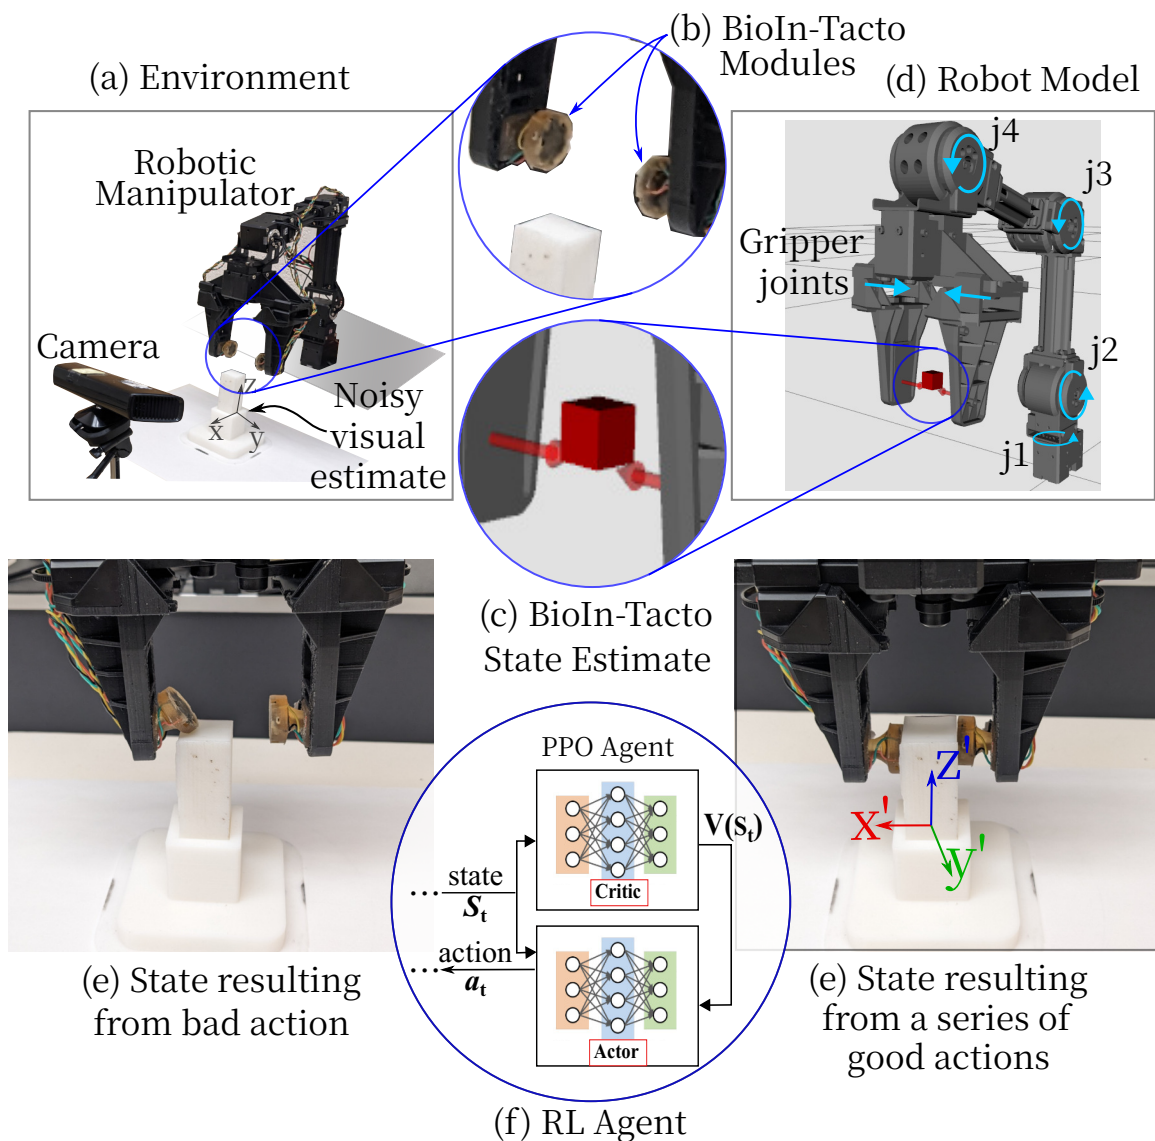


Figure 3.2: The blue arrows are the direction of the joint movement of the manipulator. The Environment provides the orientation of the sensors attached to the end effector and the pressure experienced by the sensors. This information is sent to the agent every time a grasp is attempted. The agent sends a new position for trial if unsuccessful.

Matak *et al.* [84] used a vision and tactile system for grasp planning and execution for a multi-fingered robotic hand. The tactile sensors are used to determine if the

predicted contact matches and, if not, to modify the grasp. Dong *et al.* [15, 31] use tactile sensing and reinforcement learning to insert objects, the second phase of the peg-in-the-hole task in which they explored only one dimension. They use the Gelsim [85] tactile sensor to obtain surface contact information and a force/torque sensor on a parallel jaw gripper as a baseline. Jiang *et al.* [75] devised a framework for grasping transparent objects using visual and tactile sensing. They trained a neural network to segment the image and determine the region to touch using the GelSight tactile sensor.

Due to their deformable nature, compliant tactile sensors can dynamically probe objects’ surfaces, providing rich tactile feedback that can reduce uncertainty. Alves de Oliveira *et al.* [2, 4] proposed the BioIn-Tactom, a compliant bio-inspired tactile sensing module that integrates various sensing modalities to provide rich tactile feedback for robotic systems. Da Fonseca *et al.* [79] addressed the challenge of tactile object recognition during the early phases of grasping using underactuated robotic hands, later proposing an approach that enables robots to perform manipulation tasks more effectively by inferring object orientation from tactile feedback [3].

Welyhorsky *et al.* [86] introduced a neuro-fuzzy grasp control system for a teleoperated robotic hand, enhancing dexterity and precision during grasping tasks. Zhu *et al.* [87] investigated teleoperated grasping using a robotic hand and a haptic-feedback data glove, improving the operator’s sense of touch and control over the robotic hand during manipulation tasks.

This paper aims to develop a pre-grasp approach that mitigates positional uncertainty using compliant tactile sensing modules and reinforcement learning techniques. Our approach is motivated by the need to reduce uncertainty in object perception




Figure 3.3: Tactile sensing module. Components: 1—MARG (magnetic, angular rate, and gravity) system; 2—compliant structure; 3—barometer [2, 3, 4].

and execute successful grasps even when the object’s properties are ambiguous or unknown.

3.3 Methodology

To evaluate the use of RL for contact exploration, the present work used an environment setup that consists of one 4-DoF OpenMANIPULATOR-X (model RM-X52-TNM) equipped with two BioIn-Tacto sensing modules mounted on the gripper fingers (see Fig.3.4).

3.3.1 Compliant Tactile Sensing

The pose uncertainty underscores the necessity for employing a compliant tactile sensor. To perform the experiments required, we mounted two BioIn-Tacto modules on each gripper’s jaws. Fig. 3.3 shows the BioIn-Tacto sensor used in this setup.

The tactile sensors comprise an inertial measurement unit (IMU) and a barom-

eter. We obtained the orientation of the sensor with respect to the globe using the accelerometer and the gyroscope, which is passed through the Madgwick filter.

Let: \mathbf{a} : acceleration vector, \mathbf{g} : gyroscope vector, \mathbf{q} : quaternion representing orientation, β : filter gain, Δt : time interval between sensor updates. Then, the Madgwick filter update equations are:

$$\dot{\mathbf{q}} = \frac{1}{2}\mathbf{q} \otimes \begin{bmatrix} 0 \\ \mathbf{g} \end{bmatrix} - \frac{\beta}{2}\mathbf{P}\mathbf{S}(\mathbf{q}^*)\mathbf{J}(\mathbf{q})^T\mathbf{a}, \quad (3.1)$$

$$\mathbf{q} = \mathbf{q} + \dot{\mathbf{q}}\Delta t.$$

where \otimes denotes quaternion multiplication, \mathbf{P} is a 4×4 identity matrix, and \mathbf{S} and \mathbf{J} are functions to compute the skew-symmetric matrix and Jacobian matrix, respectively. By integrating these equations, the Madgwick filter provides accurate and stable orientation estimates, enabling precise pose estimation even in the presence of noise and uncertainty. This is then transformed to obtain the sensor's orientation relative to the manipulator's base frame.

The pressure obtained is normal to the surface of the sensor. Since the sensor's surface either conforms to the object's surface, it grasps, or is distorted by a collision, the pressure determined by the sensor is proportional to the distortion it experiences in the normal direction to the surface. This pressure is normalized to the initial reading obtained during the home position so that the observation space is of the same magnitude.

The MARG sensor data and manipulator inverse kinematics determine the real-time contact normal orientation and position. The sensing module's flexible structure provides a frame for the estimation of contact orientation and position at each contact

tentative without requiring a planning approach.

