

ARTIFICIAL INTELLIGENCE

Manipulation for self-Identification, and self-Identification for better manipulation

Kaiyu Hang*, Walter G. Bircher, Andrew S. Morgan, Aaron M. Dollar

The process of modeling a series of hand-object parameters is crucial for precise and controllable robotic in-hand manipulation because it enables the mapping from the hand's actuation input to the object's motion to be obtained. Without assuming that most of these model parameters are known a priori or can be easily estimated by sensors, we focus on equipping robots with the ability to actively self-identify necessary model parameters using minimal sensing. Here, we derive algorithms, on the basis of the concept of virtual linkage-based representations (VLRs), to self-identify the underlying mechanics of hand-object systems via exploratory manipulation actions and probabilistic reasoning and, in turn, show that the self-identified VLR can enable the control of precise in-hand manipulation. To validate our framework, we instantiated the proposed system on a Yale Model O hand without joint encoders or tactile sensors. The passive adaptability of the underactuated hand greatly facilitates the self-identification process, because they naturally secure stable hand-object interactions during random exploration. Relying solely on an in-hand camera, our system can effectively self-identify the VLRs, even when some fingers are replaced with novel designs. In addition, we show in-hand manipulation applications of handwriting, marble maze playing, and cup stacking to demonstrate the effectiveness of the VLR in precise in-hand manipulation control.

INTRODUCTION

Manipulation is one of the most common actions for robots to physically interact with the world. As we see an increasing amount of real-world applications being deployed in recent years, robotic manipulation is still a challenging problem that involves many subproblems yet to be addressed, such as complex system dynamics, multimodal sensing, planning and control, and hardware design, particularly for manipulation in unstructured environments (1–3). For grasp-based manipulation, numerous tasks can be achieved purely through arm motions, such as the most common pick-and-place tasks. However, relying only on arm motions is often undesirable because it can be energy consuming, unsafe in human-centered environments, and unnecessarily complicated (4). To this end, another dimension of manipulation, in-hand dexterity, has been explored to provide smaller-scaled yet more focused manipulation of objects using only finger motions (5). This kind of dexterity is used heavily by humans, and that human dexterity has inspired a large number of robotic hand designs promising more dexterity from the finger motions to the objects in-hand reconfigurations (6). However, despite the large number of degrees of freedom featured by many hand designs, coordinating multiple fingers to manipulate objects robustly is a complex system-level problem that still remains open (7).

In this work, to effectively capture the sophisticated mechanics of in-hand manipulation and, in the meantime, keep it generalizable across various hardware platforms, we propose the concept of virtual linkage-based representations (VLRs) to represent contact and linkage-based hand-object systems. On the basis of the VLR, which can rely on parameters that are not directly accessible, we derive algorithms for the system to self-identify necessary parameters through exploratory hand-object interactions and probabilistic reasoning and then, in turn, show that the self-identified VLR endows the hand with precise control of in-hand manipulation.

Background, challenges, and beyond

As a prerequisite of in-hand manipulation, grasp planning initializes the hand-object system with stable contact configurations (8, 9). Although data-driven approaches have gained great success in generating grasps (10–12) and maintaining their stability (13, 14) in challenging scenarios, analytical methods are still necessary in modeling many aspects of grasping, such as optimality and task requirements (15, 16) and kinematic and environmental constraints (17–19). Although learning has been shown to facilitate in-hand manipulation in a number of challenging scenarios—such as highly dexterous hands (20), underactuated hands (21, 22), and even contact-rich grasps (23)—the learned models usually do not provide precise control and are not easily generalizable across platforms or tasks (24). As such, similar to grasp planning, analytical approaches are necessary for dexterous manipulation to model system dynamics (25–27), to adapt to general task requirements (28), and to deal with unseen objects (29). With the aim of precise control of in-hand manipulation, we focus this work on the analytical side of hand-object systems.

In our previous work (30), we showed that by using four cameras to fully observe relevant parameters, we could learn a regression model of the hand-object system to directly map from the hand's actuation to the object motion. Under the same setup, by additionally deriving a model predictive controller, we achieved significant improvement of the motion accuracy on the basis of the learned regression model (31). Alternatively, by using only one camera, we achieved automatic estimation of hand-object configurations (32) by assuming that the hand model, fingertip contact positions, and relevant physical properties (e.g., joint stiffness) are known a priori. In this work, rather than use methods that require more sensors to be incorporated or more prior knowledge to be available, we aim to provide an alternative for robots to self-identify necessary hand-object parameters using minimal sensing. This can be beneficial from a variety of perspectives. First, minimizing the number of sensors in a system reduces the requirement of calibrations among all system components, avoiding non-negligible accumulated errors. Second, the alleviated sensing requirement allows more hardware, especially low-cost designs, to provide more complex manipulation skills. Third,

Copyright © 2021
The Authors, some
rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim
to original U.S.
Government Works

Department of Mechanical Engineering and Material Science, Yale University, New Haven, CT, USA.

*Corresponding author. Email: kaiyu.hang@yale.edu

a self-identifiable system provides not only similar functionality to sensor-rich systems but also higher robustness in complex working scenarios.

Furthermore, beyond traditional concerns, we consider the possibility of recovering hands from damages. As illustrated in Fig. 1, when a damaged part of a hand has to be replaced while there are no identical spare parts available, using a similar but different replacement could be a good backup option. Alternatively, even if self-healing materials can help (33), it is not guaranteed that the robot will recover to the same exact embodiment. In either case, if the system relies on prior information of the hand design, it will not be able to ideally use the recovered hand, because the new embodiments are not necessarily known. In this work, we try to recover the hand’s capabilities on the basis of the VLR, which can be self-identified to enable the hand to use novel part replacements, without necessarily having the full knowledge of their embodiments.

Virtual linkage-based representations

Every hand-object system is different. Even for the same object and hand, every grasp is different in its contact locations, hand configurations, force distributions, etc., which are essential parameters determining the underlying mechanisms of in-hand manipulation. Nonetheless, hand-object systems have many common and invariant

properties effectively describing the interactions between them, despite the differences in parameter values from case to case. Next, we propose to model both hand and object in a unified framework by capturing common and essential properties across such systems, in terms of kinematic constraints.

Being the case in a variety of dexterous in-hand manipulation tasks, we assume fingertip contacts on the object will be maintained throughout the entire process. Centered around the contacts, as exemplified in Fig. 2, we can establish virtual linkages between several key points in the system. Let us denote by $\{p_1, \dots, p_N\}$ the contact points on N fingertips. First, we can create a set of virtual links between all contacts to geometrically represent the constraints between the fingers; i.e., the grasp has to be maintained. Second, suppose that we are interested in controlling the motion of a point on the object, $p^{obj} \in SE(3)$, termed as point of manipulation (POM). An additional set of links can be made between the POM and all contacts. Formally, these links are written as

$$E^C = \{\overline{p_i p_j} \mid p_i, p_j \in \{p_1, \dots, p_N, p^{obj}\} \wedge i \neq j\} \quad (1)$$

Moreover, starting from each contact, we can make virtual links to sequentially connect the joints in that finger toward the base joint

$$E_i^J = \{\overline{q_i^K p_i}\} \cup \{\overline{q_i^k, q_i^{k+1}} \mid k \leq K - 1\} \quad (2)$$

where $q_i^k \in \mathbb{R}^3$ is the k th joint, $k = 1, \dots, K$, of the i th finger, indexed from q_i^1 being the base joint.

Figure 2 shows an example of these virtual linkages with three fingers and two joints per finger. Assuming that the fingertip contact locations are fixed under local motions, all virtual links, therefore, can be considered rigid with constant lengths. As such, the angle configuration of the virtual links $E^J = \cup_i E_i^J$, which are embedded in the fingers, will uniquely describe the motion of POM via simple forward kinematics. Let θ^J be the joint configuration of E^J , we define VLRs for such systems as

$$VLR = (E^C \cup E^J, \theta^J) \quad (3)$$

which is a linkage-based kinematic representation and effectively models the formation and underlying mechanisms of in-hand




Fig. 1. A robot has broken its finger and needs to replace it with a new one. In some scenarios, the replacement is not same as the original. Using our proposed VLRs, the robot can pick any replacement candidate and self-identify it via exploratory interactive actions.




