

Improving robot manipulation with data-driven object-centric models of everyday forces

Advait Jain · Charles C. Kemp

Received: 30 March 2012 / Accepted: 4 June 2013 / Published online: 19 June 2013
© Springer Science+Business Media New York 2013

Abstract Based on a lifetime of experience, people anticipate the forces associated with performing a manipulation task. In contrast, most robots lack common sense about the forces involved in everyday manipulation tasks. In this paper, we present data-driven methods to inform robots about the forces that they are likely to encounter when performing specific tasks. In the context of door opening, we demonstrate that data-driven object-centric models can be used to haptically recognize specific doors, haptically recognize classes of door (e.g., refrigerator vs. kitchen cabinet), and haptically detect anomalous forces while opening a door, even when opening a specific door for the first time. We also demonstrate that two distinct robots can use forces captured from people opening doors to better detect anomalous forces. These results illustrate the potential for robots to use shared databases of forces to better manipulate the world and attain common sense about everyday forces.

Keywords Mobile manipulation · Haptic sensing · Service robotics

1 Introduction

Little is known about the statistics of real-world forces associated with everyday tasks in human environments. While vast quantities of everyday auditory and visual data are publicly available for use by humans and machines, publicly available haptic data is much less common. Capturing and modeling

the forces associated with everyday tasks could benefit robots by enabling them to better interact with the physical world.

For example, despite progress towards service robots that autonomously open doors and drawers, the answers to basic questions have been unclear, such as, “How hard does a robot need to pull to open most doors?”. Given the wide variation in the forces required to initially open a door (e.g., ~60 N for a spring-loaded door, and <5 N for a kitchen cabinet), a robot without common sense about everyday forces risks damaging a locked cabinet or giving up prematurely on a functioning spring-loaded door. Likewise, while opening a door for someone, distinguishing forces due to the door opening properly versus the door being in contact with the person could improve safety.

Within this paper, we present data-driven methods to inform robots about the forces that they are likely to encounter when performing specific tasks. We focus on the example task of pulling open a door. As Sect. 2.4 discusses, door-opening robots have lacked compelling ways to deal with many common situations, such as a door that is locked, blocked, or damaged. In this paper, we provide evidence that models of real-world forces can be used by robots to better handle these situations.

The models we present have the following three important characteristics:

1. *Task-specific* Each model is specific to a narrowly defined manipulation task to reduce the complexity of the model and the data requirements. For this paper, we defined the task to be smoothly and slowly pulling open a door with contact restricted to the handle. For our data, the average linear velocity of the door handle while pulling open a door in a single trial was between 4.6 and 79.4 cm/s. The mean of this average velocity across all trials with all doors was 18.2 cm/s and the median was 14.4 cm/s.

A. Jain (✉) · C. C. Kemp
Healthcare Robotics Lab, Georgia Tech, Atlanta, GA, USA
e-mail: advaitjain@gmail.com

C. C. Kemp
e-mail: charlie.kemp@bme.gatech.edu

2. *Data-driven* The models directly use forces, points of application of the forces, and kinematics captured during real-world performance of the task. With a data-driven approach, we intend to capture the natural variation that a robot will encounter. For this paper, we used data captured while opening 26 doors in 6 homes and one office in Atlanta, GA, USA.
3. *Object-centric* Each model relates the relevant state of the manipulated object to the relevant forces applied to the object to make the models independent of the robot or human manipulating the object. For this paper, the models are quasi-static. They relate the opening angle of the door to the component of the applied force that is tangential to the trajectory of the point of contact on the door handle. Our choice of object-centric representations is intended to make the models useful for distinct robots and methods of manipulation. For example, it should not matter how the robot applies the forces, whether with its left hand, its right hand, or some other part of its body.

1.1 Organization of the paper

Section 2 discusses related work and contrasts it with our approach. The next three sections present our contributions to three areas. First, Sect. 3 describes our methods of collecting force data from robots and humans, a quasi-static model for doors, and the object-centric representation that we use.

Second, Sect. 4 shows that a standard supervised learning classifier can recognize the class of a door (e.g., refrigerator or kitchen cabinet) and the specific instance of a door based on the forces during opening. Third, Sect. 5 presents a method for detecting anomalous forces using a probabilistic model of the forces captured while successfully opening different doors.

Section 6 presents a comparison of our anomalous force detection method with two baseline methods. It shows that using previously captured haptic data enables the detection of locked doors and contact between the door and an idealized rigid obstacle more quickly and with lower applied forces. Section 7 reports results of anomalous force detection from trials on two distinct robots and discusses the implications of online state estimation on the performance of our anomalous force detection method. Finally, Sect. 8 discusses limitations of our work and directions for future research, after which the paper concludes with Sect. 9.

2 Related work

2.1 Capturing haptic data

Although vast quantities of everyday visual and auditory data can be easily accessed on the web, very little haptic data from

everyday activities can be found. Researchers have looked at capturing haptic interactions in order to synthesize realistic haptic sensations for human users, for example (Pai et al. 2000; Dupont et al. 1999; MacLean 1996; Angerilli et al. 2001; Weir et al. 2004; Romano and Kuchenbecker 2011). Recent work has modeled the forces while opening a refrigerator door for realistic haptic feedback in a virtual environment (Shin et al. 2012). To date, however, this body of work has emphasized high-fidelity models of objects to convey realistic haptic sensations to people, rather than capturing haptic datasets to inform robots.

Recently, researchers have begun to collect datasets of physical properties of objects (Matheus and Dollar 2010) and of forces during some everyday activities, such as tooth brushing (Redmond et al. 2010), with the motivation of informing the design of robot hardware and software. In previous research, we captured the forces applied to door handles and the trajectories of the handles while pulling open doors with the same motivation (Jain et al. 2010). However, these previous works have not used the datasets to inform robot manipulation.

2.2 Using forces to inform robot manipulation

There is a large body of research on using realtime force and tactile signals to improve manipulation. Recent research in this domain includes work by Romano et al. (2011), Hsiao et al. (2010), Dollar et al. (2010), Prats et al. (2009), Platt Jr et al. (2011), Chitta et al. (2011). Similar in spirit to our approach of using previously captured forces, Chan et al. (2012) have characterized the forces involved in human-to-human object transfer and then presented guidelines for robot controllers for the same task. Wiste et al. (2011) have used force data from everyday activities from the literature on rehabilitation research as a guideline for the design of prosthetic hands.

2.3 Haptic recognition and anomaly detection

In contrast to our use of data-driven object-centric models, previous research on haptic recognition and anomaly detection for robot manipulation has often used data-driven models that are robot-centric or models of the dynamics of the robot arm.

2.3.1 Data-driven and robot-centric models

Researchers have used data-driven robot-centric models of haptic data in the form of tactile sensor arrays, joint torques, and joint angles to haptically recognize objects grasped by a robot hand (e.g., Navarro et al. 2012; Gorges et al. 2010; Johnsson and Balkenius 2007; Takamuku et al. 2008). Sinapov et al. (2011) demonstrated that haptic data in the

form of joint torques associated with specific behaviors for the entire arm could be used to recognize objects.

Other research has used robot-centric models to detect anomalous conditions during a manipulation task. For example, [Rodriguez et al. \(2010\)](#) have demonstrated that force data captured during an assembly operation can be used to predict failure in future trials. [Pastor et al. \(2011\)](#) have shown that a database of joint angles, joint torques, tactile sensor information, and accelerometer data can be used to predict failure as a robot plays-back a learned trajectory for two tasks, (1) flipping a box using chopsticks, and (2) hitting a pool ball with a cue stick. [Sukhoy et al. \(2012\)](#) used deviations from a data-driven model of joint torques during free-space swiping motions to detect conditions when a magnetic card gets stuck as a robot swiped it through a card reader.

For these methods, the state of the robot is intertwined with the haptic representations. For example, these methods often use ego-centric sensor data parameterized by time or the robot's state, such as joint angles. As a result, there is no direct way for different robots to share data-driven models that use robot-centric representations. Our approach includes capturing data in a way that permits transformation to an object-centric representation. This transformation can enable distinct robots and humans to share haptic data and data-driven object-centric models.

[Schneider et al. \(2009\)](#) have presented methods for object identification with bag-of-features models using haptic data in the form of readings from tactile sensor arrays on the robot's parallel jaw gripper, and the width and height at which the robot grasps the object. Although [Schneider et al. \(2009\)](#) presented results from data collected by a single robot, these models are data-driven and object-centric, and different robots with similar sensing capabilities may be able to share these haptic data.

2.3.2 Anomaly detection using joint torque sensors and arm dynamics

Previous research has used deviations from an expected torque, predicted using a model of the dynamics of the robot arm, to detect anomalous conditions (e.g., [Dixon et al. 2000](#); [De Luca and Mattone 2004](#); [Haddadin et al. 2008, 2011](#)).

Determining an accurate model of the arm dynamics can be challenging. Additionally, these approaches often detect anomalous conditions in free-space motions. Estimating expected torques in situations where the robot makes contact with its environment is more complex (e.g., [Morinaga and Kosuge 2003](#)). For example, while pulling open doors, forces along the radial direction (and hence the torques at the joints) can change without triggering an anomalous condition, as discussed in Sect. 3.1. Additionally, the expected torques will also depend on the characteristics of the object that the robot is manipulating. As a result, modeling and

representing acceptable torques would have to be more complex compared to free-space motions.

2.4 Robotic door opening

Recently, researchers have developed a number of robotic systems to operate doors between rooms (e.g., [Klingbeil et al. 2008](#); [Jain and Kemp 2009a](#); [Meeussen et al. 2010](#); [Chitta et al. 2010](#); [Kalakrishnan et al. 2011](#); [Kormushev et al. 2011](#)), and open cabinets, drawers, and appliances (e.g., [Jain and Kemp 2010](#); [Wieland et al. 2009](#); [Diankov et al. 2008](#); [Rühr et al. 2012](#); [Becker et al. 2011](#)). These efforts have focused on controllers and planners that enable a robot to open doors. They have not addressed haptically detecting contact between the door and the environment (e.g., collisions), or haptically recognizing the door identity or class.

