Single-Grasp Object Classification and Feature Extraction with Simple Robot Hands and Tactile Sensors

Adam J. Spiers, *Member, IEEE*, Minas V. Liarokapis *Member, IEEE*,
Berk Calli, *Member IEEE* and Aaron M. Dollar, *Senior Member, IEEE*

Abstract—Classical robotic approaches to tactile object identification often involve rigid mechanical grippers, dense sensor arrays and exploratory procedures (EPs). Though EPs are a natural method for humans to acquire object information, evidence also exists for meaningful tactile property inference from brief, non-exploratory motions (e.g., a haptic "glance"). In this work we implement tactile object identification and feature extraction techniques on data acquired during a single, un-planned grasp with a simple, underactuated robot hand equipped with inexpensive barometric pressure sensors. Our methodology utilizes two cooperating schemes based on an advanced machine learning technique (random forests) and parametric estimations. The available data is limited to actuator positions (one per finger) and force sensors values (8 per finger). The schemes are able to work both independently and collaboratively, depending on the task scenario. When collaborating, the results of each method contribute to the other, improving the overall result in a synergistic fashion. Unlike prior work, the proposed approach does not require object exploration, re-grasping, grasp release or force modulation and works for arbitrary object start positions and orientations. Due to these factors the technique may be integrated into practical robotic grasping scenarios without adding time or manipulation overheads.

Index Terms— Tactile Sensing, Object Classification, Object Feature Extraction, Underactuated Robot Hands

1 Introduction

THE extraction of object properties or class through both vision and haptic feedback is a natural sensory ability afforded to humans and other animals. In the field of robotics, both sensory modalities have been investigated. Though many properties of objects may be determined visually, common issues of occlusion and/or poor lighting conditions can limit the performance of vision based methods. Furthermore, other physical properties, such as stiffness, are difficult to visually ascertain, particularly without some kind of object manipulation. Regarding haptics, humans are known to make use of various 'exploratory procedures' (EPs) [1], in order to glean object properties through active manipulation of objects by one or both hands. While there have been several robotic approaches that have taken inspiration from this concept (e.g. [2]–[7]), such methods tend to rely on time-consuming palpatory motion sequences and robot hand and/or arm dexterity. In various real-world application of robotics, such as industrial pick and place, time is a critical factor that would make such exploratory procedures inappropriate. Evidence has demonstrated that meaningful (e.g. inferring more than just collision) knowledge related to objects is achieved by humans during minimal tactile object interactions, termed in [8] as a 'haptic glance'. In turn, the ability of a robot to acquire meaningful haptic knowledge of an object during purely functional actions may extend the usefulness tactile sensing into scenarios where time, computational capabilities or hardware specifications are limited.

In a similar vein to the above objectives, adaptive underactuated grippers have proposed highly practical robot grasping solutions, compared to traditional approaches. Such systems rely on simple and compliant mechanics (e.g. flexure joints) to passively adapt to a wide variety of object shapes and sizes without prior object knowledge, hand modelling, grasp planning [9], actuator regulation or sensory feedback [10], [11]. These benefits result in a low cost, easily implementable solution to grasping in unstructured scenarios.

In this work we seek to combine the benefits of simple adaptive robot grippers with methods of acquiring meaningful haptic object properties during a functional grasp via low-cost, commercially available, tactile sensors. Towards this goal we use a two-finger underactuated hand equipped with 8 *TakkTile* barometric pressure sensors per finger (as illustrated in Fig 1). To maintain consistency with popular open source underactuated hand designs (e.g. [10]), the hand does not implement joint position sensing. Our proposed method achieves classification and feature extraction via tactile sensor outputs and actuator position, sampled at several instances during grasping of an unknown object in arbitrary pose. Notably, our approach does not modify the typical open-loop actuator behavior of such hands during the grasping process. All necessary computations are also designed to be completed within a short time frame (<100ms), to allow achievement of the aforementioned goals within a normal grasp-

• *A.J. Spiers, M.V. Liarokapis, B. Calli and A.M. Dollar are with the GRAB Lab, Department of Mechanical Engineering and Materials Science, Yale University, New Haven, CT 06511. E-mail: adam.spiers@yale.edu.*

Fig 1: The adaptive underactuated hand used in this work. Each finger has a proximal pin joint and distal flexor joint, both driven by a single tendon. TakkTile sensors are embedded in the grip pads of each finger.

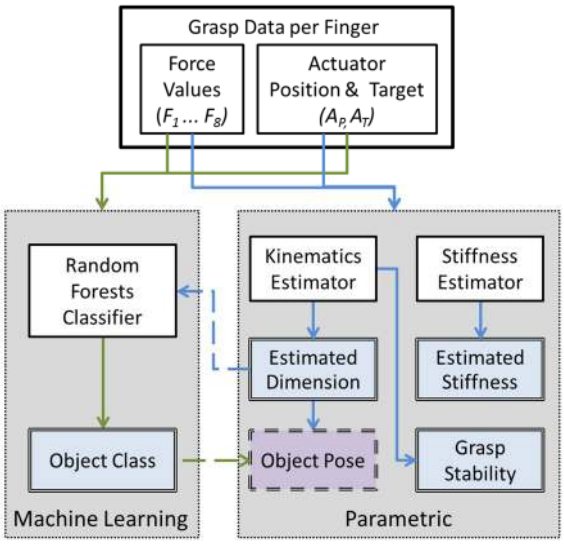

Fig 2: The hybrid approach to data processing. The machine learning and parametric approaches execute simultaneously during an object grasp. Outputs are displayed with a double outline.

ing process. Overall, the presented process is designed to be executed *during* a single, typical, functional underactuated grasp with no temporal or motor overhead. Of course, these constraints (which include uncertain hand kinematics, a single grasp action and unknown object pose) significantly reduce data breath, resolution and redundancy compared to more traditional approaches (e.g. [3]). As a result some parameter estimation is accompanied by some level of uncertainty. However, we believe that some uncertainty is acceptable, given the minimal influence of the method on the fundamental grasping activities carried out daily by thousands of real-world robots. Nevertheless, the classifier performance is excellent, given the limited features.

An example application for such a technique may be in the inspection, sorting and packaging of objects (such as fruit) as part of a production line. In this scenario, the proposed classification and feature extraction method could permit the class of object (e.g. apple, pear), stiffness (ripeness), size and pose within the hand to be determined *as* the object is being lifted from the conveyor belt (a functional and necessary manipulation action).

2 Method

Due to the limited available data and differing strengths of alternative approaches to tactile sensing, we have used a hybrid approach to data processing in this work. The structure of the approach is illustrated in Fig 2. This methodology makes use of a random forests classifier (a machine learning approach) and parametric estimators. The classifier is capable of high level recognition of different objects based on training data. Conversely, the parametric method aims to provide low-level outcomes related directly to physical properties of size and stiffness. The two schemes therefore address different aspects of the tactile sensing problem space. Additional cross communication between these two approaches leads to additional parameter determination (object pose within the grasp) and improved classification accuracy. In the latter case, parametrically determined object dimensions are used to either verify a class decision or dynamically retrain the system. This retraining process rejects the current class and improves subsequent classification accuracy. In previous work we presented some aspects of the classifier alone [12]. We now build upon this work with the parametric estimation and the collaborative hybrid framework. Performance of the classifier is subsequently improved via dynamic retraining features.

The benefits of such approaches apply to alternative use cases. For example, remote exploration or disaster response robots may encounter objects with unique and previously unseen shape, such as an unusual stone or a fragment of a larger object. While raw parameters may be measured for such an object, machine learning approaches may have limited classification success, due to object novelty compared to training data. Conversely, the machine learning scheme can provide high level object identification in structured or semi-structured environments, such as a production line, grocery store or warehouse, where object picking tasks are important. Here, encountered objects will always be part of the company's inventory, though parameter variations in size, stiffness and pose are likely.

The structure of the remainder of the paper is as follows. First, we will review related literature with focus first on biological then robotic systems, providing also a motivation for our approach. A detailed explanation of the experimental conditions (including hardware) will be described in Section 4. All methods and algorithms required to formulate the proposed schemes will then be presented in Section 5 with subsequent results for a variety of objects in section 6. Discussion, future work and conclusions will summarize the paper.

3 Related Work

Roboticists have applied tactile sensing to robot hands for many decades, inspired by nature's most versatile and dexterous end-effector, the human hand. The hand has

approximately 17,000 mechanoreceptive units that innervate its skin and provide a highly sophisticated system for understanding the environment [13]. It has been noted that motion of the hand is crucial to fully exploiting its perceptual qualities during physical interaction [1], [13], [14]. Such observations have been reflected in the active-touch approaches of many robot systems. A large proportion of such endeavors make use of complex, high-density sensor systems, such as the multimodal *BioTac* [15] sensor (used in [3], [4], [6], [16], [17]). This >$10,000 sensor is capable of providing thermal, vibratory and multiple pressure readings over an anthropomorphic finger pad.

