

Learning Haptic-based Object Pose Estimation for In-hand Manipulation Control with Underactuated Robotic Hands

Osher Azulay, Inbar Ben-David and Avishai Sintov

Abstract—Unlike traditional robotic hands, underactuated compliant hands are challenging to model due to inherent uncertainties. Consequently, pose estimation of a grasped object is usually performed based on visual perception. However, visual perception of the hand and object can be limited in occluded or partly-occluded environments. In this paper, we aim to explore the use of haptics, i.e., kinesthetic and tactile sensing, for pose estimation and in-hand manipulation with underactuated hands. Such haptic approach would mitigate occluded environments where line-of-sight is not always available. We put an emphasis on identifying the feature state representation of the system that does not include vision and can be obtained with simple and low-cost hardware. For tactile sensing, therefore, we propose a low-cost and flexible sensor that is mostly 3D printed along with the finger-tip and can provide implicit contact information. Taking a two-finger underactuated hand as a test-case, we analyze the contribution of kinesthetic and tactile features along with various regression models to the accuracy of the predictions. Furthermore, we propose a Model Predictive Control (MPC) approach which utilizes the pose estimation to manipulate objects to desired positions solely based on haptics. We have conducted a series of experiments that validate the ability to estimate poses of various objects with different geometry, stiffness and texture, and show manipulation to goals in the workspace with relatively high accuracy.

Index Terms—Pose estimation, in-hand manipulation, underactuated hands.

I. INTRODUCTION

WHILE the ability to manipulate an object within the hand is a fundamental everyday task for humans, such problem remains challenging for robots. Traditional robotic hands, such as the Shadow and the Allegro hands [1], have achieved significant accuracy and performance. However, they have complex structure while being fragile, costly and difficult to control [2]–[4]. *Underactuated hands*, on the other hand, are mechanisms that can adapt to the shape of the object through the use of compliance [5]. Consequently, they have gained popularity in recent years due to their low-cost and ability to maintain a stable grasp with open-loop control. In addition, precision manipulation with underactuated hands have been shown possible [6]–[8]. However, certain limitations hinder practical usage. While pose estimation is generally available analytically in rigid hands where the kinematics are known [9], such model is rarely available for underactuated ones



Fig. 1. Manipulating an object within an occluded cabinet using a robotic arm and an underactuated hand into a target slot. Visual perception is not available within the cabinet and, therefore, the hand must use haptic perception. We assume that the position of the hole relative to the robot is known.

[10]. Acquiring a precise analytical model for an underactuated hand is not always easy or feasible due to inherent uncertainties and fabrication inaccuracies. Hence, estimating properties such as joint stiffness, size, weight, friction and inertia is a significant challenge [11]. To cope with the lack of an analytical solution, data-based modeling was shown useful in providing accurate predictions in motion planning and control [12]–[14]. Such modeling approach is able to intrinsically estimate model parameters that can be difficult or impossible to model otherwise.

As mentioned above, exact analytical solutions are rarely available for underactuated hands. Therefore, attempts to model these hands in the context of precision manipulation usually rely on external visual feedback. Calli and Dollar [7] used a linear approximation of a two-finger hand to apply visual servoing to track desired paths. Recently, Morgan et al. [15] proposed an object agnostic manipulation using a vision-based Model Predictive Control (MPC). The work in [16] used a depth camera to estimate the pose of an object grasped and partly-occluded by the two-fingers of an underactuated hand. Extension of the work proposed the use of the depth-based 6D pose estimation to control precise manipulation of a grasped object [17]. Recent work [18] integrated allocentric visual perception along with four tactile modules, that combine pressure, magnetic, angular velocity and gravity sensors, on two underactuated fingers. These sensors were used to train a pose estimation model. More related, Sintov et al. [10]

O. Azulay, I. Ben-David and A. Sintov are with the School of Mechanical Engineering, Tel-Aviv University, Israel. e-mail: osherazulay@mail.tau.ac.il, sintov1@tauex.tau.ac.il

This work was supported by the Pazy Foundation (grant No. 283-20).

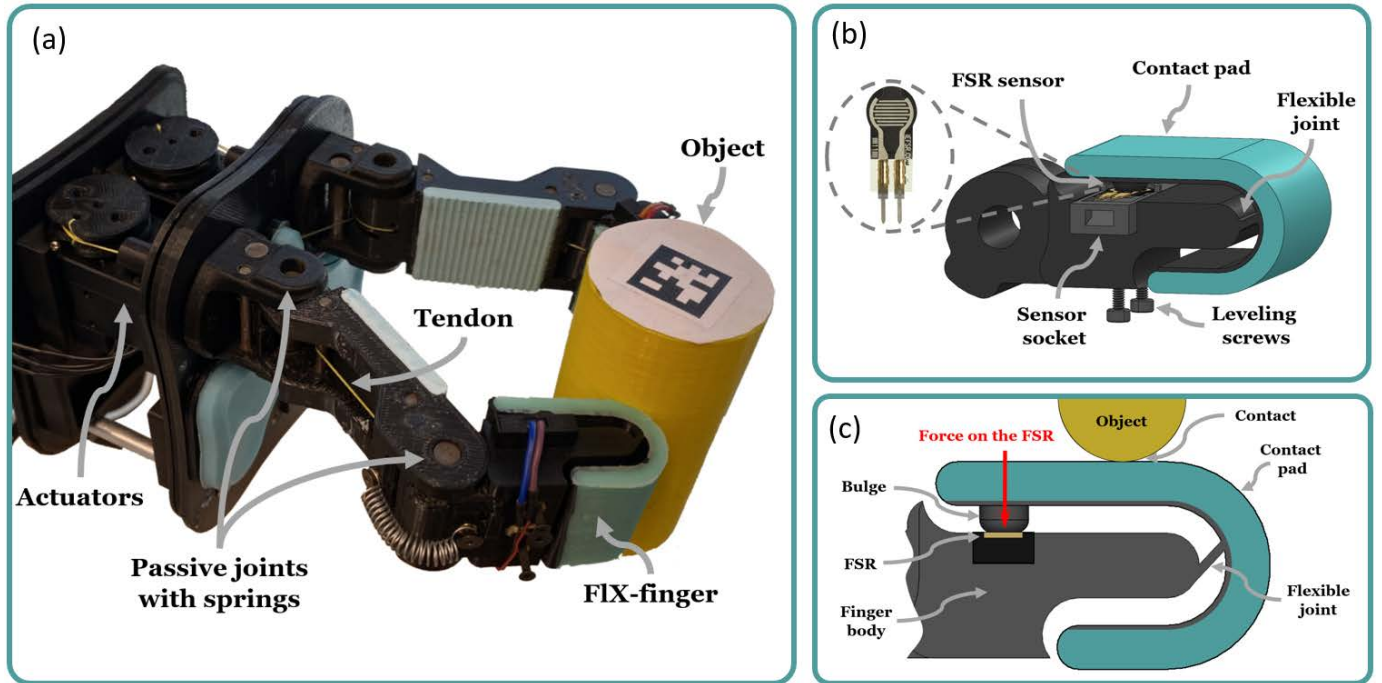


Fig. 2. (a) Structure of a tendon-based underactuated hand used in this work. The hand is comprised of two opposing fingers. Each finger has two passive joints with springs where a tendon wire runs along its length and is connected to an actuator. In addition, each finger includes a FLX-finger for contact sensing. A (b) detailed and (c) side view of the FIX-finger with the FSR sensor are also seen. When contact force is exerted on the contact pad by the object, the flexible joint and pad deform. Consequently, the bulge is pressed onto the FSR which, in turn, outputs a voltage that depends on the contact distance from the bulge across the pad along with the magnitude of the normal contact force.

proposed a data-based transition model where the state of the hand involves kinesthetic features such as actuator torques and angles along with the position of the manipulated object acquired with visual feedback. While a visual approach may exhibit good results, relying on continuous visual feedback limits the performance of various tasks in which visual uncertainty (e.g., poor lighting or shadows) or occlusion may occur. In such cases, it may be impossible to solve the task altogether. This may include operating tools at the back of a cabinet or in confined spaces (Figure 1).

Information regarding grasped objects are often acquired through haptic perception including both tactile sensing and internal sensing of joint actuators known as Kinesthetic haptics [19]. Traditionally, tactile refers to information received from touch and contact sensing, while kinesthetic refers to information sensed through movement, force or position of joints and actuators. Tactile sensing is the leading method for haptic-based object recognition [20]–[22]. State of the art in tactile sensing is focused on optical devices where an internal camera observes the deformation of a flexible surface during contact with an object [23], [24]. Recent work has explored the use of these optical tactile sensors with advanced deep networks to estimate the relative pose of an object in contact [25]. While these sensors can provide accurate performance, they can complicate the hardware.