3.3.2 Reinforcement Learning

The present paper uses an exploratory approach, leveraging the compliant sensor structure to evaluate approach attempts and contact. The robot assumes an initial grasp position that has a level of uncertainty (see Section 3.3.3). Every each approach tentative in case of contact a RL strategy is updated.

The robotic agent controls this process and chooses actions implementing a Proximal Policy Optimization (PPO) as presented in Algorithm 2. The network consists of 2 fully connected layers with 64 units each.

In the algorithm, each approach tentative starts with the *reset()* function, taking the agent back to an initial state, which is set to be above the target object at a position p_i given by the uncertainty model \mathbb{U} .

The initial observations are passed to the *step()* function, which starts looping while the *done* signal is set to False. Once the done signal changes to True, the episode ends, sequences of states and cumulative rewards are stored, and a new episode can begin. Once n_{steps} reaches *total timesteps*, the iteration is completed, and the policy is updated using PPO.

Each training iteration consists of several episodes, which vary depending on how soon the agent reaches a terminal state. We choose to work with 2048 steps per training iteration, meaning that the policy will be updated after 2048 steps have been executed over the number of episodes necessary.

Algorithm 2 is defined around five main elements: the action and observation

Algorithm 2 Approach tentative iteration

Require: $obs; p_i \in \mathbb{U}; total_timesteps$.

```
1:  $n \leftarrow 0$ 
2:  $done = False$ 
3: while  $n\_steps \leq total\_timesteps$  do ▷ Iterate episodes.
4:    $obs_{ini} \leftarrow Reset(\mathbb{U})$ 
5:    $obs_t \leftarrow obs_{ini}$ 
6:   while  $done = False$  do ▷ Iterate steps.
7:      $obs_{t+1}, r_{t+1}, done \leftarrow Step(obs_t)$ 
8:      $n\_steps \leftarrow n\_steps + 1$ 
9: Update policy with PPO.
```

spaces, the *reset()* and *step()* methods, and the reward function. The *action* space contains all actions that the agent can execute. In this case, our action is a single value $a \in [-1; 1]$ for the base actuator z-axis.

The observation space *obs* is composed of joint efforts $effs \in \mathbb{R}^4$ in *N.m*, barometer level $baro \in \mathbb{N}$, and orientation $orient \in \mathbb{R}^3$. It describes the environment to the agent and is used to compute actions.

The *reset()* method is responsible for taking the agent back to an initial valid state where the agent receives the initial observations obs_{ini} from the sensing module. In the present work, the robot is performing a fine adjustment of the approach position, which implies that the manipulator is above the target object. This means that, during this phase, it first retreats to an initial pose called *home*, making no contact with the object. A point p_i is sampled from the uncertainty model \mathbb{U} , and

the manipulator moves to the given location, hovering above the object. Then, the initial observations obs_{ini} are sent to the $step()$ function.

The $step()$ method, as shown in Algorithm 3, is responsible for sending the state S_t to the policy through the observations obs_t and getting an action a_t back, executing it and calculating the reward r_t based on the resulting observations. It is also responsible for verifying that the agent has reached its goal and ending the training episode. We condition the end of the episode to happen under two scenarios: 1) the number of steps for an episode reaches $n_{steps} = 50$ or 2) the barometer level and the MARG pose variations indicate an excess of contact pressure or wrong contact (see Eq. 3.2). This allows the manipulator to keep changing the initial position p_i at the end of every episode for generalization purposes and also avoid catastrophic contact with the target object.

Algorithm 3 Step

Require: $obs; p_i \in \mathbb{U}$, where $p_i = \theta_z$

- 1: $orient, effs, baro \leftarrow obs$ ▷ Real-time input.
 - 2: $a \leftarrow Policy(obs)$ ▷ $a \in [-1; 1]$
 - 3: $pos_{new,i} \leftarrow p_i + orient \times step \times a$ ▷ $step = 0.005m$
 - 4: Move manipulator to pos_{new} ▷ Eq. 3.3.
 - 5: $reward \leftarrow f(sensors)$. ▷ Eq. 3.2.
 - 6: Checks if the terminal state has been reached.
-

Equation 3.2 shows the reward function $f(sensors)$ is a function of the MARG and barometer levels where (a) if there was a significant change in the orientation angle of the sensors (indicating a collision); (b) negative if there was no grasp; (c) a

smaller magnitude of negative reward if there was too much pressure from one sensor but not from the other since it was close to grasp but offset; (d) finally, positive on a successful grasp where the object can be lifted and placed without any slippage.

$$f(sensors) = \begin{cases} -0.1 & \text{if } orien > threshold \text{ (a)} \\ -0.1 & \text{if } baro < threshold \text{ (b)} \\ -0.1 & baro1 > thres \oplus baro2 > thres \text{ (c)} \\ +0.5 & baro1 \oplus baro2 > threshold \text{ (d)} \end{cases} \quad (3.2)$$

3.3.3 Object Pose Uncertainty

We introduced an error term to the measurement to simulate the noise inherent in the pose estimation of an object, specifically along the x-axis. This error term follows a normal distribution characterized by its mean (μ) and standard deviation (σ). The normal distribution equation for the x-axis error is given by:

$$f(x|\mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (3.3)$$

where: $f(x|\mu, \sigma)$ is the probability density function, μ is the mean of the distribution along the x-axis, σ is the standard deviation along the x-axis, and x is the variable of interest along the x-axis.

By incorporating this error term into our measurements along the x-axis, we are able to mimic a $0.02m$ variability and uncertainty present in real-world pose estimation scenarios.

We used a four-degree-of-freedom robotic manipulator with two sets of sensors

attached on each side of the end effector and a 3D printed cuboid for the object to be grasped, as shown in Fig 3.4.

3.4 Experimental Setup

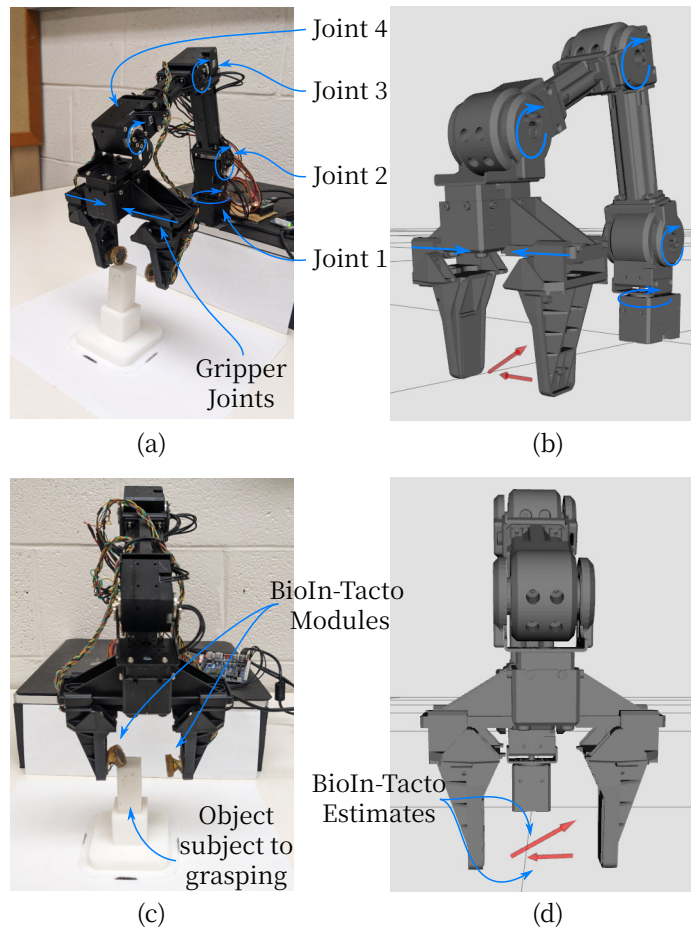


Figure 3.4: The manipulator attempts to grasp the white cuboid by exploring the area within the vicinity of the estimated location provided.