2.5 Prior research by the authors

We have previously developed methods that enable a robot to autonomously open doors and drawers without prior knowledge of the mechanism kinematics ([Jain and Kemp 2009b, 2010](#)). We use these methods in this paper to enable the robots Cody and a Willow Garage PR2 (Personal Robot 2) to open doors Fig. (1).

In collaboration with Sturm, Stachniss, and Burgard, we have shown that a robot can use a database of kinematic trajectories of mechanism handles to increase the online prediction accuracy of the kinematic state of a mechanism that it is currently opening ([Sturm et al. 2010](#)). We do not use our research from [Sturm et al. \(2010\)](#) in this paper. However, it is complementary to the current paper as it looks at kinematic data, and this paper investigates haptic data.

Lastly, as introduced in Sect. 2.1, we captured the applied forces and the trajectories of the points of contact while people pulled open doors and drawers in our previous work ([Jain et al. 2010](#)). We showed that the forces can be transformed to an object-centric representation, and we discussed the implications of such a database of forces and kinematic trajectories for assistive robot design. In the current paper, we demonstrate how robots can use such a database in a variety of ways.

3 Capturing haptic interactions

We describe how we previously captured the haptic interactions for humans and our method for capturing additional data for the two robots that we use in this paper. We use the term haptic interaction to refer to a sequence of tuples that represent the applied force and the state of the object while a task is being performed. In this paper, a haptic interaction consists of a sequence of 2-tuples that contain the opening

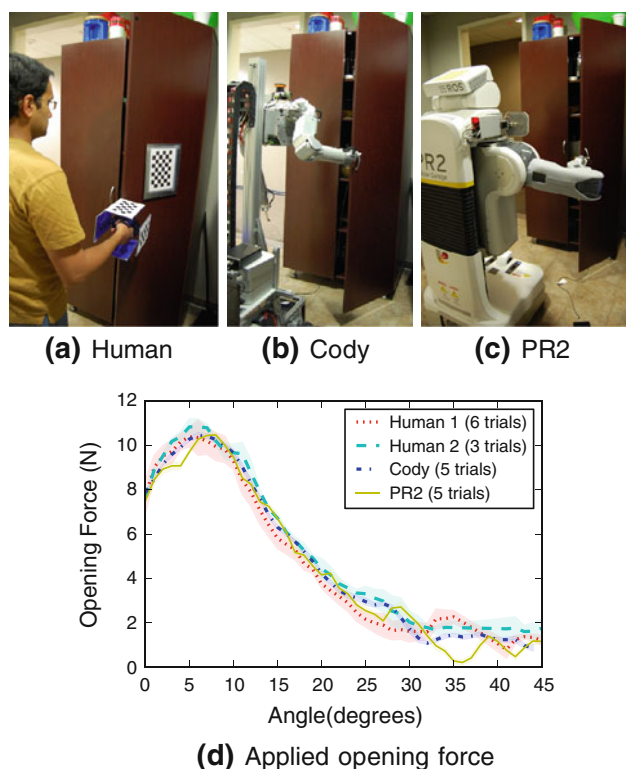


Fig. 1 When opening the same mechanism using an instrumented hook (a)–(c), two humans and two robots apply very similar opening forces to the handle in spite of their distinct bodies and control. (d) The mean and standard deviation of the opening force as a function of mechanism configuration for multiple trials

force (component of the force applied to the door handle that is tangential to the trajectory of the point of contact) and the angle through which the door has been opened.

3.1 Quasi-static model of doors

Our methods for haptic identification and anomalous force detection rely on modeling the relationship between the relevant forces, f , applied to a mechanism, m , and the mechanism's state, θ . As such, we need to make measurements to obtain estimates of the relevant applied forces, \hat{f} , and the relevant state of the mechanism, $\hat{\theta}$.

In this work, our kinematic model of doors is a single degree of freedom (DoF) rotary joint whose axis of rotation is parallel to gravity, as shown in Fig. 2. We have found that at relatively low opening speeds, the configuration-dependent forces dominate the haptic interactions (Jain et al. 2010). So, our model assumes that the relevant state consists solely of the angle of the door.

Our model also assumes that the relevant force is the component of the total force applied to the handle that is tangential to the trajectory of the point of contact. From Fig. 2, the tangential component of the applied force will open the mechanism, while other components will result in constraint




Fig. 2 Top: Diagram showing the top view of a door as it is opened using a hook. Bottom: The relevant force, f , is the component of the total force between the hook and the handle that is tangential to the trajectory of the point of contact on the door handle. We refer to this component as the opening force. The radial force, f_{rad} , will result in constraint forces at the hinge. The relevant mechanism state is the opening angle, θ




Fig. 3 The hook with a force-torque sensor at its base that we used for the three data capture systems. Left: Handheld hook used by human participants. Middle: The hook end effector mounted on Cody. Right: The hook with an adapter that the PR2 grasped with its gripper

forces at the hinges. In this paper, we will refer to this tangential component as the opening force.

3.2 Estimating the relevant applied force and the relevant mechanism state

To measure the applied force, we used a rigid hook that we 3D-printed with ABS plastic and instrumented with a six-axis force-torque sensor (ATI Nano25 with a calibration of SI-125-3) (Fig. 3).

To estimate the mechanism state, we first measured the trajectory of the door handle as a human operator (Sect. 3.3) or a robot (Sect. 3.4) opened the door, using a motion capture system or forward kinematics, respectively. We then fit a circle to this trajectory to estimate the radius and the location of the axis of rotation of the door. This procedure enabled us to estimate the angle of the door for each point of the trajectory.

We also used the estimated angle of the door to compute the component of the force measured by the force-torque

sensor that was tangential to the trajectory of the point of contact on the handle. We have previously described our method for estimating the relevant state and relevant force in Jain et al. (2010).

3.3 Capturing forces applied by humans

In previous work, we captured the forces and kinematic trajectories as human operators opened 29 doors and 15 drawers in 6 homes and one office. We created a database of estimates of the opening force applied to the handle, \hat{f} , as a function of the estimated opening angle of the door, $\hat{\theta}$. For this paper, we used this database for haptic identification (Sect. 4) and anomalous force detection (Sect. 5). Figure 1a shows part of our force and motion capture system, described in detail in Jain et al. (2010).

As in our previous work, we filtered the captured data and used trials with a low per-trial average velocity of the door handle. The mean of this average velocity across all trials with all doors was 18.2 cm/s and the median was 14.4 cm/s. For this paper, we removed from consideration all doors that had fewer than two trials in the database.

3.4 Capturing forces applied by robots

We use a feedback controller that we developed in Jain and Kemp (2010) to enable two robots to autonomously open a door without prior knowledge of the kinematics. The input to this controller is a 3D location and orientation of the door handle. Additionally, for this paper, we positioned each robot such that the handle was in its workspace and it was facing the surface of the door before running the feedback controller to autonomously open the door.

Both robots use the same feedback controller to autonomously open doors, but they have different low-level control. We use joint space impedance control on Cody (Jain and Kemp 2010) and a Cartesian space stiffness controller on the PR2 (Glaser 2010). Section 7.1 describes the two robots.

While each robot was opening a door, we recorded the trajectory of the tip of the hook (using joint encoders and forward kinematics), and the force measured by the force-torque sensor that we attached to the hook, as shown in Fig. 3. We then used the method described in Sect. 3.2 to estimate the opening force applied to the handle, \hat{f} , and the angle of the door, $\hat{\theta}$. We use haptic data from trials with the robots along with the database of haptic interactions from humans to report results of anomalous force detection in Sect. 7.

3.5 Representing a haptic interaction

We represent a haptic interaction as a sequence of tuples with the estimates of the applied opening force and the

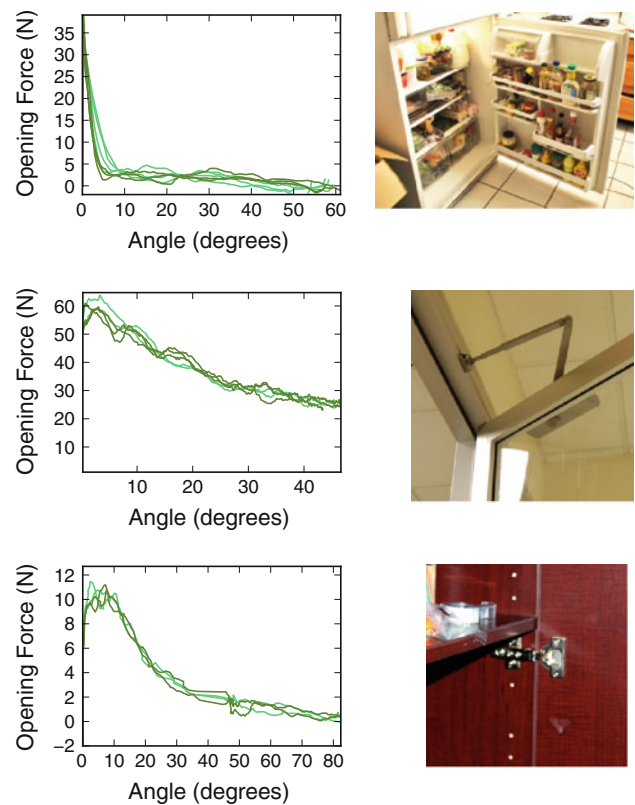


Fig. 4 Forces recorded while humans opened three doors. Left plots show applied opening forces as a function of the mechanism's configuration. *Lighter green* indicates trials with higher average velocity. Pictures on the right highlight a key mechanical element of each mechanism. *Top*: Refrigerator, 7 recordings, avg. velocities of 16.5–23.8°/s. High initial force due to magnetic door gasket (Shin et al. 2012). *Middle*: Springloaded door, 5 recordings, avg. velocities of 6.4–12.8°/s. Large forces throughout movement due to linkage at top. *Bottom*: Cabinet, 4 recordings, avg. velocities of 19.8–29.3°/s. Non-linear spring keeps it closed with max force at about 4°. Figure reproduced from our previous publication (Jain et al. 2010)

mechanism's state, $\{(\hat{f}_1, \hat{\theta}_1), (\hat{f}_2, \hat{\theta}_2), \dots, (\hat{f}_N, \hat{\theta}_N)\}$. Figure 4 shows examples of raw haptic interactions captured when people opened the pictured mechanisms. For opening doors, the handle defines the location at which the instrumented hook applies forces to the mechanism. For other mechanisms, the haptic interaction would need to explicitly include the point of application of the force relative to the manipulated object.