The abilities of underactuated hands in various tasks and their advantages over traditional rigid hands have been widely demonstrated in prior work [2], [7], [10], [12], [26], [27]. However and as discussed above, these abilities rely on extrinsic visual perception. In this paper, we investigate the

ability to track and manipulate an object within a tendon-based two-finger underactuated hand (seen in Figure 2) solely based on haptics. As opposed to prior work where the hand’s state included explicit object position acquired from vision [10], visual perception is not relied upon in this work. Hence, we explore the sole use of kinesthetic information available to the hand along with novel low-cost tactile fingers termed *FLX-fingers*. The FLX-finger is mostly 3D printed along with a simple off-the-shelf force Sensitive Resistor (FSR). While the use of FSR for tactile perception is not novel [28], we propose a flexible 3D printed mechanism that augments the sensor and provides more contact data. Therefore and unlike custom made force sensors (e.g., [29]), we use a standard off-the shelf FSR sensor and a 3D printed finger structure to acquire diverse contact data.

With the FLX-fingers, we search for a feature representation of the state of the hand while not including explicit pose information about the object. Hence, the state of the hand includes measurable features, either tactile or kinesthetic ones, that can be mapped to the true position and orientation of the object. Using such state representation, we study fundamental tasks for in-hand manipulation without relying on visual perception: learning an haptic-based pose observation model of a grasped object and controlling the motion of the hand to manipulate the object to some goal. Furthermore, we aim to understand the data requirements and learning abilities over a set of test objects. While an observation model enables somewhat estimation of the pose of the object, the accuracy of a learned model is limited and reaches saturation with some amount of data. Hence, we train a *critic* model with independent data

to estimate prediction errors. The critic is then incorporated in a Model Predictive Control (MPC) [30] to manipulate an object to a goal position while minimizing estimation errors along the motion. In such way, the controller will be able to avoid regions where the observation model tends to provide large erroneous predictions. The scope of this work is focused on the information requirement for haptic-based object pose estimation. While we do not aim for object-agnostic models, we test the ability of a trained observation model to generalize to new objects of different shapes that were not included in training.

To summarize, the contributions of this work are as follows:

- A novel 3D printed simple and low-cost tactile finger is proposed for underactuated hands.
- An observation model is proposed to map haptic sensing of an underactuated hand to the pose of an object. The haptics include sensing from the tactile fingers along with kinesthetic sensing from the actuators of the hand. With such, a study is conducted about the haptic information required for object pose estimation with an underactuated hand while including tactile and kinesthetic sensing.
- We analyze the partial-occluded pose-estimation problem where an initial visual glance of the grasp is available prior to performing the task in an occluded region.
- Models for pose estimation are shown to be somewhat generalizable.
- We propose an MPC approach to manipulate a grasped object to desired goal positions solely based on the predictions of the observation model. While an entire pose estimation is available, controlling the orientation of the object is highly dependent on the initial grasp. Hence, the MPC can only control the position of the grasped object. To include orientation, a designated motion planning with some re-grasping maneuvers is required in future work.
- A critic model is added to the MPC approach such that the controller can reason about the accuracy of the observation model and avoid potentially erroneous actions.

To the best of the author's knowledge, this is the first analysis of pose estimation with underactuated hands solely by using simple and low-cost haptic perception. Achieving haptic-based pose estimation would enable further fusion with vision for better estimation in unstructured or partly-occluded environments.

II. RELATED WORK

The presented work discusses the use of underactuated hands in occluded environments by using haptics. Related work on these topics is briefly discussed below.

Underactuated robotic hands. Underactuated hands are compliant mechanisms that exhibit complex behavior due to the passivity of the adaptive joints [10]. Compliance can be achieved in various mechanisms such as inflating soft materials [31], [32], or through passive joints with springs and tendons [2]. The work in this paper will focus on tendon-based underactuated hands seen in Figure 2a. The recent introduction of open-sourced hands fabricated through 3D

printing has enabled easy access to low-cost hardware [33]. However, precise mechanical properties such as joint stiffness, contact coefficients and joint friction are hard to extract. Consequently, accurate models of such systems are difficult to formulate. Modeling tools, introduced in several works [5], [34], [35], examine joint configurations, torques, and energy with a simplified frictional model. Nevertheless, these techniques have been shown to be sensitive to assumptions in external constraints and are generally suitable for simulations. Recent work has proposed the self-identification of necessary parameters through exploratory hand-object interactions using an external camera and particle filtering [36].

Manipulation in occluded environments. Additional sensor modalities are commonly used to augment or alternate vision [37]. Most work utilizes haptic sensors to reason about the state of contact [38], [39]. Haptic perception is generally used to learn various features of an object in uncertain environments to grasp and manipulate it. Such information may include stiffness, texture, temperature variations and surface modeling [40]. Often, haptic perception is used alongside vision to refine initial pose estimation [41]. For instance, recent work has trained a policy to leverage multimodal feedback in contact-rich manipulation tasks from tactile and vision through self-supervision [42]. Another work used 6-axis force-torque and tactile sensors on a multi-finger rigid hand to train temporal neural-network models for in-hand manipulation tasks [43]. Opposed to contact sensing, a different approach used acoustic perception to gain information regarding hand-object contact [44]. Radio frequency perception was used in a different work to retrieve objects in fully-occluded settings [45].

Contact sensing. The common approach for pose estimation during manipulation is tactile to sense contact. Such sensing has been generally achieved using simple force or pressure sensors [46]. As such, Koval et al. [47] used contact sensors and particle filtering to estimate the pose of an object during contact manipulation. In recent years, sensor arrays have become more common due to advancements in fabrication abilities and due to their effectiveness in covering large contact areas [48]. The work by Sodhi et al. [49], for instance, used data from a optical tactile sensor to estimate the pose of an object being pushed. Nonetheless, an analysis in [50] has proven that sensor complexity does not necessarily provide better performance mainly due to the way tactile data is represented. It was shown that simpler fingertip sensors yielded equivalent accuracy.

III. PROBLEM FORMULATION

We consider a two-fingered adaptive hand comprised of two opposing tendon-based fingers as seen in Figure 2a. Each finger has two compliant joints with springs where a tendon wire runs along its length and is connected to an actuator. Also, each distal link of the finger has high friction pads to avoid slipping. Let $\mathbf{x} \in \mathcal{C}$ and $\mathbf{a} \in \mathcal{U}$ be some observable state of the hand and action vector, respectively, where $\mathcal{C} \subset \mathbb{R}^n$ and $\mathcal{U} \subset \mathbb{R}^2$. Action \mathbf{a} is the change of actuator angles (i.e., pulling or releasing the tendons) over a fixed time step Δt . The true state of the hand is not entirely accessible. Hence, we consider

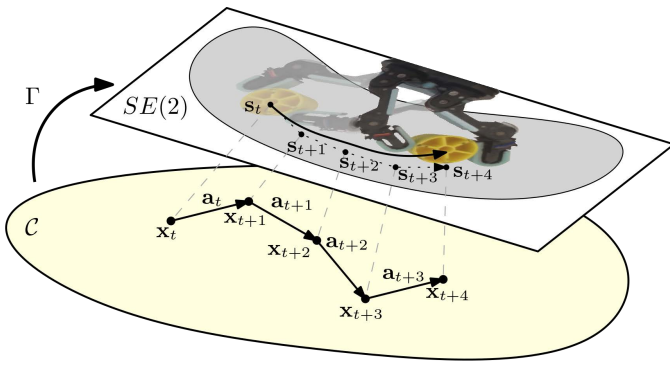


Fig. 3. Haptic-based feature space \mathcal{C} where motion is governed by a transition model such that $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{a}_t)$. The observation model maps a state \mathbf{x}_t in \mathcal{C} to the corresponding pose \mathbf{s}_t of the grasped object in $SE(2)$. Hence, a path from \mathbf{x}_t to \mathbf{x}_{t+4} is mapped (gray dashed lines) with the observation model Γ to the approximated path of the object in $SE(2)$ (dotted curve from \mathbf{s}_t to \mathbf{s}_{t+4}). The solid curve in $SE(2)$ denotes the true path of the object.

a feature state in space \mathcal{C} that may correspond to different kinesthetic features (e.g., actuator angles and loads) and tactile feedback. Also, the system is governed by a transition function $f: \mathcal{C} \times \mathcal{U} \rightarrow \mathcal{C}$ such that, given the current state \mathbf{x}_t and action \mathbf{a}_t , the next state is given by $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{a}_t)$.

As stated in the introduction and due to inherent fabrication uncertainties, neither the true state of the hand nor an analytical formulation of the transition model are known. In addition, feedback regarding the true pose of the manipulated object is not always available with vision, for instance, in an occluded environment. The goal is therefore twofold. First, we search for n state features that are easy to measure and best embed the pose $\mathbf{s} \in SE(2)$ of a grasped object. Such state representation should enable training an observation model $\Gamma: \mathcal{C} \rightarrow SE(2)$ to estimate the pose of the object. Second, we wish to explore the use of such state representation to manipulate objects without external pose feedback. Hence, we utilize a learned state transition model \tilde{f} to plan and control motion in the feature space \mathcal{C} while acquiring desired positional motion in the work-plane. The above formulation is illustrated in Figure 3. We note that it may be possible to learn a direct model for current state and action to the next pose, i.e., $h: \mathcal{C} \times \mathcal{U} \rightarrow SE(2)$. However, such model would not enable propagation as required in planning and Reinforcement Learning (RL).