Fig. 2. An example of VLRs. (A) A pear is grasped by a robotic hand. The VLRs are constructed by virtual linkages between the hand joints, contacts, and the POM on the object. VLR 1 (magenta) is an accurate model, whereas VLR 2 (yellow) is inaccurate, exhibiting erroneous contact locations on the fingertips. (B) Given the same actuation input, an inaccurate VLR will directly cause errors in modeling system motions. (C) The hand-object kinematics can be purely represented by the VLRs, without modeling the geometry of the hand or the object.

Downloaded from https://www.science.org at Columbia University on June 29, 2022

manipulation. Different from most traditional approaches, when the fingertip contacts are stable on the object, this representation does not rely on any geometrical information of the hand or the object to describe the system motion; see Fig. 2C.

Note that, as depicted in Fig. 2, VLR is a contact-centered representation—the precision of contact locations is crucial for the model to correctly reflect the physical system (Fig. 2B). This fact poses several major challenges for this representation to be useful. First, to generalize VLRs to various platforms, we need to enable the system to estimate the contact locations without assuming that there are sensors directly providing this information. Furthermore, although we assume the contacts are fixed under local motions, we should be aware that the contacts do vary when small local motions accumulate over time, in the forms of rolling or sliding. Thus, the second challenge is that we have to be able to update the virtual links and joint configurations of the VLR, to ensure it always reflects the true underlying mechanism of the hand-object system over time. Third, because the angle configurations of the distal virtual joints are contact-dependent, they are different from the joint configuration of the physical hand and hence cannot be directly acquired even if there are joint encoders. Last, beyond traditional concerns, an extra challenge is how to recover hands by replacing damaged parts. On the basis of the VLR, we will next introduce how these challenges can be addressed with self-identification.

Self-identification

System identification is often required in the development of mathematical models for dynamical systems. On different complexity levels, when the entire or part of the system mechanism, or some of the model parameters, are infeasible to be directly derived or acquired, system identification aims at modeling and analyzing the relationship between the system inputs and outputs to obtain high-fidelity approximations, which can be then used to control the system without fully investigating the underlying dynamics or interworkings of the system (34, 35). Because of the nonlinear and stochastic nature of many real-world systems, probabilistic filters have been often adopted as a major methodology to provide generalizable and robust estimations (36–38). For robotic systems that must operate autonomously, especially in uncertain environments, automatic system identification is often required and relies heavily on the notion of self.

Self-modeling, self-recognition, or sometimes called self-identification loosely and on different levels refer to the processes where the robot reasons about the existence of its own embodiment, the mapping from its motor behaviors to its embodiment movements, and how its motions interact with the world. Being a core component of such procedures, interactive perception is often required for the robot to perceive and analyze itself and the world resulted from its own actions (39, 40). Along this process, a robot can iteratively figure out its own embodiments, its mirror reflections, its kinematic structures, and even motor behaviors (41–43). In addition, by incorporating external objects into the model, interactive perception can further facilitate hand-object configuration estimation (32, 44, 45), object segmentation (46), grasp planning (47), manipulation skill learning (48), haptic property estimation (49), and even the estimation of complex articulated models (50). Next, on the basis of the proposed VLR, we explain how to achieve self-identification of hand-object systems using interactive perception.

For in-hand manipulation, the VLR alone is merely a kinematic description and not sufficient for modeling the control. We need to

incorporate necessary physical properties, denoted by Ω , into the model, so as to eventually establish the mapping from the hand’s actuation input to the object’s motion, which in its general form can be written as

$$\Gamma : (\xi_{t-1}, u_t) \mapsto \xi_t \quad (4)$$

where $\Gamma(\cdot)$ is the system motion function; $\xi_t = (\text{VLR}_t, \Omega_t)$ is the system configuration, in terms of both kinematics and physics, at time t ; and u_t is the actuation input at time t . In practice, depending on the hardware platform, Ω can include friction coefficients, joint torque limits, and object weight, as needed for describing the hand-object dynamics. Similar to the problem with VLRs, it is possible that there are parameters in Ω not directly available. Therefore, to control the hand-object system in the form of Eq. 4, the robot needs to self-identify both the VLR and Ω .

Algorithm 1. Self-identification by particle filtering

Input: Ξ_{t-1}, u_t

Output: Ξ_t

- 1: for each $\xi_{t-1}^i \in \Xi_{t-1}$ do
- 2: $\xi_t^i \leftarrow \Gamma(\xi_{t-1}^i, u_t)$ ▷Move particle forward
- 3: $(\phi_p^i, \eta_t^i) \leftarrow \Lambda(\xi_t^i)$ ▷Predict particle’s observation
- 4: $(\phi_p^*, \eta_t^*) \leftarrow \text{Sensors.Get}()$ ▷Read real sensors
- 5: $\omega^i \leftarrow \text{Importance}(\|(\phi_p^i, \eta_t^i) - (\phi_p^*, \eta_t^*)\|)$ ▷Particle importance
- 6: end for
- 7: $\Xi_t \leftarrow \text{Resample}(\Xi_{t-1} \propto \{\omega^i\})$ ▷Importance resampling
- 8: return Ξ_t

To this end, let us assume that the pose of the POM, $\phi \in \text{SE}(3)$, can be directly observed because it is the point being manipulated. Moreover, if additional sensors are available to estimate other parameters of the system—such as joint torque sensors, encoders, or tactile sensors—we assume that there is a function mapping the VLR and Ω to their expected corresponding sensor readings

$$\Lambda : \xi_t \mapsto (\phi_t, \eta_t) \quad (5)$$

where η_t denotes all observations from extra sensors. Note that extra sensors are not required in our framework and can be left out if not available or if the mapping in Eq. 5 cannot be formulated. Basically, this mapping requires the system to be able to predict its sensor readings in terms of the system configuration. For example, a tactile sensor’s output can be predicted if the contact location and the force exerted on it can be calculated from the system configuration.

We can now formalize self-identification of the hand-object system into a particle filtering framework, which is a form of sequential importance resampling (51). To figure out the unknowns in the VLR and Ω , the hand will execute exploratory interactions with the object despite not being fully controlled and, along the process, estimate the unknowns from the observations. Denoted by ξ_t^i , a hypothesis of the system configuration at time t , this process is initialized by generating a set of M hypotheses (particles), $\Xi_0 = \{\xi_0^1, \dots, \xi_0^M\}$, to construct a distribution that covers the true system configuration $\xi_t^* = (\text{VLR}_t^*, \Omega_t^*)$. Note that, although a conservative initialization, e.g., a wide distribution, can ensure the coverage of ξ_t^* , it is not preferred because it can negatively affect the estimation accuracy. Because the number of particles is always limited by the computational resources, a good initialization should provide a focused distribution around the true system configuration ξ_t^* with a relatively higher particle

density (52). In practice, if the hypothesis set is incorrectly initialized, the filter will fail, and the only way to correct it is to reinitialize (53).

Thereafter, at every time step $t - 1$, the system will reconfigure with an actuation input u_t and will observe the true sensor outputs (ϕ_t^*, η_t^*) . Meanwhile, every hypothesis $\xi_{t-1}^i \in \Xi_{t-1}$ will be moved forward to ξ_t^i via the system motion function Eq. 4, and the sensor readings of each hypothesis, (ϕ_t^i, η_t^i) , are in turn predicted by Eq. 5. On the basis of the difference between the true sensor readings (ϕ_t^*, η_t^*) and predicted readings (ϕ_t^i, η_t^i) , the likelihood ω^i (importance) of each ξ_t^i can be calculated, and the particles in Ξ_t are resampled with probabilities proportional to their importance values.

