

Material Recognition using Tactile Sensing

Emmett Kerr, TM McGinnity and Sonya Coleman^{a,b,a}

^a*Intelligent Systems Research Centre, University of Ulster, Magee Campus, Northern Ireland. e-mail: ep.kerr@ulster.ac.uk, sa.coleman@ulster.ac.uk*

^b*College of Science and Technology, Nottingham Trent University, UK. e-mail: martin.mcginnity@ntu.ac.uk*

Abstract

Identification of the material from which an object is made is of significant value for effective robotic grasping and manipulation. Characteristics of the material can be retrieved using different sensory modalities: vision based, tactile based or sound based. Compressibility, surface texture and thermal properties can each be retrieved from physical contact with an object using tactile sensors. This paper presents a method for collecting data using a biomimetic fingertip in contact with various materials and then using these data to classify the materials both individually and into groups of their type. Following acquisition of data, principal component analysis (PCA) is used to extract features. These features are used to train seven different classifiers and hybrid structures of these classifiers for comparison. For all materials, the artificial systems were evaluated against each other, compared with human performance and were all found to outperform human participants' average performance. These results highlighted the sensitive nature of the BioTAC sensors and pave the way for research that requires a sensitive and accurate approach such as vital signs monitoring using robotic systems.

Keywords: Tactile Sensing, Signal Processing, Neural Networks, Artificial Intelligence, Classification, Material classification

1. Introduction

Humans can quickly gain information about properties of an object or material simply by viewing them from different angles; for example they can estimate how it might feel to touch or how heavy it may be. This is due to our highly sophisticated visual capabilities and ability to adapt prior knowledge

learned from similarly shaped and known objects. However, there are some properties which are difficult to detect by vision alone, for example thermal conductivity, compressibility or a determination from what material an object is made. In order to learn these characteristics and remember them for future interactions, physical manipulation of an object is required. Distinguishing between objects and materials of different compressibility, temperature and texture can be achieved by performing complex manipulation tasks such as squeezing or rubbing. These complex manipulation tasks can inherently be performed extremely well by humans due to our very sophisticated tactile perception.

In order to produce the tactual perceptions required to learn about object properties, humans perform various types of movements when interacting with an object. Experimental psychologists have identified six general types of exploratory movements: pressure to determine hardness, static contact to determine thermal properties, lateral sliding movements to determine surface texture, enclosure to determine global shape and volume, lifting to determine weight, and contour following to determine exact shape (Lederman and Klatzky, 1987). Humans can complete these exploratory movements very fast and rapidly evaluate the object leading to a possible identification.

One of the main applications of an artificial tactile sensing system is expected to be in robotic applications. In order to retrieve the necessary tactual perceptions, a robot must explore an object in a similar manner to that of a human. However, in comparison to humans, these exploratory movements can be slow for a robot and typically involve the evaluation of a great volume of acquired data, thus being computationally expensive. It would be a useful skill for a robot if it could complete a preliminary evaluation of the object by quickly identifying the physical nature of the material (e.g. metal, wood, plastic etc.) through completing a small number of initial basic actions, thus removing the need for extensive manipulation. These initial basic actions could include the retrieval of thermal information upon initial contact with a material which may provide an indication of its physical nature, e.g. wood, metal or plastic. Furthermore the texture of a material can be identified by sliding a robotic hand or finger along the material, i.e. determining if it is rough or smooth.

We use an artificial fingertip to acquire data relating to the thermal properties and surface texture of materials, that are first analysed to initially identify to which group the material belongs (e.g. wood, metal, plastic etc.) and subsequently to determine the specific individual material (e.g. alu-

minium, copper, pine etc.). Building on the work presented in (Kerr et al., 2013, 2014a) and (Kerr et al., 2014b) the contribution of this work is the implementation of six further classifiers, hybrid configurations of the classifiers and a comparison of the performances of these classifiers. The classifiers implemented and compared are a one stage ANN (Kerr et al., 2014a), a two stage ANN (Kerr et al., 2014b), a one stage SVM, a two stage SVM, a hybrid ANN and SVM approach, GMM, LDA, NB, k -NN and MLP. Utilisation of these techniques to classify materials from tactile data resulted in an increase in classification accuracy and system efficiency in comparison to previous work (Kerr et al., 2014b) resulting in an approach capable of a fast initial identification of the material group and/ or individual material. These classifiers are also evaluated against human performance when identifying the same test materials as reported in (Kerr et al., 2014b). Therefore, in this paper we propose an approach for material classification using a combination of two of the three key characteristics, namely surface texture (represented through vibration data) and thermal characteristics. Although one weakness of the approach is that it does not achieve accuracies as high as that of (Xu et al., 2013), the strengths include the fact that it is computationally efficient and it demonstrates high classification accuracies across 14 materials, some of which are very similar, unlike many methods presented in the literature where the test materials can be very different. A range of common classifiers together with novel combinations of classifiers are evaluated and proven to be accurate and efficient algorithms for material classification. Furthermore, the system is tested with previously unlearned materials in order to assess its robustness.

The remainder of this paper is organised as follows: Section 2 presents an overview of related research in tactile sensing-based material classification. Section 3 describes the experimental set-up, including an overview of the equipment used and the learning algorithms used for the robot system. Section 4 presents an evaluation of the performance of the various artificial systems developed and compares their performance to that of human subjects. Section 5 discusses the artificial system and human participants' performance. Conclusions and plans for future work are given in Section 6.

2. Background and Related Research

Three of the key characteristics that are critical in performing an efficient grasp on an object are texture, compressibility and thermal properties.

Such characteristics can be obtained from tactile sensing. Materials from different classes may have similar compressibility; for example neither acrylic or medium density fibreboard (MDF) are very compressible but are from different classes of materials, i.e. plastic and wood. Drimus et al. (2011), developed a sensor and presented a technique used to classify between ten rigid and deformable household objects. The objects were squeezed in a gripper fitted with force sensors and a k -NN classification algorithm was utilised to classify the different objects, based on their compressibility. Classification accuracy of up to 92% was achieved for classification between the ten objects ranging from a rubber ball to a plastic bottle. The algorithm struggled to distinguish between similarly compressible objects even though the objects were very different, for example a “bad orange” and an empty bottle were found to have similar compressibility and therefore the algorithm failed to classify these objects correctly. Another approach which uses images derived from collected tactile sensor data is that of Pezzementi et al. (2011), which utilises computer vision algorithms to allow tactile sensor readings, representing compressibility, to be viewed as images. The tactile array sensor explores the object images representing contact made on the array are produced. Six feature descriptors including the well known Scale Invariant Feature Transform (SIFT) (Lowe, 1999) and MR-8 (Varma and Zisserman, 2005) are used to extract features from these images. The Bag-of-Words (BoW) approach (Jurie and Triggs, 2005), commonly used in computer vision methods, is then used to learn and classify the objects based on these features. Luo et al. (2015), also present an approach which utilises tactile images and a BoW approach to object recognition. An optimised BoW approach and dictionary were produced by considering different outputs of the classic BoW with different dictionary sizes for different views of objects. The classification performance of this optimised approach was compared against the performance of the classic BoW model utilising 12 objects. It was found that the proposed approach outperformed the classic BoW model particularly for classification of objects with similar features.

Exploiting other characteristics such as the object’s surface texture or its thermal properties may help to distinguish between very different materials which have similar compressibility, leading to improved classification accuracy. Related literature reports methods that are capable of classifying materials using surface texture and achieving high performance rates, i.e. is the object rough or smooth? Fishel and Loeb (2012) used the BioTAC sensors with a Bayesian exploration method for the intelligent identification of tex-

tures. Classification from a set of 117 textures was achieved. This approach related the textures to human style descriptions frequently used in psychophysical literature exploring texture discrimination (i.e. sticky/slippery, rough/smooth). The focus was to find the most useful exploratory motions on the surface of the material, relating to six general types of exploratory movements that humans make when tactually exploring objects, as outlined by experimental psychologists Lederman and Klatzky (1987). Ho et al. (2012) assessed the ability of their sensors to classify between three different materials (denim, a photo and tape) based on their surface texture. Classification accuracy of 91% was achieved between two very different materials, i.e. denim and a photo. However, the method struggled to distinguish accurately between two similar materials, namely the photo and tape. Jamali and Sammut (2010) propose a method which used an artificial silicon tactile sensor with Polyvinylidene Fluoride (PVDF) films to slide along materials and collect vibration data. A total of seven materials were tested including sponge, carpet, wood, two tiles of different roughness and two pieces of vinyl of different roughness. A Bayes classifier was trained with the Fourier coefficients of the sensor output data retrieved from the sliding motion over each material. It was shown that the authors' classifier performed well for classification between dissimilar surface textures, however, their approach was unable to distinguish between two types of tiles, as the texture of the surfaces were similar. Jamali and Sammut (2011) extended this work to evaluate multiple classifiers and added a majority voting algorithm. In a similar way to their previous work, data were retrieved from the fingertip making contact with test materials. The test materials included the same seven as in their previous work with the addition of a second type of carpet (different roughness), making eight material surfaces in total. Whilst the fingertip was being rubbed over each surface, different textures induced different intensities of vibration in the silicon fingertip. As in their previous work, the data from the fingertip were pre-processed and the Fourier coefficients of the sensor outputs were used as inputs to various classifiers that were implemented in the WEKA machine learning toolkit (Hall et al., 2009). It was found that the best performance of $95\% \pm 4\%$ was achieved using the Naive Bayes tree (NBTree). The system struggled most when trying to classify between materials of similar roughness. The approach proposed by Chathuranga et al. (2013) is another example of a system which produced promising results but still struggled to classify between similar material surfaces. The authors developed a biomimetic fingertip and used it to discriminate between

seven different materials based on surface texture. Using ANNs as classifiers, accuracies of up to 85% were achieved.