IV. LEARNING AN OBSERVATION MODEL

In this section, we discuss the possible kinesthetic and tactile features to be included in the feature state representation. This includes the design and fabrication of the FIX-finger. Further, we discuss data-based models to learn model Γ .

A. FIX-finger: A low-cost tactile finger pad

We present the design of a low-cost tactile finger termed *FIX-finger*. FIX-finger is a flexible 3D printed mechanism that augments the sensor and provides more contact data. It is mostly fabricated through 3D printing and its shape is based on the design of the distal finger link of the Model-T42 hand [33]. The FIX-finger is composed of four parts seen in Figure

2b: finger body, contact pad, FSR sensor and sensor socket. The finger body is connected to the base of the contact pad solely through a flexible joint. The entire finger body including the base of the contact pad is a one-shot 3D print. After print, we cast Urethane Rubber on the base of the contact pad to acquire a high-friction and soft surface as in [33]. The sensor socket is also printed and contains the FSR. Hence, the sensor socket with the FSR are positioned fixed to the finger body within a designated groove. At such configuration, the FSR is in contact with a small bulge on the back side of the contact pad as seen in Figure 2c. The bulge applies force on the FSR when pressure is exerted on the contact pad and the flexible joint deforms.

FSR sensors are made of polymer films that vary their electrical resistance upon changing pressure on their surface. They are simple to use and low-cost. The FSR sensor is connected to an analog pin of the Arduino Nano through a voltage divider of $4.7k\Omega$ resistor. When an object is being pressed against the contact pad, the flexible joint and pad deform and press the bulge onto the FSR. The contact distance from the bulge across the contact pad along with the magnitude of the normal contact force would define the force on the FSR and, therefore, the measured voltage. Hence, different contact locations across the pad provide different voltages. Consequently, pressure on the contact pad provides voltage signals that embed information regarding the location and magnitude of contact. To acquire somewhat repeatability of measurements, leveling screws on the back side of the finger body are used to vary the height of the sensor socket and, thus, tune the initial voltage on the FSR. By standard, the initial voltage V_o is determined with a weight of 100 grams directly placed on the contact pad when in an horizontal posture. Measurements are then taken relative to V_o . To sum-up, a standard off-the shelf FSR sensor and a 3D printed finger structure is proposed to acquire diverse contact data when grasping an object. The prototype cost of the FIX-finger is estimated at less than \$10.

In the formation of the tactile finger, we expect to acquire a force distribution along the contact pad yielding information on contact position. To validate this and acquire a characteristic force distribution of the FIX-finger, we have designed an experiment using a six degrees-of-freedom robotic arm equipped with a Force/Torque sensor and a rigid pole at its tip as seen in Figure 4. During the experiment, the pole was repeatedly pushed against the finger with predefined forces between $0 - 10 N$ and at 50 equal length (1 mm) locations along the contact pad. Results for FIX-finger with $V_o = 200 mV$ are seen in Figure 5. The results show a near linear behaviour in the work region of contact. Due to the spring-like mechanism within the finger, the signal output increases as the contact location approaches the FSR, and decreases when moving away. While the force-distribution of a single FIX-finger cannot be used to identify explicit or accurate contact position, these results indicate that contact information is embedded in the measurements. However, we care about the pose of the object and not explicit contact locations. Hence, we explore the use of two force-distributions from two FIX-fingers along with kinesthetic haptics to build an observation




Fig. 4. Experiment to characterize signal distribution on the FIX-finger.

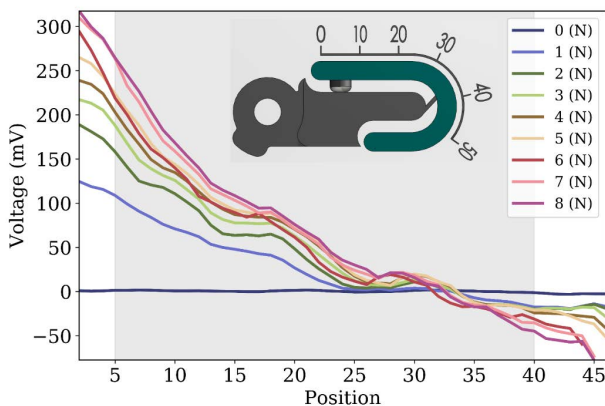


Fig. 5. Signal distribution results for the FIX-finger. The grey area denotes the main working region during manipulation. When the contact is at the back of the finger, force is reduced on the FSR and the signal becomes negative (relative to the initial load).

model as described next.

B. Training data

Training data is collected by grasping an object and manipulating it with random actions. During manipulation, ground-truth data of object poses \mathbf{s}_t is provided. Thus, the resulting data is a set of observed states, poses and actions $\mathcal{P} = \{(\mathbf{x}_1, \mathbf{s}_1, \mathbf{a}_1), \dots, (\mathbf{x}_N, \mathbf{s}_N, \mathbf{a}_N)\}$. In our work, we investigate the contribution of different available features in \mathbf{x} for learning accurate observation and transition models. These include actuator torques, actuator angles and FIX-finger tactile voltage signals. Given a state definition, we train an observation model $\mathbf{s}_t = \Gamma_\theta(\mathbf{x}_t)$ to map a feature state at time t to the true pose of the object. Vector θ consists of the trained parameters of the model. Model Γ_θ is trained using recorded input and output $\{\mathbf{x}_i, \mathbf{s}_i\}_{i=1}^N$ from \mathcal{P} . The actions in \mathcal{P} will later be used to train a transition model.

C. Observation Model

We investigate several architectures of regression models that could approximate $\mathbf{s}_t = \Gamma_\theta(\mathbf{x}_t)$ in real-time. We test Gaussian Processes (GP) [51], Fully-Connected Neural-Network (FC-NN) and Long Short-Term Memory (LSTM) based Recursive-NN (RNN). GP and FC-NN provide prediction for a single time-step, i.e., based on an instantaneous state perception. We employ local GP regression similar to [10] where the function's value at a given query state is predicted using a subset of the training points that are in the vicinity of the query point. LSTM is a time-dependant NN where we exploit internal memory of the motion. Since the force sensors may be slightly noisy, having several consecutive measurements may improve accuracy. Hence, such model is formulated by $\mathbf{s}_t = \Gamma_\theta(\mathbf{x}_t, \mathbf{x}_{t-1}, \dots, \mathbf{x}_{t-w})$ where w is the size of past state measurements to include. Observing prediction accuracy of an LSTM model would provide insights about the importance of previous states compared to single state models. Since the motion of the object is slow while prediction time is very fast (demonstrated in the experimental section), latency between the measured state and prediction is assumed to be negligible. Hence, the predicted pose sufficiently corresponds to the true pose expressed by the feature state.

D. Critic model

With the addition of data to the training of the observation model, the accuracy improves until reaching saturation at some point. Hence, the accuracy may be limited with the addition of training data. This would be demonstrated in the experiments. Inspired by [12], we use redundant training data to generate a critic model $e_t = C_\varphi(\mathbf{x}_t)$ where $e_t \in \mathbb{R}^+$ and φ consists of the trained parameters of the model. Model C_φ is trained with measured inputs \mathbf{x}_j and corresponding labels $e_j = \|\tilde{\mathbf{s}}_j - \mathbf{s}_j\|$ where $\tilde{\mathbf{s}}_j = \Gamma_\theta(\mathbf{x}_j)$ and \mathbf{s}_j is the ground-truth pose. The critic provides an estimate of the haptic observation error. With the critic, we would be able to plan and control the system to move in regions with higher observation accuracy as discussed in the next section.

V. TRANSITION MODELING AND CONTROL

Having a transition model $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{a}_t)$ is a key component for planning and control for an underactuated hand. In this section, we discuss the training of a transition model and its usage to control motion. The data described in Section IV-B is organized in the form $\{(\mathbf{x}_i, \mathbf{a}_i), (\mathbf{x}_{i+1})\}_{i=1}^N$. Unlike previous work [10] where actions were from a discrete set with a cardinality of 4-8, actions in our work are continuous within a normalized range $[0, 1]$. Similar to previous work [7], [27] and since the hand is underactuated, we focus on controlling the position of the object. The ability to reach some orientation is highly dependent on the initial grasp. Hence, including orientation, which is predicted by the proposed observation model, requires a designated motion planning while also considering re-grasping maneuvers (e.g., pick and place). The orientation control problem is left for future work.

As discussed in Sections I-II, previous work such as in [10] used a vision-based transition model. In another work [26], MPC is applied using image-based visual servoing with a linear approximation of the system [7] without the use of a transition model as in this work. While these methods and others rely on visual perception, we propose an MPC approach that is based on haptics and a data-based model. Hence, an accurate motion to various goals is possible without the need for a line-of-sight with the hand.