Iteratively, because the false hypotheses in the distribution tend to make incorrect predictions of the sensor readings, they would be associated with lower importance values and become less likely to get resampled. Eventually, as we keep the size of the hypothesis set, the false particles will be filtered out, and the mean $|\Xi_t|$ of the hypothesis set will converge toward ξ_t^* (54). One iteration of this procedure is summarized in Algorithm 1. Upon convergence, the unknowns in both the VLR and Ω are self-identified, and the system will be able to use them for manipulation control based on the system motion function Eq. 4.

RESULTS

To evaluate and challenge the VLR and its self-identification, we, in this work, instantiated the proposed system on a Yale Model O underactuated hand (55). As shown in Fig. 3, this hand has three fingers, with an abduction joint between the left finger and the right finger. Each finger has two spring-loaded joints, actuated by only one motor through a tendon. While the tendon length is changed by the motor actions, the joints are reconfigured accordingly with compliance provided by the springs in each joint.

To avoid the need for perception from external devices and extra calibration, we mounted only an in-hand camera to observe the motion of the POM. As such, different from the example in Fig. 2 where the POM is on top of the object, the POM for our hand-object systems

is defined as a point at the bottom of the grasped object, observed via an AprilTag tracker (56). Note that the tag-based POM tracking can be replaced by other vision-based approaches. In addition, as shown in Fig. 3, to collect ground truth information, such as joint configurations, there is an AprilTag attached at the back of each fingertip. These extra tags are observed by calibrated external cameras and are only used for evaluation purposes.

For achieving dexterous in-hand manipulation using the proposed VLR, this hand is challenged by three facts: (i) The hand is not able to directly obtain its joint configuration because of the lack of encoders; (ii) it is difficult to acquire contact information, such as locations, because there are no tactile sensors; and (iii) the hand is underactuated and does not have individual control over each joint. Next, we apply the proposed approach to address these challenges by self-identifying the kinematic VLR and some physical properties Ω . On the basis of the self-identified VLR, we show its in-hand manipulation ability in various real-world tasks. Last, to enable the hand to recover from damages, we evaluate how it can self-identify its VLR with novel finger replacements. In this work, the system was implemented in Python on a machine with Ubuntu 16.04 running on an AMD Ryzen Threadripper 1950X 16-core processor, which allows us to parallelize the self-identification procedure using 32 threads. In all experiments, 60,000 was set as the number of particles, and every iteration of the self-identification took about 6 s. Because every particle is independently processed (see Algorithm 1), the run time of each iteration is about proportional to the number of particles.

Self-identification of model parameters

In the first experiment, we assume that the hand model is fully known, including its geometry, kinematics, and all relevant physical properties. Our task is to establish the VLR for our hand-object system using limited sensing resources. According to the definitions in Eqs. 1 to 3, because the pose of the POM is directly observed by the in-hand camera and the finger link lengths are known, there are two parameters to be self-identified for the VLR: the contact locations and the joint configuration θ^j of the virtual links.

Note that the VLR is a linkage-based representation—the contact locations can be directly obtained from forward kinematics if we can acquire relevant linkage properties, such as the link lengths and joint configurations. Therefore, in our implementation, instead of explicitly modeling the contact locations into self-identification hypotheses, the self-identification procedure will estimate the joint configuration θ^j , and the virtual link lengths, to indirectly identify contact locations.

First, for an unknown hand-object system, self-identification was initialized by generating a set of random hypotheses representing a distribution of the unknown joint configurations and virtual links. An example is shown in Fig. 4A. At the beginning (iteration = 0), the hypothesized virtual links E^j , depicted by magenta lines, were distributed widely to ensure that the true model was




Fig. 3. Experimental setup. (A) The hand-object system instantiated on a Yale Model O hand. The hand is calibrated with three side cameras, which are used for tracking the AprilTags on the back of fingertips to collect ground truth data. (B) The underactuated mechanism of each finger. (C) Five test objects from the YCB dataset (58): apple (#13), wine glass (#23), toy block (#73), flat screwdriver (#44), and gelatin box (#9).

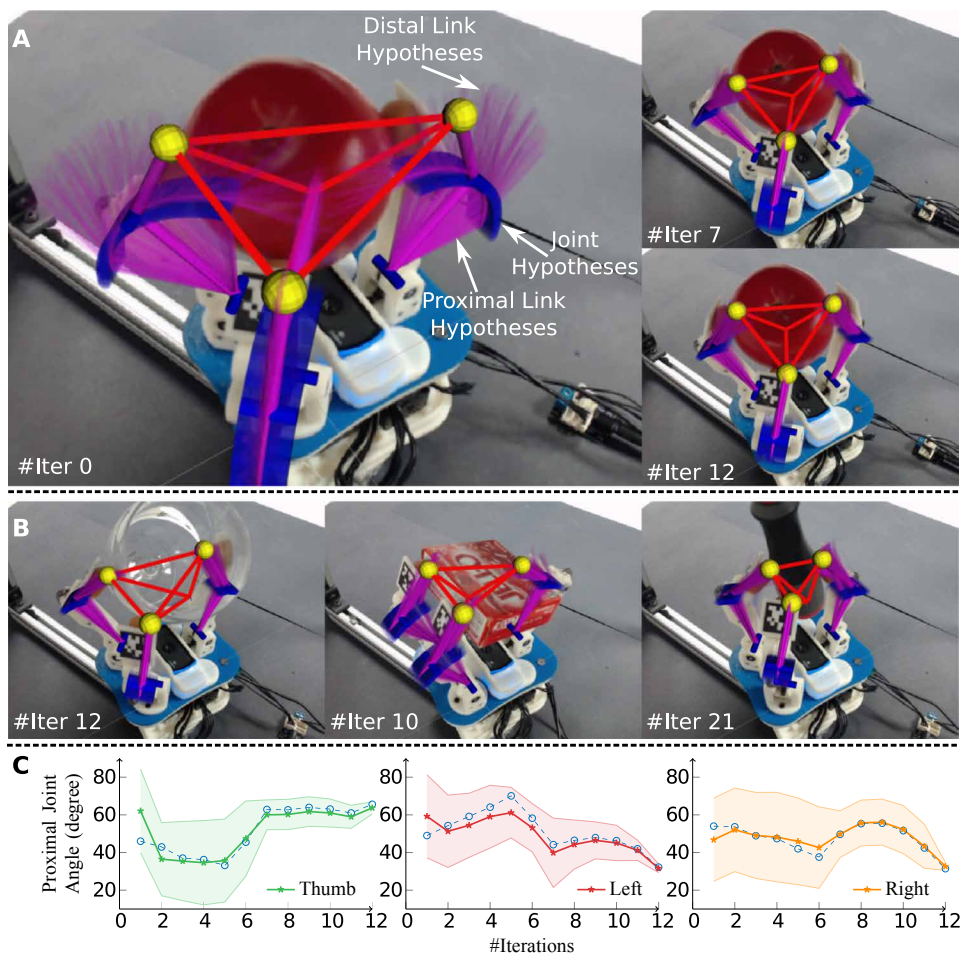


Fig. 4. Self-identification of VLR. (A) An example record of the self-identification process with an apple object at iteration 0, 7, and 12. The magenta lines represent the hypothesized virtual links E^J of different joint configurations, and the blue cylinders indicate the joints of the hypothesized virtual links. The yellow points mark the estimated contact locations, and the red lines represent the estimated virtual links E^C (Eq. 1). (B) Self-identification examples of the wine glass, gelatin box, and flat screwdriver converged at iterations 12, 10, and 21, respectively. (C) An example self-identification process recorded for the angles of proximal joints. The shaded areas show the 95% confidence interval, and the dashed lines mark the ground truth.

contained. The blue cylinders mark the joints connecting each pair of sampled proximal and distal virtual links. Because the hand's geometry is known, all sampled virtual proximal links had the same length; thus, the hypothesized distal virtual joints (blue) of each finger were distributed along an arc of the same radius. However, by definition, the length of a distal virtual link is determined by the contact locations on the fingertips; its length can actually vary as the contact location changes. Therefore, the hypothesized virtual distal links were initially sampled with different lengths in a range, to make sure that the true contact locations were included in the sampled E^J . In the figure, the yellow points depict the contact locations estimated by averaging all sampled E^J . On the basis of the estimated contacts, a complete estimation model of the VLR, although inaccurate at the beginning, can be established by making additional virtual links E^C between the contacts and the POM.

