Robotic Touch:
A Unified Representation for Tactile-Based Assessment of Geometric and Mechanical Properties of Objects

Jacob Green

Supervised by Perla Maiolino

UNIVERSITY OF OXFORD

SOFT ROBOTICS LAB
OXFORD ROBOTICS INSTITUTE

Motivation

The majority object identification techniques applied in robotics use shape as the salient feature for classification, which does not take advantage of the **rich data available from the object's material properties**. Identification of the **material properties** of objects is crucial for both **robot-environment and robot-human interaction**. Alongside the ability to identify objects, tactile-based robotic control is essential in how both humans and robots are able to grasp and manipulate objects [1]. Humans do hold eggs in the same way that they hold whisks!

Object identification using computer vision is not always possible; for example, due to occlusion or poor lighting. It is also nearly impossible to identify visually similar but materially different objects through vision alone, such as a fully inflated ball compared to a flat one.

Objective

To **identify the material stiffnesses of geometrically simple objects** using a gripper equipped with a tactile sensor.

To apply this in order to **discriminate between geometrically similar objects of different**

Hardware and Experiments

Hardware:
A Franka Emika Panda **7 degree of freedom robotic arm** equipped with a two finger gripper was used to gather data. The gripper pads were replaced by two 3D printed fingers affixed with the **CySkin tactile sensor** (skin). The skin consisted of 78 tactile elements (taxels) that produced a response according to the pressure applied. A custom ROS-based (Robot Operating System) software architecture was implemented to integrate the arm, gripper and skin for control and data collection.

Figure 1: Franka Erika Panda Arm and CySkin Tactile Sensor

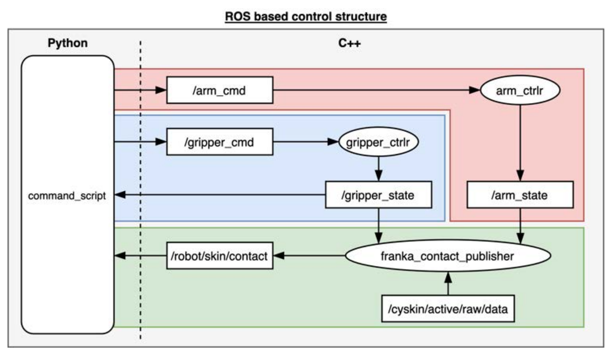

Figure 2: Software architecture for integration of the arm, gripper and skin

Experimental Setup:
Five custom test objects of different stiffnesses were designed for testing. A human-like **'squeezing' exploratory procedure** was used [2]. In each experiment, the object was placed between the gripper fingers and **incrementally compressed** a set number of times, pausing at each width.

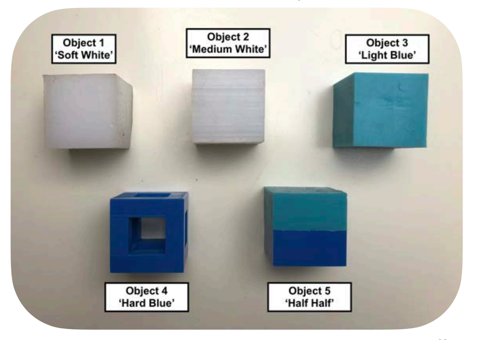

Figure 3: Test objects - Objects 1 to 4 increase in stiffness

Model-Based Approach

Our model-based approach used **Hooke's Law** to model the stiffness of the object at many locations on the object. These **local stiffnesses** were then combined to estimate the **global stiffness** of the object as a whole. An **augmented point cloud** of the object was then generated, which contained both geometric and stiffness information.

Machine Learning-Based Approach

The machine learning approach aimed to classify our test objects using a **single snapshot** of gripper and taxel information (a **haptic frame**). Redundancy of some features (taxels), dimensionality reduction was performed using **Principal Component Analysis** (PCA) projection. Several algorithms were applied and their performances compared.

Methods

Hookean Model:
Objects under test were modelled as linear springs under each taxel that followed Hooke's Law.