In this paper, we further process this raw haptic interaction into a more compact and uniform representation. We first quantize the opening angle of the door into 1° intervals. We then represent each haptic interaction by a fixed length vector, where each element in the vector is set to the mean opening force applied in the corresponding 1° interval, or set to NaN if the interval was not encountered. Within this paper, we will refer to this vector as the haptic interaction vector.




Fig. 5 Data from the PR2 opening the same mechanism 15 times (five times from three different positions relative to the mechanism handle). The mean and standard deviation of the applied opening force, \hat{f} , as a function of mechanism configuration

3.6 Sharing haptic data

We now illustrate that haptic interactions from opening doors can be insensitive to some forms of task variation. First, Fig. 5 illustrates the small variation in the haptic interaction resulting from changes in the PR2's position relative to the handle of a cabinet. For this test, we positioned the PR2 at three different positions 10 cm apart along a line parallel to the surface of the door. The relative height of the handle and its distance from the robot normal to the surface of the door were the same.

Second, Fig. 1 shows the mean and standard deviation (over multiple trials) of data from four sets of trials in which two humans and two robots opened the same mechanism. For our choice of the relevant component of the force, the configuration-dependent force due to the mechanism dominates the variation due to the operator. This observation and the results in Sect. 7 demonstrate that robots and humans can share haptic data through a common database of haptic interactions.

4 Haptic recognition

In this section, we show that standard classifiers trained with supervised machine learning can recognize the class of a door (e.g., refrigerator or kitchen cabinet), as well as the specific door. We used a dataset of 148 haptic interaction vectors, defined in Sect. 3.5, from four humans opening 26 different doors as described in Jain et al. (2010). These 148 haptic interaction vectors include data from one person per door. In our previous publication, we also demonstrated that when many different people opened the same door, the variation in the opening force as a function of the opening angle was relatively small.

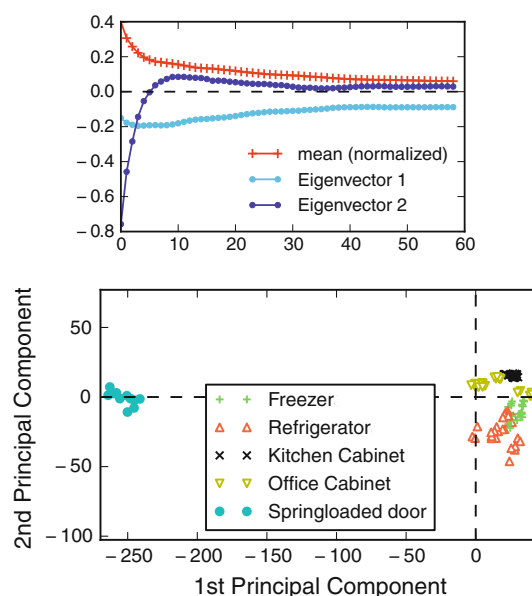


Fig. 6 Top: The mean and the first two principal components of the dataset of haptic interaction vectors. Bottom: Scatter plot showing the projection of all the haptic interaction vectors onto the first two principal components

4.1 Dimensionality reduction

To reduce the influence of noise and overfitting, we first computed a low-dimensional representation of the haptic interaction vectors with principal component analysis (PCA) using singular value decomposition. This section is an extension of the analysis from our previous publication (Jain et al. 2010). Figure 6 shows the first two principal components and a scatter plot for 148 vectors from 26 doors with the points colored by mechanism class. The first three principal components account for 99.4 % of the variance over these 148 vectors.

The scatter plot of Fig. 6 shows that even after projecting to two dimensions, there tends to be separation between different classes of doors. We use these same classes in Sect. 4.3 to present results on haptic recognition. The points corresponding to the office cabinet class that are close to the freezer class are from a cabinet which has a magnet in the door, instead of a spring in the hinge. Qualitatively, the haptic interaction vectors for this cabinet have a similar shape but lower magnitude compared to freezers, potentially because freezers are also held shut using magnets.

4.2 Recognizing a specific mechanism

We now present results on haptically recognizing a specific door after opening it. We assume that the database includes haptic interactions from previously opening the same specific mechanism.

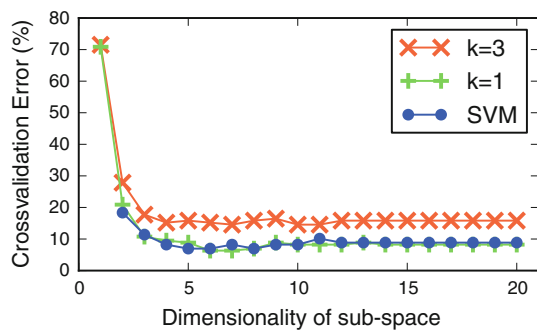


Fig. 7 Leave-one-out cross-validation error as a function of the dimensionality of the subspace for three classifiers

Figure 7 shows the leave-one-out cross-validation error for a k -nearest neighbor classifier ($k = 1$ and $k = 3$) and a support vector machine on our dataset for subspaces of different dimensionality. We used the multiclass C-SVM with a polynomial kernel of degree three as implemented in PyMMPA (Hanke et al. 2009).

The cross-validation errors with the kNN classifier ($k = 1$) and the SVM were similar. The error was nearly constant for subspaces of dimensionality ≥ 5 . Figure 8 shows the confusion matrix for the kNN classifier after projecting the data onto the first five principal components. The leave-one-out cross-validation accuracy for identifying the mecha-

nism was 89.7 %. Confusion occurred between mechanisms for which the opening forces were similar. Scaling this database up to a large size might reveal further mechanism subclasses, such as doors made by various manufacturers.

4.3 Recognizing the mechanism class

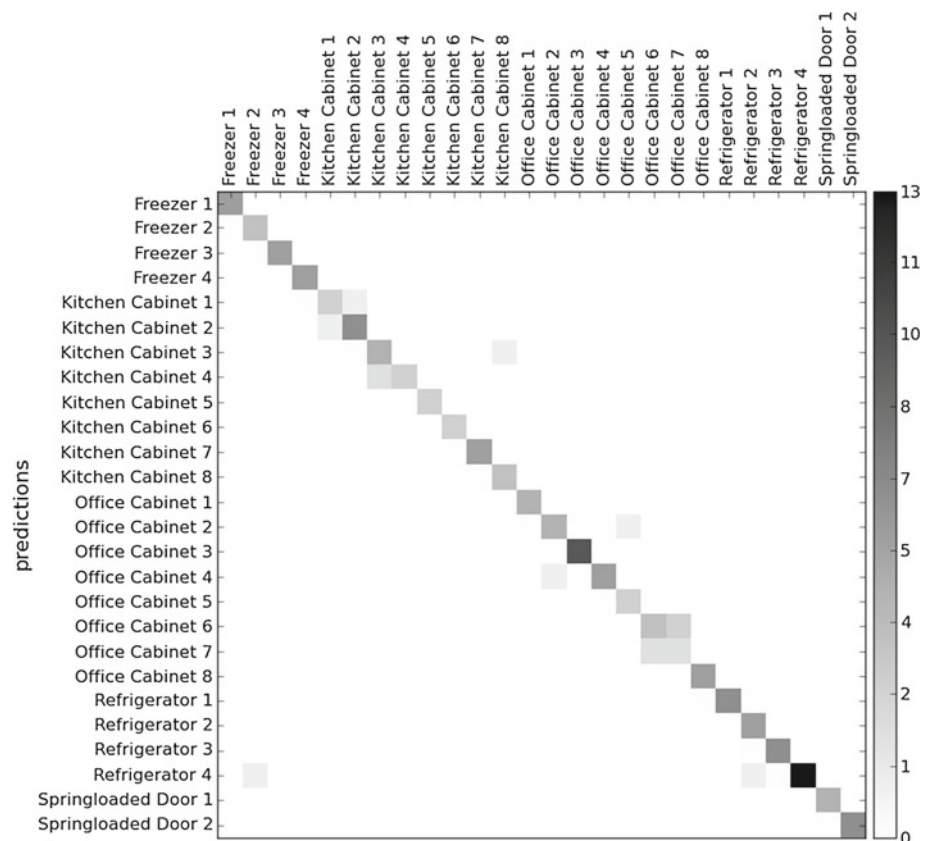
In this section we look at the problem of identifying the class of a mechanism after a human or a robot opens it. We assume that our database includes haptic interactions from opening mechanisms in the same class, but does not include the specific mechanism.

We assigned the class labels of “freezer”, “refrigerator”, “kitchen cabinet”, “office cabinet”, and “spring loaded door” to each haptic interaction vector in our database.

