Multimodal Tactile Perception of Objects in a Real Home

Tapomayukh Bhattacharjee[1], Henry M. Clever[2], Joshua Wade[3], and Charles C. Kemp[2]

***Abstract*—When operating in human environments, such as homes, robots could use tactile sensing to better perceive objects. A challenge that has not been sufficiently addressed is the influence of an object's surroundings on tactile perception. Prior research has focused on perception of objects in laboratory settings. Yet, a number of factors found in homes can affect multimodal tactile sensing. For example, the time-varying thermal characteristics of an object's surroundings, such as sunlight, HVAC, and refrigeration, can affect thermal sensing. Likewise, the placement of an object with respect to other objects and surfaces will affect force sensing and alter the way an object moves when pushed. In order to investigate these and other issues, we had a mobile robot reach out and push 47 different objects found in a real home over a three day period resulting in 1340 pushing episodes. We then characterized the performance of data-driven methods (k-NNs, SVMs, HMMs, and LSTMs) for a variety of tactile perception problems using the first two seconds of force, thermal, and motion sensing data collected by the robot. We paid particular attention to the ability of these methods to generalize what they have learned to different robot velocities, times of day, and object instances. Our results demonstrate the value of multimodal tactile sensing and data-driven methods for tactile perception from short-duration contact, and also illustrate the great diversity of real-world phenomena relevant to tactile sensing.**

***Index Terms*—Haptics and Haptic Interfaces; Force and Tactile Sensing; Perception for Grasping and Manipulation**

I. INTRODUCTION

INFERRING properties of objects or distinguishing objects in the real world by touch has many interesting challenges. In laboratory settings, we have previously seen encouraging results for haptic perception with force and motion sensing [1], as well as thermal sensing [2] separately. Our objective in this work is to use these multimodal sensing modalities together for inferring properties of objects in a real home. For this, we used data-driven approaches using signals from a multimodal sensing module at the end effector of a robot which made contact with in-situ objects in a real home where three people live. However, in a real home, factors like sunlight, HVAC, and refrigeration can affect thermal sensing, whereas an object's physical interaction with surfaces and other objects in its vicinity due to manipulation can affect force sensing. It can also change the object's movement pattern when pushed.

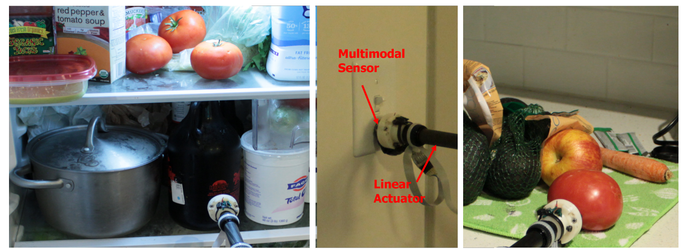

Fig. 1: *A mobile robot with a linear actuator and a multimodal tactile sensor attached at its end touching objects such as a bottle inside refrigerator, a light switch, and a fruit on a countertop.*

One of the important aspects of haptics is that sensing depends on action. For example, changing the velocity of a robot arm when pushing an object can alter the forces generated (force sensing) or the way the object moves (motion sensing). Collecting data at a different time of the day can affect thermal sensing because the same object can be at a different temperature. Our focus in this work is on generalizing the haptic perception performance across various robot velocities, times of day and object instances using force, thermal, and motion sensing modalities.

For haptic perception, we use various widely used and state-of-the-art data-driven methods. We use k-nearest neighbors (k-NNs), support vector machines (SVMs), hidden Markov models (HMMs), as well as long short-term memory networks (LSTMs). We want robots to infer various properties of an object. Inferring whether an object is soft or hard, or if the object moved during manipulation can help devise strategies such as, avoiding a hard object but pushing through a soft object which moved, to reach a goal [1]. Detecting an object which moved may also help in deciding if the robot should move the object to access a new location or avoid it [3]. We are also interested in material based haptic labels because knowing the material of an object is informative [2].

In addition to inferring object properties, we also want robots to distinguish various objects. To distinguish objects, we identified 50 tasks relevant to ADLs (Activities of daily living) and IADLs (Instrumental activities of daily living) such as 'putting a *towel* on *rack*', 'fetching a *bottle* from *refrigerator*', 'pushing a *door handle* on *door surface*' to open the door, etc. Thus, conditioned on the task, we want robots to distinguish a target object (tactile foreground) from a background object in its vicinity (tactile background) that the robot may come in contact with while performing the task. We are also interested in analyzing if robots can distinguish different object parts given

Manuscript received: September, 10, 2017; Revised January, 3, 2018; Accepted February, 10, 2018.

This paper was recommended for publication by Editors Antonio Bicchi and Yasuyoshi Yokokohji upon evaluation of the Associate Editor and Reviewers' comments. This work was supported in part by NSF Awards EFRI-1137229 and IIS-1150157, the National Institute on Disability, Independent Living, and Rehabilitation Research (NIDILRR) grant 90RE5016-01-00 via RERC TechSAge, and a Google Faculty Research Award

[1]T. Bhattacharjee is with the Paul G. Allen School of Computer Science and Engineering, University of Washington tapo@cs.washington.edu

[2]H. M. Clever, and C. C. Kemp are with the Healthcare Robotics Lab, Institute for Robotics and Intelligent Machines, Georgia Institute of Technology

[3]J. Wade is a commissioned officer in the US Navy

Work was done at the Healthcare Robotics Lab, Institute for Robotics and Intelligent Machines, Georgia Institute of Technology

Digital Object Identifier (DOI): see top of this page.

an object (such as 'chair cushion' given chair, or 'mattress' given bed).

II. Related Work

To the best of our knowledge, robotic tactile perception of real home environments remains an unexplored area of research. Researchers used data collected in households for analyzing grasping [4] by humans and organization behaviors in the kitchen [5]. Researchers also developed data acquisition devices to collect multimodal data from objects in real home bathrooms [6] and sample surfaces [7]. Gandhi et al. deployed unmanned aerial vehicles to crash into objects in indoor academic campus environments to learn a navigation policy [8]. Most of the related work on haptic perception has focused on using exploratory probing behaviors (squeezing, tapping, sliding, etc.) to extract information from contact with objects in laboratory settings (see Sections II-A and II-B). In this work, we focus on extracting information using a simple linear motion with a linear actuator on a mobile robot pushing objects in a real home with two seconds of contact. In addition, few related studies have looked into generalizing the haptic perception performance to new robot and environment conditions which are different from the conditions used during training. For this paper, we focus on related work with force and thermal sensing modalities.

