

MIT Open Access Articles

Robust proprioceptive grasping with a soft robot hand

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Homberg, Bianca S. et al. "Robust Proprioceptive Grasping with a Soft Robot Hand." *Autonomous Robots* (April 2018): 1-16 © 2018 The Author(s)

As Published: <https://doi.org/10.1007/s10514-018-9754-1>

Publisher: Springer-Verlag

Persistent URL: <http://hdl.handle.net/1721.1/114883>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution





Robust proprioceptive grasping with a soft robot hand

Bianca S. Homberg¹ · Robert K. Katzschmann¹  · Mehmet R. Dogar^{1,2} · Daniela Rus¹

Received: 5 January 2017 / Accepted: 7 April 2018
© The Author(s) 2018

Abstract

This work presents a soft hand capable of robustly grasping and identifying objects based on internal state measurements along with a combined system which autonomously performs grasps. A highly compliant soft hand allows for intrinsic robustness to grasping uncertainties; the addition of internal sensing allows the configuration of the hand and object to be detected. The finger module includes resistive force sensors on the fingertips for contact detection and resistive bend sensors for measuring the curvature profile of the finger. The curvature sensors can be used to estimate the contact geometry and thus to distinguish between a set of grasped objects. With one data point from each finger, the object grasped by the hand can be identified. A clustering algorithm to find the correspondence for each grasped object is presented for both enveloping grasps and pinch grasps. A closed loop system uses a camera to detect approximate object locations. Compliance in the soft hand handles that uncertainty in addition to geometric uncertainty in the shape of the object.

Keywords Soft robotics · Soft gripper · Proprioceptive soft robotic hand · Proprioceptive sensing · Online object identification · Learning new objects · Autonomously grasping

1 Introduction

Soft and under-actuated robotic hands have a number of advantages over traditional hard hands (Dollar and Howe 2006, 2010; Deimel and Brock 2013, 2014; Ilievski et al. 2011; Stokes et al. 2014; Shepherd et al. 2013; Brown et al.

2010). The additional compliance confers a greater intrinsic robustness to uncertainty, both for manipulating a broad range of objects and for conforming during interactions with the static environment.

Traditionally, grasping with rigid robotic hands requires detailed knowledge of the object geometry and precise location information for the object. Complex algorithms calculate the precise locations where the hand will grasp an object. With soft hands, we can grasp with a simpler, more intuitive approach handling more uncertainty.

While compliance enables intuitive grasping, the hand's specific configuration at a given time is hard to know due to the body's compliance. This is especially true when the hand is interacting with objects or the environment. This requires advanced internal sensing approaches, called *proprioception*, to infer the Soft Hand's configuration at any given moment. Knowing the configuration of the hand is crucial for decision making during the manipulation process. The hand configuration, for example, can be useful for determining whether a grasp is successful, whether a grasp is robust, and whether the object was grasped in the intended pose. The hand configuration can also be very useful in determining the shape of a grasped object, since the soft links tend to conform to the environmental constraints they interact with.

Bianca S. Homberg and Robert K. Katzschmann have contributed equally to this work.

✉ Robert K. Katzschmann
rkk@csail.mit.edu
<http://robert.katzschmann.eu/>

Bianca S. Homberg
bhomberg@csail.mit.edu
<http://bhomberg.github.io/>

Mehmet R. Dogar
m.r.dogar@leeds.ac.uk
<http://www.personal.leeds.ac.uk/~scsmrd/>

Daniela Rus
rus@csail.mit.edu
<http://danielarus.csail.mit.edu>

¹ Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

² School of Computing, University of Leeds, Leeds LS2 9JT, UK




Fig. 1 The soft robotic hand, mounted to the wrist of a Baxter robot, is picking up a sample object

In this paper we build on our previous work (Katzschmann et al. 2015; Marchese et al. 2015) and develop a soft robotic gripper called the DRL (Distributed Robotics Laboratory) Soft Hand (Fig. 1). The DRL soft hand is modular, allowing for the interchange of digits. Internal sensing from bend and force sensors provide feedback while grasping. In this paper we also evaluate two key features of this soft hand: its *proprioceptive grasping* capability and its *robustness to object pose uncertainty* during grasping.

In evaluating the proprioceptive grasping capability of this new hand, we build a model to relate the values coming from integrated bend sensors to the configuration of the soft hand. We then use this model for haptic identification of objects during grasping: The DRL soft hand is able to identify a set of representative objects of different shape, size and compliance by grasping them. We do this by building a relation between objects and the configurations the soft hand takes while grasping them. Then, given an unidentified object from our training set, the robot grasps it and uses proprioception to identify it. We also present an online identification algorithm where the hand learns new objects progressively as it encounters them by detecting measured sensor differences from grasps of known objects.

The intrinsic compliance of the DRL soft hand allows it to pick up objects that a rigid hand is not capable of picking without extensive planning and precise execution. Through experiments we show that the DRL hand is robust to a high degree of uncertainty. We perform an extensive number of experiments, in total attempting over 900 grasps of more than 100 randomly selected objects, to characterize this robustness quantitatively. We also show that the soft hand is more suc-

cessful at grasping compared to a standard electric parallel gripper, especially for delicate objects that are easily crushed and for grasping thin, flat objects that require contacting the environment. We evaluate the hand's capabilities by itself, but also for pick-and-drop tasks in an end-to-end system that integrates perception, planning, and grasping on a Baxter robot. We evaluate the DRL soft hand for a wide variety of grasping modes that include *enveloping* grasps, *pinch* grasps, *side* grasps, and *top* grasps.

In this paper we make the following contributions to soft robotic grasping:

- A modular, proprioceptive soft hand that includes integrated bend and force sensors;
- Evaluation of the proprioceptive grasping capabilities of the soft hand, which includes development of algorithms for the haptic identification of objects;
- Evaluation of the hand's robustness to object pose uncertainty during grasping, which includes an end-to-end solution to grasping that starts by visually recognizing the placement of the object, continues with planning an approach, and ends by successfully grasping the object by a Baxter robot;
- Extensive set of grasping experiments that evaluates the hand with a wide variety of objects under various grasping modes.

We present a complete account of the hardware design decisions and the grasping and sensing capabilities of this hand. Moreover, we test the practical performance of this hand in a state-of-the-art end-to-end perception-planning-grasping system. This complete account of the design decisions and the performance of the hand should inform future designers of soft hands.

In Sect. 2, we start with a discussion of related work. In Sect. 3, we present the DRL soft hand and describe the components and fabrication. In Sect. 4, we discuss the high-level system and algorithms used to control the hand and identify objects. In Sect. 5, we describe the experiments validating the construction of the hand. In Sect. 6, we conclude with a discussion of future work.

2 Related work

We build on recent developments in the fabrication of soft or underactuated hands. An overview of soft robotics is presented in Rus and Tolley (2015), Laschi et al. (2016) and Polygerinos et al. (2017). Dollar and Howe (2006, 2010) presented one of the earliest examples of underactuated and flexible grippers. Ilievski et al. (2011) created a pneumatic starfish-like gripper composed of silicone and PDMS membranes and demonstrated it grasping an egg. Deimel and

Brock (2013) developed a pneumatically actuated three-fingered hand made of reinforced silicone that was mounted to a hard robot and capable of robust grasping. More recently, they have developed an anthropomorphic soft pneumatic hand capable of dexterous grasps (Deimel and Brock 2014, 2016). Stokes et al. (2014) used a soft elastomer quadrupedal robot to grasp objects in a hard-soft hybrid robotic platform. A puncture resistant soft pneumatic gripper was developed by Shepherd et al. (2013). An alternative to positive pressure actuated soft grippers is the robotic gripper based on the jamming of granular material developed by Brown et al. (2010). The fast Pneu-net designs by Mosadegh et al. detailed in Mosadegh et al. (2014) and by Polygerinos et al. detailed in Polygerinos et al. (2013) is closely related to the single finger design used in this paper. The design and the lost-wax fabrication of the fingers of the DRL soft hand builds upon the soft gripper and arm structure proposed in Katzschmann et al. (2015), which demonstrates autonomous soft grasping of objects on a plane.