We used a kNN classifier ($k = 1$) and a 5 dimensional linear subspace (PCA) for class recognition. For a selected mechanism, we generated a training set by removing all the vectors from that mechanism from our dataset to simulate opening the mechanism for the first time. We then tested the kNN classifier for each of the vectors from the selected mechanism, and repeated this procedure for all 26 mechanisms.

Figure 9 shows the confusion matrix for our dataset. The cross-validation accuracy for identifying the class for the 26 mechanisms, given that the specific mechanism had not been encountered before, was 86.4 %. Most of the classification

Fig. 8 Confusion matrix from leave-one-out cross-validation with a kNN classifier ($k = 1$), after dimensionality reduction to five dimensions using PCA. The goal is to recognize a specific mechanism (Sect. 4.2). The cross-validation error was 9.3 %






Fig. 9 Confusion matrix from cross-validation with a kNN classifier ($k = 1$), after dimensionality reduction to five using PCA. The goal is to identify the mechanism class (Sect. 4.3). The cross-validation error was 13.6 %

errors were between the refrigerator and freezer classes. Both of these classes have similar opening forces as a function of the configuration, with the main difference being the initial force required to open them.

4.4 Summary

Haptic recognition can serve as a way for a robot to check that it is correctly performing a manipulation task. For example, by recognizing a specific mechanism, a robot could confirm that it opened the door it intended to open. Through recognition of a specific door, it might also infer its location in an environment. Recognizing the category of a mechanism could help a robot infer other properties. For example, the category of a door relates to its use, name, appearance, location, the objects found behind it and the category of the room. Haptic information in conjunction with other perceptual modalities could help make mechanism categorization more robust or support multi-modal category learning. For example, a robot could use haptic and visual features together to recognize a door while opening it, which could then be used to better detect anomalous forces, as described in the next section.

5 Anomalous force detection

5.1 A simple example

An example of the potential benefit of haptic data can be observed from Fig. 10, which shows the mean and standard deviation of the maximum opening force encountered over the first 10° of opening. The large variation in this initial




Fig. 10 Maximum force (mean and standard deviation) expected while opening mechanisms of five different classes of rotary mechanisms by an angle of 10°

opening force for different mechanism classes illustrates that knowledge of a mechanism's class can enable a robot to better select a maximum force to apply to a door before deciding that it is locked or malfunctioning. The applied opening force at which a robot would decide to stop pulling on a door's handle could vary from less than 10 N for kitchen cabinets to greater than 60 N for springloaded doors. Without using this type of information, a robot risks damaging a locked cabinet or giving up prematurely on a functioning springloaded door.

5.2 A probabilistic model of the relevant force

We now present a probabilistic model of the relevant force, f , applied while successfully operating a mechanism, m , conditioned on the relevant mechanism state, θ , its class, C , and any previous data from the specific mechanism, D_θ . D_θ is a vector of any forces previously measured at state θ while operating mechanism m .

We model the relevant applied force at a particular mechanism state θ as being normally distributed and conditionally independent of other forces, given θ , C , and D_θ . So,

$$P(f|\theta, C, D_\theta) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(f-\mu)^2}{2\sigma^2}}, \quad (1)$$

where μ is the mean and σ^2 is the variance of our Gaussian model for the relevant force at mechanism state θ . Our goal is to estimate μ and σ^2 given θ , C , and D_θ .

5.2.1 Operating a mechanism for the first time

Consider the case when a robot operates a mechanism for the first time. We assume that the robot knows the class C to which the mechanism belongs. For example, using vision and knowledge of the room type, the robot might know that the door is a kitchen cabinet. The robot also has access to a database of haptic interaction vectors, defined in Sect. 3.5, from mechanisms that are members of class C .

In this situation, the robot knows the mechanism's state θ and the mechanism's class C , but D_θ is a zero-dimensional vector, since the robot has not previously operated the specific

mechanism. We use the haptic data from mechanisms of class C to estimate μ and σ^2 with a weighted sample mean and weighted sample variance, $\hat{\mu}$ and $\hat{\sigma}^2$ (Bishop 2006). So,

$$\hat{\mu} = \frac{\sum_{m \in C} (w_m \sum_i \hat{f}_{m,\theta}^i)}{\sum_{m \in C} \sum_i w_m} \quad (2)$$

$$\hat{\sigma}^2 = \frac{\sum_{m \in C} (w_m \sum_i (\hat{f}_{m,\theta}^i - \hat{\mu})^2)}{\sum_{m \in C} \sum_i w_m} \quad (3)$$

where $\hat{f}_{m,\theta}^i$ represents the element of the i th haptic interaction vector for mechanism m corresponding with mechanism state θ . $m \in C$ represents selecting a mechanism from all mechanisms in class C . The weight for mechanism m is

$$w_m = \frac{1}{\# \text{ of haptic interaction vectors for } m}. \quad (4)$$

5.2.2 Operating a mechanism for the n th time

Now, consider the case when a robot operates a mechanism that it has previously operated. In this situation, the robot knows the mechanism's state θ , the mechanism's class C , and D_θ , which is a vector with previous forces for this specific mechanism at state θ . In this case, we make a maximum a posteriori (MAP) estimate (Bishop 2006) of μ and σ^2 with the following equation:

$$(\hat{\mu}, \hat{\sigma}^2) = \underset{(\mu, \sigma^2)}{\operatorname{argmin}} \left\{ -\log P(\mu, \sigma^2 | \theta, C, D_\theta) \right\}. \quad (5)$$

We use Bayes' rule to obtain

$$P(\mu, \sigma^2 | \theta, C, D_\theta) = \frac{P(D_\theta | \mu, \sigma^2, \theta, C) P(\mu, \sigma^2 | \theta, C)}{P(D_\theta | \theta, C)}. \quad (6)$$

Assuming that the forces from the previous $n - 1$ operations of the mechanism were independently drawn from $N(\mu, \sigma^2)$, then

$$P(D_\theta | \mu, \sigma^2, \theta, C) = \prod_i \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(D_\theta^i - \mu)^2}{2\sigma^2}}. \quad (7)$$

We model $P(\mu, \sigma^2 | \theta, C)$, the prior distribution over (μ, σ^2) given the state θ and class C , as a normal distribution with mean $(\mu_\mu, \mu_{\sigma^2})$ and covariance matrix *diagonal* $(\sigma_\mu^2, \sigma_{\sigma^2}^2)$. We estimate $(\mu_\mu, \mu_{\sigma^2})$ and $(\sigma_\mu^2, \sigma_{\sigma^2}^2)$ by first computing the sample means $\hat{\mu}_{\theta,m}$ and sample variances $\hat{\sigma}_{\theta,m}^2$ of the measured opening forces at state θ for each mechanism m in class C . We then compute $(\hat{\mu}_\mu, \hat{\mu}_{\sigma^2})$ by concatenating the sample means of $\hat{\mu}_{\theta,m}$ and $\hat{\sigma}_{\theta,m}^2$ over all the mechanisms in class C into a vector. Likewise, we compute $(\hat{\sigma}_\mu^2, \hat{\sigma}_{\sigma^2}^2)$ by concatenating the sample variances of $\hat{\mu}_{\theta,m}$ and $\hat{\sigma}_{\theta,m}^2$.

Equation 5 now simplifies to

$$(\hat{\mu}, \hat{\sigma}^2) = \underset{(\mu, \sigma^2)}{\operatorname{argmin}} \sum_i \left(\log \sigma + \left(\frac{D_\theta^i - \mu}{\sigma} \right)^2 \right) + \left(\frac{\mu - \hat{\mu}_\mu}{\hat{\sigma}_\mu} \right)^2 + \left(\frac{\sigma^2 - \hat{\mu}_{\sigma^2}}{\hat{\sigma}_{\sigma^2}} \right)^2. \quad (8)$$

We find approximate solutions for Eq. 8 using the implementation of the BFGS optimization algorithm from SciPy (Jones et al. 2001) with seed estimates $\mu = \sum_{i=1}^{n-1} D_\theta^i / (n - 1)$ and $\sigma^2 = \hat{\mu}_{\sigma^2}$. The robot is opening the mechanism for the n th time and has data from $n - 1$ previous operations of that specific mechanism.

5.3 Detecting anomalous forces

We detect an anomaly if the force measured at the current state of the mechanism, θ , equals or exceeds a threshold force, i.e.

$$\hat{f}_\theta \geq f_\theta^{thresh}. \quad (9)$$

We do not investigate the potential for a lower bound on the force, although a low force could also be indicative of an anomaly. Additionally, our detector only uses the opening force at the current angle of the door. As discussed in Sect. 8.2.4, features that use information over time or across multiple configurations of the mechanism could also be beneficial for anomaly detection.

We present three methods of determining f_θ^{thresh} . The first method uses our probabilistic model of expected forces (Sect. 5.2) to detect when forces are unlikely given the mechanism's state θ , the mechanism's class C , and any previous operation of the mechanism D_θ . For this method,

$$f_\theta^{thresh} = \hat{\mu} + n\hat{\sigma}, \quad (10)$$

where the parameter n serves as a detector threshold that adjusts the sensitivity and specificity of the anomaly detector.

The other two detectors are baseline methods that use no prior information about the mechanism or its class. The first baseline detector sets f_θ^{thresh} equal to a constant c . The second baseline detector sets $f_\theta^{thresh} = r \cdot \hat{f}_{initial}$, where r defines a fixed ratio of the initial opening force. We have used this method in our previous work (Jain and Kemp 2010). Like n , c and r serve as detector thresholds that adjust the sensitivity and specificity of these anomaly detectors.

5.4 Performance measures for anomalous force detection

We evaluated the performance of anomalous force detection methods using: (1) the increase in the magnitude of the force from the onset of contact between the door and an obstacle and the detection of anomalous force; (2) how much time

passes between the manually labeled onset of contact and the detection of anomalous force; and (3) the false positive rate (i.e. the frequency with which the detector incorrectly reports a force as being anomalous). Lower force implies less risk of damage to the robot and the environment, and less risk of injury to people. Faster detection implies that the robot can more efficiently respond to the event that caused an anomalous force, such as by trying a new strategy or stopping and asking for assistance. A lower false positive rate implies that the robot will be less likely to falsely detect an anomalous force and thereby unnecessarily change its approach or give up.