A. Single Modality : Force or Thermal

Researchers have used force sensing to classify objects based on their material property, shape property, functional property etc. [1], [9], [10], [10]–[14]. For a more detailed survey, refer to [1], [3]. Though thermal sensing is relatively unexplored in robotics compared to force sensing, there have also been studies on hardware development [15], [16] and material recognition using only thermal sensing. For a detailed overview of such material recognition studies, please refer to [2], [6], [17], [18].

B. Multiple Modalities : Force and Thermal

1) Hardware Development: Researchers developed multi-modal sensors using various techniques such as a capacitive and thermal sensor [16], conductive rubber based force sensors and temperature sensors [19], a polymer-based tactile and thermal sensor [20], capacitive tactile sensors with temperature-dependent semiconductors [21], and a tactile and thermal sensor using single pressure-conductive rubber sheet [22].

2) Object Recognition: Researchers have also used force and thermal sensing together for object recognition. Using force and thermal sensing, Takamuku et al. [23] successfully classified 5 materials and Engel et al. [24], [25] achieved 90% accuracy over 50 trials for recognizing 5 materials. Caldwell et al. [26] used force and thermal sensing to infer texture, stiffness and object profile, temperature and thermal properties of 7 materials using exploratory behaviors. Xu et al. [27] used multimodal sensor feedback to identify various materials with a BioTAC sensor [28]. Chu et al. [28] attached two BioTAC sensors to a robotic gripper and assigned 24 adjectives to 60 objects using four exploratory behaviors. They used discrete HMMs to construct a feature vector of likelihoods and used binary SVM classifiers on the feature vector for classification [29]. Schmitz et al. [30] used power grasping of objects and multiple modalities for object recognition with deep learning. Hoelscher et al. [31] used both thermal and force features for recognizing 49 objects and found thermal features to be the most informative.

III. Relevant Haptic Labels

A. Inferring Object Properties

1) Compliance and Mobility: We want robots to infer an object's compliance and mobility through physical interaction. Specifically, we perform three kinds of classification tasks: (a) 'Hard' vs. 'Soft', (b) 'Moved' vs. 'Unmoved', and (c) 'Hard-Unmoved' vs. 'Soft-Unmoved' vs. 'Hard-Moved' vs. 'Soft-Moved'.

2) Material: We are also interested in material property based haptic labels. We sampled 67 object parts relevant to ADLs and IADLs in the bedroom, kitchen and bathroom of a household. These object parts are made of 14 different materials. Our objective is to perform binary classification between every binary combination of these 14 materials (inspired by 'tactile foreground' and 'tactile background' [6]) thus leading to 91 such comparisons. The 14 materials are 'Foam', 'Aluminum', 'Wood', 'Steel', 'Dry Wall', 'Plastic', 'Glass', 'Paper', 'Porcelain', 'Granite', 'Cardboard', 'Fiberglass', 'Vegetable Matter', and 'Fabric'.

B. Distinguishing Objects

1) Tactile Foreground vs. Tactile Background: During a manipulation task while moving a target object (tactile foreground), a robot may come in contact with another object (tactile background) in its immediate vicinity [6]. We have identified 50 such comparisons for the objects relevant to ADLs and IADLs in a household environment, such as 'rack vs. towel', 'book cover vs. book spine', 'TV vs. TV remote', 'kitchen faucet vs. kitchen backsplash', 'chair cushion vs. chair frame', 'socket vs dry wall' etc. Our objective is to see if our algorithms can distinguish objects in each of the 50 foreground-background pairs.

2) Object Parts given Object: Finally, we analyze if robots can distinguish different object parts given an object. This could be useful in assistive manipulation scenarios such as assisting a human sitting on a chair or lying in a bed. Knowing whether the robot made contact with 'chair cushion or chair frame given its manipulating near a chair' could be useful in devising manipulation or control strategies used by the robot. Similarly, knowing if the robot is in 'contact with bed frame or the mattress or pillow on the bed given its performing a task near the bed' could provide relevant and useful information to the robot's manipulation task. We identified 13 such comparisons involving two to three object parts given an object in the bedroom, kitchen, and bathroom of a household.

IV. Generalization Tasks

We focus on generalizing the classification results using force and thermal sensing modalities across different velocities, times of day, and instances.

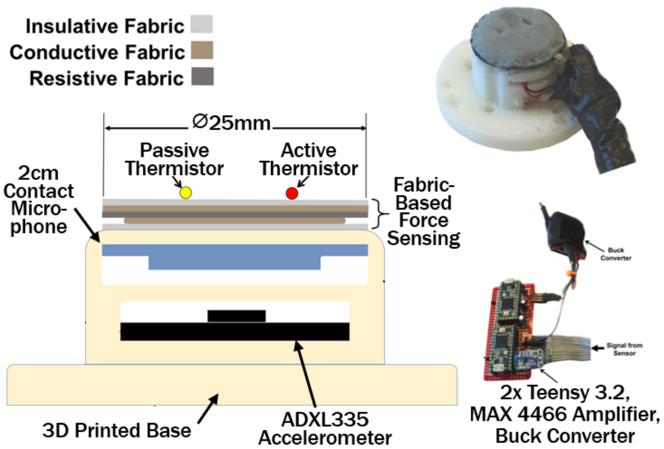

Fig. 2: *Multimodal sensor module and Teensy 3.2 microcontroller Hardware for analog-to-digital conversion (ADC) and buck converter to power active thermal sensor.*

A. Generalizing Across Velocity

We want to see if results obtained using data-driven methods trained on data collected at one velocity from various objects can generalize to data collected from the same set of objects at another velocity. We varied the robot-arm velocity to two distinct velocities.