To the best of our knowledge, configuration estimates of soft robots so far have been acquired primarily through exteroceptive means, for example motion tracking systems (Marchese et al. 2014) or RGB cameras (Marchese et al. 2014). Various sensor types that can measure curvature and bending have been studied, but few have been integrated into a soft robot. Park et al. (2010, 2012) have shown that an artificial skin made of multi-layered embedded microchannels filled up with liquid metals can be used to detect multi-axis strain and pressure. Danisch et al. (1999) described a fiber optic curvature sensor, called Shape Tape, that could sense bend and twist. Weiß and Worn (2005) have reported on the working principle of resistive tactile sensor cells to sense applied loads. Biddiss and Chau (2006) described the use of electroactive polymeric sensors to sense bend angles and bend rates in prostheses. Kusuda et al. (2007) developed a bending sensor for flexible micro structures like Pneumatic Balloon Actuators. Their sensor used the fluid resistance change of the structure during bending. Other recent work in this area include that by Vogt et al. (2013) and Chossat et al. (2014). Chuah and Kim (2014) presented a new force sensor design approach that mapped the local sampling of pressure inside a composite polymeric footpad to forces in three axes.

Previous studies on haptic recognition of objects focus on hands with rigid links (Allen and Roberts 1989; Caselli et al. 1994; Johnsson and Balkenius 2007; Takamuku et al. 2008; Navarro et al. 2012). Paolini et al. (2014) presented a method which used proprioception to identify the pose of an object in a rigid hand after a grasp. Tactile and haptic sensors have also been used in manipulation to sense the external environment in Hsiao et al. (2007), Jain et al. (2013), Javdani et al. (2013) and Koval et al. (2013).

Liarokapis et al. (2015) presented a method to identify objects using force sensors in the context of a hybrid hard/soft

underactuated hand powered by cables. Farrow and Correll (2015) placed a liquid metal strain sensor and a pressure sensor in a soft pneumatic actuator and used the data to estimate the radius of a grasped object. Bilodeau et al. (2015) presented a free-hanging starfish-like gripper that is pneumatically actuated and has embedded strain sensors made of liquid metal. The sensor was used for determining if a grip is established. Building on the fiber-optical sensor in Danisch et al. (1999) and Zhao et al. (2016) presented a soft prosthetic hand that has integrated stretchable optical waveguides for active sensation experiments. Shih et al. (2017) showed a custom sensor skin for a soft gripper that can model and twist convex-shaped objects.

This paper improves the soft proprioceptive hand we presented in Homberg et al. (2015) with new capabilities and a new set of experiments. Particularly, our contributions in this paper over Homberg et al. (2015) and other previous work include:

- A new capability to sense force/contact through the integration of a force sensor in each finger of the soft hand;
- Force-controlled grasping experiments using the new force sensors;
- Addition of a fourth finger for improved grasping capability;
- All experiments previously presented were again conducted with the new hand using also the force sensors;
- An algorithm that allows the hand to identify new objects as it encounters them;
- New set of experiments to test this online object identification approach;
- Incorporation of the DRL Soft Hand into an end-to-end autonomous grasping pipeline and extensive experiments to measure its grasping performance under object pose uncertainty.

3 Device

The gripper used in this paper, the DRL soft hand, is an improved version of the gripper used in Homberg et al. (2015). Our objective was to develop modular, interchangeable fingers that can be slipped onto a 3D-printed interface. We designed each finger with several key goals in mind:

- Internal state sensing capability
- Force contact sensing
- Constant curvature bending when not loaded
- Partially constant curvature bending under loaded conditions
- Highly compliant and soft in order to be inherently safer

Most notably, a resistive bend sensor was embedded into each finger by affixing it beneath the finger's inextensible




Fig. 2 Create wax core mold using 3d printed model (a). For each finger create a wax core by pouring wax into the wax core mold. Create mold assembly for finger base (c) using wax core, insert part (f) and white lid. Cast the first layer of the finger using mold assembly (c) Melt the wax core out of the rubber piece and remove the insert piece. Re-insert

the rubber piece into the base mold. Glue the sensors onto the constraint layer (d). Place the constraint layer on top of the rubber piece (e). Pour a second layer of softer rubber into the mold. Remove the finger and plug the hole at the finger tip with solid tubing




Fig. 3 A cutaway view of the finger, showing the internal air channels, sealed sections, and inserted sensors and constraint layer

constraint layer, as can be seen in Fig. 3. Bending the resistive strip changes the resistance of the sensor. The resistive change can be correlated with the curvature of the finger. A force sensor was added on top of the constraint layer, also visible in Fig. 3. When the finger contacts an object, the resistive sensor's resistance changes, allowing us to detect the contact.

The combined hand, which we refer to as the DRL soft hand, is modular. Fingers easily attach and detach via 3D-printed interface parts. We can combine fingers to create different configurations of soft hands with different numbers of fingers. The primary configuration discussed in this paper is a four-fingered hand, an improved version of the previous three-fingered design (Homberg et al. 2015). The added finger directly opposes the thumb of the hand, allowing for a better enveloping of the object and an increased payload capability due to the firmer grasp at the center and the additional contact force. The four-fingered design allows for additional grasping options when compared to the previous design, such as a two finger pinch on small objects.




Fig. 4 Views of an individual finger and the entire composed hand

3.1 Fabrication

The fabrication of a single finger is based on a lost-wax casting process (Katzschmann et al. 2015; Marchese et al. 2015). As described in Homberg et al. (2015) and Fig. 2, the process has an added step where the bend and force sensors are added to the stiff constraint layer.

Figure 3 shows an image of the inside of the finished finger; the constraint layer and the sensors are visible.

The updated DRL finger is streamlined at 1.8 cm wide by 2.7 cm tall by 10 cm long, contains both bend and force sensors, and is not prone to popping or leaks. Various views of

a completed finger can be seen in the left column of Fig. 4. The new version benefits from shaping of the internal air channels and eternal finger shape to avoid all sharp corners which can be places of stress on the rubber. While the old version of the finger often broke intermittently, sometimes after only light use, the new version of the finger lasts several months and many hundreds of grasps before succumbing to rubber fatigue.

3.2 Actuation

Each finger is connected via a tube attached along the arm to a pneumatic piston. The actuation system is described in Katzschmann et al. (2015); Marchese et al. (2015).

3.3 Sensing

There are two sensors in each finger: the Flexi-force force sensor at the tip of the finger and the Bendshort-2.0 flex sensor from iCubeX. Both sensors are resistive sensors: as the sensor is pressed or bent, the resistance of the sensor changes.

(1) *Force sensor* The force sensor has a range of 4.5N but has an op-amp circuit to lower the range and increase the sensitivity. In order to get accurate results, we place a small metal piece behind the active area of the sensor. This prevents the bending of the finger from affecting the resistance of the sensor so that any sensed measurement comes just from the contact of the finger with an object.

(2) *Bend sensor* The sensors embedded in each finger are resistive bend sensors. The resistance of a sensor changes as it is bent.

3.4 Resistive sensor characterization

Due to the construction of the sensor, the relative change in resistance increases as the curvature of the sensor increases. Thus, the sensor has better accuracy and resolution as its diameter decreases. The diameter we refer to is the diameter of a circle tangent to the bend sensor at every point, for some constant curvature bend of the sensor. This relation between diameter of the finger and sensor value is shown in Fig. 5, where sensor values versus finger curvatures are plotted for the unloaded case.

Due to the inherent changes in variance for the sensor values, we are able to distinguish objects more accurately for objects with a smaller diameter.

4 Control

In this section we discuss the high level algorithms governing control for the finger and overall DRL soft hand. Implementation details are discussed in the next section.




Fig. 5 The diameter of the finger versus the sensor values

4.1 Architecture

For the hand-specific control, there are three sets of components: the physical fingers, the electronics, and the control software. The fingers are actuated via pneumatic tubing. The pneumatic tubing is pressurized by a piston, and the piston is driven by a linear actuator. Motor controllers, one per linear actuator, set the position of the linear actuators, setting the volume of the air in each finger. Additionally, each finger has a bend and a force sensor. Each of the sensors are connected to filtering and buffering electronics and then measured using an Arduino board.

