Inferring Object Properties with a Tactile-Sensing Array Given Varying Joint Stiffness and Velocity

Tapomayukh Bhattacharjee*, James M. Rehg[†] and Charles C. Kemp[‡]

*Institute of Robotics and Intelligent Machines,
Georgia Institute of Technology,
Atlanta GA 30332, USA*

**tapomayukh@gatech.edu*

[†]rehg@cc.gatech.edu

[‡]charlie.kemp@bme.gatech.edu

Received 2 April 2017

Accepted 13 July 2017

Published 23 October 2017

Whole-arm tactile sensing enables a robot to sense contact and infer contact properties across its entire arm. Within this paper, we demonstrate that using data-driven methods, a humanoid robot can infer mechanical properties of objects from contact with its forearm during a simple reaching motion. A key issue is the extent to which the performance of data-driven methods can generalize to robot actions that differ from those used during training. To investigate this, we developed an idealized physics-based lumped element model of a robot with a compliant joint making contact with an object. Using this physics-based model, we performed experiments with varied robot, object and environment parameters. We also collected data from a tactile-sensing forearm on a real robot as it made contact with various objects during a simple reaching motion with varied arm velocities and joint stiffnesses. The robot used 1-nearest-neighbor (1-NN) classifiers, hidden Markov models (HMMs), and long short-term memory (LSTM) networks to infer two object properties (hard versus soft and moved versus unmoved) based on features of time-varying tactile sensor data (maximum force, contact area, and contact motion). We found that, in contrast to 1-NN, the performance of LSTMs (with sufficient data availability) and multivariate HMMs successfully generalized to new robot motions with distinct velocities and joint stiffnesses. Compared to single features, using multiple features gave the best results for both experiments with physics-based models and a real-robot.

Keywords: Haptics; tactile sensing; hidden Markov models; long short-term memory networks; k -nearest-neighbors; object categorization; physics-based models.

1. Introduction

Manipulation in unstructured environments with high clutter is difficult due to a variety of factors, including uncertainty about the state-of-the-world, a lack of

This is an Open Access article published by World Scientific Publishing Company. It is distributed under the terms of the Creative Commons Attribution 4.0 (CC-BY) License. Further distribution of this work is permitted, provided the original work is properly cited.

noncontact trajectories, and reduced visibility for line-of-sight sensors.¹ Tactile sensing is well-matched to these challenges, since it benefits from contact and uses sensors that move with the manipulator into the clutter. When contact occurs with tactile sensors, the robot has an opportunity to acquire information. By fully covering the robot's manipulator with tactile sensors, the robot is likely to have more opportunities to acquire useful information through contact. However, with a typical serial manipulator, a robot cannot independently control the pose of each of the sensors. In addition, contact may not be anticipated.

Within this paper, we address the problem of haptic perception based on contact²⁻⁴ with a tactile-sensing forearm on a humanoid robot. The time-varying signals from a tactile-sensing array depend on the robot's actions, including the joint stiffnesses and joint velocities of the robot. A key problem for data-driven approaches is how to infer object properties based on these signals when the robot's actions are different from those used during training. In other words, after a robot has learned about an object using one action, it would ideally be able to infer the same properties of the object when making contact with it using a different action. In this paper, we consider an example of this type of problem. *Specifically, we focus on a robot inferring object properties with a tactile-sensing array when the robot's joint stiffness and joint velocity differ from those used during training.*

The type of action we consider in this work is a short compliant movement of a robot's forearm akin to movements that occur during reaching. Figure 1 shows one such example when the robot, Cody, came into contact with a cylindrical object made of polystyrene foam, while trying to reach a goal configuration. We intentionally do not have the robot use exploratory behaviors, and instead investigate inference from short duration contact (i.e., 5 s and 1.2 s). Our choice of action is inspired by the potential for robots to infer useful properties of the world based on *incidental contact*. By incidental contact, we mean contact that is not central to

Fig. 1. Force data from forearm skin sensor mounted on the robot, Cody. The forearm came into contact with a cylindrical object made of polystyrene foam, while trying to reach a goal configuration. The red arrows show the forces acting on the skin.

the robot's current actions and may occur unexpectedly or unintentionally.^a For example, while manipulating in cluttered environments, incidental contact is more likely to occur and can be common with some approaches to robot control.^{1,5,6} Incidental contact will typically not involve active exploration and interrogation of the contact, since the robot will be directing its resources elsewhere.^{3,4} Each such contact event is an opportunity for sensing. Unlike deliberate probing, during which the robot has more control over its actions to optimize its sensing, the sensing is opportunistic during these motions and contact events can be brief and simple, which could make inference challenging.

1.1. Opportunities and challenges of haptic perception

Inferring mechanical properties of objects from contact could be beneficial in a number of ways. For example, we have shown that haptically recognizing leaves versus trunk while a robot reaches into artificial foliage can be used to haptically map the environment and plan paths to goals.^{3,4} Rather than recognizing a particular object type, detecting an object's properties could be informative in novel environments. Detecting that an object has *moved* due to a robot's actions might be used by the robot to make better decisions, such as moving the object away in order to access a new location, or avoiding the object, so as not to alter the environment further. Likewise, detecting that an object is *hard* or *soft* has implications for the robot's ability to compress the object and the consequences of collisions with the object.

One of the challenges of haptics is that sensing depends on action. For example, haptic perception of surface roughness depends on contact speed⁷ and contact force.⁸ Many existing tactile systems use carefully controlled exploratory behaviors to reduce variability of the actions and gather information about the object (see Sec. 2). The signals produced from a robot's tactile sensors depend on the mechanics of both the object and the robot. Physical interaction depends on the dynamics of impact between a robot and an object. For example, *depending upon how stiff or compliant a robot is or how fast a robot moves, the physical interaction will vary.* As such, a key issue for data-driven approaches to tactile perception of contact is the extent to which the performance of perceptual classifiers can generalize when a robot's actions differ from those used to collect training data.

Classifying an object based on its compliance or mobility can become challenging when robot joint stiffness or velocity changes. For example, interaction forces are a function of the robot joint stiffness as well as object stiffness. The forces generated by contact between a stiff robot and a hard object are not the same as the forces generated by contact between a compliant robot and a hard object. Hence, perceiving an object as hard/soft is challenging if the robot stiffness changes, because it is difficult to distinguish the interaction between a stiff robot and a soft object and a compliant robot and a hard object. Similarly, it may be difficult to determine if

^aThis description supersedes our previous descriptions from Refs. 3 and 4.

an object has moved based on tactile sensing if the robot actions change. A movement in the robot's arm after contact can be due to the actual movement of the object (e.g., sliding motion) or because of the compliance in the structure of the object (the robot arm pushes into the object and the object deforms). A compliant robot moving with a low velocity could push into a soft object resulting in robot arm movement after initial contact, which could be similar to the actual movement of a hard object (without any deformation) generated by a stiff robot impacting with a high velocity. In this work, we focus on addressing these challenges by:

- classifying an object into one of the four categories: (1) *hard-unmoved* (HU), (2) *hard-moved* (HM), (3) *soft-unmoved* (SU), and (4) *soft-moved* (SM), and
- investigating the potential of data-driven methods to generalize the performance of haptic perception to different robot conditions such as robot stiffness and velocity used to collect training data.

1.2. Methods

We used univariate and multivariate hidden Markov models (HMMs) as well as long short-term memory (LSTM) networks to infer object properties (Sec. 3) and compared the results with 1-nearest-neighbor (1-NN) used in our previous work.² In Sec. 4, we present an idealized physics-based lumped element model of a robot with a compliant joint making contact with an object. Using this model, we performed experiments with physics-based simulations of varied robot, object, and environment parameters in Sec. 5. In Sec. 6, we present experiments with a real robot for which we varied the robot arm's velocity and joint stiffness to values distinct from those used during training. Our multivariate HMM-based method performed well in experiments with both physics-based simulations (Sec. 5.5) and with a real robot (Sec. 6.5) for classifying objects into the four categories. Our LSTM networks performed better when a larger amount of data was available with the physics-based simulation compared to experiments with the real robot. Section 7 discusses the methods and their limitations and Sec. 8 provides the conclusion of our work.

2. Related Work

Object categorization is a well-studied task. In this work, we focus on inferring properties of objects based on haptic sensing. Although there have been multiple studies on haptics-based compliance discrimination, most have used specific exploratory behaviors using end effectors to extract information from the environment. Lederman and Klatzky discussed in detail the various factors and exploratory behaviors that humans use for inferring haptic properties of objects.⁹ However, studies of discrimination tasks using information from incidental contact with large-area tactile sensing are lacking. Also, studies on inferring properties of objects and *generalizing the performance across various conditions such as robot stiffnesses and velocities* are rare.

Researchers have also used haptics and tactile sensing to infer properties of the world for purposes other than categorizing objects. Silvera-Tawil *et al.* presented a comprehensive review of state-of-the-art methods in tactile sensing for robots in socially interactive scenarios.¹⁰ Wu *et al.* used tactile arrays for recognizing human-intended directions in two dimensions using support vector machine (SVM) classifiers.¹¹ Muscari *et al.* developed algorithms for reconstructing force and shape distributions using capacitive tactile-sensing arrays.¹² Hughes and Correll presented a soft robotic artificial skin for texture recognition and localization.¹³ Javaid *et al.* used pressure sensors to classify human activities during assistive tasks to help older adults.¹⁴ Matheus and Dollar used a load cell and a custom-built device to infer static friction properties of some “Objects of Daily Living”.¹⁵ Boonvisut and Çavuşoglu used exploratory behaviors and vision to collect deformation data for identifying boundary constraints of deformable objects required to estimate soft tissue properties.¹⁶ Researchers have used tactile sensors to classify events such as slips between fingertips and objects and slips between objects and external surfaces.^{17–19} Researchers have also used haptic sensing for texture perception,^{20–22} tactile servoing,²³ contour following,²⁴ and human–robot collaborative tasks.^{25,26}

In the following subsections, we review the existing literature that addresses object categorization using haptics.