B. Generalizing Across Time of Day

We want to see if results obtained using data-driven methods trained on data collected at one time of a day could generalize to data collected at another time on another day. We collected the data from the same set of objects at two different times over a period of three days.

C. Generalizing Across Instance

We want to see if results obtained using data-driven methods trained on a set of objects in one specific category with a particular velocity can generalize to another object in the same category with the same velocity. We focused this analysis on the 'Hard vs. Soft' and 'Moved vs. Unmoved' classification problems. We used a 'leave-one-object-out' cross-validation scheme to test for generalization, meaning that data from the same object does not overlap both training and testing sets.

V. METHODS

A. Experimental Setup

Our experimental setup consists of a multimodal sensor attached at the end of a linear actuator on a mobile robot (a 2-DOF Parallax Arlo robot base), and the real-home environment which includes the objects.

1) The Multimodal Sensor: Figure 2 shows the multimodal tactile sensor module with force, audio, acceleration, active and passive thermal sensing modalities. Individual sensors are attached to a 3D printed base. Two Teensy 3.2 microcontrollers and a buck converter on a separate module provide power to and read data from the sensors using the microcontrollers' built in analog-to-digital converters (ADCs). We did not perform any analysis with the audio and acceleration sensing modalities in this paper.

The fabric-based force sensor is based on the design in [32]. It reads force data at 1kHz. We used a voltage dividing reference resistor, $R_{\mathrm{ref}} = 1\,\mathrm{k}\Omega$. Because the fabric-based electrodes are small compared to the touched objects, we assumed the contact would cover the entire sensor's area. In this work we used the fabric-based force sensor to detect the start of contact with a force threshold of 0.1 N.

The active thermal sensor uses a self-heated $10\,\mathrm{k}\Omega$ B57541G1103F NTC thermistor [33], with adjustable voltage input from a digitally-controlled buck converter (see Fig. 2). We use a closed-loop temperature controller to heat the thermistor to $55\,^{\circ}\mathrm{C}$ prior to contact with the object. The sensor reads data at 100 Hz. Once the fabric-based force sensor detects contact, this closed-loop temperature controller turns off by holding the desired buck converter voltage, V_{des}, constant. This is necessary to ensure that the temperature controller does not cancel potentially-informative temperature changes resulting from contact.

The passive thermal sensor uses a second $10\,\mathrm{k}\Omega$ B57541G1103F NTC thermistor [33] similar to the active thermal sensor reading data at 100 Hz. The passive thermal sensor is powered by a constant 3.3 V input from the microcontroller rather than a higher voltage buck converter.

2) The Environment: The environment was the household of one of the co-authors of this paper where three people live. It consisted of objects in the bedroom, bathroom, and the kitchen of the household, relevant to ADLs and IADLs. There were 12 objects with 22 object parts in the bedroom, 9 objects with 15 object parts in the bathroom, and 26 objects with 30 object parts in the kitchen, as shown in Fig. 3.

B. Experimental Procedure

Our experimental procedure consisted of a teleoperation phase and an autonomous phase. First, in the teleoperation phase, we used a joystick controller to manually position the mobile robot at a location feasible for reaching a specific object before starting the trials. Then in the autonomous phase, for each experimental trial, the robot autonomously reached towards the object to push it. We define a trial as a single manipulation behavior of linear actuator extending and pushing an object and then bringing the actuator back. We set the robot-arm linear actuator to move at a specific velocity ('slow : 3 cm/s' or 'fast : 6 cm/s') until it reached a contact force threshold of 5 N. We programmed the robot-arm to be in contact with the object for 5 seconds or until the robot arm pushes the object for 2 cm, whichever is earlier. After the contact duration, the linear actuator fully retracted and waited for 20 seconds for the thermal sensors in the multimodal sensor module to reach a consistent initial condition prior to beginning the next trial. Note that for objects with large surface area, we sampled the space such that at every trial, the robot collected data from a different part of the object surface. For objects with small surface area, we collected data from the same spot over multiple trials. The waiting period of 20 seconds also ensured that the object was at the same initial temperature for the next trial.