On the software side for the hand, there is a middle-level controller enabling us to command the hand using primitive actions such as “close the hand” or “open the hand”. This middle-level controller communicates with the low-level motor controllers via serial communication. It also receives sensor values from the Arduino board on the hand through `rosserial`.¹

On the robot side, the two key pieces of hardware are the hand cameras and the robot arm, to which the hand is attached via 3D printed interface parts. For the robot software, we implemented the grasping and object recognition pipeline using a set of ROS nodes (Quigley et al. 2009). One main ROS node coordinates the overall behavior. One ROS node reads the camera input streams and performs object detection using basic image processing in OpenCV (Bradski 2000). One strength of the DRL soft hand is its ability to grasp unknown objects with uncertain pose estimation. This vision system serves to detect approximate poses of objects even if they are completely unknown to the robot. A suite of ROS nodes run for the MoveIt planner Sucan and Chitta 2018. One object in the codebase interfaces with the MoveIt planner to coordinate calls to plan motions to different locations. For

¹ <http://wiki.ros.org/rosserial>.

side grasps, the motion planning node finds a grasp plan given a potential object location using an intermediate way point. The planner considers 16 potential directions by which to approach the object. Along the direction, it first considers an offset pre-grasp location which is offset far enough to be simple to plan to without getting too close to the object. For top grasps, the motion planner is called to find a plan to a pose where the hand is vertically above the object. For top grasps of small objects, first the fingers are half closed to allow the hand to approach closer to the table without the fingers hitting the table and being unable to bend to grasp due to excessive friction. Another object handles the control of the soft hand, opening, closing, and grasping. A separate node sends specific commands via serial to the motor controllers.

4.2 Finger control

The value measured from the force sensor is an approximate force. Due to noise after the hardware low pass filter, we buffer the output in software and consider an average of the past five data samples. If the average of the data samples crosses a certain threshold, we consider this to be a contact between the fingertip and an external object.

In grasping an object, we keep increasing the volume of air in each finger until we detect a point of contact. Since the grasp criterion for each finger is independent, it does not matter if an object is oddly shaped; the fingers will each close the correct amount. If no contact is detected, the fingers simply keep closing until their maximum closure (Algorithm 1).

Algorithm 1: Active Grasping Algorithm

```

while not each finger either completely closed or in contact do
  for each finger which is not either completely closed or in
  contact do
    | Move delta towards closing finger.
  end
end
end

```

4.3 Grasping

We incorporated the DRL soft hand into a complete, end-to-end grasping system. This grasping system demonstrates the versatility of the soft hand by showing its robustness to uncertainty in the location of the object and the minimal need for grasp planning.

The grasping system pipeline consists of three phases: location, approach, and grasp (Algorithm 2). A successful execution of a grasp means that all steps of Algorithm 2 are executed without failure. To be more specific, this entails that the arm motion is successfully planned and executed, and

the object is then successfully grasped, correctly identified, dropped into the bin.

Algorithm 2: Full System Algorithm

```

while true do
  Move arms to object spotting location (cameras facing down).
  Plan grasp to object.
  Execute grasp to object.
  if grasp execution is unsuccessful then
    Repeat planning and execution steps until successful
    execution is achieved or the maximum number of
    attempts is exceeded.
  end
  Grasp object (using force sensor feedback to detect grasp).
  Move to fixed side location.
  Record sensor values and perform object identification.
  Drop object into bin.
end
end

```

4.4 Object identification

Once trained, the DRL soft hand is able to identify the grasped objects based on a single grasp. We first characterize the relation between hand configurations and sensor readings. Then, we present a data-driven approach to identify an object based on sensor readings.

(1) *Modeling the sensor noise* The DRL hand has different configurations as it interacts with the environment and grasps objects. We define a configuration of the DRL hand as a vector $\mathbf{q} = [q_1, q_2, q_3, q_4]$, where each $q_i \in \mathbb{Q}$ represents the way finger i is bent. \mathbb{Q} is the configuration space of a finger: that is, the space of all different shapes our soft finger can achieve. For a given configuration of the hand, we get bend sensor readings $\mathbf{s} = [s_1, s_2, s_3, s_4]$, where each s_i represents the bend sensor reading for finger i and a force value $\mathbf{f} = [f_1, f_2, f_3, f_4]$, where each f_i represents the force sensor reading for finger i .

The sensor readings are noisy. We represent the sensor reading given a hand configuration as a probability distribution, $p(\mathbf{s}, \mathbf{f} | \mathbf{q})$. Given the configuration of a finger, the sensor values of that finger is independent of the configuration of the other fingers. Therefore, the sensor model of the whole hand can be expressed in terms of the sensor model for each finger:

$$p(\mathbf{s}, \mathbf{f} | \mathbf{q}) = \prod_{i=1}^4 p(s_i, f_i | q_i) \quad (1)$$

We model $p(s_i | q_i)$, the bend sensor noise for a finger, in a data-driven way by placing the finger at different configurations and collecting the sensor value data. In Sect. 3.4 we present experiments for such a characterization, where we

use constant curvature configurations of the unloaded finger. The force value will depend not just on the configuration of the finger but also on the interface of the finger with the environment, as the sensor values differ based on where exactly the sensor is pressed. In order to model the desired probability $p(s_i, f_i | q_i)$, we also need to take into account the interaction with the environment. In the absence of any external interaction, the finger is constructed such that f_i will always be equal to 0.

Note that when the finger is loaded in a grasp, the resulting finger configurations and the corresponding sensor readings have significant variation due to the highly compliant nature of the fingers. Therefore, to identify objects during grasping, we use data collected under the grasp load.

(2) *Object identification through grasping* We use the sensors on the hand to predict the hand configuration, which we then use to identify the grasped object.

The grasping configuration for an object can be different for different types of grasps. In this work we focus on two types of grasps: *enveloping grasps* and *pinch grasps*. For a given object, o , we represent the configuration at the end of an enveloping grasp as \mathbf{q}_o^{envel} ; and we represent the configuration at the end of a pinch grasp as \mathbf{q}_o^{pinch} .

For given sensor readings \mathbf{s} and \mathbf{f} and a grasp type $g \in \{envel, pinch\}$, we define the object identification problem as finding the object with the maximum likelihood:

$$o^* \leftarrow \operatorname{argmax}_{o \in \mathbb{O}} p(\mathbf{q}_o^g | \mathbf{s}, \mathbf{f}) \quad (2)$$

where \mathbb{O} is the set of known objects and o^* is the predicted object. Applying Bayes' rule, we get:

$$o^* \leftarrow \operatorname{argmax}_{o \in \mathbb{O}} \frac{p(\mathbf{s}, \mathbf{f} | \mathbf{q}_o^g) p(\mathbf{q}_o^g)}{p(\mathbf{s}, \mathbf{f})} \quad (3)$$

Since the denominator is a constant over all objects o , we see:

$$o^* \leftarrow \operatorname{argmax}_{o \in \mathbb{O}} p(\mathbf{s}, \mathbf{f} | \mathbf{q}_o^g) p(\mathbf{q}_o^g) \quad (4)$$

Assuming a uniform prior over finger configurations, the above formulation becomes:

$$o^* \leftarrow \operatorname{argmax}_{o \in \mathbb{O}} p(\mathbf{s}, \mathbf{f} | \mathbf{q}_o^g) \quad (5)$$

In our experiments we use a trained dataset to build an empirical model of $p(\mathbf{s}, \mathbf{f} | \mathbf{q}_o^g)$ for different objects and grasp types. Then, we identify the object for a new grasp (Eq. 5) using a k-nearest neighbor algorithm.

3) *Trained Object Identification*: Algorithm 3 uses an initial trained dataset. We train using a dataset of sensor values for

repeated grasps of known objects. We use the same dataset as for clustering, but with the originally known identities of each of the objects. We use this training set to identify objects as they are grasped in a separate testing phase. After each new grasp, the five nearest neighbors of the new point in the original training data are determined. We calculate the distance via the Euclidean metric on the 4-dimensional point comprised of the four sensor values, one per finger. The object is identified based on the most common identity of the five nearest neighbors, using the *KNeighborsClassifier* from *scikit-learn* (Pedregosa et al. 2011). The identification algorithm runs in less than 0.01 seconds. This algorithm is flexible: it was used on three-fingered and four-fingered versions of the hand without modification. Given a number of fingers, D , and a number of objects to identify, N , the running time of the algorithm grows as $O(D \log(N))$. The number of classes we are able to successfully distinguish is limited by the sensor resolution: with noisy sensors, object clusters must be relatively far apart in order to be distinguishable. Higher fidelity sensors, additional sensors in each finger, or the use of additional sensing sources (e.g., vision) would enable this technique to work with more classes of objects.