The thermal characteristics of a material is another modality that can be assessed and used to identify a material/object. Although one can generally tell the difference between materials that feel hot and cold, one would struggle to discriminate between similar materials on the basis of their actual temperatures. What human fingers detect thermally is not the absolute temperature of the material alone, but the effect of thermal conductivity and diffusivity. This is due to the fact that the human finger is at a consistent temperature, approximately 34° Celsius, and therefore will be different from the ambient temperature of most objects encountered. The nerves of the finger detect the flow of heat to/from the source. This suggests that human temperature sensing may be emulated by thermal conduction sensing (Monkman and Taylor, 1993). An approach which analyses the thermal properties of various materials is described in (Kerr et al., 2013). The materials are classified into initial groupings and further identified individually via the data gathered from test materials using the Syntouch® BioTAC™ sensor. PCA was applied to identify relevant features from the static temperature and thermal conductivity data that were used to train an ANN. The system was compared with humans and was found to outperform human performance. A success rate of 73% was obtained by the system when classifying the material group (i.e. plastic, wood, metal etc.) and 61% was achieved when identifying the individual material. When similar experiments were conducted using humans, results showed classification rates of 61% for group identification and 51% for individual material identification. However, even though the system achieves better performance than the humans, the classification rates are poor which is potentially caused by the fact that only thermal conductivity data are being used. Additionally, the system is computationally expensive due to the large number of input features necessary to accurately represent the data and the substantial number of neurons in the hidden layer. Bhattacharjee et al. (2015) present another approach that focusses on material classification. Their approach uses a tactile sensor that is capable of measuring solely thermal properties and they focus on classification within a short period of time, for instances where longer term contact is not viable. Three classification algorithms are considered, SVM+PCA, k -NN+PCA and Hidden Markov Models (HMM). It was found that SVM+PCA performed best with classification rates of 84% and up to 98% with longer contact times across 11 materials with varying initial conditions and 3-fold cross valida-

tion. Although the classification rates achieved are impressive, the materials tested varied widely in thermal conductivity and the purpose of the sensor is solely to determine thermal properties; it therefore cannot measure material properties such as texture or compressibility, thus limiting the applications for which it can be used.

Rather than classifying materials by analysing just one property, a combination of compressibility, texture and/ or thermal properties can be used to achieve higher rates of classification accuracy. Kerr et al. (2014a), used a combination of two of these properties, extending their work in (Kerr et al., 2013) by considering surface texture (vibration) as an additional modality alongside the thermal properties of a material. Additionally, the authors redesigned the system to improve efficiency by selecting a reduced number of principal components, reducing the number of neurons in the hidden layer and hence training the system faster than previously. The one stage ANN system in (Kerr et al., 2014a) obtained a classification rate of 83.81% for material group classification and 79.05% for individual material identification. Conversely, the humans achieved classification rates of 73.00% and 60.90% respectively. A two stage ANN approach to the same material classification experiment is presented in (Kerr et al., 2014b). The output from the network in the first stage is used as an input to the second stage, which is a set of networks, one for each material group. This second stage is used to identify each individual material, within a specified group. Despite not being as effective as the one stage ANN method presented in (Kerr et al., 2014a), this two stage approach obtained a classification rate of 83.81% for material group identification and 70.48% for individual material identification. Bhattacharjee et al. (2016) continued their earlier research and presented a method for distinguishing between contact with inanimate objects and humans, using a novel portable handheld device developed by the authors consisting of three tactile sensing modalities: a force sensor to detect contact, a heat-transfer sensor that is actively heated, and a small thermally-isolated temperature sensor (Wade et al., 2015). Data was collected from the arms of 10 human subjects from 3 different locations on the right arm and 80 objects consisting of 8 similar objects from 10 different bathrooms (Bhattacharjee et al., 2016). The effect of varying durations of contact made with the human subjects and the objects was also evaluated. A SVM was used to learn to classify between an object and human in the first instance and it achieved an average classification accuracy of 98.75% when contact was held for 3.65 seconds, 93.13% for 1.0 second of contact and 82.13% for 0.5 seconds of contact. High

classification accuracies were also achieved when generalising to new contact locations within the same bathroom with an average of 92.14% for 3.65 seconds of contact, 91.43% for 1.0 second and 84.29% for 0.5 seconds. However, when generalising to new environments, i.e. similar objects in different bathrooms, accuracies were lower achieving 84.00% for 3.65 seconds of contact, 71.00% for 1.0 second and 65.00% for 0.5 seconds. Like many other methods it struggled to identify similar objects especially when the environment changes. Hoelscher et al. (2015) use a BioTAC fingertip sensor to collect several features from a range of 49 objects to develop a system capable of identifying material and recognising objects. The objects are made of materials including plastic, metal, stone ceramic, paper etc. They collect data by making static and lateral contact on each material with the BioTAC sensor. The authors then compare seven different methods of extracting features from their processed data, including temporal data, PCA reduced dimensionality raw data, pressure features, electrode features, physically motivated features, temperature features and mean features. The authors compare two generative (NB and Gaussian) and two discriminative (SVM and Random Forests) classifiers to evaluate identification of the materials and objects from the collected data. The best performing method consisted of an SVM classifier with dimensionality-reduced mean values of filtered data, achieving a classification accuracy of material and object identification of over 97%. Xu et al. (2013) presented an algorithm which considered three key properties of the material to enable classification; compliance (compressibility), texture and thermal conductivity. To collect data for these three key properties a BioTAC fingertip was mounted on a Shadow Robot Hand. Based on a Bayesian exploration approach in the authors previous work (Fishel and Loeb, 2012) three exploratory movements were selected and used to explore ten test materials. A classification accuracy of 99% using 10 materials was achieved with the only failure being that a damp sponge was classified as a feather. The approach in (Xu et al., 2013) may have obtained 99% classification, however this was while assessing 10 substantially different materials, varying from a brick to a light feather. Therefore, the materials have very different properties allowing for more straight forward classification. Furthermore, all three tactile properties were used making the system both slow and computationally expensive.

It may be concluded that using just one material property does not allow for desirable classification of materials and combining modalities can lead to high computation cost. The ideal scenario would be to design a system that

can perform classification with accuracy similar to that of (Xu et al., 2013), using similar materials, with reduced computational burden and hence run in real-time (or close to). An overview of the best performing machine learning methods for the application of material identification using tactile sensing can be seen in Table 1. It demonstrates that the biggest difficulty in classification of materials is between similar materials or materials of similar roughness. Although, this is to be expected, it is still an ongoing challenge within the field.

3. Methodology

3.1. Experimental set-up

3.1.1. Tactile Sensor

The BioTAC fingertip is a tactile sensor that is shaped like a human fingertip and is filled with an incompressible liquid, giving it similar compliance to a human fingertip (Lin et al., 2009; Kerr et al., 2013). The sensor is capable of detecting a full range of information: micro vibrations, forces and temperature, similar to a human finger. Figure 1 shows a cross section view of the BioTAC fingertip.




Figure 1: Cross Section View of BioTAC Fingertip Tactile Sensor (Syntouch, 2013)

There are various sensors on the fingertip that enable retrieval of the aforementioned sensory information. An array of 19 electrodes measures the force applied on the fingertip by reading the impedance between each electrode and four common excitation electrodes. As the impedance over the sensing electrode increases, the measured voltage decreases, demonstrating

Table 1: Table comparing tactile sensing-based material identification methods

Research Group	Machine Learning Method	Sensor Type	Accuracy achieved (%)	Description
Xu et al. (2013)	Probability Matrices	Multi-modal tactile sensor (BioTAC)	99.00	10 very different materials were used. Continuous exploration and three modalities required therefore high computational costs.
Hoelscher et al. (2015)	SVM	Multi-modal tactile sensor (BioTAC)	97.00	Struggled to distinguish between similar materials. Required training from all materials in the group
Drimus et al. (2011)	KNN	Single Mode piezoresistive tactile sensor array	92.0	Classification across just 10 objects. Struggled to identify between similar objects.
Chathuranga et al. (2013)	ANN	Multi-modal tactile sensor	85.0	Classification across just 7 materials, 6 of which were fabrics.
Ho et al. (2012)	ANN	Single Mode electro-conductive soft skin sensors	81.0	Reported to be a very computationally expensive approach.
Jamali and Sammut (2011)	NBTree	Single mode silicon tactile sensor with PVDF films	95.0	Struggled to distinguish between materials of similar roughness.
Bhattacharjee et al. (2015)	SVM	Single mode thermistor-based tactile sensor	98.0	Longer contact with materials required for maximum accuracy and test materials varied widely
Bhattacharjee et al. (2016)	SVM	Multi-modal Hand-held tactile sensor	98.8	Classification accuracies achieved across already learned data, accuracy drops to 84.0% when tested against new data from the same materials.

contact in that area of the finger. Having a 19 electrode array allows the exact point of contact on the fingertip to be derived. Furthermore, it enables multiple points of contact to be derived at any instance. Inside the core of the fingertip there is a pressure sensor located in the rigid core. This sensor reads pressure changes in the conductive fluid located between the hard core and the elastomeric skin. The vibration sensor generated two values, a DC pressure signal (PDC) that is obtained using a low pass filter and an AC pressure signal (PAC) that is obtained using a band pass filter. The PAC data represent the vibration strength recorded when the BioTAC finger makes contact with materials or objects. Static temperature (TDC) and thermal flow (TAC) data can also be collected from using the BioTAC sensor as there is a thermistor at the tip of the finger. TAC data represent the rate that heat is transferring from the fingertip to an object. This is made possible by the thermistor which is located in the tip of the fingertip.