A. Learning a Transition Model

Given a feature state representation as discussed above, we now learn a transition model to use in the control. A straightforward parameterization of $\tilde{f}_\phi(\mathbf{x}_t, \mathbf{a}_t)$, where ϕ is the a vector of network weights, can be difficult when the sampling frequency is high and the consecutive states are too similar [52]. Therefore, instead of learning a direct transition function, we train a model to predict the change from state \mathbf{x}_t given an action \mathbf{a}_t over time step Δt . The next state prediction is, consequently, given by

$$\tilde{\mathbf{x}}_{t+1} = \mathbf{x}_t + \tilde{f}_\phi(\mathbf{x}_t, \mathbf{a}_t). \quad (1)$$

Several data-based models are benchmarked for training model \tilde{f}_ϕ including GP [10] and FC-NN. To capture deviations that could violate the standard Markov assumption on the state space, we also learn a time-dependent LSTM-based RNN such that $\tilde{\mathbf{x}}_{t+1} = \mathbf{x}_t + \tilde{f}_\phi(\mathbf{x}_t, \mathbf{x}_{t-1} \dots, \mathbf{x}_{t-k}, \mathbf{a}_t)$ where $k > 0$ is the number of preceding states. We note that the data used for learning the transition model includes sliding of the object along the contact pads relative to the fingers in some states and actions. Hence, sliding is modeled and taken into account in planning and control.

B. Haptic Servoing

Previous work by Calli and Dollar [7] have proposed the use of Visual Servoing (VS) to control the hand for reaching a goal position and tracking a path. VS utilizes a simple control rule based on linear approximation of the hand-object system and is only able to control the position of the object (without orientation). The control maps the desired object velocity to the required actuator velocity and is given by [7]

$$\mathbf{a}_t = \begin{bmatrix} \frac{1}{K_x} & \frac{1}{K_y} \\ -\frac{1}{K_x} & \frac{1}{K_y} \end{bmatrix} \mathbf{v}_t, \quad (2)$$

where $\mathbf{v}_t = \mathbf{p}_g - \mathbf{p}_t$, $\mathbf{p}_g \in \mathbb{R}^2$ is the desired goal position, \mathbf{p}_t is the current position extracted from \mathbf{s}_t and $K_x, K_y > 0$ are constant scalars related to the hand. Position \mathbf{p}_t is originally measured in real-time using visual perception. Let map $\tilde{\Gamma}_\theta : \mathcal{C} \rightarrow \mathbb{R}^2$ be a position estimation function taken from the pose estimation map Γ_θ . Our proposed Haptic Servoing (HS) approach approximates \mathbf{p}_t such that $\mathbf{v}_t = \mathbf{p}_g - \tilde{\Gamma}_\theta(\mathbf{x}_t)$. Hence, the system is driven by measuring the haptic state rather than having visual perception.

C. Haptic-based Planning and Control

Since we have no explicit feedback in $SE(2)$, the above HS control relies solely on pose predictions of the learned observation model. However, the model includes some errors

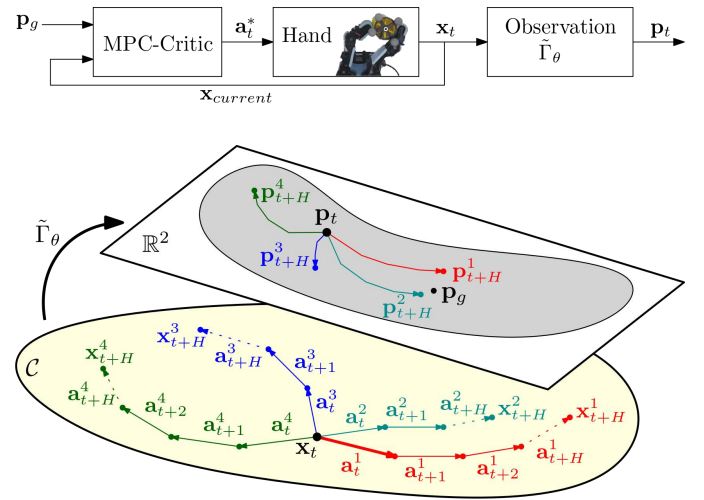


Fig. 6. (Top) Control scheme and (bottom) iteration of the shooting method in \mathcal{C} within an MPC control to find the next action to drive the system to goal position \mathbf{p}_g . Random action sequences are applied from the current state $\mathbf{x}_t = \mathbf{x}_{current}$ and used to propagate states in \mathcal{C} with the transition model. The resulted paths are inputted to observation model $\tilde{\Gamma}_\theta$ to estimate their positions. The first action of the path that reaches closer to the goal while having a low critic error is chosen. In this example, four sequences are propagated (red, cyan, blue and green). While the cyan path (ending with \mathbf{p}_{t+H}^2) reached closer to the goal than the red one (ending with \mathbf{p}_{t+H}^1), it has higher critic error. Hence, the first action \mathbf{a}_t^1 of the red sequence is exerted on the hand. This process is repeated until reaching the goal position.

and uncertainties over the state-space. While HS can be useful to reach a near-by position (as will be demonstrated in the experiments), it does not reason about these possible erroneous predictions of the observation model. Consequently, inaccurate predictions may deviate the motion from the desired goal. Our objective is, therefore, to formulate a controller that can reason about positioning inaccuracies without visual feedback. We propose a Model Predictive Control (MPC) with a random-sampling shooting method. As discussed in [15], MPC can be advantageous for manipulation. In MPC, the next action is optimized based on some cost function while considering the future outcome of candidate paths. The first action of the path with the lowest cost is exerted.

Given the current hand state \mathbf{x}_t , prediction horizon H and an action sequence $A_{t,H} = \{\mathbf{a}_t, \dots, \mathbf{a}_{t+H}\}$, the predicted path in \mathcal{C} acquired by propagating model (1) for H steps with actions $A_{t,H}$ is $\sigma_t = \{\mathbf{x}_t, \tilde{\mathbf{x}}_{t+1}, \dots, \tilde{\mathbf{x}}_{t+H}\}$. The cost of the path is defined to be

$$J(\sigma_t) = \sum_{\mathbf{x}_\tau \in \sigma_t} \left[w_1 \|\tilde{\Gamma}_\theta(\mathbf{x}_\tau) - \mathbf{p}_g\|^2 + w_2 C_\varphi(\mathbf{x}_\tau) \right] \quad (3)$$

where $\mathbf{p}_g \in \mathbb{R}^2$ is the goal position and $w_1, w_2 > 0$ are pre-defined weights. The first component in the sum of (3) estimates the distance of states in the path to the goal. This component prioritize paths that steers more towards the goal. The second component, on the other hand, evaluates the accuracy of the pose approximated by the observation model. We aim to exert the path that takes the object closer to the goal while having high certainty about its pose along the path. Therefore, at time t , we search for the set of actions $A_{t,H}^*$ that is the solution to

$$A_{t,H}^* = \arg \min_{A_{t,H}} J(\sigma_t). \quad (4)$$

Finding the optimal solution of (4), i.e., optimal sequence of actions $A_{t,H}^*$, could be slow and not suitable for real-time control. Hence, we search for a near-optimal solution in a random-sampling shooting setting. Algorithm 1 presents the *MPC-Critic* control with the shooting method. Figure 6 presents the control scheme and illustration of the shooting method within the MPC-Critic control. In each iteration and current state $\mathbf{x}_t = \mathbf{x}_{current}$, m random action sequences $\{A_{t,H}^{(1)}, \dots, A_{t,H}^{(m)}\}$ are generated and their costs are evaluated according to (3). The first action of the lowest cost sequence is exerted. The process is repeated until reaching within distance ϵ to the goal \mathbf{p}_g .