Artificial tactile perception efforts generally focus on either deriving physical object properties or the higher level discrimination of an object's class. In [13], it was considered that an object's properties contribute to manipulation actions while the object class enables the execution of object specific strategies or plans. Aspects of human tactile object perception will now be discussed.

3.1 Human Haptic Data Acquisition

The ability to characterize and identify objects without reliance on vision is beneficial in a number of scenarios. In humans, such common tasks as reaching for computer mouse or cup while reading from a computer screen discretely employs complex tactile perceptual methods [8]. Such methods facilitate object understanding (e.g. determining the object class and its pose relative the hand) and subsequent motor action (orientating the hand to facilitate appropriate grasping) with limited physical interaction [18]. Studies on more elaborate exploration of objects have demonstrated the fast and accurate ability of humans to identify a large number of objects and properties through touch alone. This is via the use of *exploratory procedures* (EPs) [1], [14], stereotypical patterns of active hand motions that expose particular physical properties of haptic objects. For example, rubbing an object permits textural perception, while squeezing an object exposes its stiffness. In medicine, such interaction permits identification of tissue type and underlying structure [19]. It was observed in two finger palpation by surgeons that EPs were often combined [20], permitting multiple feature extraction with increased efficiency.

Investigations have also been made into the capabilities of humans to extract meaningful haptic information with limited active finger/hand motion, which is akin to single-grasp robotic approach taken in our work. In [8], perceptual accuracy was considered for a 'haptic glance', a brief and restrained contact between fingertips and an object. Similar investigations have been considered with reduced sensory and motor [14], [18], [21] capabilities. Lederman et al. noted that minimal haptic information is often informative enough to lead to object/feature identification and appropriate subsequent manipulation strategies [18]. On these lines, the work presented in this paper explores what may be achieved by robots via minimal active touch sensing and minimal motor control. This is realized as data acquired during a non-exploratory, 'functional' grasp with adaptive fingers. In [21] it was observed that adaptive 'molding' of the human hand around objects facilitates improved haptic identification. Such 'molding' is fundamental to adaptive grippers, whose use in similar scenarios will be described next.

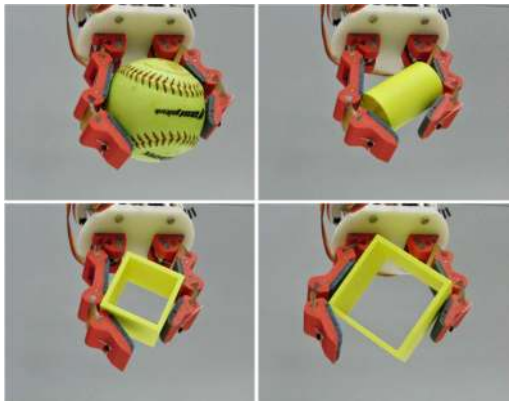

Fig 3: A T42 OpenHand robotic gripper securely grasps a variety of object shapes and sizes using open loop motor control with compliant adaptive fingers.

3.2 Adaptive hands and Tactile Sensing

As previously stated, adaptive underactuated grippers permit grasping of a wide variety of objects with little control or planning effort (Fig 3) [22]. In particular, compliant flexure joints permit out-of-plane finger adaptation to various conditions while maintaining grasp stability [23]. The transmission mechanisms employed in such designs as [10], [11] have similarities to the mechanics and resulting adaptive behavior of the human finger [24].

Despite the benefits of adaptive grippers, there has been relatively limited use of such systems in haptic applications. This is likely to be due to kinematic uncertainty of the fingers, after encountering unknown objects in arbitrary poses. The authors of [25] determined contact with underacted grasper based on motor current models. Grasp force regulation and some object shape distinction was achieved in [26] through tactile sensors and closed loop control. The use of tactile contact sensing to further enhance grasping performance through individual finger control was proposed in [23]. In [27], tactile sensing of finger contacts with an object during workspace exploration led to more optimal object/hand positioning prior to grasping. Closer to feature extraction, underactuated fingers, equipped with joint sensors in [28], re-constructed the contours of immobile rigid objects based on finger positions, while physically exploring a workspace.

3.3 Tactile Feature Extraction in Robotics

Tactile perception techniques in robotics generally fall into categories of feature extraction and object classification, both of which are explored by our hybrid methodology. The majority of techniques rely on active tactile sensor motion via fully actuated robotic systems with predictable parameters. Such approaches have aimed to expose object parameters such as texture [6], [29], stiffness [4], [7], [30], [31] surface contours [28], [32], [33] and thermal properties [7]. In some cases, a direct subset of human inspired exploratory procedures (EPs) were implemented [3], [6], [7], [34]. In [35], a series of non-human

EPs were executed with a parallel jaw gripper improve force regulation during manipulation. In [6], human inspired EPs facilitated information gathering from a BioTac sensor mounted on an anthropomorphic robot finger. Limitations in finger motor resolution limited textural / vibration sensing, compared to precision positioning platforms. Four EPs were implemented in early work by Dario et al. in [7]. A series of six EPs (variations of pushing and sliding) permitted attribution of properties to 34 'haptic adjectives' [3], interpreted by a classifier.

Though EPs permits significant extension of the spatial and dynamic range of tactile sensors, the procedures are often associated with significant time overheads. The relatively extensive repertoire of 6 motions in [3] takes over 85 seconds, while the reduced motions of [35] resulted in a grasping process of 30 seconds. In industrial processes, this time demand can be considered excessive.

3.4 Robotic Tactile Classification

Tactile data is often vast, interconnected and noisy. Machine learning approaches have been used to relate such complex data to object class. Following a process of training, such systems aim to identify objects from new data. Machine learning approaches have been used both for high-level object class distinction in addition to classifying specific feature properties. In [36] pressure data acquired from gripping vegetables facilitated categorization into three classes of ripeness. Unfortunately, the gripping method destroyed the vegetables, via the combination of open loop control with a fully actuated gripper. In [33], a classifier determines surface features (edge, face, empty space) in order to construct an overall spatial object model. Active sliding of a 6-axis force/torque sensor generated data for neural network based distinction into classes of materials in [37]. The popular vision based object recognition technique 'bag-of-features' was applied to tactile based classification in [38]. This method constructed a 'vocabulary' of tactile images based on a several grasp locations. Tactile array 'images' gathered during squeezing and releasing objects are used as the basis of a *k nearest neighbors'* approach in [39], though only slight object pose variations were implemented. Unsupervised learning techniques have also been applied to this problem space. Incremental online learning was applied to tactile and joint sensor data in [40] to improve classifier's efficiency. In [41] spatial and temporal data was obtained for a variety of hands following data acquisition via a sequence of five squeezing actions followed by releasing of an object. In this work an unsupervised hierarchical feature learning methodology was employed and a 1-vs-all classifier obtained. Reinforcement learning techniques were also utilized in [6], [17], to cluster *BioTac* data resulting from exploratory finger motion in order to report stiffness, texture and thermal properties of objects. The work of Chu et al. [3] classified data resulting from 6 EPs into adjectives via machine learning approaches. Classification of the fullness of plastic bottles was achieved in [30] based on a single grasp. Unlike our approach, the bottles do not vary in size or orientation and closed loop force and velocity control of the robot gripper is employed. These control modes were required for achieving 'safe' container grasping with the PR2 parallel gripper. Such a concern is avoided in our setup via the adaptive gripper.

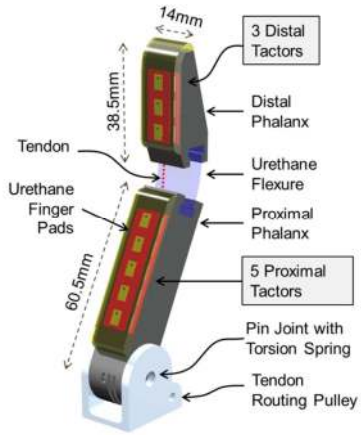

Fig 4: An underactuated prototype Reflex Hand finger, equipped with barometric Takktile sensors. The finger has no position sensors.

Note that approaches such as [3], [39] only permit object classification once an object has been released. Presumably, any actions related to object classification (for example sorting) would then require re-grasping.

In previous work we presented some aspects of the classifier alone [12], but we now build upon this work with the parametric methods and collaborative efforts.

3.5 Literature Overview

The review of existing work has demonstrated trends in tactile identification that favor dense sensory data and extended exploration of objects. It has also been illustrated that though humans make use of EPs, useful haptic object knowledge is also often extracted via minimal, non-exploratory, active haptic interaction [8], [18]. In this work we strive for an equivalent robotic approach based on such minimal interaction. By negating motion, processing, hardware and time overheads, we hope for a solution that is practically implementable. Added to this, we distinguish ourselves from previous work via the robustness of the system to perturbations in object pose (position and orientation) within the grasp of the hand.