2.1. Material property-based classification

In this work, we do not explicitly model material properties, but the features that we extract from the interactions between the robot arm and environmental objects are a direct consequence of material properties. Drimus *et al.* classified hard and deformable objects based on haptic feedback from a novel tactile sensor consisting of a flexible, piezo-resistive rubber.²⁷ They represented tactile information from a palpation procedure as a time series of features and used a k -nearest-neighbor (k -NN) classifier to categorize the objects.²⁷ Our classification scheme considers both compliance and mobility characteristics and uses information from incidental contact sensed with large-area tactile sensors.

Chu *et al.* presented research that uses discrete HMMs to construct a feature vector of likelihoods and used binary SVM classifiers to classify those vectors and automatically assign 24 adjectives to 60 objects.²⁸ In contrast to our work, their research focused on classifying data using both static and dynamic features from four deliberate exploratory procedures with sophisticated BioTac²⁹ robotic fingers from *Syntouch*.

Kaboli *et al.* used multi-modal tactile features to distinguish texture and weight of objects using sliding and nonsliding exploratory behaviors.³⁰ Kim and Kesavadas presented a method for estimating the material properties of objects (steel, aluminum, wood, silicon rubber) using an active tapping procedure.³¹ Takamuku *et al.* estimated the hardness properties of objects through tapping and squeezing behaviors.³² Hosoda and Iwase obtained haptic data using a bionic hand to grip an object.

They used a recurrent neural network to classify objects based on haptic cues learned from dynamic interactions.³³ Charniya and Dudul classified the material type and surface properties by developing a sensor that used a lightweight plunger probe to detect surface properties. They also used an optical mouse sensor to obtain surface images and implemented a radial basis function neural network for classification.³⁴ Liarokapis *et al.* used random forests to classify size and stiffness of objects as well as distinguish object types using a single force closure grasp with an underactuated robotic hand and force sensors.³⁵ Schmitz *et al.* used power grasping of objects and multiple modalities for object recognition with deep learning.³⁶ Kiwatthana and Kaitwanidvilai used system identification and *K*-means clustering techniques to classify different cans using proprioceptive feedback.³⁷ Hoelscher *et al.* used BioTac fingers and Schunk F/T sensors with multiple classifiers and feature extraction methods for object recognition. They concluded that simple, dimensionally-reduced features performed better than more elaborate features.³⁸

In summary, although there have been many studies on material property-based classification, most have focused on carefully-controlled specific exploratory behaviors using the robot's end-effector. These studies have not looked at whether the classification performance can generalize to different robot behaviors.

2.2. Shape-based classification

Many researchers have either used tactile images from touch sensors or analyzed object shape deformation behaviors for object categorization. Schneider *et al.* applied a "bag-of-words" approach and unsupervised clustering techniques to categorize objects.³⁹ Pezzementi *et al.* identified the principal components of features, then clustered them, and constructed per-class histograms as a class characteristic.⁴⁰ Gorges *et al.* introduced passive joints in the hand for better adaptability to different object shapes and used a Bayes classifier to classify the objects.⁴¹ Babu *et al.* used "C4.5" algorithm to generate a decision tree and a naive Bayes classifier to categorize shapes of objects using a tactile-sensor array.⁴²

Other researchers have analyzed deformation behavior to classify objects. They have used vision and haptic sensors^{43,44} or finite element models^{45,46} and volumetric models such as superquadrics,⁴⁷ polyhedral models⁴⁸ or wrapping polyhedra.⁴⁹

To summarize, shape-based classification methods have used tactile images or deformation behaviors to classify objects after exploring or grasping them using multi-fingered robot hands. Again, most of these studies used exploratory behaviors and have not looked at whether the performance can generalize to different robot behaviors used to collect training data.

2.3. Functional property-based classification

This group of studies focused on functional property-based classification or classification based on how objects behave when they are moved. Sinapov *et al.* used the acoustic properties of objects during specific interaction schemes and the behavioral

interactions performed with them, such as grasping, shaking, dropping, pushing, and tapping, to classify 36 different household objects.⁵⁰ Berquist *et al.* monitored the changes in the joint torques of a robot while it performed five exploratory procedures — lifting, shaking, crushing, dropping, and pushing — on several objects and demonstrated that the robot could learn to recognize objects based solely on the joint-torque information.⁵¹ Jain and Kemp used data-driven object-centric models to haptically recognize specific doors as well as classes of doors (refrigerator versus kitchen cabinet).⁵² Griffith *et al.* used multiple exploratory behaviors and employed clustering techniques to categorize containers and noncontainers. After extracting visual and acoustic features from interactions with objects, they employed unsupervised clustering techniques to form several categories.⁵³ Sinapov *et al.* combined proprioceptive and auditory feedback and used a behavior-grounded relational classification model to recognize categories of household objects.⁵⁴ Sinapov and Stoytchev extended their previous work by using auditory, proprioceptive, and visual modalities to cluster novel and unlabeled objects for object individuation based on the robot's sensorimotor experience of handling those novel objects.⁵⁵

To summarize, functional property-based classification methods have used multimodal feedback from robot behaviors and actions to distinguish between different object categories. But, similar to other work, most of them used exploratory behaviors and have not looked at if the results can generalize to different robot behaviors.

3. Data-Driven Methods

For this work, we chose HMMs and LSTM networks as our data-driven methods to infer object properties from contact. To compare our results, we used k -NNs from our previous work² as the baseline data-driven method.

3.1. Hidden Markov models

HMMs have a long history of success for classifying time series such as human speech.⁵⁶ HMMs are known to have rich mathematical structure for modeling nonstationary signals and can work effectively in practical applications.⁵⁶ In this study, we used HMMs to classify an object as being in one of four categories: (1) *hard-unmoved* (HU), (2) *hard-moved* (HM), (3) *soft-unmoved* (SU), and (4) *soft-moved* (SM).

A HMM is a state-based data modeling tool that assumes the states are hidden and the system is a Markov process. The hidden states (HS_i) are inferred using observations (O_i) at time t . The components of an HMM include N , the number of states in the model; $\mathbf{A}$, the state transition probabilities; $\mathbf{B}$, the observation probabilities; and π , the initial state probabilities. We use notation from Ref. 56. Equation (1) shows the HMM model (λ). Equation (2) shows the state transition probabilities:

$$\lambda = (\mathbf{A}, \mathbf{B}, \pi), \quad (1)$$

$$\mathbf{A} = \{a_{ij}\} = \{P(x_t = HS_j | x_{t-1} = HS_i)\}, \quad 1 \leq i, j \leq N, \quad (2)$$

where i, j are state indexes and

$$\begin{aligned} a_{ij} &\geq 0, \\ \sum_{j=1}^N a_{ij} &= 1. \end{aligned} \quad (3)$$

Note, for a left-right HMM, as time increases, the state index increases or stays the same. Therefore, there is an additional constraint given by

$$a_{ij} = 0, \quad j < i. \quad (4)$$

Equation (5) represents the initial state probabilities:

$$\pi = \{\pi_i\} = \{P(x_1 = HS_i)\}, \quad 1 \leq i \leq N. \quad (5)$$

For a discrete HMM with M distinct observation symbols, Eq. (6) shows the observation probabilities:

$$\mathbf{B} = \{b_j(k)\} = \{P(O_t = v_k | x_t = HS_j)\}, \quad 1 \leq j \leq N, \quad 1 \leq k \leq M, \quad (6)$$

where $\mathbf{V} = \{v_1, v_2, \dots, v_M\}$ are the individual symbols. However, for a continuous HMM (which we use in this work) with observation vector $\mathbf{O}$ and multivariate Gaussian emissions with mean vector $\boldsymbol{\mu}$ and covariance matrix $\mathbf{U}$, Eq. (7) shows the emission probabilities:

$$\mathbf{B} = \{b_j(\mathbf{O})\} = \{P(\mathbf{O} | x_t = HS_j)\} = \{\mathcal{N}[\mathbf{O}, \boldsymbol{\mu}_j, \mathbf{U}_j]\}, \quad 1 \leq j \leq N. \quad (7)$$

Figure 2 (Left) shows the schematic of one left-right HMM with $N = 10$ states.

3.2. Long short-term memory networks

LSTM networks have been successfully used for modeling time-series in many applications such as machine translation, generating cursive writing, and speech

Fig. 2. Schematic of a left-right HMM with 10 states. The observations are modeled using Gaussian distributions (Left). Schematic of a stacked LSTM network with 50 cells in each layer, dropout in-between, and a fully connected layer as the output layer (Right).

recognition.^{57–59} For mathematical details, please refer to Ref. 60, where this was first introduced. LSTMs have also been successfully used for haptic perception such as to estimate forces during robot-assisted dressing simulations,⁶¹ and during robot-assisted surgery.⁶² Note, Gao *et al.*⁶³ used LSTMs for haptic perception using visual and tactile features but the performance was lower compared to other deep-learning methods.

For our applications, we use an LSTM structure inspired by Gers *et al.*⁶⁴ in which each memory cell has an input gate, a forget gate, and an output gate. The specific structure that we used for our haptic classification tasks consists of two layers with 50 cells each. To reduce overfitting, we also added a dropout layer in-between the two layers, which helps in regularization.⁶⁵ We added a fully connected output layer. Section 5.4 shows the details of the implementation. The LSTM has a total of 31,004 parameters. The 31,004 parameters correspond to the weights of the stacked LSTM layers, the individual cells, as well as the fully connected layer for the 4-category classification output.

Figure 2 (Right) shows the schematic of the LSTM structure.