The environment, that is, the pressure experienced by the and orientation of the sensors as well as the pseudo-position of the object to be grasped are provided to the

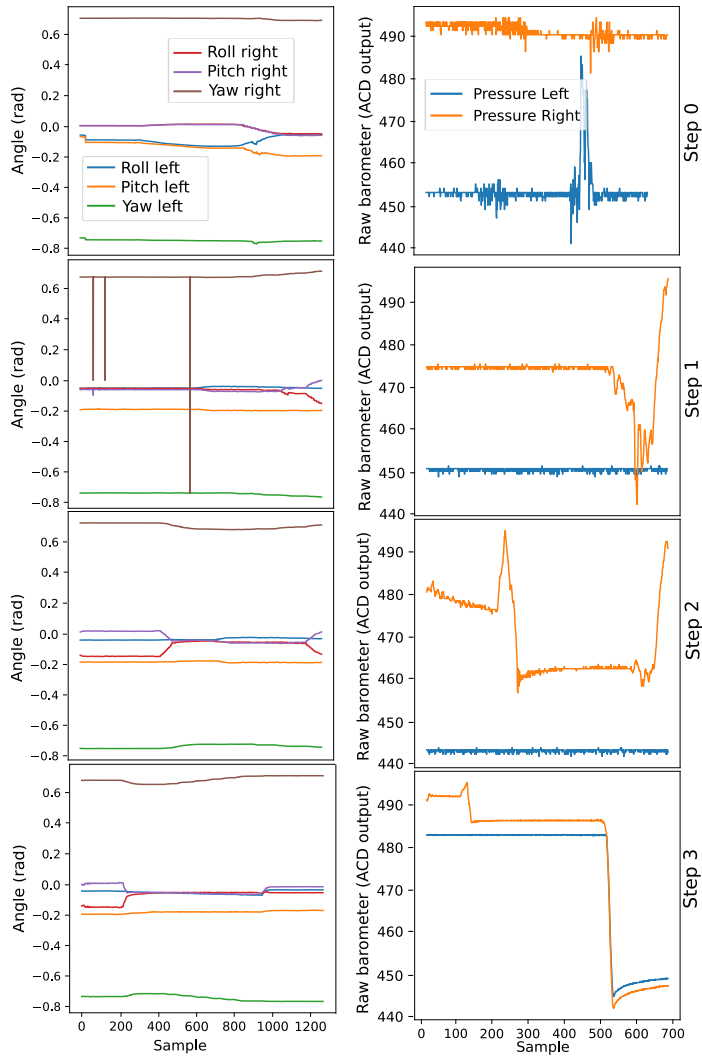


Figure 3.5: The pressure and the orientation of each sensor on the end effector during each step in the episode.

reinforcement learning agent, the agent then provides the manipulator the estimated position of the object. Subsequently the results of the attempted grasp based on the estimate are provided to the agent. If successful, the environment is reset. Conversely, if unsuccessful the agent tries to estimate again and the cycle repeats as shown in Algorithm 4.

Fig. 3.5 shows a graph of each sensor output during each step on an episode. The first row is the attempt at trying to grasp the object but doesn't interact with it at all. For steps 1 and 2, the right sensor experienced a collision with the object and that provided feedback that could be used to correct the orientation. Finally on step 3, the object is grasped successfully.

Algorithm 4 Grasp exploration

```
1: for each episode do:
2:   Move manipulator to the home position
3:   for each step do:
4:     attempt grasp using position from agent
5:     calculate reward
6:     update policy
7:   if successful grasp then:
8:     end episode
```

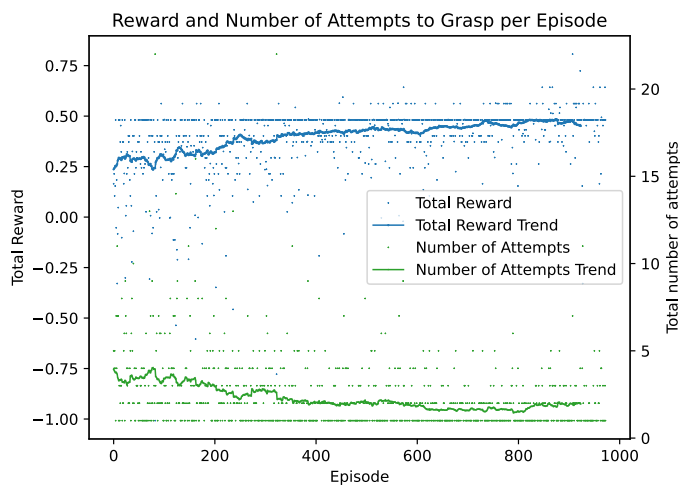


Figure 3.6: The trend line is calculated with a 50 moving average. The trend line of the reward increasing with the episodes and the number of steps per episode reducing. As the number of episodes increase the rewards and number of steps reach an asymptote.

3.5 Results

The results presented here show our evaluation of the learning rate when the grasp approach is affected by uncertainty. Fig. 3.5 shows an example of the agent learning behavior of our system, which corresponds to a positive reward in the last step in which there was a successful grasp and negative rewards for the rest.

The agent learning effect can also be observed in Fig. 3.6, which shows the total reward increasing while proportioning the number of attempts to grasp reducing successfully. The 50-episode moving average trend shows that as the model iterates, it could grasp the object with fewer attempts and, consequently, in a shorter time span on average despite the number of attempts varying from one attempt to the next. Although more than 20 steps were taken at some points, most of the episodes of high density occurred at the start, when the model had the fewest iterations.

However, we still start at an average of 5 steps, and it drops to an average of close to 2 steps. As the number of episodes increases, the density of the number of steps is mainly below the 5-step line, showing that the system improves the search pattern, minimizing the number of attempts to grasp.

3.6 Conclusion and Future Work

Our investigation introduces a method that leverages compliant tactile sensors and reinforcement learning to enhance the robotic grasping approach under object pose uncertainty. The results underscore the substantial potential of this approach, demonstrating a remarkable improvement in grasp accuracy and efficiency and a notable reduction in the number of attempts needed for a successful grasp. By fusing tactile feedback with adaptive learning, we significantly advance robotic capabilities, paving the way for systems that can more effectively comprehend and interact with their environment. Future work will involve applying the method to various objects with different materials, shapes, and scenarios to assess its versatility and effectiveness.

Chapter 4

Human Example Pretraining

Influences on Reinforcement

Learning for Obscured Object

Trajectory Guidance

4.1 Introduction

With the widespread use of robotics in assembly lines, the push for more general-purpose robots has become reinvigorated over the last decade. However, there are open questions that need to be addressed to substitute human efficacy in tasks like unscrewing, plugging and unplugging cables, and machine disassembly for repair in unstructured environments. In particular, understanding the environment and performing as well as humans do. The human ability to use the multitude of senses

together or as needed provides much-needed flexibility in trajectory identification of visually obscured objects. Aspects of this challenge could be reduced to the peg-in-hole problem, which has been a widely researched assembly task [88, 89, 90].

Insertion and extraction of peg-in-hole is a stepping stone to a variety of manipulation tasks that require local exploration of the environment. Especially for fixed cameras where the targets are occluded by other objects. This exploration provides information on object interaction boundaries between the target and objects in the environment. Tactile information can take the reigns in such instances. Many works exploring the peg-in-hole problem tend to use some form of contact sensing capability ranging from force/torque [91, 92] to vision [36]. Apart from giving information that the camera cannot obtain from occlusion, it can also provide information that is not available using a camera, and this information can be leveraged to make localized corrections. This is especially true for compliant grippers for which the dynamics of the hand can not be predicted based on the known geometry and kinematics distorted by unknown object properties.

Sensors such as e-skin [24], visiflex [25], GelSight Fin Ray [26], and BioIn-Tacto [4] are examples of the foray into compliment tactile sensors that are being developed for the characteristics of the human fingertips that can conform to the shape of the object being grasped and also give local information not distorted by the conformation. These developments benefit manipulation tasks that rely on fine motor movements with environmental uncertainty to determine contact states. These contact states can guide reinforcement learning algorithms [93].

Lu *et al.* [94] presents a human-centric model of manufacturing in which humans work with emphatic machines in a symbiotic relationship, with one of its fundamental

components being human-machine understanding and, in particular, understanding the instruction, action and goal. One of the challenges in explaining the process of achieving the goal is its impreciseness with the variability of the environment and the task. One method of reducing impreciseness while reducing the training space is by human examples. In this section, we explore the effects of pretraining the machine learning model using human-based imperfect examples in speeding up the ability of the agent to infer tactile signals and consequently speed up the success in object extraction in peg-in-hole tasks using reinforcement learning. Viability in this domain can give reason to explore similar techniques in applications like prosthesis control, where pretraining is commonly done for everyday prosthesis use.

4.2 Literature Review

As robots take over more human tasks, the need for effective human-robot knowledge transfer increases. Lee *et al.* [95] developed a human-robot knowledge transfer framework. Humans and robots have different modes of perception, communication and reasoning. Therefore, knowledge has to be transferred using learning, comprehension and translation. For dis/assembly tasks, they categorize assembly knowledge into assembly design, plan, and action. They suggested that an agent can consider a demonstration as a knowledge transfer using the assembly plan.

Using interactions of the manipulator with the object and the object with the environment is one strategy used to extrapolate local information and consequently increase the likelihood of a successful peg-in-hole insertion as done by Kim *et al.* [36]. They used a vision-based compliment tactile sensor attached to the phalanges to

develop a reinforcement policy based on the estimate of the misalignment of the first attempt. The tactile sensor provides a deformation image that is used to determine the deformation angle of the sensor based on a trained convolution neural network. They use the tactile inferred angle and proprioception as factors that they explore to determine the contact line of the object and the hole boundary. This is used to determine the following action for insertion. They use a tactile feedback controller to move the robot down and then rock or pivot the object during the exploration phase to determine the contact line. More than 95% of their attempts are successful in all six object and hole shape combinations in which the clearance of the object and hole was an average of 2.25mm.