Algorithm 1: MPC-Critic($\mathbf{p}_g, H, \epsilon$)

```

1 Get current hand state  $\mathbf{x}_{current}$ ;
2 do
3    $\mathbf{x}_t \leftarrow \mathbf{x}_{current}$ ;
4   Sample random  $\{A_{t,H}^{(1)}, \dots, A_{t,H}^{(m)}\}$ ;
5   for  $i \leftarrow 1$  to  $m$  do
6      $\sigma_t \leftarrow \{\mathbf{x}_t\}$ ;
7     for  $\mathbf{a}_\tau \in A_{t,H}^{(i)}$  do
8        $\mathbf{x}_t \leftarrow \tilde{f}_\phi(\mathbf{x}_t, \mathbf{a}_\tau)$ ;
9        $\sigma_t \leftarrow \sigma_t \cup \{\mathbf{x}_t\}$ ;
10     $J_i \leftarrow J(\sigma_t)$ ; // sol. (3)
11   $j \leftarrow \arg \min(\{J_1, \dots, J_m\})$ ;
12   $a_t^* \leftarrow$  first action in  $A_{t,H}^{(j)}$ ;
13  Execute  $a_t^*$ ;
14  Get current hand state  $\mathbf{x}_{current}$ ;
15 while  $\|\tilde{\Gamma}_\theta(\mathbf{x}_{current}) - \mathbf{p}_g\| < \epsilon$ ;

```

VI. EXPERIMENTS

To validate the proposed approach and analyze the accuracy of various state formulations, we have designed and built an experimental system. We use only the two opposing fingers of the three-finger OpenHand Model-O underactuated hand [33]. The experimental system consists of the hand and an automated reset mechanism as seen in Figure 7. Hence, training data \mathcal{P} is collected in episodes where, at each episode, the robot grasps the object, performs in-hand manipulation with random actions until it drops, and then repeats. Once dropped, a thin string that runs through a hole in the center of the object pulls the object into the reach of the fingers toward a new grasp. The system is operated by the Robot Operating System (ROS). During manipulation, data stream of various features is available in 10 Hz and recorded including instantaneous actions, actuator angles and torques, and tactile signals from the FIX-finger. In addition, the pose of the object relative to the hand base is recorded using fiducial markers and cameras. We test eight prismatic objects. Six PLA objects with different cross-sections seen in Figure 8: circular (15 mm and 10 mm radius), elliptical, square, crescent and arbitrary curved. Two additional cylinders of 15 mm radius were used: one flexible (printed with Thermoplastic Polyurethane) and one wrapped with sandpaper for higher friction. Open-source FIX-finger model, fabrication instructions, datasets, open-source




Fig. 7. Automated experimental setup based on the OpenHand model-O underactuated hand.



Fig. 8. Six prismatic test objects.

code and CAD models of the test objects are available online¹.

A. Pose estimation

We begin by observing pose estimation accuracy for the cylinder of radius 15 mm. The collected training data comprises of $N = 92,700$ points recorded by randomly applying actions to the hand in multiple episodes. An additional 10,000 test points were collected in separated episodes and were not included in the training set in any way. Hyper-parameters optimization of the three tested observation models yielded the following architectures. The FC-NN is composed of 3 layers with 300 neurons and a ReLU activation each. A dropout of 30%, early stopping and an L2 regularizer with factor 10^{-5} were included to reduce overfitting. For the LSTM based RNN, the optimization yielded a network of 2 layers, 256 hidden neurons each and a sequence of $w = 5$ past states. Next, the output is fed into a 256 neurons FC layer with ReLU activation. For both NN models, we used the Adam optimizer along with the MSE loss function. GP with a radial basis kernel used 100 nearest neighbors for local regression.

In our experiments, nine possible feature combinations for a state representation are analyzed. The combinations are

¹https://github.com/osheraz/haptic_pose_estimation

TABLE I
POSE ESTIMATION ACCURACY FOR A CYLINDER (15 MM) WITH VARIOUS FEATURE COMBINATIONS AND REGRESSION METHODS

Comb.	Features				GP		FC-NN		LSTM	
	Actuator angles	Actuator loads	FIX fingers	Initial pose	Position RMSE (mm)	Orientation RMSE (deg)	Position RMSE (mm)	Orientation RMSE (deg)	Position RMSE (mm)	Orientation RMSE (deg)
1		✓			22.1 ± 1.7	20.1 ± 1.8	20.6 ± 2.8	19.8 ± 2.7	18.8 ± 2.3	18.1 ± 2.4
2			✓		10.5 ± 1.4	8.7 ± 1.2	10.6 ± 1.4	6.8 ± 1.9	10.1 ± 1.3	6.8 ± 1.8
3		✓	✓		8.5 ± 1.0	8.0 ± 1.1	8.6 ± 1.2	5.9 ± 0.8	7.7 ± 1.6	6.1 ± 0.8
4	✓				6.6 ± 1.5	8.8 ± 3.4	6.5 ± 1.2	6.1 ± 1.4	5.8 ± 0.6	6.0 ± 0.8
5	✓	✓			6.0 ± 0.5	7.3 ± 1.8	5.7 ± 1.0	6.0 ± 1.3	5.3 ± 0.6	5.5 ± 1.0
6	✓		✓		5.1 ± 0.6	7.0 ± 1.0	4.4 ± 0.5	5.2 ± 0.6	5.1 ± 0.8	4.9 ± 0.8
7	✓	✓	✓		4.6 ± 1.2	5.5 ± 1.8	4.3 ± 0.4	5.1 ± 0.8	4.0 ± 0.2	5.0 ± 0.5
8	✓	✓		✓	4.3 ± 1.0	5.3 ± 2.0	4.3 ± 0.7	5.0 ± 0.9	4.2 ± 0.8	4.8 ± 1.0
9	✓	✓	✓	✓	4.1 ± 1.0	4.9 ± 1.7	3.9 ± 0.6	4.4 ± 0.8	3.1 ± 0.6	3.0 ± 0.6

listed in Table I. For comparison, we have also included two feature combinations where the initial pose of the object at the moment of grasp is known. Such scenario can occur when a camera has an initial view prior to manipulating the hand in a confined space. In addition, Table I presents accuracy results for the object pose estimation during in-hand manipulation. The accuracy for translation and orientation is the RMSE between predicted and ground truth ones. Figure 9 shows a trajectory example of a grasped object during in-hand manipulation and the predicted poses along it. Actuator angles clearly provide meaningful information that is essential for accurate estimation. Moreover, the results show that the angles alone can provide relatively good accuracy. Actuator loads provide somewhat accuracy improvement for GP and FC-NN. However, including also tactile information from the FIX-finger exhibits better accuracy for all models. While the FIX-fingers cannot provide accurate predictions by their own (Comb. 2), the results indicate that accuracy is limited by the resolution exhibited in Figure 2c. The results also show that including sequential past states with LSTM improves accuracy compared to single state models (GP and FC-NN) over all feature combinations. Also, having a glance at the ground truth of the initial pose (Comb. 9), if available, exhibits accuracy improvement. The average velocity of the object is 0.62 mm/sec and states are sampled at 10Hz while the LSTM model, for instance, provides a prediction in 9 milliseconds. Hence, an instantaneous sampled state sufficiently corresponds to the predicted pose.

Table II presents additional accuracy results for the remaining objects. The results show that the position estimation exhibits similar or slightly higher errors than the 15 mm cylinder. Results were fairly similar between the flexible and rigid cylinders. Additionally, sliding was not visually observed for the high friction cylinder and, therefore, the model was able to produce more accurate estimations. Furthermore, orientation is shown to be harder to predict as the complexity of the object increases (e.g., crescent and arbitrary objects) and the available haptics do not provide sufficient information. It is assumed that higher resolution tactile sensors such as in [24], [25] would better extract shape features and yield lower orientation errors. Nevertheless and for all objects, having an initial glance at the grasp pose enables accurate orientation prediction.

We next analyze the accuracy with regards to data size. Recall that the total amount of data recorded was $N = 92,700$




Fig. 9. An example of the (left) position and (right) orientation estimation with three observation models over a test trajectory of the cylindrical object. The gray shading illustrates the workspace of the hand approximated with the collected data.




Fig. 10. Pose estimation accuracy with regards to the percentage of total training data used with GP, FC-NN and LSTM-based RNN.

points. We note that it took approximately 12 hours for automated collection of all data. The accuracy with regards to data size is observed when considering feature combination 7. Hence, we arranged all data sequentially with no shuffling and repeatedly trained the model for a varying portion of the data. The mean prediction accuracy over the test data with the three regression methods can be seen in Figure 10. The accuracy improves with the increase of data while the error reaches near saturation at about 60% of the data (corresponds to approximately 7 hours of data collection). The performance provided by the haptics of the hand indicates the ability to perform coarse in-hand manipulation tasks which is experimented in the next sub-section.

Lastly, we study generalization properties of the LSTM-based observation model for different objects that were not included in the training set. The generalization is analyzed by training a model on one object and evaluating it on test data from the remaining ones. This is repeated for all objects. Figure 11 shows the cross transfer accuracy across the various objects for feature combination 7. Note that the diagonals of the tables are the object-wise accuracy results presented previously. The study shows that position predictions have a higher transferability than orientation predictions. In general, both position and orientation generalize better when trained over circular objects of varying sizes and textures. Particularly,