Using the BioTAC fingertip, two actions were performed on test materials to classify them. The first action was a press action where the fingertip was pressed into the material at a constant force. The second action was a slide action where the fingertip was dragged across the surface of the material at a constant force. Both actions aim to replicate two of the actions outlined by experimental psychologists Lederman and Klatzky (1987), that a human completes when first investigating an unknown material, namely static contact to determine thermal properties and lateral sliding movements to determine surface texture. The press action enables collection of the TAC and TDC data whereas the slide action enables the collection of PDC, TDC, PAC and TAC; these data are used for classification in Section 3.2.

3.1.2. Materials Classified

Consistent with previous work (Kerr et al., 2014a,b), we use 14 materials with the slide and press actions. We have intentionally selected some materials to be similar to determine if the approach can readily classify the material type and then identify the specific material. The 14 materials are Acrylic (Ac), Rough Acrylic (AcR), Copper (Cr), Aluminium (Al), Rough Copper (CrR), Rough Aluminium (AlR), Redbrick (R), Glossy Cardboard (GC), Plain Cardboard (PC), Soft Foam (S), Carpet (Ct), Doormat (D), MDF (M) and Pine (P). Figure 2 presents pictures of each of the materials used.

Analogous with (Kerr et al., 2014a,b), the materials used in the experiments were divided into six groups corresponding to the six different material




Figure 2: Samples of the 14 materials used in the experimental set-up.

Table 2: Table showing the individual materials and the groups to which they belong.


Material	Group
Acrylic (Ac)	Plastic (P)
Rough Acrylic (AcR)	Plastic (P)
Copper (Cr)	Metal (Me)
Aluminium (Al)	Metal (Me)
Rough Copper (CrR)	Metal (Me)
Rough Aluminium (AlR)	Metal (Me)
Redbrick (R)	Masonry (Ma)
Glossy Cardboard (GC)	Cardboard (C)
Plain Cardboard (PC)	Cardboard (C)
Soft Foam (S)	Fabrics (F)
Carpet (Ct)	Fabrics (F)
Doormat (D)	Fabrics (F)
MDF (M)	Wood (W)
Pine (P)	Wood (W)

types. The individual materials and the groups to which they belong are shown in Table 2. The artificial system was then evaluated on its ability to classify each material individually.

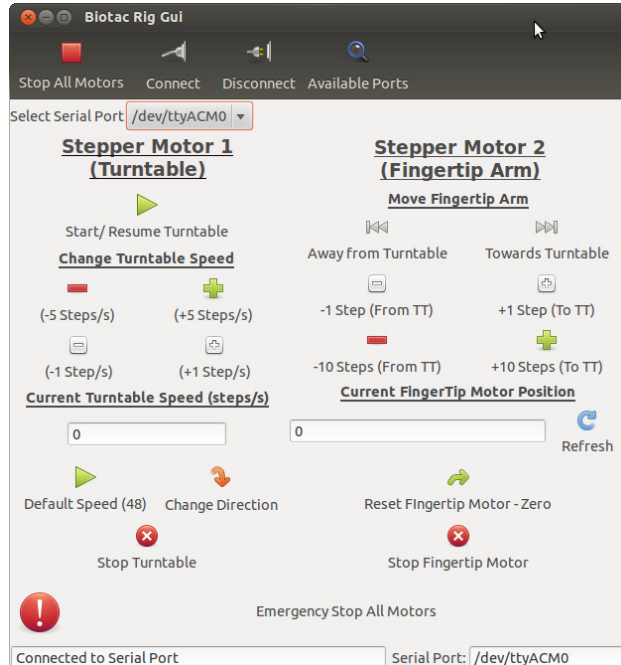
3.2. Data Collection

Analogous with work carried out in (Kerr et al., 2013; Xu et al., 2013), the BioTAC fingertip was powered on and left to rest for 15-20 minutes to allow it to reach its steady state temperature (approximately 31°C, 10°C above ambient). A rig including a motorised arm and turntable was designed and built for the experiment. Two stepper motors (one for the arm and one for the turntable) were connected to an Arduino UnoTM board and a Graphical User Interface (GUI) was developed using Python to control the motors. An image of the experimental rig and GUI can be seen in Figure 3(a) and (b) respectively. The fingertip was placed at the end of the motorised arm so that it could be moved towards or away from the material. Additionally, the material was placed on the motorised turntable to allow for the material to be moved below the finger to replicate a sliding action.

To collect data corresponding to a thermal exploratory movement, the fingertip is pressed onto a material using a constant force of 2N, measured



(a)



(b)

Figure 3: a) Image showing the experimental rig; b) Screen shot of the developed GUI.

by an ATI Nano17 6-axis F/T Sensor. To allow time for the heat flow to stabilise, it is necessary to collect 20 seconds of data for the press action. For the slide action, the finger tip was pressed on the material using a constant force of 1.59N (also measured with the aforementioned force sensor), and moved along the material surface for approximately 5 cm. The BioTAC data are recorded at 4400Hz and in a sequence, with the PAC values and the individual electrode values alternating (i.e. PAC data are recorded at 2200Hz). All other values such as TDC (static temperature), TAC (thermal flow rate) and PDC (static vibration) are given at the end of each sequence of values. An example of this sequence can be seen in Figure 4. All values are sampled with 12-bit resolution. These values are extracted and individual datasets for PAC, PDC, TAC and TDC data are obtained using MATLAB.

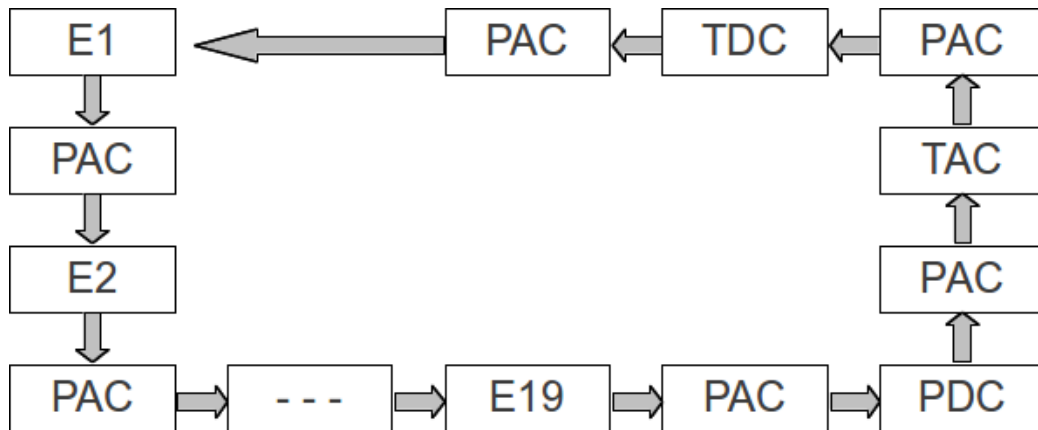


Figure 4: Diagram showing the Sequence of Data from the BioTac Fingertip

3.3. Pre-Processing

To obtain high quality thermal conductivity (TAC and TDC) time series data, we use only the data collected after the contact force has reached its maximum, this was found to occur approximately 8 seconds after contact. The TAC and TDC datasets, collected during the press action, each contain 800 data points. Additionally a set of 220 data points for each of TAC, TDC, PAC and PDC is obtained during the slide action. In line with other work presented in the literature, initial evaluation of using three modalities (i.e. thermal properties, vibration data and impedance values to represent compressibility) was conducted. It was hypothesised that due to the vast

quantity of electrode data collected from the BioTAC fingertip, the processing time would increase greatly when identifying materials. Initial visual analysis of the electrode values (voltage related to impedance) and initial experimentation was conducted to quantify the additional time it would take to include analysis of compressibility when classifying the materials in this experiment. As demonstrated in Figure 5, it is clear that all materials behave relatively similar during the press action with the exception of softer and rougher materials such as the softfoam and doormat. However, these differences are also clearly identified when analysing the vibration data (Kerr et al., 2014a) and therefore it was felt that any additional benefit from including compressibility data in this experiment was minimal to none. In fact, preliminary results indicated a decrease in accuracy when the electrode data were considered as part of the algorithm. Furthermore, the computation time to include the compressibility data was analysed. It took approximately 6 times longer to train one fold of a SVM network when the vast compressibility data was included in comparison to when only thermal and vibration data are considered. Therefore, considering the decrease in accuracy in preliminary investigations with compressibility data being considered, together with the increase in computation time, we use only thermal and vibration data in the experiments in this paper.

Initially we trained the classifiers with raw TAC, TDC, PAC and PDC data collected directly from the BioTAC sensor, however these resulted in poor classification accuracy such as 68% training accuracy for the ANN with 56% testing accuracy. Hence, to improve the characteristics of the data we applied Principal Component Analysis (PCA) to reduce the dimensionality of the inputs. We applied PCA using both Eigen-decomposition of the sample’s covariance and Singular Value Decomposition (SVD). With the sensor data, using the Eigen-decomposition based approach generated the best inputs for the classifiers, evidenced by the training and testing accuracies over multiple experiments. Hence, the presented approaches and results use PCA with Eigen-decomposition. In (Kerr et al., 2013), PCA was applied to a combined set of TAC and TDC data (1600 values), however in this paper we apply PCA to each individual modality (PAC, PDC, TAC, TDC) first. In this way, the principal components matrix for each material is formed from different modalities tailored for either the press or slide action experiment. As described in Section 3.4, the principal components matrices are used to train the numerous classifiers and hybrid combinations of classifiers evaluated.