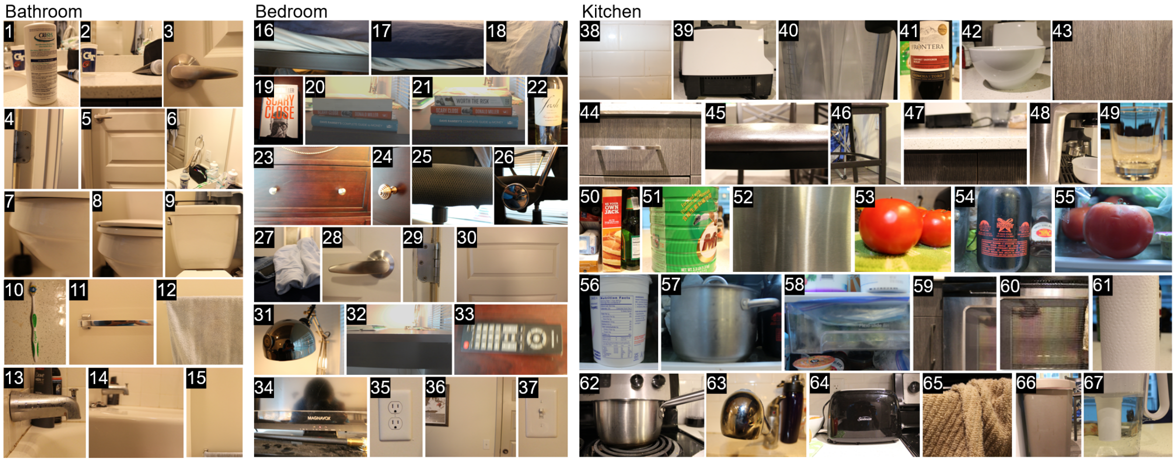

Fig. 3: *Objects touched in Bathroom: 1. Bottle 2. Countertop 3. Door Knob 4. Door Hinge 5. Door Surface 6. Mirror 7. Toilet Bowl 8. Toilet Seat 9. Toilet Tank 10. Toothbrush 11. Towel Rack Frame 12. Towel Rack Towel 13. Tub Faucet 14. Tub Wall 15. Wall. Objects Touched in Bedroom: 16. Bed Frame 17. Bed Mattress 18. Bed Pillow 19. Book Cover 20. Book Pages 21. Book Spine 22. Bottle 23. Cabinet Drawer 24. Cabinet Knob 25. Chair Cushion 26. Chair Frame 27. Clothes 28. Door Knob 29. Door Hinge 30. Door Surface 31. Lamp 32. Table 33. TV Remote 34. TV 35. Wall Socket 36. Wall Surface 37. Wall Switch. Objects touched in kitchen: 38. Backsplash 39. Blender Base 40. Blender Container 41. Bottle 42. Bowl 43. Cabinet Surface 44. Cabinet Handle 45. Chair Cushion 46. Chair Frame 47. Countertop 48. Expresso Machine 49. Cup 50. Food Box 51. Food Can 52. Fridge 53. Fruit 54. Fridge Bottle 55. Fridge Fruit 56. Fridge Jug 57. Fridge Pot 58. Fridge Surface 59. Oven Frame 60. Oven Window 61. Paper Towel 62. Pot 63. Sink Faucet 64. Toaster 65. Towel 66. Trash can 67. Water Pitcher.*

C. Data Collection and Preprocessing

We collected data from all the objects over a period of three days. Our first day's data collection started at 11.50 AM and continued until 6.50 PM, the second day from 9 AM until 10.45 PM, and the third day from 9.40 AM until 10.40 PM. During the data collection, we collected data from the same object at different times of the day (either 'Morning - early Afternoon' or 'Afternoon - late Night').

For each object part at one specific time, once we manually positioned the robot's mobile base at a particular location, the robot autonomously reached to touch a given object 5 times with 'slow' velocity (3 cm/s) and then 5 times with 'fast' velocity (6 cm/s). If the object moved during 'slow' velocity trials, we reset the object back to the original position before the 'fast' velocity trials. We did this to ensure uniform initial conditions for the first trial in each of the 'slow' and 'fast' group of trials. We collected 20 trials (5 trials x 2 velocities x 2 times) for each of the 67 object parts, totaling 1340 trials. For each, we collected time-series vectors of force (f), motion (relative position of the robot arm after onset of contact) (m), active thermal (h), and passive thermal (t) modalities. We sampled force and motion signals at 1000 Hz to match the frequency of active and passive thermal signals (100 Hz). Note, when we used combinations of features, we normalized each feature by subtracting the mean and dividing by the variance of the feature across all the data.

Two experimenters (co-authors) independently labeled the objects for material-based labels, 'hard' or 'soft' labels, and 'moved' or 'unmoved' labels. We convened after the experiments and found no disagreements. For material-based labels, we noted both the 'surface material' (i.e. paper label) and the 'dominant material' (i.e. glass bottle) for each object, but used the dominant material label for analysis. We also measured the stiffness of the objects using a compression spring. The compression spring was unable to measure the stiffness for very hard objects with $k > 100,000$ N/m. Both the experimenters tended to classify objects with $k < 10,000$ N/m as 'soft' and objects with $k > 10,000$ N/m as 'hard'. The label ratio was skewed towards hard objects, at 83%, and towards unmoved objects, at 75%. Note, this is a property of both the home and the objects (objects relevant to ADLs and IADLs in the bedroom, bathroom, and kitchen of that specific home).

D. Data-driven Haptic Perception

We used four data-driven algorithms for classifying the objects according to haptic labels. We selected these data-driven machine-learning algorithms based on their wide usage and suitability for time series. We used various combinations of the force, motion, active and passive thermal time-series vectors as the features for performance analysis. For each algorithm, we used the same set of corresponding parameters with which we found success after exhaustive grid search in their corresponding parameter space in [2], [3] and [6].

1) Support Vector Machines: We implemented binary support vector machines (SVMs) using the scikit-learn package [34] in Python. We used a linear kernel. To produce feature vectors for training, we truncated each feature to 2.0 seconds after the onset of contact (detected using the fabric-based force sensor). For the active thermal modality, we estimated the slope of the raw temperature data using first central difference. We concatenated the features to obtain a $200*(n+1)$-dimensional feature vector where 'n' is the number of sensing modalities. One feature vector corresponds to the estimated slope of the active heat-transfer data.

2) k-Nearest Neighbors: We also implemented 1-nearest neighbor (k-NN with k=1) using the scikit-learn package [34] in Python. We used the same concatenated feature vector. To reduce the effect of noise and overfitting, we computed a low-dimensional representation of the training data with principal component analysis (PCA) before classification with 1-NN [1], using three principal components for dimensionality reduction. Three principal components could account for more than 95% of the variance of the data.

TABLE I: Variability in the Mechanics of Interactions

Object	Slide	Slide + Contact	Roll/ Tip	Does Not Move
Blender Container	0	1	0	19
Chair Cushion (Bedroom)	0	0	5	15
Chair Cushion (Kitchen)	0	0	1	19
Food Box	9	0	11	0
Fruit (on Countertop)	18	2	0	0
Bottle (in Fridge)	0	12	0	8
Fruit (in Fridge)	11	5	4	0
Jug (in Fridge)	0	3	0	17
Lamp	15	0	5	0
Pot	13	0	6	1
Toaster	0	5	0	15
Water Pitcher	15	0	0	5

3) Hidden Markov Models: We used a multivariate continuous left-right HMM with 10 hidden states and $n{+}1$ dimensional Gaussian emissions for each category using the GHMM toolkit [35] in Python. With the same truncated features as Section V-D1, we set a uniform prior to all the states. We used a spherical covariance matrix structure for initialization. We chose these specifications based on our previous results [1], [11]. We trained these HMMs with the standard Baum-Welch algorithm. For testing, we ran the Viterbi algorithm to find the HMM with the most probable state sequence given the observations, classifying the category as being the category associated with this HMM [36].