Dong *et al.* [15] explore methods to align the object and the hole in peg-in-holes tasks without using prior knowledge of the object or the shape or contact locations. They compare supervised and reinforcement learning, curriculum training, and the tactile representation of a light-based tactile sensor, comparing tactile flow to RGB representation. They determined the RGB sensors with tactile flow representation and curriculum training reinforcement learning policy performed the best compared to either when everything else was constant and

- force/torque is used instead of a vision-based tactile sensor,
- curriculum training is not used,
- RGB representation is used instead of tactile flow.

Simulation to real-life transfer is another commonly used practice used to leverage the generally needed high amount of episodes required for reinforcement learning

[96, 97, 98, 99, 77]. Xie *et al.* [100] use simulation to real-life transfer learning with human demonstration for new objects to reduce the transfer cost. They use an RGB camera to segment the peg and the hole using human guidance to create masks to train a segmentation network. Subsequently, they use a deep neural network to determine the position and rotation features of the segmentation masks. Finally, they use an Advantage Actor-Critic network to create an insertion policy based on the rotation and position features. They take into account the features of occlusion using the Long Short-Term Memory network to account for the history that is otherwise limited by a single frame during occlusion. They obtained an accuracy of greater than 90% in both, simulation and real-world applications for hole tolerances less than 1mm. They addressed the occlusion problem by considering the time series of the frames. The best result for unseen objects was 65.6%, which they suggest remains a big challenge, especially in dynamic environments.

This challenge of occlusion is significantly amplified in object extraction since there is no initial bearing to extrapolate from either. Therefore, we explore using tactile sensing to determine the trajectory the manipulator needs to follow to remove the peg entirely from a hole.

4.3 Materials and Methods

We use a four-degree-of-freedom manipulator with two tactile sensors attached to each phalanx to perform object extraction from a peg-in-hole setting where most of the peg is occluded. The method involves three stages, as shown in figure 4.1. First, we perform teleoperation to perform the peg-in-hole object extraction task, recording




Figure 4.1: The process is broken down into three steps. We first record the actions and the environment using teleoperation for object extraction in the peg-in-hole task. We use the recorded data to train a neural network. We use the trained neural network as a foundation for the Q-learning reinforcement learning setup for the object extraction.

the inputs and behavior. Next, we train a neural network based on the data collected during teleoperation. Finally, we use a reinforcement learning agent with the trained neural network to perform object extraction.

4.3.1 Sensor and Manipulator Setup

We use a four-degree-of-freedom robotic manipulator with the BioIn-Tacto compliant tactile sensor attached to each phalanx of the manipulator as shown in 4.2. Due to the four-degree freedom limitation, the manipulator could only move at any point in the XZ and YZ plane but only on a curved trajectory in the XY plane with respect to the end-effector frame while maintaining the orientation. The manipulator can be controlled linearly in the X and Z directions and can rotate in the Y and Z axis of the end effector frame \mathbb{E} . If the new end effector pose \mathbb{F} , is the current end-effector




Figure 4.2: This is the manipulator and sensor setup we are using for this experiment. The four degrees of freedom manipulator with two tactile sensors, one attached at the end of each phalanx. The red arrows represent the joint motions for the four degrees of freedom manipulator. The manipulator is controlled with respect to the end effector frame \mathbb{E} in the directions shown with the black, green, blue and white arrows. The positions with respect to the base frames \mathbb{O} are converted to the positions with respect to the end-effector frame \mathbb{E} . Resulting in motions on the x, y and z axis as well as rotation about the y and z axis with respect to \mathbb{E} . We also obtained the position of the tactile with respect to \mathbb{O} sensor based on the data it provided.

pose that shifted in only in either the individual X, Y, or Z directions or rotated only in the Y or Z direction.

The tactile sensor provides magnetic field, angular velocity, linear acceleration in the three dimensions, and pressure perpendicular to the surface of the sensor. We obtain the orientation of the sensor using the angular velocity and linear acceleration.

We obtain the orientation using the Madgwick filter using equation 3.1. We use the scaled pressure, the end effector pose and the change in orientation in quaternion for each tactile sensor. The pressure reading is scaled in a range of 0 – 1 from 0 – 400 by dividing by 400, and the change in orientation is obtained using

$$\mathbf{Q}_{\Delta} = \mathbf{Q}_t \otimes \mathbf{Q}_{t-1}^{-1} \quad (4.1)$$

Where \otimes denotes quaternion multiplication, \mathbf{Q}_{t-1} is the orientation of the sensor at $t-1$ in quaternion, \mathbf{Q}_t is the orientation of the sensor at t in quaternion and \mathbf{Q}_{Δ} represents the change in orientation between time t and $t-1$ to prevent the distortion of error accumulation from the integrator since the most relevant information is based only on the recent information.

4.3.2 3D printed Object Set used for Extraction




Figure 4.3: These are pegs and the respective holes used in this peg-in-hole setup. **Top right:** a slanted, curved and vertical peg. **Left:** the respective holes for each of the pegs. **Bottom right:** the cross-section of respective holes.

We 3D printed the three objects for the pulling task with varying levels of complexity to compare the cross-compatibility of training. A vertical peg, a slanted peg and a curved peg with its respective holes as shown in figure 4.3.

4.3.3 Teleoperation and Data collection for Pretraining

The first step is collecting data using teleoperation. Here, we configure the manipulator to be controlled by a user using the keyboard. Peg one was attempted to be picked three times, since it is symmetric. Whereas pegs two and three are picked three times in each orientation of 0° , 45° , 90° and 135° . The user pressed the keys on the keyboard corresponding to the direction in which they wanted to move the end-effector to guide the peg out of the hole. The corresponding input, the end-effector pose, the pressure on each tactile sensor, and the orientation of each tactile sensor are recorded.

4.3.4 Pretraining the Neural Network

The user keyboard inputs are encoded using a one-hot encoder. We use a three dense layer deep neural network and a final softmax layer. Each layer contained 64, 32, and 16 units, respectively. It was trained based on the data collected using teleoperation. We explore performance on no pretraining, pretraining with the same object, training with a different object, training with the other two objects and pretraining with all three objects.

The pertaining is used to reduce the amount of time it takes for the early sessions of reinforcement learning sessions. Since it is governed by the randomness of the initial Q-network and the learning rate α , the discount factor γ and the window of exploration required based on the action space in the bellman equation:

$$Q(s, a) = (1 - \epsilon)Q(s, a) + \alpha \left[r + \gamma \max_{a'} Q(s', a') - Q(s, a) \right] \quad (4.2)$$

Where s is the state, a is the action, ϵ is the exploration parameter, r is the immediate

reward, s' is the next state, and a' is the next action.

4.3.5 Object extraction using Reinforcement Learning

We use the deep Q-learning reinforcement learning algorithm [101] due to the large and continuous state space. All teleoperation actions took less than thirty-two steps to complete, so a batch size of thirty-two steps was chosen. The network is updated every four steps since the task can be completed in less than thirty-two steps based on the teleportation, and similarly updating the target network is every eight steps. The discount factor was 0.75 since the previous information was not very significant. We increased the exploration capability by setting the epsilon greedy parameter to 200 for the session in which we did not use any pretrained network so that the agent could explore more, whereas, for the pre-trained network, we used the greedy parameter of 50.

Equation 4.3 shows that the reward *reward* is a function of the barometer levels and height of the end-effector where (a) if there was a large change in pressure indicating there was a lot of friction experienced from the walls of the base if there was a large change in the pressure of the sensors (indicating a collision); (b) a function of the change in height of the end effector if the friction was not large; (c) finally, positive reaching the designated height.

$$reward = \begin{cases} -0.5 & \text{if } threshold\ pressure < baros \text{ (a)} \\ f(\Delta end - effector\ height) & \text{if } baros < threshold\ pressure \text{ (b)} \\ +1 & \text{end} - effector \text{ if } height > threshold\ height \text{ (c)} \end{cases} \quad (4.3)$$

4.4 Results

The results shown here explore the effects of tactile information used for object extraction in the peg-in-hole task for which the training phase can be sped up with imperfect human examples. In particular, whether pretraining reduces the training episodes needed, the effects of training time on pretraining using other objects, and the effects of pretraining on the complexity of the trajectory required.

Figure 4.4 shows a variety of behaviour of the manipulator based on the environment, object, and the pretraining methodology. Object one, pre-trained on object two, took only three episodes, whereas object three, pretraining on object three, took more than 150 episodes, with a series of episodes having the same number of steps and rewards when the updated network has an equally bad next step.