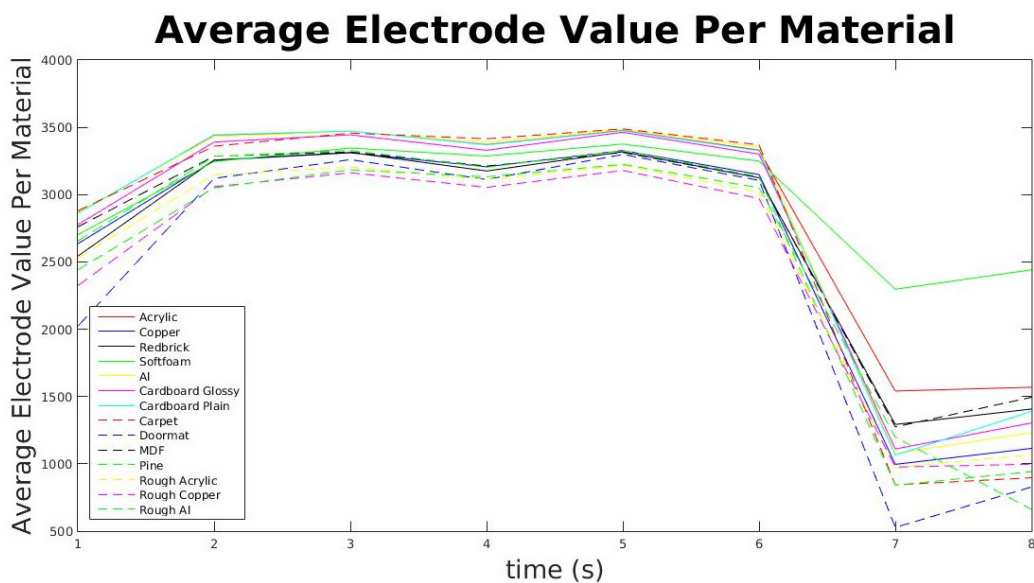


Figure 5: Graph showing Initial Analysis of Compressibility Data from the BioTac Fingertip

3.4. Classifiers

The data collected from the fingertip during the press and slide action experiments were normalised and PCA using Eigen-decomposition was applied to extract features. The optimal number of principal components (PCs) for each dataset was determined empirically as three. Hence in total there are 18 inputs to a classifier (six modalities and three PCs per modality, 12-bit values) for each experiment. To obtain the classification accuracy of each classifier in each experiment, the classified test data form an output matrix which is compared with a matrix containing the actual material labels (6 classes for classification of materials into the groups and 14 classes for classification of the individual materials). Average classification accuracy is obtained using 5 fold cross validation. The MATLAB package (MATLAB, 2013) and R package (R Core Team, 2013) were used to implement the range of classifiers. Each approach used for classification is described in more detail in Section 3.4.1 to Section 3.4.10.

3.4.1. One Stage Support Vector Machine (SVM)

Support Vector Machine is a useful and well established technique used in machine learning to learn and classify data. The main objective of a

SVM is to separate a set of inputs by identifying a line, or hyperplane in n -dimensional space. Each of the input points is known to belong to one of two classes. If it is possible for the algorithm to find a suitable hyperplane that will minimise future misclassification, then the data are easily classified based on what side of the hyperplane they reside. However, it is common that it is not possible to linearly separate the data. In this instance the SVM algorithm is required to map the input data to a higher dimensional space through the use of a kernel function (KF). Although there are many types of KFs, five are evaluated in this work and their performances compared. Within each KF there are various parameters that can be changed to optimise the performance of the classifier. One common parameter is the *soft margin* (sometimes known as the *box constraint*), C . The purpose of the soft margin is to modify the SVM algorithm to tolerate training errors by intuitively tolerating a few outliers on the wrong side of the hyperplane. Deciding on the size of the margin requires parameter tuning and optimisation. The KFs evaluated for the classification of materials and their respective parameters that were optimised are:

- Linear KF, Quadratic KF and Multilayer Perceptron (MLP) KF: Soft margin value (C) is optimised.
- Polynomial KF: Two parameters are optimised in this KF; the soft margin value (C) and the order of polynomial (p). Various combinations of values were evaluated in order to find the best performing combination of parameters.
- Radial Basis Function (RBF) KF: Two parameters are optimised in this KF; the soft margin value (C) and the gamma parameter for the Gaussian RBF (γ). The Gamma parameter is a positive number representing the scaling factor of the RBF. Various combinations of values were evaluated in order to find the best performing combination of parameters.

SVM is more commonly used as a binary classifier. However, binary classification is not suitable for the material classification task outlined in this work; the standard SVM algorithm has been adapted to enable multi-classification for material classification. Two common ways of utilising binary SVM classification techniques to allow for multi-class classification are by either having a sequence of numerous one-versus-one binary classifiers or by

having a multiple one-versus-all classifier. The approach utilised in this work is a series of one-versus-all classifiers in which the system is trained with each class classified against the samples of all the other classes, to develop a SVM multi-class classifier.

Two separate experiments were conducted: one for the classification of the materials into their six respective groups and one for the classification of the fourteen individual materials. All KFs were evaluated and their parameters optimised for both experiments. For each optimised KF implemented and evaluated, three different methods for calculating the separating hyperplane were also implemented and evaluated. The three standard methods evaluated were Sequential Minimal Optimization (SMO), Quadratic Programming (QP) and Least-Squares (LS) method. Table 3 outlines the best performance achieved from each KF for material group classification, the best performing method for calculating the separating hyperplane and the respective optimised parameters. It can be seen that the two best performing KFs with a classification rate of 86.19% are the quadratic KF and the polynomial KF. Although the two KFs performed equally well in terms of classification, the polynomial KF was found to be marginally more efficient than the quadratic KF at achieving the highest classification rate. The quadratic KF took 0.49 seconds to train for one fold of data whereas the polynomial KF took 0.47 seconds. Therefore, the Polynomial KF with a soft-margin value of 0.01 and a polynomial order of 2 was used in the SVM implemented in this approach.

Table 3: Table comparing the SVM KF classification accuracies.

Kernel Function	Best Performing Method	Optimised Parameters	Classification Accuracy
Linear	QP	$C = 1.00$	79.524%
Quadratic	SMO	$C = 0.01$	86.190%
MLP	LS	$C = 0.01$	57.143%
Polynomial	SMO	$C = 0.01$ $p = 2.00$	86.190%
RBF	QP	$C = 0.10$ $\gamma = 1.00$	66.667%

The SVM system was used for both experiments with the only difference being the number of groups/ classes (g) that the SVM had to classify. For the first experiment, i.e. classifying the materials into their six groups, $g = 6$ and for the second experiment involving 14 individual materials, $g = 14$ as described in Section 3.1.2.

3.4.2. Two Stage SVM

In contrast to having two separate experiments for the classification of the materials into their groups and the classification of individual materials, as described in Section 3.4.1, this section explains a two stage approach to material classification. This approach utilises the multi-class SVM classification algorithm explained in Section 3.4.1 to firstly classify the materials into their groups and then uses this information to further classify the materials individually using a mixture of binary and multi-class SVM classification in the second stage.

The same SVM design as outlined in Section 3.4.1 is used in the first stage of this two stage approach. For the second stage, 6 SVM systems are trained, one for each material group, some of which have different numbers of outputs/ groups for the individual materials within them, requiring the need for binary classification in some cases and multi-class classification in others. These SVM systems are saved and later used for testing the data input at the second stage. Similar to the development of the SVM in the one stage approach, five different KFs with a range of parameters and hyperplane calculation methods were evaluated for the second stage SVM classification system. The three equally best performing KFs in the second stage were the Quadratic KF, the Polynomial KF and the RBF KF, with the most efficient proving to be the Polynomial KF. Therefore, the Polynomial KF was used in the second stage SVM.

Upon completion of material group classification in the first stage, the results of the classification are used as the inputs to the second stage. A model of the two stage SVM approach can be seen in Figure 6.