4) Long Short-term Memory Networks: Finally, we also implemented long short-term memory (LSTM) networks [37] with the same truncated features as in Section V-D1. In this paper, we used an LSTM structure where each memory cell has an input gate, a forget gate, and an output gate. We implemented a stacked-LSTM structure of 2 layers with 50 cells each. We also added a dropout layer in between the two layers, which helps in regularization. We added a dense output layer which was fully connected. The LSTM has a total of 31,004 parameters. We initialized the parameters with uniform distribution, used 'softsign' activation functions for the hidden layers, and 'softmax' activation function [38] for the fully connected output layer. Our dropout probability was 0.2. We used 'RMSprop' [38] as the optimizer and 'categorical_crossentropy' [38] as our loss function because our task is a classification task. We used 'MinMaxScaler' function [38] to scale multivariate features for LSTMs. These parameters and choice of functions are similar to our implementation in [3] with which we found the best results after an exhaustive search in the parameter space.

VI. VARIABILITY IN THE REAL WORLD

In this section, we characterize real-world variability that could affect tactile perception. For example, due to various sources of real-world variability, force, motion and thermal signals from the same object can differ considerably across different trials which poses a challenge for tactile perception. We characterize thermal variability by qualitative observations in the real home environment. We characterize variability in mechanics of interactions by analyzing all experimental videos for the different types of object movement following robot-object interactions. Figure 4 shows examples of varied multimodal data from the robot touching some objects.

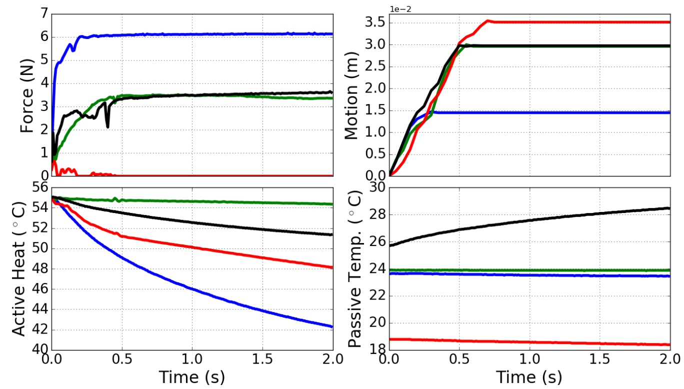

Fig. 4: *This figure shows force, motion, active heat, and passive temperature data from four objects (wall switch (blue), fruits inside fridge (red), fruits outside (green), and heated lamp (black)). For hard and unmoved objects, the magnitude of forces go up higher than softer objects or objects which moved. Note, the fruit inside the refrigerator rolled and eventually the fruit broke contact with the sensor. Also, the heat-transfer from fruit is different when outside vs. when inside refrigerator (different initial temperature). Similarly, the heat-transfer from heated lamp (made of aluminum) is different from that of an aluminum block at normal room temperature [2].*

A. Thermal Variability

We collected data from an object at different times of the day. This often resulted in an object having different initial temperatures across trials. For example, one of the sets of data captured in the bathroom was immediately following a morning shower of a resident. This caused all surfaces of the bathroom to be hotter than usual, many of them wet. The towel was non-uniformly damp and the faucet, that the robot pushed, had formed a layer of condensation. Another example of thermal variability was due to the heated lamp and the refrigerator in the kitchen. The same fruit (tomato) outside had a different heat-transfer than when kept inside the refrigerator. It also contained objects that formed an increasing deposit of condensation and became warmer the longer the refrigerator door was kept open. Yet another thermal variation was the change of sunlight throughout the day between data collecting sessions on a particular object or object part.

B. Variability in Mechanics of Interactions

While some physical interactions such as the robot pushing a wall could be considered predictable, others such as pushing a stove pot showed interesting phenomena that would be nontrivial to capture in a laboratory setting. Note, the linear actuator always pushed (approached) the object in a direction perpendicular to the object's surface, but the resulting movement of the object depended on its immediate environment. We found that 12/67 object parts showed variability in the mechanics of interactions across trials. Table I provides a characterization of mechanical variability, defined by the following movements:

- Slide : Movement in a direction parallel to the linear actuator's movement direction and normal to gravity
- Slide + Contact : Movement direction normal to gravity but not necessarily parallel to the direction of linear actuator movement due to contact with other objects in the vicinity
- Roll / Tip : Movement when the object rolls, tips over, appears at risk for toppling or falling suddenly, or twists / bends without making contact with other objects.

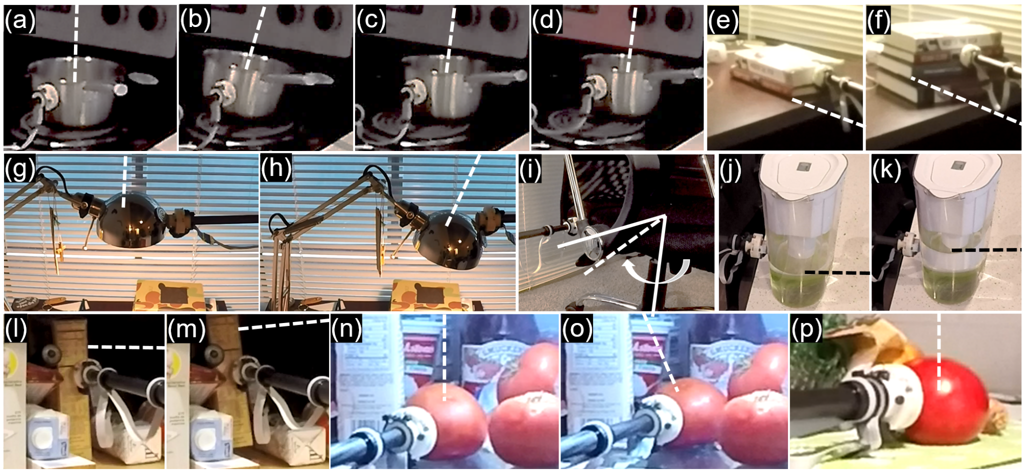

Fig. 5: *Variability in mechanics of interactions. Dashed line shows the normal to an object's orientation. The pot (a) slid, (b) rolled/tipped due to geometrical features, (c) tipped at the edge of burner, and (d) did not move due to the backboard. Stack of books slid on (e) wooden surface, and (f) on other book surface. A lamp (g) moved backwards and (h) tipped down due to its internal mechanics. (i) A swivel chair twisted about a center point. (j,k) A pitcher did not move when water level increased. A food box (l) slid back in a kitchen cabinet and (m) rolled/tipped due to another object. A fruit in the fridge (n) slid and (o) rolled/tipped due to contact with another object. (p) A fruit on kitchen countertop slid on a towel and changed direction due to contact with a carrot behind it.*

