
Tactile Exploration of Object Properties

Taktile Untersuchung von Objekteigenschaften

Bachelor-Thesis von Janine Hölscher aus San Bartolome de Tirajana, Spanien

Tag der Einreichung:

1. Gutachten: Tucker Hermans PhD

2. Gutachten: Prof. Jan Peters PhD



Tactile Exploration of Object Properties
Taktile Untersuchung von Objekteigenschaften

Vorgelegte Bachelor-Thesis von Janine Hölscher aus San Bartolome de Tirajana, Spanien

1. Gutachten: Tucker Hermans PhD
2. Gutachten: Prof. Jan Peters PhD

Tag der Einreichung:

Erklärung zur Abschlussarbeit gemäß § 22 Abs. 7 APB der TU Darmstadt

Hiermit versichere ich, Muster Mustermann, die vorliegende Master-Thesis / Bachelor-Thesis ohne Hilfe Dritter und nur mit den angegebenen Quellen und Hilfsmitteln angefertigt zu haben. Alle Stellen, die Quellen entnommen wurden, sind als solche kenntlich gemacht worden. Diese Arbeit hat in gleicher oder ähnlicher Form noch keiner Prüfungsbehörde vorgelegen.

In der abgegebenen Thesis stimmen die schriftliche und elektronische Fassung überein.

Datum:

Unterschrift:

Thesis Statement pursuant to § 22 paragraph 7 of APB TU Darmstadt

I herewith formally declare that I have written the submitted thesis independently. I did not use any outside support except for the quoted literature and other sources mentioned in the paper. I clearly marked and separately listed all of the literature and all of the other sources which I employed when producing this academic work, either literally or in content. This thesis has not been handed in or published before in the same or similar form.

In the submitted thesis the written copies and the electronic version are identical in content.

Date:

Signature:

Zusammenfassung

Roboter können eine immer größere Bandbreite von Aufgaben immer unabhängiger ausführen. Dafür wird es zunehmend wichtig, dass sie ihre Umgebung wahrnehmen können. Taktile Sensoren sind besonders hilfreich für Manipulationsaufgaben, da Erkundungs- und Manipulationsbewegungen gleichzeitig ausgeführt werden können. Sie sind besonders nützlich, um Objekte und Materialien wieder zu erkennen und sie entsprechend ihrer Eigenschaften zu behandeln. In dieser Arbeit verwenden wir den BioTac finger, einen multimodalen, taktilen Sensor. Dieser ist an einem PA10 Roboterarm montiert und kann Temperatur, Verformung seiner elastischen Haut und Druck messen.

Wir vergleichen verschiedene Ansätze, um Objekte anhand ihrer Oberflächeneigenschaften wieder zu erkennen. Der Roboter hat Erkundungsbewegungen an 49 Objekten ausgeführt, die aus verschiedenen Materialien gemacht sind. Anhand dieses Datensatzes vergleichen wir unterschiedliche Methoden, um Eigenschaften der Objekte zu bestimmen, die sich auf die verschiedenen Sensoren des Fingers beziehen. Für die Wiedererkennung der Objekte verwenden wir verschiedene überwachte Klassifizierungsmethoden, wie beispielsweise Support Vector Machines oder Random Forests. Außerdem kombinieren wir die Informationen aus verschiedenen Erkundungsbewegungen, um eine genauere Beschreibung der Objekte zu erhalten. Des Weiteren erwägen wir, ob die Objekteigenschaften genügend Informationen enthalten, um auch andere Objekte desselben Materials zuordnen zu können. Für diese Frage haben wir sowohl überwachte als auch nichtüberwachte Klassifizierungsverfahren verwendet. Wir suchen auch nach der bestmöglichen Kombination von Erkundungsbewegungen und Methoden zur Eigenschaftsbestimmung und Klassifizierung. Auf diese Weise erreichen wir eine korrekte Klassifikation von 97.55% der Testobjekte.

Abstract

Perceiving their environment becomes more and more important for robots that can perform a large variety of tasks independently. Tactile sensors are especially helpful for manipulation tasks, because exploration and manipulation movements complement each other. They are especially useful to identify and recognize objects and materials and handle them due to their properties. In this work we use the BioTac finger, a multimodal, tactile sensor. It is mounted on a PA10 robot arm and measures temperature, fingertip deformation and pressure.

We compare different approaches to recognize objects considering their surface properties. The robot performed exploration movements on 49 objects belonging to different material classes. We compare various methods to determine meaningful features from the collected raw data. These are taking different sensor signals into account. For the object recognition, we use different common supervised classification methods such as Random Forests or Support Vector Machines. In addition, we consider combining information gained from different exploratory movements. We also determine whether the features provide sufficient information to describe materials appropriately. For this classification task we use both supervised and unsupervised methods. Furthermore, we consider the combination of the proposed exploration, feature calculation and classification methods and compare their performance. This way we reach a classification accuracy of 97.55%.

Foremost, I would like to thank my advisor Tucker Hermans, who always motivated me and who found answers to all of my questions. I am also grateful to Jan Peters and the IAS group for giving me the opportunity to write this thesis.

Furthermore, I would like to thank my friends, who supported me during the last months. I am especially thankful to Julian Gross for his last-minute support, Jens Krüger, who showed me the wonderful world of Photoshop and Fiona Paul for advice on language matters.

List of Figures and Tables

1	Schematic diagram showing the structure and the sensors of a BioTac finger.	2
2	The PA10 robot equipped with a SHUNK force torque sensor and a BioTac finger.	3
4	49 objects grouped by their material. From top to bottom: plastic, metal, paper, fabric, ceramic, stone, wood, sponge. The last two rows of objects do not belong to any material class.	5
5	Electrode data samples for a slide movement on the blue plastic box.	6
6	Temperature signals over time for static contac on the metal box and on the beige sponge.	8
7	Number of features before and after applying PCA and the percentage of variance preserved for all feature calculation methods.	8
8	Structure of a random tree with a depth of three.	9
9	Objects examined in the pilot study.	10
10	Classification accuracy for each object with a varying number of training samples per object, feature calculation method F2 and linear SVM.	10
11	Classification accuracy for each feature calculation method, averaged over all movements and classification methods, with 9 training samples per object.	11
12	Classification accuracy for each feature calculation method and each movement, averaged over all classification methods, with 9 training samples per object.	12
13	Classification accuracy for each feature calculation method and classification method, averaged over all movements, with 9 training samples per object.	12
14	Classification accuracy for all possible combinations of finger movements, feature calculations and classification methods.	13
15	Classification accuracy for each movement with feature calculation F4, classification methods Naive Bayes and linear SVM and 9 training samples per object.	14
16	Classification accuracy for each movement combination with an increasing number of movements, feature calculation F4, classification methods Naive Bayes and linear SVM and 9 training samples per object.	14
17	Classification accuracy for the concatenation of all movements with feature calculation F4, classification method linear SVM and an increasing number of training samples per object.	15
18	Classification accuracy for each of the 23 sensor signals averaged over all movements, with classification method linear SVM and 9 training samples per object.	15
19	BioTac electrode layout. The red circled electrodes are among the the 5 best ranked features.	16
20	Classification accuracy for each of the 23 sensor signals for static contact, with classification method linear SVM and 9 training samples per object.	16
21	Confusion matrix for all 49 objects, with concatenation of all movements, feature calculation F4, classification method linear SVM and 9 training samples per object.	17
22	Classification accuracy for objects of each material class separately with concatenation of all movements, feature calculation F4, classification method linear SVM and an increasing number of training samples per object.	18
23	Classification accuracy for each material with concatenation of all movements, feature calculation F4, classification method linear SVM and an increasing number of training samples per object.	18
24	Confusion matrix for each material with leave one out classification for each object, concatenation of all movements, feature calculation F4 and classification method linear SVM.	19
25	Confusion matrix for each cluster and each object, concatenation of all movements, feature calculation F4 and classification method K-means.	19
26	Objects examined with additional movements.	20
27	Classification accuracy for each object, with concatenation of all movements, feature calculation F4, classification method linear SVM and an increasing number of training samples per object.	20

