Multimodal Material Identification through Recursive Tactile Sensing

A. Gómez Eguíluz[1] and I. Rañó[1] and S.A. Coleman[1] and T.M. McGinnity[1,2]

[1] *A. Gómez Eguíluz, I. Rañó, S.A. Coleman and T.M. McGinnity are with the Intelligent Systems Research Centre, Ulster University, Northern Ireland, UK* {Gomez_Eguiluz-A, i.rano, sa.coleman}@ulster.ac.uk

[2] *T.M. McGinnity is also with the College of Science & Technology, Nottingham Trent University, UK* martin.mcginnity@ntu.ac.uk

Abstract

Tactile sensing has recently been used in robotics for object identification, grasping, and material identification. Although human tactile sensing is multimodal, existing material recognition approaches use vibration information only. Moreover, material identification through tactile sensing can be solved as an continuous process, yet state of the art approaches use a batch approach where readings are taken for at least one second. This work proposes a recursive multimodal (vibration and thermal) tactile material identification approach. Using the frequency response of the vibration induced by the material and a set of thermal features, we show that it is possible to accurately identify materials in less than half a second. We conducted an exhaustive comparison of our approach with commonly used vibration descriptors and machine learning algorithms for material identification such as k-Nearest Neighbour, Artificial Neural Network and Support Vector Machines. Experimental results show that our approach identifies materials faster than existing techniques and increase the classification accuracy when multiple sensor modalities are used.

Keywords: Recursive Material Classification, multimodal classification, Robotic Tactile Sensing, supervised learning

1. Introduction

Touch lies at the core of many human skills like grasping, temperature detection, and material identification, among others. While substantial research in robotics has focused on object identification through vision, human visual perception is often complemented with tactile sensing. In fact, experimental analysis of human visual and tactile representations show that they are closely linked, since the same brain regions are activated during visual and tactile exploration of objects [1]. Efforts to endow robots with a sense of touch have

recently attracted significant research interest [2] [3] leading to successful applications in object classification from grasping [4] and palpation [5], object state identification [6], grasping improvement and adaptation [7] [8] [9], and material identification. Until recent times, tactile sensing research was performed with ad hoc hardware making it impossible to perform a comparison between the different approaches to material recognition found in the literature. This paper proposes a novel multimodal approach to material identification and presents a comparison with state of the art approaches.

Existing research on material identification rely on batch surface recognition approaches, i.e. a whole sliding movement of the sensor over the material surface has to be performed for the identification to occur. Although this limits the identification speed, several excellent works provided unimodal (vibration only) batch approaches to material identification using different types of tactile sensors and techniques. In a pioneering work in tactile sensing for surface recognition [10] a finger with a microphone was used to detect the vibration induced by 3D printed textured surfaces. The authors defined a set of features to characterise the vibration signal captured by the microphone such as the modal frequency and power, and the average vibration amplitude. Using these features and the Fast Fourier Transform (FFT) of the raw signal projected, through Principal Component Analysis (PCA), as classification inputs the authors compare the classification performance of the k-Nearest Neighbour (kNN) algorithm. Although this work focuses on texture identification, not real material identification, it established the methodological approach of using features for tactile material recognition, which was followed by later works. The work presented in [11] uses a three-axial force sensor to classify 10 different paper types through two different techniques. In one approach, the most likely texture was selected by finding the minimum euclidean distance between the mean frequency spectrum of the evaluation set and training examples. A second approach used a set of five distinctive features (friction coefficient, mean, variance, kurtosis and spectrum slope of the vibration signal) as inputs to an Artificial Neural Network (ANN). The paper shows that both methods were proven to be equally accurate indicating the usefulness of the defined features. An accelerometer based vibrotactile sensor attached to a fingertip was used in [12] to measure the changes in the norm of the acceleration vector and to create spectrotemporal histograms as features for classification. A set of five exploratory movements were performed to gather data from the material surfaces, and Support Vector Machines (SVM) and kNN were used to successfully identify 20 materials.

A tactile sensor measuring the strain applied on the finger surface is presented in [13]. After segmenting the input signal, removing the average and band-pass filtering, the authors extract a set of features consisting of five peaks identified on the smoothed FFT profile. Combining these features with the average strain readings, the authors compare different machine learning approaches, which successfully differentiate between nine materials with high accuracy. Another comparison of machine learning algorithms for material identification based on tactile sensing is presented in [14], where two kernel methods (SVM and Regularised Least Square), and one neural network were used to

classify pairs of materials based on the raw strain measurements of the sensor. Although the authors concluded that SVM showed the best trade-off between classification accuracy and computational complexity, they observed that the raw sensor signals did not provide a good discrimination performance compared with other works. [15] presents a tactile micro-sensor able to differentiate surfaces with spatial periods within a 40 μm difference. This sensor was used to classify textiles through a robotic finger that slid across the materials for two seconds. Using wavelet transforms, the evolution over time of features like the peak power were obtained and fed into a kNN classifier. The work in [16] presents a texture based material classification through SVMs, with a set of temporal domain features obtained from one second time windows as input. Specifically, the components of the feature vector were obtained from the accelerometer in the finger, and whitened individually before feeding the classifier. This work splits the whole set of readings into short time intervals to obtain more training data, but the approach is still batch based since the authors do not exploit sequential information for classification. Another accelerometer based fingertip texture recognition approach is presented in [17], where seven different fabrics are classified based on a mixture of temporal (acceleration variance) and frequency (power spectra) features. Since these features do not provide enough discriminative power on their own, the authors use a neural network with the FFT coefficients over a given frequency range as the input. Recently, [18] explored real time classification of eight materials using a soft three axis tactile sensor with a new set of features. Specifically, the mean value of the three dimensional vibration signal, and the Frobenius norm of the covariance matrix were used to train and classify a cascade of binary SVM classifiers, grouping materials together in each classification step. Although their approach is fast and accurate, it becomes computationally expensive for large number of materials, since the depth of the cascade, and the number of classifiers to train grows quickly. A flexible tactile skin attached to the fingers of a DLR Hand-II was used in [19] to discriminate 6 tubes made of different materials by stroking at them with the thumb and index fingers. The authors present a comparative study of different classification techniques and conclude that a convolutional deep learning network outperforms other classifiers such as kNN and SVM, especially when repositioning the tactile skin. In another recent work [20] a custom made tactile sensor was used for material identification using only the heat transfer from the sensor to an object. The authors showed that, although changes on the experiment set-up (initial conditions, ambient temperature and contact duration) have an impact on the performance, a multi-class SVM can classify eleven materials with high accuracy.

Nearly all the works mentioned so far rely on vibration signals obtained from ad hoc sensors. However the appearance of the SynTouch BioTAC commercial fingertip has made tactile sensing widely accessible and provided the opportunity of performing multimodal sensing. Using this new tactile sensor, a series of works [21] [22] [23] implement Bayesian exploratory movements to classify materials using vibration and temperature signals. Bayesian inference was used to implement different exploratory movements in order to increase the classifica-

3

tion certainty. The classification itself was based on features extracted from the sequences and the external force needed to slide the fingertip across the material, i.e. information about the friction coefficient of the material and the BioTAC's rubber skin was included. The combination of several exploratory movements enables a high material identification rate with a large range of materials. One of the first multimodal approaches to material identification is presented in [24] where raw temperature and vibration signals were projected through PCA and used as inputs to an ANN. The authors show that this classifier outperforms humans in similar experimental conditions. The approach presented in [25] identifies among 49 objects with high accuracy using multimodal data from 5 different object explorations. The authors performed a thorough analysis of tactile features found in the literature and concluded that simple descriptors, such as average values of the filtered signal, outperforms more sophisticated feature extraction techniques.