Assumptions:
Simplifying assumptions about the behaviour of the object, the gripper fingers and the lycra that covered them included:

- Force was applied over equal area for every taxel.
- The deformation of the lycra and the fingers was was negligible compared to that of the object.
- Objects deform in the **linear part of a force-displacement curve**.

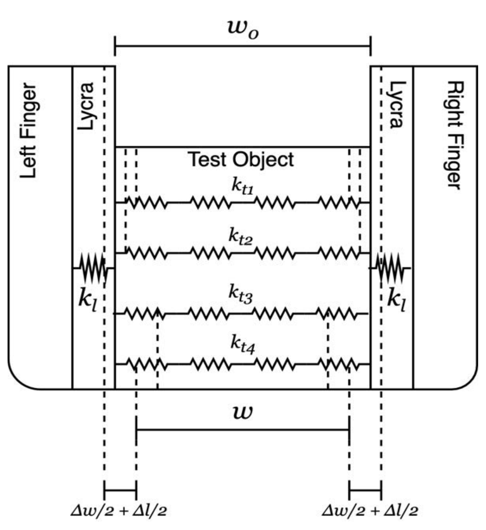

Figure 4: Hookean stiffness model

Modelling Equation:
This resulted in a final static modelling equation for the stiffness of the object at each taxel in contact in terms of the taxel's response, r_t, and the gripper displacement, Δw:

$$r_t = k_t \cdot \Delta w + b$$

Therefore, the **object's stiffness at a taxel location (local) is equal to the gradient, k_t, of the taxel response plotted against the gripper displacement**. The local estimates can then be combined to estimate the stiffness of the entire object (global).

Machine Learning Models:
Three supervised and one unsupervised algorithms were implemented to classify data based on a **single haptic frame** (consisting of the set of taxel responses as well as the gripper width at one instant in time). They were:

- Random Forests (RFs)
- K-Nearest Neighbours (KNNs)
- Support Vector Machines (SVMs)

Supervised

- K-Means Clustering

Unsupervised

These algorithms have been previously shown to perform well when applied to haptic data [2], [3], [4]

Given the ability to classify based on a single haptic frame, accuracy could then be further increased by taking a **majority vote** from a series of haptic frames from a grasp.

PCA Projection:
Given haptic data often has a high degree of redundancy [2], [3], PCA projection was chosen as a **dimensionality reduction** technique to project the data onto the three principal axes that showed the highest variance. This also allowed haptic frames for different objects to be **visualised in three dimensions**.

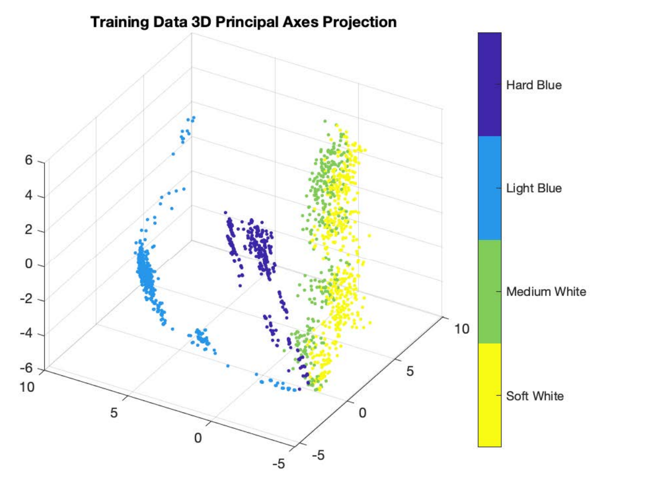

Figure 8: PCA Projection of training data - observably separable

Results

Filtering of raw data:
A series of threshold based filters were applied to the data for each taxel. These had the effect of:

- Removing data points from when the gripper was moving (dynamic points).
- Removing 'sparse' data from taxels that did not make good contact.