The agent successfully learnt to extract all three pegs using just tactile information. Without pretraining, it took an average of 178 episodes to complete the extraction of the object. This is substantially improved upon as shown in Figure 4.5 with pretraining. It took an average of 48 episodes with pretraining on all objects. Significantly fewer episodes were required to obtain a successful extraction after pretraining. Some outliers can be attributed to the oscillation from fewer random motions which

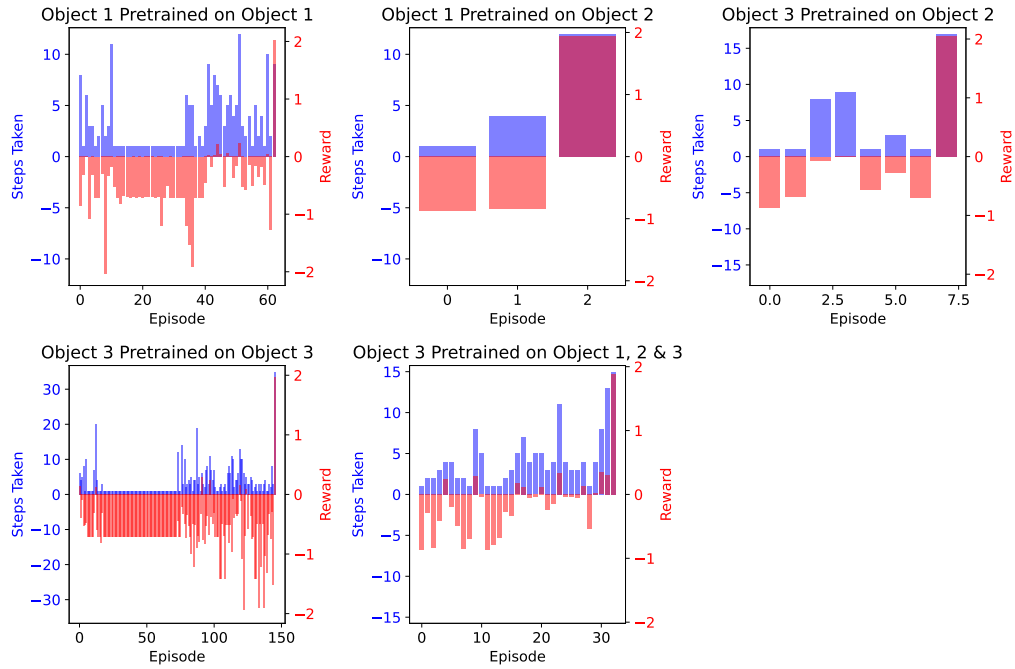


Figure 4.4: This figure shows the reward that corresponds to the steps for each episode for some trials until the last episode, where the attempt was successful. Object one pretrained on object one and took more than 60 episodes.

may result from a local trough that is cyclical.

Figure 4.6 shows that the effects of pretraining on all three objects were substantial. However, for the complicated objects, there were some attempts that required more episodes after pretraining compared to no pretraining. This could be explained by the fact that pretraining on another object and reducing the exploratory phase results in a longer recovery period from learning the wrong actions.

Figure 4.7 shows the effects of pretraining on different objects. For object one, the type of pretraining did not affect the number of episodes as much as it did for objects two and three. Since object one was simpler than objects two and three,

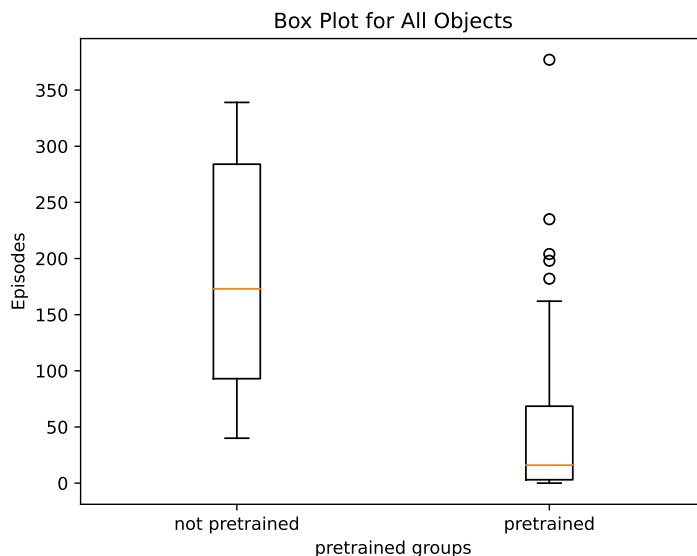


Figure 4.5: This figure shows a comparison between the number of episodes required to complete an extraction on all objects with pretraining or no pretraining.



Figure 4.6: This figure shows a comparison between box plots of the number of episodes taken for a successful extraction with any form of pretraining and no pretraining for each object.

the training in any of the objects may have captured the behaviour required for a successful episode more easily. For object two, pretraining on objects one and three together had a smaller range of the total number of episodes required to obtain a

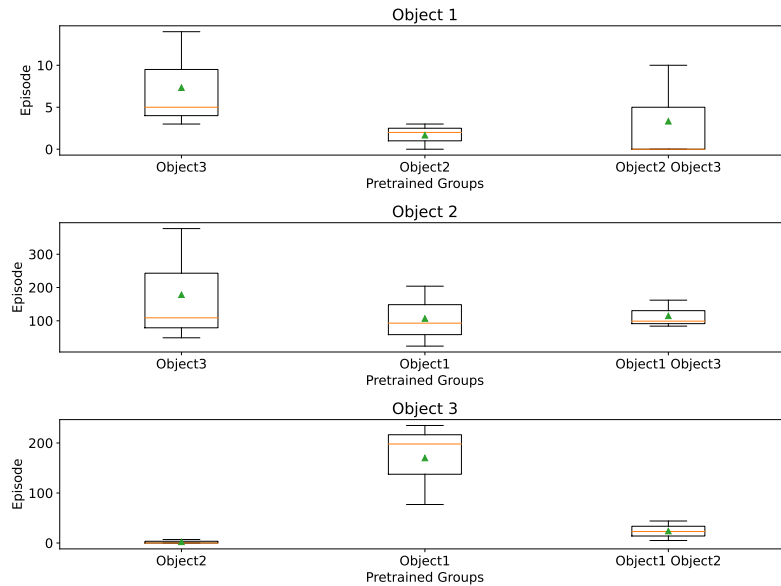


Figure 4.7: This figure shows box plots for the number of episodes required to successfully extract the object based on the pretraining on other objects. This explores the generalizability of the pretraining.

successful extraction. However, pretraining on object 3 had a worse average and upper bound, this could be because there was less transferability from object two knowledge compared to object one and may have distorted the behaviour corresponding to the training. On the other hand, pretraining on object two performed substantially better than pretraining on object one to extract object three.

4.5 Conclusion and Future works

In general, pretraining resulted in fewer episodes being required to successfully extract an object. However, the type of pretraining also influenced the number of episodes. In particular, more complicated trajectories were sensitive to the limits of pretraining on

a simple trajectory. Moreover, the pretraining on more complicated objects assisted the simpler objects to be extracted faster.

Future works should explore more complicated extraction objects to determine the extent of pretraining and to identify the dimension of pretraining required to cover a large subset of instances that can be experienced by the manipulator.

A holistic system should also be explored for insertion and extraction tasks combined to determine the crossover between the two tasks and if prior knowledge in insertion will assist in extraction.

Chapter 5

Conclusion and Future Work

The work presented in this thesis explores the function of tactile sensing in robotic manipulation. In Chapter 2, we use the data from the tactile sensors to estimate the object’s pose. In Chapter 3 we use feedback from tactile sensors to correct for the deviance in object pose estimates obtained from the camera during grasp and in Chapter 4 we use a neural network pretrained on human examples and tactile feedback for a Q-learning agent to extract a peg from a hole as a form of disassembly.

We demonstrated that tactile sensors can be beneficial in various aspects of robotic manipulation. They can either support vision based sensors or provide data that vision based sensors can not. The sensors in the phalanges provide localized data at the point of contact with the objects being interfaced. Reinforcement learning is also useful in adapting to dynamic environments and uncertainty.

In Chapter 2, we evaluated the effectiveness of a sliding window sampling strategy for pose estimation using tactile data collected from an underactuated gripper equipped with multiple sensors. By employing an LSTM network, we leveraged the

temporal nature of the data to improve angle prediction accuracy for objects under grasp manipulation.

The results demonstrated that the use of temporal data from the tactile sensors provided the lowest error. In particular, the window size of 40 samples captured sufficient tactile information to yield the lowest MAE of 0.0375 radians. This performance was consistent across various metrics, including MSE, R^2 , and EXP. Additionally, our model outperformed traditional regression models, such as ridge and linear regression, which achieved significantly lower R^2 and EXP scores, highlighting the advantage of utilizing temporal data for pose estimation.

The successful integration of different sensor readings, synchronized and down-sampled to match the camera’s frame rate, was crucial for our LSTM model’s performance. The preprocessing steps ensured that the diverse sensor data were aligned correctly, facilitating effective model training and evaluation.

This study’s findings underscore the potential of using LSTM networks for time-series analysis in tactile sensing applications, particularly in robotics. Future work could explore extending this approach to more complex and varied object shapes and further incorporating additional sensors to enhance pose estimation’s robustness and accuracy.

In Chapter 3 we develop a system to reduce the the number of grasp attempts based on the uncertainty of the object pose obtained from cameras. We do this by using reinforcement learning to and tactile sensing.

The tactile sensors obtain local information based on each grasp attempt, which then helps the reinforcement learning agent update its knowledge of the environment. These attempts can help update the policy and get an improved estimate of the

uncertainty distribution, which could enable the grasp attempt to be more likely to be successful.