4 Experimental Setup

4.1 Underactuated Robot Hand

The robot hand used in this study (illustrated in Fig 2) consists of two prototype fingers of the *Reflex Hand* (manufactured by *Right Hand Robotics, Boston, USA*) mounted on a modified *model T42* base from the Yale OpenHand project [10]. Each *Reflex Hand* finger (shown in Fig 4) consist of two phalanges with a distal urethane flexure joint and a proximal pin joint with torsional spring. The benefits of this arrangement are described in [10]. Each finger is actuated by a single tendon, wound via a pulley attached to a Dynamixel MX-28 actuator. Unlike other models of the Reflex Hand fingers, the fingers used in this work (like similar open source designs [10][11]) do not

Fig 5: The three sets of objects used in this work. Product logos have been obscured for copyright considerations.

feature position sensing.

4.2 Tactile Sensors

A row of 'TakkTile' force sensors [42] are embedded in the compliant, high-friction grip pads of each link of the robot finger. The grip pads are cast from 'VytaFlex 40' material from Smooth-On Inc. The robust and inexpensive TakkTile sensors are based on urethane encased MEMS barometers, mounted on a printed circuit board with 8mm separation. A 5 sensor 'Takkstrip' may be purchased for $150. Each robot finger features 8 sensors mounted on two such strips; 3 sensors on the distal phalanx and 5 on the proximal phalanx. Each sensor outputs a single pressure value at 100Hz, with resolution of <0.01N [42]. Prior to the experiments, the embedded tactile sensors were calibrated using a series of weights (10g to 100g) applied via a passive linearly constrained motion platform (weighing 10.1g) and custom finger clamp. Calibration illustrated primarily linear responses. This allowed linear sensor equalization for the parametric mode. Sensor 3 of the left finger showed significantly reduced sensitivity and was negated from the parametric processes.

4.2 Objects

The experiments were conducted with two sets of 'model objects' and a set of 'everyday' (i.e. household) objects (Fig 5). The model objects were custom fabricated to constrain parameter variation and validate the parametric

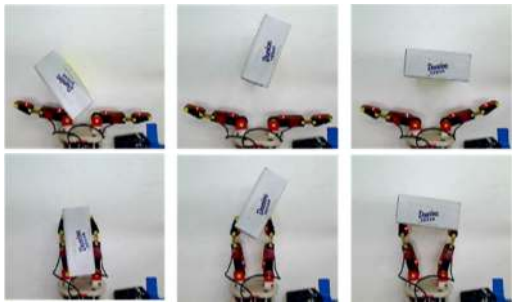

Fig 6: Examples of object pose variation and resulting grasp

Table I: Sizes (mm) of set 1 model objects.

Objects	Small	Medium	Large
Cylinders	50	70	90
Rectangles	50	70	90

Table I: Stiffness (N/m) of set 2 model objects.

Objects	*Softest (Green)*	*Soft (White)*	*Hard (Black)*	*Hardest (Yellow)*
Cylinders	156	346	2,182	97,286
Rectangles	156	346	2,182	51,111

Table III: Characteristics of the everyday objects.

Objects	Dimensions (mm)	Stiffness Side 1 (kN/m)	Stiffness Side 2 (kN/m)
Coffee Can	102x139	67.2	N/A
Soup Can	66x101	49.4	N/A
Sugar Box	38x89x175	4.73	26.87
Apple (plastic)	75	10.6	N/A
Peach (plastic)	59	8.79	N/A
Windex Bottle	80x105x270	9.87	5.0
Mustard Bottle	50x85x 175	4.69	2.98
Cracker Box	60x160x230	2.6	3.0
Bleach Bottle	50x93x250	3.2	3.2
Gelatin Box	28x85x73	3.1	4.7
Cracker Box	60x160x230	2.6	3.0

estimation methods. The everyday objects were selected from the *YCB object set* [43] to represent a diverse range of size, shape, stiffness and weight parameters. The set is a recent benchmarking standard for robotic manipulation that facilitates replication of test equipment and procedures between research groups. The model objects consisted of two sets, each of which contained circular and cubic objects. The first set were fabricated from 3D printed ABS (wall thickness 4mm) to maintain stiffness but varied in size. The second set maintained the same size but varied stiffness, by the use of different foam materials. Characteristics of the model object and everyday sets are presented in, Table I, II and III. Set 1 stiffness corresponds to the 'Hardest' measure of set II. Stiffness in all cases was measured via a load cell (10mN accuracy) mounted on a linear actuator (0.01mm resolution).

5 METHODS

5.1 Data Collection

Data was collected by grasping each object with the robot hand twenty times, in various positions and orientations. As our method currently involves no post-grasp manipulation (e.g. lifting), the hand was mounted (via clamps) to a table. During each trial (grasp), the actuators were commanded to move 270deg over 3.25seconds with a constant target velocity. The final target position was maintained for 250ms at the end of the motion before the actuators returned to 0deg, releasing the object. Actuator target and actual positions, plus force sensor data were measured at 100Hz. Images from a webcam, mounted over the hand, were also logged for validation purposes. All logging and control was performed via ROS.

In each trial, the object was placed on the table surface, in an arbitrary pose (position and planar orientation) within the workspace of the gripper. For some classifier tests,

planar orientation of the object was constrained to perturbations of ±45deg (these will be discussed in Section 6.1). The same object surface rested on the table in all cases. An example of such pose variation is illustrated in Fig 6, from the logged webcam data. An additional 7 empty grasps (with no object present) were also recorded. An example of the actuator and force data resulting from a single grasp is demonstrated in the first three plots of Fig 7. Other features of this figure will be further described in subsequent sections, as will the methods used to process the data. The objects were not constrained after placement. It was observed that objects placed in a pose with a horizontal offset would be 'pushed' into the center of the hand during grasping.

5.2 Machine Learning Scheme

In this section we present the machine learning scheme, which aims to identify objects from the data acquired during a single grasp. To achieve this goal an ensemble classifier was employed based on the random forests classification technique.

5.2.1. Random Forests Classifier

Random forests (RF) were originally proposed by Ho [44] and Breiman [45] and are an ensemble classifier based on different decision trees. The output is the most popular class between the decisions of the individual classifiers. The RF technique provides high classification accuracy, does not overfit and handles multiclass problems (such as distinguishing between multiple object). Furthermore, the method is fast and efficient when dealing with large databases and has the capability to handle high numbers of input variables. A diagram of the RF classification procedure for n trees is presented in Fig 8. Each tree of the RF is constructed from a different out-of-bag (*oob)* sample set from the classifier training data. This training data comprises two thirds of the recorded grasps data. The remaining data is used for validation. Classification accuracy, and comparison of RF technique to various state-of-the-art classifiers, will be presented in Section 5.

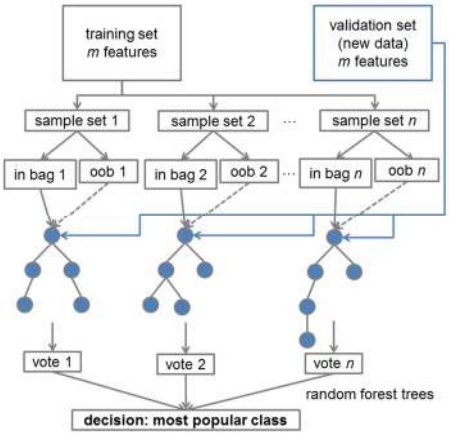

Fig 8: The random forests (RF) classifier with n trees.

5.2.2. Feature Selection

The feature space used for discriminating between the objects is defined by the actuator and force sensor data at two different time instances of the grasping process. The first instance (t_1) is taken when the sum of actuator target positions (A_T) exceeds the sum of actual actuator positions (A_P) by a given threshold, (T_{Stall} = 20deg) via $|A_T - A_P| > T_{Stall}$. This deviation indicates a stall in actuator motion due to finger /object interaction. The second instance (t_2) occurs when actuator target positions (A_T) have reached steady state (this occurs at time t=3.25s). Here, the object is being held with constant tendon exertion. These instances are indicated as 'Deviation' and 'Target Motion End' on Fig 7. Actual actuator positions (2 values) and force sensor readings (16 values) are extracted at these two instances, giving a feature space of 36 variables. This raw data is obtained without a-priori information regarding the robot model or actual joint angles, making the machine learning methodology model-free. Another beneficial characteristic of the classifier is that it does not necessarily require calibrated force values, as classification is