Contents

1	Introduction	1
1.1	Related Work	1
2	Experimental Setup	2
2.1	Hardware	2
2.2	Movements	3
2.3	Objects	4
3	Data Evaluation	6
3.1	Data Processing	6
3.2	Feature Calculation	6
3.3	Classification Methods	8
4	Results	10
4.1	Pilot Study	10
4.2	Main Experiment	10
4.2.1	Object Classification	11
4.2.2	Movement Combination	13
4.2.3	Feature Ranking	15
4.2.4	Materials	16
4.2.5	Clustering	19
4.3	Additional Movements	20
5	Conclusion and Future Work	21
	Bibliography	22

1 Introduction

In our modern world, robots become more and more important as support in every possible situation. Due to technical advances, they do no longer only help out in clearly defined industrial processes, but they also have arrived in our everyday lives. Robots can be found in uncertain, continually changing environments. Additionally, the variety of tasks one robot can perform increases and so does the complexity of these tasks. These processes become too complex to be controlled manually. For this reason robots need to become more and more autonomous in gaining information about their environment and deciding for performing actions based on this information. That is why different kinds of sensing technologies become important. So far, mostly visual sensors are used. These can be helpful for scene recognition, object detection, segmentation [1] or identification of a position depending on the environment. But they also have limits: Important parts of a scene can be occluded and certain features can be difficult to discern from a distance or a specific angle.

Tactile sensing can be used to fill some of these gaps and collect additional information, because the approach different. While cameras can operate without any movements, touch is an interactive sense and exploratory movements have to be made. That is why tactile sensing is naturally helpful for all sorts of manipulation tasks. Both exploration and manipulation can be executed with the same devices. Properties determined by tactile sensing match the information needed for manipulation.

Tactile information can be useful in a large variety of tasks. It can be used to localize objects or to determine their shape and texture information. This way objects can be identified or it can be determined what kind of manipulation would be suitable. Some parts of these properties can be generalized and objects can be classified depending on their material or function and objects can be treated in an appropriate way for this class .

As our world is designed by human beings for human beings, many manipulation tasks require human-like interaction and tactile sensing. The BioTac is a finger-sized sensing device, which provides some of the human sensing abilities. It can measure temperature, vibrations and forces applied to its skin. In this work we want to determine the possibilities and boundaries of the BioTac regarding the recognition of objects and their materials.

1.1 Related Work

Previous studies show the large variety of tasks robots are able to perform using tactile sensing. Romano et al. [2] demonstrated grasping of objects with a gripper with pressure arrays on its fingertips. By evaluating the pressure sensor data, the minimal force was determined which is necessary such that the object neither slips nor becomes damaged. Chitta et al. [3] used a gripper with 22 individual capacity sensor cells on each finger to determine the amount of liquid in beverage containers. To achieve this, they grasped the container and rolled it from side to side. Pape et al. [4] used reinforcement learning to let a robot learn tactile exploration skills autonomously. This robot is able e.g. to find out wether it moves without contact or taps on a surface.

We want to concentrate on the classification of objects. On the one hand, this can be achieved by detecting features of the object's shape, such as edges [5]. On the other hand, their different surface structures can be exploited. We focus on the latter approach. In the past, different methods to measure these surface features have been proposed. Dallaire et al. [6] used a turntable with a three-axis accelerometer to distinguish between 28 disks made of different materials. A similar sensor built into an artificial fingernail was used by Sinapov et al. [7] to discriminate objects via five different exploratory scratching movements. The BioTac finger is a multimodal sensor other researchers have already achieved satisfying results with. Fishel and Loeb [8] were able to distinguish between 117 textures with a 95.4% success rate using the BioTac. They provided an active learning method to choose between 36 different sliding movements with variation in speed and force.

Apart from sliding there are other stereotypical exploratory procedures human beings perform, such as static contact or enclosing an object[9]. Chu et al. [10] focused on testing those finger movements regarding the object's surface structure, like applying pressure on an object. They did not only recognize ojects, but also used the collected data to learn and recognize labels describing human sensations when touching the object. Lepora et al. [11] proposed an active perception approach that controls the sensor movements. The process decides for the exploratory movement that most improves the knowledge about the examined object.

There are many different approaches on how to gain meaningful features from collected raw data. It is possible to distinguish objects with simple spatial features, as provided by Xu et al. [12]. They were able to classify 10 objects with three features regarding temperature, compliance and texture of the object's surface. Tanaka et al. [13] simply use the mean value over time for each sensor. More specific information can be determined with feature calculations regarding temporal changes of features and explicit fits to the expected changes of each sensor signal [10].

2 Experimental Setup

In this section, we explain our experiments in a detailed way. First of all, we point out the hardware setup and the way the objects are presented to the finger. Then, we describe the movements the finger performs to capture the objects. Finally, we present the objects we examine and the criteria for choosing them.

2.1 Hardware

We used SynTouch’s BioTac, a fingertip-sized, multimodal, tactile sensor. It consists of a rigid core enveloped by an elastic skin, which is filled with an incompressible conductive fluid. Both the skin’s surface structure, which is similar to a fingerprint and its elasticity imitate a human fingertip. There are no further sensors or other electrical components placed in the skin. That is why the sensor is very robust and force up to 50N can be applied. The BioTac provides 3 sensor modalities:

When the elastic skin deforms because of pressure appliance, 19 electrodes attached to the core measure the changes of impedance due to the conductive fluid. As the BioTac’s core temperature is approximately 10 °C higher than room temperature, heat flow into a touched object can be measured by the thermistor. The pressure transducer measures vibrations when the BioTac slides over a surface [8] [14].

The collected data consists of 5 types of tactile signals:

- 19 electrode impedances $E1 - 19$
- low-frequency fluid pressure P_{DC}
- high-frequency fluid vibrations P_{AC}
- temperature T_{DC}
- temperature changes T_{AC}

This results in a total of 23 data channels. The P_{AC} channel is sampled at 2200 Hz, the others at 100 Hz.

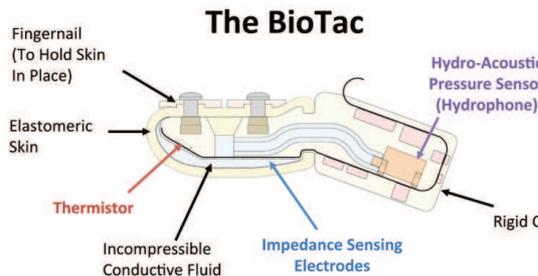


Figure 1: Schematic diagram showing the structure and the sensors of a BioTac finger.

The BioTac finger is mounted to a SHUNK force torque sensor. This sensor is fixed on a Mitsubishi PA10, a robot arm with seven degrees of freedom. It is controlled with SL, a robotics simulation and control engine. The desired end effector velocities of the robot are achieved by calculation of the inverse manipulator Jacobian. BioTac Data is internally sampled and included in SL data bags, which are transmitted to the control PC via USB connection. Before the experiment starts, the BioTac has to warm up for 15-20 minutes to reach its core temperature [12]. For our classification experiments the examined objects are placed on a table in front of the robot. To avoid that objects are moved by the finger, some of them are fixed on the table. This occurs mostly for objects with a small weight and a smooth surface. The complete setup can be seen in Figure 2.