After initialization, following Algorithm 1, the hand has to interact with the object via exploratory actions, so as to iteratively filter out the false hypotheses. Fortunately, because the hand adopted

in our experiment can passively guarantee stability (55), our hand-object system was explored simply by random actions sampled in a small range. After each action, all the hypothesized VLR models were moved forward by the motion function Eq. 4, and their corresponding sensor outputs were predicted by Eq. 5. Meanwhile, the true motion of the POM was observed by the in-hand camera and compared with the predictions to obtain the importance of each hypothesis, and then the distribution was resampled. Iteratively, false hypotheses were removed from the distribution, and the self-identification finally converged at iteration 12. In our implementation, convergence is defined by a threshold on the average of the 95% confidence intervals of the joint angles, which was set to 2° in all experiments. Note that, although the motion function Eq. 4 assumes fixed contacts under local motions, the contact changes cannot be prevented through the manipulation. Fortunately, because the system maintains a distribution of contacts represented by the set of hypotheses, the contact locations are updated and tracked over iterations.

To quantitatively evaluate the performance of self-identification, we have attached an AprilTag to the back of each fingertip. On the basis of the tag readings and inverse kinematics, we were able to calculate the ground truth joint angles of all proximal joints. However, because the angles of distal joints are determined by unknown contact locations, it was impossible to collect ground truth for their angles. Figure 4C shows

an example record of the self-identified angles of all proximal joints against their ground truth over 12 iterations. At iteration 0, we can see that the estimations were off from the truth and that the confidence intervals were large. As the self-identification iterates, the estimates converge toward the truth, and the confidence intervals shrink, implying that the system was able to precisely self-identify and track the angles. Because the Yale Model O hand is nonredundant, this indicates that the self-identification of other joint angles and contact locations was also accurate; otherwise, it would have been impossible to only accurately estimate the proximal joint angles. We repeated this experiment five times on each of the four objects shown in Fig. 4. Because VLRs do not rely on geometrical information of the object, the objects were just arbitrarily grasped by the hand without any prior shape information. Statistically, the estimation errors averaged $1.78^\circ \pm 0.92^\circ$, and the self-identification process took 15.1 ± 3.6 iterations to converge. Here and hereafter, all statistical results in this work are reported in the form of mean \pm SD.

In addition to purely self-identifying the kinematic VLR, we further challenged our system by introducing uncertainties into the physical properties Ω . For the adopted underactuated hand, one crucial physical parameter is its joint stiffness, which affects the share of workloads across the joints and determines the hand configuration; see Fig. 3. As such, we replaced the spring of the distal joint in the thumb with a novel spring and tried to self-identify both the VLR and the stiffness of that joint simultaneously. For this, in addition to sampling joint angles and the lengths of virtual links, the hypothesis set also included a distribution of the joint stiffness. Because this extra variable directly affects the output of the motion function Eq. 4, it can be self-identified in the same manner as the VLR. An example record of this experiment is shown in Fig. 5. On average, our experiments reported a precision of 0.0026 ± 0.0012 N·m/rad over five repeated trials.

In-hand manipulation

Once the VLR and physical properties Ω have been self-identified, the hand-object system will be able to use the motion function Eq. 4 to explicitly map the system actuation inputs to the motion of the POM. On another hand, as will be described later, this mapping can be inverted for the system to plan and control the manipulation; i.e., given a desired motion trajectory of the POM, the system determines the actuation inputs to achieve it. For testing how the VLR-based in-hand manipulation performs under different task requirements, we designed three experiments to evaluate its capability in terms of position control in \mathbb{R}^3 , orientation control in $SO(3)$, and full pose control in $SE(3)$.

For evaluating position control, the hand was initialized with a grasp on the apple object, followed by the iterations of self-identification to acquire the VLR model. Thereafter, on the basis of the VLR, the hand was positionally controlled to translate the POM through predefined waypoints, without any imposed orientation constraints. As shown in Fig. 6A, by traveling through the waypoints, the hand-object system completed a task of writing a seven-letter English word at the bottom of the object. The maximum scale of the English letters was 12 mm in both horizontal and vertical

directions, and the accuracy of our VLR-based position control averaged 0.42 ± 0.34 mm. In another task, Fig. 6B, the hand was tasked to play a marble maze. Initially, the maze was grasped by the hand without a marble in it. Once the VLR was self-identified, we placed the marble into the maze; controlled the orientations, which were precalculated, of the maze to move the marble through the desired path; and finally solved it. In this process, there was not any positional requirement imposed, and the orientation control reported a precision of $1.20^\circ \pm 1.38^\circ$.

Last, we designed a cup-stacking task to evaluate the in-hand manipulation with respect to the control in $SE(3)$. As shown in Fig. 7, the start configuration of this task has five cups on the table, and the goal is to pick up four of them in the order of decreasing size and finally stack them into the target blue cup. To make this task more challenging, we randomized each grasp by tilting the object placement with another small object underneath and commanded the hand to always grasp the cups by approaching them vertically with respect to the table. As such, the objects were all grasped with unknown poses, and the hand had to manipulate them to align in both position and orientation, to successfully stack them. However, because the hand is underactuated, it was not possible to independently control the six dimensions in $SE(3)$ at the same time to perfectly align the cups. In our implementation, we weighted the controlled dimensions and prioritized the rotational dimensions as necessary. Therefore, the performance in this task was not as good as the above experiments, achieving a positional accuracy of 2.1 ± 0.92 mm and a rotational accuracy of $5.16^\circ \pm 1.83^\circ$.

Hand recovery with novel designs

So far, we have always assumed that the hand geometry and kinematics are known. As previously mentioned and as depicted in Fig. 1, we also would like to recover damaged hands, particularly by replacing the damaged parts. However, in practice, it is possible that the available replacements are different from the original design. In this case, it would be beneficial if the hand-object system can still establish its VLR, with the help of self-identification. For this, we designed an experiment that assumed that there were one, two, and

three fingers broken and needed replacement. To make it challenging, we provided three finger replacements of very different geometries. As shown in Fig. 8A, the available replacements are different in both lengths and geometries. However, we assumed that they have the same kinematic structure as the original finger, i.e., each finger has two spring-loaded joints, actuated by only one motor through a tendon.

To this end, in addition to the parameters already included in the previous hypothesis set, the system needs to initialize and maintain an extra distribution of unknown virtual link lengths, which also have to be self-identified to establish the VLR. On another hand, to optimize computational resource usage, if some parameters are already known to the system, e.g., the lengths of the original fingers or the stiffness of the




Fig. 5. Self-identification of both VLR and joint stiffness. (A) A record of the self-identification process with the toy block object. The colors for visualization are the same as in Fig. 4. (B) The self-identification process recorded for the stiffness of the distal joint in the thumb. The shaded area shows the 95% confidence interval, and the dashed line marks the ground truth.




Fig. 6. In-hand manipulation applications based on VLR and self-identification. (A) The POM of the object was controlled to move through waypoints to write the English letters: S, C, I, E, N, C, E. (B) The hand played a marble maze by controlling the orientation of the POM. Complete videos of these experiments can be found in movie S1.



Fig. 7. A cup-stacking task requiring the hand to align both the position and the orientation of cups. The hand was installed on an arm; the task was to pick up all cups from the table and stack them into the target blue cup. At the bottom of each cup, there was an AprilTag for the in-hand camera to observe the POM.

original joints, those parameters are not being represented by the hypotheses distribution. An example is shown in Fig. 8C, in which the thumb of the hand was replaced with the novel finger NF1. At iteration 0 of self-identification, we can see that the hypotheses of NF1's distal joint (blue) were distributed in a cloud, rather than along an arc of known radius, indicating the length of the proximal link was also being self-identified. Last at iteration 17, in addition to the contact locations and joint configurations, the length of NF1's proximal link was also self-identified. This can be seen by the hypotheses of its distal joint (blue), which eventually converged to the same point at iteration 17.