6 Evaluation with human data

This section shows that captured haptic interactions can be used to improve manipulation in two ways. First, by using captured haptic interactions, an anomalous force detector can reduce the increase in force from the onset of a contact until its detection. Second, captured haptic interactions enable the detector to report contact with an obstacle faster. The first improvement corresponds to an increase in the safety of the system by lowering the excess force applied to the door, and the second corresponds to an improvement in the efficiency.

Cameras and other non-contact line-of-sight sensors are not well matched to many of these detection problems, which naturally occur as anomalous forces. For example, a door gently touching and deforming a curtain may not be an anomalous condition. Also, there may not be any clear visual cue for a locked door.

6.1 Modeling locked doors and contact with an idealized fixed rigid obstacle

Our database of haptic interactions consists of collision-free trials only. In order to compare the performance of these three anomaly detection methods using this database, we simulated locked doors, and contact between the door and an idealized fixed rigid obstacle.

We modeled these situations as a force that increases monotonically with time while the configuration of the mechanism remains constant. The actual rate of increase of the force over time would depend on the robot control algorithm. For example, with an impedance controlled robot, we would expect the force to increase at a rate that depends on the stiffness at the end effector.

This model enabled us to simulate contact at any configuration of the mechanism and investigate how well our anomaly detection methods could detect lower magnitude anomalous forces while avoiding false alarms.

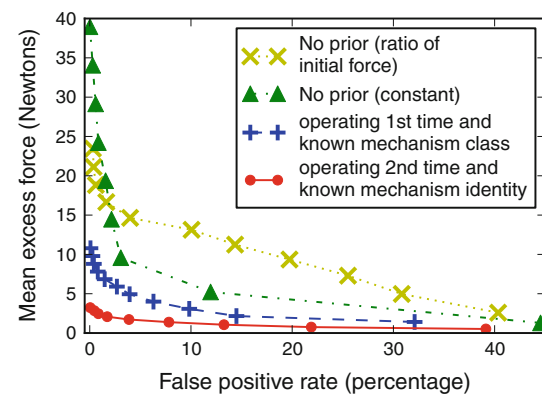


Fig. 11 Graph illustrating the trade-off different anomalous force detection methods make between the mean excess force before reporting contact (y axis) and the false positive rate, which is the percentage of times that a collision-free trial is incorrectly reported as having an anomalous force (x axis). These results use the collision-free dataset of human trials with simulated contact with an idealized rigid obstacle, described in Sect. 6.1

6.2 Comparison of anomaly detection methods

For this comparison, we simulated contact at discrete angles (1° intervals) for all operations of all doors. Given our model of contact with an idealized rigid obstacle, each of the three detectors will eventually detect the contact, since the magnitude of the applied force will continue to increase over time. As such, we focus on the increase in the magnitude of the force from the onset of contact until its detection, and the false positive rate. This increase in the magnitude of the force is the excess force applied to the mechanism before the detector reports contact with an obstacle. False positives correspond to the detector incorrectly reporting contact for a collision-free trial.

Let us assume that the door is at a configuration θ when it makes contact. For a given value of f_θ^{thresh} , the excess force before contact is detected will be $f_\theta^{thresh} - \hat{f}$, where \hat{f} is the estimated opening force at configuration θ for one collision-free trial of pulling open the mechanism door.

We compute the average value of this excess force over all the configurations of all the trials of all the mechanisms to obtain the mean excess force before an anomalous force is detected. This average value gives us the y coordinate of a point in the plot of Fig. 11. The x axis is the percentage of the configurations for which $f_\theta^{thresh} < \hat{f}$. Since our database contains only collision-free trials, $f_\theta^{thresh} < \hat{f}$ is a false positive for anomaly detection.

Figure 11 shows the performance of the three detectors for different values of n , r , and c on 148 trials from 26 different rotary mechanisms belonging to 5 different classes. Each point in the plot represents a different value of f_θ^{thresh} , obtained by changing the value of the parameters n , r , and c for the different methods of detecting anomalous forces.

For all detectors, increasing the free parameter (n , r , or c) decreases the false positive rate, but increases the unnecessary force applied prior to detection.

The blue curve in Fig. 11 shows the performance of our data-driven detector when opening a particular mechanism of a known class for the first time. For it, we computed f_{θ}^{thresh} for each trial in our dataset by ignoring all data from that specific door and computing the sample mean and variance over all the other mechanisms in the same class, as described in Sect. 5.2.1. The red curve shows the performance of our data-driven detector when opening a particular mechanism of a known class for the second time. For it, we used the data from all the other mechanisms in the same class, along with a single collision-free trial from the mechanism, as described in Sect. 5.2.2 with $n = 1$. We computed our results using all possible initial trials for a mechanism. For example, 5 total trials for a particular mechanism in our dataset would result in 20 tests, due to there being 4 possible initial trials for each of the 5 trials.

Figure 11 shows that using captured haptic data can reduce the excess force for a given false positive rate. It shows that knowledge of the mechanism class enables the detection of an anomalous force with a lower excess force, since for any false positive rate, the blue curve is below the green and yellow curves. Additionally, it shows that even a single trial with the specific mechanism decreases the excess force further, as evidenced by the fact that the red curve is below the blue curve for all false positive rates.

7 Experiments with two robots

7.1 The robots

We used two different mobile manipulators for this paper. Figure 1 shows the two mobile manipulators, and Fig. 3 shows how we mounted the hook end effector to them. One mobile manipulator is a PR2 robot from Willow Garage, which has two 7 DoF compliant arms with low gear ratios and current control for the motors at the joints (Wyrobek et al. 2008). The second robot, Cody, has two compliant 7 DoF arms from Meka Robotics with series elastic actuators for torque control at each DoF.

For low-level 1 kHz control on the PR2, we use low stiffness gains with an open source controller (Glaser 2010) that is similar to Cartesian stiffness control (Salisbury 1980). On Cody, we use joint-space impedance control running at 1 kHz (Jain and Kemp 2009b). Although the two robots have different lower-level control and actuation, we use equilibrium point control as described in Jain and Kemp (2010) on both robots to autonomously open doors with a hook end effector.

7.2 Online estimation of the relevant force and mechanism state

To detect anomalous forces using the methods of Sect. 5.3, a robot needs to generate online estimates of the state of the mechanism and the opening force while operating the mechanism. We estimate the radius of the trajectory of the handle, r , and the location of the axis of rotation, (c_x, c_y) , which we then use to compute an estimate of the mechanism state and opening force. We denote (r, c_x, c_y) with β .

We assume that a perception algorithm gives the robot an initial estimate of the radius of the trajectory of the handle, r_p . As an example, the perception algorithm could compute r_p based on the estimated width of the door and the location of the handle (Rusu et al. 2008). In addition, the robot estimates the pose of the tip of the hook using forward kinematics while operating the mechanism, giving it an estimate of the trajectory of the handle, $((x_1, y_1), (x_2, y_2), \dots, (x_n, y_n))$ which we denote as T_n . The number of points in the trajectory of the handle, n , increases with time.

Given r_p and T_n we compute a maximum likelihood estimate (Bishop 2006) of β as

$$\hat{\beta} = \underset{\beta}{\operatorname{argmax}} P(T_n, r_p | \beta), \quad (11)$$

We assume that the observed trajectory, T_n , and the perception algorithm's estimate of the radius, r_p , are conditionally independent given β . So,

$$P(T_n, r_p | \beta) = P(T_n | \beta) P(r_p | \beta). \quad (12)$$

Next, we assume that the perception algorithm's estimate of the radius is normally distributed around the true radius of the mechanism, with a variance of σ_r^2 . So,

$$P(r_p | \beta) = \frac{1}{\sqrt{2\pi}\sigma_r} e^{-\frac{(r_p - r)^2}{2\sigma_r^2}}. \quad (13)$$

We compute $P(T_n | \beta)$ using the assumptions detailed in Sturm et al. (2009), which include assuming that the measurements of the points along the trajectory of the handle, (x_i, y_i) , are conditionally independent given β and have a Gaussian error with a variance of σ_{pos}^2 . These assumptions result in

$$P(T_n | \beta) = \prod_{i=1}^n \left(\frac{1}{\sqrt{2\pi}\sigma_{pos}} e^{-\frac{(r - \sqrt{(c_x - x_i)^2 + (c_y - y_i)^2})^2}{\sigma_{pos}^2}} \right). \quad (14)$$

Equation 11 then simplifies to

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \frac{(r - r_p)^2}{\sigma_r^2} + \sum_{i=1}^n \frac{(r - \sqrt{(c_x - x_i)^2 + (c_y - y_i)^2})^2}{\sigma_{pos}^2}, \quad (15)$$

which we optimize using the implementation of the BFGS algorithm from SciPy (Jones et al. 2001). We then use $\hat{\beta}$ to compute the current state of the mechanism. For our tests, we set $\sigma_r = 10$ cm, and $\sigma_{pos} = 1$ cm. We set these values conservatively to reflect uncertainty due to a variety of factors. For example, perceptual uncertainty when estimating the width of the door would influence σ_r , and uncertainty due to joint encoder resolution and our use of a hook with a layer of compliant rubber would influence σ_{pos} .

7.3 Using data captured from humans

For our results with robots in Sects. 7.4 and 7.5, we used the entire dataset of human trials. This included data from humans opening the same refrigerator and cabinets that the robots opened. As a result, although we investigated cases where the robot opens a door for the first or second time, the database used to model the prior for the class of mechanism ($P(\mu, \sigma^2|\theta, C)$ from Sect. 5.2.2) included haptic interactions from humans opening the same door.