Figure 2: The PA10 robot equipped with a SHUNK force torque sensor and a BioTac finger.

2.2 Movements

According to Xu et. al [12] there are three different typical exploratory movements made by humans that require only cutaneous information: applying pressure, static contact and lateral sliding. In contrast to these movements, others such as enclosure or hefting also need proprioceptive information (positions and forces applied on forces and muscles [9]). We decided to execute static contact and lateral sliding to different directions, as well as a combination of both movements: finger rotation in one position. In the following paragraph we describe each of the movements in detail.

- Static contact:

First of all, contact to the object must be established. The robot end effector with the mounted BioTac moves down until contact is measured. To standardize the pressure applied on the object, the baseline value of BioTac P_{DC} pressure signal before contact is established. When the signal exceeds a certain threshold value, the movement stops. Then, contact is maintained for 15 seconds. Experiments show (see Figure ??) that this timespan is necessary to reach a thermal equilibrium with the object[12]. Afterwards, the finger is moved back to its initial position.

- Lateral sliding:

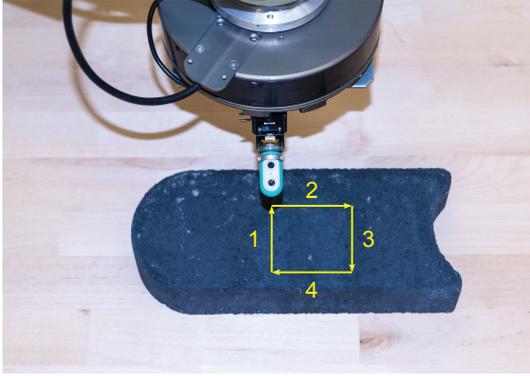
For the sliding movements contact is established in the same way as for the static contact. Subsequently, four sliding movements to different directions are executed. Each movement has a distance of 10 cm and a velocity of 2.5 cm/s. It was shown that it is good for determining a surface's roughness [8]. As it can be seen in Figure 3a, the finger moves into the positive direction of the y-axis, the positive direction of the x-axis, the negative direction of the y-axis and the negative direction of the x-axis one after another. Finally, the finger moves back up.

In the following chapters, these movements are referred to as backward, left, forward and right.

- Rotation:

We differ between two rotation movements. For both we use the same procedures for contact establishment as mentioned before. The first movement is an end effector rotation around the y-axis. It is called roll in the next chapters. The second movement is a rotation around the z-axis and will be referred to as rotate. Because the finger sensor is not located in the center of the end effector coordinate system, compensation movements are necessary to ensure, that the finger stays in the same position. Before the actual movement starts, the finger is rotated in the opposite direction to extend the movement. Each rotation takes 20 seconds with a velocity of 0.01m/s. Afterwards, the initial finger position is headed for.

This yields a total of 7 movements we look at separately.



(a) Arrows indicate the sliding movements the finger executes on the object surface.



(b) The robot's end effector's right-hand coordinate system.

We choose a different starting position for each movement execution. On the one hand, we want to ensure the same conditions for each temperature measurement and avoid influences from former contact. On the other hand, we want to avoid overfitting and create more robust datasets capturing most of each surface's variations.

2.3 Objects

The objects we chose to analyse have to fulfill several requirements to ensure similar conditions for the data collection of each object: First of all, each object has to stand on a table without any further support. Further, they all needed to have a straight surface to simplify the interaction with the sensor. The object's surface has to be greater than 15 x 15 cm to ensure permanent contact with the BioTac finger during data collection and to allow some variation in the position of the measurements. The surfaces should not have any additional structures that vary from the typical objects material's structure. As the BioTac's silicone skin can be damaged easily, objects with sharp edges had to be avoided. Furthermore, the objects should belong to one of the following material classes: wood, ceramic, stone, plastic, sponge, paper, metal or fabric. We found 5 objects with different characteristics for each class, which satisfied all the conditions. The objects' features also vary in terms of thermal conductivity, texture roughness and compliance. Additionally, we chose 9 objects that do not belong to any of the classes or which have very specific surface properties. This results in a set of 49 objects, which is shown in Figure 4. Most of them are household items or can be found in the garden.



Figure 4: 49 objects grouped by their material. From top to bottom: plastic, metal, paper, fabric, ceramic, stone, wood, sponge. The last two rows of objects do not belong to any material class.

3 Data Evaluation

In this section we describe how we convert collected raw datasets into specific features. Furthermore, we explain the classification methods we used.

3.1 Data Processing

We extract the BioTac data for each sensor from the transmitted SL-datafiles for each timestep and create separate datasets for every executed movement, every object and every iteration of this measurement. We then extracted data for time intervals specific for every movement. This way, we avoided noisy sensor signals due to contact establishment in the beginning of each movement. This is especially important for sliding movements (see Figure 5), because the elastic skin is pushed to the opposite direction of the current movement. This is the case after a first contact establishment as well as after a direction change.

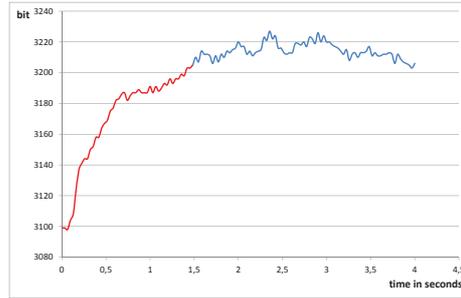


Figure 5: Electrode data samples for a slide movement on the blue plastic box.

Our measurements took place in a time period of three days. During this time, the environment changed, which results in shifts in the measured data (Figure ??). For this reason we subtract the mean of the first 50 timesteps from channel to avoid big shifts.

3.2 Feature Calculation

In order to create features for each collected data sample we chose 7 different ways to extract features from the raw data. This collection of feature calculation methods allows us to compare their performance in the classification of the objects they represent. The first two feature calculation methods are special methods designed for the BioTac sensors. The third and fourth method use unmanipulated raw data and the last three use only one of the 3 sensor modalities each.

- F_1 : Simple BioTac features
Xu et. al [12] provide one simple feature per sensor modality:

- $f_1 : \log\left(\frac{\Delta_{jointangle}}{\Delta E18}\right)$
- $f_2 : \log\left(\text{var}\left(P_{AC,filtered}\right)\right)$
- $f_3 : \Delta T_{AC}$

As we only had access to the BioTac data, we were not able to use the angle and we took only the electrode data for f_1 instead.

- F_2 : Temporal BioTac features
Chu et al. [10] suggest a more complex featureset, which also takes the signal changes over time into credit.
 1. P_{DC} features:
 - f_1 : maximum P_{DC}
 - f_2 : mean P_{DC}
 - f_3 : Smooth P_{DC} data with a Hanning window and determine the greatest change in the signal over time

2. P_{AC} features: Convert the data into an energy spectral density (ESD)
 - f_4 : area under the ESD curve
 - f_5 : weighted average over ESD
 - f_6 : spectral variance
 - f_7 : spectral skewness
 - f_7 : spectral kurtosis
3. T_{DC} and T_{AC} features:
 - f_8 : area under the T_{AC} curve
 - f_9 : time constant of an exponential fit of T_{DC} over time.
4. Electrode features:
 - $f_{10} - f_{22}$: Use Principal Component Analysis (PCA) and fit fifth order polynomials to the first two principal components. The features are the polynomial coefficients.
5. robot features:
 - f_{23} : mean of the aperture gripper distance
 - f_{24} : minimum of the aperture gripper distance
 - f_{25} : range of the gripper's vertical position

As we only had access to the BioTac data, we did not use $f_{23} - f_{25}$.