We further challenged our system by replacing more fingers with novel shapes; see Fig. 8D. Again, without knowing the geometries of the replacements, the VLRs were successfully self-identified using the same framework. By definition, because the proximal virtual links have the same lengths as their corresponding physical links, we were able to evaluate the precision of its self-identification against the ground truth. One example is shown in Fig. 8E, which was recorded during one run on the Novel Hand 3 (NH3). We repeated this test five times on each novel hand, and the precision of link length estimation for each hand was 2.4 ± 2.1 mm (NH1), 2.9 ± 1.7 mm (NH2), and 3.4 ± 2.3 mm (NH3), respectively. We can see that the precision was decreasing as the number of novel fingers increased. However, comparing with the ground truth lengths of

these three links, which are 40, 50, and 50 mm, respectively, these errors are considered in an acceptable range for in-hand manipulation. In addition, using the ground truth tags on the back of the fingertips, the precision of joint configuration estimation was evaluated to be $2.58^\circ \pm 1.20^\circ$ (NH1), $3.04^\circ \pm 1.49^\circ$ (NH2), and $3.21^\circ \pm 1.26^\circ$ (NH3), which were all a little worse than when the hand model was fully known. Because the dimensionality of the self-identification framework is proportional to the number of unknown parameters, the decrease of the system performance is attributed to the limited number of particles. Because of the computational complexity involved, it is infeasible to always increase the number of particles to accommodate more unknown parameters. Therefore, the granularity of the estimation would decrease when the system's dimensionality increases, such as by introducing novel finger replacements or increasing the number of links in fingers.

Despite more unknowns being involved, because the lengths of novel virtual links, joint configurations, and contact locations were self-identified, we can see that the VLRs can be established for the novel hands under the same framework. This shows that our system provides the possibility for hand recovery when some parts, e.g., fingers, have to be replaced by novel designs. On the basis of the self-identified VLR, a novel hand will be able to manipulate the object with precise control, in the same way as if the hand model was fully known.

DISCUSSION

We have shown that the VLR representation can be self-identified and that it can be used to precisely control in-hand manipulation through challenging tasks. In particular, because self-identification seeks to estimate the VLR by actively figuring out the mapping from actuation to the system's motion, it is able to estimate the underlying system parameters with very limited sensing capability. For example, in our experiments, we show that the system can self-identify its joint configuration, contact locations, and even novel finger replacements using only an in-hand camera. Because the VLR is a virtual linkage-based representation, it does not constrain itself to particular kinematic structures or actuation types. For example, although our experiments were conducted only with a three-fingered hand, VLRs can be constructed with more fingers or more joints in each finger.

To successfully deploy the proposed approach in various scenarios, we next discuss some important aspects for implementation in practice, as well as the limitations of our system. Moreover, we expand our discussion by looking into other potential applications of our approach, in addition to hand-object systems.

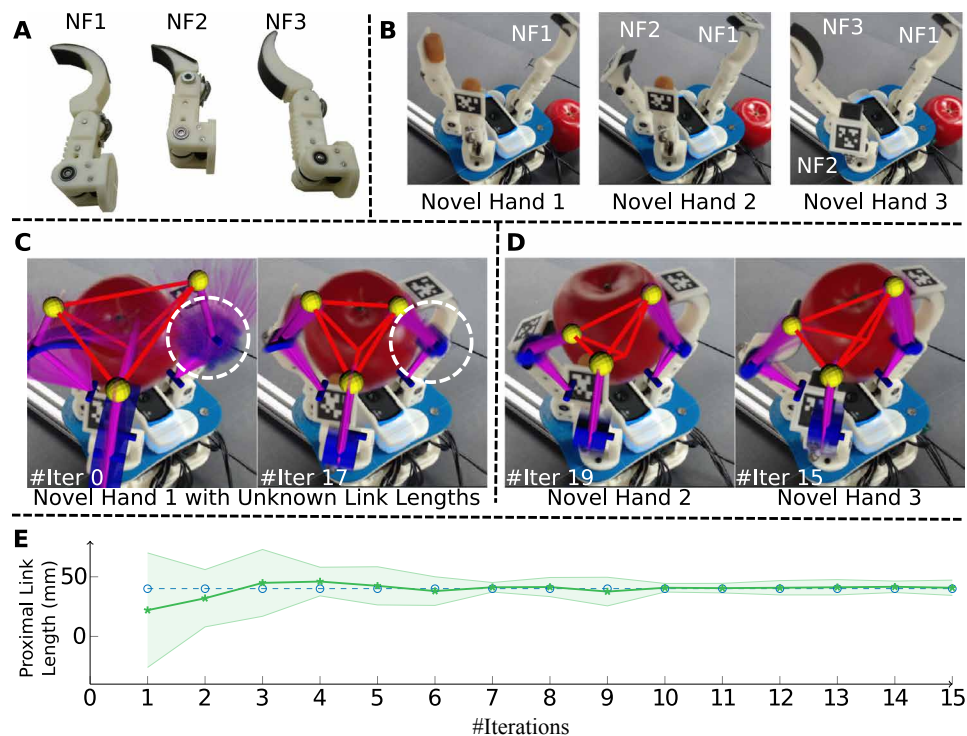


Fig. 8. Self-identification for hands with novel finger replacements. (A) The three novel replacements. (B) The three novel hands with one, two, and three fingers replaced. (C) For the novel finger, the lengths of both its proximal and distal links were sampled and self-identified. The colors for visualization are the same as in Fig. 4. (D) The self-identification results for NH2 and NH3. (E) The self-identification recorded for the length of the proximal link in NF1. The shaded area shows the 95% confidence interval, and the dashed line marks the ground truth.

Stability of exploratory interactions

As previously mentioned, self-identification is based on the more general concept of interactive perception. For this, the hand needs to actively interact with the object, so as to iteratively collect information and estimate the unknown parameters. In this process, it is crucial to guarantee that the interactions will not break the system formation; i.e., the object has to be stably grasped. In our experiments, although it was challenging to use an underactuated hand without any encoders or tactile sensors, we could take the advantage of underactuation to passively guarantee the stability and simply execute random actions for exploratory interactions.

However, in practice, this is not a common benefit for most other hands. For example, most fully actuated robotic hands will have to require active stabilization. It is worthwhile noting that, even for fully actuated hands, there will still be many unknown parameters, such as the exact contact positions, joint stiffness, and novel links. In such cases, before the self-identification process is finished, the system cannot be accurately controlled for explicitly providing stable hand-object configurations. To enable exploratory interactions, a potential solution is to mimic underactuation by introducing compliance into the grasp controller. For example, based on tactile or joint torque feedback, an impedance controller can be implemented to actively adjust contact forces (13).

Observation function and sensing modalities

Once stability is guaranteed, the system will make use of available sensing modalities to compare its observations against the predictions made by the VLR hypotheses, so as to iteratively filter out those

that are unlikely. In our experiments, being the minimal sensing requirement, we used only an in-hand camera to track the motion of the POM. This was challenging because there was limited information directly available about the hand-object configuration, even the hand joint configuration had to be self-identified. In practice, this can be substantially improved if more sensors become available, because we can potentially obtain more evidence to reason about the likelihoods of VLR hypotheses using Eq. 5. We observed from the experiments that, as the number of unknowns in the VLR increased, e.g., when joint stiffness or link lengths were unknown, the accuracy of self-identification accordingly decreased. This was a result of the decreased sampling resolution. Because the core of self-identification is particle filtering, more unknowns will increase the dimensionality of the hypothesis, which, in turn, requires exponentially more particles to be sampled to achieve the same sampling resolution. In practice, limited by the computational power, it is not always possible to increase the number of particles. However, again, if there were more sensors, the number of unknowns can be decreased, and the hypothesis set can be distributed such that it can focus more

on necessary dimensions with higher sampling resolution.