Algorithm 3: Trained Object Identification Algorithm

```

Import previously recorded grasp data, 10 data points per item
for all objects to be grasped do
  Grasp item.
  Record sensor values.
  Calculate Euclidean distances to all recorded points
  Find the 5 nearest neighbors.
  Output the identity of the object-based voting from the 5
  nearest neighbors.
end

```

4) *Online Object Identification*: In Algorithm 4, the robot identifies objects online as it grasps new and old objects. Initially, the hand is trained to identify the empty grasp as a baseline via ten grasps. This allows the robot to have a known starting point of what the sensor values are when it has not grasped an object. As the hand grasps objects, the algorithm decides as it grasps each object whether the object is a known object from the objects it has already learned or a new object it has not yet seen. If the object is identified as a known object, it adds the data from that grasp with the label of the identified object. If the object is identified as a new object, it creates a new label and adds the data from that grasp with the new label.

Essentially, when grasping an object, the algorithm considers the distance between the values of the sensors for that object and all of the other objects currently in the dataset. Based on the average distance to all data points from a label and the number of data points with that label, a score is calculated for each label. If the label with the highest score has a

score higher than a fixed *cutoff score*, then the grasped object is labelled with that label. This cutoff score was empirically determined based on the sensor variability across identical grasps. The score for a label is equal to $1/(avg_{dist} \cdot n)$ where avg_{dist} is the average distance from the current sensor values to all data points (in the 4D space of the sensor values for the four fingers) with the given label and n is the number of data points with that label (Algorithm 4). Given a number of fingers, D , and a number of past grasps M , the running time of the algorithm grows as $O(D \cdot M)$. Further work should consider an adaptive cutoff score and post-processing to re-balance learned classes. Such methods could allow the same algorithm to adapt to sensors with different levels of noise.

Algorithm 4: Online Object Identification Algorithm

```

Train with three grasps of the empty grasps, record data
for all objects to be grasped do
  Grasp item.
  Record sensor values.
  Calculate Euclidean distances to all recorded points.
  Calculate the average distance to all currently identified
  objects.
  Calculate a score for each object equal to  $1/(avg_{dist} \cdot n)$ .
  if the highest score > the cutoff score then
    Add the data as a data point from an existing object with
    that label.
  else
    Add the data as a new object with a new label.
  end
end
end

```

5 Experiments and results


We performed experiments to evaluate the DRL soft hand's capability in three different aspects, presented in separate sections below:

- The basic grasping capability of the soft hand (Sect. 5.1),
- The proprioceptive capability of the hand applied to autonomous object identification and force-controlled grasping (Sect. 5.2),
- The grasping performance of the soft hand under object pose uncertainty within an end-to-end system (Sect. 5.3).

In general, our goal with these experiments has been to produce an exhaustive characterization of the capabilities of this soft hand. Therefore we tested the hand with a high number of different objects (more than 100 objects), used many different grasping modes (enveloping grasps, pinch grasps, side grasps, and top grasps), used different proprioceptive modalities (finger curvature and contact forces), and different end-to-end setups. During the experiments, in total, we



(a)



(b)

Fig. 6 Objects grasped. **a** ~100 objects grasped by the DRL soft hand and **b** the six objects the DRL soft hand failed to pick up

performed more than **900 grasping attempts** with the DRL soft hand, of which more than 600 were successful.

5.1 Basic grasping capability

To first evaluate the general grasping capability of the soft hand, we tested it with a wide variety of objects. The full set of 100 objects can be seen in Fig. 6. Some grasps of these objects can be seen in Fig. 7.

During the experiments, each object was placed at a known pose and a grasp was attempted by the Baxter robot using the DRL soft hand. For this set of experiments, our goal was to focus on the grasping capability of the DRL soft hand, and therefore we took the liberty to implement different grasping strategies for different objects. (Section 5.3 presents our set of experiments where we evaluated the soft hand within an end-to-end system using *autonomous* perception and planning.) Some objects were grasped via enveloping grasps. Others were picked up via a top grasp with two or three fingers in a pinch grasp. The flat objects, e.g. the CD and the piece of paper, were grasped off of the table as was shown in Fig. 8b. All objects were positioned in the orientations as they are in Fig. 6. The DRL soft hand was able to successfully grasp and pick up 94 of 100 objects.

We made three key observations during these experiments.




Fig. 7 Various objects grasped by the DRL soft hand. **a** Aquaphor, **b** lemonade bottle, **c** squash, **d** mug, **e** ring and **f** marker




Fig. 8 Rigid gripper squashing a cup and soft gripper picking up a thin object. **a** Cup squashed by rigid gripper and **b** gripper performs a compliant grasp to pick up a thin object off a table

First, **the DRL soft hand was capable of grasping a wide variety of objects with different sizes and shapes.** The objects in Fig. 6 were chosen to explore the extents of the grasping capability of the soft hand and it, thanks to its soft compliance, easily adapted to different shapes and sizes.

Second, **the DRL soft hand was capable of grasping objects that require simultaneous compliant interaction with the environment.** Specifically, we tested grasping a CD and a piece of paper off of a table, again using both the DRL soft hand and the default rigid gripper. The default gripper was unable to pick up a CD or piece of paper. Our soft hand was reliably able to pick up the CD and the piece of paper. Figure 8b shows how the soft gripper smoothly interacts with the environment to pick up the CD.

Third, **the DRL soft hand was qualitatively better at grasping a compliant object when compared with the rigid gripper.** Specifically, we tested grasping a soft paper

Table 1 Identification percentages for each of the tested objects. Dashes represent that an object was not used in a particular test due to it not being the right shape for grasping in that orientation

Object	Enveloping grasp (%)	Pinch grasp (%)
Zip tie container	100	–
Cup	60	–
Egg	100	–
Tennis ball	100	80
Plastic lemonade bottle	100	100
Glass lemonade bottle	100	–
Aquaphor	100	–
Plush hedgehog	80	100
Plastic bin	100	100
Wood block	100	100
Safety goggles	–	80
Eggbeater	–	40
Empty grasp	100	100

cup using the DRL soft hand and the default rigid parallel grippers of the Baxter robot. When the default gripper picked up the cup (Fig. 8a), it crushed it; the soft gripper was able to pick it up repeatedly without crushing.

The DRL soft hand was **not** able to pick up six of the 100 objects. These objects can be seen in Fig. 6b. The hand was not able to pick them up primarily because they were too heavy or too flat for the finger to gain any traction on them. The gripper had trouble picking up a spoon, a pair of scissors, and a propeller because they were not high enough – the fingers were unable to make a good connection. The gripper was unable to pick up an amorphous bag of small objects because of the combination of its odd shape and heaviness, the fingers did not get a solid grasp below the center of mass and the bag deformed to slip out of the fingers. The fish tail could be grasped, but slipped due to its weight. The screw was simply too small to be reliably grasped.

5.2 Proprioceptive grasping

We performed our second set of experiments to evaluate the proprioceptive capability of the DRL soft hand. First, we performed experiments to identify objects based on finger curvature after grasping. Second, we performed experiments to perform force-controlled grasping of objects.