All these works achieve good material recognition ratios using one or several batch readings and combinations of exploratory movements. However, some materials might have a characteristic texture or special thermal properties and might therefore be identified faster than others, i.e. using shorter readings. Moreover, combining these two modalities might improve the identification accuracy, for instance, for materials with similar textures but different thermal properties. This paper extends our previous works [26, 27] presenting a multimodal recursive identification approach to material identification through tactile sensing. The contribution of this paper is threefold. First, we present an approach that allows fast and very accurate material identification using vibration signals, and can be extended to other sensing modalities. Second, we show that including temperature information significantly reduces the time needed to identify the correct material. Finally, we perform a comparative study of the proposed approach with state of the art tactile material identification techniques such as kNN, ANN and SVM, comparing also the standard features to characterise the vibration signal with the selected feature vector. the rest of the paper is organised as follows. Section 2 presents the methodology used for the proposed continuous material identification approach, including the processing of the raw data, probabilistic modelling of the readings of each material, and the recursive classification technique. This section also introduces a brief description of the state of the art selected features to compare with the proposed approach. Experimental results and a comparison with state-of-the-art approaches are presented in Section 3. Section 4 concludes the paper and highlights future research.

2. Tactile signal processing and material identification

We use the SynTouch BioTAC finger tip [28] as our experimental platform; it provides different types of sensing modalities with different sampling frequencies integrated in a single sensing unit. The sensor is composed of a conductive fluid that separates the wrinkle rubber skin and the rigidcore in which the sensing elements are located. An array of impedance sensors provides information about
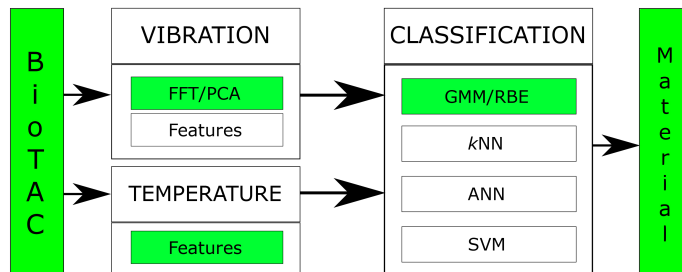
4

Figure 1: Scheme of the pipeline followed during the experiments. Green coloured boxes represent the proposed approach, while plain boxes represent the features and classifiers used for comparison.

the deformation of the rubber skin upon contact with a surface. A small heating system, jointly with a thermistor, constitutes the thermal sensing subsystem, to generate temperature and heat flow measurements. The BioTAC measures the temperature of the core and thermal flux from the device to the surface. Finally, a pressure sensor in the core provides pressure and vibration measures from the finger when it touches and slides over a surface. The interaction of the object surface and the rubber wrinkles produces a vibration in the skin which is transferred to the fluid, and measured by the pressure sensor. The pressure signal is then low pass filtered to obtain a pressure measurement, and band pass filtered to generate the vibration signal.

Figure 1 shows at a glance the classification pipeline of the proposed approach including the comparisons we performed with state of the art techniques. While the coloured boxes represent the proposed approach, the plain boxes show features and classifiers commonly found in the literature. Moreover, we added our thermal features to standard approaches to perform a fair comparison. In a nutshell, we propose using the PCA projected Fast Fourier Transform (FFT) of the vibration signal, and the thermal power transferred to the material in a recursive Bayesian estimation algorithm to classify the materials. The rest of this section will provide details on how signals are processed and classified, both in the proposed approach and for the comparison performed.

2.1. Signal Processing

This section presents the processing conducted on the BioTAC vibration and thermal signals, shown in the first column of Figure 1. Standard approaches to vibration based material identification rely on features extracted from the signal in temporal or frequency domains. As these features typically discard frequency information, which is highly relevant to identify a surface material, our approach uses a combination of the Fourier coefficients of the vibration signal to classify the surfaces. To perform multimodal identification we present novel thermal features, one of them grounded in the physical process of heat transfer (i.e. using thermal power loss in the BioTAC) computed from the heat

5

flux and impedance readings. In our comparison, we test a set of six features found in the literature, extracted from the temporal and frequency domains (see Section 2.1.1).

2.1.1. Signal processing and vibration features

The vibration signal induced in the liquid gel by the interaction of the fingertip and the material is a combination of oscillatory signals with a frequency spectrum dependent on the material. Therefore, the Fourier Transforms (FT) of the vibration signal can be used to characterise the material texture. Specifically, we use the Fast Fourier Transform (FFT) algorithm to convert the vibration signal, $\nu(t)$ into the frequency domain $\rho(\omega)$. We compute the FFT for small non-overlapping windows of duration Δt although overlapping windows could also be used. The existing literature [10] points out that good discrimination results are obtained by using the frequency response in the range between $2\ Hz$ and $500\ Hz$, so we use this range for consistency.

For a Δt window, the FFT is a high dimensional vector of complex numbers $\rho(w) \in \mathbb{C}^{d_{\Delta t}}$, where the dimension $d_{\Delta t}$ depends on the selected time interval Δt and the sampling period. We used Principal Component Analysis (PCA) over $\rho(\omega)$ to obtain a lower dimensional feature vector $\bar{\rho} \in \Re^d$, where $\bar{\rho}$ is the result of projecting $\rho(\omega)$ and $d \ll d_{\Delta t}$. The complex nature of $\rho(\omega)$ requires special care when processing the FFT through PCA. Although PCA can be used with complex vectors, we found that the module of the FFT carries the discriminative information for material identification. However, we performed the centring process over the complex vectors, i.e. the complex mean was subtracted before obtaining the projection matrix of the FFT modulus. Applying PCA to the centred vectors $\rho - \mu_\rho$, with mean complex values of the FFTs $\mu_\rho \in \mathbb{C}^{d_{\Delta t}}$, helps to cancel the random noise in the FFT, as white noise's phase is random while its amplitude is fixed.

Many existing works in vibration based tactile material identification use features obtained from the temporal signal instead of the FFT. These features encode general properties of the vibration readings like its total energy, or statistical properties (central moments of the signal), and reduce the dimension of the vector used to characterise the material. This process eliminates relevant information of the surface and, in some cases, is not very efficient computationally, e.g. obtaining frequency domain features requires the computation of the FFT anyway. For comparison purposes we defined a vector $\mathbf{x}$ of six features, $\mathbf{x} \in \Re^6$, found in the related literature as discussed below, that combine independent information in the frequency and temporal domains.