The total number of grasp attempts tended towards 2 steps from an average of 5 steps after training. This approach enables the manipulator to improve its performance over time, minimizing the number of attempts needed to achieve a successful grasp, even in pose uncertainty. The other benefit of this strategy is the ability to adapt to the changing uncertainty due to environmental factors.

Finally, after obtaining a grasp, we attempt to extract the object in the peg-in-hole problem in 4. We use three different objects, a vertical, a slanted, and a curved peg, to have variability in the trajectory and to determine the transferability and adaptability of the system. The methodology is divided into three main stages: data collection through teleoperation, pre-training a neural network using the teleoperation data, and applying a reinforcement learning (RL) agent with the pre-trained network to perform object extraction.

The pre-training stage uses a dense layer deep neural network to encode user inputs and sensor data, significantly reducing the time required for the RL agent to learn the task. The RL agent employs a deep Q-learning algorithm, optimizing its actions based on a reward function that accounts for pressure changes and end-effector height. Experiments with different peg shapes demonstrate the benefits of pre-training, showing that it reduces the number of episodes needed for successful extraction and improves the generalizability of the learned behaviors across various objects.

Future works should explore joining the in hand pose estimation, grasping and object extraction into one continuous pipeline. It would benefit from shared feature

information that can be useful in various combinations of sub-tasks.

Moreover, increasing the dimensions and degrees of freedom of the manipulator used would be useful in identifying the complexities of scaling.

Finally, exploring real life applications such as USB connector insertions would be illustrate the feasibility of this work.

Bibliography

- [1] Vinicius Prado da Fonseca, Thiago Eustaquio Alves de Oliveira, and Emil M. Petriu. the orientation of objects from tactile sensing data using machine learning methods and visual frames of reference. *Sensors (Basel, Switzerland)*, 19, 5 2019.
- [2] Thiago Eustaquio Alves de Oliveira, Ana-Maria Cretu, and Emil M. Petriu. Multimodal bio-inspired tactile sensing module. *IEEE Sensors Journal*, 17(11):3231–3243, 2017.
- [3] Vinicius Prado da Fonseca, Thiago Eustaquio Alves de Oliveira, and Emil M. Petriu. Estimating the orientation of objects from tactile sensing data using machine learning methods and visual frames of reference. *Sensors*, 19(10):2285, 2019.
- [4] Thiago Eustaquio Alves de Oliveira and Vinicius Prado da Fonseca. Bioin-tacto: A compliant multi-modal tactile sensing module for robotic tasks. *HardwareX*, 16:e00478, 2023.
- [5] A. Gasparetto and L. Scalera. A brief history of industrial robotics in the 20th century. *Advances in historical studies*, 8(1):24–35, 2019.

- [6] Christos Iavazzo, Xanthi-Ekaterini D Gkegke, Paraskevi-Evangelia Iavazzo, and Ioannis D Gkegkes. Evolution of robots throughout history from hephaestus to da vinci robot. *Acta medico-historica adriatica: AMHA*, 12(2):247–258, 2014.
- [7] Irati Zamalloa, Risto Kojcev, Alejandro Hernández, Inigo Muguruza, Lander Usategui, Asier Bilbao, and Victor Mayoral. Dissecting robotics-historical overview and future perspectives. *arXiv preprint arXiv:1704.08617*, 2017.
- [8] Ravi Kumar, Jørgen Dale, Jayant Singh, and Jing Zhou. Design and development of an anthropomorphic gripper for service robotics and prosthetic applications. In *2023 11th International Conference on Control, Mechatronics and Automation (ICCMA)*, pages 233–238. IEEE, 2023.
- [9] Rongrong Liu, Florent Nageotte, Philippe Zanne, Michel de Mathelin, and Birgitta Dresch-Langley. Deep reinforcement learning for the control of robotic manipulation: A focussed mini-review. *Robotics*, 10(1), 2021.
- [10] Weiqiang Dou, Guoliang Zhong, Jinglin Cao, Zhun Shi, Bowen Peng, and Liangzhong Jiang. Soft robotic manipulators: Designs, actuation, stiffness tuning, and sensing. *Advanced Materials Technologies*, 6(9):2100018, 2021.
- [11] Hannah B Helbig and Marc O Ernst. Haptic perception in interaction with other senses. In *Human haptic perception: Basics and applications*, pages 235–249. Springer, 2008.
- [12] Shuang Li, Jiayi Jiang, Philipp Ruppel, Hongzhuo Liang, Xiaojian Ma, Norman Hendrich, Fuchun Sun, and Jianwei Zhang. A mobile robot hand-arm teleoper-

- ation system by vision and imu. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10900–10906, 2020.
- [13] Mary B. Alatiser and Gerhard P. Hancke. A review on challenges of autonomous mobile robot and sensor fusion methods. *IEEE Access*, 8:39830–39846, 2020.
- [14] Shaoxiong Wang, Yu She, Branden Romero, and Edward Adelson. Gelsight wedge: Measuring high-resolution 3d contact geometry with a compact robot finger. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6468–6475. IEEE, 2021.
- [15] Siyuan Dong, Devesh K Jha, Diego Romeres, Sangwoon Kim, Daniel Nikovski, and Alberto Rodriguez. Tactile-rl for insertion: Generalization to objects of unknown geometry. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6437–6443. IEEE, 2021.
- [16] Guoguang Du, Kai Wang, Shiguo Lian, and Kaiyong Zhao. Vision-based robotic grasping from object localization, object pose estimation to grasp estimation for parallel grippers: a review. *The Artificial intelligence review*, 54(3):1677–1734, 2021.
- [17] Andrea Sipos and Nima Fazeli. Simultaneous contact location and object pose estimation using proprioception and tactile feedback. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3233–3240, 2022.
- [18] Yiming Li, Tao Kong, Ruihang Chu, Yifeng Li, Peng Wang, and Lei Li. Simultaneous semantic and collision learning for 6-dof grasp pose estimation. In

2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 3571–3578, 2021.

- [19] Felix von Drigalski, Shohei Taniguchi, Robert Lee, Takamitsu Matsubara, Masashi Hamaya, Kazutoshi Tanaka, and Yoshihisa Ijiri. Contact-based in-hand pose estimation using bayesian state estimation and particle filtering. 05 2020.
- [20] Maria Bauza, Antonia Bronars, and Alberto Rodriguez. Tac2pose: Tactile object pose estimation from the first touch. *The International Journal of Robotics Research*, 42(13):1185–1209, 2023.
- [21] Sanghoon Yang, Won Dong Kim, Hyunkyu Park, Seojung Min, Hyonyoung Han, and Jung Kim. In-hand object classification and pose estimation with sim-to-real tactile transfer for robotic manipulation. *IEEE Robotics and Automation Letters*, 9(1):659–666, 2024.
- [22] Tarik Kelestemur, Robert Platt, and Taskin Padir. Tactile pose estimation and policy learning for unknown object manipulation. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 742–750, 2022.
- [23] Leif P. Jentoft, Qian Wan, and Robert D. Howe. Limits to compliance and the role of tactile sensing in grasping. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6394–6399, 2014.
- [24] Wennan Xiong, Hui Feng, Haosen Liwang, Dan Li, Wanbing Yao, Dilinazha Duolikun, Yunlei Zhou, and YongAn Huang. Multifunctional tactile feedbacks

- towards compliant robot manipulations via 3d-shaped electronic skin. *IEEE Sensors Journal*, 22(9):9046–9056, 2022.
- [25] Alfonso J. Fernandez, Huan Weng, Paul B. Umbanhowar, and Kevin M. Lynch. Visiflex: A low-cost compliant tactile fingertip for force, torque, and contact sensing. *IEEE Robotics and Automation Letters*, 6(2):3009–3016, 2021.
- [26] Sandra Q. Liu and Edward H. Adelson. Gelsight fin ray: Incorporating tactile sensing into a soft compliant robotic gripper. In *2022 IEEE 5th International Conference on Soft Robotics (RoboSoft)*, pages 925–931, 2022.
- [27] Ana-Maria Cretu, Thiago Eustaquio Alves De Oliveira, Vinicius Prado Da Fonseca, Bilal Tawbe, Emil M Petriu, and Voicu Z Groza. Computational intelligence and mechatronics solutions for robotic tactile object recognition. In *2015 IEEE 9th international symposium on intelligent signal processing (WISP) proceedings*, pages 1–6. IEEE, 2015.
- [28] Naveen Kuppuswamy, Alejandro Castro, Calder Phillips-Grafflin, Alex Alspach, and Russ Tedrake. Fast model-based contact patch and pose estimation for highly deformable dense-geometry tactile sensors. *IEEE Robotics and Automation Letters*, 5:1811–1818, 4 2020.
- [29] Jiaqi Jiang, Guanqun Cao, Aaron Butterworth, Thanh-Toan Do, and Shan Luo. Where shall i touch? vision-guided tactile poking for transparent object grasping. *IEEE/ASME Transactions on Mechatronics*, 28(1):233–244, 2023.