7.4 Effect of online estimation on the performance of anomalous force detection

We investigated how noise in online estimates can impact the performance of our data-driven anomalous force detector. For this evaluation, we recorded the trajectory of the hook and the applied force while robots opened doors. We collected a total of fifteen collision-free trials consisting of the PR2 opening an office cabinet five times, Cody opening a different office cabinet five times, and Cody opening a refrigerator five times.

We then generated multiple haptic interaction vectors, simulating online estimation of the mechanism state and a noisy initial estimate of the radius of the door (r_p), as follows: for each of the 15 trials, we computed multiple values of r_p in Eq. 15 by sampling from a Gaussian with mean equal to the true radius of the mechanism and a standard deviation of 10 cm. We then used $\hat{\beta}$ from Eq. 15 to generate simulated online estimates of the state of the mechanism and the opening force, which is the component of the force tangential to the trajectory of the point of contact on the handle. The online estimates of the state of the mechanism improved as the robot opened the door through a larger angle.

Figure 12 shows the trade-off between the mean excess force and the false positive rate, analogous to the results from Sect. 6.2. Errors in the estimates of the mechanism configuration (due to uncertainty about the radius and location of the axis of rotation) resulted in poorer performance when the robot operates a mechanism for the first time (higher force on average before detection of contact with an idealized rigid obstacle). When the robot operates a mechanism with a known identity for the second time, it has the opportunity to use its previous estimates of the mechanism's radius.




Fig. 12 Effect of errors in kinematic state estimation on performance of anomalous force detection using data captured with Cody and the PR2

We would also expect an improvement in performance if a robot were to use methods for kinematic estimation that yield more accurate initial estimates of the radius of the door and the state of the mechanism, such as in Rusu et al. (2008), Sturm et al. (2009, 2010).

7.5 Detecting anomalous forces due to collisions and a locked door

We also investigated the performance of our data-driven anomalous force detector with real collisions and a real locked door. We performed six trials with two robots, Cody and the PR2, shown in Fig. 13. We either placed an obstacle in front of the door or locked the door. We processed data collected from these trials off-line for anomalous force detection using the detector based on our probabilistic model of applied forces (Eq. 10). We set n in Eq. 10 as the least value that resulted in zero false positives over our entire dataset of human trials. Further, we assumed that the robot has an accurate estimate of the radius of the mechanism (as opposed to being drawn from a Gaussian around the true radius as described in Sect. 7.2). Table 1 shows the time (detection time) and additional force (excess force) between the onset of contact with the obstacle and detection of an anomalous force, or between the force on a locked door first exceeding the expected force and detection of an anomalous force.

Figure 13 illustrates the performance of our data-driven anomalous force detector. The yellow curve is the opening force applied by the robot during the trial. The dashed green and blue lines are the minimum forces at which our methods would report an anomalous force if the robot were operating the mechanism for the first time (with knowledge of the mechanism class), or the second time, respectively. The three numbered circles (1–3) in Fig. 13 represent the following: (1) the manually labeled onset of contact with an obstacle or when the force on a locked door handle exceeds the expected force; (2) the point at which the robot would have detected an anomalous force if it were operating the mechanism for the second time; and (3) the point at which it would have detected

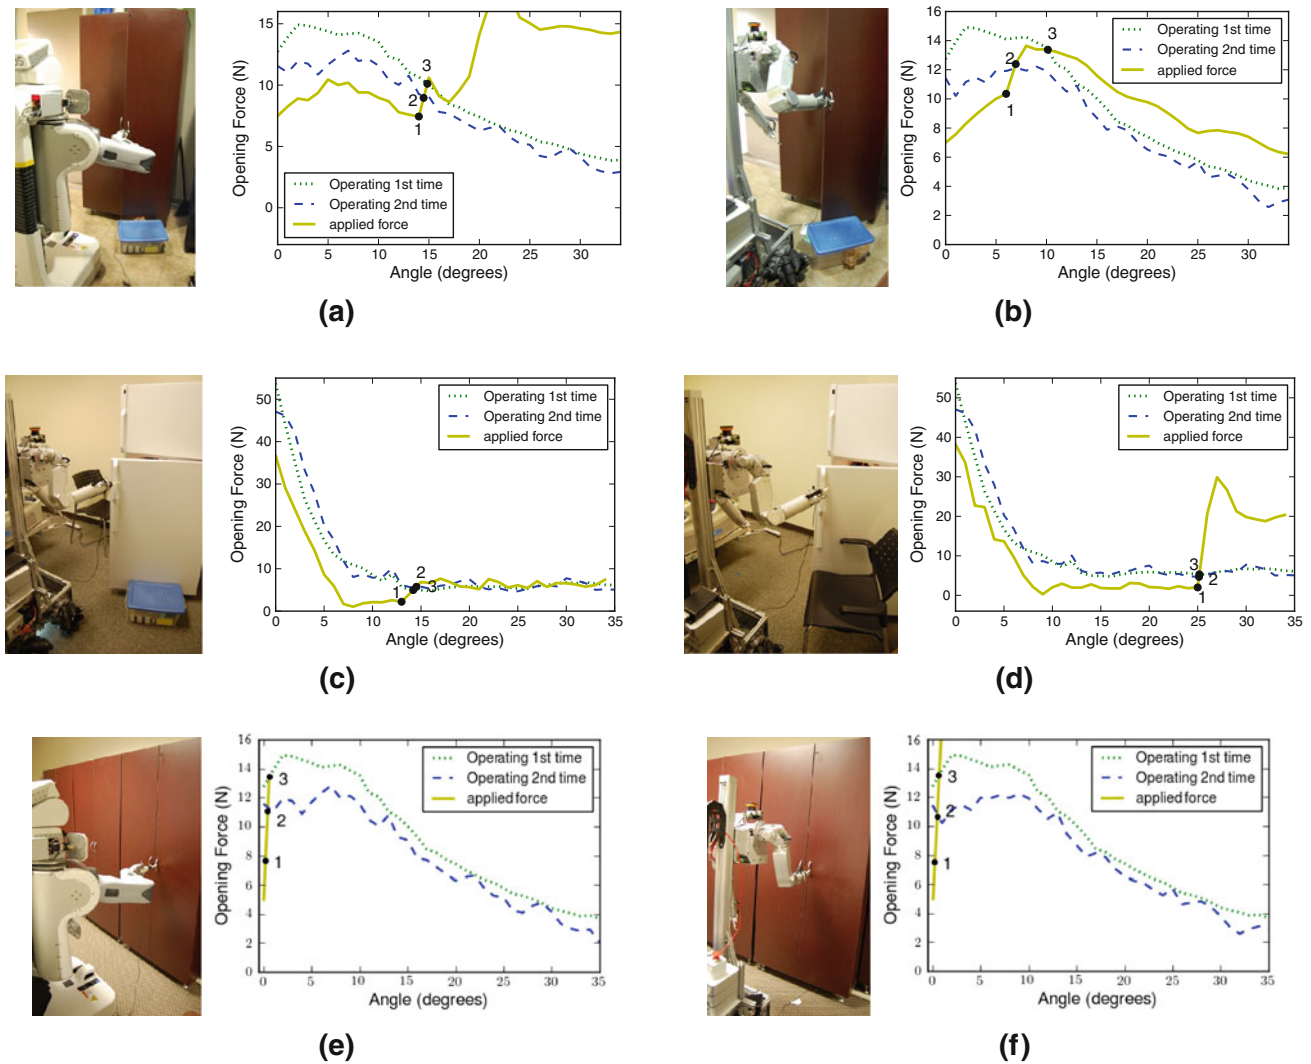


Fig. 13 Six trials with two different robots to test the anomaly detection method. Each pair of images shows a picture of the mechanism and the obstacle that we placed in the path of the door, and a graph that gives details of the trial. In the graph, the *yellow curve* is the opening force applied by the robot during the trial. The *dashed green* and *dashed blue* curves are the minimum forces at which the robot would report an anomalous force if it were operating the mechanism for the first or the second time, respectively. The three numbered *circles* indicate the

manually labeled onset of the anomalous force, and the points when the two anomaly detection methods would have reported an anomalous force, (a) Trial 1—Cabinet door makes contact with a box, (b) Trial 2—Cabinet door makes contact with a box, (c) Trial 3—Refrigerator door makes contact with a box, (d) Trial 4—Refrigerator door makes contact with a chair, (e) Trial 5—Locked cabinet door, (f) Trial 6—Locked cabinet door

Table 1 Performance of anomalous force detector on trials with Cody and the PR2

Robot	Trial	Detection time (secs)		Excess force (N)		
		Open door the		Open door the		
		1st time	2nd time	1st time	2nd time	
Trial 1	PR2	Cabinet door makes contact with a box	1.1	0.3	2.6	1.9
Trial 2	Cody	Cabinet door makes contact with a box	2.4	0.6	3.1	2.1
Trial 3	Cody	Refrigerator door makes contact with a box	1.5	1.8	2.6	3.5
Trial 4	Cody	Refrigerator door makes contact with a chair	0.5	0.5	3.7	3.7
Trial 5	PR2	Locked cabinet door	25.5	6.1	5.7	3.4
Trial 6	Cody	Locked cabinet door	1.4	0.7	6.2	3.4
		Mean (std)	5.4 (9.0)	1.7 (2.0)	4.0 (1.4)	3.0 (0.7)

an anomalous force if it were operating the mechanism for the first time.

For trials 3 and 4, the robot has a slightly lower force threshold for detecting anomalous conditions when it is operating the refrigerator for the first time than when it is operating it the second time. This also shows up as similar or better performance while opening the first time compared to the second time for the results in rows 3 and 4 in Table 1. We believe that a lower threshold when opening it for the first time is due to the database of human trials having data from only four refrigerators, resulting in the mean and variance not being representative of the mechanism class.