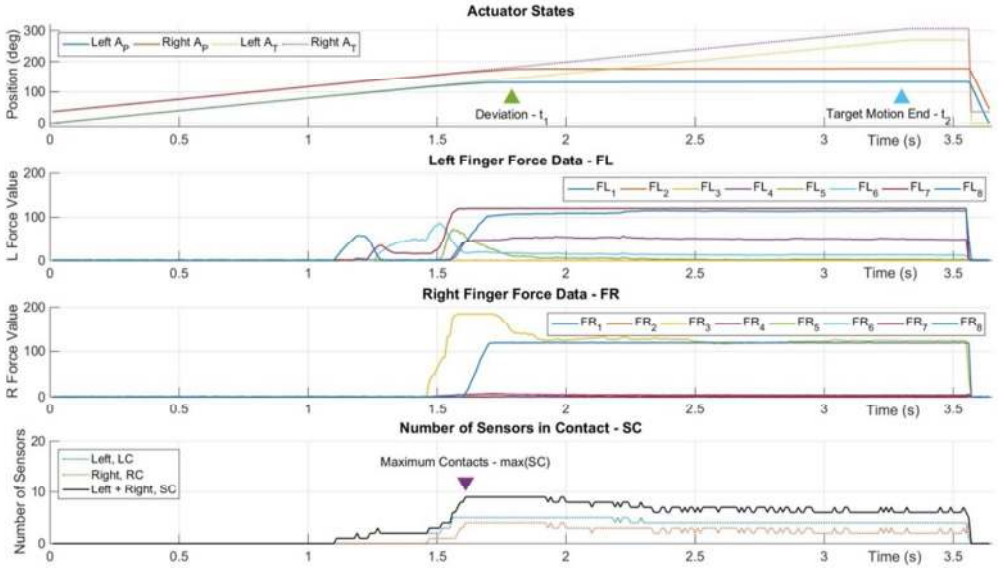

Fig 7: Actuator position and force sensor data during a single object grasp. Triangular markers denote events identified by the methodology. The bottom plot shows the number of sensors in contact during a grasp. The stable grasp start may occur before or after actuator deviation, depending on grasp conditions.

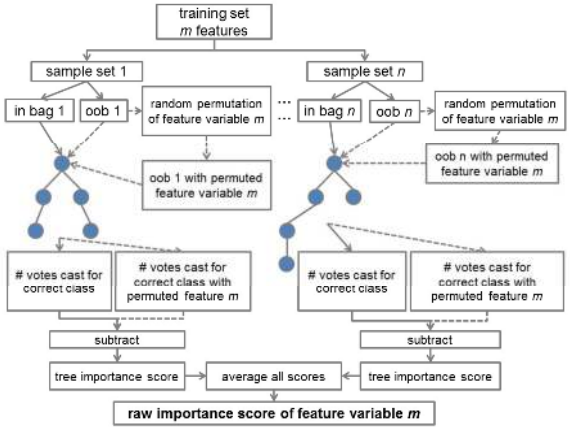

Fig 9: RF feature importance calculation procedure.

based on the differentiation in input data.

5.2.3. Features Importance Calculation

The RF technique has an inherent capability of computing the importance scores for all feature variables and comparing these relatively. Such a calculation is useful for optimizing future hand designs, by minimizing the number of sensors required to achieve a certain level of classification accuracy. When fewer sensors are used, their locations on the fingers become more critical.

Importance calculation is based on manipulations of a subset of the training data, that is called out-of-bag (*oob*) samples. These *oob* samples are given as an input to all decision trees and the number of correct votes counted. Then, the *oob* samples values of a feature variable m are randomly permuted. The modified samples are once again "fed" to each of the n decision trees. Importance of the feature variable m is then calculated by the operation $I_m = V_P - V_U$. Where Vp is the number of correct votes cast with the m-variable permuted *oob* data and V_U is the same metric with untouched *oob* data. The overall/raw importance score (I_m) for each feature variable m, is the average of the importance scores computed for all trees of the RF. The process is described in Fig 9.

In this work, we normalize importance scores to facilitate comparisons of the different feature variables, even for different classification problems.

5.3 Parametric Method

The parametric method estimates physical object parameters of size and stiffness based on data acquired during the grasp. Size is related to a contact polygon constructed from force and actuator data. Additionally, some measure of grasp stability is provided. The parametric method makes use of several processes, as illustrated in Fig 10. These processes will now be described in more detail.

5.3.1 Forward Kinematics Estimator

To estimate the size and shape of an object, the parametric method relies on knowledge of the kinematic position of the robot fingers once a secure grasp has been made. Predicting the kinematic behavior of mechanically compliant underactuated fingers is non-trivial. In addition to the complexities of modelling flexure joints [46], multiple joint position solutions exist for each actuator position. Actual finger kinematics result from finger and transmission dynamics, which is modulated during different stages of an adaptive grasp by interaction with unknown objects [22]. In the case of the fingers used in this work, the inclusion of tactile sensors permits interaction detection on each phalange. Based on this data, a computationally efficient kinematics estimator (Fig 11) was constructed that uses force sensor information to switch between different grasping 'modes', as illustrated in Fig 12. In each mode, a different set of transmission gains (G) converts actuator position (A_P) to motion of the proximal and distal joint (θ_P and θ_D). A_P is also equivalent to tendon length from fingertip anchor to actuator. For simplicity, θ_D is considered as a pin joint in the kinematic structure of the finger. Mode selection is based on F_P and F_D, which are the sum of individual force sensor values on the proximal and distal phalanges respectively. The force value thresholds required to halt link motion on each joint are defined independently as T_D and T_P. These values were determined experimentally and are higher than the threshold used for contact detection (T_C). This allows the kinematics to deal with the common case of a single finger pushing an object into the center of the hand (as illustrated in the left example of Fig 4), prior to a grasp being made. The different modes may be explained as follows:

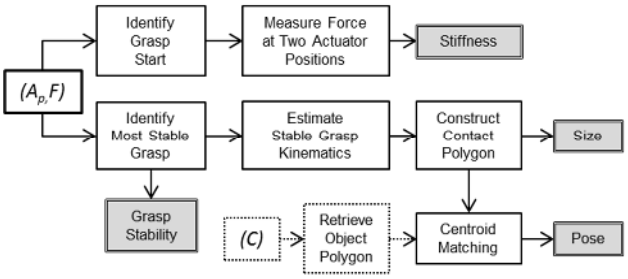

Fig 10: Processes within the parametric estimator. Outputs are shaded boxes. Classifier dependent components (such as object class, C) have dotted lines.

- **Mode 1 - Pre-contact**: 'Free motion' of the finger prior to object contact. Actuator motion generates a large change in θ_P and small change in θ_D.
- **Mode 2 – Proximal Contact**: If $F_P > T_P$ motion of the θ_P stops. Actuator motion is transferred to θ_D.
- **Mode 3 – Proximal & Distal Contact**: If $F_D > T_D$ motion of both joints stop.
- **Mode 3a – Distal Only Contact:** Mode 3a in Fig 12 denotes a distal only grasp (no proximal contact). This is also recognized by $F_D > T_D$ (same as Mode 3) but without precedence by a proximal contact.

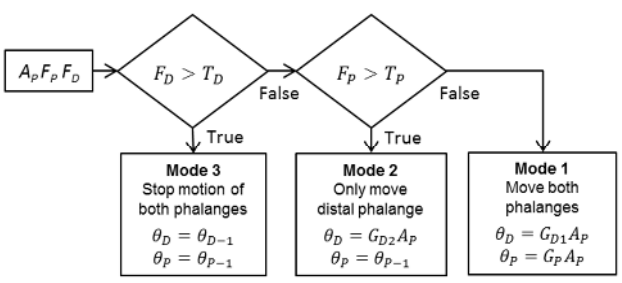

Fig 11: Kinematic estimation via actuator positions (AP) with proximal (FP) and distal (FD) force sensor values.

The forward kinematics method is iterative and as such, determination of joint angles for a specific instance necessitates calculation of all joint angles up until that instance. Due to a lack of typical kinematic matrix operations, this process has very little computational overhead. Kinematics estimation for a complete grasp (e.g. the data in Fig 8) takes less than 30ms on an Intel i7 3.6Ghz PC.

Motion gains and threshold values were determined via a calibration process in which joint angles were visually observed from ten overhead video frames recorded during a single empty grasp and a single grasp of a rigid 70mm cylinder. Joint angles were determined by locating the spatial centroid of small yellow markers attached to the finger phalanges. These markers are visible in Fig 6 and Fig 13. Interpolation of these co-ordinates with synchronized actuator data led to linear transmission models and three gains. These gains (for both fingers) are (proximal) G_P *= (3.58, 3.38),* (distal during mode 1) G_{D1} *= (0.26, 0.11)* (distal during mode 2) G_{D2} *= (5.29, 5.06)*. A limitation of this method is that different gains are required for different orientations of the hand relative to gravity. We predict that performing calibration for a small number of orientations would allow interpolated gain estimation.