- The filtered average energy [21] is computed from the band pass filtered FFT as:

$$x_1 \quad = \quad \frac{1}{n}\sum_{i=1}^{n}|\rho_i|^2, \tag{1}$$

  where $\rho_i = \rho(\omega_i)$ is the Discrete Fourier Transform (DFT) of the vibration,

6

and $\omega_1 = \omega_m$, $\omega_n = \omega_M$ are the limit frequencies of the band pass filter, and n is the number of samples within the frequency range $[\omega_m, \omega_M]$.

- The high frequency to low frequency energy ratio [17] is another frequential feature computed as:

$$x_2 \quad = \quad \frac{\sum\limits_{i=1}^{k} |\rho_i|^2}{\sum\limits_{i=k+1}^{n} |\rho_i|^2} \tag{2}$$

where $\rho_i = \rho(\omega_i)$, $\omega_k <= \omega_{th}$, $\omega_{k+1} > \omega_{th}$ and ω_{th} is a threshold that defines what low and high frequencies are (as in [17] we set $\omega_{th} = 100Hz$).

- The spectral centroid [21] is computed from the FFT as the weighted frequency as follows:

$$x_3 \quad = \quad \frac{\sum\limits_{i=1}^{n} |\rho_i|^2 \omega_i}{\sum\limits_{i=1}^{n} |\rho_i|^2} \tag{3}$$

where ρ_i is defined as above for the corresponding frequency ω_i.

- The signal temporal variance, used as a feature in [16] and [17], can be computed as:

$$x_4 \quad = \quad \frac{1}{n} \sum\limits_{i=1}^{n} (\nu_i - \mu_\nu)^2 \tag{4}$$

where $\nu_i = \nu(t_i)$ is the temporal signal of the vibration, μ_ν is its average value, and n is the number of vibration samples considered, which are also used to compute the temporal Skewness and Kurtosis of the vibration signal.

- The temporal Skewness of the signal [16], a statistic of the distribution of vibration values, is computed as:

$$x_5 \quad = \quad \frac{\frac{1}{n} \sum\limits_{i=1}^{n} (\nu_i - \mu_\nu)^3}{\left[\frac{1}{n-1} \sum\limits_{i=1}^{n} (\nu_i - \mu_\nu)^2 \right]^{\frac{3}{2}}} \tag{5}$$

- The temporal Kurtosis of the signal, used in [16], is a higher central moment of the vibration computed as:

$$x_6 \quad = \quad \frac{\frac{1}{n} \sum\limits_{i=1}^{n} (\nu_i - \mu_\nu)^4}{\left[\frac{1}{n} \sum\limits_{i=1}^{n} (\nu_i - \mu_\nu)^2 \right]^{2}} - 3 \tag{6}$$

7

For simplicity, the rest of this section will refer to the vibration signal $\bar{\rho}$, the PCA projected FFT, as the proposed vibration feature vector. Nevertheless, when we perform the comparison $\bar{\rho}$ is substituted by **x** as an input for the classifiers.

2.1.2. Temperature

The BioTAC thermal system has a heating device that maintains the core temperature typically higher than room temperature. It then measures the temperature at the core, which when in contact with a material decreases due to the heat flux leaving the BioTAC finger through the contact surface. The thermal energy lost depends on the temperature difference between the finger and the external material, the contact area, and the thermal conductivity of the material. The thermal power lost (thermal energy per unit of time $\frac{\partial E}{\partial t}$) is defined as the integral of the heat flux over the contact surface:

$$\frac{\partial E}{\partial t} = \oint_S \vec{\phi} \cdot \vec{dS} \tag{7}$$

where $\vec{\phi}$ is the heat flux, and the integral is computed on the contact area S between the two materials. We can assume the flux $\vec{\phi}$ always leaves the finger, and its modulus – measured by the BioTAC sensor – increases with the temperature difference and the thermal conductivity of the material (although the thermal conductivity of the finger might play a role). Furthermore, assuming friction is too weak to generate thermal energy, and the thermal conductivity of the air compared with the material is small, the thermal flux will be directed towards the surface normal. Since the contact area is typically small, we can approximate the power loss as the product of the average flux modulus $\bar{\phi}$ by the contact area. Since all objects to be identified are at room temperature, the temperature difference depends only on the temperature of the core, and therefore $\frac{\bar{\phi}A}{T}$ is a measure of the thermal conductivity of the material, where A is the contact area, and T is the finger core temperature. To estimate this conductivity we need to compute the contact area, which can be estimated from the readings of the 19 electrodes placed in the core under the BioTAC's skin. The electrodes measure of impedance is related to the thickness of the fluid between the core and the skin at their corresponding locations. The skin deformation, upon contact, makes the readings in the surrounding electrodes decrease, and, therefore, electrodes with a negative value w.r.t. their resting level indicate contact. We approximate the contact area for each electrode i as a circle of radius r_i equal to half the distance between the electrode and its nearest neighbour. Hence, the contact area of the fingertip and the material is approximated as a weighted average of small circular areas as:

$$A = \sum_i \lambda_i \pi r_i^2 \tag{8}$$

where $\lambda_i \in [0,1]$ is a scale factor that depends on the value returned by electrode i. This scaling factor λ_i is a piece wise linear function of the average impedance

value $\bar{e}_i$ of each electrode during the time interval Δt. At the resting level (or above) λ_i is zero, and it increases to 1 for decreasing impedances down to a fixed minimum threshold e_m (in our case $e_m = -400$), and is 1 for values below that threshold. Besides using the power loss per unit of temperature we experimentally found that two other thermal features improved the material identification accuracy. For each sequence of Δt readings we performed a linear regression of the thermal flux as a function of time and used the slope and the regression error as additional features. Therefore, we obtained a three dimensional heat based feature vector θ for each Δt time window.

2.2. Classification

Because our aim is to classify materials in a continuous manner we implemented a Recursive Bayesian Estimation algorithm to maintain over time a probability of the materials to identify. Although recursive estimation algorithms are widely used in robotics they have not been used for material identification, a problem treated as batch classification in the literature. Let us denote by M the discrete random variable encoding the N materials to be identified, i.e. $\{m^1, m^2, \cdots, m^N\}$, $P \in \Re^d$ is the random vector characterising the vibration signal (FFT projected through PCA, or the six dimensional vector of selected vibration features), and $\Theta \in \Re^3$ the vector of thermal features defined in Section 2.1. We first estimate the likelihood functions of the feature vectors for each material and sensing modality, $p(P = \bar{\rho}|M = m^j)$ and $p(\Theta = \theta|M = m^j)$, using mixtures of Gaussian distributions. Therefore Gaussian Mixture Models (GMM) of the form $\sum_{i=1}^{K_j} \alpha_i^j \mathcal{N}\left(\mu_i^j, \Sigma_i^j\right)$ are obtained for each material, where K_j is the number of Gaussians in the mixture for material m^j, α_i^j is the weight of the i-th Gaussian, and $\mathcal{N}(\mu, \Sigma)$ denotes a normal distribution with mean μ and covariance Σ. These GMMs are obtained independently from the vibration and thermal feature vectors. The parameters of the GMMs are obtained using the Expectation-Maximisation (EM) algorithm, where the number of Gaussians was chosen individually for each material and modality based on the decrease of the likelihood of the training data, i.e. increasing the number of Gaussian components to find the point at which the growth in the likelihood started to slow down.