4. Physics-Based Model

We modeled contact between a tactile-sensing robot forearm and an object using a lumped element model. We developed a physics-based model that can model the mechanics of a variety of robot–object physical interaction phenomena. In this work, the tactile-sensing robot arm moves by actuating its shoulder joint towards a goal angle. The arc formed by the contact point during the motion is approximately a straight line for small angle movements and a large radius. We modeled the robot arm’s motion at the contact point as a linear movement. We modeled the object as a deformable object and the contact surface as flat.

4.1. Robot and object model

We modeled the robot’s arm trajectory using equilibrium point control.⁶⁶ Figure 3 (Left) shows a robot-arm of mass m_{arm} making contact with an object of mass m_{obj} . x_{arm} is the position of the robot-arm, x_{obj} is the position of the object, and x_{uobj} is the position of the undeformed object. During contact, if δ_{obj} is the object deformation, note that

$$x_{\text{obj}} = x_{\text{arm}} = x_{\text{uobj}} + \delta_{\text{obj}}, \quad (8)$$

otherwise, if the arm is not in contact with the object,

$$x_{\text{obj}} = x_{\text{uobj}} > x_{\text{arm}}. \quad (9)$$

x_{eq} is the equilibrium point of the actuator spring with stiffness k_{act} and actuator damping b_{act} . k_{obj} is the object stiffness. F_{fr} is the frictional force between the object and the environment.

Fig. 3. Lumped element model of our system at the onset of contact between the robot-arm and the object (Left). Free-body diagrams of the robot-arm and object in contact (Right).

Figure 3 (Right) shows free-body diagrams for the system depicted in Fig. 3 (Left). F_{act} is the force applied on the robot-arm by the actuator, F_{sur} is the force applied by the robot-arm to the surface of the object when in contact. F_{arm} and F_{obj} are the net forces acting on the robot-arm and the object, respectively.

The force applied by the actuator to the robot-arm, F_{act} , is given by

$$F_{act} = k_{act}(x_{eq} - x_{arm}) + b_{act}(\dot{x}_{eq} - \dot{x}_{arm}). \quad (10)$$

The net force on the arm, F_{arm} , is therefore calculated as

$$F_{arm} = F_{act} - F_{sur}. \quad (11)$$

The net force on the object, F_{obj} , is given by

$$F_{obj} = F_{sur} - F_{fr}. \quad (12)$$

We model the joint as frictionless. Hence, the position of the robot-arm can be calculated by Eq. (13), where $\dot{x}_{0_{arm}}$ is the initial velocity and $x_{0_{arm}}$ is the initial position of the arm:

$$\begin{aligned} \dot{x}_{arm} &= \int (F_{arm}/m_{arm}) dt + \dot{x}_{0_{arm}}, \\ x_{arm} &= \int \dot{x}_{arm} dt + x_{0_{arm}}. \end{aligned} \quad (13)$$

The position of the undeformed object is calculated using Eq. (14), where $\dot{x}_{0_{obj}}$ is the initial velocity and $x_{0_{obj}}$ is the initial position of the undeformed object:

$$\begin{aligned} \dot{x}_{uobj} &= \int (F_{obj}/m_{obj}) dt + \dot{x}_{0_{obj}}, \\ x_{uobj} &= \int \dot{x}_{uobj} dt + x_{0_{obj}}. \end{aligned} \quad (14)$$

F_{sur} is calculated using

$$F_{sur} = k_{obj}(x_{arm} - x_{uobj}). \quad (15)$$

Please note that the frictional force F_{fr} is calculated differently depending on whether the object is moving or not as shown in Eq. (16). If the applied force overcomes static friction, the object starts moving. μ_s and μ_k are the coefficients of static and kinetic friction, respectively and $g = 9.81 \text{ m/s}^2$:

$$F_{fr} = \begin{cases} F_{surf}, & \text{stationary,} \\ \mu_s(m_{obj}g), & \text{just before motion,} \\ \mu_k(m_{obj}g), & \text{in motion.} \end{cases} \quad (16)$$

4.2. Contact area model

For an object in the shape of a cube with edge length L_{obj} , and surface area A_{obj} , let us assume that due to a robot-arm's applied force F_{surf} , there is a positive axial compression given by ΔL_{obj} in the x -direction (see Fig. 4). Let the cube be made up of a homogeneous material with Poisson's ratio ν_{obj} (Ref. 67) and due to the applied force, let the elongations in the y - and z -directions be $\Delta L'_{obj}$. Before deformation, the surface area is given by

$$A = L^2. \quad (17)$$

After deformation, the surface area becomes

$$A' = (L + \Delta L')^2. \quad (18)$$

Therefore, the increase in the surface area, due to the applied force, becomes

$$\Delta A = A' - A = (L + \Delta L')^2 - L^2, \quad (19)$$

which leads to

$$\frac{\Delta A}{A} = \left(1 + \frac{\Delta L'}{L}\right)^2 - 1. \quad (20)$$

Let the strains in x -, y -, and z -axis due to the applied force F_{surf} be given by $d\epsilon_x = \frac{dx}{x}$, $d\epsilon_y = \frac{dy}{y}$, and $d\epsilon_z = \frac{dz}{z}$, respectively. We assume that when the object is compressed,

Fig. 4. Due to applied force in x -direction, an object undergoes axial compression. This results in elongation in the other orthogonal axes directions, thus increasing the surface contact area proportional to the object material's Poisson's ratio.

the transverse strain is positive:

$$\nu_{\text{obj}} = -\frac{d\varepsilon_y}{d\varepsilon_x} = -\frac{d\varepsilon_z}{d\varepsilon_x}. \quad (21)$$

Therefore, using the above relations, we have from Ref. 68:

$$\left(1 - \frac{\Delta L}{L}\right)^{-\nu_{\text{obj}}} = \left(1 + \frac{\Delta L'}{L}\right). \quad (22)$$

Combining Eqs. (22) and (20), we have

$$\frac{\Delta A}{A} = \left(1 - \frac{\Delta L}{L}\right)^{-2\nu_{\text{obj}}} - 1. \quad (23)$$

Therefore, due to the applied force and resultant deformation, the new surface area is given by

$$A' = A + \Delta A = A \left(1 - \frac{\Delta L}{L}\right)^{-2\nu_{\text{obj}}}, \quad (24)$$

where $\Delta L = \frac{F_{\text{surf}}}{k_{\text{obj}}}$.

5. Experiments with Physics-Based Simulations

Using our physics-based model from Sec. 4, we generated simulated data and used the data for our experiments. The simulations with the physics-based model enabled us to generate data with a wide variety of robot, object, and environment conditions, which would be challenging with experiments with a real robot (Sec. 6). We focused on whether our algorithms could classify objects into four different categories and whether the classification performance could generalize to robots having varying mechanical characteristics. By *varying mechanical characteristics*, we mean motion of the robot in which the robot arm's joint stiffness and velocity are varied.

5.1. Experimental setup

We modeled the objects as cubes of 10 different volumes. We varied the volumes of the cubes linearly. The edge-length varied from $l_{\text{obj}} = 0.01$ m to $l_{\text{obj}} = 0.2$ m. We chose 13 different materials to model the objects for our simulations. We chose materials representative of everyday household objects and the human body (see Table 1). We calculated the mass of the object, m_{obj} , using Eq. (25), where ρ_{obj} is the density of the object. To calculate the friction coefficients, we modeled the objects resting on two kinds of surfaces made of wood and glass found on commonly-used table tops (see Table 1).

We labeled objects as hard if the calculated stiffness (k_{obj}) was greater than 100,000 N/m which is the stiffness of a medium-stiff environment.⁹³ Note, stiffness is a property of object structure as well as its material and is distinct from the Young's modulus of a material which is a material property. For example, a flat thin

Table 1. Material selection for simulations.

General material	Example of a common object type	Young's modulus (E) (Refs. 78, 83, 89 and 90) N/m^2	Density (ρ_{obj}) (Refs. 78, 83–87) kg/m^3	Friction coefficients ^{69–82}				Ratio (ν)* (Refs. 88, 91 and 92)
				Wood surface		Glass surface		
				Static (μ_s)	Kinetic (μ_k)	Static (μ_s)	Kinetic (μ_k)	
ABS plastic	Beverage containers	1.60E+09	1110.00	0.5	0.4	0.3	0.25	0.41
Glass	Beverage containers	6.25E+10	2250.00	0.17	0.14	0.95	0.35	0.2
Pine wood	Furniture	1.68E+10	520.00	0.5	0.364	0.36	0.14	0.38
Ceramic	Countertops	7.00E+10	2250.00	0.5	0.4	0.42	0.35	0.18
Steel	Appliances	2.00E+11	7850.00	0.61	0.35	0.13	0.12	0.27
Polymer foam	Mattresses	2.50E+05	54.00	0.12	0.11	0.37	0.3	0.27
Light foam	Pillows	3.30E+04	16.00	0.12	0.11	0.37	0.3	0.28
Soft foam	Seat cushions	1.00E+04	48.06	0.12	0.11	0.55	0.5	0.30
Open cell foam	Sponges	2.20E+04	72.08	0.12	0.11	0.55	0.5	0.30
Natural rubber	Footwear	2.00E+06	925.00	0.9	0.7	0.87	0.7	0.50
Neoprene rubber	Clothing and bags	1.35E+06	1240.00	0.9	0.7	0.87	0.7	0.49
Fat (tissues)	Human body	1.90E+03	919.60	0.91	0.6	0.45	0.3	0.41
Muscle (thigh)	Human body	1.28E+04	1060.00	0.91	0.6	0.45	0.3	0.30

Note: *Ratio = Poisson's ratio.

rectangular block of plastic and a long slender cylindrical plastic bottle may have same Young’s modulus (both are made of plastic, assuming the material is homogeneous) but they may have completely different stiffness due to different structural properties. If the calculated stiffness of the object (k_{obj}) was less than or equal to 100,000 N/m, we labeled it as “soft” (see Sec. 7 for more detailed discussion on the “hardness” of an object). We labeled objects as “moved” or “unmoved” based on whether x_{obj} was greater than 0 or equal to 0, respectively, at the end of the simulation. We ran simulations with the physics-based model for 5 s. The mass of the robot, m_{arm} , was 1.167 kg based on the model of the forearm of the robot “Cody” used in our experiments. During the simulations, the robot-arm model came in contact with objects in the shape of a cube of various edge-lengths. Each object of mass m_{obj} is made up of a single material from Table 1:

$$m_{\text{obj}} = \rho_{\text{obj}} l_{\text{obj}}^3. \quad (25)$$

For each object, we calculated the stiffness as

$$k_{\text{obj}} = 2l_{\text{obj}} E_{\text{obj}}, \quad (26)$$

where E_{obj} is the Young’s modulus of the material of the object. We derived 26 for an object under both cantilever and axial loading.