- F_3 : Raw Data
We concatenate the values for all 23 data channels over time to produce one single vector, where every sensor signal at every timestep is a feature.
- F_4 : Raw Data Means
We calculate the mean of each signal channel over time, which results in a total of 23 features.
$$f_i = \frac{1}{N} \sum_{t=1}^N signal_{i,t}$$
- F_5 : pressure data
Dallaire et al.[6] used these feature calculations for a Micro Electro-Mechanical Systems (MEMS) accelerometer: We calculated the features for the P_{AC} signal.
 - f_1 : variance
 - f_2 : skewness
 - f_3 : kurtosis
 - f_4 : fifth moment
 - f_5 : sum of the variation over time
 - f_6 : number of times 20 uniformly separated thresholds are crossed
 - f_7 : sum of higher half of amplitude spectrum
- F_6 : electrode data
The electrode data is filtered by a first-order Butterworth filter and then we calculate the Euclidean Norm to combine all signals into a single signal:
$$f(t) = \left(\sum_i^N signal_i^2 \right)^{1/2}$$
 - f_1 : mean of $f(t)$
 - f_2 : variance of $f(t)$

Chitta et al. [3] designed these features for tactile sensor cells directly attached to a gripper.
- F_7 : temperature data
Based on our experience with the former mentioned feature calculations we developed a set of features considering only temperature data mainly for the static contact movement. Distinguishing values for the analog derivative of temperature (T_{AC}) can be seen after about 6 seconds of contact (Figure 6a). After an initial pike the temperature signal (T_{DC}) remains nearly constant. We take the mean for both T_{AC} and T_{DC} value for the

remaining 9 seconds. Another distinguishing feature for the T_{AC} signal is the change after the initial peak, which we examine in f_1 .

- $f_1 = \Delta T_{AC}$
- f_2 : mean of T_{AC}
- f_3 : mean of T_{DC}

To avoid numerical problems, we normalize each feature calculation.

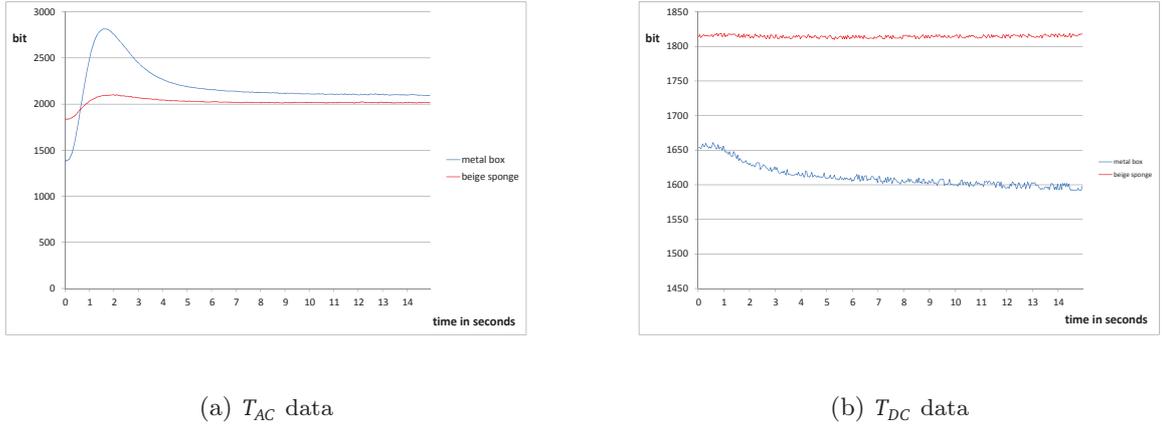


Figure 6: Temperature signals over time for static contact on the metal box and on the beige sponge.

In the following chapters, the features calculation methods will be referred to as F1-F7 in the same order we used in this paragraph.

We also had to reduce the dimensionality for some featuresets, because some classification methods (e.g. Naive Bayes) can only be used with more samples per class than feature dimensionality per sample. We did so applying Principal Component Analysis (PCA) [15]. In Figure 7 we show the results of the feature reduction for each feature set.

featureset	dimensions	dimensions after PCA	preserved variance
F1	3	3	1
F2	22	8	0,9883
F3	-	8	0,5412
F4	23	8	0,9783
F5	7	7	1
F6	2	2	1
F7	3	3	1

Figure 7: Number of features before and after applying PCA and the percentage of variance preserved for all feature calculation methods.

3.3 Classification Methods

We use several common, supervised classification methods to train decision models based on our training data sets.

- Naive Bayes Classifier:

Naive Bayes is a classifier based on Bayes' theorem, which describes the relation between prior and posterior probabilities that a data sample X was collected examining object O_i .

$$P(O_i|X) = \frac{P(X|O_i) * P(O_i)}{\sum_{j=1}^N (P(X|O_j)P(O_j))} \quad (1)$$

The class-conditional probabilities can be modeled with a Gaussian distribution function [16].

$$p(X|O_i) = \frac{1}{\sigma_i \sqrt{2\pi}} e^{-\frac{1}{2} \left(\frac{X-\mu_i}{\sigma_i} \right)^2} \quad (2)$$

- Support Vector Machines:

A support vector machine is a distribution-free, discriminative model. It attempts to find a single decision boundary that maximizes the (SVM) margin between the classes. The classification margin is defined as the sum of the distances of the closest samples of each class, called support vectors. The most simple form is a linear SVM. If data samples are not linearly separable, they can be mapped into high-dimensional feature spaces, where they are separable. A SVM model can do so efficiently with the so called kernel trick. This way, the classification is reduced to a quadratic programming problem. [17]

$$f(x) = \sum_{j=1}^N \alpha_j K(x_j, x) + b \quad (3)$$

α_i is the weight associated with support vector x_i and b is a constant offset. In order to deal with overlapping classes, it is possible to use the Soft Margin method that allows some data points to be misclassified. This is done introducing a slack variable ξ , which describes the distance from the margin boundary for each data sample and is penalized with a constant value C .

We use a linear SVM and a radial basis SVM [6] with a kernel that takes the form:

$$K(x_i, x_j) = \exp(-\lambda \|x_i - x_j\|^2) \quad (4)$$

where λ is the radius. For this SVM, we also used the soft margin method. In order to identify optimal parameters for the kernel function and the penalty for misclassified samples, we executed a grid search [18].

- Random Forest:

A random forest is a set of binary decision trees. These trees are independently trained classifiers. A decision tree's inner nodes represent splitting conditions. The thresholds for these conditions are set when the tree is trained. The leaves contain probabilities for the training samples to reach this node. Based on the decisions on each node, a test sample can traverse the tree and the probability distributions in the leaf assigns the most probable label. With different random feature subsets for the training of each tree, randomization can be achieved. A classification result is the average classification of all trees in the forest. [19] We trained 100 trees with a maximal depth of 3 nodes.

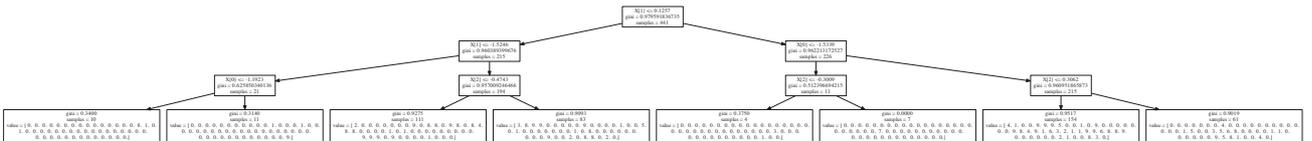


Figure 8: Structure of a random tree with a depth of three.