5.3.3 Stable Grasp Location and Analysis

The hand kinematics may be combined with force data to establish the spatial co-ordinates of contact points on the finger pads for any instance of grasp data. A suitable time instance for such a task would be when maximum stability has been achieved for a grasp. For the purpose of this work we associate grasp stability with number of contacts of the hand with the object. Counting the number of sensors whose values exceed a given contact threshold *(T_C)* at each time instance *(t)* of the grasp produces the following 'sensors-in-contact' array *SC(t)*:

$$SC(t) = \sum_{i=1}^{n=8} LC_i(t) + \sum_{i=1}^{n=8} RC_i(t) \tag{1}$$

$$where\ LC_i = \begin{cases} 1, & FL_i > TS \\ 0, & FL_i \le TS \end{cases},\quad RC_i = \begin{cases} 1, & FR_i > TS \\ 0, & FR_i \le TS \end{cases} \tag{2}$$

Where *LC* and *RC* are binary array of the sensors in contact, derived from the force values *FL* and *FR*. The *SC* array may be searched to locate the first instance of *max(SC)*, the maximum sensors in contact. This instance is indicated on Fig 7 as 'Maximum Contacts'. The process is highly efficient and may be computed at the 'Target Motion End', t_2, with no requirement to release the grasp.

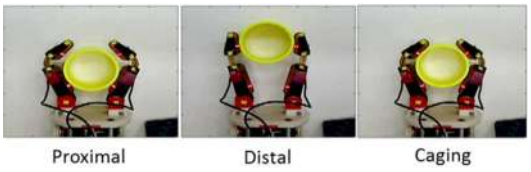

Fig 13: Examples of the three grasp types.

The value *of max(SC)* may be used to give an indication of grasp stability *(G_S)*, when combined with grasp type determination (proximal, distal or caging, as illustrated in Fig 13). Grasp type may be determined by counting the number of proximal (C_D) and distal (C_P) contacts.

$$C_P = \sum_{i=1}^{n=5} LC_i(t) + RC_i(t),\ \ C_D = \sum_{i=6}^{n=3} LC_i(t) + RC_i(t) \tag{3}$$

Each grasp type is represented by a gain (G_T) associated with the stability of that grasp. The number of sensors in contact *SC*, is scaled by G_T to give the G_S score.

$$G_S = G_T\, SC(t),\ \ where\ G_T = \begin{cases} 1, & C_D > 0, \\ 1.2, & C_P > 0, \\ 2, & C_P > 1 \wedge C_D > 1 \end{cases} \tag{4}$$

The highest gain (2) is associated with a caging grasp, where the object is secured by both proximal and distal links. The proximal only grasp is rated slightly higher (1.2) than a distal only grasp (1), due to a tendency for an object that slips from a distal grasp to become secured by a proximal grasp. Example G_S scores will be provided for different example cases in Section 6.

5.3.4 Grasp Polygon Construction

The kinematics estimator also permits the location of contact points to be established at *max(SC)*. At this point, the co-ordinates of in-contact sensors may be determined from finger joint angles (established in Section 5.3.1) and the arrays LC and RC. These co-ordinates may then be connected to construct convex *grasp polygons*, as will be later illustrated and discussed in Section 6.2.1 and Fig 17.

The grasp polygon provides various aspects of the object and grasp. One easily extractable feature is the grasp aperture (polygon width), which may be used as a simple dimensional output. We use this metric for verification of classifier results. Further results pertaining to this will be given in Section 6.

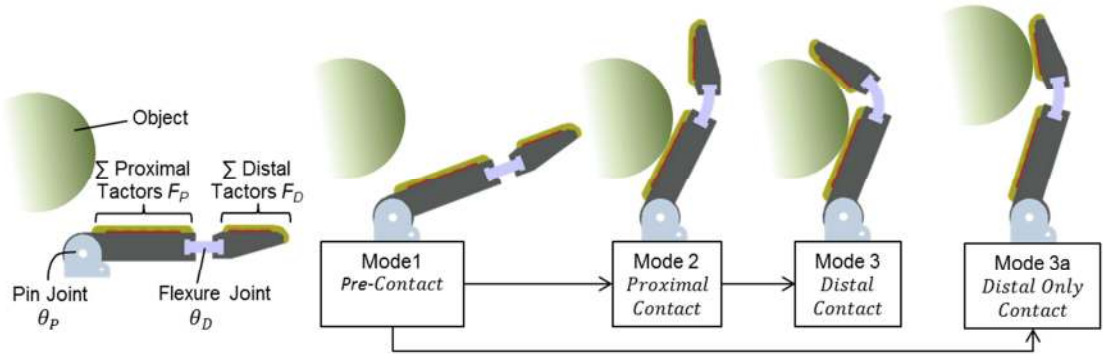

Fig 12: Kinematic mode progression based on finger adaptation to object contact

5.3.5 Stiffness Estimation

Estimation of object stiffness is achieved by observing the change in all measured forces between two intervals in the grasp. Due to the limits of available data, it is not possible to fulfill the stiffness equation $S = \Delta F / \Delta x$, where stiffness, S, is measured by observing a change in force ΔF for a change in surface displacement Δx. This is because the change in finger position after a grasp has been established cannot be accurately measured. However, we have found that a metric related to object stiffness (S_M) may be achieved by considering the average reaction force per tactile sensor at two instances of the grasping process. These intervals are the same as those used for feature selection (defined in Section 5.2.2), i.e. the actuator stall instance (t_1) and the steady state condition of A_T (t_1). This capability results from the open-loop nature of the grasp, which maintains some consistency between conditions.

$$F(t) = \sum_{i=1}^{n=8} LC_i(t)FL_i(t) + RC_i(t)FR_i(t) \quad (5)$$

$$S_M(t) = \frac{|F(t_1)|}{SC(t_1)} + \frac{|F(t_2)|}{SC(t_2)} \quad (6)$$

Where *LC*, *RC* and *SC* were defined in equations (1) and (2). *FL* and *FR* are force sensors measurements on left and right fingers. By combining *FL*, *FR* with *LC*, *RC* the values of non-contact sensors are set to zero, reducing noise. Other methods of stiffness estimation that incorporated difference in actuator error and time were evaluated, but led to large error margins with variations in grasp type and pose, particularly for stiffer objects. Results for stiffness estimation of model object set 2 will be presented in Section 6.

5.3.6 Pose Estimation

Pose estimation is a feature extraction process that allows the location of objects within a grasp to be estimated. Such a technique is useful for determining subsequent manipulation of an object. For example, placing a grasped object at a target location will require different hand positioning depending on the location of the object in the hand. The technique is facilitated by collaboration with the random forests classifier, which provides the estimator with the object class (C). This leads to recall of a simple object polygon for known objects, based on size and shape. For example, the apple object's polygon would be a circle with a diameter of 75mm. The centroid of the object polygon is then matched to the Y (distal) component of the centroid of the grasp polygon. Currently this method functions only for circular objects.

6 Results

6.1 Classification Accuracy

For the training of the random forest classifiers, a 10-fold cross-validation procedure [47] was used to assess their efficiency and avoid overfitting. The classification accuracies are reported in Table IV. These results were computed by averaging multiple rounds of the cross-validation method. In all cases, two thirds of the acquired grasping data for each set (20 grasps per object) was used for training, with the remainder used for validation. All classification accuracy results are given in Table IV.

Table IV: Classification results for all experiments

Objects	Case	Accuracy
Model 1	Size & Shape	93.57% (SD: 3.25%)
Model 2	Shape & Stiffness	93.01% (SD: 3.02%)
Everyday	Constrained Orientations	100% (SD: 0%)
Everyday	Free Orientations	94.32% (SD: 3.09%)

6.1.1. Model Objects Classification Results

The first classification problem involved discrimination between the model objects of set 1 and 2 (see Section 3.2). The trained classifier is slightly better at discriminating between objects with different shapes and sizes rather than objects of different shapes and stiffness.

6.1.2. Everyday Objects Classification Results

The second classification problem involved discrimination between the various everyday objects. This was first attempted for a constrained orientations case, in which the objects were positioned in an arbitrary manner but orientations were constrained within ±45 degrees of the principal axis. In the unconstrained case, objects were positioned with arbitrary positions and orientations (±180 degrees). As expected the classification accuracy is higher for the constrained orientations case, though the free orientation result also demonstrates excellent accuracy.

6.1.3. Comparison of Various Classifiers

To test the suitability of the RF machine learning approach to this problem, the classification of everyday objects in constrained orientations was repeated with a number of alternative state-of-the-art classifiers.