TABLE II
POSE ESTIMATION ACCURACY OF THE OBSERVATION MODEL FOR VARIOUS OBJECTS

Object cross-section	Comb.	GP		FC-NN		LSTM	
		Position RMSE (mm)	Orientation RMSE (deg)	Position RMSE (mm)	Orientation RMSE (deg)	Position RMSE (mm)	Orientation RMSE (deg)
Circular ($r = 10\text{ mm}$)	7	5.1 ± 2.1	7.0 ± 4.6	4.6 ± 0.5	7.2 ± 1.1	4.4 ± 0.5	7.1 ± 1.1
	9	4.8 ± 2.0	6.5 ± 3.2	4.0 ± 0.8	6.0 ± 2.0	4.0 ± 0.8	5.3 ± 2.0
Circular ($r = 15\text{ mm}$, Flexible)	7	5.5 ± 2.1	5.8 ± 2.8	4.7 ± 0.5	4.1 ± 0.6	4.6 ± 0.4	4.0 ± 0.5
	9	3.9 ± 2.4	5.1 ± 1.7	3.9 ± 0.4	4.7 ± 0.6	3.6 ± 0.3	3.9 ± 0.5
Circular ($r = 15\text{ mm}$, High friction)	7	4.0 ± 1.1	5.5 ± 2.2	4.1 ± 0.4	5.3 ± 0.7	3.9 ± 0.4	5.0 ± 0.6
	9	3.5 ± 1.0	4.4 ± 1.5	3.2 ± 0.4	4.9 ± 0.6	2.9 ± 0.5	3.9 ± 0.4
Square ($20 \times 20\text{ mm}$)	7	5.6 ± 2.4	10.7 ± 6.7	5.0 ± 0.6	11.6 ± 1.3	4.6 ± 0.6	11.3 ± 2.0
	9	5.1 ± 2.0	6.7 ± 4.8	4.6 ± 0.6	7.5 ± 1.6	4.4 ± 0.6	6.1 ± 1.0
Elliptical ($r_1 = 40\text{ mm}$, $r_2 = 20\text{ mm}$)	7	7.5 ± 2.0	30.1 ± 11.2	7.4 ± 1.1	26.1 ± 1.4	6.2 ± 0.7	17.4 ± 3.6
	9	3.8 ± 1.7	5.7 ± 1.8	4.3 ± 0.5	6.0 ± 0.5	4.0 ± 0.7	5.6 ± 0.6
Crescent ($r_{out} = 50\text{ mm}$, $r_{in} = 40\text{ mm}$)	7	6.2 ± 4.3	27.3 ± 7.1	7.6 ± 1.5	28.2 ± 5.1	7.2 ± 0.5	21.6 ± 4.0
	9	5.8 ± 4.0	6.8 ± 6.2	5.1 ± 1.0	9.6 ± 1.6	4.5 ± 0.8	7.3 ± 1.8
Arbitrary	7	5.9 ± 2.6	33.1 ± 4.1	6.5 ± 0.5	39.2 ± 6.1	5.0 ± 0.2	28.6 ± 3.4
	9	5.8 ± 2.5	9.0 ± 6.6	5.5 ± 0.6	11.2 ± 1.9	4.4 ± 0.4	7.7 ± 3.0

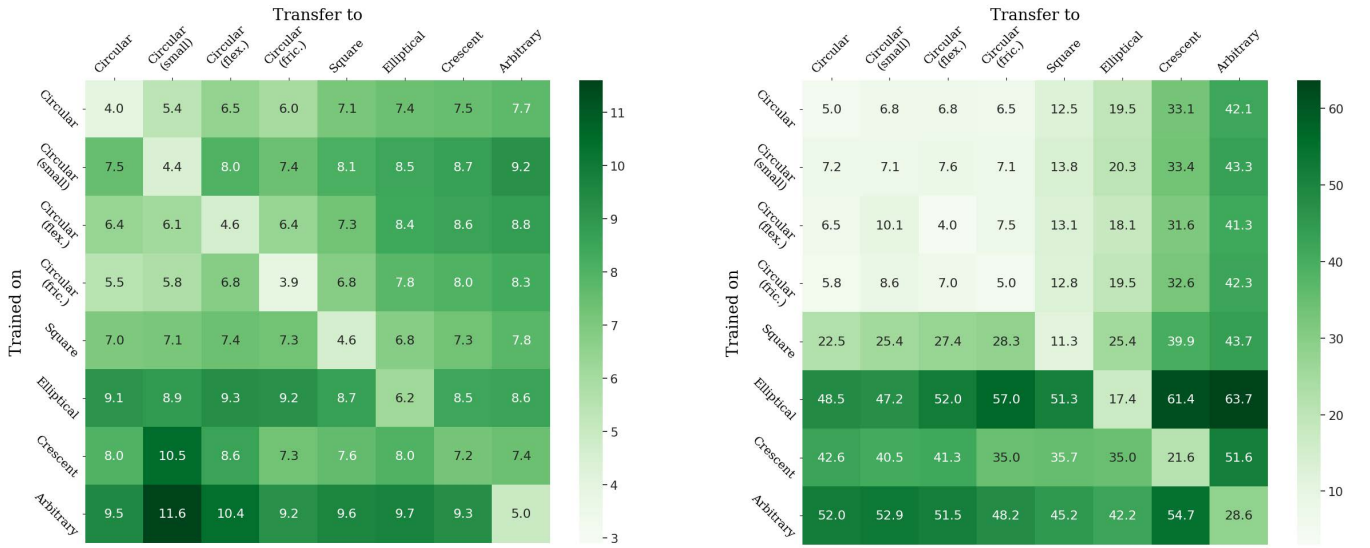


Fig. 11. Cross transfer pose estimation accuracy between the tested objects with feature combination 7. Transfer scores include (left) position RMSE in mm and (right) orientation errors in $degrees$.



Fig. 12. Cross transfer pose estimation accuracy between the tested objects when including the initial grasp pose in the state (feature combination 9). Transfer scores include (left) position RMSE in mm and (right) orientation errors in $degrees$.

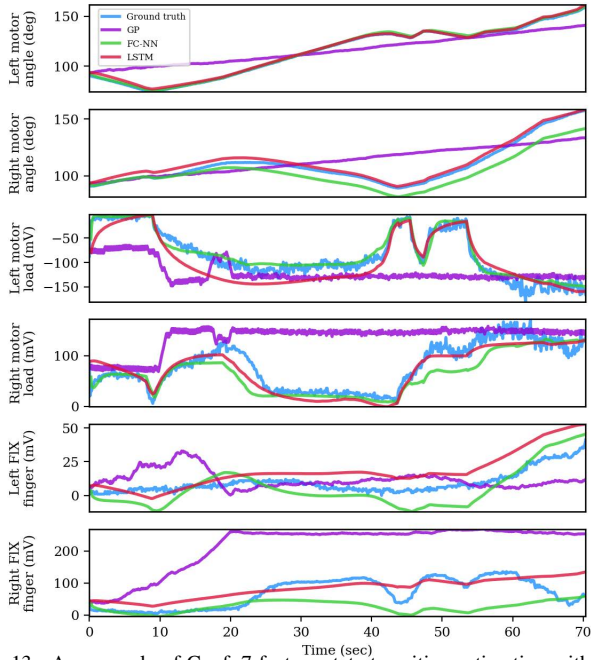


Fig. 13. An example of Conf. 7 feature state transition estimation with three transition models over a test trajectory in \mathcal{C} .

training over a circular object is sufficient for acquiring a relatively good position accuracy for all objects. The orientation, however, has more difficulty in generalization to more complex objects. On the other hand, Figure 12 presents the cross transfer accuracy when also including the initial pose to the state (feature combination 9). Clearly, the inclusion of the initial pose significantly improves generalization for both position and orientation predictions, and yields low errors for all objects.

B. Closed loop planning

In this section, we test the MPC closed loop control approach discussed in Section V-C. We first use the collected data to train transition model (1) with feature combination 7. As discussed in Section V-A, the data collected for the observation model is used to train a model $\Delta\tilde{\mathbf{x}}_t = \hat{f}_\phi(\mathbf{x}_t, \mathbf{a}_t)$ and compare between GP, FC-NN and LSTM with $k = 2$ past states. The hyper-parameters in ϕ were optimized to reduce transition error. We note that the transitions of the model are made in feature space \mathcal{C} while the pose accuracy is computed after using the LSTM observation model presented above. Figure 13 shows a usage example of the three transition models in the feature state space over a test trajectory. Tracking along the ground truth signals is best performed by the LSTM model. Additionally, Figure 14 presents the mean error (over the test data) of the three models in an open-loop fashion for an horizon of up to 5.0 seconds. Note that the prediction error is composed of the transition accuracy on top of inherent errors of the observation model. The LSTM model evidently provides better accuracy and is, therefore, used in the following experiments.

We now observe the performance of the MPC-Critic over a large set of goals in the workspace of the hand. The workspace of the hand is crescent-shaped as illustrated in Figure 9.




Fig. 14. (left) Position and (right) orientation transition accuracy of the three tested models. Predictions are performed in open-loop for an horizon of 50 steps (5.0 seconds).