On another hand, having more sensing modalities can raise new challenges at the same time. Under a system configuration, one has to formulate all adopted sensing into the observation function Eq. 5, which potentially depends on other unknown information. For example, assuming we use a hand that has no fingertip tactile sensors but torque sensors in all its joints. For deriving its observation function, we have to obtain a mapping from the system configuration, $\xi = (\text{VLR}, \Omega)$, to its corresponding joint torques as required by Eq. 5. However, without knowing the contact forces at the fingertips, it is generally infeasible to analytically calculate such mappings. In such cases, there are two options for implementing the self-identification. First, although less sensing information can decrease the system's ability of estimating the likelihoods of hypotheses, we can safely leave some sensing modalities out and still achieve the same functionality, as long as the POM can be observed. Nevertheless, if some sensing modalities are preferred to be involved for planning or control purposes, we can include them by partially composing the observation function in a data-driven manner. This can be done offline, and the learned mapping does not even need to be bijective, because the iterative self-identification will eventually filter out false hypotheses, even if some were incorrectly estimated because of some ambiguities given by the learned model.

Point contact model and interaction control

Recall that the VLR representation is contact centered, and the contacts are modeled as point contacts to establish the virtual linkages.

Although this is not accurately modeling either the geometries nor the dynamics of contact, this is an effective model to describe the local interactions at contacts. First, imagine that the object was locally reconfigured by the hand because of some local contacts rolling. As long as the motion was sufficiently local, it can always be modeled by a small rotation around a point, which is the contact point being self-identified in the VLR. Second, if the local motion was a small contact translation, then it does not matter whether the contact was a point or some other geometry, again, as long as the motion was sufficiently small. For the same reasons, the point contact model is effective when the local motion is a mix of rolling and translation.

In practice, this indicates that, to ensure that the system motion function Eq. 4 is valid, we need to control the hand-object interaction so that the above assumption is held. In cases where the contact interactions are more complex, e.g., the local geometry of the object is complex, the self-identification process should limit its step size in the motion function Eq. 4 and increase the observation frequency for Eq. 5. Intuitively, this follows the same rule as in many other exploratory tasks that, if there are more uncertainties, one needs to be more careful by taking actions slowly while observing more often.

Limitations

Hand-object interactions can be achieved in various forms; e.g., assuming stable and fingertip-only contacts, we have already shown example tasks that require in-hand object translation or reorientation, which can be precisely controlled based on VLR self-identification. However, because the VLR is a contact-centered representation and its self-identification is an iterative process requiring continuous updates of the model, it is not able to deal with cases where contacts are not stable. For example, as a common practice for largely reorienting an object in-hand, palm contacts will be needed to provide stable support, whereas the fingers can manipulate the object without dropping it (23). However, in this case, it is impossible to define the VLR, because the contacts and virtual linkages can be made, broken, and remade constantly through the process. Similarly, VLRs are also not suitable in tasks with finger gaiting actions, which can suddenly change the VLR and break the consistency of the iterative process. In addition, because VLRs require explicit key points to be defined for establishing virtual linkages, they can only be used by hands with rigid links and are not feasible for modeling hands with soft links or continuum actuators (57).

Beyond hand-object systems

Although we showed only how VLRs can be used to self-identify hand-object systems and precisely control the in-hand manipulation, this linkage-based representation is not limited to hand-object systems. On various scales, interactions in the physical world can be modeled on the basis of some form of linkages. For example, when a robot grasps the handle of a door or pushes a shopping cart using two arms, virtual linkages can be established between its arms and the external objects. Same as the VLRs of hand-object systems, those virtual linkages can effectively describe the mapping from the robot's actuation inputs to the motion of the robot-door or robot-cart systems.

In scenarios where parameters are unknown, they can be iteratively self-identified on the basis of the motion and observation functions. For example, the width of the door and the location of the hinge can be self-identified by locally wiggling the door while

observing the door's motion as evidence. Nonetheless, although the linkage-based representation can model a variety of robot-object interactions, we need to keep in mind that it can be self-identified only if the contacts are stable, because of the same reasons as discussed for hand-object systems.

MATERIALS AND METHODS

In addition to the introduction of the VLR and its self-identification algorithm in their general forms, we next describe the technical details of the hand-object system implemented in this work. We will first derive the hand-object motion model and observation model and then explain how to inverse the motion model, so as to enable precise planning and control for in-hand manipulation.

Hand-object motion model

The Yale Model O hand used in this work is an underactuated hand. As shown in Fig. 3, this hand has three fingers, each of which has two spring-loaded joints driven by only one motor through a tendon. To enable self-identification, we need a motion function for this hand in the form of Eq. 4. Although the hand can passively provide stability, because its motors can only control the tendon length in each finger, the hand-object motion is actually indirectly determined by equilibrium of the grasp. Intuitively, given certain tendon lengths and their local changes (actuation inputs), the motion function should output a new hand-object configuration, so that the resulted contact forces provide a new equilibrium. Instead of directly deriving a motion function on the basis of forces and kinematics, which is difficult due to underactuation, we, in this work, find the equilibrium by modeling the system's energy.

Recall that there is a spring in each joint. While the springs provide torques at the joints, they at the same time store elastic potential energy. When the tendon lengths change, this spring-based system will reconfigure itself until it reaches the lowest potential energy possible, which is equivalent to the force equilibrium. Therefore, the motion function can be derived by finding hand-object configurations that minimize the elastic potential energy.

Concretely, for the i th finger, we denote by θ_{pi} and θ_{di} the joint angles, by k_{pi} and k_{di} the spring constants, and by r_{pi} and r_{di} the joint radii, for the proximal and distal joints, respectively. The energy stored in the i th finger can be calculated by

$$U_i = \frac{1}{2}(k_{pi}\theta_{pi}^2 + k_{di}\theta_{di}^2) \quad (6)$$

Independent of the absolute tendon length in the finger, if the motor's motion has changed the tendon length by Δl_i , the changes of joint angles are constrained by

$$\Delta l_i = r_{pi}\Delta\theta_{pi} + r_{di}\Delta\theta_{di} \quad (7)$$

In addition, assuming the fingertip contacts are fixed under local motions, there is a constraint imposed by the grasp that the lengths of all virtual links in E^C (see Eq. 1) should remain constant

$$\forall e^i \in E^C: \|e_{t-1}^i\| = \|e_t^i\| \quad (8)$$

where the subscript t indicates the virtual link e^i at different time t .

Under these constraints, an actuation input Δl will not change the virtual linkages of the VLR or the physical properties Ω , and the

system motion will only affect the virtual joint configuration. Therefore, we can instantiate the motion function Γ (Eq. 4) by solving the corresponding joint configuration using energy minimization

$$\theta^* = \arg \min_{\theta^j} \sum_i U_i \text{ s.t. Eqs. 7 and 8} \quad (9)$$

where θ^j is the virtual joint configuration of the VLR composed by θ_{pi} and θ_{di} of all fingers, as defined by Eq. 3. As such, given any actuation input Δl , the motion function will directly output the new virtual joint configuration θ^* to stabilize the hand-object system with equilibrium.

In addition to the motion function, we need to derive an observation function to relate the system configuration to its observation. Note that the only sensing modality available was from the in-hand camera, and the motion of the POM is completely determined by the configuration of the VLR. Therefore, the observation function Λ , as defined by Eq. 5, is simply a forward kinematics function of θ^j .

In-hand manipulation

Once the VLR is self-identified, given an actuation input, we can use the motion function Γ to calculate the system's reconfiguration and use the observation function Λ to predict the motion of the POM, denoted as $\Delta\phi$. Therefore, by chaining Γ and Λ , we can obtain another function, $\Upsilon : \Delta l \mapsto \Delta\phi$, that directly maps the system's actuation input to the motion of the POM, which we want to control. However, to derive a controller for in-hand manipulation, we have to inverse this mapping so that the desired actuation inputs Δl can be computed to achieve a given goal motion $\Delta\phi$.