1) Object identification using finger curvature: We first tested the trained object identification algorithm described in Algorithm 3. To characterize the hand's capabilities in different grasping modes, we performed experiments both for enveloping grasps and pinch grasps. For enveloping grasps we used ten objects and the empty grasp (Table 1, Fig. 10), and for pinch grasps we used seven objects and the empty grasp

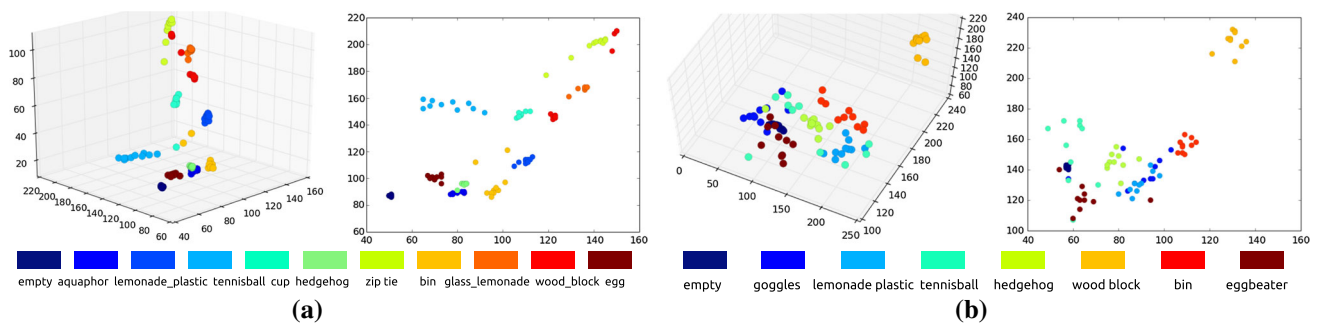


Fig. 9 The 4-finger data for **a** enveloping grasps and **b** pinch grasps. In both **a** and **b**, the first 3D plot uses the curvature values from the first three fingers and the second 2D plot uses the third and fourth fingers. There were ten grasps of each object. These grasps were labeled with true object ids. The true ids of these objects are shown via color, as seen

in the keys of each subfigure. Using this labeled set as training data, the soft hand predicted the identity of grasped objects in further unlabeled grasping experiments, as shown in Table 1, with 100% accuracy for most of the objects (Color figure online)

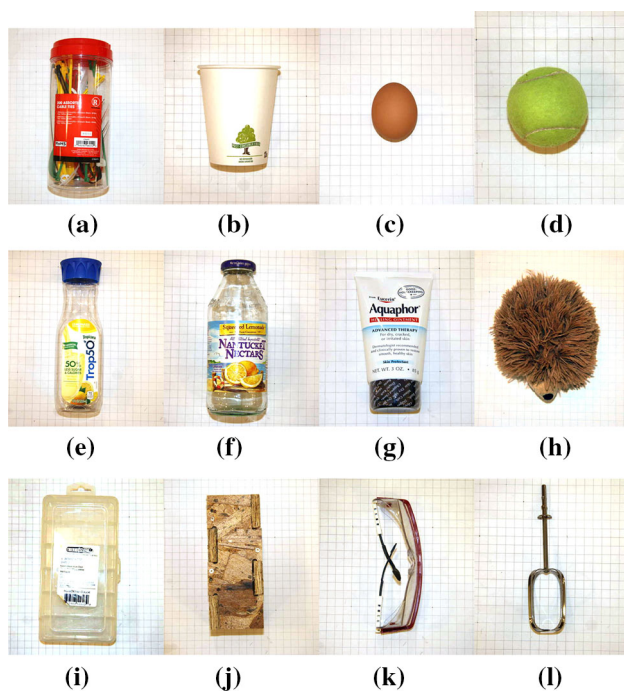


Fig. 10 The test objects used in the grasping experiments. **a** Zip tie, **b** cup, **c** egg, **d** tennis ball, **e** lemonade plastic, **f** lemonade glass, **g** aquaphor, **h** hedgehog, **i** Bin, **j** wood block, **k** goggles, **l** eggbeater

(Table 1, Fig. 10). For each grasp type, we first performed ten grasps of an object and labelled it with its object id. Then we performed additional unlabeled grasps (55 for enveloping grasps and 40 for pinch grasps) and used Algorithm 3 to identify the objects based on proprioception. In Fig. 9, we present the distribution of the 4-dimensional proprioceptive data and the labels Algorithm 3 assigned to each grasped object. 94.5% of tests (52/55 trials) identified the objects correctly for enveloping grasps; the breakdown per object is shown in Table 1. For pinch grasps, 87.5% of tests (35/40 trials) identified the objects correctly; again, the breakdown per object can be seen in Table 1. This includes correctly

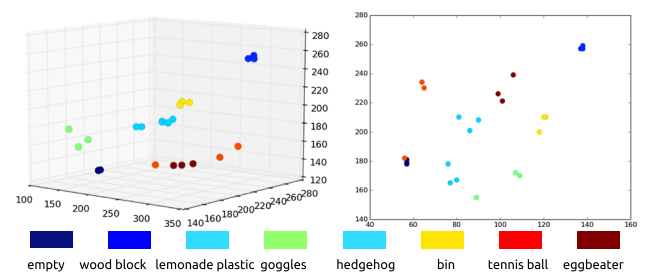


Fig. 11 The predicted IDs of objects at the end of the online object identification. The first 3D plot uses the curvature values from the first three fingers during the grasp of an object. The second 2D plot uses the third and fourth fingers. The predicted ids of objects are shown in color. Except the lemonade and the hedgehog, which were clustered as the same object, all other objects were clustered distinctly and correctly by the system. The key at the bottom shows the corresponding object name for each color (Color figure online)

identifying the empty grasp when the robot did not actually pick up an object.

We also tested the online object identification algorithm outlined in Algorithm 4 by grasping the same objects used in the pinch grasp tests. We started by training the empty grasp, then picked up the wood block, plastic lemonade bottle, goggles, hedgehog, bin, tennis ball, and eggbeater in that order. We trained the empty grasp with three iterations and then picked up each of the other objects three times. Except for the hedgehog, for which all grasps were identified as the previously-grasped lemonade bottle, the algorithm correctly identified each object as a new object the first time and as itself for subsequent grasps, for an identification success rate of 85.7%, on par with the trained results for pinch grasps. Notably, once the system identified an object correctly as a distinct object, it successfully matched future grasps of the object correctly in all cases. 6/7 objects were identified correctly as distinct objects. The plot of the identified data points can be seen in Fig. 11.

Table 2 Number of fingers which sensed contact with the object for each of the tested objects

Object	Average number of fingers that sensed force (out of 4)
Glass lemonade bottle	2
Plastic lemonade bottle	.33
Ziptie container	2.33
Wood block	2.66
Tennis ball	0
Squash	0
Hedgehog	.66

**Fig. 12** The DRL soft hand with stiff backdrop inserted for object identification grasps

(2) *Proprioceptive force-controlled grasping* We also performed experiments to test the accuracy of the proprioceptive force-controlled grasping algorithm and the force sensors. First, we calibrated the force sensors and identified a threshold that was high enough so that we did not get any false positive contact signals. Afterwards, three grasps each of seven objects were tested. Table 2 shows how many fingers stopped grasping due to force contact before the grasp was completed.

We made two key observations during the proprioceptive grasping experiments using the finger curvature sensors and the force sensors.

First, during these experiments, **we had to place a stiff pad behind the rubber bands** to provide a backdrop (Fig. 12) and allow the primary conforming of the hand to come from the fingers rather than the rubber bands. This observation will play an important role in designing future versions of the hand.

Second, **while the finger curvature sensors provided valuable data** which resulted in an impressive object identification performance, **the data from force sensors resulted in mixed performance** as shown in Table 2. Many objects were simply too small so maximum closure was required

before the force sensors were activated. For others, the fingers grasped them at an angle that did not activate the sensor.

5.3 Grasping under object pose uncertainty within an end-to-end system

In our third set of experiments we evaluated the grasping performance of the DRL soft hand under object pose uncertainty within an end-to-end system.

Soft hands, due to their intrinsic compliance, have the advantage of being robust to uncertainties in object pose during grasping. We use the *capture area* as a measure of the degree to which the DRL Soft Hand is robust to object position uncertainty. We define the capture area as the size of the region within which an object can move and can still be grasped robustly by the hand.

We performed this evaluation in three steps of increasing system complexity:

- (1) We evaluated the DRL soft hand's grasping performance under object pose uncertainty. This provided us with a baseline for the following two cases.
- (2) We evaluated the same grasping performance under object pose uncertainty, but in an end-to-end system that consists of the DRL soft hand, a perception component to detect objects, and a planning component to move the hand to a detected object. To measure the extent to which this system can tolerate uncertainty, we injected artificial uncertainty into the system.
- (3) We evaluated the general grasping performance of the end-to-end system with a wide variety of initial object poses, but without artificial uncertainty.

For all three test types, we tested both side grasps and top grasps with an appropriate object for each. We present the details of these three sets of experiments below.