Having models of the likelihood functions $p(\bar{\rho}|m^j)$ and $p(\theta|m^j)$ for the materials would already allow to identify the materials using a maximum likelihood (ML) classifier. However, given a set of prior material probabilities $p(m^j)$, one could also estimate, using the Bayes rule, the posterior probabilities of each material given the current vibration and thermal readings $p(m^j|\bar{\rho}, \theta)$, and classify according to that probability (maximum a posteriori, MAP). For uninformative priors, $p(m^j) = \frac{1}{N}$ for all j, the classification result of ML and MAP approaches are identical. However, if we take the posterior probability as the prior for the next set of readings, i.e. the material prior $p(m_{k+1}^j)$ at time step $k+1$ will be the posterior from the previous iteration $p(m_k^j|\bar{\rho}_k, \theta_k)$, we can expect to get a more accurate estimate of the probability of all materials given the current sequence

of readings. Therefore, when a material is presented to the robot for identification, the initial prior probabilities are distributed evenly among all materials. Then we obtain from the BioTAC data stream the vibration and thermal feature vectors, $\bar{\rho}_k$ and θ_k, at time intervals Δt, and update the material probabilities using:

$$p(m_k^j|\bar{\rho}_k, \theta_k) \quad = \quad \frac{p(\bar{\rho}_k, \theta_k|m_k^j)p(m_{k-1}^j|\bar{\rho}_{k-1}, \theta_{k-1})}{p(\bar{\rho}_k, \theta_k)} \tag{9}$$

where the normalisation constant $p(\bar{\rho}_k, \theta_k)$ is:

$$p(\bar{\rho}_k, \theta_k) \quad = \quad \sum_{i=1}^{N} p(\bar{\rho}_k, \theta_k|m_k^i)p(m_{k-1}^i|\bar{\rho}_{k-1}, \theta_{k-1}) \tag{10}$$

and, we assume the vibration and thermal features are conditionally independent, therefore $p(\bar{\rho}_k, \theta_k|m_k^j) = p(\bar{\rho}_k|m_k^j)p(\theta_k|m_k^j)$, with each individual likelihood function given by the corresponding GMM model for material m^j. It is worth noting that, although the proposed approach relies on vibration and thermal information, it could be easily extended to more sensing modalities assuming they are conditionally independent.

In each iteration the algorithm generates an updated probability $p(m_k^j|\bar{\rho}_k, \theta_k)$ for all materials, $j = 1, \cdots, N$, and the one with the highest posterior can be considered the one presented to the robot. Alternatively, one could define a confidence threshold to decide for a material only if the posterior probability is high enough. Instead of classifying the perceived texture from a batch reading, the proposed algorithm generates increasingly confident updates of the material probability.

Besides comparing our proposed input vector, FFT-PCA, to state of the art features, we also performed a comparison with commonly used machine learning algorithms for material identification, specifically kNN, ANN, and SVM. Because these learning techniques are used with batch readings, while our approach is recursive, performing a fair comparison, is not straight forward. Moreover, once the classifiers are trained, in the case of our approach the core computational load is obtaining the FFT-PCA vector (which can be done in real time). In contrast, some classification mechanisms (e.g. kNN or SVM) can hardly be applied in real time because of the extensive computation required for large training datasets.

2.3. A note on sensor fusion

The combination of vibration and thermal information in our recursive Bayesian estimation framework is straightforward assuming both modalities are conditionally independent. The likelihood function of the combined data is simply the product of the individual likelihood functions. In the case of the compared classifiers (SVM, ANN, and kNN) we have several ways in which we can combine the data. We could, for instance, classify each modality independently and then use another classification level to make a final decision on the material. Our

approach was instead to define a single feature vector $\mathbf{X}$ joining together the vibration and thermal features $\mathbf{X} = [\bar{\rho}\,\theta]$ (alternatively $\mathbf{X} = [\mathbf{x}\,\theta]$), i.e. performing the sensor fusion at the feature level. This is simpler than a combination of classifiers as there is no need to evaluate several classifiers to make the final decision. However, when using signals from several sources we face the problem of the different numerical scales in each component of the vector provided as an input to the classifier. To avoid this issue, we implemented data whitening as a pre-processing step for the learning techniques used, i.e. the training and testing datasets were transformed to have zero mean and covariance equal to the identity matrix.

3. Results of the Recursive material identification and comparison

Our experimental setup, shown in Figure 2, consists of a turntable moved by a step motor through a set of reduction gears. The motor is controlled by an Arduino board, running code to set the turning speed of the motor to rotate at 4 seconds per lap. The fingertip is attached to a worm drive bar, that can be moved up and down by a second motor also controlled by the Arduino board. When the bar is driven down, the BioTAC touches the material on the turn table. To collect the training and testing datasets, we first wait until the thermal system has reached a stable temperature, then set the speed of the turntable and move the fingertip down until it touches the material surface. After a few seconds we collect readings continuously storing all the information provided by the BioTAC sensor running at 4.4 KHz, the maximum communication rate. This means that the sampling frequency of the vibration signal is 2.2 KHz, while the absolute temperature, impedance and thermal flux have frequencies
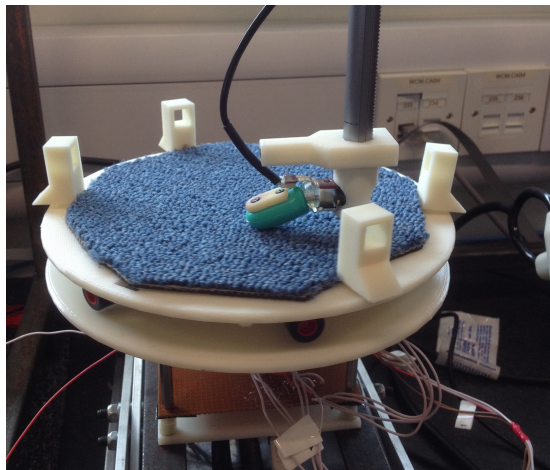

Figure 2: Experimental setup consisting of a BioTAC tactile sensor attached to a worm drive bar and a turntable that rotates a disk with the material attached.

Figure 3: Materials used in the experiments

of 100 Hz. It is worth noting that the temperature of the finger takes time to reach a steady state. We allowed the finger to warm up until the temperature was stable before to gathered any of the data-sets.

We selected a total of 34 materials for our identification experiments, which were purposely selected to include different material groups and several materials in each group. Table 1 lists all the materials we used, also shown in Figure 3. The selection contains materials with very different textures such as cork and glass from a mirror, but also some fabrics and materials which are alike. Moreover, for some of the fabrics we used both sides of the cloth, like materials 2 and 3, and materials 30 and 31, which have different textures and this increases the number of materials. We also include genuine and synthetic leather, materials 16 and 18 respectively, while materials 13 and 14 were obtained using both sides of a padded envelope, material 13 being the side of the bubbles, i.e. plastic, and material 14 the paper side. Other pairs correspond to both sides of the same object such as materials 6-7, 22-23 and 28-29, although their surfaces were clearly different. We collected continuous sequences of 10 minutes of data for each material, which were used for all the experiments presented in this paper to compare the different classification techniques with the two defined feature inputs. A main advantage of our algorithm is that it minimises the time needed to identify a material, as the classification is performed over several small data windows, while batch approaches decide on a single, typically larger window. Therefore, we cannot compare the time interval required for successful material identification for the proposed approach and the assessed machine learning techniques. We measured first the average number of iterations the recursive Bayesian estimation takes to successfully classify a material as a performance measure. As we will see, this time measurement was used to perform a fair comparison between the recursive identification and kNN, ANN, and SVMs.