Because the forward mapping involves a complex constrained optimization problem, it is not easy to analytically inverse it. Instead, we derived a Jacobian-based controller to numerically calculate desired actuation inputs and iteratively drive the POM toward the goal. Concretely, the Jacobian matrix, $J \in \mathbb{R}^{6 \times 3}$, is obtained by

$$J = \begin{bmatrix} \frac{\partial \Upsilon}{\partial l_1} & \frac{\partial \Upsilon}{\partial l_2} & \frac{\partial \Upsilon}{\partial l_3} \end{bmatrix} \quad (10)$$

where l_i denotes the tendon length in the i th finger and each column $\frac{\partial \Upsilon}{\partial l_i} \in \text{SE}(3)$ represents the local motion of the POM determined by the actuation input from the i th finger.

On the basis of this numerical Jacobian matrix, we have locally inverted the motion function Υ , and the control input is acquired by

$$\Upsilon^{-1}; J^T \cdot \Delta\phi \mapsto \Delta l \quad (11)$$

However, because the hand-object system is highly nonlinear, it is unlikely that this Jacobian-based actuation input Δl can accurately achieve the motion $\Delta\phi$ in one step. Instead, similar to implementing a proportional controller, we use this inverted mapping to iteratively generate actuation inputs while updating the desired motion $\Delta\phi$ by tracking the POM using the in-hand camera, until the motion has achieved a predefined precision.

Furthermore, dependent on the application requirements, the Jacobian matrix can be truncated to partially control the motion of the POM in its subdimensions. For example, in the hand-writing experiment, the POM was only positionally controlled, whereas in the marble maze experiment, the POM was only controlled for reorientations. However, because the hand is underactuated, in most cases, it is not able to control the POM to achieve high precision in

all dimensions of $\text{SE}(3)$, such as in the cup-stacking experiment. In our implementation, we weighted between the positional and rotational dimensions in $\Delta\phi$ to obtain a skewed $\Delta\phi'$

$$\Delta\phi' = [\Delta\phi_X, \Delta\phi_Y, \Delta\phi_Z, \alpha\Delta\phi_{\text{roll}}, \alpha\Delta\phi_{\text{pitch}}, \alpha\Delta\phi_{\text{yaw}}] \quad (12)$$

where $\alpha \in \mathbb{R}^+$ is the weight applied on the rotational dimensions. By substituting $\Delta\phi'$ for the $\Delta\phi$ in Eq. 11, we will be able to prioritize the precision for certain dimensions in our controller. Note that this skewed input only affects the controller when we have to trade-off between positional and rotational controls; it will have no effect when the system is able to achieve both, i.e., when $\Delta\phi = 0$. For the cup-stacking task, we empirically set $\alpha = 9.5$ to prioritize the orientation control, to ensure that the upper cup can be stacked successfully with an appropriate balance between position and orientation.

SUPPLEMENTARY MATERIALS

robotics.sciencemag.org/cgi/content/full/6/54/eabe1321/DC1

Movie S1. Method summary and experiment recordings.

Reference (58)

REFERENCES AND NOTES

1. J. Bütepage, S. Cruciani, M. Kocik, M. Welle, D. Kragic, From visual understanding to complex object manipulation. *Ann. Rev. Control Robot. Auton. Syst.* **2**, 161–179 (2019).
2. M. A. Lee, Y. Zhu, K. Srinivasan, P. Shah, S. Savarese, L. Fei-Fei, A. Garg, J. Bohg, Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks, in *IEEE International Conference on Robotics and Automation (ICRA)*, 2019, pp. 8943–8950.
3. W. Yuan, S. Dong, E. H. Adelson, Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors* **17**, 2762 (2017).
4. R. R. Ma, A. M. Dollar, On dexterity and dexterous manipulation, in *Proceedings of the 15th International Conference on Advanced Robotics (ICAR)*, 20 to 23 June 2011, pp. 1–7.
5. A. M. Okamura, N. Smaby, M. R. Cutkosky, An overview of dexterous manipulation, in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 24 to 28 April 2000, vol. 1, pp. 255–262.
6. C. Piazza, G. Grioli, M. Catalano, A. Bicchi, A century of robotic hands. *Annu. Rev. Control Robot. Auton. Syst.* **2**, 339–364 (2019).
7. A. Billard, D. Kragic, Trends and challenges in robot manipulation. *Science* **364**, eaat8414 (2019).
8. J. Bohg, A. Morales, T. Asfour, D. Kragic, Data-driven grasp synthesis—A survey. *IEEE Trans. Robot.* **30**, 289–309 (2014).
9. M. A. Roa, R. Suárez, Grasp quality measures: Review and performance. *Auton. Robots* **38**, 65–88 (2015).
10. S. Levine, P. Pastor, A. Krizhevsky, D. Quillen, Learning hand-eye coordination for robotic grasping with large-scale data collection, in *International Symposium on Experimental Robotics (Springer)*, 2016, pp. 173–184.
11. L. Shao, F. Ferreira, M. Jorda, V. Nambiar, J. Luo, E. Solowjow, J. A. Ojea, O. Khatib, J. Bohg, Unigrasp: Learning a unified model to grasp with multifingered robotic hands. *IEEE Robot. Autom. Lett.* **5**, 2286–2293 (2020).
12. J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, K. Goldberg, Learning ambidextrous robot grasping policies. *Sci. Robot.* **4**, eaau4984 (2019).
13. K. Hang, M. Li, J. A. Stork, Y. Bekiroglu, F. T. Pokorny, A. Billard, D. Kragic, Hierarchical fingertip space: A unified framework for grasp planning and in-hand grasp adaptation. *IEEE Trans. Robot.* **32**, 960–972 (2016).
14. Y. Bekiroglu, J. Laaksonen, J. A. Jorgensen, V. Kyrki, D. Kragic, Assessing grasp stability based on learning and haptic data. *IEEE Trans. Robot.* **27**, 616–629 (2011).
15. K. Hang, J. A. Stork, N. S. Pollard, D. Kragic, A framework for optimal grasp contact planning. *IEEE Robot. Autom. Lett.* **2**, 704–711 (2017).
16. Y. Lin, Y. Sun, Grasp planning to maximize task coverage. *Int. J. Robot. Res.* **34**, 1195–1210 (2015).
17. C. Borst, M. Fischer, G. Hirzinger, Calculating hand configurations for precision and pinch grasps, in *IEEE International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2002).
18. A. Kimmel, R. Shome, Z. Littlefield, K. Bekris, Fast, anytime motion planning for prehensile manipulation in clutter, in *IEEE International Conference on Humanoid Robots (HUMANOIDS)* (IEEE, 2018), pp. 1–9.
19. J. A. Haustein, K. Hang, D. Kragic, Integrating motion and hierarchical fingertip grasp planning, in *IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2017), pp. 3439–3446.