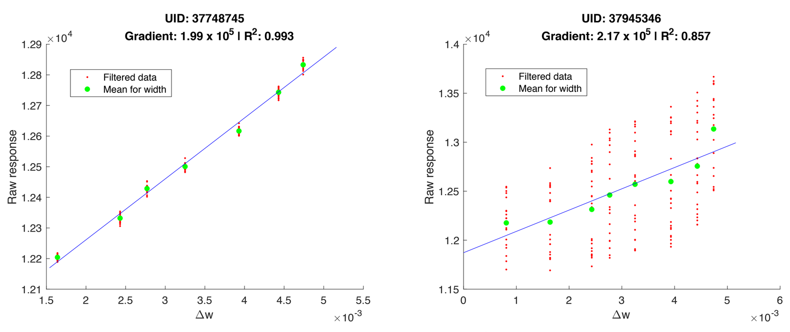

Figure 5: Response v. Gripper displacement plots for two typical taxels (Object 1)

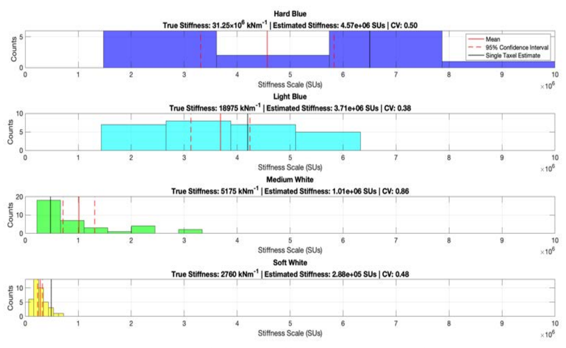

Figure 6: Histograms of local stiffness estimates for Objects 1-4 (solid red line = global estimate)

Local and Global Stiffness Estimation:

- The **gradient of the least squares regression line** fitted to the response-gripper displacement plots was taken as the **local stiffness** estimate at that taxel in arbitrary stiffness units (SUs).
- The mean of all local stiffnesses was used as the estimate of the **global stiffness** for the whole object.
- Relative stiffnesses between the all objects were found to be accurate, and approximately correlated with the actual stiffnesses for the three deformable objects.

Augmented Point Clouds:
Point clouds were generated using known coordinates of taxel-object interactions that were found using the **robot's kinematics**. Two methods for interpolation were applied to estimate the stiffness at locations between taxels:

- **Mean vertex interpolation** - triangles were drawn between taxels and then resampled. The stiffness assigned to these points was equal to the mean stiffness from the triangles' vertices.
- **Gaussian Process Regression (GPR)** - a GPR was trained using the known stiffnesses at the locations of taxels and then used to predict stiffnesses at unknown locations.

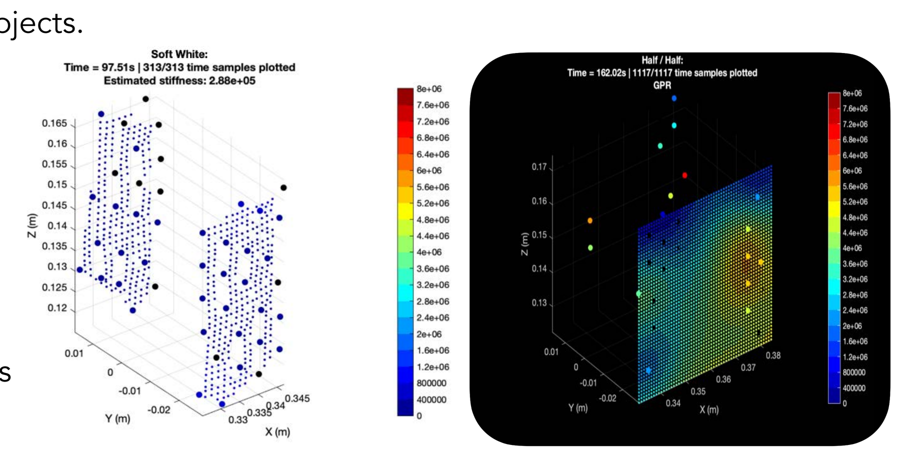

Figure 7: Augmented point clouds with mean vertex (L - Object 1) and GPR interpolation (R - Object 5)

Data Visualisation:
The data for each of the four uniform test objects can be visualised as a series of haptic frames, making up a **static haptic video**. Each frame corresponds to a snapshot of the haptic information at a single time instance.