- Does Not Move : No movement

Fig. 5 (a)-(f) detail examples of notable pushing phenomena resulting from contact interactions with the underlying surface. For example, the pot sometimes tipped on the electric stove burner and sometimes slid. For the tipping cases, four times the pot snagged on some geometric feature of the burner, and two times the pot tipped off of the burner. Other examples include motion of books on a bedside table, and pushing a blender container and toaster until they reached a wall behind the counter and no longer moved.

Examples of mechanical variability in the object itself are shown in Fig. 5 (g)-(k). The tensioning springs in the balanced-arm lamp appeared to affect how the lamp moved when pushed. Similarly, the presence of ball bearings in a swivel chair caused it to twist upon pushing the cushion. Another example is the water level in a water pitcher. On different days, the pitcher had different levels of water in it. Sometimes, the pitcher could not be moved by the robot's motion with slow velocity when the water level was high.

In Fig. 5 (l)-(p), we consider variability from contacting other objects. During some pushes of a food box, it slid back normally. In others, the food box tipped over a short object behind it. When pushing a fruit in a fridge, the fruit usually slid back normal to gravity, but sometimes changed direction due to contact with other objects or rolled / tipped due to the underlying surface. The fruit on the kitchen countertop never tipped/rolled during a push. It rested on a high friction surface (a rough towel). The bottle and jug in the fridge also showed variability from contacting objects during pushes.

VII. RESULTS

We present results of haptic classification with various combinations of force (f), motion (m), active thermal (h), and passive thermal (t) features. For generalization across velocity, we trained on one velocity and tested on another (across both the times of day) and vice-versa and reported the average percentage accuracy. For generalization across time of day, we trained on one time of day and tested on another (across both the velocities) and vice-versa and reported the average percentage accuracy. For generalization across instance, we used the 'leave-one-object-out' cross-validation scheme. We did not implement LSTMs for binary material recognition tasks and for distinguishing objects as the amount of training data was less. Table II shows the best average results.

A. Compliance and Mobility

In this section, we present the classification results for 'Hard vs. Soft' (HS), 'Moved vs. Unmoved' (MU) and 4-category 'Hard Unmoved vs. Hard Moved vs. Soft Unmoved vs. Soft Moved' (4C) labels.

1) Generalizing Across Velocity: Irrespective of the type of categorization, multivariate HMMs usually showed the best results (Table II) whereas 1-NN with single features showed the worst results (61% for HS, 23% for MU). This is similar to our results with laboratory objects [1], [3], showing that the results from these algorithms can apply to situations with different mechanical interactions. Univariate HMMs with temperature feature also showed the worst result (13% for 4C) as temperature modality alone did not sufficiently inform the compliance or mobility of an object. LSTMs showed promise (84% for HS, 92% for MU, and 81% for 4C).

2) Generalizing Across Time of Day: Irrespective of the type of categorization, SVMs with both force and motion features gave the best results (Table II). Similar to the generalization results for velocity, the worst results were for 1-NN with single features (54% for MU, 27% for 4C). Univariate HMMs with heat-transfer feature also performed the worst (50% for HS). This is intuitive because we do not expect heat-transfer feature to be informative about the compliance or mobility of an object. LSTMs again showed promise (84% for HS, 92% for MU, and 81% for 4C).

3) Generalizing Across Instance: Again, irrespective of whether its the 'Hard vs. Soft' or 'Moved vs. Unmoved' categorization, multivariate HMMs performed the best (Table II) and 1-NNs performed the worst (67% for HS, 27% for 4C). Interestingly for 'Moved vs. Unmoved', LSTMs did the worst

TABLE II: Best Algorithm Performance with 2.0s of Contact

Categories	Generalize Across	Algorithm	Features	Accuracy (%)
Hard vs. Soft	Velocity	HMM	f+m+h+t	**86**
	Time	SVM	f+m	**87**
	Instance	HMM	f+m+h+t	**92**
			Maj. Classifier	83
			Rand. Guess	50
Moved vs. Unmoved	Velocity	HMM	f+m+h+t	**93**
	Time	SVM	f+m	**94**
	Instance	HMM	f+m	**100**
			Maj. Classifier	76
			Rand. Guess	50
Hard Unmoved vs. Hard Moved vs. Soft Unmoved vs. Soft Moved	Velocity	HMM	f+m+h+t	**83**
	Time	SVM	f+m	**84**
			Maj. Classifier	66
			Rand. Guess	25
Material Recognition	Velocity	SVM	f+m+h+t	**82**
	Time	SVM	f+m+h+t	**86**
			Maj. Classifier	50
			Rand. Guess	50
Tactile Foreground vs. Tactile Background	Velocity	SVM	f+m+h+t	**91**
	Time	SVM	f+m+h+t	**83**
			Maj. Classifier	50
			Rand. Guess	50
Object Parts given Object	Velocity	SVM	f+m+h+t	**78**
	Time	SVM	f+h+t	**71**
			Maj. Classifier	42
			Rand. Guess	42

Note 'f' = force, 'm' = motion, 'h' = active thermal (heat-transfer), and 't' = passive thermal (temperature) feature. 'Maj. Classifier' is Majority Classifier and 'Rand. Guess' is Random Guess Classifier.

with just the motion feature (53%), probably because of the lack of sufficient data.

B. Material

1) Generalizing Across Velocity: SVMs with all the features of force, motion, heat-transfer, and temperature showed the best results (Table II). This showed that all these modalities with complementary information were informative for getting material specific information. The worst results were with force or temperature features as these features alone could not extract relevant information for material recognition.