The methods used were a Linear Discriminant Analysis (LDA), a Naïve Bayes classifier, a Neural Network (NN), a binary Support Vector Machines classifier (SVM) and the random forests (RF) technique. The SVM classifier was trained using different kernels (linear, RBF etc.) and the best results were acquired. The NN classifier was constructed using a single hidden layer with fifteen hidden units trained with the Levenberg-Marquardt backpropagation algorithm. RF forests were grown with ten trees for processing speed (two times faster) and one hundred trees for accuracy. All classifiers were compared for the task of discriminating between the everyday life objects with constrained orientations (±45deg) and arbi-

Table V: Comparison of different classifier accuracy on everyday life objects with constrained orientations

Classifier	Classification Accuracy
LDA	89.74%
NB	90.61%
NN	95.65%
SVM	92.30%
RF (10 trees)	98.04%
RF (100 trees)	100%

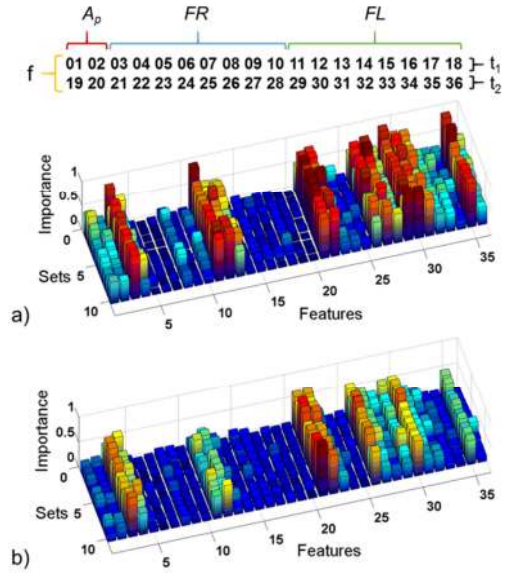

Fig 14: Feature variables importance bar plots for discrimination of everyday objects with a) constrained orientations, b) free orientations.

trary positions. The classification accuracies for the different techniques are reported in Table V.

Random forests outperform all other classification methods, but all methods provide high classification accuracies. Thus, the proposed method could have been used with other classifiers with similar results.

6.1.4. Sensor Placement Optimization

Feature variable importance was implemented to determine optimal force sensor locations for different quantities of force sensors. The feature variables importance scores for all 36 features (18 features at two time instances - Section 4.2.2) are presented in Fig 14, for cases of everyday objects with constrained and free orientations. The height of the bars represents the importance scores of the features for 10 sets. Each set is a distinct random splitting of the training data. It may be observed that all scores are robust along the different training data splitting. To evaluate this approach, the most important feature variables were selected and the classifier was re-trained after removal of redundant features. Three different cases were examined, performing retraining for the 4, 6 and 8 most important sensors of the hand, rather than using the initial 16 force sensors. The classification results are reported in Table VI. These results illustrate a significant redundancy in the initial feature space, with respect to the chosen classification method. In this respect, future hands could be constructed with less sensors, in optimal positions Designs suggested by this approach, which all feature asymmetrical sensor placement, are shown in Fig 15.

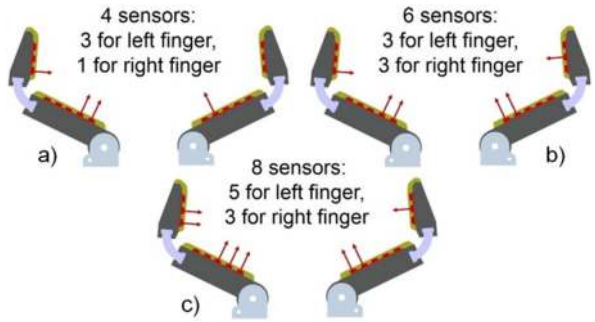

Fig 15: Optimal sensor placement for 6 and 8 sensor setups based on features variables importance.

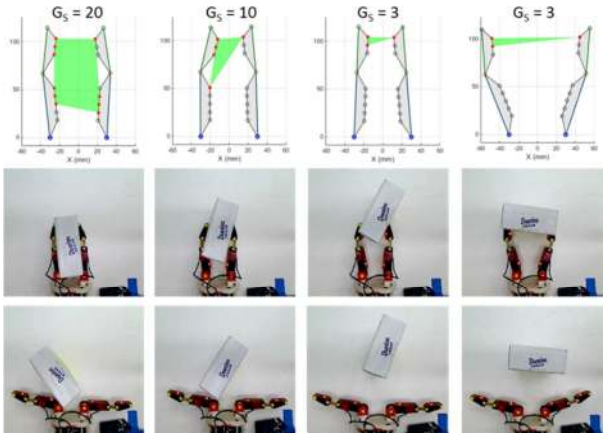

Fig 16: (Top row) Grasp polygon construction and finger kinematics estimation. G_S shows grasp stability score. (Middle row) corresponding video frame of the most stable grasp instance, based on *max(SC)*. (Bottom row) Object starting pose.

Table VI: Effect of feature variables selection (sensor reduction) on classification accuracy for everyday objects.

Sensors	16 (All)	8	6	4
Accuracy	94.32%	95.78%	93.14%	94.67%
SD	3.09%	3.7%	2.67%	3.79%

6.2 Parametric Results

The parametric method resulted in a number of outputs for each grasp. To provide reliable benchmarks, results are presented for the model objects.

6.2.1 Dimension Estimation

Estimation of dimension was performed for the 50, 70 and 90mm diameter cylindrical model objects in set 1 (Fig 5) and an empty grasp. These were based on the width of grasp polygons, derived from estimated finger kinematics. Examples of determined grasp polygons for an everyday object (box of sugar) are illustrated in Fig 15. Each example is also annotated with a measure of grasp stability G_s, as defined in Section 5.3.3.

Fig 17 provides histograms of estimated dimensions (based on the width of the contact polygon) for all grasps (with unconstrained object positions). Clear distinctions are illustrated between the different objects sizes with Gaussian distribution and median values following a trend consistent with the actual diameters of the objects, though a linearly increasing offset between estimated and actual values may be observed. It may be seen that sizable error bounds are present, particularly for the smaller, 50mm diameter cylinder. A further breakdown of dimension estimations into grasp types Fig 17 illustrates persistent errors for the 50mm object with the caging grasp (Fig 13). Conversely, little variance is indicated overall for proximal grasps. It is likely that the uncertainty of finger motion in different grasps is being reflected here, with the kinematics estimator faring better for finger interaction

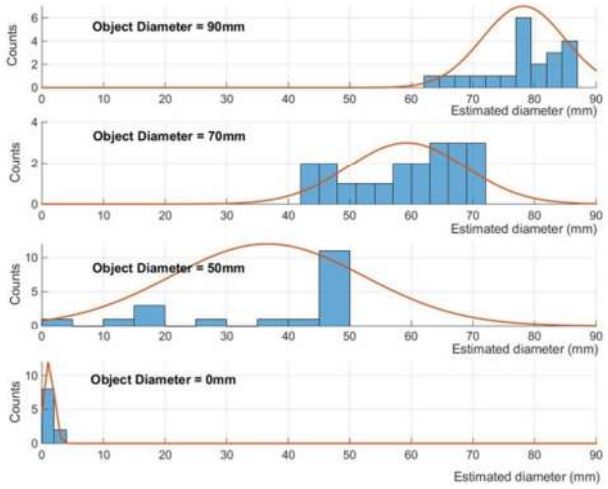

Fig 17: Histogram of parametrically estimated diameters of three stiff cylinder objects and an empty grasp. Gaussian distribution curves are also shown.

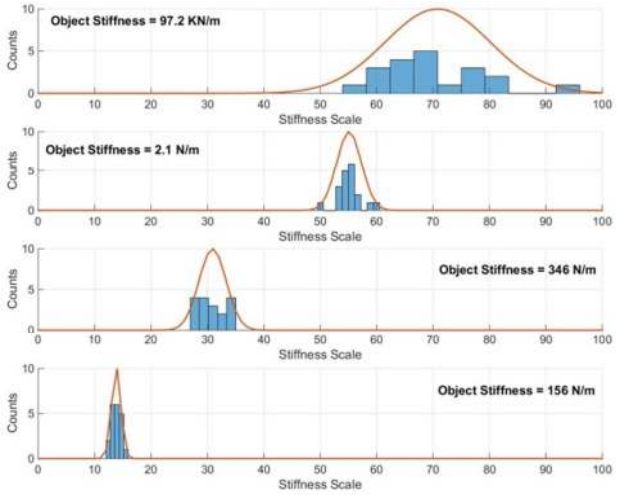

Fig 19: Histogram of estimated stiffness (scale 0-100) with Gaussian distributions for various stiffness cylinders. Data has been grouped across grasp types.

for certain cases. Indeed, in some grasps it was observed that a distal only grasp could lead to further object motion (such as pulling the object towards the palm). Of course, the grasp does not always occur at the widest point of the object, which is represented by the negative offset in some many cases.

6.2.2 Stiffness Estimation

Parametric stiffness estimations for set 2 of the model objects (Fig 5) are illustrated in Fig 19 for all cases. The outputs of the stiffness metric, S_M were scaled between values of 0-100. The results indicate excellent distinction between three different stiffness objects (in the range 156 to 2100 N/m) with some larger distribution of errors for the most rigid object (where sensor output would often saturate during the grasp). Inspection of S_M distribution in Fig 20 illustrates that the results do not overlap within each grasp type. Therefore, by determining the grasp type G_S via equation (4), it is possible to automatically categorize each result into an appropriate stiffness value.