TABLE III
ROLL-OUT RESULTS

		Goal error (mm)	Path len. (mm)	Success rate
Circular 15mm	OL	6.7 ± 2.0	15.4 ± 9.5	40%
	HS	6.9 ± 2.9	19.0 ± 11.8	60%
	MPC w/o critic	5.0 ± 2.5	36.8 ± 20.9	58%
	MPC-Critic	4.7 ± 2.8	41.6 ± 29.8	70%
	VS	-	30.6 ± 16.8	73%
	MPC-Vis	-	41.2 ± 15.5	82%
Elliptical	OL	7.8 ± 1.2	14.5 ± 1.0	38%
	HS	7.1 ± 3.0	17.9 ± 10.7	56%
	MPC w/o critic	7.0 ± 1.8	29.4 ± 14.2	60%
	MPC-Critic	6.5 ± 2.3	33.5 ± 20.1	74%
	VS	-	36.6 ± 10.5	78%
	MPC-Vis	-	37.5 ± 13.4	80%
Crescent	OL	8.2 ± 1.1	16.0 ± 5.6	40%
	HS	8.8 ± 3.1	28.7 ± 12.5	60%
	MPC w/o critic	7.5 ± 2.0	31.2 ± 12.5	58%
	MPC-Critic	7.3 ± 1.0	36.5 ± 22.4	70%
	VS	-	37.5 ± 14.2	70%
	MPC-Vis	-	39.5 ± 17.2	76%

Furthermore, previous work has shown that some goals in the workspace, mostly near the boundaries, are harder to reach due to the risk of getting stuck, slide or drop of the object [14]. The MPC-Critic is implemented according to Algorithm 1 with horizon $H = 10$ and weight parameters $w_1 = 0.8$ and $w_2 = 0.2$. We benchmark the MPC-Critic with three baseline approaches: Open-loop (OL), HS and MPC without critic. For the OL, we incorporate the MPC-Critic in an off-line manner where a set of actions is planned to reach the goal from the mean of all initial grasp locations in the data. Once the object is grasped, the planned actions are rolled-out in open-loop. The HS is implemented as described in Section V-B where $K_x = K_y = 200$. The MPC without critic considers only the first component in the sum of (3) with $w_1 = 1$ and, therefore, aims to shorten the path to the goal. Furthermore, we compare our results to control when visual feedback is available: VS as described in Section V-B and MPC with a vision-based transition model. For the latter, termed *MPC-Vis*, we use a transition model proposed in [10] where the state comprises of the camera perceived pose and actuator torques. *MPC-Vis* is implemented similar to MPC without critic. The experiment is conducted by randomly sampling 50 poses in the workspace of the hand and attempting to reach them after grasping the object in an arbitrary initial grasp. A roll-out attempt is ended if the object falls or the observation model predicts reaching to the goal. Therefore, a roll-out success is considered if the object reaches within 4 mm to the goal at the end of the attempt. This goal threshold is chosen based to the accuracy limit of combination 7 with LSTM exhibited in the pose estimation results. For the non-visual methods, vision is used solely to estimate and report the true error of reaching a goal.

Table III presents averaged results for roll-outs with all




Fig. 15. Four roll-out examples where the circular, elliptical and crescent objects are manipulated between several intermediate goals (green markers) using the MPC-Critic. The rolled-out path and position estimation are seen in red and cyan curves, respectively.

methods of the circular (15mm), elliptical and crescent objects. The table shows the mean distance to the goal over all trials after finishing the roll-outs. We first compare between the four non-visual methods. MPC-Critic exhibits the lowest error. In addition, the results show that MPC-Critic is able to reach farther goals with higher success rate. The critic has sparse data in regions which are harder to cross due to high probability of dropping the object or actuator overload. Hence, the critic has lower accuracy in those regions such that the MPC-Critic inherently avoids them. In addition, the goals were randomly sampled and the system is not guaranteed to be able to reach them. Figure 15 shows roll-out demonstrations with the MPC-Critic where the object must move between several intermediate goals. The accuracies of the visual-based methods, MPC-Vis and VS, are equal to the success tolerance and, therefore, not reported. However and as expected, they exhibit slightly higher success rates than MPC-Critic due to more accurate pose feedback. Nevertheless, the results validate the proposed approach and show feasible manipulations with solely haptic feedback.

We have also conducted a peg-in-a-hole demonstration in which the objects are manipulated within a confined cabinet into slots. The positions of the slots relative to the hand are assumed to be known. While not in our scope, further work is required in order to sense the target hole, given the general location, using the available haptic feedback. Images of the demonstration are seen in Figure 1. Videos of the demonstration and other experiments are included in the supplementary video.

VII. CONCLUSIONS

In this letter, we have investigated the ability to track and manipulate an object within a two-finger underactuated hand solely based on haptics without relying on visual perception. We explored the state representation of an observation model using the kinesthetic information available to the hand along with novel low-cost tactile fingers. This observation is valuable since vision is often limited, i.e., working in an occluded environment. Adding an initial pose glance at the moment of grasping has shown to provide an additional accuracy improvement. Additionally, we proposed a data-driven MPC approach which utilized learnt observation and transition models to reach desired goals within the workspace of the hand. The MPC approach reasons about the accuracy of the observation model using a critic and attempts to avoid erroneous regions. Results have shown the ability to estimate the pose of several test objects and manipulate them to desired goals with less than

5 mm accuracy. Moreover, generalization to various object was shown feasible. The results validate that in-hand manipulation is feasible with solely haptic feedback. However, the accuracy achieved by the system may not be sufficient in more complex tasks such as inserting a key into a lock. These may require higher resolution hardware. An interesting extension to the work may include more intelligent exploration methods that can increase data quality and be more sample-efficient.

REFERENCES

- [1] J. Bae, S. Park, J. Park, M. Baeg, D. Kim, and S. Oh, "Development of a low cost anthropomorphic robot hand with high capability," in *IEEE/RSJ Int. Conf. on Intel. Rob. and Sys.*, 2012, pp. 4776–4782.
- [2] A. M. Dollar and R. D. Howe, "The highly adaptive sdm hand: Design and performance evaluation," *The International Journal of Robotics Research*, vol. 29, no. 5, pp. 585–597, 2010.
- [3] Y. Bai and C. K. Liu, "Dexterous manipulation using both palm and fingers," pp. 1560–1565, 2014.
- [4] R. Michalec and A. Micaelli, "Stiffness modeling for multi-fingered grasping with rolling contacts," in *IEEE-RAS International Conference on Humanoid Robots*, 2010, pp. 601–608.
- [5] L. U. Odhner and A. M. Dollar, "Dexterous manipulation with underactuated elastic hands," in *IEEE Int. Conf. on Rob. and Aut.* IEEE, May 2011, pp. 5254–5260.
- [6] —, "Stable, open-loop precision manipulation with underactuated hands," *Int. J. of Rob. Res.*, vol. 34, no. 11, pp. 1347–1360, Sep 2015.
- [7] B. Calli and A. M. Dollar, "Vision-based precision manipulation with underactuated hands: Simple and effective solutions for dexterity," in *IEEE/RSJ Int. Conf. on Intel. Rob. and Sys.*, 2016, pp. 1012–1018.
- [8] B. Calli, K. Srinivasan, A. Morgan, and A. M. Dollar, "Learning modes of within-hand manipulation," in *IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 3145–3151.
- [9] R. Ozawa, S. Arimoto, S. Nakamura, and J.-H. Bae, "Control of an object with parallel surfaces by a pair of finger robots without object sensing," *IEEE Trans. on Robotics*, vol. 21, no. 5, pp. 965–976, 2005.
- [10] A. Sintov, A. S. Morgan, A. Kimmel, A. M. Dollar, K. E. Bekris, and A. Boularias, "Learning a state transition model of an underactuated adaptive hand," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1287–1294, April 2019.
- [11] J. Borras and A. M. Dollar, "A parallel robots framework to study precision grasping and dexterous manipulation," in *IEEE Int. Conf. on Rob. and Aut.* IEEE, May 2013, pp. 1595–1601.
- [12] A. Sintov, A. Kimmel, K. Bekris, and A. Boularias, "Motion planning with competency-aware transition models for underactuated adaptive hands," in *IEEE Int. Conf. on Rob. & Aut.*, Paris, 2020, pp. 7761–7767.
- [13] A. Kimmel*, A. Sintov*, B. Wen, A. Boularias, and K. Bekris, "Belief-space planning using learned models with application to underactuated hands," in *Int. Symposium on Robotics Research*, Hanoi, 2019.
- [14] A. S. Morgan, W. G. Bircher, and A. M. Dollar, "Towards generalized manipulation learning through grasp mechanics-based features and self-supervision," *IEEE Transactions on Robotics*, vol. 37, no. 5, pp. 1553–1569, 2021.
- [15] A. S. Morgan, K. Hang, and A. M. Dollar, "Object-agnostic dexterous manipulation of partially constrained trajectories," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5494–5501, 2020.
- [16] B. Wen, C. Mitash, S. Soorian, A. Kimmel, A. Sintov, and K. E. Bekris, "Robust, occlusion-aware pose estimation for objects grasped by adaptive hands," in *IEEE Intl. Conf. on Rob. & Aut.*, 2020.