1) *Grasping under Object Pose Uncertainty*: The first tests considered the range over which top and side grasps would successfully grasp the object. The hand was centered over a 10x10cm grid and performed repeated grasps of an object placed at the center. Over multiple trials, the object was moved to different positions in the grid and we recorded whether or not the object was successfully grasped. For each grasp, the hand moves to the same location; the purpose of this test is to see how much uncertainty in object location the gripper can handle while still successfully grasping the object.

For side grasps, the object used was the lemonade bottle. See Fig. 13a to see the configuration of the test setup and the approach angle of the hand. Figure 14a shows the grid used in the test with dots at points which were tested. For each location, two trials were performed. For the 41 locations




Fig. 13 Objects in grasp configuration. **a** Lemonade bottle in test configuration and **b** foam block in test configuration

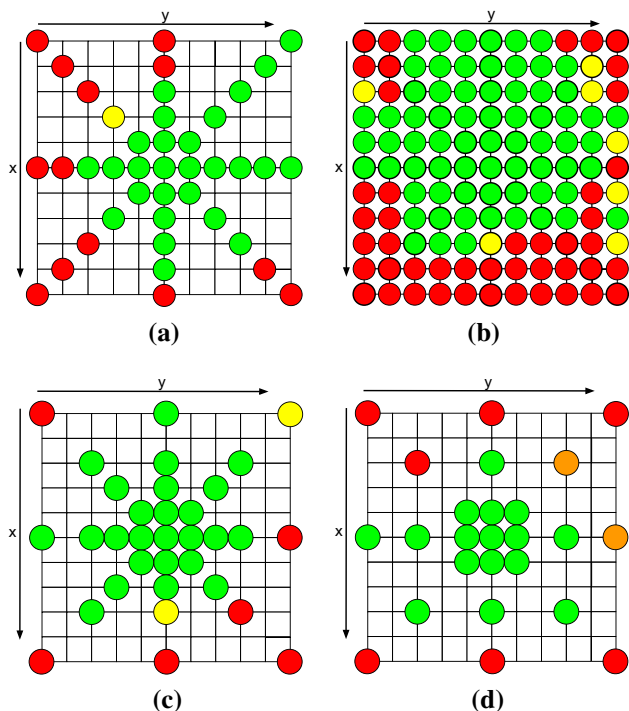


Fig. 14 Top row shows the results with a hard-coded grasp location (Sec. V-C.1). Bottom row shows the results using a full perception-planning-grasping pipeline (Sec. V-C.2). Grasp success rates are shown over a 10×10 grid. ●: two successful grasps; ●: one successful grasp and one failed grasp; ●: one successful grasp and two failed grasps; ●: two failed grasps. **a** Tests of the pre-programmed side grasps, **b** Tests of the pre-programmed top grasps, **c** Tests of the end-to-end side grasps and **d** Tests of the end-to-end top grasps (Color figure online)

shown in the grid, 82 grasps were attempted, 55 of which were successful.

For top grasps, the object used was a foam block covered with black electrical tape. The test setup can be seen in Fig. 13b. The object was placed with its longer dimension along the hand's opening as seen in the figure. The hand descended on a centered position and closed three fingers in one of the types of top grasp. Figure 14b shows the grid used in the test with dots at points which were tested. For each location, two trials were performed. For the 121 locations shown on the grid, 242 grasps were attempted of which 145

were successful. Again, we observed that the soft hand was able to grasp the object reliably even when the object was significantly away from the center, the exact value changing between 3cm and 5 cm depending on the axis, since the object is asymmetric. Other asymmetries in the plot are due to different dynamic interactions between the fingers and the block during grasping.

In Table 3, we present the capture areas of the DRL Soft Hand as a general measure of how robust it is to object pose uncertainty during grasping. In the total area of 100 cm^2 within which the object position was varied, the capture area shows the size of the region for which the grasps were robustly successful. The values in the table are found by measuring the areas spanned by the green dots in Fig. 14. If a data point is missing on a particular grid point, we assumed that the capture region is convex (Dogar and Srinivasa 2010). We are only using this assumption for calculating the capture region as a means of representing the raw data in Fig. 14 with a summarizing metric. As Table 3 shows, the DRL Soft Hand was robust to an uncertainty region of 53 cm^2 during side grasps and 66 cm^2 during top grasps.

2) Grasping under Object Pose Uncertainty in an End-to-End System: For the second test, we used the same grid as before and the same location of the grid on the table, but rather than using a hard-coded location, we used the full perception-planning-grasping pipeline to detect objects, plan motions, and grasp. We allowed the robot to detect the object with the vision system while it was in the centered location. Then, while the robot planned its motion, we moved the object to the testing location. The robot planned its motion and grasp with the object in the original location, so this test examines what uncertainty in the object location the whole grasping system can handle. This takes into account not just the uncertainty from the object's different location versus anticipated location (as was tested in the previous test) but also the uncertainty from the vision system and motion planner.

For side grasps, we again used a lemonade bottle to ensure a fair comparison between this test and the previous test. Figure 14c shows the grid used in the test with dots at points which were tested. For each location, two trials were performed. For the 33 locations shown on the grid, 66 grasps were attempted, 52 of which were successful. The difference in offset in the NW/SE axis offset versus the pre-programmed scenario is most likely due to an offset in the vision system, sending the hand to a different location, on average, than the correct location. The average location that the hand went to is approximately $(.56m, .28m)$ while the ground truth location for a centered grasping pose is approximately $(.55m, .26m)$, measured in the robot coordinate system. Specifically, this means that the robot was aiming more to the right and bottom of the image, shifting the pattern of successful grasps to the top left compared to the pre-programmed position, as expected. The size of the capture area, shown in Table 3,

Table 3 Capture areas during grasping under uncertainty

Pre-programmed side grasp (Fig. 14a)	Pre-programmed top grasp (Fig. 14b)	End-to-end side grasp (Fig. 14c)	End-to-end top grasp (Fig. 14d)
53 cm ²	66 cm ²	50 cm ²	41 cm ²

reduces to 50 cm² for side grasps with the end-to-end system. This is expected as the uncertainty of the vision system and the motion planner also affects grasping performance during these tests.

For top grasps, we again used the foam block with tape as a test object for grasping. Figure 14d shows the grid used in the test with dots at points which were tested. For each location, two trials were performed, except for the orange dots, where after one success and one failure an additional trial was run. For the 25 locations shown on the grid, 52 grasps were attempted, 34 of which were successful. Again, there is a slight offset, in the same direction, between the system with vision tests and the preset grasping location, again due to an offset in the vision system object detection. The extent of the grasp area is roughly similar to the prior test, though at the edges some more objects were knocked away during the trajectory than with the perfectly straight down trajectory in the previous trials. The size of the capture area, shown in Table 3, reduces to 41 cm² for top grasps with the end-to-end system. Again, this is expected due to the uncertainty of the vision system and the motion planner.

(3) Grasping performance within an end-to-end system For the third test, we placed a single object at various locations over the table to identify where the system could pick up the object and where it failed. We executed the complete Algorithm 2 during these experiments. We considered a trial a failure if the motion plan failed, if the motion plan was not executed successfully, or if the object was not grasped successfully. Points tested on the table were 10 cm apart; we tested 11 points in the y-axis of the Baxter coordinate system and 5 points in the x-axis. We only performed grasps with the left hand; the range would increase if the right hand were used as well.

For side grasps, we used a dark cylindrical object. The results can be seen in Fig. 15a. Again, we use dots to show the data at each point; the locations represented by the dots were spread uniformly 10 cm apart in both dimensions. Two grasps were attempted at each grid point. For the 55 grid points, 110 grasps were attempted, 40 of which were successful.

Here, there are more potential causes of failure: sometimes, the vision system did not identify the object. Often, errors came from the vision system reporting an inaccurate location for the object – often due to objects far away from the camera being elongated due to perspective – so the arm knocked over the object on its trajectory over or failed to grasp it since the object was outside the successful uncer-




Fig. 15 Grasp success for grasps over the whole table. Points are 10 cm away from each other. The colors are coded as follows. ●: two successful grasps; ●: one successful grasp and one failed grasp; ●: two failed grasps; ●: one out of workspace and one successful grasp; ●: one out of workspace and one failed grasp; ●: one successful grasp and one vision failure; ●: one vision failure and one failed grasp; ●: two vision failures; ●: one vision error and one out of workspace of left arm; ●: out of workspace of left arm. **a** Side grasps, **b** top grasps (Color figure online)

tainty range. Asymmetries in the vision spotting are due to the fact that the right hand had a two-finger gripper while the left hand had a three-finger gripper which blocked more area.