If a material was correctly classified into a specific group in the first stage then this activates the second SVM to classify the individual material within this material group. However, if the material group is incorrectly classified in the first stage then it is recorded as an incorrect classification in the final output matrix and is not classified in the second stage to improve efficiency. Likewise, if the material is classified incorrectly in the second stage in comparison with ground truth (i.e. it has been correctly classified

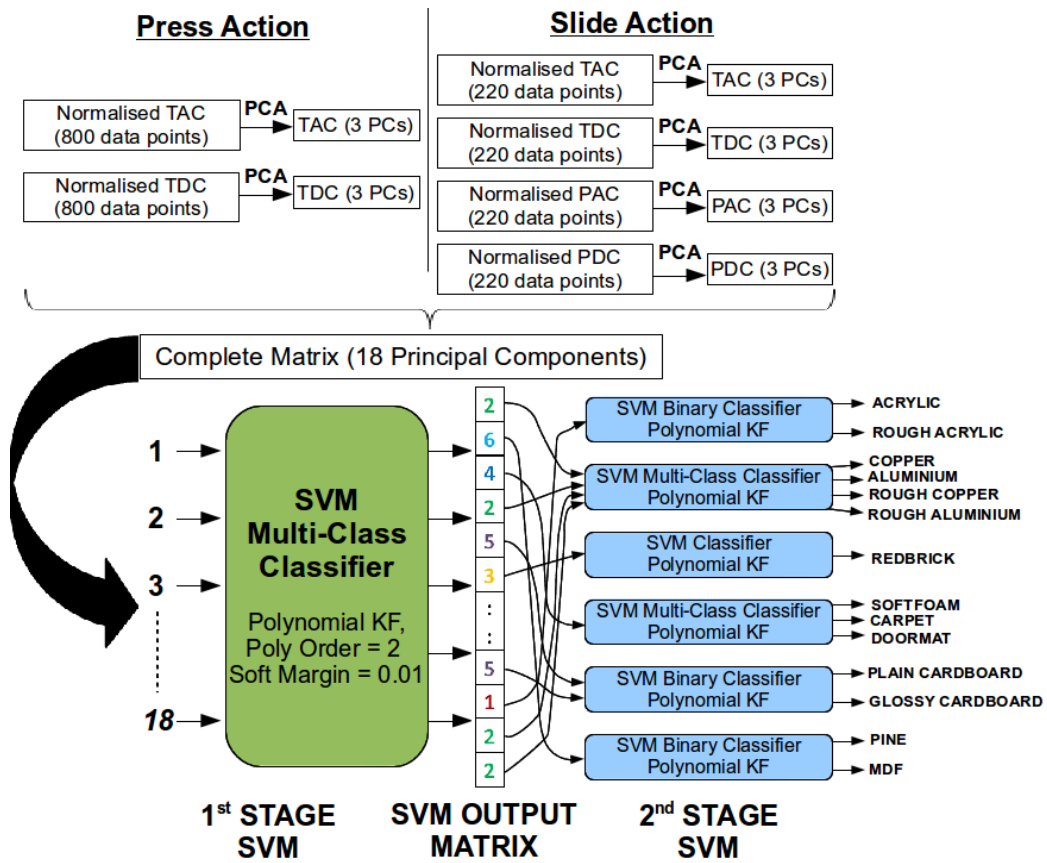


Figure 6: Diagram showing the 2-Stage SVM approach used for material classification

in its group but incorrectly classified individually), then it is recorded as an incorrect classification in the output matrix.

3.4.3. One Stage ANN

We use a back propagation ANN as illustrated in Figure 7. During training, the number of neurons in the hidden layer and the number of training epochs were varied and evaluated. This determined that the ideal structure for the ANN contained 75 neurons and required 1500 epochs for training.

The network structure illustrated in Figure 7 was applied in both experiments differing numbers of outputs (n). In the first experiment, $n = 6$ as there are six material groups and in the second experiment $n = 14$ corresponding to fourteen individual test materials, as described in Section 3.1.2.

3.4.4. Two Stage ANN

The two stage approach (Kerr et al., 2014b), used a series of back propagation ANNs, each with one hidden layer, where the output of the first ANN is the input of the second to allow for group classification first followed by individual classification. This differs from the one stage ANN approach outlined Section 3.4.3 as the two networks work in series rather than two individual networks, one classifying materials into groups and one classifying individual materials. The two stage approach structure can be seen in Figure 8.

The ANNs used in the two stage approach are similar in structure to the ANN used in the one stage approach explained in Section 3.4.3. The first ANN is used to classify the materials into groups and the second stage of the ANN is comprised of six individual ANNs for each material group (i.e. plastic, metal, masonry, fabrics, paper and wood). These ANNs were all trained individually with their respective outputs relating to how many materials there are in each group, for example plastic has two outputs for the two plastic materials whereas metal had four outputs for the four metal materials in that group. After training these networks were saved and then used for classifying the resulting output from the first ANN. In order to avoid false positives being identified, a threshold was empirically identified and set for testing the materials within the group. If the output of the neuron fired for the material sample being tested is not greater than 0.3, then it was considered to be a failed classification.

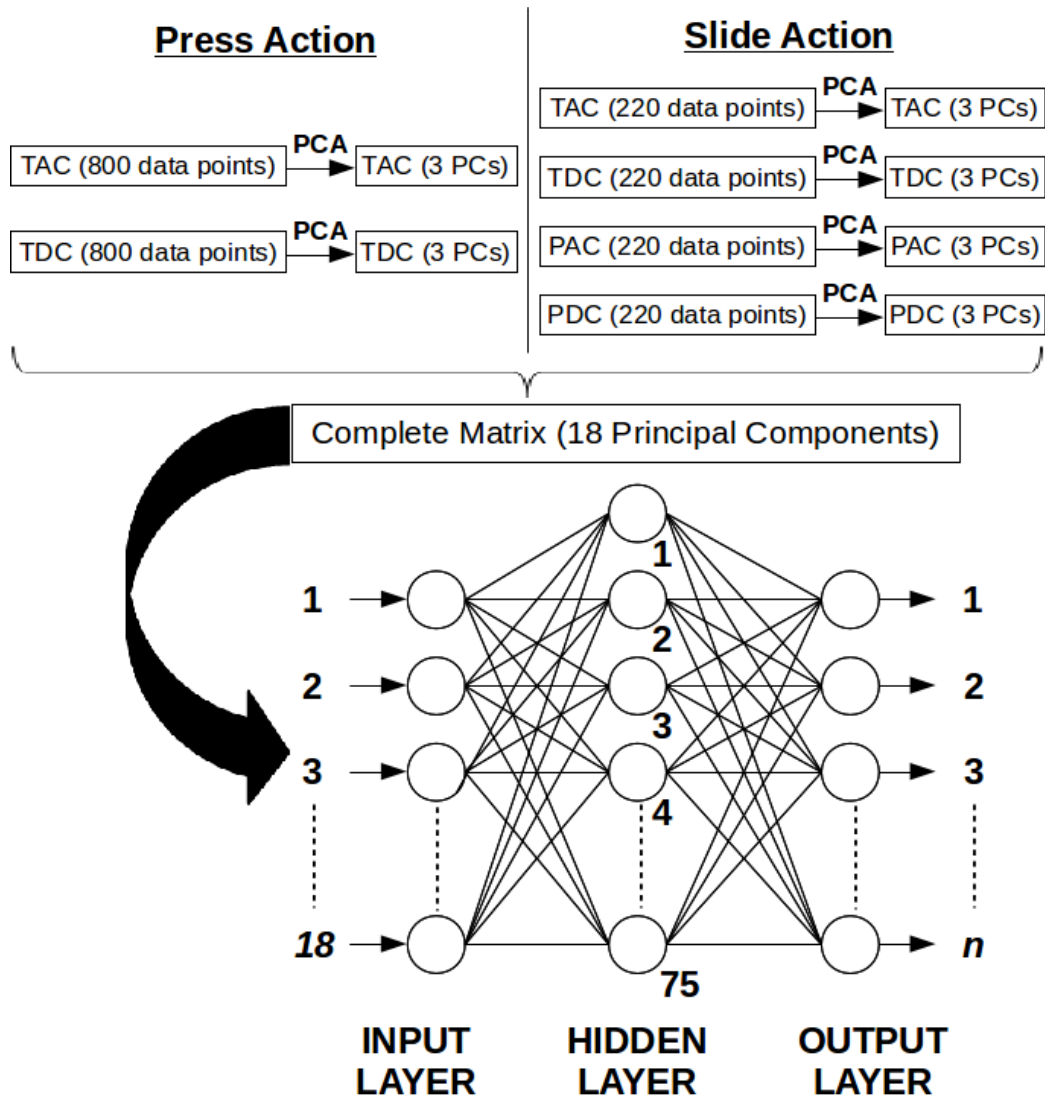


Figure 7: Diagram showing the ANN used for material classification

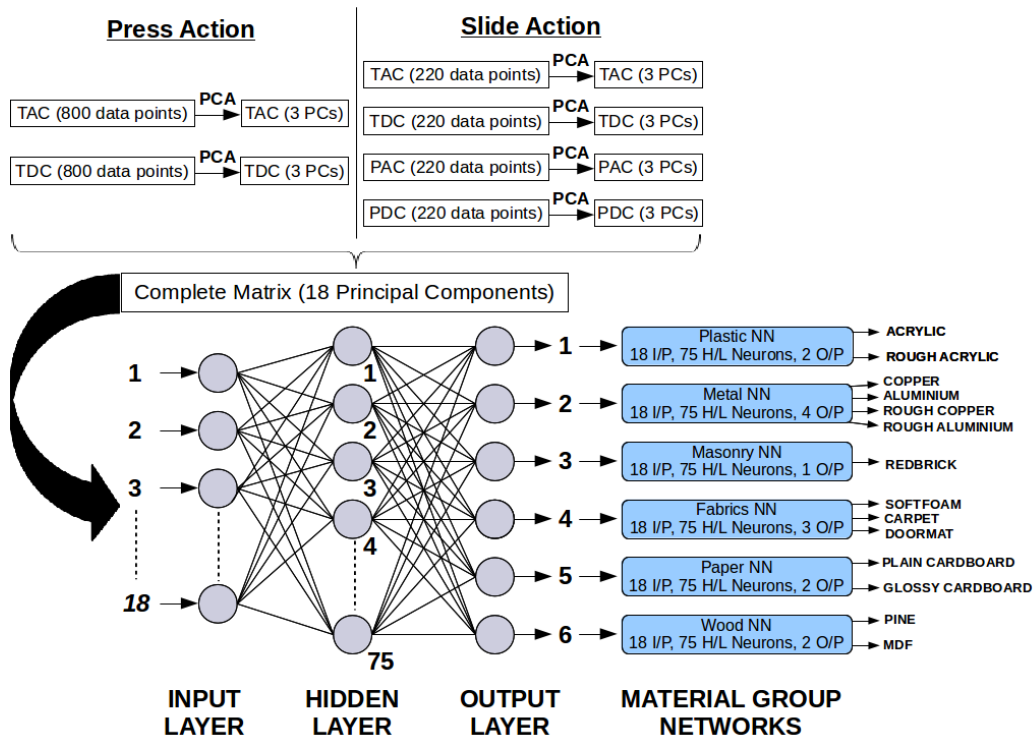


Figure 8: Diagram showing the two stage ANN used for material classification

3.4.5. SVM and ANN Hybrid System

A hybrid approach of the SVM classifier for material group classification and separate ANN classifiers for the individual material classification within each group has been developed. The structure of the hybrid system can be seen in Figure 9.