We also use an unsupervised clustering method:

- K-means Clustering:

The goal of k-means clustering is to partition the data into k clusters, where each data point belongs to the cluster with the nearest mean. Therefore the following formula has to be minimized:

$$J = \sum_{i=1}^k \sum_{x_j \in C_i} \|x_j - \mu_i\|^2 \quad (5)$$

where C_i are the clusters, μ_i is the mean for all elements belonging to a cluster and k is the number of clusters. This can be obtained with iterative refinement by an expectation-maximization (EM) algorithm: During the expectation step, new means are calculated for each cluster depending on the assigned data points. In the maximization step, every data point is assigned to the cluster with the nearest mean[16].

We have used the python machine learning library scikit-learn [20] for random forest and clustering methods and the SVM library libSVM [21]. Feature calculations are implemented in MATLAB.

4 Results

We have collected data for three different datasets, which vary in terms of objects and executed movements. We also tested different amounts of feature calculation and classification methods.

The pilot study was a small experiment testing the experimental setup and more generally the whole classification process. Our main experiment repeated this experiment with small adjustments regarding the data collection process and a larger set of objects, movements, feature calculations and classification methods. The third experiment serves to evaluate the potential of two more finger movements.

4.1 Pilot Study

Our first dataset consists of measurements for a subset of five objects (Figure 9), which have significantly different features regarding compliance, texture roughness and thermal conductivity. We executed one sliding movement (backward) with a velocity of 3 cm/s for each object ten times and we used the force torque sensor to determine whether contact between the finger and an object was established.



Figure 9: Objects examined in the pilot study.

We used the two feature calculation methods made for the BioTac. As classification methods we chose a linear SVM and Naive Bayes. We randomly split the collected data into training and validation data. As Figure 10 shows, we used training sets with different sizes. With six training samples and one validation sample per object, we obtain correct classification for all objects except from metal. With nine training samples we obtain a classification accuracy of 98%.

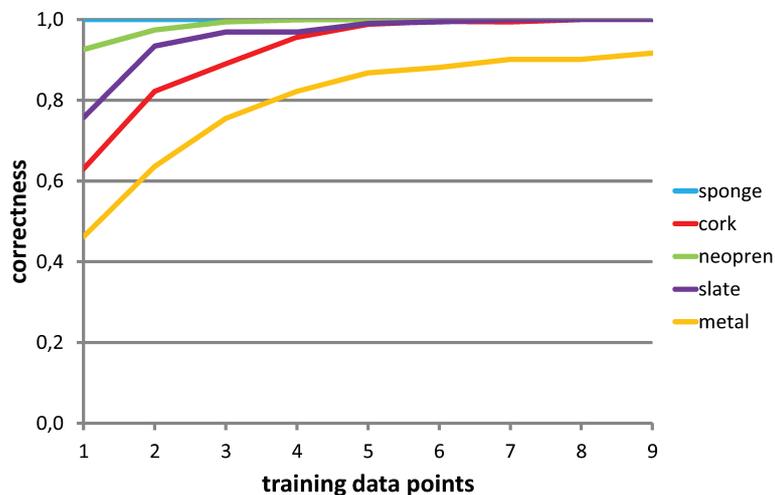


Figure 10: Classification accuracy for each object with a varying number of training samples per object, feature calculation method F2 and linear SVM.

4.2 Main Experiment

After the positive results from our first study, we enlarged our experiment: We used the whole set of 49 objects and executed the four sliding movements and the static contact movement. We did so ten times for each object, which results in a total of 2450 data sets. In contrast to the pilot study, we adjusted the experimental setup to the proceedings described in chapter three. We tuned the sliding velocity to 2.5 cm/s and introduced the method to establish contact with an object using BioTac P_{DC} data. Furthermore, we used a slightly steeper angle between the finger and the object in order to allow the finger to move more freely and avoid contact of the table and the screws fixing the finger.

4.2.1 Object Classification

After manipulating the data and calculating feature values to represent each dataset with all seven proposed calculation methods, we execute the following procedure for each possible combination of feature calculation and movement: First, we randomly choose one test sample for each object. The remaining nine featuresets are used as training samples. This results in 49 test samples and 441 training samples. Then, we apply each classification method on the created data sets. We repeat this split procedure 100 times and calculate mean and variance for each result. This way, we obtain an average overall classification accuracy of 50.97%

In Figure 11, we display the average results for each feature calculation method. We observe that F4, the raw data means, produces an accuracy of 73.83%, which is more than ten percent higher than the next best method F2, the temporal feature calculation, with 62.43% followed by F3, the raw data and F7, the temperature features. The simple BioTac features, pressure features and vibration features obtain classification accuracies of less than 50%.

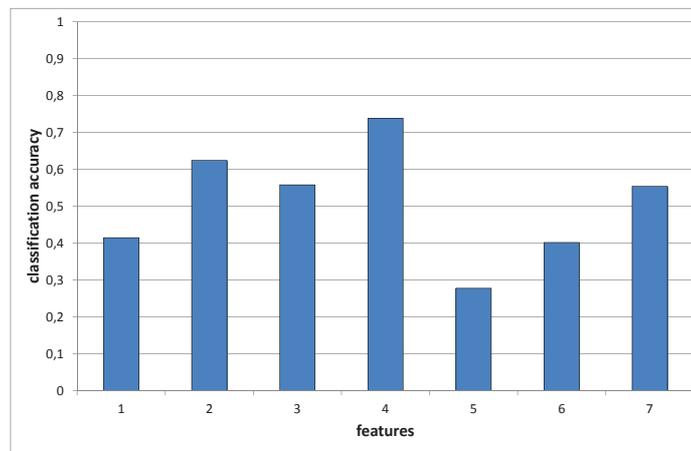


Figure 11: Classification accuracy for each feature calculation method, averaged over all movements and classification methods, with 9 training samples per object.

In Figure 12 it can be seen that F4 obtains not only the highest accuracy results on average, but also when we look at each movement separately. We observe similar results for each classification method in Figure 13. Taking a closer look to Figure 12, it is striking that for all feature calculation methods except for raw data (F4) and temperature data (F7) the most meaningful features are produced when applied to the static contact datasets.

For the comparison of the classification methods (Figure 13), it is noticeable that the most performant methods are Naive Bayes and the linear SVM. Only for F5 the random forest method achieves with 31.63% slightly better results than Naive Bayes with 28.71%. The linear SVM obtains the highest classification accuracy for F2 and F3, whose dimensionality had to be reduced the most drastically to meet the constraints the classification algorithms impose (see table 7).

To reduce the dimensionality of our data analysis, we proceed with a smaller dataset: We choose the best feature calculation method F4 and the two best classification methods Naive Bayes and linear SVM for further evaluation.

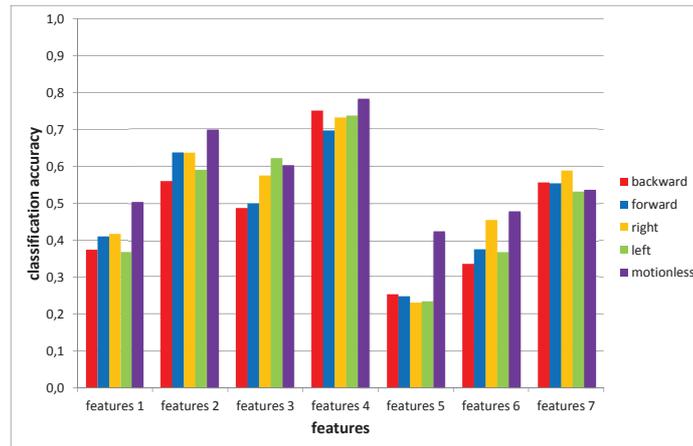


Figure 12: Classification accuracy for each feature calculation method and each movement, averaged over all classification methods, with 9 training samples per object.