12

Table 1: List of materials used in all our classification experiments

Idx	Material	Idx	Material
1	Synthetic Green fabric	18	Genuine leather
2	Synthetic Pink fabric 1	19	Linen
3	Synthetic Pink fabric 2	20	Mirror
4	Cardboard box	21	Normal paper
5	Cardboard disk	22	Ping pong paddle 1
6	Carpet	23	Ping pong paddle 2
7	Rubber	24	Plastic
8	Baize	25	Plastic dish
9	Can of drink	26	Rough fabric
10	Copper	27	Slate stone
11	Cork	28	Sponge 1
12	100% Cotton	29	Sponge 2
13	Padded envelope 1	30	Leopard fabric 1
14	Padded envelope 2	31	Leopard fabric 2
15	Aluminium	32	Watercolour paper
16	Synthetic leather	33	Wood
17	Floor tiles	34	Peach skin fabric

3.1. RBE approach to material identification using vibration

To illustrate the advantage of a multimodal approach to tactile material identification we first performed the task using vibration information only. As explained in Section 2, we compute the FFT of the vibration signal, band pass filter it and project it into a lower dimensional space. The time domain signal was split into non-overlapping windows of $\Delta t = 0.25$ seconds, which sets the frame rate for the continuous material classification to 4 Hz. This window corresponds to a sequence of 550 pressure samples, which, after selecting only the $2 - 500 \ Hz$ frequency range, became 124 FFT spectral components. It is worth noting that, given the short time interval selected, the frequency interval between spectral components is around 4 Hz, that is our actual minimum frequency in the band pass filter. As stated in Section 2.1.1 we projected the 124 spectral components to a 16 dimensional space, computing the complex average with the training dataset for all 34 materials, while projecting their amplitude vectors, therefore the input to the recursive classification will be $\bar{\rho} \in \Re^{16}$. The number of reduced dimension was chosen to keep 97% of the total variance of the original vectors.

The FFT-PCA vectors were then used to model the likelihood functions of the data for each material as GMMs (trained using EM algorithm), and used in our RBE approach. Obviously, in the case of vibration only material identification, the thermal part of the likelihood function is not considered. To evaluate the proposed approach we performed a 10-fold cross validation on the 10 minutes of data sequences recorded for each material. We split the

13

dataset into 10 groups of 60 seconds sequential readings using alternatively 9 of the groups for training and one for testing. Our experiments showed that once a material is correctly identified by the RBE approach, the rest of the sequence is properly classified. The evaluation procedure used 60 seconds of test time series of every material starting from the first sample and stopping once the recognition was successful, storing the number of required iterations. The first sample was then discarded (creating a shorter test sequence) and the process was repeated until the sequence was too short for the algorithm to identify the material (typically sequences of 0.25 seconds). This process was repeated for each testing sequence in the 10-fold cross validation, and the average number of iterations, as well as its standard deviation, taken to identify each material was computed. Because the RBE algorithm always correctly identifies the material, the ratio of successful material classification cannot be considered as a performance measure. Therefore, as a first step to compare our approach to other state of the art classifiers we evaluated how long it takes our approach to detect a material based on the maximum posterior probability $p(m_k^j|\bar{\rho}_k)$ (or alternatively using the feature vector $p(m_k^j|x_k)$). We computed the average number of iterations and its standard deviation, and translated into seconds multiplying by Δt (0.25). These time values will be used in the next section to compare the proposed RBE approach with the rest of the learning techniques.

Table 2 shows the average time (in seconds) required to correctly identify each material over the 10 trials of the 10-fold crossvalidation. The results illustrate that, when using the FFT-PCA components to model the material vibrations, most of the time the algorithm identifies the materials within less than 0.5 seconds, the average time (0.36 seconds) is just above one iteration, while the standard deviation (0.2 seconds) is just below one iteration. Although, obviously the recognition time is a discrete random variable with positive skewness, if we attend to the probabilities of a normal distribution, we can conclude that, with nearly 99% probability, the time required for a successful identification is less than one second, i.e. average time plus three times its standard deviation. The table also shows the average time taken for identification and corresponding standard deviations when the six dimensional feature vector x is used instead. The 10-fold cross-validation results show that the time needed to recognise the correct material is on average 0.67 seconds (more than two iterations), while the standard deviation is also just above two iterations (0.51 seconds). Following the same reasoning high probability of successful material identification is above 2 seconds, i.e. $\mu + 3\sigma$, when modelling the vibration response through the feature vector x. Therefore, Table 2 quantifies the (expected) benefit of using the whole frequency spectrum instead of a set of features. Moreover, it shows why previous material identification works require over one second long vibration readings. As it can be seen from the table, using the PCA-FFT features not only reduces the average identification time, but also the standard deviation, making this input vector less sensitive to variations in the signal. The improvement in the average detection time when using PCA-FFT is above 45%, with improvement of 81% for material 10, and over 60% for materials 17, 23, 24, and

Table 2: Average time (in seconds) needed for material classification using the FFT of the vibration signals and 10-fold cross validation. First number is the average time and second is the standard deviation.

Material	1	2	3	4	5
FFT+PCA	0.39+0.21	0.35+0.19	0.35+0.19	0.43+0.28	0.39+0.25
Features	0.74+0.62	0.64+0.56	0.56+0.42	0.68+0.41	0.72+0.64
Material	6	7	8	9	10
FFT+PCA	0.3+0.12	0.33+0.14	0.44+0.26	0.34+0.16	0.3+0.2
Features	0.34+0.20	0.64+0.51	0.90+0.46	0.79+0.59	1.58+1.29
Material	11	12	13	14	15
FFT+PCA	0.3+0.11	0.35+0.18	0.33+0.16	0.34+0.17	0.3+0.16
Features	0.48+0.34	0.54+0.53	0.76+0.60	0.49+0.28	0.56+0.48
Material	16	17	18	19	20
FFT+PCA	0.44+0.25	0.46+0.34	0.32+0.14	0.49+0.35	0.29+0.11
Features	0.74+0.59	1.24+1.29	0.41+0.28	0.77+0.63	0.41+0.28
Material	21	22	23	24	25
FFT+PCA	0.39+0.25	0.38+0.22	0.36+0.19	0.38+0.22	0.38+0.21
Features	0.42+0.28	0.55+0.41	1.06+0.70	1.01+0.83	1.15+0.79
Material	26	27	28	29	30
FFT+PCA	0.27+0.07	0.42+0.25	0.29+0.13	0.32+0.14	0.27+0.08
Features	0.55+0.41	1.06+0.70	1.01+0.83	0.60+0.44	0.30+0.16
Material	31	32	33	34	**Avg.**
FFT+PCA	0.27+0.08	0.46+0.3	0.52+0.34	0.4+0.26	**0.36+0.2**
Features	0.35+0.25	0.29+0.12	0.98+0.70	0.65+0.63	**0.67+0.51**

25. That in turn means that using the whole frequency spectrum information (from 2Hz to 500Hz) makes material identification more than twice as fast for these materials. Interestingly, the watercolour paper (material 32) seems to be the only one identified faster using the state of the art features than using the FFT.