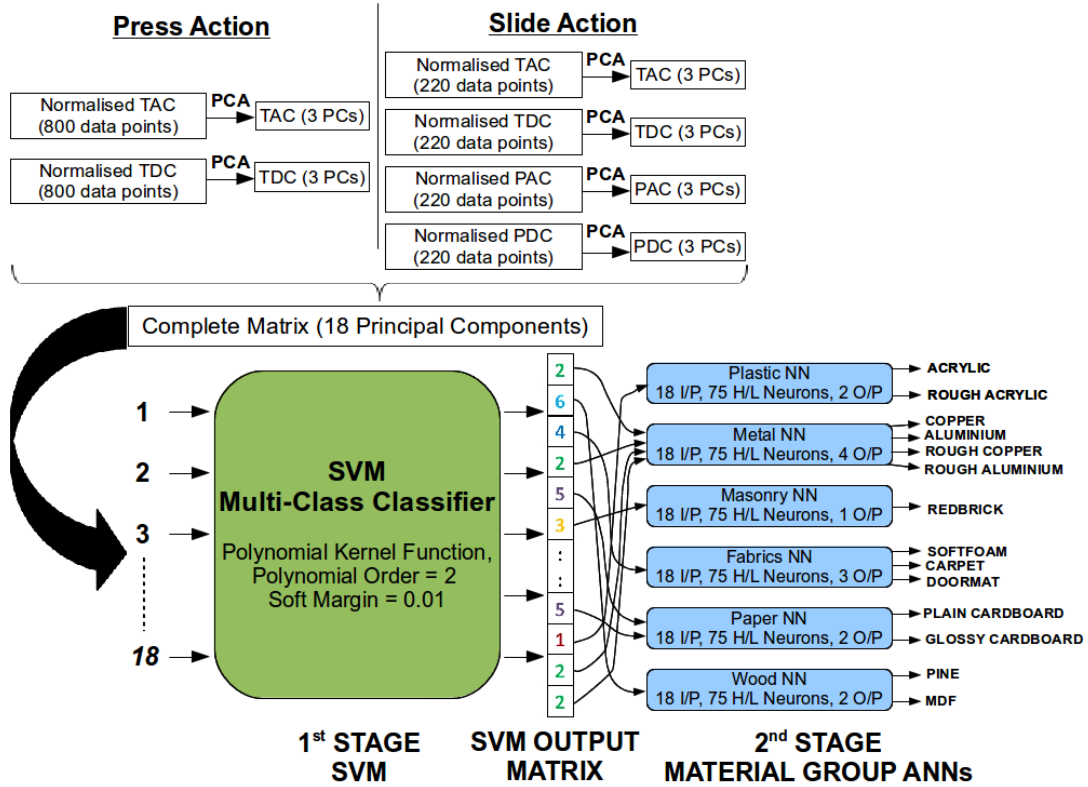


Figure 9: Diagram showing the SVM+ANN hybrid approach used for material classification

A SVM classifier, as described in Section 3.4.2, is used first to classify the materials into their material group (i.e. using the polynomial KF with a polynomial order value of 2 and a soft margin value of 0.01). Upon classification into their groups an output matrix is formed and the values in this matrix initiates the second stage of the hybrid approach with six ANNs for individual material classification within each of the material groups. The ANNs are constructed as described in Section 3.4.4 (i.e. with one hidden layer of 75 neurons and trained for 1500 epochs). As with the other two stage ap-

proaches outlined in this section, in order to improve efficiency only materials correctly classified into their group are classified in the second stage.

3.4.6. Gaussian Mixture Models (GMMs)

GMM based on the Maximum Likelihood (ML) estimation using Expectation Maximisation (EM) was also used for classification. The Gaussian probability density function (PDF) is a bell shaped curve defined by two parameters, mean and variance. The Gaussian distribution is commonly used for approximating a class model shape in a selected feature space. It is assumed that the class model is truly a model of one basic class, however if the actual PDF is multi-modal it fails. GMM is a mixture of several Gaussian distributions and is therefore able to represent different subclasses within a class. The resulting probability density function is defined as a weighted sum of the Gaussian distributions that make up the GMM. The PDF of the GMM can be used to calculate the likelihood of any new input data within each class and identify the class for which the PDF generates the maximum likelihood and therefore classify the data. The EM algorithm is an iterative method for calculating the maximum likelihood distribution in GMMs (Dempster et al., 1977).

GMMs are calculated to represent each material (cluster centres are initialised randomly and the k -means algorithm is used for calculating convergence). The EM algorithm is then used to calculate the maximum likelihood and the associated class label is selected. These classified labels are used to form an output matrix. As the cluster centres are initialised randomly 10 runs of each classification (i.e. material groups and individual materials) were completed and an average classification accuracy calculated.

3.4.7. Linear Discriminant Analysis (LDA)

LDA finds a linear combination of features that characterises or separates two or more classes of objects or events by creating the linear combination which yields the largest mean differences between the classes (Martinez and Kak, 2001). It is a generalisation of Fisher's linear discriminant and it is a widely used method in statistics, pattern recognition and machine learning. LDA can be used for binary or multi-class classification. In this work it is used for multi-class classification to identify materials individually and into their groups.

3.4.8. Naive Bayes

Naive Bayes is a family of algorithms used for training classifiers capable of binary or multi-class classification. Rather than the classifier taking all features of a class into account collectively to describe the class, all naive Bayes classifiers assume that the value of a particular feature is completely independent of any other feature in the class meaning they can be trained very efficiently during a supervised learning task (Rish, 2001). In this work it is used firstly to classify materials individually and secondly classify them into their groups.

3.4.9. k -Nearest Neighbour (k -NN)

k -NN is a non-parametric method which can be used for classification or regression. In this work it is used for multi-class classification of materials. When used for classification, an object is assigned to the class which is most common amongst its k nearest neighbours following a majority vote. The number of neighbours is defined by the user and is normally a small positive integer, for example if $k=1$ then the object is assigned to the class of the single nearest neighbour or if $k=3$ then the most popular class amongst the three closest neighbours is the class that the object will be assigned to. A range of k values was tested to identify the most accurate number of neighbours to take into account for maximum classification accuracy. The k value that proved to be most effective for classifying the materials individually and into their groups was $k=5$.

3.4.10. Multi-Layer Perceptron (MLP)

Multi-layer Perceptron (MLP) is a feed forward ANN which uses back-propagation to train the network. A MLP consists of multiple layers of adaptive weights that are trained to map sets of input data onto a set of appropriate outputs (Bishop, 1995). A range of numbers of layers and neurons on each layer was evaluated in order to find the optimal structure for the MLP. The structure of the MLP used in this work to classify the materials into their groups and individually was a network consisting of 10 layers containing 19 neurons each.

4. Results

In order to evaluate the artificial system's performance, analysis comparing the algorithms was conducted. In addition, the artificial system was

evaluated against human performance. In this section, the results of the human evaluation experiments are presented. The initial analysis of the extracted data collected from the fingertip during the experiments is also presented. Furthermore, the performances of the aforementioned classifiers are presented, evaluated against each other and against human performance.

4.1. Human Evaluation

To evaluate each artificial system, they are compared against human performance in identifying the the same set of 14 materials. The participants were evaluated on identifying the materials both individually and identifying which group each material belonged to. Initially they were evaluated by only being allowed to press into the materials with their finger and assess its thermal properties and secondly they were evaluated using both a press and slide action (to evaluate the surface texture of the material). The findings from these experiments can be found in (Kerr et al., 2013) and (Kerr et al., 2014a) respectively.

4.2. System Evaluation

A comparison of the results obtained for all of the classification systems presented and the average classification accuracies of the human participants are shown in Table 4. As specified in Section 3.4, the results for the artificial system are calculated by computing the average of the classification rates using 5-fold cross validation.

As shown in Table 4, SVM has outperformed the other approaches for both the classification of the materials into their groups and for classification of the individual materials. The two stage SVM has proven to be capable of maintaining the group classification accuracy for the classification of the individual materials in its second stage. As the SVM only has to use binary classification in three out of the six groups and the most materials to classify between in the other groups is four, then it was possible to achieve such classification accuracy in the second stage using SVM.

The second best performing algorithm is the SVM and ANN hybrid approach. The SVM and ANN hybrid approach using the quadratic KF achieved 81.89% classification. Using the polynomial KF within the hybrid approach proved to be marginally more accurate, achieving 82.86% accuracy. The Gaussian Mixture Models (GMM) approach was also a highly performing classifier with 86.14% and over 80% for material groups. However, it can be seen in Table 4 that aside from the fact that they perform

Table 4: Table comparing the experimental results.