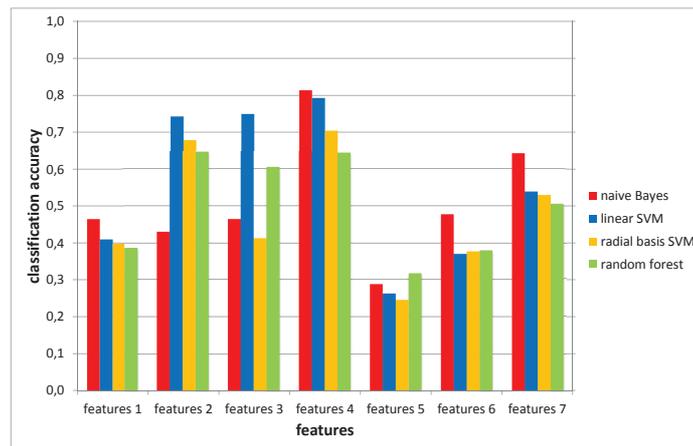


Figure 13: Classification accuracy for each feature calculation method and classification method, averaged over all movements, with 9 training samples per object.

	backward	forward	right	left	motionless
naive Bayes					
F1	0,4312	0,4633	0,4665	0,4082	0,5610
F2	0,3473	0,4235	0,4651	0,3843	0,5586
F3	0,4376	0,4214	0,4463	0,5088	0,5510
F4	0,8100	0,7484	0,8363	0,7886	0,8729
F5	0,2416	0,2447	0,2769	0,2188	0,4663
F6	0,3814	0,4667	0,5276	0,4696	0,5549
F7	0,6533	0,6118	0,6910	0,6080	0,6363
linear SVM					
F1	0,3586	0,4147	0,3955	0,3694	0,4922
F2	0,6637	0,7641	0,7441	0,7312	0,8151
F3	0,6561	0,6790	0,7708	0,8351	0,7971
F4	0,7967	0,7657	0,7755	0,8124	0,7924
F5	0,2194	0,2382	0,1884	0,1927	0,4241
F6	0,3151	0,3284	0,4441	0,3539	0,4353
F7	0,5278	0,5686	0,5753	0,5167	0,5055
radial basis SVM					
F1	0,3561	0,3831	0,3992	0,3618	0,5000
F2	0,6149	0,7163	0,6704	0,6469	0,7557
F3	0,3424	0,3845	0,4810	0,4782	0,3661
F4	0,7135	0,6610	0,6951	0,7124	0,7422
F5	0,2310	0,2306	0,1839	0,1933	0,3918
F6	0,3161	0,3345	0,4547	0,3435	0,4288
F7	0,5282	0,5308	0,5622	0,5129	0,5051
random forest					
F1	0,3500	0,3682	0,4063	0,3471	0,4649
F2	0,5947	0,6527	0,6686	0,6237	0,6894
F3	0,5180	0,5310	0,6069	0,6714	0,6965
F4	0,6629	0,5920	0,6265	0,6202	0,7133
F5	0,3004	0,2547	0,2698	0,3412	0,4163
F6	0,3243	0,3651	0,3955	0,3133	0,5041
F7	0,5310	0,4918	0,5335	0,4798	0,4947

Figure 14: Classification accuracy for all possible combinations of finger movements, feature calculations and classification methods.

4.2.2 Movement Combination

We use three different methods to combine the features obtained by different movements on the same object.

- Concatenation:
We concatenate the features determined for different movements and use this vector for classification.
- Voting:
We classify for each movement separately and decide for the classification result that occurs most often.
- Probabilities:
We classify for each movement separately and sum up the probabilities for each class.

We combine the linear SVM classification with the first two methods. (As SVMs produce no probability distributions, the third method is not possible.) The Naive Bayes approach is combined with the last two methods, as the dimensionality of the feature vectors produced by the first method is too high. We first sort the movements according to their classification accuracy. In Figure 15 it can be seen that the most performant movement is static contact with an accuracy rate of 83.51% and the least performant is the forward sliding movement with 75.31%. We add the movement features to our data set in this order. Figure 16 displays the results of this procedure. The SVM concatenation approach obtains the highest classification accuracy for every movement combination. It results in 97.55% when all movements are combined. It is noticeable that the results for the combination of two movements via voting are lower than those for single movements.

In the next sections we focus on the results of the linear SVM-Concatenation combination. Figure 17 shows the classification accuracy of this approach depending on the number of training samples. With four training samples the classification accuracy is more than 90%.

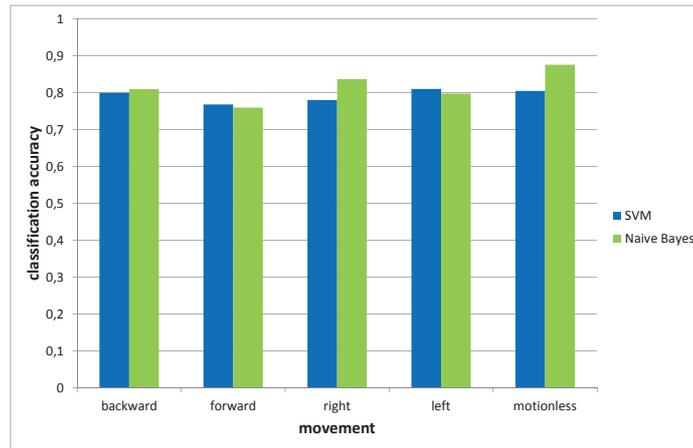


Figure 15: Classification accuracy for each movement with feature calculation F4, classification methods Naive Bayes and linear SVM and 9 training samples per object.

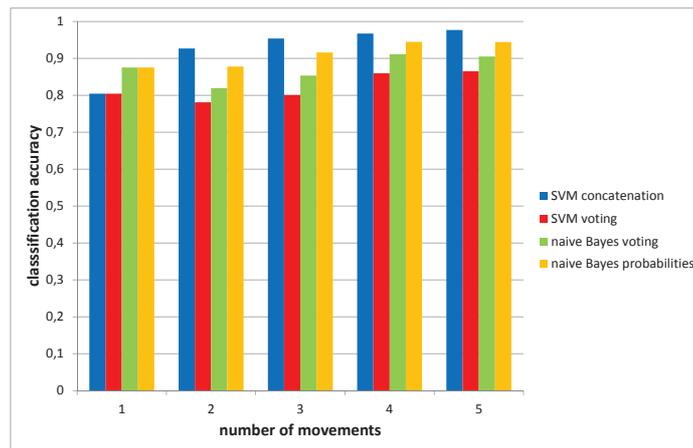


Figure 16: Classification accuracy for each movement combination with an increasing number of movements, feature calculation F4, classification methods Naive Bayes and linear SVM and 9 training samples per object.

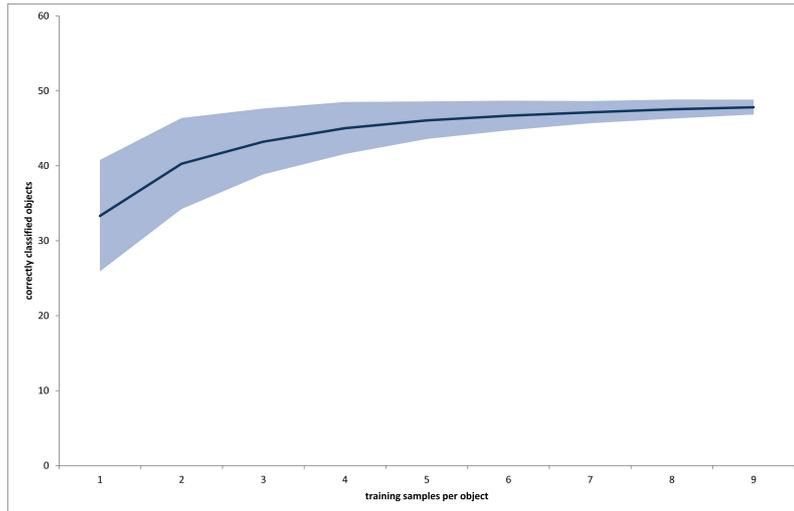


Figure 17: Classification accuracy for the concatenation of all movements with feature calculation F4, classification method linear SVM and an increasing number of training samples per object.