8 Discussion

We now discuss the broader implications and limitations of this paper, and directions for future research.

8.1 Broader implications

Machine intelligence has benefitted greatly from large collections of sensory data. Web-based databases of user-generated content, such as videos from YouTube, images from Flickr, and 3D models from Google 3D Warehouse, have begun to support the development of robots and technology relevant to robots (Klank et al. 2009; Lai and Fox 2009; Kollar and Roy 2009; Kuffner 2010; Waibel et al. 2011). More generally, research has shown that large datasets can lead to performance gains and make computationally simple methods effective (Torralba et al. 2008; Halevy et al. 2009).

Our results in this paper suggest that humans and robots have the potential to share haptic data to improve robot manipulation. Additionally, by storing associated contextual information, such as where the interactions occurred and the appearance of manipulated objects, robots could anticipate haptic interactions.

In general, this type of common sense knowledge would help robots behave more intelligently. In the future, robots might use these data in numerous ways, including selecting better postures prior to manipulation, detecting when mechanisms are in need of repair, and anticipating when a human will require assistance. Moreover, these data could be used by humans to rationally design robots with the kinematic and force capabilities necessary to perform real-world tasks.

Enabling humans to easily capture the haptic interactions, such as with a wearable system, could potentially accelerate the accumulation of this type of data. With motion capture capabilities and sensors continuing to improve in quality and decrease in cost, there is the potential for the robotics community to accumulate large datasets in a practical manner. Likewise, robots in the field could potentially record their haptic interactions and upload them to an online database in

order to produce a continuously evolving source of haptic knowledge for various manipulation tasks.

8.2 Limitations and future work

8.2.1 *More tasks, more data, more robots*

In this paper, we have presented results with real-world data for one manipulation task: slowly and smoothly pulling open doors. Scaling up our approach to more manipulation tasks, more mechanisms and trials, and more robots is an important area for future inquiry. These could include twisting door knobs, turning keys, pushing buttons on appliances, and inserting a cell phone charger into a wall socket. It could also include tasks relevant to activities of daily living such as bed baths, shaving, grooming, and manipulating a person's body (King et al. 2010).

Data-driven object-centric models of haptic interactions for these tasks may enable robots to efficiently detect anomalous conditions without excessive force. For example, a robot may be able to haptically detect that it is attempting to insert the incorrect key or that the door is not completely shut.

For this paper, we tested our methods on a dataset of forces from 148 trials of human operators opening 26 doors in 6 homes and one office, and 21 trials from two robots opening three doors in one office. Our results are promising, and suggest the potential for scaling up to more trials, mechanisms, and robots, but actually doing so remains an open area for inquiry. Ideally, the robotics community will begin to collect large-scale datasets of forces from everyday activities to facilitate progress, much like the computer vision community has collected and used large datasets of images and video.

8.2.2 *Haptic data from different sensors*

In this paper, humans and two robots used a hook instrumented with the same six-axis force-torque sensor at the base while pulling open doors. Additionally, we restricted contact between the hook and the door to be at the handle. In general, other sensors, such as joint torque sensors and tactile sensors, might be used to record the forces.

We have shown that while slowly pulling open a door, the component of the force tangential to the trajectory of the point of contact on the door handle primarily depends on the mechanism and not on the control method used to open the door. However, different sensors will have varying accuracy and noise levels. For example, using joint torque sensing to estimate the force at the end effector will be affected by the dynamics of the arm and friction and flexibility in the joints. In this paper, we do not discuss methods for combining data from distinct sensors with varying accu-

racy and amounts of noise into a common database of haptic interactions.

Robot-centric models, discussed in Sect. 2.3.1, do not provide a direct way for different robots to share information, but they make it easier for a robot to use new sensors by representing the haptic interaction directly in terms of the robot's state and sensors. Our method requires additional effort to transform the measurements from different sensors into an object-centric representation, but it offers the potential for multiple robots to share the transformed data and models of haptic interactions. Sharing data can be important for acquiring and using large datasets.

8.2.3 Identifying useful object-centric representations

For our current work, we identified the relevant applied forces and the relevant mechanism state, and found a useful low-dimensional model by using our task knowledge, modeling the kinematics, and experimenting with various models. For example, we found that a quasi-static model was sufficient for low-speed door opening and that we did not need to include the angular velocity or angular acceleration of the door in the relevant mechanism state. Similar models might work for other tasks with 1 DoF kinematics, such as twisting a door knob or turning a key. More generally, methods to automate aspects of the modeling process would be desirable and machine learning might be able to autonomously discover appropriate low-dimensional representations.

8.2.4 Haptic detection of events

In this paper, we used relatively simple models of the forces to focus on the primary goal of demonstrating the value of data-driven object-centric models of forces for autonomous robot manipulation. Our probabilistic model of the relevant force assumes that the opening forces are independent given the door angle. For anomaly detection, we use a configuration-dependent threshold on the force. Weakening the assumption that the forces are independent given the mechanism configuration and using additional features for anomalous force detection are directions for future research.

Incorporating features such as the rate of change of the opening force with the angle of the door might improve performance. Recent work on haptically detecting events during manipulation has demonstrated the use of high-frequency information to detect events such as collisions while placing an object on a table (Romano et al. 2011), and motion of liquid inside a container when the container is shaken (Chitta et al. 2011). Modeling the relationship between the mechanism state and features based on high-frequency haptic signals might improve the performance of our methods, such as when detecting a collision.

9 Conclusion

We have demonstrated that humans and robots can capture and share haptic data in spite of variations in their bodies and control methods. We recorded relevant forces, locations of force application, and mechanism states while humans and two different robots pulled open doors at low speeds. We then represented the haptic interactions with an object-centric representation that could be used by distinct robots. We demonstrated that these data could be used to haptically recognize mechanism classes, haptically recognize specific mechanisms, and haptically detect anomalous forces using data-driven object-centric models.

More generally, we have presented a method for building probabilistic models of haptic interactions that are data-driven, task-specific, and object-centric. These models can be shared by different robots for improved manipulation performance. We have used pulling open doors, an important task for service robots, as an example to demonstrate our method.

10 Supplementary material

The dataset and code associated with this paper are part of the supplementary materials.

A video showing the custom force and motion capture system that we used to record data from human trials can be viewed at the following web address: <http://www.youtube.com/watch?v=MJW77v76cJE>

Acknowledgments We thank Hai Nguyen for his technical input. This work benefitted from discussions with Tapomayukh Bhattacharjee and Kelsey Hawkins. We thank Mrinal Rath and Jason Okerman for their assistance with data collection. We gratefully acknowledge support from NSF awards IIS-1150157 and CNS-0958545, Willow Garage, and DARPA Maximum Mobility and Manipulation (DARPA M3) Contract W911NF-11-1-603. We also thank the anonymous reviewers for their valuable feedback and suggestions.

References

- Angerilli, M., Frisoli, A., Salsedo, F., Marcheschi, S., & Bergamasco, M. (2001). *Haptic simulation of an automotive manual gearshift*. In: ROMAN.
- Becker, J., Bersch, C., Pangercic, D., Pitzer, B., Rühr, T., Sankaran, B., et al. (2011). *The pr2 workshop-mobile manipulation of kitchen containers*. In: IROS workshop on results, challenges and lessons learned in advancing robots with a common platform.
- Bishop, C. (2006). *Pattern recognition and machine learning*. New York: Springer.
- Chan, W., Parker, C., Van der Loos, H., & Croft, E. (2012). *Grip forces and load forces in handovers: implications for designing human-robot handover controllers* (pp. 9–16). In: Proceedings of the seventh annual ACM/IEEE international conference on human-robot Interaction, ACM.