5.2. Experimental procedure

For the experiments, we simulated the robot moving using varying stiffness and velocity. Specifically, we trained the algorithms with three of the four possible combinations of stiffness and velocity conditions (low-velocity-low-stiffness, low-velocity-high-stiffness, high-velocity-low-stiffness, and high-velocity-high-stiffness) and tested with the other combination to find out how well the results generalized to different robot conditions. We repeated this procedure for each of the four conditions and reported the mean classification accuracy for all the conditions. We set the robot’s stiffness (k_{act}) to a low value (543 N/m) or a high value (2050 N/m) based on the stiffness values identified in Ref. 94 and the velocity ($\dot{x}_{\text{arm}}$) to a low value (0.005 m/s) or a high value (0.02 m/s). These values correspond with the values used in our experiments with a real robot in Sec. 6.2. We set the damping (b_{act}) based on a critically damped system as

$$b_{\text{act}} = 2\eta\sqrt{k_{\text{act}} m_{\text{arm}}}, \quad (27)$$

where $\eta = 1$ is the damping ratio of a critically-damped system.⁹⁵

5.2.1. Data collection

The simulated robot made contact with a set of solid cube objects (see Sec. 5.1) made of materials given in Table 1 while performing a simple, goal-directed reaching motion as shown in Sec. 4.1. We simulated the robot–object interactions on both a wooden and a glass table. We simulated the robot movement using an

equilibrium-point control similar to that of our experiments with the real robot. We actuated the simulated robot by commanding a goal location and the robot moved according to a PD controller. The final goal location was inside the object.

Our algorithm classified the objects into the four categories “hard-unmoved”, “hard-moved”, “soft-unmoved”, and “soft-moved”. There were 1835 simulation trials with 544 “hard-unmoved”, 365 “hard-moved”, 496 “soft-unmoved” and 430 “soft-moved” trials. Note that for each of the trials in which an object moved, we fixed it to generate one additional “unmoved” trial, similar to our experiments with the real robot.

5.2.2. Preprocessing and feature selection

We truncated the data to begin at the estimated onset of contact (whenever the force exceeds 0 N) between the robot and the object. We collected three time-varying features at each time-instant for 5 s at 100 Hz. Our “force” feature is F_{surf} (see Sec. 4.1), “contact area” feature is A' (see Sec. 4.2), and “motion” feature is x_{arm} (see Sec. 4.1). Note, if the robot loses contact with the object at any time-instant ($F_{\text{surf}} = 0$), A' becomes A . However, to match the scenario of experiments with the real robot (see Sec. 6.2.2), we make the “contact area” feature 0 for those time-instants during preprocessing. We expected these three features (force, contact area, and motion) to be informative about the object’s softness and mobility. These features are similar to the features selected in our experiments with a real-robot (see Sec. 6.2.2).

The frequency of the signal from tactile-sensing forearm on the real robot is 100 Hz. We modeled an analog-to-digital anti-aliasing filter for the tactile-sensing forearm using a low-pass Butterworth filter of order 6 and cut-off frequency of 200 Hz. We also modeled the tactile-sensing forearm noise as Gaussian with zero mean and a signal-to-noise ratio of 0.5% for force measurements and zero mean and a signal-to-noise ratio of 0.05% for contact area measurements. We modeled the joint encoder noise as Gaussian with zero mean and a signal-to-noise ratio of 0.05%. The signal-to-noise ratio (SNR) values correspond with the sensor resolution values on the real robot and we used Eq. (28) to compute them:

$$\text{SNR} = \frac{\sigma_{\text{signal}}^2}{\sigma_{\text{noise}}^2}. \quad (28)$$

5.3. Implementation

We performed the simulations in MATLAB/Simulink with the “ode15s (stiff/NDF)” solver⁹⁶ of maximum order “5” using the “variable-step” solver type and the “full perturbation” Jacobian method as well as the “adaptive” zero-crossing option. The resultant simulation trials are of variable length due to the “variable-step” solver type. To make each simulation trial a vector of uniform length, first, we interpolated the data to a very high sampling rate of 100,000 Hz. Then, we applied a low-pass

Butterworth filter to the signal. Finally, we resampled the data to 100 Hz to match the frequency of the tactile-sensing forearm, and added Gaussian noise to match sensor noise.

We implemented continuous univariate and multivariate HMMs as well as LSTMs to model the temporal trends of features for different categories of objects. We modeled each of these four object categories: *hard-unmoved*, *hard-moved*, *soft-unmoved*, and *soft-moved*, using an HMM for each category (λ_{HU} , λ_{HM} , λ_{SU} , λ_{SM}).

We used the GHMM toolkit⁹⁷ to model the HMMs and implemented them in Python. We trained the models with the standard Baum–Welch algorithm, which uses expectation maximization. For testing, we ran the Viterbi algorithm which estimates the maximum likelihood with which a model can describe a given test data. These are standard methods for modeling sequential data (see Ref. 56 for details). We ran the Viterbi algorithm on the given test data for each of the trained HMM models and assigned the category associated with the model that returned the highest likelihood.

For LSTMs, we used the “Keras” library⁹⁸ with the “Theano” backend.⁹⁹ We trained our model for 10 epochs and used a batch-size of 5.

For comparison purposes, we implemented a 1-NN classifier using the Scikit-learn package¹⁰⁰ in Python. We used principal component analysis (PCA) to reduce the dimensionality of the concatenated features similar to Ref. 2. We used the Amazon EC2 cloud computing service¹⁰¹ to run the experiments on a $\text{c3.4} \times \text{large}$ system (high performance compute-optimized instances) with 30 GB of memory, 16 vCPUs and multiple $\text{c4.8} \times \text{large}$ systems with 60 GB of memory, 32 vCPUs. All of the systems were 64-bit Ubuntu 14.04 Linux platforms.

5.4. Algorithm parameters

For the univariate and multivariate HMMs, we analyzed the performance with 10 states. We set a uniform prior to all the states and initialized the emission matrix with Gaussian distributions with means and standard deviations obtained from the training data. For multivariate HMMs using multiple features (force, area, and/or motion), we used a spherical covariance matrix for initialization. Also, for multivariate HMMs, we scaled each feature (f) to a scaled feature (S_f) according to Eq. (29) to normalize the values:

$$S_f = (f - \text{mean}(f)) / \text{std}(f). \quad (29)$$

For LSTMs, we initialized the parameters with a uniform distribution, used “softsign” activation functions for the hidden layers, and “softmax” activation function⁹⁸ for the fully connected output layer. Our dropout probability was 0.2. We used “RMSprop”⁹⁸ as the optimizer and “categorical_crossentropy”⁹⁸ as our loss function because our task is a classification task. We used the “MinMaxScaler” function⁹⁸ to scale multivariate features for LSTMs.

To compare this with our previous work,² we used 1-NN. Before using 1-NN, we applied PCA representing more than 95% variance in the training data to reduce the dimensionality as described in our previous work.² We used PCA to reduce the effect of noise. As with the HMMs, we used Eq. (29) to scale the multivariate features for 1-NNs.

5.5. Results

As seen in Table 2, multivariate HMMs with contact force and motion as the features performed the best (82.13%) when compared to our previous method, 1-NN,² which failed to generalize (best performance was 64.58% with area and motion features) across different robot-arm stiffnesses and velocities. Univariate HMMs also failed to generalize. LSTMs showed the best performance with force, area, and motion features (80.46%).

Note that multivariate features with HMMs and LSTMs showed better performance compared to univariate features. Figure 5 shows the confusion matrix from multivariate HMMs with force and motion features. The algorithm classified “hard-unmoved” category well. But, there is some confusion between “soft-moved”,

Table 2. Summary of algorithm performance for simulations (ranked based on performance).

Algorithm	Features	Ranked accuracy (%)
MM	$f + m$	82.13
LSTM	$f + m + a$	78.02
LSTM	$m + a$	70.64
LSTM	m	68.64
LSTM	$f + m$	66.73
LSTM	$f + a$	64.68
1-NN	m	64.41
1-NN	$m + a$	62.72
HMM	$f + m + a$	59.34
LSTM	f	58.52
HMM	m	57.87
1-NN	a	57.6
HMM	$m + a$	48.01
1-NN	$f + m$	40.93
1-NN	f	40.71
1-NN	$f + m + a$	40.54
1-NN	$f + a$	40.44
HMM	a	39.4
HMM	f	39.07
HMM	$f + a$	36.89
LSTM	a	36.17
Majority classifier		29.65
Random guess		25.0

Note: “f” = force, “a” = contact area, and “m” = motion feature.

Fig. 5. Classification into four categories using multivariate HMMs with 10 states for experiments with physics-based models. The figure shows the result with force and motion features from the robot moving with varying stiffness and velocity. The numbers in the figure represent the number of trials.

“soft-unmoved” and “hard-moved” categories. Section 7 summarizes a possible reason for this.