3.1.1. Evaluating the classification accuracy against standard Machine Learning methods

In order to perform a comparison between our proposed RBE classification mechanism and the most used algorithms in the material identification literature (kNN, ANN, and SVM) we provided the same amount of information to each classification technique. To do so we used the results from the classification time already presented for both the FFT-PCA vector $\bar{\rho}$ and the feature vector $\mathbf{x}$. Our results indicate that the RBE algorithm typically takes approximately one second to achieve successful identification of a material when using the vibration model based on $\bar{\rho}$, and more than two seconds when using the vibration

feature vector $\mathbf{x}$ ($\mu + 3\sigma$ seconds). Therefore, from the same datasets of the 34 materials, we generated new FFT-PCA vectors from one second length data streams, and new feature vectors from 2.25 seconds length data streams to train and evaluate the kNN, ANN and SVM classifiers. In the case of the RBE approach to material identification this means that when the algorithm takes more than four iterations, i.e. 1 second (nine iterations, i.e. 2.25 seconds, in the case of the feature vector) to identify the material we consider it as a classification failure. This provides a fair comparison between the two methods as they are all required to classify the surfaces in the same amount of time with the same amount of information (although in the case of the RBE it is seen in shorter time windows). As taking a longer temporal sequence of vibration readings changes the dimension of the FFT vector $\rho(\omega)$, although we kept the same amount of variance (97% of the original signal), the training vectors for these algorithms have a higher dimension, specifically the $\bar{\rho}$ used to train and evaluate the kNN, ANN, and SVM have dimension 49, after band pass filtering (between 2Hz and 500Hz) and PCA projection.

For the KNN algorithm, a systematic search procedure using Euclidean distance and 10-fold cross-validation was implemented to identify the optimal value of k (from $k = 1$ to $k = 21$), which was found to be $k = 5$ for both the FFT-PCA vector and the state of the art vibration feature vector. Similarly, the topology of the ANNs was found running 10-fold cross-validation tests with two hidden layer networks and different numbers of units on each layer. The number of outputs corresponds to the number of materials to be identified, i.e. 34 outputs encoding the materials as a binary output vector, while the number of inputs was 49 for the FFT-PCA input vectors and 6 for the feature vector. We experimentally found that the best trade-off between the network size and the identification accuracy was achieved for networks with 65 units in both hidden layers for the FFT-PCA vector input, and 34 units in both hidden layers for the feature vector $\mathbf{x}$, a reflection of the higher dimensionality of the former input vector. The input vectors of the training and testing datasets were pre-processed through whitening (i.e. scaled to have zero mean and identity covariance matrix), and the networks were trained using the Levenberg-Marquardt algorithm. Finally, we designed a One-Vs-One strategy multiclass approach for the classification using SVMs, as we found experimentally to perform better than the One-Vs-All alternative. Although it provided better identification accuracy, the One-Vs-One approach needs to perform $N(N-1)/2$ binary classifications for each input pattern, which in our case (i.e. 34 materials) corresponds to 561 binary classifications. Therefore, this classification approach is the most computationally expensive one.We used the standard Radial Basis Function (RBF) kernel for all the SVMs, while its scale parameter was chosen through a grid search.

As already mentioned we compared the classification accuracy of these three techniques, typically used in tactile material identification, with the proposed RBE approach "with a recognition timeout". As expected, all techniques worked better when the PCA projected Fourier transform was used as an input (kNN 98%, ANN 85.5%, SVM 92.2%, and RBE 96.7%), instead of the feature vector

(kNN 73.2%, ANN 64.9%, SVM 75.2%, and RBE 93.8%). It is worth noting that besides the performance decrease, the feature vectors were collected over more than twice the time, i.e. worse identification rate and much slower identification time. Interestingly, the recursive material identification is the best performing classifier when the state of the art feature vector is used (93.8%), since it actually performs nine different classifications with much lower certainties, but combines the material probabilities in an optimal way. All in all, the kNN algorithm outperforms the rest with a 98% average accuracy over the 10-fold cross-validation, followed by the RBE 96.7%, although the latter can typically identify the materials faster, meaning that the material is often identified in less than one second.

3.2. Multimodal Material Identification.

This section presents the experimental results of the multimodal approach to material identification using vibration information and the thermal features (θ_k) described in Section 2.1. We again classify the 34 materials listed in Table 1 and compare the recursive identification against the commonly used machine learning techniques. Similarly to the vibration only approach, we found that, if enough time is given, our recursive approach always identifies the materials correctly, in a variable number of iterations. This occurs for both multimodal approaches FFT-PCA and thermal features, and vibration and thermal features. It is worth mentioning that processing, training, and testing procedures are identical to those presented in the vibration only section, i.e. $\Delta t = 0.25$ seconds, the vibration is band pass filtered, likelihood models are mixtures of Gaussians, and the priors are uninformative. Moreover, the thermal features were computed with the same time window, but because of the differences in sampling frequencies, a 0.25 seconds time window corresponds to 25 impedance and heat flux readings. To estimate the thermal power loss we averaged the heat flux and impedance readings during the measuring interval, while the heat flux slope and error are computed with all the samples. The number of Gaussian functions in the mixture model for the thermal features was typically 2, except for materials 4, 8, 12 and 23 which we estimated as 3 using the change in the training data likelihood. Like in the case presented in Section 3.1, we performed 10-fold cross-validation on 10 minutes of data sequences to identify the average time required by the RBE to successfully identify the materials.

Table 3 shows the average identification time results across the ten iterations for both sets of input vectors (FFT-PCA vs. features both enhanced with thermal information). Just like in the case of the vibration only identification, using the PCA projected Fourier coefficients results in faster material classification, although in this case the time difference is not as large as before, i.e. the improvement of the classification using features is much bigger when thermal information is included. Indeed, including thermal information results in a 40% reduction in the average identification time when combined with the feature vector $\mathbf{x}$ (0.4 seconds vs. 0.67 seconds in table 2), and 22% improvement when combined with the FFT-PCA vector $\bar{\rho}$ (0.28 seconds vs. 0.36 seconds in table 2). This brings the multimodal average identification time (and the

Table 3: Temperature for recognition improvement. Average time needed using 10-fold cross validation. First number is the average time and second is the standard deviation.