2) Generalizing Across Time of Day: Again, SVM with all the features gave the best results (Table II) which shows the relevance of all the features for material recognition problems. And again, univariate HMMs (60%) and 1-NNs (56%) gave the worst performance with just force and temperature features, probably because of similar reasons outlined above.

C. Tactile Foreground vs. Tactile Background

1) Generalizing Across Velocity: SVMs with all features showed the best results (Table II) for this recognition task. Similar to the binary material recognition results in Section VII-B, these showed the importance of all the features in recognizing two objects. The worst results were for 1-NN with temperature feature alone (55%). This is because passive temperature sensing is a function of the ambient temperature and not any object characteristics. This showed that heat-transfer modality (active thermal sensing) was an informative feature for object categorization task, which agrees with our previous result in [6].

2) Generalizing Across Time of Day: Here, again SVMs with all the features gave the best results (Table II) whereas 1-NNs with temperature feature (55%) and HMMs with force and motion features (55%) performed the worst.

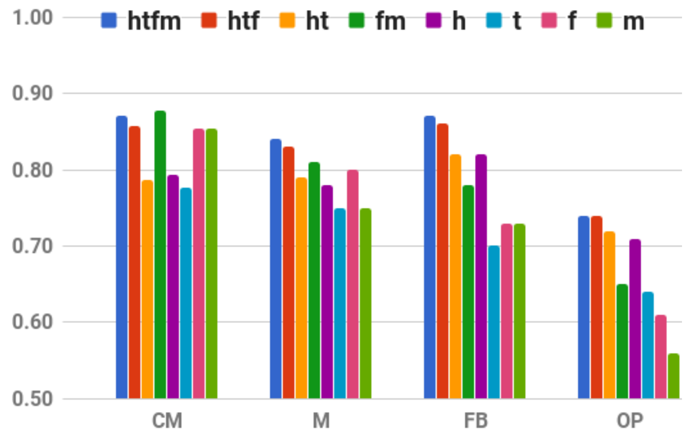

Fig. 6: *This figure shows the overall results. Here, 'CM' = Compliance and Mobility based haptic labels, 'M' = Material-based haptic labels, 'FB' = Tactile Foreground vs. Tactile Background, and 'OP' = Distinguishing Object Parts given Object. Each bar in the chart represents the average best performance using a particular feature combination across all data-driven algorithms and across all generalization schemes (time of day, velocity, and instance (if available)). For the features, 'f' = force data, 'm' = motion data, 'h' = active heat-transfer data, and 't' = passive temperature data.*

D. Object Parts given Object

1) Generalizing Across Velocity: SVMs with all the features showed the best results (Table II) showing the relevance of all the features to recognize an object. 1-NN with heat-transfer and temperature feature (47%) showed the worst results. This is probably because without using complementary modalities of force and thermal sensing, it was difficult to extract useful information.

2) Generalizing Across Time of Day: SVMs with all the features or using features that captured the complementarity of force and thermal sensing performed the best (Table II), whereas 1-NNs with just the force and motion features (44%) gave the worst results.

E. Overall Results

Figure 6 shows the overall results. Irrespective of algorithms used, on average, force and motion features were especially informative for classifying compliance and mobility based haptic labels. All four features (force, motion, active heat, and passive temperature) features were informative for binary material recognition, distinguishing tactile foreground vs tactile background, and distinguishing object parts given object. This shows that the complementary information from force, motion, and thermal sensing modalities are informative for inferring material properties as well as distinguishing objects. However, passive temperature modality did not seem to be very informative for these tasks.

VIII. CONCLUSIONS

We demonstrated the usefulness of the complementary capabilities of force and thermal sensing modalities for haptic perception tasks in a real home. Specifically, we were interested in inferring compliance and mobility based haptic labels, material based haptic labels as well as distinguishing objects in relevant tasks. We identified various sources of variability in haptic perception with objects in the real world due to its surroundings and the environment. Our objective was to analyze if the haptic perception results from these objects in the real world can generalize to different speeds (which may affect force sensing) or different times of day (which may affect thermal sensing) or across different object instances. We collected data using a mobile robot with a multimodal sensor module attached at the end of a linear actuator from 67 object parts in the bedroom, bathroom, and kitchen of a house over a period of three days. We implemented widely used and state-of-the-art data-driven algorithms such as 1-NNs, SVMs, HMMs, and LSTMs for generalization tasks. Our results showed the importance of using features from multiple sensing modalities in inferring haptic properties of objects or distinguishing objects. SVMs and HMMs showed the best results for different haptic label based classification tasks whereas 1-NNs failed to generalize. LSTMs showed promise for our problem but may need large amounts of data to give meaningful results.

REFERENCES

[1] T. Bhattacharjee, J. M. Rehg, and C. C. Kemp, "Haptic classification and recognition of objects using a tactile sensing forearm," in *IEEE International Conference on Intelligent Robots and Systems (IROS)*, October 2012, pp. 4090 – 4097.

[2] T. Bhattacharjee, J. Wade, and C. Kemp, "Material recognition from heat transfer given varying initial conditions and short-duration contact," in *Proceedings of Robotics: Science and Systems*, Rome, Italy, July 2015.

[3] T. Bhattacharjee, J. M. Rehg, and C. C. Kemp, "Inferring object properties with a tactile-sensing array given varying joint stiffness and velocity," *International Journal of Humanoid Robotics*, pp. 1–32, 2017.

[4] I. M. Bullock, T. Feix, and A. M. Dollar, "The yale human grasping dataset: Grasp, object, and task data in household and machine shop environments." *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 251–255, 2015.

[5] E. Cha, J. Forlizzi, and S. S. Srinivasa, "Robots in the home: Qualitative and quanitative insights into kitchen organization." in *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, 2015, pp. 319–326.

[6] T. Bhattacharjee, J. Wade, Y. Chitalia, and C. C. Kemp, "Data-driven thermal recognition of contact with people and objects," in *IEEE Haptics Symposium (HAPTICS)*. IEEE, 2016, pp. 297–304.