Modalities Utilised	Run Time (to train one fold)	Material Group	Individual Materials
Human Participants - Thermal Only (Kerr et al., 2013)	N/A	61.00%	50.90%
Human Participants - Thermal and Vibration	N/A	79.76%	69.64%
One Stage ANN - Thermal Only (Kerr et al., 2013)	4.94secs	73.00%	60.90%
One Stage ANN - Thermal and Vibration	5.47secs	83.81%	79.05%
Two Stage ANN - Thermal and Vibration	7.66secs	83.81%	70.48%
One Stage SVM - Thermal and Vibration	0.48secs	86.19%	72.38%
Two Stage SVM - Thermal and Vibration	0.55secs	86.19%	86.19%
Hybrid SVM and ANN - Thermal and Vibration	2.66secs	86.19%	82.86%
GMM - Thermal and Vibration	13.35secs	86.14%	80.24%
LDA - Thermal and Vibration	1.52secs	78.57%	80.95%
NB - Thermal and Vibration	1.68secs	71.43%	80.95%
k -NN - Thermal and Vibration	0.95secs 29	71.43%	64.29%
MLP - Thermal and Vibration	5.45secs	71.43%	66.67%

with slightly less accuracy in comparison with the two stage SVM algorithm, like the other systems, the GMM system also takes significantly longer to train (taking 13.35secs to train for one fold of data) than the SVM system (taking 0.48secs to train for one fold of data). Interestingly the LDA and NB classifiers both performed better when classifying the materials individually rather than into their groups. This is contrary to what was expected as there are 14 outputs when classifying into individual materials compared with only six outputs for group classification. As a consequence, LDA and NB both performed marginally better than GMM for classifying the individual materials with 80.95% accuracy however performed poorer than GMM for the classification of material groups with 78.57% accuracy for LDA and 71.43% accuracy for NB. k -NN was found to perform optimally when $k = 5$ with a fast training time of 0.95sec for one node, however in comparison with the other classifiers it is still lacking in accuracy for classifying materials individually and into their groups. MLP was found to be relatively slow to train and was out performed by the majority of the other classifiers. The training times presented are based on training one fold of data for each classifier and the training was completed in Matlab on a desktop PC with a Intel® Xeon® CPU E5-16070 @ 3.00GHz 4 processor and 16Gb of RAM. It is demonstrated in Table 4 that the novel two stage SVM and hybrid SVM and ANN approaches have outperformed all standard approaches such as GMM, LDA or singular SVM and ANN approaches. These architectures have proven to be more accurate than “off-the-shelf” architectures with the novel two stage SVM system being the best performing and most efficient system for the task of material classification when compared with others.

In order to assess robustness, completely unseen objects and materials were presented to the system. Seven previously completely unseen objects in total were presented namely, a plastic bottle, a ceramic cup, a glass cup, a metal cup, a metal fork, a plastic fork and a paper plate. Three datasets for each object were tested. There is no “any other material” option in the system so every material will be classified as one of the 6 materials outlined in Table 2. Interestingly, materials that the system had not previously been trained on, i.e. ceramic and glass, were identified as a material with similar properties. The glass was identified as plastic in the majority of cases and the ceramic was classed as metal in all cases. In some instances the system struggled to identify the materials on which the system had been trained. The paper plate was only correctly identified as paper in 1 out of the 3 instances. However, this may be due to the thin plastic film coating on

paper plates that makes them water retardant, as shown in Figure 10(a). This made it difficult to identify the plate as paper and indeed it was also incorrectly identified as plastic in both other instances. The plastic fork was also correctly identified as plastic in 1 of the three instances. This may be due to the fact that this specific metal fork was manufactured to be “metal like”, i.e. look like a metal fork as shown in Figure 10(b). This means that there is a fine film on the surface of the fork making it difficult to identify as plastic, similar to the difficulty faced with identifying the paper plate. Indeed, in both other cases the fork was identified as being metal.



(a)



(b)

Figure 10: a) Image showing two of the previously unlearned objects; a) the paper plate; b) the “metal-like” plastic fork

However, the metal cup was identified as the correct material in at least 2 of the 3 instances and both the plastic bottle and metal fork were correctly identified as their respective material in all 3 instances. This illustrates the

robustness of the system and with further training for more materials the system is capable of correctly identifying previously learned materials.

5. Discussion

From the results of the human evaluation carried out in (Kerr et al., 2014a), it was found that the average accuracy was 79.76% when identifying the material groups and 69.64% when identifying individual materials. A detailed analysis of these findings can be seen in the confusion matrices in Figure 11(b) and 12(b). It can be seen from Table 4 that all of the artificial systems presented have outperformed the human participants (when performance is averaged across all test materials), in some cases by over 6% for material group classification and by over 16% for individual material classification. This is due to the fact that is easier for a human to determine what group a material would belong to (i.e. identify that the material is definitely metal) than it is for them to identify a specific material within that group (i.e. is it copper or aluminium). It can be seen from the confusion matrix in Figure 11(a) and 11(b) that both the artificial system and the human participants successfully classified the soft foam in every instance and did not get it confused with any other material. It can also be seen from Figure 11(a) that where the artificial system classified the pine material correctly on 60% of the instances, the highest misclassification was with MDF, another material with similar properties and in the same material group. Likewise, for classifying MDF, the largest misclassification was with pine. This misclassification can also be witnessed as a common misclassification amongst the human participants, with 50% of the participants correctly classifying MDF but in almost 42% of the instances the human participants misclassified MDF as pine as shown in Figure 11(b). It is shown in Figure 11(a) that a common misclassification of the artificial system was a confusion between acrylic and glossy cardboard, with glossy cardboard being classified correctly for 40% of the instances however in just over 33% of the instances it was misclassified as another material with similar texture and thermal conductivity properties: acrylic.

Interestingly, the identification of the rough materials was evidently challenging for the human participants, with only 2 out of 12 being able to identify all three rough materials. For example, one of the most challenging materials for humans to identify was the rough copper, and it was commonly misclassified as redbrick, as shown in Figure 11(b). This is due to the fact

	Ac	Cr	R	S	P	Ct	M	GC	PC	D	AI	CrR	AIR	AcR
Ac	60.0	0.0	0.0	0.0	0.0	0.0	0.0	20.0	6.7	0.0	0.0	0.0	13.3	0.0
Cr	0.0	100.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
R	0.0	0.0	66.7	0.0	20.0	0.0	6.7	6.7	0.0	0.0	0.0	0.0	0.0	0.0
S	0.0	0.0	0.0	100.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
P	0.0	0.0	13.3	0.0	60.0	0.0	20.0	6.7	0.0	0.0	0.0	0.0	0.0	0.0
Ct	0.0	0.0	0.0	0.0	0.0	80.0	0.0	6.7	6.7	0.0	0.0	0.0	6.7	0.0
M	6.7	0.0	0.0	0.0	20.0	0.0	73.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0
GC	33.3	0.0	0.0	0.0	13.3	0.0	0.0	40.0	0.0	0.0	0.0	0.0	13.3	0.0
PC	6.7	0.0	0.0	0.0	0.0	6.7	0.0	6.7	73.3	0.0	0.0	0.0	6.7	0.0
D	6.7	0.0	20.0	0.0	6.7	0.0	0.0	0.0	6.7	53.3	0.0	0.0	6.7	0.0
AI	0.0	13.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	80.0	0.0	6.7	0.0
CrR	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	93.3	6.7	0.0
AIR	13.3	0.0	0.0	0.0	0.0	0.0	0.0	6.7	0.0	6.7	6.7	0.0	66.7	0.0
AcR	6.7	0.0	0.0	0.0	0.0	0.0	0.0	13.3	0.0	0.0	0.0	0.0	13.3	66.7

(a)

	Ac	Cr	R	S	P	Ct	M	GC	PC	D	AI	CrR	AIR	AcR
Ac	100.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Cr	0.0	58.3	0.0	0.0	0.0	0.0	0.0	16.7	0.0	0.0	25.0	0.0	0.0	0.0
R	0.0	0.0	58.3	0.0	8.3	0.0	0.0	0.0	8.3	0.0	0.0	8.3	0.0	16.7
S	0.0	0.0	0.0	100.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
P	8.3	0.0	0.0	0.0	41.7	0.0	25.0	8.3	0.0	0.0	16.7	0.0	0.0	0.0
Ct	0.0	0.0	0.0	0.0	0.0	100.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
M	0.0	8.3	0.0	0.0	41.7	0.0	50.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
GC	16.7	8.3	0.0	0.0	0.0	0.0	8.3	66.7	0.0	0.0	0.0	0.0	0.0	0.0
PC	0.0	0.0	0.0	0.0	0.0	0.0	8.3	0.0	83.3	0.0	0.0	0.0	8.3	0.0
D	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	100.0	0.0	0.0	0.0	0.0
AI	8.3	8.3	8.3	0.0	16.7	0.0	0.0	0.0	0.0	0.0	58.3	0.0	0.0	0.0
CrR	0.0	0.0	25.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	41.7	16.7	16.7
AIR	0.0	0.0	8.3	0.0	0.0	0.0	8.3	0.0	0.0	0.0	0.0	25.0	41.7	16.7
AcR	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	25.0	0.0	75.0

(b)

Figure 11: a) Confusion matrix showing the percentage accuracy of the artificial system (using 1 stage SVM) for classifying all the materials individually; b) Confusion matrix showing the percentage accuracy of the human participants for classifying all the materials individually, collated from the findings in (Kerr et al., 2014a).

that the redbrick is perceived as cold and rough by the human participants, in a similar manner to how metals are perceived. Therefore, this is one of the instances where the artificial system outperforms the human participants, as the system can analyse thermal conductivity much more accurately than the human participants and therefore not misclassify it as a masonry material. This is shown in Figure 11(a) where the artificial system classifies the rough copper in just over 93% of the instances with the only misclassification being with another rough metal: aluminium. This common misclassification is also evident when classifying the materials into their groups, as shown in Figure 12(a) and 12(b). The artificial system outperforms the human participants when classifying materials in both metal and masonry groups. The majority of instances of misclassification for the human participants when classifying masonry objects is confusing it with a metal object, as both are perceived as cold.