Material	1	2	3	4	5
FFT+PCA	0.28+0.08	0.28+0.1	0.29+0.11	0.28+0.1	0.31+0.14
Features	0.32+0.16	0.39+0.29	0.39+0.28	0.29+0.11	0.59+0.54
Material	6	7	8	9	10
FFT+PCA	0.26+0.06	0.27+0.08	0.33+0.18	0.32+0.17	0.26+0.03
Features	0.29+0.12	0.30+0.14	0.50+0.31	0.64+0.50	0.41+0.35
Material	11	12	13	14	15
FFT+PCA	0.25+0.02	0.28+0.09	0.29+0.13	0.26+0.06	0.25+0.01
Features	0.26+0.04	0.31+0.15	0.71+0.60	0.35+0.17	0.39+0.34
Material	16	17	18	19	20
FFT+PCA	0.38+0.25	0.32+0.17	0.30+0.14	0.30+0.14	0.26+0.5
Features	1.11+1.55	0.50+0.48	0.39+0.17	0.53+0.47	0.31+0.15
Material	21	22	23	24	25
FFT+PCA	0.26+0.04	0.27+0.07	0.29+0.11	0.27+0.07	0.31+0.14
Features	0.27+0.09	0.31+0.17	0.41+0.27	0.34+0.19	0.53+0.44
Material	26	27	28	29	30
FFT+PCA	0.25+0.01	0.32+0.17	0.25+0.01	0.27+0.07	0.26+0.06
Features	0.25+0.01	0.48+0.43	0.27+0.07	0.32+0.18	0.29+0.12
Material	31	32	33	34	**Avg.**
FFT+PCA	0.25+0.01	0.29+0.13	0.30+0.12	0.26+0.06	**0.28+0.09**
Features	0.25+0.01	0.26+0.05	0.50+0.40	0.27+0.08	**0.40+0.28**

standard deviation) for the features commonly used in the state of the art close to the time of the vibration only case when using FFT-PCA. However, using the Fourier coefficients to characterise vibration generates a faster identification on average, 0.28 seconds slightly above one single iteration. Attending to the standard deviations of the identification times, one could say that, if combined with thermal features, to have good identification accuracies ($\mu + 3\sigma$) the multimodal approaches need to see data for around 0.5 seconds (two iterations) in the case of the Fourier coefficients, and 1 second (four iterations) for the vector $\mathbf{x}$. These values will be used as the baseline for our comparison with other learning approaches in the following section.

3.2.1. Comparison of the multimodal recognition against commonly used machine learning approaches

We compared the multimodal RBE approach with three machine learning techniques (kNN, ANN and SVM) measuring the identification accuracy for the 34 material dataset. We performed 10-fold cross validation using 10 minutes of data sequences for each of the four algorithms, finding the optimal parameters

Table 4: Summary of material identification results

	KNN	ANN	SVM	RBE
Vibration - Features	73.2%	54.2%	53.7%	81.7%
Vibration - FFT+PCA	98.0%	85.5%	92.0%	96.6%
Multimodal - Features	76.6%	82.9%	82.4%	93.8%
Multimodal - FFT+PCA	89.4%	94.6%	97.5%	98.6%

for the kNN, ANN and SVM as described in Section 3.1.1, and, as stated in Section 2.3, the datasets were whitened to have zero mean and identity covariance. The kNN algorithm achieved the optimal accuracy for $k = 5$ using the two tested multimodal combinations (FFT-PCA and feature vector combined with thermal features). The topology of the ANN were selected to have two hidden layers with 65 units on each layer for the Fourier coefficients combined with the thermal features, and 30 hidden units for the combined vibration-thermal feature vector. Finally, all the One-Vs-One SVM classifiers had RBF kernels with a parameter tuned through a grid search. The input data streams for all classifiers were chosen to have a length according to the same principle used in Section 3.1.1, i.e. three iterations long for the Fourier coefficients, and five for the vibration features (see Table 2). This was the same number of iterations (deadline) in which the RBE algorithm was required to successful identify the material.

Table 4 presents, in the last two rows, the classification accuracy results for the multimodal case, while, for comparison purposes, the first two rows summarise the results for the vibration only material identification presented in Section 3.1.1. Although the length of the perceived time series used as input for each row is different, they all follow our selected criteria of the average time taken by the RBE to identify the materials. This means that the amount of information fed into all the classification techniques is the same for each row when evaluating the performance. As the table shows, including thermal information improves the identification accuracy regardless of the technique used and the way of characterising the vibration (Fourier coefficients or vibration features). This points out that the thermal feature vector we propose is an excellent measure of the thermal properties of the material. The only exception is the kNN classifier, as the performance decreases from 98% to 89.4% when the Fourier coefficients are used in combination with the thermal features. Moreover, the performance increase in the case of the vibration features for the kNN classifier (from 73.2% to 76.6%) is relatively low compared with the other approaches evaluated. This is due to the data whitening preprocessing step necessary to perform the combination, as the distances between clusters are changed during the whitening. Another conclusion that can be drawn from Table 4, something already mentioned in Section 3.1.1, is that using the full spectral range always results in better classification. This points out the difficultly of identifying materials using the vibration features found in the literature, as they typically

Table 5: Iterative estimation of material 17 using vibration only approach with FFT-PCA.

Iteration	0	1	2	3	4	5	6	7	8
$P(m_{17})$	0.03	0.03	0.01	0.01	0.05	0.61	0.98	0.99	1.00
$1 - P(m_{17})$	0.97	0.97	0.99	0.99	0.95	0.39	0.02	0.01	0.00
$\max P(m_i), i \neq 17$	0.03	0.88	0.95	0.99	0.95	0.39	0.02	0.01	0.00

Table 6: Iterative estimation of material 17 using Multi-modal iterative estimation.

Iteration	0	1	2	3	4
$P(m_{17})$	0.03	0.98	0.99	0.99	1.00
$1 - P(m_{17})$	0.97	0.02	0.01	0.01	0.00
$\max P(m_i), i \neq 17$	0.03	0.01	0.01	0.01	0.00

extract only statistical measures of the vibration signal, completely ignoring the spectral components which actually characterise the interaction between the BioTAC fingertip and the material. All in all, the best identification results are obtained when the multimodal recursive Bayesian estimation algorithm is used (98.6%). This is the result of setting a classification deadline to this approach, yet when longer time sequences are used, the RBE algorithm always identifies the material correctly.

To further illustrate the improvement of the RBE algorithm when multimodal information is provided, Tables 5 and 6 show the evolution of the posterior probabilities for the recognition of material 17 for the same test sequence using only vibration and multimodal approaches. In both sequences the FFT-PCA was used to characterise the vibration and the initial prior probabilities (i.e. iteration 0) were uninformative. Both tables show the probability of the correct material $P(m_{17})$, the combined probability of all the wrong materials $1 - P(m_{17})$, and the highest probability of a wrong material $\max P(m_i), i \neq 17$. While Table 5 shows that it takes 5 iterations to classify material 17 with a probability of 0.61 using vibration only approach, Table 6 shows the correct material is identified in the first iteration with a probability of 0.98. As it can be clearly seen, including thermal information makes the identification correct from the first iteration on, reducing in this case the identification time from 1.25 seconds to 0.25 seconds i.e. from 5 to 1 iteration. It is worth mentioning that once the algorithm has detected the correct material with probability 1.00 it maintains that probability value as the prior for any other material in the next iteration will be zero.