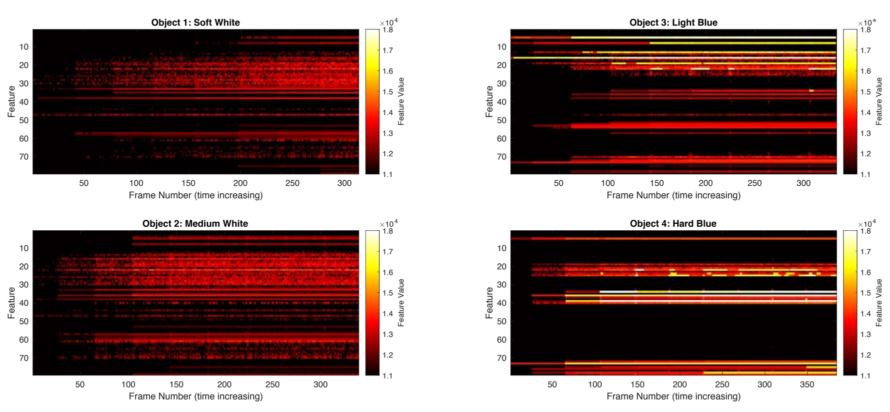

Figure 9: Static haptic videos for Objects 1-4

PCA Projection onto 3 Principal Axes:
The data gathered was then split into training and test sets. A Random Forest was trained, which is able to analyse feature importance and uses a form of inbuilt dimensionality reduction. **Notable redundancy was found** within the feature space and PCA projection was then used for the remaining three models.

- Training and test sets were **normalised separately.**
- The **test set data was then projected** onto the three principal axes of the training set.
- 3 principal components explained **60% of the variance** seen in the training data.

Algorithm	Test Set Accuracy
SVM	89.1%
KNN	92.3%
RFs	100%
KMC	72.7%

Table 1: Algorithm accuracy comparison

Comparison:

- As seen in previous studies [2], [4], RFs were seen to be the most accurate in classifying haptic data. However, as opposed to such studies, we were able to demonstrate **extremely high accuracy based on just one haptic frame**, rather than from several through a grasp.
- KNNs and SVMs were also seen to perform well using a highly reduced feature set projected onto 3 principal components. Increasing the number of components was seen to increase their accuracy but only up to a point still lower than RFs.
- KMC generally performed well but got **confused between the two softest objects**.

Conclusions

We've shown that the stiffnesses of objects can be found and the objects classified using **only squeeze-based haptic information**. Model-based techniques were also used to create **augmented point clouds** containing geometric and stiffness information, and machine learning models were able to **accurately classify objects using grasp data from a single instance in time**.

Key characteristics of each approach include:

Model-Based

- Applicable to unseen objects.
- Surface stiffness distribution estimation possible.
- Used to create augmented point clouds integrable with computer vision

ML-Based

- Accurate classification based on a single haptic frame.
- Higher accuracy possible with more PCs and majority voting.

References

[1] S. J. Lederman and R. L. Klatzky. "Haptic perception: A tutorial". In: *Attention, Perception & Psychophysics71.7 (2009)*, pp. 1439–1459.
[2] A. J. Spiers et al. "Single-Grasp Object Classification and Feature Extraction with Simple Robot Hands and Tactile Sensors". In: *IEEE Transactions on Haptics 9.2 (2016)*, pp. 207–220.
[3] I. Bandyopadhyaya et al. "Tactile sensing based softness classification using machine learning". In: *2014 IEEE International Advance Computing Conference (IACC)*, 2014, pp. 1231–1236.
[4] Z. Flintoff, B. Johnston, and M. Liarokapis. "Single-Grasp, Model-Free Object Classification using a Hyper-Adaptive Hand, Google Soli, and Tactile Sensors". In: *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 1943–1950.