6. Experiments with a Real Robot

We also performed a set of experiments with a real robot. Similar to our experiments with the physics-based models in Sec. 5, our experiments with the real robot varied arm stiffness and velocity.

6.1. Experimental setup

We used the robot, Cody, for our experiments. Cody, as shown in Fig. 1, is a statically stable (wheeled) mobile humanoid robot weighing approximately 160 kg. The components of the humanoid robot are: two Meka A1 arms, a Segway omni-directional base and a Festo 1-DOF (degree of freedom) linear actuator for a spine to adjust the torso height. The two seven-DOF anthropomorphic arms contain series elastic actuators. When we control these arms, each joint simulates a low-stiffness, visco-elastic, torsion spring. We control the robot’s arms by changing the equilibrium angles of these simulated springs.

Cody has a force-sensitive skin covering its forearm. Meka Robotics and the Georgia Tech Healthcare Robotics Lab developed the forearm tactile skin sensor, which is based on Stanford’s capacitive-sensing technology, as described by

Ulmen *et al.*¹⁰² The skin consists of a capacitive pressure-sensor array. We refer to the elements of this array as taxels (tactile pixels). There are 384 taxels on the entire skin, and these are distributed in a 24×16 array, with each taxel being $9 \text{ mm} \times 9 \text{ mm}$ in size. The array of taxels reports the measured force applied to each taxel at 100 Hz.

6.2. Experimental procedure

We conducted experiments to determine whether the performance of our algorithm could perform well on data collected with different robot arm stiffness and velocity than the training data. For the experiments, we selected two velocity settings, low = 5 deg/s and high = 20 deg/s, and two arm stiffness settings, low = 2.01 Nm/rad and high = 20.1 Nm/rad.

6.2.1. Data collection

The robot made contact with a set of objects on a wooden table while performing a simple, goal-directed reaching motion. We actuated the robot's shoulder joint only and it pushed into soft and hard objects in moved and unmoved conditions with varied arm stiffness (compliance) and velocities. We performed experiments with the eight objects shown in Fig. 8 (seven in both moved and unmoved conditions, one (heavy iron bucket) in the unmoved condition only). We selected large objects that have mostly uniform material properties and vary widely in their mass, friction, and stiffness. We actuated the robot's shoulder joint by commanding a goal angle in the joint space. The robot arm tried to reach the goal using a joint PD controller. We selected the final goal angle in joint space such that the equivalent point in the Cartesian space was inside the object. Thus, the robot would come in contact with the object before reaching the goal. When the robot incidentally came in contact with the object, it pushed against it and tried to reach the goal as shown in Fig. 6. For each object, we collected haptic data by commanding the same goal angle for the arm and recording the sensor readings from the taxels of the forearm skin at approximately 100 Hz.

Fig. 6. Sequence of images that illustrates the data collection for our experiments on inferring mechanical properties of objects (foliage). Each image shows a picture of the robot Cody and a visualization of the data from the forearm skin sensor as a 24×16 image (darker pixels correspond to larger forces). The leftmost picture shows a noncontact situation, the middle picture corresponds to the situation just after the onset of contact, and the rightmost picture shows the situation when the robot has pushed the foliage.

We labeled each of these objects as either soft or hard. For objects that could be pushed aside by the robot's motion, we fixed them with a clamp or a heavy weight so that we could obtain both moved and unmoved conditions. We repeated the experiments for four trials with each of the stiffness and velocity settings. We collected data for a total of 240 trials $((7 \text{ objects} \times 2 \text{ stiffness} \times 2 \text{ velocities} \times 2 \text{ conditions} \times 4 \text{ trials}) + (1 \text{ object} \times 2 \text{ stiffness} \times 2 \text{ velocities} \times 1 \text{ condition} \times 4 \text{ trials}))$. Each object category had 60 trials. Our experiments with the heavy iron bucket were only with the unmoved condition because it could not be moved by the robot's motion.

6.2.2. Preprocessing and feature selection

After recording the time-series data using the forearm taxel array, we truncated them to begin at the estimated onset of contact between the robot and object. We estimated the onset of contact whenever the force exceeded a threshold. We represented the data at every time-step as a grayscale image, as shown in Fig. 6. We converted this image to a binary image, representing the taxels in contact by applying a threshold to each taxel. We used two thresholds. One threshold (0.01 N) was for objects which were less likely to harm the robot-arm. For the other set of objects, we used a larger threshold (0.1 N) to account for the extra covering that we put over the skin sensor to prevent damage to the skin. This is equivalent to biasing the skin sensor. Then, we computed the connected components on the binary image to segment the contact regions. For the connected component with the largest area, we computed three features.

Figure 7 depicts the feature collection procedure. The first feature is the maximum force ($F_{\max}$) measured by a taxel in the contact region at each time-step. This is analogous to measuring the highest pressure. In our initial tests, the maximum force performed better than total force (F_{total}) or mean force (F_{mean}).

For the second feature, we estimated the area of contact (a) between the arm and object (contact region) as the number of taxels in the connected component.

Fig. 7. Schematic representation of the experimental protocol.

Fig. 8. Set of objects for experiments with varying stiffness and velocity (Left). Classification into four categories using multivariate HMMs for experiments using the robot “Cody”. The figure shows the results with force, area, and motion features from the robot moving with varying stiffness and velocity. The numbers in the figure represent the number of trials (Right).

For the third feature, we estimated the distance that the centroid of the connected component traveled in the world frame from its position at the onset of contact (d). We held the robot’s torso fixed throughout the trials and used the forward kinematics from the robot’s torso to the center of the contact location on the robot’s forearm to estimate the 3D positions and distance traveled.

Similar to the experiments with physics-based models, we expected these three features to be informative about the object’s softness and mobility. For example, with increasing force applied to a soft, unmoved object, we would expect the contact area to increase. Likewise, we would expect the 3D position of the contact area to travel when encountering moved and soft objects. When making contact with a hard and unmoved object, we would expect the maximum force to increase. Our algorithms used the values of maximum force, the number of taxels in the contact region, and the contact motion for each trial during the first 1.2s time window after the onset of contact.

6.3. Implementation

We implemented HMMs, LSTMs, and 1-NN similar to our experiments with physics-based models in Sec. 5.3. We performed experiments with the real robot using a 32-bit Ubuntu 10.04 system with 8 CPU(s) and Intel(R) Core(TM) i7-2600 CPU processor with 3.40 GHz. We used Python to send commands to the robot’s real-time PC using ROS-Diamondback.¹⁰³

6.4. Algorithm parameters

We used 10 states for the univariate and multivariate HMMs. We set a uniform prior for all states and initialized the emission matrix with Gaussian distributions with means and standard deviations obtained from the data.

For LSTMs, similar to experiments with physics-based models, we initialized the parameters with a uniform distribution, and used the same structure, scaling function, activation functions and dropout probability.

To compare this with our previous work,² we extracted the features and converted them to a low-dimensional representation of these feature vectors using PCA. We used a dimensionality of three which represented greater than 95% of the variance in the data.

Also, we used the same evaluation procedure as in the experiments with the physics-based model (Sec. 5.4). We trained on three conditions of robot stiffness and velocity and tested on one. We repeated this procedure for each of the four conditions. For multivariate HMMs, we scaled each feature (f) to a scaled feature (S_f) according to Eq. (29) to normalize the values.

6.5. Results

We used both univariate HMMs, multivariate HMMs, and LSTMs for classification to model the temporal trends of all combinations of the three-feature vectors: maximum force ($F_{\max}$), contact area (a) and contact motion (d). Table 3 presents the results for classification into four categories: (1) *hard-unmoved*, (2) *hard-moved*,

Table 3. Summary of algorithm performance for experiments (ranked based on performance).

Algorithm	Features	Ranked accuracy (%)
HMM	$f + m + a$	83.75
HMM	$f + m$	71.25
HMM	$m + a$	60.41
HMM	$f + a$	52.08
LSTM	$f + m$	43.33
LSTM	f	40.42
HMM	f	40.41
HMM	a	40.41
LSTM	$f + a$	40.0
1-NN	$f + a$	39.17
LSTM	a	39.17
1-NN	m	37.92
LSTM	$f + m + a$	35.83
1-NN	$f + m + a$	35.0
LSTM	$m + a$	34.58
HMM	m	32.5
LSTM	m	31.67
1-NN	$m + a$	28.33
1-NN	f	27.5
1-NN	a	26.25
1-NN	$f + m$	25.42
Majority classifier		25.0
Random guess		25.0

Note: “f” = force, “a” = contact area, and “m” = motion feature.

(3) *soft-unmoved*, and (4) *soft-moved*. Our previous method (from Ref. 2) performed poorly, and the highest accuracy was only 37.92% with a single feature (motion) and only 35% with three features (force, area, and motion) using a dimensionality of three. Note that 12 trials could not be captured up until the time window of 1.2 s because of the varying velocity conditions. In those cases, we extrapolated the data with the mean value for that particular trial to obtain a consistent time window of 1.2 s. With HMMs using a single feature (force or area), the accuracy improved only slightly to 40.41%. However, using multivariate HMMs, the accuracy improved to 83.75% with all three features. This provides evidence that multivariate HMMs can be used to generalize the data-driven inference results to testing data that differ from training data due to varying robot conditions. Figure 8 (Right) shows the resulting confusion matrix. From the figure, we see that some “soft-moved” and “soft-unmoved” trials were categorized as “hard-moved”. However, the algorithm categorized “hard-moved” and “hard-unmoved” categories well. Note, the results with LSTMs did not perform well. With access to more data, LSTMs might match the better results we obtained with physics-based models shown in Sec. 5.5.