[7] A. Burka, S. Hu, S. Helgeson, S. Krishnan, Y. Gao, L. A. Hendricks, T. Darrell, and K. J. Kuchenbecker, "Proton: A visuo-haptic data acquisition system for robotic learning of surface properties," in *Multisensor Fusion and Integration for Intelligent Systems (MFI), 2016 IEEE International Conference on*. IEEE, 2016, pp. 58–65.

[8] D. Gandhi, L. Pinto, and A. Gupta, "Learning to fly by crashing." *arXiv preprint arXiv:1704.05588*, 2017.

[9] A. Drimus, G. Kootstra, A. Bilberg, and D. Kragic, "Classification of rigid and deformable objects using a novel tactile sensor," in *in Proceedings of 15th International Conference on Advanced Robotics (ICAR)*, 2011, pp. 427–434.

[10] V. Sukhoy, R. Sahai, J. Sinapov, and A. Stoytchev, "Vibrotactile recognition of surface textures by a humanoid robot," in *Proceedings of the 2009 Humanoids Workshop: Tactile Sensing in Humanoids Tactile Sensors and Beyond*, 2009.

[11] 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 International Conference on Humanoid Robots (Humanoids)*, October 2013.

[12] 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 *Proceedings of the 3rd IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) Workshop on Robots in Clutter: Perception and Interaction in Clutter*, 2014.

[13] 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 *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 2015, pp. 5073–5080.

[14] K. Matheus and A. M. Dollar, "Benchmarking grasping and manipulation: Properties of the objects of daily living," in *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, 2010, pp. 5020–5027.

[15] R. A. Russell, "Thermal sensor for object shape and material constitution," *Robotica*, vol. 6, no. 01, pp. 31–34, 1988.

[16] D. Siegel, I. Garabieta, and J. M. Hollerbach, "An integrated tactile and thermal sensor," in *IEEE International Conference on Robotics and Automation (ICRA)*, vol. 3. IEEE, 1986, pp. 1286–1291.

[17] H. R. Nicholls and M. H. Lee, "A survey of robot tactile sensing technology," *The International Journal of Robotics Research*, vol. 8, no. 3, pp. 3–30, 1989.

[18] M. H. Lee and H. R. Nicholls, "Review article tactile sensing for mechatronicsa state of the art survey," *Mechatronics*, vol. 9, no. 1, pp. 1–31, 1999.

[19] Y. J. Yang, M. Y. Cheng, S. C. Shih, X. H. Huang, C. M. Tsao, F. Y. Chang, and K. C. Fan, "A 32x32 temperature and tactile sensing array using pi-copper films," *The International Journal of Advanced Manufacturing Technology*, vol. 46, no. 9-12, pp. 945–956, 2010.

[20] P. Dario, D. De Rossi, C. Domenici, and R. Francesconi, "Ferroelectric polymer tactile sensors with anthropomorphic features," in *IEEE International Conference on Robotics and Automation*, vol. 1. IEEE, 1984, pp. 332–340.

[21] F. Castelli, "An integrated tactile-thermal robot sensor with capacitive tactile array," in *IEEE Industry Applications Conference*, vol. 3. IEEE, 1995, pp. 1970–1975.

[22] J. I. Yuji and K. Shida, "A new multifunctional tactile sensing technique by selective data processing," *IEEE Transactions on Instrumentation and Measurement*, vol. 49, no. 5, pp. 1091–1094, 2000.

[23] S. Takamuku, T. Iwase, and K. Hosoda, "Robust material discrimination by a soft anthropomorphic finger with tactile and thermal sense," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2008, pp. 3977–3982.

[24] J. Engel, N. Chen, C. Tucker, C. Liu, S. Kim, and D. Jones, "Flexible multimodal tactile sensing system for object identification," in *5th IEEE Conference on Sensors*. IEEE, 2006, pp. 563–566.

[25] J. Engel, J. Chen, Z. Fan, and C. Liu, "Polymer micromachined multimodal tactile sensors," *Sensors and Actuators A: Physical*, vol. 117, no. 1, pp. 50–61, 2005.

[26] D. G. Caldwell and J. O. Gray, "Dynamic multi-functional tactile sensing," in *RoManSy 9*. Springer, 1993, pp. 187–198.

[27] D. Xu, G. E. Loeb, and J. A. Fishel, "Tactile identification of objects using bayesian exploration," in *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2013, pp. 3056–3061.

[28] *Syntouch BioTac Sensor*, Syntouch Inc. [Online]. Available: https://www.syntouchinc.com/wp-content/uploads/2017/01/BioTac_Product_Manual.pdf

[29] 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," *Robotics and Autonomous Systems*, vol. 63, pp. 279–292, 2015.

[30] A. Schmitz, Y. Bansho, K. Noda, H. Iwata, T. Ogata, and S. Sugano, "Tactile object recognition using deep learning and dropout," in *Humanoid Robots (Humanoids), 2014 14th IEEE-RAS International Conference on*. IEEE, 2014, pp. 1044–1050.

[31] J. Hoelscher, J. Peters, and T. Hermans, "Evaluation of tactile feature extraction for interactive object recognition," in *Humanoid Robots (Humanoids), 2015 IEEE-RAS 15th International Conference on*. IEEE, 2015, pp. 310–317.

[32] T. Bhattacharjee, A. Jain, S. Vaish, M. D. Killpack, and C. C. Kemp, "Tactile sensing over articulated joints with stretchable sensors," in *World Haptics Conference (WHC), 2013*. IEEE, 2013, pp. 103–108.

[33] "EPCOS (TDK) B57541G1103F NTC Thermistor 10k Bead, Glass," 2016, http://www.digikey.com/product-detail/en/epcos-tdk/B57541G1103F/495-4599-ND/3712554. Accessed: 2016-10-01.

[34] 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," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[35] "General Hidden Markov Model Library," http://ghmm.org/.

[36] L. R. Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," in *Readings in Speech Recognition*, A. Waibel and K. F. Lee, Eds., Kaufmann, San Mateo, CA, 1990, pp. 267–296.

[37] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[38] F. Chollet, "Keras," https://github.com/fchollet/keras, 2015.