6.3. Collaborative Results

The following two brief results are based on collaboration between the parametric estimator and machine learning approach. Sharing data between processes in this way is demonstrated to be highly beneficial.

6.3.1 Dynamic Classifier Retraining

After the parametric method has estimated the dimension of the object (a process that takes less than 15ms after kinematic computation), the measured dimension is passed to the machine learning algorithm. This parameter provides an extra level of validation of the classification decision. For example, say the classifier were to mistakenly classify a soup can as a coffee can (from Fig 5). However, the estimated object dimension from the parametric method indicates that the grasped object is too small to be a coffee can, given the known dimensions from Table III. In such a case, the classification decision is verified as false, and the classifier is dynamically retrained, excluding all the objects that have dimensions that are significantly different from the estimated dimensions. Essentially, this is an a-posteriori filtering and correction of the classification decisions that increases the machine learning approach's efficiency. Due to the efficiency of the method, the retraining (which takes less than 65ms) may also be completed during an object grasp.

This synergistic approach was tested on a variety of objects with real and simulated errors. In all cases, retraining of the classifier led to subsequently improved classification accuracy to 100%.

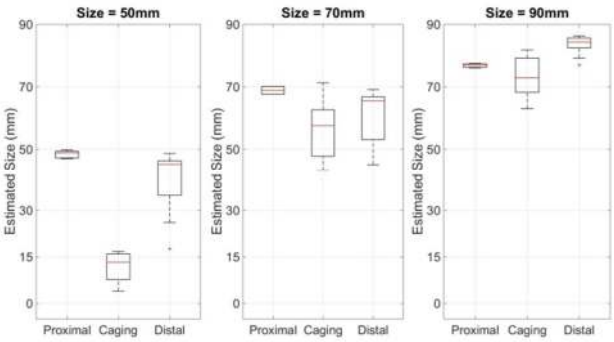

Fig 18: Diameter estimation of stiff cylinders, categorized by grasp type.

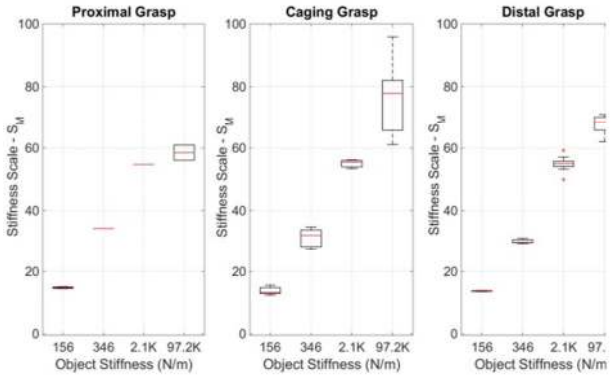

Fig 20: Stiffness estimation results, categorized by grasp type

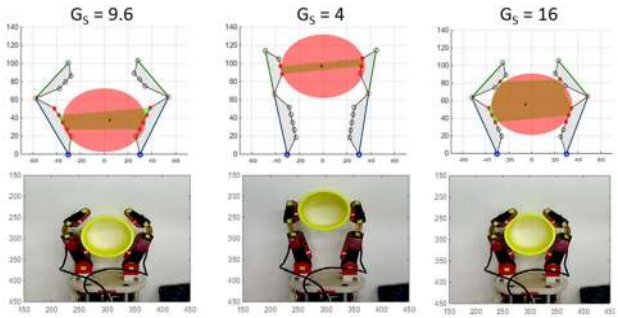

Fig 21: Pose Estimation for a rigid cylinder. The object class polygon (circle) overlaps the contact polygon. G_S scores are also provided

6.3.2 Pose Estimation

Collaboration between elements also occurs by passing of class data from the machine learning approach to the parametric estimator, as described in Section 5.3.6. This enables estimation of pose for round (orientation free) objects. Sample results of this approach are illustrated in Fig 21 for proximal, distal and caging grasps of a 70mm rigid cylinder. Accurate kinematic estimation of finger pose is also shown in Fig 21. The geometric model of the cylinder was reported to the parametric approach by the classifier.

8 Conclusion

In this paper we have presented a hybrid methodology for performing tactile classification and feature extraction during a single grasp with a simple underactuated robot hand. Such robot hands provide highly practical and easily implementable grasping solutions for robotics. Similarly, our work has aimed to provide a system with low motion, time, complexity and cost overheads for haptic sensing applications in practical robotics.

Promising results have been presented, showing high classification accuracy and the ability to extract features (with error bounds) of an object's dimension, stiffness and pose. While more accurate parameter identification has been carried out by other robotic approaches, these have tended to focus on a single parameter as part of a more extensive exploration process. Though the systems machine learning and parametric methods are capable of working independently, a novel collaborative technique allows each method to contribute its outcomes to the other. This improves classification accuracy and allows estimation of object pose.

The various aspects of the proposed methodology are highly suited to dynamic, semi-structured environments where the time or dexterity necessary for detailed haptic object exploration is not available.

Funding

This work has been partially supported by the NSF NRI grant IIS-1317976.

Acknowledgment

Raymond R. Ma is thanked for assistance in robot hand design and experimental setup. Kevin Gemmell is acknowledged for object stiffness measurement help.

References

[1] S. J. Lederman and R. L. Klatzky, "Extracting object properties through haptic exploration," *Acta Psychologica*, vol. 84, pp. 29–40, 1993.

[2] A. M. Okamura, M. L. Turner, and M. R. Cutkosky, "Haptic Exploration of Objects with Rolling and Sliding," *2013 IEEE International Conference on Robotics and Automation (IROS)*

[3] V. Chu, I. McMahon, L. Riano, C. G. McDonald, J. M. Perez-Tejada, M. Arrigo, N. Fitter, J. C. Nappo, T. Darrell, and K. J. Kuchenbecker, "Using robotic exploratory procedures to learn the meaning of haptic adjectives," in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 3048–3055.

[4] Z. Su, J. a. Fishel, T. Yamamoto, and G. E. Loeb, "Use of tactile feedback to control exploratory movements to characterize object compliance," *Frontiers in Neurorobotics*, vol. 6, pp. 1–9, 2012.

[5] A. Bierbaum, M. Rambow, T. Asfour, and R. Dillmann, "A potential field approach to dexterous tactile exploration of unknown objects," *2008 8th IEEE-RAS International Conference on Humanoid Robots, Humanoids 2008*, pp. 360–366, 2008.

[6] J. A. Fishel and G. E. Loeb, "Bayesian exploration for intelligent identification of textures.," *Frontiers in neurorobotics*, vol. 6, 2012.

[7] P. Dario, P. Ferrante, G. Giacalone, L. Livaldi, B. Allotta, G. Buttazzo, and a. M. Sabatini, "Planning And Executing Tactile Exploratory Procedures," *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1992.

[8] R. L. Klatzky and S. J. Lederman, "Identifying objects from a haptic glance.," *Perception & psychophysics*, vol. 57, no. 8, pp. 1111–1123, 1995.

[9] M. Ciocarlie, C. Goldfeder, and P. Allen, "Dimensionality reduction for hand-independent dexterous robotic grasping," *IEEE International Conference on Intelligent Robots and Systems*, pp. 3270–3275, 2007.

[10] R. R. Ma, L. U. Odhner, and A. M. Dollar, "A modular, open-source 3D printed underactuated hand," *Proceedings - IEEE International Conference on Robotics and Automation*, 2013.

[11] L. U. Odhner and A. M. Dollar, "Stable, open-loop precision manipulation with underactuated hands," *The International Journal of Robotics Research*, 2015.

[12] M. V. Liarokapis, B. C. Calli, A. Spiers, and A. M. Dollar, "Unplanned, Model-Free, Single Grasp Object Classification with Underactuated Hands and Force Sensors," *IEEE/RSJ International Conference on Intelligent Robots and Systems*.

[13] R. S. Johansson and J. R. Flanagan, "Coding and use of tactile signals from the fingertips in object manipulation tasks.," *Nature reviews Neuroscience*, vol. 10, no. 5, pp. 345–59, May 2009.

[14] S. J. Lederman and R. L. Klatzky, "Hand movements: A window into haptic object recognition," *Cognitive Psychology*, vol. 19, no. 3, pp. 342–368, Jul. 1987.

[15] N. Wettels, J. a. Fishel, Z. Su, C. H. Lin, and G. E. Loeb, "Multi-modal synergistic tactile sensing," *IEEE/RAS International Conference on Humanoid Robotics*, vol. 105, no. 2, pp. 2–4, 2009.

[16] N. Wettels and G. E. Loeb, "Haptic Feature Extraction from a Biomimetic Tactile Sensor :," pp. 2471–2478, 2011.