- [17] A. S. Morgan, B. Wen, J. Liang, A. Boularias, A. M. Dollar, and K. Bekris, "Vision-driven compliant manipulation for reliable, high-precision assembly tasks," in *Robotics: Science and Systems*, 2021.
- [18] V. Fonseca, T. E. Alves de Oliveira, and E. Petriu, "Estimating the orientation of objects from tactile sensing data using machine learning methods and visual frames of reference," *Sensors*, p. 2285, 2019.
- [19] J. Carter and D. Fournay, "Research based tactile and haptic interaction guidelines," in *Guidel. On Tactile and Hap. Inter.*, 01 2005, pp. 84–92.
- [20] G. Rouhafzay and A. Cretu, "Object recognition from haptic glance at visually salient locations," *IEEE Transactions on Instrumentation and Measurement*, vol. 69, no. 3, pp. 672–682, 2020.
- [21] S. Pohongkam and J. Srinonchat, "Tactile object recognition for humanoid robots using new designed piezoresistive tactile sensor and dcnn," *Sensors*, vol. 21, no. 18, 2021.
- [22] M. Liu, Y. Zhang, J. Wang, N. Qin, H. Yang, K. Sun, J. Hao, L. Shu, J. Liu, Q. Chen, P. Zhang, and T. H. Tao, "A star-nose-like tactile-olfactory bionic sensing array for robust object recognition in non-visual environments," *Nature Communications*, vol. 13, no. 79, 2022.
- [23] C. Chorley, C. Melhuish, T. Pipe, and J. Rossiter, "Development of a tactile sensor based on biologically inspired edge encoding," in *International Conference on Advanced Robotics*, 2009, pp. 1–6.
- [24] E. Donlon, S. Dong, M. Liu, J. Li, E. Adelson, and A. Rodriguez, "Gelslim: A high-resolution, compact, robust, and calibrated tactile-sensing finger," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1927–1934, 2018.
- [25] N. Lepora and J. Lloyd, "Optimal deep learning for robot touch: Training accurate pose models of 3d surfaces and edges," *IEEE Robotics and Automation Magazine*, vol. 27, no. 2, Apr. 2020.
- [26] B. Calli and A. M. Dollar, "Vision-based model predictive control for within-hand precision manipulation with underactuated grippers," in *IEEE International Conference on Robotics and Automation*, May 2017, pp. 2839–2845.
- [27] B. Calli, A. Kimmel, K. Hang, K. Bekris, and A. Dollar, "Path planning for within-hand manipulation over learned representations of safe states," in *International Symposium on Experimental Robotics*, Buenos Aires, Argentina, 2018.
- [28] Z. Kappasov, J.-A. Corrales, and V. Perdereau, "Tactile sensing in dexterous robot hands — review," *Robotics and Autonomous Systems*, vol. 74, pp. 195–220, 2015.
- [29] J. T. Muth, D. M. Vogt, R. L. Truby, Y. Mengüç, D. B. Kolesky, R. J. Wood, and J. A. Lewis, "Embedded 3d printing of strain sensors within highly stretchable elastomers," *Advanced Materials*, vol. 26, no. 36, pp. 6307–6312, 2014.
- [30] J. H. Lee, "Model predictive control: Review of the three decades of development," *International Journal of Control, Automation and Systems*, vol. 9, no. 3, p. 415, Jun 2011.
- [31] V. Begoc, S. Krut, E. Dombre, C. Durand, and F. Pierrot, "Mechanical design of a new pneumatically driven underactuated hand," in *Proceedings 2007 IEEE International Conference on Robotics and Automation*, 2007, pp. 927–933.
- [32] S. Abondance, C. B. Teeple, and R. J. Wood, "A dexterous soft robotic hand for delicate in-hand manipulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5502–5509, 2020.
- [33] R. R. Ma and A. M. Dollar, "Yale openhand project: Optimizing open-source hand designs for ease of fabrication and adoption," *IEEE Rob. & Aut. Mag.*, vol. 24, pp. 32–40, 2017.
- [34] A. Rocchi, B. Ames, Z. Li, and K. Hauser, "Stable simulation of underactuated compliant hands," in *ICRA*, May 2016, pp. 4938–4944.
- [35] I. Hussain, F. Renda, Z. Iqbal, M. Malvezzi, G. Salvietti, L. Seneviratne, D. Gan, and D. Prattichizzo, "Modeling and Prototyping of an Underactuated Gripper Exploiting Joint Compliance and Modularity," *IEEE Rob. and Aut. Let.*, vol. 3, no. 4, pp. 2854–2861, Oct 2018.
- [36] K. Hang, W. G. Bircher, A. S. Morgan, and A. M. Dollar, "Manipulation for self-identification, and self-identification for better manipulation," *Science Robotics*, vol. 6, no. 54, 2021.
- [37] A. Bicchi, M. Bergamasco, P. Dario, and A. Fiorillo, "Integrated tactile sensing for gripper fingers," in *Proc. Int. Conf. on Robot Vision and Sensory Control*, I. F. S. Publications Ltd. Zuerich, Switzerland: I. F. S. Publications Ltd., 1988.
- [38] S. Funabashi, A. Schmitz, T. Sato, S. Somlor, and S. Sugano, "Robust in-hand manipulation of variously sized and shaped objects," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 257–263.
- [39] A. Melnik, L. Lach, M. Plappert, T. Korthals, R. Haschke, and H. Ritter, "Using tactile sensing to improve the sample efficiency and performance of deep deterministic policy gradients for simulated in-hand manipulation tasks," *Frontiers in Robotics and AI*, vol. 8, 2021.
- [40] Z. Su, J. Fishel, T. Yamamoto, and G. Loeb, "Use of tactile feedback to control exploratory movements to characterize object compliance," *Frontiers in neurobotics*, vol. 6, p. 7, 07 2012.
- [41] J. Bimbo, L. D. Seneviratne, K. Althoefer, and H. Liu, "Combining touch and vision for the estimation of an object's pose during manipulation," in *IEEE/RSJ Int. Conf. on Intel. Rob. & Sys.*, 2013.
- [42] M. A. Lee, Y. Zhu, K. Srinivasan, P. Shah, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg, "Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks," in *International Conference on Robotics and Automation (ICRA)*, 2019, pp. 8943–8950.
- [43] S. Funabashi, S. Ogasa, T. Isobe, T. Ogata, A. Schmitz, T. P. Tomo, and S. Sugano, "Variable in-hand manipulations for tactile-driven robot hand via cnn- lstm," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, p. 9472–9479.
- [44] M. Du, O. Y. Lee, S. Nair, and C. Finn, "Play it by ear: Learning skills amidst occlusion through audio-visual imitation learning," in *Robotics: Science and Systems*, 2022.
- [45] T. Boroushaki, I. Perper, M. Nachin, A. Rodriguez, and F. Adib, "Rfusion: Robotic grasping via rf-visual sensing and learning," in *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*, 2021, p. 192–205.
- [46] J. Tegin and J. Wikander, "Tactile sensing in intelligent robotic manipulation—a review," *Indus. Robot: An Int. Jou.*, vol. 32, 2005.
- [47] M. C. Koval, M. R. Dogar, N. S. Pollard, and S. S. Srinivasa, "Pose estimation for contact manipulation with manifold particle filters," in *IEEE/RSJ Int. Conf. on Intel. Rob. and Sys.*, 2013, pp. 4541–4548.
- [48] J. Bimbo, S. Luo, K. Althoefer, and H. Liu, "In-hand object pose estimation using covariance-based tactile to geometry matching," *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 570–577, 2016.
- [49] P. Sodhi, M. Kaess, M. Mukadam, and S. Anderson, "Learning tactile models for factor graph-based estimation," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2021.
- [50] A. S. Vazquez, R. Fernandez, A. Lopez, E. Valero, I. Payo, and A. Adan, "In-hand object localization: Simple vs. complex tactile sensors," in *SENSORS, 2014 IEEE*, 2014, pp. 1710–1713.
- [51] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, The MIT Press, 2005.
- [52] A. Nagabandi, G. Kahn, R. S. Fearing, and S. Levine, "Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning," in *IEEE Int. Conf. Rob. & Aut.*, 2018, pp. 7559–7566.

VIII. BIOGRAPHY

Osher Azulay received his B.Sc. and M.Sc. in Mechanical Engineering from Ben-Gurion University at 2019 and 2020, respectively. He is currently a Ph.D. student at the School of Mechanical Engineering at Tel-Aviv University. His research interests include in-hand manipulations, robot learning and tactile sensing.



Inbar Ben-David received her B.Sc. in Mechanical Engineering from Ben-Gurion and MBA in Business Administration from the College of Management Academic Studies at 2017 and 2019, respectively. She is currently a M.Sc student at the School of Mechanical Engineering at Tel-Aviv University. Her research interests include robot design and genetic optimization.



Avishai Sintov received his B.Sc. M.Sc. and Ph.D. in 2008, 2012 and 2016, respectively, from the department of Mechanical Engineering at Ben-Gurion University of the Negev. He is currently an assistant professor at the School of Mechanical Engineering at Tel-Aviv University. Previously, he was a post-doc fellow at the University of Illinois at Urbana-Champaign (2016-2018) and Rutgers University (2018-2019). His research interests include Human-Robot Collaboration dynamic manipulations, grasping, regrasping, constrained systems, machine learning and robot design.



ing and robot design.