7. Discussion and Limitations

Our data and models will be available publicly as a part of our “open-access-haptic-database (OAHB)”. From the overall results, we learned that:

- It is feasible to use data-driven methods to infer object properties from contact during a reaching motion;
- Classification results using HMMs and LSTMs (with sufficient data availability) with multiple features can generalize well to different robot behaviors, such as robot-arm stiffness and arm velocity used to collect training data.

We used the physics-based model to do experiments with objects with a wide variety of stiffnesses. Collecting haptic data from robots touching real-world objects is challenging. The physics-based model we used in this work can help collect data from a wide variety of objects and robot settings (varied robot stiffnesses, velocities, etc.) by leveraging the widely-available material properties in online databases. In addition to collecting the real-world data, the results from our physics-based model matched quite well with the results from our real-world data. For example, for both the experiments with physics-based models as well as the real robot, we note that there is confusion between “hard-moved”, “soft-unmoved” and “soft-moved” categories. This is probably because, given the features we have, the algorithm finds it difficult to disambiguate between sliding motion in hard objects and motion due to deformation of soft objects.

Note that for our experiments with the real robot using multivariate HMMs, there was very little confusion between hard and soft objects in the “unmoved” condition. However, for experiments with the physics-based models, there was some confusion

between hard objects and soft objects. This could be because objects in the hard category for the experiments with the real robot were much stiffer than objects in the soft category. We used a compression spring to compress the objects used in the real-robot experiment and found that the soft objects had stiffnesses ranging between 630 N/m and 1500 N/m whereas the hard objects had stiffnesses ranging between 8000 N/m and 100,000 N/m and above. This is not the case for experiments with the physics-based models. We experimented with a wide variety of stiffnesses and used a stiffness threshold of 100,000 N/m that differentiated hard and soft categories. Thus, the stiffnesses of the objects in the hard and soft categories were much closer (e.g., our simulation labeled an object with stiffness 100,100 N/m as hard but labeled an object with stiffness 99,900 N/m as soft) for experiments with the physics-based models and this led to more confusion.

For our experiments with physics-based models, we used stiffness as a criterion for labeling objects as hard or soft. In material sciences community, “hardness” is defined as a measure of how resistant an object is to deformation when a compressive force is applied.^{104,105} We are interested in forces only in the elastic deformation range of a material and for elastic deformation ranges, this highly correlates with stiffness of an object^{104,105} which is a function of the object material as well as structural properties. This is related to deformation of the object, as stiffness is monotonically related to deformation. Similarly, we labeled the objects for experiments with the real robot depending on how an object deforms in macroscopic scale due to applied forces. The human labeling of hard versus soft objects for experiments with the real robot was coarse. However, note that when consistent force was applied, all the objects labeled soft showed larger macroscopic deformation when compared to objects labeled hard.

Finally, we intentionally made the problem harder by not giving the perceptual classifiers information about the robot’s joint stiffness or arm velocity. This is because robots might not have good stiffness estimates for all contact locations on their bodies. Likewise, robots might have uncharacterized compliant coverings or components like soft robots. Similarly, accurate velocity estimates may not be attainable. Also, although the performance with LSTMs (with sufficient data) and multivariate HMMs generalized well to new robot behaviors used to collect training data, there are some limitations to the results presented here. In this work, our objective was to see if our algorithms can infer haptic properties using simple motions (e.g., linear) without haptic exploratory behaviors. Note that many motions are locally linear and therefore, this simple type of motion may be applicable in many scenarios. However, there are many factors which are relevant to real-world incidental contact like the impact dynamics, non-ideal contact (partial, non-normal) due to different robot trajectories and object shapes and sizes, etc. In this paper, we focused on two aspects — “stiffness” and “velocity” — which affect the way contact occurs. Other aspects of non-ideal contact merit consideration in the future.

8. Conclusion

We developed algorithms to infer object properties using haptic information obtained from contact between a robot's tactile-sensing forearm and objects in the robot's environment. We showed that using our algorithms and relevant tactile-sensing features, haptic inference can be generalized to data collected using different robot-arm velocities and stiffnesses. Our algorithms classified objects into four categories: (1) *hard-unmoved*, (2) *hard-moved*, (3) *soft-unmoved*, and (4) *soft-moved*.

We developed an idealized physics-based model and generated simulated data under varying robot stiffness and velocity. We also performed experiments with a real robot. We used univariate and multivariate HMMs and LSTM networks to classify objects under these conditions and compared the results with our previous method of PCA + 1-NN.² Our results showed that HMMs are a useful tool to model robot-object interactions. Multivariate HMMs consistently performed better in all cases with varying robot velocity and compliance parameter values and outperformed our previous technique using PCA + 1-NN² (see Fig. 2 and Table 3). With the availability of more data, the classification performance using LSTMs can also generalize to data collected using different robot actions. Also, for HMMs and LSTMs, classification results using a combination of relevant features such as force, area, and motion generalize better than using single-features.

Acknowledgments

We gratefully acknowledge the support from DARPA's Maximum Mobility and Manipulation (M3) Program, Contract W911NF-11-1-603, NSF Emerging Frontiers in Research and Innovation (EFRI) 1137229, and NSF Career Award 1150157. We thank Joshua Wade for his help with selecting the materials for our simulations, as well as Ariel Kapusta and Zackory Erickson for their valuable feedback. We also thank Mark Cutkosky and the Stanford Biomimetics and Dexterous Manipulation Lab for their contributions to the forearm tactile skin sensor.

References

1. A. Jain, M. D. Killpack, A. Edsinger and C. Kemp, Reaching in clutter with whole-arm tactile sensing, *Int. J. Robot. Res.* **32**(4) (2013) 458–482.
2. T. Bhattacharjee, J. M. Rehg and C. C. Kemp, Haptic classification and recognition of objects using a tactile sensing forearm, in *IEEE Int. Conf. Intelligent Robots and Systems (IROS)* (October 2012), pp. 4090–4097.
3. T. Bhattacharjee, A. Kapusta, J. M. Rehg and C. C. Kemp, Rapid categorization of object properties from incidental contact with a tactile sensing robot arm, in *IEEE-RAS Int. Conf. Humanoid Robots (Humanoids)*, October 2013, pp. 219–226.
4. T. Bhattacharjee, P. M. Grice, A. Kapusta, M. D. Killpack, D. Park and C. C. Kemp, A robotic system for reaching in dense clutter that integrates model predictive control, learning, haptic mapping, and planning, in *Proc. 3rd IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS) Workshop on Robots in Clutter: Perception and Interaction in Clutter* (2014), pp. 1–7.

5. T. Bhattacharjee, A. Jain, S. Vaish, M. D. Killpack and C. C. Kemp, Tactile sensing over articulated joints with stretchable sensors, in *IEEE World Haptics Conf. (WHC)* (2013), pp. 103–108.
6. P. M. Grice, M. D. Killpack, A. Jain, S. Vaish, J. Hawke and C. C. Kemp, Whole-arm tactile sensing for beneficial and acceptable contact during robotic assistance, in *13th Int. Conf. Rehabilitation Robotics (ICORR)* (2013), pp. 1–8.
7. S. J. Lederman, R. L. Klatzky, C. L. Hamilton and G. I. Ramsay, Perceiving surface roughness via a rigid probe: Effects of exploration speed and mode of touch, *Haptics-e Electron. J. Haptics Res.* **1**(1) (1999) 1–20.
8. S. J. Lederman, R. L. Klatzky, C. L. Hamilton and M. Grindley, Perceiving surface roughness through a probe: Effects of applied force and probe diameter, in *Proc. ASME DSCD-IMECE*, Vol. 69 (2000), pp. 1065–1071.
9. S. J. Lederman and R. L. Klatzky, Extracting object properties through haptic exploration, *Acta Psychol.* **84**(1) (1993) 29–40.
10. D. Silvera-Tawil, D. Rye and M. Velonaki, Artificial skin and tactile sensing for socially interactive robots: A review, *Robot. Auton. Syst.* **63** (2015) 230–243.
11. H. Wu, H. Liu and D. Liu, Two-dimensional direction recognition using uniaxial tactile arrays, *IEEE Sens. J.* **13**(12) (2013) 4897–4903.
12. L. Muscare, L. Seminara, F. Mastrogiovanni, M. Valle, M. Capurro and G. Cannata, Real-time reconstruction of contact shapes for large area robot skin, in *2013 IEEE Int. Conf. Robotics and Automation (ICRA)* (2013), pp. 2360–2366.
13. D. Hughes and N. Correll, Texture recognition and localization in amorphous robotic skin, *Bioinsp. Biomim.* **10**(5) (2015) 055002.
14. M. Javaid, M. Zefran and A. Yavolovsky, Using pressure sensors to identify manipulation actions during human physical interaction, in *2015 24th IEEE Int. Symp. Robot and Human Interactive Communication (RO-MAN)* (2015), pp. 670–675.
15. K. Matheus and A. M. Dollar, Benchmarking grasping and manipulation: Properties of the objects of daily living, in *2010 IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)* (2010), pp. 5020–5027.
16. P. Boonvisut and M. C. Çavuşoğlu, Identification and active exploration of deformable object boundary constraints through robotic manipulation, *Int. J. Robot. Res.* **33**(11) (2014) 1446–1461.
17. B. Heyneman and M. R. Cutkosky, Slip classification for dynamic tactile array sensors, *Int. J. Robot. Res.* **35**(4) (2015) 404–421.
18. C. Schürmann, M. Schöpfer, R. Haschke and H. Ritter, A high-speed tactile sensor for slip detection, in *Towards Service Robots for Everyday Environments* (Springer, 2012), pp. 403–415.
19. A.-V. Ho and S. Hirai, Slip perception via soft robotic skin made of electroconductive yarn, in *Mechanics of Localized Slippage in Tactile Sensing* (Springer, 2014), pp. 113–154.
20. M. Kaboli, R. Walker, G. Cheng *et al.*, In-hand object recognition via texture properties with robotic hands, artificial skin, and novel tactile descriptors, in *2015 IEEE-RAS 15th Int. Conf. Humanoid Robots (Humanoids)* (2015), pp. 1155–1160.
21. V. Sukhoy, R. Sahai, J. Sinapov and A. Stoytchev, Vibrotactile recognition of surface textures by a humanoid robot, in *IEEE Int. Conf. Robotics and Automation*, Vol. 1 (1989), pp. 145–150.
22. R. D. Howe and M. R. Cutkosky, Dynamic tactile sensing: Perception of fine surface features with stress rate sensing, *IEEE Trans. Robot. Autom.* **9**(2) (1993) 140–151.
23. Q. Li, C. Schürmann, R. Haschke and H. Ritter, A control framework for tactile servoing, in *Proc. of Robotics: Science and Systems* (Berlin, Germany, 2013), Doi: 10.15607/RSS.2013.IX.045.