From the results of the human evaluation reported in (Kerr et al., 2014a), it was evident that the introduction of texture analysis by sliding their finger along the material has increased the accuracy of the human participants for classifying materials, particularly when identifying the material groups. It should be noted that due to the nature of some materials, humans can intuitively take other characteristics into account, for example compressibility (i.e. soft foam is much more compressible than the other materials). The artificial system was trained to analyse only the thermal properties and surface texture of the material and not compressibility. This may have given the human participants a slight advantage in identifying some materials. This is evident when comparing the accuracy of the artificial system for classifying fabrics as a material group (80%) with the accuracy of the human participants (100%), as shown in Figure 12(a) and 12(b). The human participants have no misclassification of the fabrics when identifying their material group. It can be assumed that this is due to the fact that the human participants could identify the compressibility of the objects like soft foam and carpet whereas the artificial system was not trained to analysis this modality. However, even though the artificial system did not consider a key modality for humans, namely compressibility, overall it still outperformed humans.

6. Conclusion and Future Work

A range of classifiers and combinations of classifiers that were implemented, tuned and evaluated for the classification of materials into their

	P	Me	Ma	C	F	W
P	90.0	0.0	0.0	3.3	3.3	3.3
Me	0.0	100.0	0.0	0.0	0.0	0.0
Ma	0.0	0.0	73.3	6.7	0.0	20.0
C	13.3	0.0	0.0	76.7	6.7	3.3
F	0.0	0.0	4.4	4.4	80.0	11.1
W	0.0	0.0	10.0	10.0	0.0	80.0

(a)

	P	Me	Ma	C	F	W
P	87.5	12.5	0.0	0.0	0.0	0.0
Me	10.4	68.8	10.4	4.2	0.0	6.3
Ma	8.3	16.7	58.3	8.3	0.0	8.3
C	8.3	8.3	0.0	75.0	0.0	8.3
F	0.0	0.0	0.0	0.0	100.0	0.0
W	4.2	12.5	0.0	4.2	0.0	79.2

(b)

Figure 12: a) Confusion matrix showing the percentage accuracy of the artificial system (using 2 stage SVM) for classifying all the materials into groups; b) Confusion matrix showing the percentage accuracy of the human participants for classifying all the materials into groups, collated from the findings in (Kerr et al., 2014a).

respective groups and the classification of individual materials has been presented. This classification was based on a combination of the thermal and surface texture properties of the materials. It is proven that a two stage SVM approach performs best and is the most efficient method in completing the material classification task. It was found that the artificial system's average performance across all test materials exceeded that of human participants' average performance, with the two stage SVM approach outperforming humans by over 16%. However, as the system does require training it is still slower at classification than the human participants due to humans' inherent sophisticated tactile sensing.

There are two immediate future directions of this work. Firstly, the presented tactile based methods for material classification could be used in collaboration with vision based methods. If an estimate of the physical size of an object is determined using image processing, then the volume of the object can also be calculated. The tactile sensing methods presented in this paper can identify the material from which the object is made and armed with the volume of the object, an estimate of the object's weight can be calculated. The incorporation of vision would further strengthen the identification of objects and materials. To train a robotic system on a vast array of objects and incorporate vision, a significant dataset of tactile data, static image and video data should be collected for numerous materials, objects, clothing and human skin. A deep learning structure could then be utilised for the robotic system to learn the vast dataset of materials and objects. This information could be used to help determine the effort and grasping method required to successfully lift previously unlearned objects. This would provide a useful skill to assistive robots in a home environment when learning the affordances of objects they have not previously encountered. Secondly, the success of the challenging task of material identification presented in this paper proved invaluable in providing familiarity with the use of the BioTAC fingertip and the analysis of the data it presents. It became apparent that the high sensitivity of the BioTAC sensor means that it may be capable of measuring biomedical parameters, namely human vital signs such as a human pulse. Therefore, this also directs the research in this paper towards the automated collection and analysis of human vital signs based on tactile perceptions.

References

- Bhattacharjee, T., Wade, J., Chitalia, Y., and Kemp, C. (2016). Data-driven thermal recognition of contact with people and objects. In *2016 IEEE Haptics Symposium, HAPTICS 2016, Philadelphia, PA, USA, April 8-11, 2016*, pages 297–304.
- Bhattacharjee, T., Wade, J., and Kemp, C. (2015). Material recognition from heat transfer given varying initial conditions and short-duration contact. In *Proceedings of Robotics: Science and Systems*, Rome, Italy.
- Bishop, C. (1995). *Neural Networks for Pattern Recognition*. Oxford University Press.
- Chathuranga, D., Ho, V., and Hirai, S. (2013). Investigation of a biomimetic fingertip’s ability to discriminate fabrics based on surface textures. In *Advanced Intelligent Mechatronics (AIM), 2013 IEEE/ASME International Conference on*, pages 1667–1674.
- Dempster, A., Laird, N., and Rubin, D. (1977). Maximum likelihood from incomplete data via the em algorithm. *JOURNAL OF THE ROYAL STATISTICAL SOCIETY, SERIES B*, 39(1):1–38.
- Drimus, A., Kootstra, G., Bilberg, A., and Kragic, D. (2011). Classification of rigid and deformable objects using a novel tactile sensor. In *Advanced Robotics (ICAR), 2011 15th International Conference on*, pages 427–434.
- Fishel, J. and Loeb, G. (2012). Bayesian exploration for intelligent identification of textures. *Frontiers in Neurorobotics*, 6(4).
- Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., and Witten, I. (2009). The WEKA data mining software: an update. *SIGKDD Explorations*, 11(1):10–18.
- Ho, V., Araki, T., Makikawa, M., and Hirai, S. (2012). Experimental investigation of surface identification ability of a low-profile fabric tactile sensor. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 4497–4504.

- Hoelscher, J., Peters, J., and Hermans, T. (2015). Evaluation of tactile feature extraction for interactive object recognition. In *Proceedings of 15th IEEE-RAS International Conference on Humanoid Robots*, pages 310–317, Seoul. IEEE.
- Jamali, N. and Sammut, C. (2010). Material classification by tactile sensing using surface textures. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 2336–2341.
- Jamali, N. and Sammut, C. (2011). Majority voting: Material classification by tactile sensing using surface texture. *Robotics, IEEE Transactions on*, 27(3):508–521.
- Jurie, F. and Triggs, B. (2005). Creating efficient codebooks for visual recognition. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, volume 1, pages 604–610.
- Kerr, E., McGinnity, T., and Coleman, S. (2013). Material classification based on thermal properties - a robot and human evaluation. In *2013 IEEE International Conference on Robotics and Biomimetics, December 12-14 2013, Shenzhen, China*, pages 1048–1053.
- Kerr, E., McGinnity, T., and Coleman, S. (2014a). Material classification based on thermal and surface texture properties evaluated against human performance. In *2014 IEEE International Conference on Control, Automation, Robotics and Vision, ICARCV 2014, Singapore, Dec 2014*.
- Kerr, E., McGinnity, T., and Coleman, S. (2014b). Tactile approach to material classification - evaluated with human performance. In *Irish Machine Vision and Image Processing 2014, Northern Ireland*, pages 175–180.
- Lederman, S. and Klatzky, R. (1987). Hand movements: A window into haptic object recognition. *Cognitive Psychology*, 19(3):342–368.
- Lin, C., Erickson, T., Fishel, J., Wettels, N., and Loeb, G. (2009). Signal processing and fabrication of a biomimetic tactile sensor array with thermal, force and microvibration modalities. In *Robotics and Biomimetics (ROBIO), 2009 IEEE International Conference on*, pages 129–134.

- Lowe, D. (1999). Object recognition from local scale-invariant features. In *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, volume 2, pages 1150–1157.
- Luo, S., Liu, X., Althoefer, K., and Liu, H. (2015). Tactile object recognition with semi-supervised learning. In Liu, H., Kubota, N., Zhu, X., Dillmann, R., and Zhou, D., editors, *Intelligent Robotics and Applications*, volume 9245 of *Lecture Notes in Computer Science*, pages 15–26. Springer International Publishing.
- Martinez, A. and Kak, A. (2001). Pca versus lda. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 23(2):228–233.
- MATLAB (2013). *version 8.10.0 (R2013a)*. The MathWorks Inc., Natick, Massachusetts.
- Monkman, G. J. and Taylor, P. (1993). Thermal tactile sensing. *Robotics and Automation, IEEE Transactions on*, 9(3):313–318.
- Pezzementi, Z., Plaku, E., Reyda, C., and Hager, G. (2011). Tactile-object recognition from appearance information. *Robotics, IEEE Transactions on*, 27(3):473–487.
- R Core Team (2013). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Rish, I. (2001). An empirical study of the naive bayes classifier. Technical report.
- Syntouch (2013). The syntouch website.
- Varma, M. and Zisserman, A. (2005). A statistical approach to texture classification from single images. *Int. J. Comput. Vision*, 62(1-2):61–81.
- Wade, J., Bhattacharjee, T., and Kemp, C. (2015). A handheld device for the in situ acquisition of multimodal tactile sensing data. *CoRR*, abs/1511.03152.
- Xu, D., Loeb, G., and Fishel, J. (2013). Tactile identification of objects using bayesian exploration. In *IEEE International Conference on Robotics and Automation (ICRA) 2013*.