4.2.3 Feature Ranking

In order to determine which features carry most information, we used a wrapper approach [22]. Every feature is used alone for classification with a linear SVM. The resulting classification accuracy decides about the feature's ranking. As featureset F4 consists of the means for every BioTac signal, each feature represents one sensor channel. Figure 18 shows the results averaged over all five movements. It is striking that all four non-electrode sensors are among the five least performing signals. The derivative sensors T_{AC} and P_{AC} have the poorest performance (16%) by a wide margin. The performance accuracies for the electrode sensors are between 20.9% and 28.6%. In Figure 19 the position of each electrode in the finger is shown. The ten best performing electrodes are colored red. It is noticeable that most of them are placed in the right half of the finger. To preclude the possible explanation that the skin is bent to different directions unevenly during the sliding movements, we compare the averaged results with those for the static contact. The performance results for the static contact provided in Figure 20 confirm the average results: The performance of the electrodes on the right side is higher than on the left side of the finger. This points to a slightly uneven table in the experimental setup.

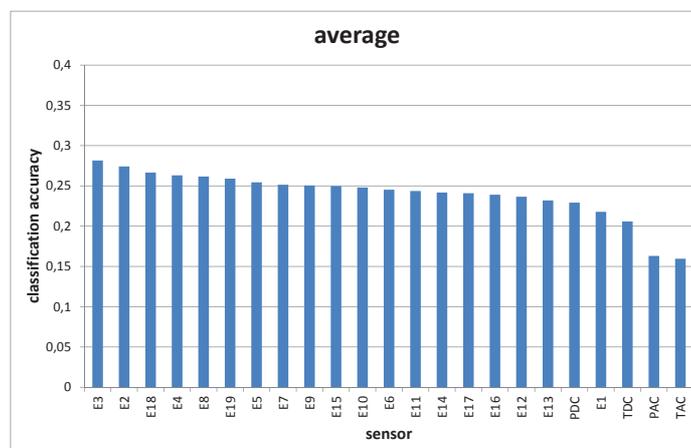


Figure 18: Classification accuracy for each of the 23 sensor signals averaged over all movements, with classification method linear SVM and 9 training samples per object.

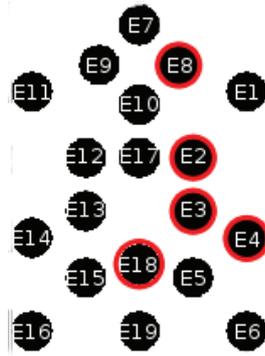


Figure 19: BioTac electrode layout. The red circled electrodes are among the the 5 best ranked features.

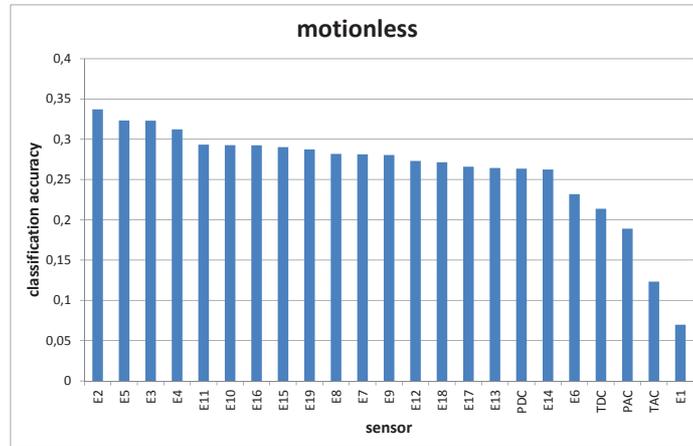


Figure 20: Classification accuracy for each of the 23 sensor signals for static contact, with classification method linear SVM and 9 training samples per object.

4.2.4 Materials

We take a closer look at the accuracy rates for each object in order to determine the possible reasons for missclassifications. The confusion matrix (Figure 21) shows that most missclassifications happen, when an object is mixed up with another one made of the same material. Especially the materials fabric, ceramic and wood are in question.

The same missclassifications can be seen, when we classify the objects of each material class separately. The nine objects that do not belong to any material class are left out. The results are displayed in Figure 22. It is shown that the former mentioned materials obtain the lowest classification accuracies. This observation suggests, that the objects belonging to these classes have a low in-class variance. The objects made of paper, sponge or plastic are perfectly distinguishable with nine training samples per object.

We also classify the objects of each material class separately. To get our training and validation data set, we proceed as before, but label all objects made from the same material identically. The results shown in Figure 23 correspond to our previous observations. Those materials, whose objects are the most difficult to distinguish, such as wood or ceramics, achieve the highest material classification accuracies. But a low in-class variance is no requirement for high material accuracy as the material class sponge shows. The features objects are easy to distinguish (perfect results for the in-class classification), but they differ even more from other objects' features.

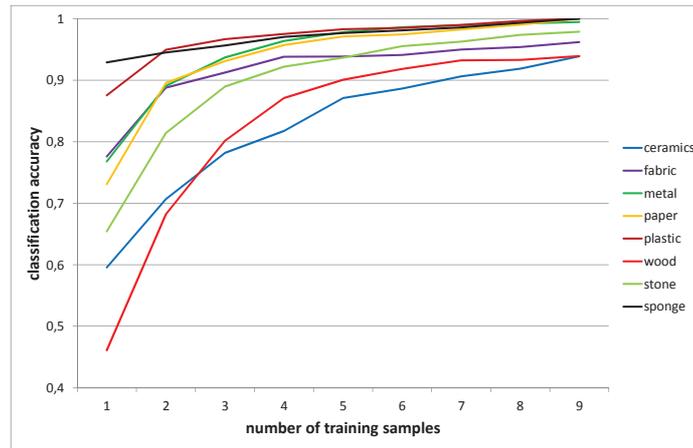


Figure 22: Classification accuracy for objects of each material class separately with concatenation of all movements, feature calculation F4, classification method linear SVM and an increasing number of training samples per object.

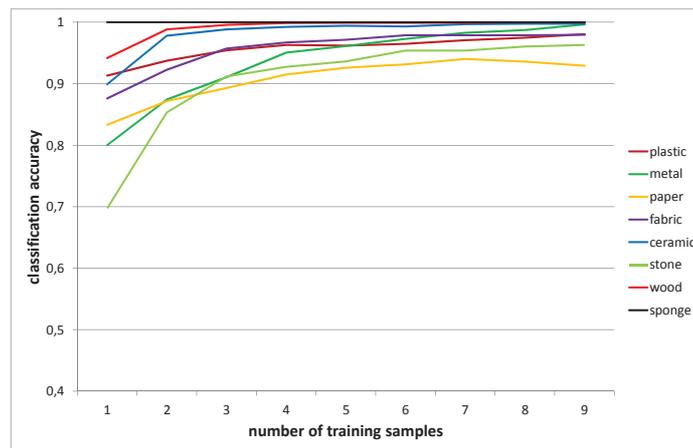


Figure 23: Classification accuracy for each material with concatenation of all movements, feature calculation F4, classification method linear SVM and an increasing number of training samples per object.