24. U. Martinez-Hernandez, G. Metta, T. J. Dodd, T. J. Prescott, L. Natale and N. F. Lepora, Active contour following to explore object shape with robot touch, in *World Haptics Conf. (WHC)* (2013), pp. 341–346.
25. D. J. Agravante, A. Cherubini, A. Bussy, P. Gergondet and A. Kheddar, Collaborative human–humanoid carrying using vision and haptic sensing, in *2014 IEEE Int. Conf. Robotics and Automation (ICRA)* (2014), pp. 607–612.
26. S. Calinon, P. Evrard, E. Gribovskaya, A. Billard and A. Kheddar, Learning collaborative manipulation tasks by demonstration using a haptic interface, in *Int. Conf. Advanced Robotics (ICAR)* (2009), pp. 1–6.
27. A. Drimus, G. Kootstra, A. Bilberg and D. Kragic, Classification of rigid and deformable objects using a novel tactile sensor, in *Proc. 15th Int. Conf. Advanced Robotics (ICAR)* (2011), pp. 427–434.
28. V. Chu, I. McMahon, L. Riano, C. G. McDonald, Q. He, J. M. Perez-Tejada, M. Arrigo, T. Darrell and K. J. Kuchenbecker, Robotic learning of haptic adjectives through physical interaction, *Robot. Auton. Syst.* **63** (2015) 279–292.
29. C.-H. Lin, T. W. Erickson, J. A. Fishel, N. Wettels and G. E. Loeb, Signal processing and fabrication of a biomimetic tactile sensor array with thermal, force and micro-vibration modalities, in *IEEE Int. Conf. Robotics and Biomimetics (ROBIO)* (IEEE, 2009), pp. 129–134.
30. M. Kaboli, P. Mittendorfer, V. Hügel and G. Cheng, Humanoids learn object properties from robust tactile feature descriptors via multi-modal artificial skin, in *2014 14th IEEE–RAS Int. Conf. Humanoid Robots (Humanoids)* (2014), pp. 187–192.
31. Y. S. Kim and T. Kesavadas, Material property recognition by active tapping for fingertip digitizing, in *Proc. Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems* (March 2006), pp. 133–139.
32. S. Takamuku, G. Gomez, K. Hosoda and R. Pfeifer, Haptic discrimination of material properties by a robotic hand, in *Proc. IEEE 6th Int. Conf. Development and Learning (ICDL)* (2007), pp. 1–6.
33. K. Hosoda and T. Iwase, Robust haptic recognition by anthropomorphic bionic hand through dynamic interaction, in *Proc. Int. Conf. Intelligent Robots and Systems (IROS)* (2010), pp. 1236–1241.
34. N. N. A. Charniya and S. V. Dudul, Sensor for classification of material type and its surface properties using radial basis networks, *IEEE Sens. J.* **8** (2008) 1981–1991.
35. M. V. Liarokapis, B. Calli, A. J. Spiers and A. M. Dollar, Unplanned, model-free, single grasp object classification with underactuated hands and force sensors, in *2015 IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)* (2015), pp. 5073–5080.
36. A. Schmitz, Y. Bansho, K. Noda, H. Iwata, T. Ogata and S. Sugano, Tactile object recognition using deep learning and dropout, in *2014 14th IEEE–RAS Int. Conf. Humanoid Robots (Humanoids)* (IEEE, 2014), pp. 1044–1050.
37. N. Kiwatthana and S. Kaitwanidvilai, Development of smart gripper for identification of grasped objects, in *Asia-Pacific Signal and Information Processing Association, 2014 Annual Summit and Conference (APSIPA)* (IEEE, 2014), pp. 1–5.
38. J. Hoelscher, J. Peters and T. Hermans, Evaluation of tactile feature extraction for interactive object recognition, in *2015 IEEE–RAS 15th Int. Conf. Humanoid Robots (Humanoids)* (IEEE, 2015), pp. 310–317.
39. A. Schneider, J. Sturm and C. Stachniss, Object identification with tactile sensors using bag-of-features, in *Proc. Int. Conf. Intelligent Robots and Systems (IROS)* (2009), pp. 243–248.
40. Z. Pezzementi, E. Plaku, C. Reyda and G. D. Hager, Tactile object recognition from appearance information, *IEEE Trans. Robot.* **27** (2011) 473–486.

41. N. Gorges, S. E. Navarro, D. Goger and H. Worn, Haptic object recognition using passive joints and haptic key features, in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)* (May 2010), pp. 2349–2355.
42. D. Babu, S. Bhattacharjee, I. Bandyopadhyaya and J. Roychowdhury, Machine learning-based shape classification using tactile sensor array, in *Advanced Computing, Networking and Informatics*, Vol. 1, Smart Innovation, Systems and Technologies, Vol. 27 (Springer, 2014), p. 47.
43. N. Ueda, S. Hirai and H. T. Tanaka, Extracting rheological properties of deformable objects with haptic vision, in *Proc. 2004 IEEE Int. Conf. Robotics and Automation* (April 2004), pp. 3902–3907.
44. M. Bjorkman, Y. Bekiroglu, V. Hogman and D. Kragic, Enhancing visual perception of shape through tactile glances, in *2013 IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)* (IEEE, 2013), pp. 3180–3186.
45. B. Frank, R. Schmedding, C. Stachniss, M. Teschner and W. Burgard, Learning the elasticity parameters of deformable objects with a manipulation robot, in *Proc. Int. Conf. Intelligent Robots and Systems (IROS)* (2010), pp. 1877–1883.
46. B. Frank, R. Schmedding, C. Stachniss, M. Teschner and W. Burgard, Learning deformable object models for mobile robot navigation using depth cameras and a manipulation robot, in *Proc. Workshop on Advanced Reasoning with Depth Cameras at Robotics, Science and Systems Conference (RSS)* (2010), pp. 1–6.
47. P. K. Allen and K. S. Roberts, Haptic object recognition using a multi-fingered dexterous hand, in *Proc. Int. Conf. Robotics and Automation (ICRA)* (1989), pp. 342–347.
48. S. Caselli, C. Magnanini and F. Zanichelli, Haptic object recognition with a dexterous hand based on volumetric shape representations, in *Proc. 1994 Int. Conf. Multisensor Fusion and Integration for Intelligent Systems* (1994), pp. 280–287.
49. E. Faldella, B. Fringuelli, D. Passeri and L. Rosi, A neural approach to robotic haptic recognition of 3D objects based on Kohonen self-organizing feature map, *IEEE Trans. Ind. Electron.* **44** (1997) 267–269.
50. J. Sinapov, M. Weimer and A. Stoytchev, Interactive learning of the acoustic properties of household objects, in *Proc. 2009 IEEE Int. Conf. Robotics and Automation (ICRA)* (2009), pp. 2518–2524.
51. T. Berquist, C. Schenck, U. Ohiri, J. Sinapov, S. Griffith and A. Stoytchev, Interactive object recognition using proprioceptive feedback, in *Proc. 2009 IROS Workshop on Semantic Perception for Mobile Manipulation (IROS)* (2009), pp. 1–7.
52. A. Jain and C. C. Kemp, Improving robot manipulation with data-driven object-centric models of everyday forces, *Auton. Robots* **35**(2–3) (2013) 143–159.
53. S. Griffith, J. Sinapov, V. Sukhoy and A. Stoytchev, A behavior grounded approach to forming object categories: Separating containers from non-containers, *IEEE Trans. Auton. Mental Develop.* **4**(1) (2012) 54–69.
54. J. Sinapov, T. Bergquist, C. Schenck, U. Ohiri, S. Griffith and A. Stoytchev, Interactive object recognition using proprioceptive and auditory feedback, *Int. J. Robot. Res.* **30** (2011) 1250–1262.
55. J. Sinapov and A. Stoytchev, Grounded object individuation by a humanoid robot, in *2013 IEEE Int. Conf. Robotics and Automation (ICRA)* (2013), pp. 4981–4988.
56. L. R. Rabiner, A tutorial on hidden Markov models and selected applications in speech recognition, in *Readings in Speech Recognition*, eds. A. Waibel and K. F. Lee (Morgan Kaufmann, 1990), pp. 267–296.
57. I. Sutskever, O. Vinyals and Q. V. Le, Sequence-to-sequence learning with neural networks, in *Advances in Neural Information Processing Systems* (Curran Associates, 2014), pp. 3104–3112.