4. Conclusions and Future Work

This paper presents a multimodal tactile based continuous material identification approach. While state of the art approaches to material identification mostly rely on vibration information, we show that including thermal features enhances the recognition accuracy and the time required for successful identification. Moreover, standard tactile identification techniques typically require a sequence of at least one second to classify materials. Using recursive Bayesian

estimation a system endowed with tactile sensors can identify materials in an average of 0.28 seconds with a very small deviation from that time lapse. This high detection speed is again due to the use of multimodal information. Hence thermal sensing enables a faster identification than a vibration only approach. We also eliminated the need of several long exploratory movements for surface identification found in the related literature.

Including thermal information brings the average material identification time very close to the used window size. Faster identification could be achieved by reducing the size of the window, yet the selection of $\Delta t = 0.25$ was empirically found to provide an excellent time-recognition trade-off. We showed that in comparison with other techniques used in the existing literature, our approach performs better than kNN, ANN and SVM when using the same data set and, despite needing more evidence to achieve perfect identification accuracy, a 1 second exploratory movement was enough to correctly classify 34 materials with high accuracy. Therefore, shorter time will possibly imply worse accuracy. Additionally, the system can be easily extended to other sensing modalities as shown in Section 2.2, and it is scalable to a greater number of materials as the computational burden of the RBE is low compared with the computation of the FFT. Our aim in the future is to include other sensing modalities, specifically vision, to enhance material recognition, for instance, generating visual texture based priors.

[1] A. Amedi, K. von Kriegstein, N. van Atteveldt, Beauchamp, M. Naumer, Functional imaging of human crossmodal identification and object recognition, Experimental Brain Research 166 (3-4) (2005) 559–571.

[2] H. Yousef, M. Boukallel, K. Althoefer, Tactile sensing for dexterous in-hand manipulation in roboticsa review, Sensors and Actuators A: physical 167 (2) (2011) 171–187.

[3] Z. Kappassov, J.-A. Corrales, V. Perdereau, Tactile sensing in dexterous robot hands, Robotics and Autonomous Systems 74 (2015) 195–220.

[4] N. Gorges, S. Navarro, D. Göger, H.Wörn, Haptic object recognition using passive joints and haptic key features, in: IEEE International Conference on Robotics and Automation, 2010, pp. 2349–2355.

[5] A. Scheneider, J. Sturm, C. Stachniss, M. Reisert, H.Burkhardt, W. Burkhardt, Object identification with tactile sensors using bag-of-features, in: IEEE/RSJ International Conference on Intelligent Robots and Systems, 2009, pp. 243–248.

[6] S. Chitta, J. Sturm, M. Piccoli, W. Burgard, Tactile sensing for mobile manipulation, IEEE Transactions on Robotics 3 (27) (2011) 558–568.

[7] J. Romano, K. Hsiao, G. Niemeyer, S. Chitta, K. Kuchenbecker, Human-inspired robotic grasp control with tactile sensing, IEEE Transactions on Robotics 6 (27) (2011) 1067–1079.

[8] K. Hang, M. Li, J. Stork, Y. Bekiroglu, A. Billard, D. Kragic, Hierarchical fingertip space for synthesizing adaptable fingertip grasps, in: ICRA Workshop: Autonomous Grasping and Manipulation: An Open Challenge, 2014.

[9] A. Gómez-Eguíluz, I. Rañó, S. Coleman, T. McGinnity, Reliable object handover through tactile force sensing and effort control in the shadow robot hand, in: IEEE International Conference on Robotics and Automation, 2017.

[10] J. Edwards, J. Lawry, C. Melhuish, Extracting textural features form tactile sensors, Bioinspiration & Biomimetics 3 (3) (2008) 035002.

[11] F. D. Boissieu, C. Godin, B. Guilhamat, D. David, C. Serviere, D. Baudois, Tactile texture recognition with a 3-axial force mems integrated artificial finger, in: Robotics: Science and Systems, 2009.

[12] J. Sinapov, V. Sukhoy, R. Sahai, A. Stoytchev, Vibrotactile recognition and categorization of surface textures by a humanoid robot, IEEE Transactions on Robotics 3 (27) (2011) 488–497.

[13] N. Jamali, C. Sammut, Majority voting: material classification by tactile sensing using surface texture, IEEE Transactions on Robotics 3 (27) (2011) 508–521.

[14] S. Decherchi, P. Gastaldo, R. Dahiya, M. Valle, R. Zunino, Tactile-data classification of contact materials using computational intelligence, IEEE Transactions on Robotics 3 (27) (2011) 635–639.

[15] C. Oddo, M. Controzzi, L. Beccai, C. Cipriani, M. Carrozza, Roughness encoding for discrimiation of surfaces in artificial active-touch, IEEE Transactions on Robotics 3 (27) (2011) 522–533.

[16] P. Dallaire, P. Giguère, D. Émond, B. Chaib-draa, Autonomous tactile perception:a combined improved sensing and bayesian nonparametric approach, Robotics and Autonomous Systems (62) (2014) 422–435.

[17] D. Chathuranga, V. Ho, S. Hirai, Investigation of a biomimetic fingertip's ability to discriminate fabrics based on surface textures, in: IEEE/ASME International Conference on Advance Intelligent Mechatronics: Mechatronics for Human Wellbeing, 2013, pp. 1667–1674.

[18] D. Chathuranga, Z. Wang, Y. Noh, T. Nanayakkara, S. Hirai, Robust real time material classification algorithm using soft three axis tactile sensor: Evaluation of the algorithm, in: IEEE/RSJ International Conference on Intelligent Robots and Systems, 2015.

[19] S. Baishya, B. Bäuml, Robust material classification with a tactile skin using deep learning, in: IEEE/RSJ International Conference on Intelligent Robots and Systems, 2016.

[20] T. Bhattacharjee, J. Wade, C. Kemp, Material recognition from heat transfer given varying initial conditions and short-duration contact, in: Robotics: Science and Systems, 2015.

[21] J. Fishel, G. Loeb, Bayesian exploration for intelligent identification of textures, Frontiers in Neurorobotics (6) (2012) 1–20.

[22] D. Xu, G. Loeb, J. Fishel, Tactile identification of objects using bayesian exploration, in: IEEE International Conference on Robotics and Automation, 2013, pp. 3056–3061.

[23] Z. Su, J. Fishel, T. Yamamoto, G. Loeb, Use of tactile feedback to control exploratory movements to characterize object compliance, Frontiers in Neurorobotics (6) (2012) 1–9.

[24] E. Kerr, T. McGinnity, S. Coleman, Material classification based on thermal and surface texture properties evaluated against human performance, in: International Conference on Control, Automation, Robotics & Vision, 2014.

[25] J. Hoelscher, J. Peters, T. Hermans, Evaluation of tactile feature extraction for interactive object recognition, in: IEEE-RAS 15th International Conference on Humanoid Robots, 2015.

[26] A. Gómez-Eguíluz, I. Raño, S. Coleman, T. McGinnity, Continuous material identification through tactile sensing, 2016, international Joint Conference on Neural Networks.

[27] A. Gómez-Eguíluz, I. Raño, S. Coleman, T. McGinnity, A multi-modal approach to continuous material identification through tactile sensing, in: IEEE/RSJ International Conference on Intelligent Robots and Systems, 2016.

[28] J. Fishel, G. Lin, G. Loeb, Syntouch llc biotac product manual, v. 16, Tech. rep. (2013).