20. V. Kumar, E. Todorov, S. Levine, Optimal control with learned local models: Application to dexterous manipulation, in *IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2016), pp. 378–383.
21. M. Liarokapis, A. M. Dollar, Combining analytical modeling and learning to simplify dexterous manipulation with adaptive robot hands. *IEEE Trans. Autom. Sci. Eng.* **16**, 1361–1372 (2019).
22. A. Sintov, A. S. Morgan, A. Kimmel, A. M. Dollar, K. E. Bekris, A. Boularias, Learning a state transition model of an underactuated adaptive hand. *IEEE Robot. Autom. Lett.* **4**, 1287–1294 (2019).
23. O. M. Andrychowicz, B. Baker, M. Chociej, R. Józefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, J. Schneider, S. Sidor, J. Tobin, P. Welinder, L. Weng, W. Zaremba, Learning dexterous in-hand manipulation. *Int. J. Robot. Res.* **39**, 3–20 (2020).
24. O. Kroemer, S. Niekum, G. D. Konidaris, *A Review of Robot Learning for Manipulation: Challenges, Representations, and Algorithms* (2019); <http://arxiv.org/abs/1907.03146>.
25. J. C. Trinkle, R. P. Paul, Planning for dexterous manipulation with sliding contacts. *Int. J. Robot. Res.* **9**, 24–48 (1990).
26. L. Han, Y. S. Guan, Z. X. Li, Q. Shi, J. C. Trinkle, Dexterous manipulation with rolling contacts, in *IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 1997), vol. 2, pp. 992–997.
27. S. S. Srinivasa, M. A. Erdmann, M. T. Mason, Using projected dynamics to plan dynamic contact manipulation, in *IEEE International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2005), pp. 3618–3623.
28. K. Hertkorn, M. A. Roa, C. Borst, Planning in-hand object manipulation with multifingered hands considering task constraints, in *IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2013), pp. 617–624.
29. B. Sundaralingam, T. Hermans, Relaxed-rigidity constraints: In-grasp manipulation using purely kinematic trajectory optimization in *Proceedings of Robotics: Science and Systems* (Cambridge, MA, 2017).
30. A. S. Morgan, K. Hang, W. G. Bircher, A. M. Dollar, A data-driven framework for learning dexterous manipulation of unknown objects, in *IEEE International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2019), pp. 8273–8280.
31. A. Morgan, K. Hang, A. Dollar, Object-agnostic dexterous manipulation of partially constrained trajectories. *IEEE Robot. Autom. Lett.* **5**, 5494–5501 (2020).
32. K. Hang, W. G. Bircher, A. S. Morgan, A. M. Dollar, Hand-object configuration estimation using particle filters for dexterous in-hand manipulation. *Int. J. Robot. Res.* **39**, 1760–1774 (2020).
33. E. Roels, S. Terryn, J. Brancart, R. Verhelle, G. Van Assche, B. Vanderborght, Additive manufacturing for self-healing soft robots. *Soft Robot.* **7**, 711–723 (2020).
34. J. Noël, G. Kerschen, Nonlinear system identification in structural dynamics: 10 more years of progress. *Mech. Syst. Sig. Process.* **83**, 2–35 (2017).
35. K. J. Åström, B. Wittenmark, *Adaptive Control* (Courier Corporation, 2013).
36. Z. Jiang, W. Zhou, H. Li, Y. Mo, W. Ni, Q. Huang, A new kind of accurate calibration method for robotic kinematic parameters based on the extended Kalman and particle filter algorithm. *IEEE Trans. Indust. Electron.* **65**, 3337–3345 (2018).
37. E. N. Chatzi, A. W. Smyth, The unscented Kalman filter and particle filter methods for nonlinear structural system identification with non-collocated heterogeneous sensing. *J. Struct. Control* **16**, 99–123 (2009).
38. R. B. Gopaluni, T. B. Schön, A. G. Wills, Particle filter approach to nonlinear system identification under missing observations with a real application. *IFAC Proc. Vol.* **42**, 810–815 (2009).
39. J. Bohg, K. Hausman, B. Sankaran, O. Brock, D. Kragic, S. Schaal, G. S. Sukhatme, Interactive perception: Leveraging action in perception and perception in action. *IEEE Trans. Robot.* **33**, 1273–1291 (2017).
40. K. Y. Goldberg, R. Bajcsy, Active touch and robot perception. *Cognition Brain Theory* **7**, 199–214 (1984).
41. K. Gold, B. Scassellati, Using probabilistic reasoning over time to self-recognize. *Robot. Auton. Syst.* **57**, 384–392 (2009).
42. R. Kwiatkowski, H. Lipson, Task-agnostic self-modeling machines. *Sci. Robot.* **4**, eaau9354 (2019).
43. J. Bongard, V. Zykov, H. Lipson, Resilient machines through continuous self-modeling. *Science* **314**, 1118–1121 (2006).
44. M. C. Koval, N. S. Pollard, S. S. Srinivasa, Pose estimation for planar contact manipulation with manifold particle filters. *Int. J. Robot. Res.* **34**, 922–945 (2015).
45. C. Corcoran, R. Platt, A measurement model for tracking hand-object state during dexterous manipulation, in *IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2010), pp. 4302–4308.
46. D. Schiebener, A. Ude, T. Asfour, Physical interaction for segmentation of unknown textured and non-textured rigid objects, in *IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2014), pp. 4959–4966.
47. R. Platt, L. Kaelbling, T. Lozano-Perez, R. Tedrake, Efficient planning in non-gaussian belief spaces and its application to robot grasping. *Robot. Res.* **2017**, 253–269 (2017).
48. S. Levine, N. Wagener, P. Abbeel, Learning contact-rich manipulation skills with guided policy search, in *IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2015), pp. 156–163.
49. H. Culbertson, J. Unwin, K. J. Kuchenbecker, Modeling and rendering realistic textures from unconstrained tool-surface interactions. *IEEE Trans. Haptics* **7**, 381–393 (2014).
50. R. Martin Martin, O. Brock, Online interactive perception of articulated objects with multi-level recursive estimation based on task-specific priors, in *IEEE International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2014), pp. 2494–2501.
51. S. Thrun, W. Burgard, D. Fox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)* (MIT Press, 2005).
52. D. Crisan, A. Doucet, A survey of convergence results on particle filtering methods for practitioners. *IEEE Trans. Sig. Process.* **50**, 736–746 (2002).
53. F. Gustafsson, Particle filter theory and practice with positioning applications. *IEEE Aerospace Electron. Syst. Mag.* **25**, 53–82 (2010).
54. S. T. Tokdar, R. E. Kass, Importance sampling: A review. *Wiley Interdiscip. Rev. Comput. Stat.* **2**, 54–60 (2010).
55. L. U. Odhner, L. P. Jentoft, M. R. Claffee, N. Corson, Y. Tenzer, R. R. Ma, M. Buehler, R. Kohout, R. D. Howe, A. M. Dollar, A compliant, underactuated hand for robust manipulation. *Int. J. Robot. Res.* **33**, 736–752 (2014).
56. E. Olson, AprilTag: A robust and flexible visual fiducial system, in *IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2011), pp. 3400–3407.
57. R. Deimel, O. Brock, A novel type of compliant and underactuated robotic hand for dexterous grasping. *Int. J. Robot. Res.* **35**, 161–185 (2016).
58. B. Calli, A. Singh, J. Bruce, A. Walsman, K. Konolige, S. Srinivasa, P. Abbeel, A. M. Dollar, Yale-CMU-Berkeley dataset for robotic manipulation research. *Int. J. Robot. Res.* **36**, 261–268 (2017).

Funding: This work was supported by the National Science Foundation under grant IIS-1734190. **Author contributions:** K.H. proposed the self-identification framework, developed and implemented all algorithms, performed all experiments, and wrote the manuscript. W.G.B. designed and fabricated the experimental hand and its replacement parts and made some figures for the manuscript. A.S.M. helped with the design and fabrication of the experiment setup and edited the supplementary video. A.M.D. provided funding and supervised the project. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper or the Supplementary Materials. Additional information can be addressed to K.H.

Submitted 2 August 2020
Accepted 23 April 2021
Published 19 May 2021
10.1126/scirobotics.abe1321

Citation: K. Hang, W. G. Bircher, A. S. Morgan, A. M. Dollar, Manipulation for self-identification, and self-identification for better manipulation. *Sci. Robot.* **6**, eabe1321 (2021).

Manipulation for self-identification, and self-identification for better manipulation

Kaiyu HangWalter G. BircherAndrew S. MorganAaron M. Dollar

Sci. Robot., 6 (54), eabe1321. • DOI: 10.1126/scirobotics.abe1321

View the article online

<https://www.science.org/doi/10.1126/scirobotics.abe1321>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)

Science Robotics (ISSN 2470-9476) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. The title *Science Robotics* is a registered trademark of AAAS.

Copyright © 2021 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works