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

[18] S. J. Lederman and R. L. Klatzky, "Haptic identification of common objects: effects of constraining the manual exploration process.," *Perception & psychophysics*, vol. 66, no. 4, 2004.

[19] E. B. Vander Poorten, E. Demeester, and P. Lammertse, "Haptic feedback for medical applications , a survey," in *Actuator*, 2012.

[20] A. Spiers, S. Baillie, T. Pipe, and R. Persad, "Experimentally driven design of a palpating gripper with minimally invasive surgery considerations," *2012 IEEE Haptics Symposium*, pp. 261–266, Mar. 2012.

[21] R. L. Klatzky, J. M. Loomis, S. J. Lederman, H. Wake, and N. Fujita, "Haptic identification of objects and their depictions.," *Perception & psychophysics*, vol. 54, no. 2, pp. 170–178, 1993.

[22] L. Birglen, T. Laliberté, and C. M. Gosselin, "Underactuated Robotic Hands," *Springer Tracts in Advanced Robotics*, vol. 40, no. 40, p. 244, 2008.
[23] L. P. Jentoft, S. Member, Q. Wan, and R. D. H. Fellow, "Limits to Compliance and the Role of Tactile Sensing in Grasping," in *2014 IEEE International Conference on Robotics and Automation*.
[24] A. Dollar . R. Balasbramanian, "A Framework for Studying Underactuation in the Human Hand," in *Annual Meeting of the Americal Society of Biomechanics*, 2010, pp. 5–6.
[25] B. Belzile and L. Birglen, "A compliant self-adaptive gripper with proprioceptive haptic feedback," *Autonomous Robots*, vol. 36, pp. 79–91, 2014.
[26] H. Tsutsui, Y. Murashima, N. Honma, and K. Akazawa, "Robot hand with soft tactile sensors and underactuated control," in *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, 2013, pp. 4148–4151.
[27] A. M. Dollar, L. P. Jentoft, J. H. Gao, and R. D. Howe, "Contact sensing and grasping performance of compliant hands," *Autonomous Robots*, vol. 28, no. 1, pp. 65–75, 2010.
[28] L. P. Jentoft and R. D. Howe, "Determining object geometry with compliance and simple sensors," *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3468–3473, 2011.
[29] A. Drimus, M. Borlum Petersen, and A. Bilberg, "Object texture recognition by dynamic tactile sensing using active exploration," in *2012 IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication*, 2012, pp. 277–283.
[30] S. Chitta, M. Piccoli, and J. Sturm, "Tactile object class and internal state recognition for mobile manipulation," in *2010 IEEE International Conference on Robotics and Automation*, 2010, pp. 2342–2348.
[31] R. Andrecioli and E. D. Engeberg, "Grasped object stiffness detection for adaptive force control of a prosthetic hand," *Proceedings of the IEEE RAS and EMBS International Conference on Biomedical Robotics and Biomechatronics*, pp. 526–531, 2012.
[32] A. M. Okamura and M. R. Cutkosky, "Feature Detection for Haptic Exploration with Robotic Fingers," *The International Journal of Robotics Research*, 2001. pp 925-938
[33] U. Martinez-Hernandez, G. Metta, T. J. Dodd, T. J. Prescott, L. Natale, and N. F. Lepora, "Active contour following to explore object shape with robot touch," *2013 World Haptics Conference, WHC 2013*, pp. 341–346, 2013.
[34] I. McMahon, V. Chu, L. Riano, C. G. McDonald, Q. He, J. M. Perez-Tejada, M. Arrigo, N. Fitter, J. C. Nappo, and T. Darrell, "Robotic Learning of Haptic Adjectives Through Physical Interaction," *IROS workshop on Advances in Tactile Sensing and Touch based Human-Robot Interaction*, 2012.
[35] J. M. Romano, K. Hsiao, G. Niemeyer, S. Chitta, and K. J. Kuchenbecker, "Human-inspired robotic grasp control with tactile sensing," *IEEE Transactions on Robotics*, vol. 27, no. 6, pp. 1067–1079, 2011.
[36] I. Bandyopadhyaya, D. Babu, A. Kumar, and J. Roychowdhury, "Tactile sensing based softness classification using machine learning," in *2014 IEEE International Advance Computing Conference (IACC)*, 2014, pp. 1231–1236.
[37] H. K. Lam, U. Ekong, H. Liu, B. Xiao, H. Araujo, S. H. Ling, and K. Y. Chan, "A study of neural-network-based classifiers for material classification," *Neurocomputing*, vol. 144, pp. 367–377, Nov. 2014.
[38] A. Schneider, J. Sturm, C. Stachniss, M. Reisert, H. Burkhardt, and W. Burgard, "Object identification with tactile sensors using bag-of-features," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 243–248.
[39] A. Drimus, G. Kootstra, A. Bilberg, and D. Kragic, "Design of a flexible tactile sensor for classification of rigid and deformable objects," *Robotics and Autonomous Systems*, vol. 62, no. 1, pp. 3–15, Jan. 2014.
[40] H. Soh and Y. Demiris, "Incrementally learning objects by touch: online discriminative and generative models for tactile-based recognition.," *IEEE transactions on haptics*, vol. 7, no. 4, pp. 512–25, Jan. 2014.
[41] M. Madry, L. Bo, D. Kragic, and D. Fox, "ST-HMP: Unsupervised Spatio-Temporal feature learning for tactile data," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 2262–2269.
[42] Y. Tenzer, L. P. Jentoft, and R. D. Howe, "The Feel of MEMS Barometers: Inexpensive and Easily Customized Tactile Array Sensors," *IEEE Robotics and Automation Magazine*, no. c, 2014.
[43] A. Dollar. B. Calli, A. Walsman, Arjun Singh, Siddhartha Srinivasa, Pieter Abbeel, "Benchmarking in Manipulation Research: The YCB Object and Model Set and Benchmarking Protocols," *arXiv:150203143*.
[44] Tin Kam Ho, "Random decision forests," in *Proceedings of 3rd International Conference on Document Analysis and Recognition*, 1995, vol. 1, pp. 278–282.
[45] L. Breiman, "Random forests," *Machine learning*, pp. 5–32, 2001.
[46] L. U. Odhner and A. M. Dollar, "The smooth curvature model: An efficient representation of Euler-Bernoulli flexures as robot joints," *IEEE Transactions on Robotics*, vol. 28, no. 4, pp. 761–772, 2012.
[47] S. Theodoridis and K. Koutroumbas, *Pattern Recognition*, Elselvier Academic Press, 2009.

Adam J. Spiers received the BSc degree in Cybernetics and Control Engineering and an MSc degree in Engineering and Information Sciences from the University of Reading (UK). His PhD degree in Mechanical Engineering is from Bristol University (UK). He conducted 3 years of postdoctoral research in medical robotics and haptics at Bristol Robotics Laboratory in the University of the West of England (UK) He is currently a postdoctoral associate with the GRAB Lab, Department of Mechanical Engineering and Materials Science, Yale University (USA). His current research focuses on human motion, robot grasping, prosthetics and haptic interfaces. He is a member of the IEEE.

Minas V. Liarokapis received the MSc degree in computer engineering from the University of Patras (Greece), the MSc degree in information technologies in medicine and biology from the National Kapodistrian University of Athens (Greece) and the PhD degree in mechanical engineering from the National Technical University of Athens (Greece). He is currently a postdoctoral associate with the GRAB Lab, Department of Mechanical Engineering and Materials Science, Yale University (USA). His research focuses on human robot interaction, robot grasping and manipulation, robot hands design, non-invasive brain machine interfaces, machine learning and pattern recognition. He is a member of the IEEE.

Berk Calli received the Bachelor of Science and Master of Science degrees in Mechatronic Program of Sabanci University, Turkey. His master thesis was on integrated visual servoing and force control for robotic manipulation. He completed his PhD at Delft University of Technology in 2015. His PhD thesis was on optimizing the viewpoint of robot's vision sensor for improving the performance of grasp synthesis algorithms. He is working at Yale University GRAB Lab on vision based manipulation, dexterous manipulation and manipulation benchmarking. He is a member of the IEEE.

Aaron M. Dollar received the BS degree in mechanical engineering from the University of Massachusetts at Amherst, SM, and PhD degrees in engineering sciences from Harvard, and conducted two years of postdoctoral research at the MIT Media Lab. He is currently the John J. Lee associate professor of mechanical engineering and materials science at Yale. His research topics include human and robotic grasping and dexterous manipulation, mechanisms and machine design, and assistive and rehabilitation devices including upper-limb prosthetics and lower-limb orthoses. He received the 2013 DARPA Young Faculty Award, 2011 AFOSR Young Investigator Award, the 2010 Technology Review TR35 Young Innovator Award, and the 2010 US National Science Foundation (NSF) CAREER Award. He is a Senior member of the IEEE.