Furthermore, we determine to which material class each object is assigned by classification. Therefore we use all 40 data sets belonging to one of the material classes as training samples. Subsequently, we choose one object's data as training set. (If this object belongs to the 40 former mentioned objects, its data is removed from the training set.) This way, we perform a classification of each object. The results can be seen in Figure 24.

The resulting class sponge contains all objects we assigned to this class and few other objects. In contrast, the paper class has more false positives than objects actually made of paper. The ceramics class that achieved good results in the last section has now an accuracy of 58%. The objects that do not belong to any material class, are mostly assigned to the material wood.

4.3 Additional Movements

In order to compare more exploratory movements, we collected a smaller dataset of 5 objects, which can be seen in Figure 26 and added the two movements roll and rotate.



Figure 26: Objects examined with additional movements.

For the classification, we use as before a linear SVM with features $F4$. We perform the classification with a changing number of training samples. It can be seen in Figure 27 that the most effective exploration movement is the roll movement. It provides a classification accuracy of 98.8% for nine training samples. The {rotation movement also provides good results with an accuracy of 94.2%. With both rotation movements it is possible to classify over 90% of the samples correctly, when the classifier is trained with just one training sample.

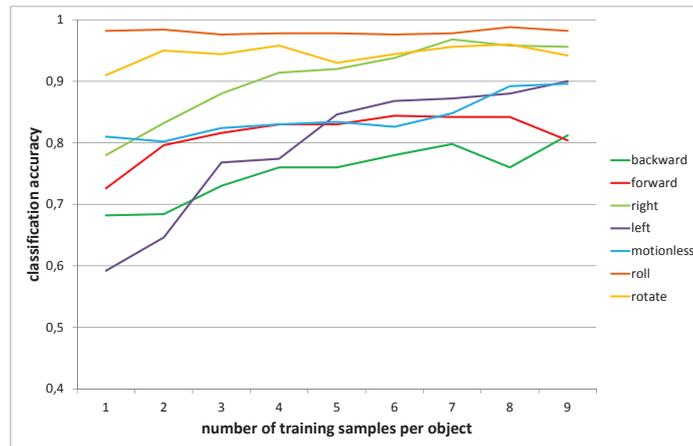


Figure 27: Classification accuracy for each object, with concatenation of all movements, feature calculation $F4$, classification method linear SVM and an increasing number of training samples per object.

5 Conclusion and Future Work

To summarize, we examined 49 objects with four different methods of classification and each of those classification methods was tested with seven different feature calculation methods. It is noticeable that the feature calculations especially designed for the BioTac sensor seem to perform more poorly than the method which uses the mean calculation directly based on the raw data for the chosen materials. It occurs that the mean method describes the features of different objects best to distinguish them during the classification process.

The feature ranking shows that the BioTac Sensor is very sensitive to noise. Furthermore, the rotation movements achieved better results than the other movements and should be repeated on a larger testset. Both results suggest that a refinement of the measurement process is necessary in the future.

To look further, on the one hand the feature calculation could be extended with additional measures to capture more information about the examined object. On the other hand the limitations of classification based on the provided feature information could be reached when separating different objects of similar material. Moreover, our material classification results suggest that the sensor modalities thermal conductivity, texture roughness and compliance are not sufficient to distinguish between materials.

For further evaluation of the methods' sufficiency they have to be applied on a specific task. This way, they can be adapted to special needs in order to increase the classification performance.

For better usability, the classification process should be provided in real time. In that way further information can be used to directly support manipulation tasks. For the same reason, autonomous selection of exploratory movements depending on manipulation tasks is eligible.

With these improvements another application is classification in cluttered scenes to differentiate between objects. With more exploratory movements it is also possible to create a spatial map of those objects. This allows us to assign meaningful labels to parts of the scene.

All these improvements have the potential to support a large variety of manipulation tasks and allow robots to interact with their environment more autonomously.

Bibliography

- [1] Marianna Madry, Liefeng Bo, Danica Kragic, and Dieter Fox. St-hmp: Unsupervised spatio-temporal feature learning for tactile data. *IEEE International Conference on Robotics and Automation (ICRA)*, 2014.
- [2] Joseph M Romano, Kaijen Hsiao, Günter Niemeyer, Sachin Chitta, and Katherine J Kuchenbecker. Human-inspired robotic grasp control with tactile sensing. *Robotics, IEEE Transactions on*, 27(6):1067–1079, 2011.
- [3] S. Chitta, J. Sturm, M. Piccoli, and W. Burgard. Tactile sensing for mobile manipulation. *Robotics, IEEE Transactions on*, 27(3):558–568, June 2011.
- [4] Leo Pape, Calogero M Oddo, Marco Controzzi, Christian Cipriani, Alexander Förster, Maria C Carrozza, and Jürgen Schmidhuber. Learning tactile skills through curious exploration. *Frontiers in neurorobotics*, 6, 2012.
- [5] Pierre Payeur, Codrin Pasca, A-M Cretu, and Emil M Petriu. Intelligent haptic sensor system for robotic manipulation. *Instrumentation and Measurement, IEEE Transactions on*, 54(4):1583–1592, 2005.
- [6] Patrick Dallaire, Philippe Giguère, Daniel Émond, and Brahim Chaib-Draa. Autonomous tactile perception: A combined improved sensing and bayesian nonparametric approach. *Robotics and Autonomous Systems*, 62(4):422–435, 2014.
- [7] Jivko Sinapov, Vladimir Sukhoy, Ritika Sahai, and Alexander Stoytchev. Vibrotactile recognition and categorization of surfaces by a humanoid robot. *Robotics, IEEE Transactions on*, 27(3):488–497, 2011.
- [8] Jeremy A Fishel and Gerald E Loeb. Bayesian exploration for intelligent identification of textures. *Frontiers in neurorobotics*, 6, 2012.
- [9] Susan J Lederman and Roberta L Klatzky. Hand movements: A window into haptic object recognition. *Cognitive psychology*, 19(3):342–368, 1987.
- [10] V. Chu, I McMahon, L. Riano, C.G. McDonald, Qin He, J. Martinez 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. *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3048–3055, May 2013.
- [11] Nathan F Lepora, Uriel Martinez-Hernandez, and Tony J Prescott. Active touch for robust perception under position uncertainty. *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3020–3025, 2013.
- [12] Danfei Xu, Gerald E Loeb, and Jeremy A Fishel. Tactile identification of objects using bayesian exploration. *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3056–3061, 2013.
- [13] Daisuke Tanaka, Takamitsu Matsubara, Kentaro Ichien, and Kenji Sugimoto. Object manifold learning with action features for active tactile object recognition. *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2014.
- [14] Biotac product manual.
- [15] Christopher M Bishop et al. *Pattern recognition and machine learning*, volume 1. springer New York, 2006.
- [16] Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [17] Christopher JC Burges. A tutorial on support vector machines for pattern recognition. *Data mining and knowledge discovery*, 2(2):121–167, 1998.
- [18] Chih-Wei Hsu, Chih-Chung Chang, Chih-Jen Lin, et al. *A practical guide to support vector classification*, 2003.
- [19] A Criminisi, J Shotton, and E Konukoglu. Decision forests for classification, regression, density estimation, manifold learning and semi-supervised learning. *Microsoft Research Cambridge, Tech. Rep. MSRTR-2011-114*, 5(6):12, 2011.
- [20] 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*, 12:2825–2830, 2011.

-
-
- [21] Chih-Chung Chang and Chih-Jen Lin. Libsvm: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27, 2011.
- [22] Ron Kohavi and George H John. Wrappers for feature subset selection. *Artificial intelligence*, 97(1):273–324, 1997.