- Chitta, S., Cohen, B., & Likhachev, M. (2010). *Planning for autonomous door opening with a mobile manipulator*. In: ICRA.
- Chitta, S., Sturm, J., Piccoli, M., & Burgard, W. (2011). Tactile sensing for mobile manipulation. *Robotics IEEE Transactions*, 99, 1–11.
- De Luca, A., & Mattone, R. (2004). *An adapt-and-detect actuator fdi scheme for robot manipulators*. In: ICRA.
- Diankov, R., Srinivasa, S., Ferguson, D., & Kuffner, J. (2008). *Manipulation planning with caging grasps*. In: Humanoids.
- Dixon, W., Walker, I., Dawson, D., Hartrant, J. (2000). Fault detection for robot manipulators with parametric uncertainty: a prediction-error-based approach. *IEEE Transactions on Robotics and Automation*.
- Dollar, A., Jentoft, L., Gao, J., & Howe, R. (2010). Contact sensing and grasping performance of compliant hands. *Autonomous Robots*, 28(1), 65–75.
- Dupont, P. E., Schulteis, C. T., Millman, P., Howe, R. D. (1999). Automatic identification of environment haptic properties. Presence.
- Glaser, S. (2010). *Teleop controllers ROS package*. Robot operating system: Willow Garage.
- Gorges, N., Navarro, S., Goger, D., & Worn, H. (2010). *Haptic object recognition using passive joints and haptic key features*. In: ICRA.
- Haddadin, S., Albu-Schaffer, A., De Luca, A., & Hirzinger, G. (2008). *Collision detection and reaction: A contribution to safe physical human-robot interaction*. In: IROS.
- Haddadin, S., Albu-Schaffer, A., Haddadin, F., Rosmann, J., & Hirzinger, G. (2011). Study on soft-tissue injury in robotics. *IEEE Robotics and Automation Magazine*, 18(4), 20–34.
- Halevy, A., Norvig, P., & Pereira, F. (2009). The unreasonable effectiveness of data. *Intelligent Systems IEEE*, 24(2), 8–12.
- Hanke, M., Halchenko, Y., Sederberg, P., Hanson, S., Haxby, J., & Pollmann, S. (2009). Pymvpa: A python toolbox for multivariate pattern analysis of fmri data. *Neuroinformatics*, 7(1), 37–53.
- Hsiao, K., Chitta, S., Ciocarlie, M., & Jones, E. (2010). *Contact-reactive grasping of objects with partial shape information*. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).
- Jain, A., Kemp, C. C. (2009a). *Behavior-based door opening with equilibrium point control*. In: RSS workshop: Mobile manipulation in human environments.
- Jain, A., & Kemp, C. C. (2009b). *Pulling open novel doors and drawers with equilibrium point control*. In: Humanoids.
- Jain, A., & Kemp, C. C. (2010). *Pulling open doors and drawers: Coordinating an omni-directional base and a compliant arm with equilibrium point control*. In: ICRA.
- Jain, A., Nguyen, H., Rath, M., Okerman, J., & Kemp, C. C. (2010). *The complex structure of simple devices: A survey of trajectories and forces that open doors and drawers*. Proceedings of the IEEE RAS/EMBS international conference on biomedical robotics and biomechanics (BIOROB).
- Johnsson, M., & Balkenius, C. (2007). Neural network models of haptic shape perception. *Robotics and Autonomous Systems*, 55(9), 720–727.
- Jones, E., Oliphant, T., Peterson, P., et al. (2001) SciPy: Scientific tools for python. <http://www.scipy.org>.
- Kalakrishnan, M., Righetti, L., Pastor, P., & Schaal, S. (2011). *Learning force control policies for compliant manipulation*. In: IROS.
- King, C. H., Chen, T. L., Kemp, C. C., (2010). *Towards an assistive robot that autonomously performs bed baths for patient hygiene*. In: Intelligent robots and systems, 2010 (IROS 2010). IEEE/RSJ international conference on, pp. 319–324.
- Klank, U., Zia, M., Beetz, M., & (2009). *3d model selection from an internet database for robotic vision*. In: Robotics and Automation, (2009). ICRA'09 (pp. 2406–2411). IEEE: IEEE International Conference on.
- Klingbeil, E., Saxena, A., & Ng, A. Y. (2008). Learning to open new doors. In R. S. S. Workshop (Ed.), *Intelligence in human environments: Robot manipulation*.
- Kollar, T., Roy, N., & (2009) Utilizing object-object and object-scene context when planning to find things. In: *Robotics and automation, (2009). ICRA'09* (pp. 2168–2173). IEEE: IEEE International Conference on.
- Kormushev, P., Calinon, S., Caldwell, D. (2011). *Imitation learning of positional and force skills demonstrated via kinesthetic teaching and haptic input*. Advanced Robotics.
- Kuffner, J. (2010). Cloud enabled humanoid robots. Humanoids: What's next? applications, challenges and perspectives workshop.
- Lai, K., Fox, D. (2009). 3D laser scan classification using web data and domain adaptation. In: *Proceedings of of robotics: Science and systems (RSS)*.
- MacLean, K. (1996). *The haptic camera: A technique for characterizing and playing back haptic properties of real environments*. In: Proceedings of ASME dynamic systems and control division.
- Matheus, K., & Dollar, A. M. (2010). *Benchmarking grasping and manipulation: Properties of the objects of daily living*. In: IROS.
- Morinaga, S., & Kosuge, K. (2003). *Collision detection system for manipulator based on adaptive impedance control law*. In: ICRA.
- Meeussen, W. et al. (2010). *Autonomous door opening and plugging in with a personal robot*. In: ICRA.
- Navarro, S., Gorges, N., Worn, H., Schill, J., Asfour, T., Dillmann, R. (2012) Haptic object recognition for multi-fingered robot hands. In: *Haptics symposium (HAPTICS), 2012 IEEE*, pp. 497–502.
- Pai, D., Lang, J., Lloyd, J., & Woodham, R. (2000). *Acme, a telero-botic active measurement facility* (pp. 391–400). VI: Experimental robotics.
- Pastor, P., Kalakrishnan, M., Chitta, S., Theodorou, E., & Schaal, S. (2011). *Skill learning and task outcome prediction for manipulation*. In: ICRA.
- Platt, Jr. R., Permenter, F., Pfeiffer, J. (2011). Using bayesian filtering to interpret tactile data during flexible materials manipulation. IEEE Transactions on Robotics (submitted).
- Prats, M., Martinet, P., Lee, S., Sanz, P. (2009). *Compliant physical interaction based on external vision-force control and tactile-force combination*. Multisensor fusion and integration for intelligent systems pp. 221–233.
- Redmond, B., Aina, R., Gorti, T., Hannaford, B. (2010). Haptic characteristics of some activities of daily living. In: *Haptics symposium, 2010 IEEE, IEEE*, pp. 71–76.
- Rodriguez, A., Bourne, D., Mason, M., Rossano, G., & Wang, J. (2010). Failure detection in assembly: Force signature analysis. *IEEE conference on automation science and engineering (CASE)*.
- Romano, J., & Kuchenbecker, K. (2011). Creating realistic virtual textures from contact acceleration data. *IEEE Transactions on Haptics*, 5(2), 109–119.
- Romano, J., Hsiao, K., Niemeyer, G., Chitta, S., & Kuchenbecker, K. (2011). Human-inspired robotic grasp control with tactile sensing. *IEEE Transactions on Robotics*, 27, 1067–1079.
- Rühr, T., Sturm, J., Pangercic, D., Beetz, M., & Cremers, D. (2012). *A generalized framework for opening doors and drawers in kitchen environments*. In: ICRA.
- Rusu, R., Marton, Z., Blodow, N., Dolha, M., & Beetz, M. (2008). Towards 3D Point cloud based object maps for household environments. *Robotics and Autonomous Systems*, 56(11), 927–941.
- Salisbury, J. K. (1980). *Active stiffness control of a manipulator in cartesian coordinates*. In: Proceedings of the IEEE conference on decision and control.
- Schneider, A., Sturm, J., Stachniss, C., Reiser, M., Burkhardt, H., & Burgard, W. (2009). *Object identification with tactile sensors using bag-of-features*. In: IROS, IEEE.
- Shin S, Lee I, Lee H, Han G, Hong K, Yim S, Lee J, Park Y, Kang B, Ryoo D, et al. (2012). Haptic simulation of refrigerator door. In: *Haptics symposium (HAPTICS), 2012 IEEE, IEEE*, pp. 147–154.
- Sinapov, J., Bergquist, T., Schenck, C., Ohiri, U., Griffith, S., & Stoytchev, A. (2011). Interactive object recognition using proprio-

- ceptive and auditory feedback. *The International Journal of Robotics Research*, 30, 1250–1262.
- Sturm, J., Pradeep, V., Stachniss, C., Plagemann, C., Konolige, K., Burgard, W. (2009). *Learning kinematic models for articulated objects*. In: Proceedings of the international conference on artificial intelligence (IJCAI).
- Sturm, J., Jain, A., Stachniss, C., Kemp, C., & Burgard, W. (2010). *Operating articulated objects based on experience*. In: IROS.
- Sukhoy, V., Georgiev, V., Wegter, T., Sweidan, R., & Stoytchev, A. (2012). *Learning to slide a magnetic card through a card reader*. In: ICRA.
- Takamuku, S., Fukuda, A., & Hosoda, K. (2008). *Repetitive grasping with anthropomorphic skin-covered hand enables robust haptic recognition*. In: IROS.
- Torralba, A., Fergus, R., & Freeman, W. (2008). 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24, 1225.
- Waibel, M., Beetz, M., Civera, J., D'Andrea, R., Elfring, J., Galvez-Lopez, D., et al. (2011). Roboearth. *Robotics and Automation Magazine IEEE*, 18(2), 69–82.
- Weir, D., Pehkin, M., Colgate, J., Buttolo, P., Rankin, J., & Johnston, M. (2004). *The haptic profile: capturing the feel of switches*. In: HAPTICS.
- Wieland, S., Gonzalez-Aguirre, D., Vahrenkamp, N., Asfour, T., & Dillmann, R. (2009). *Combining force and visual feedback for physical interaction tasks in humanoid robots*. In: Humanoids.
- Wiste, T., Dalley, S., Varol, H., & Goldfarb, M. (2011). Design of a multigrasp transradial prosthesis. *Journal of Medical Devices*, 5(031), 009.
- Wyrobek, K., Berger, E., Van der Loos, H., & Salisbury, J. (2008). *Towards a personal robotics development platform: Rationale and design of an intrinsically safe personal robot*. In: ICRA.



Charles C. Kemp is an Assistant Professor at the Georgia Institute of Technology in the Wallace H. Coulter Department of Biomedical Engineering at the Georgia Institute of Technology and Emory University. He is also an Adjunct Assistant Professor in the School of Interactive Computing. He received a Ph.D. in Electrical Engineering and Computer Science (EECS) from the Massachusetts Institute of Technology (MIT) in 2005, an M.Eng. in EECS from MIT in 1998, and

a B.S. in Computer Science and Engineering from MIT in 1997. He is a member of the Center for Robotics and Intelligent Machines at Georgia Tech (RIM@GT) and the Health Systems Institute, which houses his lab, the Healthcare Robotics Lab. Charlie's research interests include autonomous mobile manipulation, human-robot interaction, assistive robotics, healthcare robotics, bio-inspired approaches to robotics, and AI.



Advait Jain is a Ph.D. student at the Georgia Institute of Technology in the School of Interactive Computing. He is a member of the Center for Robotics and Intelligent Machines at Georgia Tech (RIM@GT), the Health Systems Institute, and the Healthcare Robotics Lab. He received a Bachelor of Technology degree in Computer Science and Engineering from the Indian Institute of Technology (IIT), Delhi, India, in 2007.