For top grasps, we again used the same foam block as in the previous top grasp tests. Since the block was shorter than the cylindrical object used in the side grasp tests, the locations determined for grasps were much more accurate throughout the range of the table. This led to the increased success rate for top grasps versus side grasps. The results can be seen in Fig. 15b. Two grasps were attempted at each grid point. For the 55 grid points, 110 grasps were attempted, 51 of which were successful. Again, we use dots to show the data at each point; the locations represented by the dots were spread uniformly 10 cm apart in both dimensions.

We made two key observations during these final set of experiments.

First, **the soft hand proved robust against uncertainty in object pose during grasping**. Specifically, we observed that the soft hand was able to grasp the object reliably even when the object was more than three centimeters away from the intended grasp pose.

Second, we observed that **it was possible to integrate our soft hand easily with an existing robot platform** and perform end-to-end sense-plan-grasp operations. This was important for us, as it showed that we achieved our design goal of building a modular hand that can be integrated with existing robot platforms and these platforms' perception and planning frameworks.

Some potential methods to further improve the grasp success rate are to

- increase the width of the hand by linearly actuating the base distance between the thumb and the other fingers;
- adjust the surface shape of the finger to conform better to objects;
- vary the type of grasp depending on object type and location;
- employ a more complicated grasping strategy such as push grasp;
- improve localization accuracy of the visual detection system.

6 Conclusions and future work

This paper presents a composed robotic system with a soft gripper which can successfully grasp objects under uncertainty and identify a set of objects based on data from internal bend sensors. Internal sensing addresses one of the primary disadvantages of soft hands: the final configuration of the fingers and the final pose of the object are unknown. This system allows us to maintain the positive aspects of soft hands, including increased compliance leading to greater ability to pick up various objects with arbitrary shapes with no need for complicated grasp planning. The resulting data from the internal sensing, assumed to be independent for each finger, is sufficient when used to identify objects within a trained set of objects and to learn the identity of objects online.

We aim to improve the soft hand in future work. Additional sensors are needed for more accurate feedback while grasping. With additional sensor data, we will be able to create a more robust and accurate prediction of the configuration of the fingers, the identity of the grasped object, and the pose of the grasped object. These additional data will enable the system to identify when objects are not grasped robustly and enable them to re-grasp accordingly.

Additionally, the data provided by the sensors has the potential to enable more capabilities. The proprioceptive feedback intrinsic to the DRL soft fingers is necessary for the in-hand manipulation of objects, extending pick and place operations to complex manipulation. This data will be useful for enabling robots to use tools, picking up objects intended for use with a certain grasp and orientation, identifying the object and confirming that the orientation of the object is correct, and then planning the interaction of the grasped object with the environment to robustly use tools.

Moving robots from experimental settings to real-world settings will require not just an excellent soft manipulator, but also the base and integration necessary to allow the robot to use the soft manipulator in varying, complex environments. We plan to mount the DRL soft hand on a mobile platform for manipulation, allowing the robot to interact with objects throughout a natural human environment, updating the vision and motion planning systems to accommodate the more complex environment.

All of these manipulation skills are crucial to enable grasping robots to leave the laboratory and the automated factory to engage in work alongside humans in factories, homes, and workplaces. Anywhere robots will need to interact in human environments, robots will need to be able to have the dexterity and flexibility of grasping that humans do. We envision a future where soft hands enable that fluidity of interaction.

Acknowledgements This work was done in the Distributed Robotics Laboratory at MIT with support from The Boeing Company and the National Science Foundation, grant numbers NSF IIS1226883, and NSF CCF1138967. We are grateful for this support. The authors declare no competing financial interests.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Allen, P. K., & Roberts, K. S. (1989). Haptic object recognition using a multi-fingered dextrous hand. In *IEEE international conference on robotics and automation. Proceedings* (pp. 342–347). IEEE.
- Biddiss, E., & Chau, T. (2006). Electroactive polymeric sensors in hand prostheses: Bending response of an ionic polymer metal composite. *Medical Engineering and Physics*, 28(6), 568–578. <https://doi.org/10.1016/j.medengphy.2005.09.009>.
- Bilodeau, R. A., White, E. L., & Kramer, R. K. (2015). Monolithic fabrication of sensors and actuators in a soft robotic gripper. In *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 2324–2329). IEEE.
- Bradski, G. (2000). The OpenCV Library. Dr. Dobb's Journal of Software Tools. <https://github.com/opencv/opencv/wiki/CiteOpenCV>.
- Brown, E., Rodenberg, N., Amend, J., Mozeika, A., Steltz, E., Zakin, M. R., et al. (2010). Universal robotic gripper based on the jamming of granular material. *Proceedings of the National Academy of Sciences*, 107(44), 18809–18814.
- Caselli, S., Magnanini, C., & Zanichelli, F. (1994). Haptic object recognition with a dextrous hand based on volumetric shape representations. In *Multisensor fusion and integration for intelligent systems* (pp. 280–287).
- Chossat, J.-B., Shin, H.-S., Park, Y.-L., & Duchaine, V. (2014). Design and manufacturing of soft tactile skin using an embedded ionic liquid and tomographic imaging. *Journal of Mechanisms and Robotics*, 7, 021008.
- Chuah, M. Y., & Kim, S. (2014). Enabling force sensing during ground locomotion: A bio-inspired, multi-axis, composite force sensor using discrete pressure mapping. *IEEE Sensors Journal*, 14(5), 1693–1703.
- Danisch, L., Englehart, K., & Trivett, A. (1999). Spatially continuous six degree of freedom position and orientation sensor. *Sensor Review*, 19(2), 106–112.
- Deimel, R., & Brock, O. (2013). A compliant hand based on a novel pneumatic actuator. In *2013 IEEE international conference on robotics and automation (ICRA)* (pp. 2047–2053). IEEE.
- Deimel, R., & Brock, O. (2016). A novel type of compliant and under-actuated robotic hand for dexterous grasping. *The International*