58. A. Graves, Generating sequences with recurrent neural networks, preprint (2013), arXiv: 1308.0850.
59. A. Graves, A.-R. Mohamed and G. Hinton, Speech recognition with deep recurrent neural networks, in *2013 IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)* (IEEE, 2013), pp. 6645–6649.
60. S. Hochreiter and J. Schmidhuber, Long short-term memory, *Neural Comput.* **9**(8) (1997) 1735–1780.
61. Z. Erickson, A. Clegg, W. Yu, G. Turk, C. K. Liu and C. C. Kemp, What does the person feel? Learning to infer applied forces during robot-assisted dressing, in *2017 IEEE Int. Conf. Robotics and Automation (ICRA)* (2017).
62. A. I. Aviles, S. Alsaleh, P. Sobrevilla and A. Casals, Exploring the effects of dimensionality reduction in deep networks for force estimation in robotic-assisted surgery, in *SPIE Medical Imaging* (International Society for Optics and Photonics, 2016), Vol. 9786, pp. 1–8.
63. Y. Gao, L. A. Hendricks, K. J. Kuchenbecker and T. Darrell, Deep learning for tactile understanding from visual and haptic data, in *2016 IEEE Int. Conf. Robotics and Automation (ICRA)* (2016), pp. 536–543.
64. F. A. Gers, J. Schmidhuber and F. Cummins, Learning to forget: Continual prediction with LSTM, *Neural Comput.* **12**(10) (2000) 2451–2471.
65. V. Pham, T. Bluche, C. Kermorvant and J. Louradour, Dropout improves recurrent neural networks for handwriting recognition, in *2014 14th Int. Conf. Frontiers in Handwriting Recognition (ICFHR)* (IEEE, 2014), pp. 285–290.
66. R. Shadmehr, The equilibrium point hypothesis for control of movement (Johns Hopkins University, Baltimore, MD, 1998).
67. H. M. Westergaard, *Theory of Elasticity and Plasticity* (Harvard University Press, Cambridge, 1952).
68. Poisson's ratio, https://en.wikipedia.org/wiki/Poisson%27s_ratio, retrieved on 2 March 2017.
69. Coefficients of Friction, http://www.roymech.co.uk/Useful_Tables/Tribology/co_of_frict.htm, retrieved on 17 November 2015.
70. Friction and Coefficients of Friction, http://www.engineeringtoolbox.com/friction-coefficients-d_778.html, retrieved on 17 November 2015.
71. Coefficient of Friction, Rolling Resistance and Aerodynamics, <http://www.tribology-abc.com/abc/cof.htm>, retrieved on 17 November 2015.
72. Coefficient of Friction Reference Table, <http://www.engineershandbook.com/Tables/frictioncoefficients.htm>, retrieved on 17 November 2015.
73. I. Grigoriev and E. Meilikhov, *Handbook of Physical Quantities* (CRC Press, Boca Raton, FL, 1997), p. 1548.
74. J. W. Jewett and R. A. Serway, *Physics for Scientists and Engineers with Modern Physics* (Cengage Learning EMEA, 2008).
75. A. I. H. Committee, *Friction, Lubrication, and Wear Technology* (ASM International, 1992).
76. P. J. Blau, *Friction Science and Technology: From Concepts to Applications* (CRC Press, 2008).
77. P. Blau, Appendix: Static and kinetic friction coefficients for selected materials, in *ASM Handbook*, Vol. 18, 4th edn. (ASM International, Materials Park, OH, 2002).
78. P. Hasgall, F. Di Gennaro, C. Baumgartner, E. Neufeld, M. Gosselin, D. Payne, A. Klingenberg and N. Kuster, ITIS database for thermal and electromagnetic parameters of biological tissues, <http://www.itis.ethz.ch/database>, retrieved on 17 November 2015.

79. M. Farvid, T. Ng, D. Chan, P. Barrett and G.F. Watts, Association of adiponectin and resistin with adipose tissue compartments, insulin resistance and dyslipidaemia, *Diabetes Obes. Metab.* **7**(4) (2005) 406–413.
80. J. Mendez and A. Keys, Density and composition of mammalian muscle, *Metab. Clin. Exp.* **9**(2) (1960) 184–188.
81. S. Derler and L.-C. Gerhardt, Tribology of skin: Review and analysis of experimental results for the friction coefficient of human skin, *Tribol. Lett.* **45**(1) (2012) 1–27.
82. N. Veijgen, Skin friction — A novel approach to measuring *in vivo* human skin, Ph.D. thesis (2013).
83. M. F. Ashby, The CES EduPack database of natural and man-made materials (2008), url: www.grantadesign.com/education/edupack.
84. Wood Densities, http://www.engineeringtoolbox.com/wood-density-d_40.html, retrieved on 17 November 2015.
85. Seats and Cushions, <http://www.usafoam.com/seat&cushion/seat&cushion.html>, retrieved on 17 November 2015.
86. Glass.Fibre.Reinforced.Products, <http://www.amiantit.com/media/pdf/brochures/Glass.Fibre.Reinforced.Products/files/Glass.Fibre.Reinforced.Products.pdf>, retrieved on 17 November 2015.
87. W. Witkiewicz and A. Zieliński, Properties of the polyurethane (PU) light foams, *Adv. Mater. Sci.* **6**(2) (2006) 35–51.
88. CES EduPack database of natural and man-made materials (2016), Granta Design, url: www.grantadesign.com/education/edupack.
89. J. T. Iivarinen, R. K. Korhonen, P. Julkunen and J. S. Jurvelin, Experimental and computational analysis of soft tissue stiffness in forearm using a manual indentation device, *Med. Engrg. Phys.* **33**(10) (2011) 1245–1253.
90. B. C. W. Kot, Z. J. Zhang, A. W. C. Lee, V. Y. F. Leung and S. N. Fu, Elastic modulus of muscle and tendon with shear wave ultrasound elastography: Variations with different technical settings, *PLoS One* **7**(8) (2012) e44348.
91. M. Birch and P. Srodon, Biomechanical properties of the human soft palate, *Cleft Palate Craniofa. J.* **46**(3) (2009) 268–274.
92. B. A. Todd and J. G. Thacker, Three-dimensional computer model of the human buttocks, *in vivo*, *J. Rehabil. Res. Develop.* **31**(2) (1994) 111.
93. L. Daneshmend, V. Hayward and M. Pelletier, Adaptation to environment stiffness in the control of manipulators, in *Experimental Robotics I* (Springer, 1990), pp. 150–165.
94. T. L. Chen, T. Bhattacharjee, J. L. McKay, J. E. Borinski, M. E. Hackney, L. H. Ting and C. C. Kemp, Evaluation by expert dancers of a robot that performs partnered stepping via haptic interaction, *PLoS One* **10**(5) (2015) e0125179.
95. S. S. Rao, *Vibration of Continuous Systems* (John Wiley and Sons, 2007).
96. Choose a Solver — MATLAB and Simulink, <http://www.mathworks.com/help/simulink/ug/types-of-solvers.html>, retrieved on 17 November 2015.
97. General Hidden Markov Model Library, <http://ghmm.org/>, retrieved on 12 January 2012.
98. F. Chollet, Keras (2015), <https://github.com/fchollet/keras>, retrieved on 15 February 2017.
99. Theano Development Team, Theano: A Python framework for fast computation of mathematical expressions, arXiv: 1605.02688.
100. F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot and E. Duchesnay, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* **12** (2011) 2825–2830.

101. Elastic Compute Cloud (EC2) — Cloud Server and Hosting, <https://aws.amazon.com/ec2/>.
102. J. Ulmen and M. Cutkosky, A robust, low-cost and low-noise artificial skin for human-friendly robots, in *2010 IEEE Int. Conf. Robotics and Automation (ICRA)* (2010), pp. 4836–4841.
103. ROS Diamondback, <http://wiki.ros.org/diamondback>.
104. D. Tabor, *The Hardness of Metals* (Oxford University Press, 2000).
105. W. C. Oliver and G. M. Pharr, An improved technique for determining hardness and elastic modulus using load and displacement sensing indentation experiments, *J. Mater. Res.* **7**(6) (1992) 1564–1583.

Tapomayukh Bhattacharjee received the B.Tech. degree from the Department of Mechanical Engineering, National Institute of Technology, Calicut, India, and the M.Sc. degree from the Department of Mechanical Engineering, Korea Advanced Institute of Science and Technology, Daejeon, Korea. He is currently a Robotics Ph.D. Student in the Georgia Institute of Technology, Atlanta. He also worked as a Visiting Scientist with the Interaction and Robotics Research Center, Korea Institute of Science and Technology, Seoul, Korea. His research interests include haptic perception, control systems for robotic manipulation, human–robot interaction, machine learning, and teleoperation systems. He is a member of IEEE, IEEE RAS, and IEEE CSS.

James M. Rehg is a Professor in the School of Interactive Computing at the Georgia Institute of Technology, where he is Co-Director of the Computational Perception Lab and is the Associate Director for Research in the Center for Robotics and Intelligent Machines (RIM@GT). He received his Ph.D. from CMU in 1995 and worked at the Cambridge Research Lab of DEC (and then Compaq) from 1995–2001, where he managed the computer vision research group. He received an NSF CAREER

Award in 2001 and a Raytheon Faculty Fellowship from Georgia Tech in 2005. His research interests include computer vision, medical imaging, robot perception, machine learning, and pattern recognition. Dr. Rehg is currently leading a multi-institution effort to develop the science and technology of behavior imaging — the capture and analysis of social and communicative behavior using multi-modal sensing, to support the study and treatment of developmental disorders such as autism.

Charles C. Kemp (Charlie) is an Associate Professor at the Georgia Institute of Technology in the Department of Biomedical Engineering with adjunct appointments in the School of Interactive Computing and the School of Electrical and Computer Engineering. He earned a doctorate in Electrical Engineering and Computer Science (2005), an M.Engg., and B.Sc. from MIT. In 2007, he founded the Healthcare Robotics Lab (<http://healthcare-robotics.com>). His lab focuses on mobile robots for intelligent physical assistance in the context of healthcare. He has received a 3M Non-Tenured Faculty Award, the Georgia Tech Research Corporation Robotics Award, a Google Faculty Research Award, and an NSF CAREER Award. He was a Hesburgh Award Teaching Fellow in 2017. His research has been covered extensively by the popular media, including the New York Times, Technology Review, ABC and CNN.