- [30] Jingang Jiang, Zhiyuan Huang, Zhuming Bi, Xuefeng Ma, and Guang Yu. State-of-the-art control strategies for robotic pih assembly. *Robotics and Computer-Integrated Manufacturing*, 65:101894, 2020.
- [31] Siyuan Dong and Alberto Rodriguez. Tactile-based insertion for dense box-packing. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7953–7960. IEEE, 2019.
- [32] Shiyu Jin, Xinghao Zhu, Changhao Wang, and Masayoshi Tomizuka. Contact pose identification for peg-in-hole assembly under uncertainties. In *2021 American Control Conference (ACC)*, pages 48–53, 2021.
- [33] Jingzhou Song, Qingle Chen, and Zhendong Li. A peg-in-hole robot assembly system based on gauss mixture model. *Robotics and Computer-Integrated Manufacturing*, 67:101996, 2021.
- [34] Dong-Hyuk Lee, Myoung-Su Choi, Hyeonjun Park, Ga-Ram Jang, Jae-Han Park, and Ji-Hun Bae. Peg-in-hole assembly with dual-arm robot and dexterous robot hands. *IEEE Robotics and Automation Letters*, 7(4):8566–8573, 2022.
- [35] Hanwen Kang, Yaohua Zang, Xing Wang, and Yaohui Chen. Uncertainty-driven spiral trajectory for robotic peg-in-hole assembly. *IEEE Robotics and Automation Letters*, 7(3):6661–6668, 2022.
- [36] Sangwoon Kim and Alberto Rodriguez. Active extrinsic contact sensing: Application to general peg-in-hole insertion. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 10241–10247, 2022.

- [37] Viral Rasik Galaiya, Mohammed Asfour, Thiago Eustaquio Alves de Oliveira, Xianta Jiang, and Vinicius Prado da Fonseca. Exploring tactile temporal features for object pose estimation during robotic manipulation. *Sensors*, 23(9):4535, May 2023.
- [38] Baude Billard and Danica Kragic. Trends and challenges in robot manipulation. *Science*, 364(6446):eaat8414, 2019.
- [39] Frank L. Hammond, Jonathan Weisz, Andrés A. de la Llera Kurth, Peter K. Allen, and Robert D. Howe. Towards a design optimization method for reducing the mechanical complexity of underactuated robotic hands. In *2012 IEEE International Conference on Robotics and Automation*, pages 2843–2850, 2012.
- [40] Caner Sahin, Guillermo Garcia-Hernando, Juil Sock, and Tae-Kyun Kim. A review on object pose recovery: From 3d bounding box detectors to full 6d pose estimators. *Image and Vision Computing*, 96:103898, 2020.
- [41] Julian Zimmer, Tess Hellebrekers, Tamim Asfour, Carmel Majidi, and Oliver Kroemer. Predicting grasp success with a soft sensing skin and shape-memory actuated gripper. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7120–7127, 2019.
- [42] Nicholas Wettels, Veronica J Santos, Roland S Johansson, and Gerald E Loeb. Biomimetic tactile sensor array. *Advanced robotics*, 22(8):829–849, 2008.
- [43] 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, 2018.
- [44] Mike Lambeta, Po-Wei Chou, Stephen Tian, Brian Yang, Benjamin Maloon, Victoria Rose Most, Dave Stroud, Raymond Santos, Ahmad Byagowi, Gregg Kammerer, et al. 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, 2020.
- [45] Alex Alspach, Kunimatsu Hashimoto, Naveen Kuppaswamy, and Russ Tedrake. Soft-bubble: A highly compliant dense geometry tactile sensor for robot manipulation. In *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)*, pages 597–604. IEEE, 2019.
- [46] 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. IEEE, 2015.
- [47] Sohee John Yoon, Minsik Choi, Bomin Jeong, and Yong-Lae Park. Elongatable gripper fingers with integrated stretchable tactile sensors for underactuated grasping and dexterous manipulation. *IEEE Transactions on Robotics*, 38(4):2179–2193, 2022.

- [48] Naghmeh Bandari, Javad Dargahi, and Muthukumaran Packirisamy. Tactile sensors for minimally invasive surgery: A review of the state-of-the-art, applications, and perspectives. *IEEE Access*, 8:7682–7708, 2020.
- [49] Yu She, Shaoxiong Wang, Siyuan Dong, Neha Sunil, Alberto Rodriguez, and Edward Adelson. Cable manipulation with a tactile-reactive gripper. *The International Journal of Robotics Research*, 40(12-14):1385–1401, 2021.
- [50] Thomas Bi, Carmelo Sferrazza, and Raffaello D’Andrea. Zero-shot sim-to-real transfer of tactile control policies for aggressive swing-up manipulation. *IEEE Robotics and Automation Letters*, 6(3):5761–5768, 2021.
- [51] Hanwen Zhang, Zeyu Lu, Wenyu Liang, Haoyong Yu, Yao Mao, and Yan Wu. Interaction control for tool manipulation on deformable objects using tactile feedback. *IEEE Robotics and Automation Letters*, 2023.
- [52] Felix Von Drigalski, Shohei Taniguchi, Robert Lee, Takamitsu Matsubara, Masashi Hamaya, Kazutoshi Tanaka, and Yoshihisa Ijiri. Contact-based in-hand pose estimation using bayesian state estimation and particle filtering. *Proceedings - IEEE International Conference on Robotics and Automation*, pages 7294–7299, 2020.
- [53] Daniel Fernandes Gomes, Zhonglin Lin, and Shan Luo. Geltip: A finger-shaped optical tactile sensor for robotic manipulation. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9903–9909, 2020.
- [54] Branden Romero, Filipe Veiga, and Edward Adelson. Soft, round, high resolution tactile fingertip sensors for dexterous robotic manipulation. In *2020*

- IEEE International Conference on Robotics and Automation (ICRA)*, pages 4796–4802, 2020.
- [55] Camill Trueeb, Carmelo Sferrazza, and Raffaello D’Andrea. Towards vision-based robotic skins: a data-driven, multi-camera tactile sensor. In *2020 3rd IEEE International Conference on Soft Robotics (RoboSoft)*, pages 333–338, 2020.
- [56] Andrea Sipos and Nima Fazeli. Simultaneous contact location and object pose estimation using proprioception and tactile feedback. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3233–3240, 2022.
- [57] John Lloyd and Nathan F. Lepora. Goal-driven robotic pushing using tactile and proprioceptive feedback. *IEEE Transactions on Robotics*, 38(2):1201–1212, 2022.
- [58] Bruno Monteiro Rocha Lima, Luiz Claudio Sampaio Ramos, Thiago Eustaquio Alves de Oliveira, Vinicius Prado da Fonseca, and Emil M. Petriu. Heart rate detection using a multimodal tactile sensor and a z-score based peak detection algorithm. *CMBES Proceedings*, 42, May 2019.
- [59] Thiago Eustaquio Alves de Oliveira, Vinicius Prado da Fonseca, Emanuil Huluta, Paulo FF Rosa, and Emil M Petriu. Data-driven analysis of kinaesthetic and tactile information for shape classification. In *2015 IEEE International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications (CIVEMSA)*, pages 1–5. IEEE, 2015.

- [60] Jingjing Ji, Yuting Liu, and Huan Ma. Model-based 3d contact geometry perception for visual tactile sensor. *Sensors*, 22(17), 2022.
- [61] Umer Shah, Rajkumar Muthusamy, Dongming Gan, Yahya Zweiri, and Lakmal Seneviratne. On the design and development of vision-based tactile sensors. *Journal of Intelligent & Robotic Systems*, 102, 08 2021.
- [62] Sudharshan Suresh, Maria Bauza, Kuan-Ting Yu, Joshua G. Mangelson, Alberto Rodriguez, and Michael Kaess. Tactile slam: Real-time inference of shape and pose from planar pushing. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, page 11322–11328. IEEE Press, 2021.
- [63] David Álvarez, Máximo A. Roa, and Luis Moreno. Tactile-based in-hand object pose estimation. In Anibal Ollero, Alberto Sanfeliu, Luis Montano, Nuno Lau, and Carlos Cardeira, editors, *ROBOT 2017: Third Iberian Robotics Conference*, pages 716–728, Cham, 2018. Springer International Publishing.
- [64] Osher Azulay, Inbar Ben-David, and Avishai Sintov. Learning haptic-based object pose estimation for in-hand manipulation with underactuated robotic hands. *arXiv*, 2022.
- [65] Satoshi Funabashi, Tomoki Isobe, Fei Hongyi, Atsumu Hiramoto, Alexander Schmitz, Shigeki Sugano, and Tetsuya Ogata. Multi-fingered in-hand manipulation with various object properties using graph convolutional networks and distributed tactile sensors. *IEEE Robotics and Automation Letters*, 7(2):2102–2109, 2022.