- Journal of Robotics Research*, 35(1–3), 161–185. <https://doi.org/10.1177/0278364915592961>.
- Deimel, R., Brock, O., et al. (2014). A novel type of compliant underactuated robotic hand for dexterous grasping. *Robotics: Science and Systems*, 35, 161–185.
- Dogar, M. R., & Srinivasa, S. S. (2010). Push-grasping with dexterous hands: Mechanics and a method. In *2010 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 2123–2130). IEEE.
- Dollar, A. M., & Howe, R. D. (2006). A robust compliant grasper via shape deposition manufacturing. *IEEE/ASME Transactions on Mechatronics*, 11(2), 154–161.
- Dollar, A. M., & Howe, R. D. (2010). The highly adaptive sdm hand: Design and performance evaluation. *The International Journal of Robotics Research*, 29(5), 585–597.
- Farrow, N., & Correll, N. (2015). A soft pneumatic actuator that can sense grasp and touch. In *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 2317–2323). IEEE.
- Homborg, B., Katzschmann, R., Dogar, M., & Rus, D. (2015). Haptic identification of objects using a modular soft robotic gripper. In *Proceedings of the 2015 IEEE/RSJ international conference on intelligent robots and systems*.
- Hsiao, K., Kaelbling, L. P., & Lozano-Perez, T. (2007). Grasping pomdps. In *2007 IEEE international conference on robotics and automation* (pp. 4685–4692). IEEE.
- Ilievski, F., Mazzeo, A. D., Shepherd, R. F., Chen, X., & Whitesides, G. M. (2011). Soft robotics for chemists. *Angewandte Chemie*, 123(8), 1930–1935.
- Jain, A., Killpack, M. D., Edsinger, A., & Kemp, C. (2013). Reaching in clutter with whole-arm tactile sensing. *The International Journal of Robotics Research (IJRR)*, 32, 458–482.
- Javdani, S., Klingensmith, M., Bagnell, J. A., Pollard, N. S., & Srinivasa, S. S. (2013). Efficient touch based localization through submodularity. In *2013 IEEE international conference on robotics and automation (ICRA)* (pp. 1828–1835). IEEE.
- Johnsson, M., & Balkenius, C. (2007). Experiments with proprioception in a self-organizing system for haptic perception. *Proceedings of TAROS, 2007*, 239–245.
- Katzschmann, R. K., Marchese, A. D., & Rus, D. (2015). Autonomous object manipulation using a soft planar grasping manipulator. In *Soft robotics* (Vol. 2, No. 4, pp. 155–164). Mary Ann Liebert, Inc. 140 Huguenot Street, 3rd Floor New Rochelle, NY 10801 USA.
- Koval, M. C., Dogar, M. R., Pollard, N. S., & Srinivasa, S. S. (2013). Pose estimation for contact manipulation with manifold particle filters. In *2013 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 4541–4548). IEEE.
- Kusuda, S., Sawano, S., & Konishi, S. (2007). Fluid-resistive bending sensor having perfect compatibility with flexible pneumatic balloon actuator. In *IEEE 20th international conference on micro electro mechanical systems, MEMS* (pp. 615–618).
- Laschi, C., Mazzolai, B., & Cianchetti, M. (2016). Soft robotics: Technologies and systems pushing the boundaries of robot abilities. *Soft Robotics*, 1(1), eaah3690.
- Liarokapis, M., Calli, B., Spiers, A., & Dollar, A. (2015). Unplanned, model-free, single grasp object classification with underactuated hands and force sensors. In *Proceedings of the 2015 IEEE/RSJ international conference on intelligent robots and systems*.
- Marchese, A., Katzschmann, R., & Rus, D. (2015). A recipe for soft fluidic elastomer robots. *Soft Robotics*, 2, 7–25.
- Marchese, A. D., Komorowski, K., Onal, C. D., & Rus, D. (2014). Design and control of a soft and continuously deformable 2D robotic manipulation system. In *Proceedings of IEEE international conference on robotics and automation*. IEEE.
- Marchese, A. D., Onal, C. D., & Rus, D. (2014). Autonomous soft robotic fish capable of escape maneuvers using fluidic elastomer actuators. *Soft Robotics*, 1(1), 75–87.
- Mosadegh, B., Polygerinos, P., Keplinger, C., Wennstedt, S., Shepherd, R. F., Gupta, U., et al. (2014). Pneumatic networks for soft robotics that actuate rapidly. *Advanced Functional Materials*, 24(15), 2163–2170.
- Navarro, S. E., Gorges, N., Worn, H., Schill, J., Asfour, T., & Dillmann, R. (2012). Haptic object recognition for multi-fingered robot hands. In *2012 IEEE haptics symposium (HAPTICS)* (pp. 497–502). IEEE.
- Paolini, R., Rodriguez, A., Srinivasa, S. S., & Mason, M. T. (2014). A data-driven statistical framework for post-grasp manipulation. *The International Journal of Robotics Research*, 33(4), 600–615. <https://doi.org/10.1177/0278364913507756>.
- Park, Y.-L., Chen, B.-R., & Wood, R. J. (2012). Design and fabrication of soft artificial skin using embedded microchannels and liquid conductors. *IEEE Sensors Journal*, 12(8), 2711–2718.
- Park, Y.-L., Majidi, C., Kramer, R., Bérard, P., & Wood, R. J. (2010). Hyperelastic pressure sensing with a liquid-embedded elastomer. *Journal of Micromechanics and Microengineering*, 20(12), 125029.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.
- Polygerinos, P., Correll, N., Morin, S. A., Mosadegh, B., Onal, C. D., Petersen, K., et al. (2017). Soft robotics: Review of fluid-driven intrinsically soft devices; Manufacturing, sensing, control, and applications in human-robot interaction. *Advanced Engineering Materials*, 19, 1700016. <https://doi.org/10.1002/adem.201700016>.
- Polygerinos, P., Lyne, S., Wang, Z., Nicolini, L. F., Mosadegh, B., Whitesides, G. M., & Walsh, C. J. (2013). Towards a soft pneumatic glove for hand rehabilitation. In *2013 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 1512–1517). IEEE.
- Quigley, M., Conley, K., Gerkey, B. P., Faust, J., Foote, T., Leibs, J., Wheeler, R., & Ng, A. Y. (2009). Ros: An open-source robot operating system. In *ICRA workshop on open source software*.
- Rus, D., & Tolley, M. (2015). Design, fabrication, and control of soft robots. *Nature*, 521, 467.
- Shepherd, R. F., Stokes, A. A., Nunes, R., & Whitesides, G. M. (2013). Soft machines that are resistant to puncture and that self seal. *Advanced Materials*, 25(46), 6709–6713.
- Shih, B., Drotman, D., Christianson, C., Huo, Z., White, R., Christensen, H. I., & Tolley, M. T. (2017). Custom soft robotic gripper sensor skins for haptic object visualization. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 494–501). Vancouver, BC, Canada: IEEE.
- Stokes, A. A., Shepherd, R. F., Morin, S. A., Ilievski, F., & Whitesides, G. M. (2014). A hybrid combining hard and soft robots. *Soft Robotics*, 1(1), 70–74.
- Sucan, I. A., & Chitta, S. (2018). MoveIt! <http://moveit.ros.org> <http://moveit.ros.org/about/>. Accessed 16 Apr 2018.
- Takamuku, S., Fukuda, A., & Hosoda, K. (2008). Repetitivegrasping with anthropomorphic skin-covered hand enables robust haptic recognition. In *IEEE/RSJ international conference on intelligent robots and systems, IROS 2008* (pp. 3212–3217). IEEE.
- Vogt, D., Park, Y.-L., & Wood, R. J. (2013). Design and characterization of a soft multi-axis force sensor using embedded microfluidic channels. *IEEE Sensors Journal*, 13(10), 4056–4064.
- Weiß, K., & Worn, H., (2005). The working principle of resistive tactile sensor cells. In *2005 IEEE international conference mechatronics and automation* (Vol. 1, pp. 471–476). IEEE.
- Zhao, H., O'Brien, K., Li, S., & Shepherd, R. F. (2016). Optoelectronically innervated soft prosthetic hand via stretchable optical waveguides. *Science Robotics*, 1(1). <https://doi.org/10.1126/scirobotics.aai7529>.



Bianca S. Homberg is currently a robotics engineer at Iron Ox. Prior to that she was a software engineer at Google Robotics and a Masters of Engineering student in the Distributed Robotics Lab at CSAIL, MIT. She holds a Bachelor of Science from the Massachusetts Institute of Technology (MIT) in Electrical Engineering and Computer Science, and Mathematics.



Mehmet R. Dogar leads the robotic manipulation research at the School of Computing at the University of Leeds. Previously, he was a postdoctoral researcher at CSAIL, MIT. He received his Ph.D. from the Robotics Institute at Carnegie Mellon University. His research focuses on autonomous robotic manipulation. He envisions a future where robots autonomously perform complex manipulation tasks in human environments.



Robert K. Katzschmann is a Ph.D. Candidate with the Department of Mechanical Engineering at the Massachusetts Institute of Technology (MIT) and a researcher in soft robotics at the Distributed Robotics Laboratory at CSAIL. Robert's research focus is on the design, fabrication and control of soft-bodied robotic systems such as robotic arms, grippers and fish. Prior to his Ph.D., he worked from 2012 to 2013 as systems engineer at Auris Surgical Robotics developing a robot for minimally invasive surgeries and endoscopies. His master's work in 2012 was on dynamic motion planning and trajectory generation of robotic manipulator arms at Stanford University.

His master's work in 2012 was on dynamic motion planning and trajectory generation of robotic manipulator arms at Stanford University.



Daniela Rus is the Andrew and Erna Viterbi Professor of Electrical Engineering and Computer Science and Director of the Computer Science and Artificial Intelligence Laboratory at MIT. She serves as the Director of the Toyota-CSAIL Joint Research Center and is a member of the science advisory board of the Toyota Research Institute. Rus' research interests are in robotics, mobile computing, and data science. Rus is a Class of 2002 MacArthur Fellow, a fellow of ACM, AAAI and IEEE, and a member of the National Academy of Engineering. She earned her Ph.D. in Computer Science from Cornell University.

She earned her Ph.D. in Computer Science from Cornell University.