- [66] David Álvarez, Máximo A. Roa, and Luis Moreno. Visual and tactile fusion for estimating the pose of a grasped object. In Manuel F. Silva, José Luís Lima, Luís Paulo Reis, Alberto Sanfeliu, and Danilo Tardioli, editors, *Robot 2019: Fourth Iberian Robotics Conference*, pages 184–198, Cham, 2020. Springer International Publishing.
- [67] Snehal Dikhale, Karankumar Patel, Daksh Dhingra, Itoshi Naramura, Akinobu Hayashi, Soshi Iba, and Nawid Jamali. Visuotactile 6d pose estimation of an in-hand object using vision and tactile sensor data. *IEEE Robotics and Automation Letters*, 7(2):2148–2155, 2022.
- [68] Stanford Artificial Intelligence Laboratory et al. Robotic operating system.
- [69] Vinicius Prado da Fonseca, Xianta Jiang, Emil M Petriu, and Thiago Eustaquio Alves de Oliveira. Tactile object recognition in early phases of grasping using underactuated robotic hands. *Intelligent Service Robotics*, 15(4):513–525, 2022.
- [70] Hassen Nigatu, Yun Ho Choi, and Doik Kim. Analysis of parasitic motion with the constraint embedded jacobian for a 3-prs parallel manipulator. *Mechanism and Machine Theory*, 164:104409, 2021.
- [71] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris

Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.

- [72] Matthew T. Mason. Toward robotic manipulation. *Annual Review of Control, Robotics, and Autonomous Systems*, 1(1):1–28, 2018.
- [73] Wenhan Hao, Chen Zhu, and Michael Meurer. Camera calibration error modeling and its impact on visual positioning. In *2023 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, pages 1394–1399, 2023.
- [74] Tomáš Hodaň, Jiří Matas, and Štěpán Obdržálek. On evaluation of 6d object pose estimation. In Gang Hua and Hervé Jégou, editors, *Computer Vision – ECCV 2016 Workshops*, pages 606–619, Cham, 2016. Springer International Publishing.
- [75] Jiaqi Jiang, Guanqun Cao, Aaron Butterworth, Thanh-Toan Do, and Shan Luo. Where shall i touch? vision-guided tactile poking for transparent object grasping. *IEEE/ASME Transactions on Mechatronics*, 28(1):233–244, 2022.
- [76] Bruno Monteiro Rocha Lima, Thiago Eustaquio Alves de Oliveira, and Vinicius Prado da Fonseca. Classification of textures using a tactile-enabled finger in dynamic exploration tasks. In *2021 IEEE Sensors*, pages 1–4, 2021.
- [77] Bruno Monteiro Rocha Lima, Vinicius Prado da Fonseca, Thiago Eustaquio Alves de Oliveira, Qi Zhu, and Emil M Petriu. Dynamic tactile explo-

- ration for texture classification using a miniaturized multi-modal tactile sensor and machine learning. In *2020 IEEE International Systems Conference (SysCon)*, pages 1–7. IEEE, 2020.
- [78] Somchai Pohtongkam and Jakkree Srinonchat. Object recognition for humanoid robots using full hand tactile sensor. *IEEE Access*, 11:20284–20297, 2023.
- [79] Vinicius Prado da Fonseca, Xianta Jiang, Emil M Petriu, and Thiago Eustaquio Alves de Oliveira. Tactile object recognition in early phases of grasping using underactuated robotic hands. *Intelligent Service Robotics*, 15(4):513–525, 2022.
- [80] Willow Mandil, Vishnu Rajendran, Kiyanoosh Nazari, and Amir Ghalamzan-Esfahani. Tactile-sensing technologies: Trends, challenges and outlook in agri-food manipulation. *Sensors*, 23(17), 2023.
- [81] Chao Wang, Xuehe Zhang, Xizhe Zang, Yubin Liu, Guanwen Ding, Wenxin Yin, and Jie Zhao. Feature sensing and robotic grasping of objects with uncertain information: A review. *Sensors*, 20(13):3707, 2020.
- [82] Kilian Kleeberger, Richard Bormann, Werner Kraus, and Marco Huber. A survey on learning-based robotic grasping. *Current Robotics Reports*, 1:239–249, 12 2020.
- [83] Stephen James, Andrew Davison, and Edward Johns. Transferring end-to-end visuomotor control from simulation to real world for a multi-stage task. *Conference on Robot Learning*, 07 2017.

- [84] Martin Matak and Tucker Hermans. Planning visual-tactile precision grasps via complementary use of vision and touch. *IEEE Robotics and Automation Letters*, 8(2):768–775, 2022.
- [85] I. Taylor, S. Dong, and A. Rodriguez. Gelslim3.0: High-resolution measurement of shape, force and slip in a compact tactile-sensing finger, 2021.
- [86] Maxwell Welyhorsky, Vinicius Prado Da Fonseca, Qi Zhu, Bruno Monteiro Rocha Lima, Thiago Eustaquio Alves De Oliveira, and Emil M Petriu. Neuro-fuzzy grasp control for a teleoperated five finger anthropomorphic robotic hand. In *2022 IEEE International Systems Conference (SysCon)*, pages 1–5. IEEE, 2022.
- [87] Qi Zhu, Vinicius Prado da Fonseca, Bruno Monteiro Rocha Lima, Maxwell Welyhorsky, Miriam Goubran, Thiago Eustaquio Alves de Oliveira, and Emil M Petriu. Teleoperated grasping using a robotic hand and a haptic-feedback data glove. In *2020 IEEE International Systems Conference (SysCon)*, pages 1–7. IEEE, 2020.
- [88] Cristian C Beltran-Hernandez, Damien Petit, Ixchel G Ramirez-Alpizar, and Kensuke Harada. Variable compliance control for robotic peg-in-hole assembly: A deep-reinforcement-learning approach. *Applied Sciences*, 10(19):6923, 2020.
- [89] Hikaru Unten, Sho Sakaino, and Toshiaki Tsuji. Peg-in-hole using transient information of force response. *IEEE/ASME Transactions on Mechatronics*, 2022.

- [90] Michelangelo Nigro, Monica Sileo, Francesco Pierri, Katia Genovese, Domenico D Bloisi, and Fabrizio Caccavale. Peg-in-hole using 3d workpiece reconstruction and cnn-based hole detection. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4235–4240. IEEE, 2020.
- [91] Shiyu Jin, Xinghao Zhu, Changhao Wang, and Masayoshi Tomizuka. Contact pose identification for peg-in-hole assembly under uncertainties. In *2021 American Control Conference (ACC)*, pages 48–53, 2021.
- [92] Jingzhou Song, Qingle Chen, and Zhendong Li. A peg-in-hole robot assembly system based on gauss mixture model. *Robotics and Computer-Integrated Manufacturing*, 67:101996, 2021.
- [93] Yongzhi Wang, Lei Zhao, Qian Zhang, Ran Zhou, Liping Wu, Junqiao Ma, Bo Zhang, and Yu Zhang. Alignment method of combined perception for peg-in-hole assembly with deep reinforcement learning. *Journal of Sensors*, 2021:1–12, 2021.
- [94] Yuqian Lu, Hao Zheng, Saahil Chand, Wanqing Xia, Zengkun Liu, Xun Xu, Lihui Wang, Zhaojun Qin, and Jinsong Bao. Outlook on human-centric manufacturing towards industry 5.0. *Journal of Manufacturing Systems*, 62:612–627, 2022.
- [95] Regina Kyung-Jin Lee, Hao Zheng, and Yuqian Lu. Human-robot shared assembly taxonomy: A step toward seamless human-robot knowledge transfer. *Robotics and Computer-Integrated Manufacturing*, 86:102686, 2024.

- [96] Arik Lämmle, Philipp Tenbrock, Balázs Bálint, Frank Nägele, Werner Kraus, József Váncza, and Marco F. Huber. Simulation-based learning of the peg-in-hole process using robot-skills. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9340–9346, 2022.
- [97] Shir Kozlovsky, Elad Newman, and Miriam Zacksenhouse. Reinforcement learning of impedance policies for peg-in-hole tasks: Role of asymmetric matrices. *IEEE Robotics and Automation Letters*, 7(4):10898–10905, 2022.
- [98] Cristian C. Beltran-Hernandez, Damien Petit, Ixchel G. Ramirez-Alpizar, and Kensuke Harada. Variable compliance control for robotic peg-in-hole assembly: A deep-reinforcement-learning approach. *Applied Sciences*, 10(19), 2020.
- [99] Thiago Eustaquio Alves de Oliveira, Vinicius Prado da Fonseca, Bruno Monteiro Rocha Lima, Ana-Maria Cretu, and M Petriu. End-effector approach flexibilization in a surface approximation task using a bioinspired tactile sensing module. In *2019 IEEE International Symposium on Robotic and Sensors Environments (ROSE)*, pages 1–6. IEEE, 2019.
- [100] Liang Xie, Hongxiang Yu, Kechun Xu, Tong Yang, Minhang Wang, Haojian Lu, Rong Xiong, and Yue Wang. Learning a simulation-based visual policy for real-world peg in unseen holes. *Review of Scientific Instruments*, 94(10), 2023.
- [101] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.