Robotics and Autonomous Systems 62 (2014) 422–435

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

journal homepage: www.elsevier.com/locate/robot

Autonomous tactile perception: A combined improved sensing and Bayesian nonparametric approach

Patrick Dallaire, Philippe Giguère *, Daniel Émond, Brahim Chaib-draa

Department of Computer Science and Software Engineering, Pavillion Adrien-Pouliot, 1065 avenue de la Médecine, Laval University, Quebec City, QC G1V 0A6, Canada

HIGHLIGHTS

- We present a new robust tactile sensor aimed at surface identification.
- We evaluate the performances of 7 features for surface identification.
- We use a Pitman–Yor process to autonomously learn a perception model.
- The model recognized all surfaces perfectly without providing the number of surfaces.
- Identification success rate on unseen data is over 90%.

ARTICLE INFO

Article history:
Received 15 August 2012
Received in revised form
30 August 2013
Accepted 26 November 2013
Available online 2 January 2014

Keywords:
Tactile sensing
Surface and texture identification
Bayesian nonparametric methods
Accelerometer
Machine learning

ABSTRACT

In recent years, autonomous robots have increasingly been deployed in unknown environments and required to manipulate or categorize unknown objects. In order to cope with these unfamiliar situations, improvements must be made both in sensing technologies and in the capability to autonomously train perception models. In this paper, we explore this problem in the context of tactile surface identification and categorization. Using a highly-discriminant tactile probe based upon large bandwidth, triple axis accelerometer that is sensitive to surface texture and material properties, we demonstrate that unsupervised learning for surface identification with this tactile probe is feasible. To this end, we derived a Bayesian nonparametric approach based on Pitman–Yor processes to model power-law distributions, an extension of our previous work using Dirichlet processes Dallaire et al. (2011). When tested against a large collection of surfaces and without providing the actual number of surfaces, the tactile probe combined with our proposed approach demonstrated near-perfect recognition in many cases and achieved perfect recognition given the right conditions. We consider that our combined improvements demonstrate the feasibility of effective autonomous tactile perception systems.

1. Introduction

In the last 20 years, autonomous exploration of unknown environments or objects by robots has been extensively studied. In general, these studies [1,2] used range sensors (sonar, or laser) or cameras to gather information about the environment. However, it is not always possible to use these sensing modalities, in which case one should consider exploring the environment more directly, possibly via tactile sensing. This strategy is seen in animals, where tactile sensing is a fundamental mechanism allowing them to navigate their environment blindly or perform surface material categorization [3].

Over the years, many types of tactile probes have been developed to mimic the sense of touch, whether for surface texture recognition [4] or surface feature recognition [5–8]. As they are, by their very nature, immune to numerous problems that plague vision-based sensing, such as illumination changes, occlusion, or the high-dimensionality output (millions of pixels) of cameras, tactile-based systems have the potential to offer much more robust methods for surface recognition. Consequently, tactile perception could be used in challenging environments where vision systems are difficult to operate. In outdoor setting for example, large and varying illumination changes are present, complicating the use of computer vision.

One of the crucial aspects involved in the development of any artificial perception system is that it must be trained before it can perform recognition. With access to a training database where sensor data samples have been hand-labeled, any standard supervised learning techniques can be used. Ideally, we are looking for minimal human supervision during this learning process, either

* Corresponding author. Tel.: +1 418 656 2131x3076; fax: +1 418 656 2324.
E-mail addresses: dallaire@damas.ift.ulaval.ca (P. Dallaire), philippe.giguere@ift.ulaval.ca (P. Giguère), daniel.emond.2@ulaval.ca (D. Émond), Brahim.Chaib-Draa@ift.ulaval.ca (B. Chaib-draa).

0921-8890/$ – see front matter
http://dx.doi.org/10.1016/j.robot.2013.11.011

to reduce labor costs or in order to develop systems exhibiting a higher level of autonomy. This implies that useful information required for supervised learning, such as data labels and number of classes, will not be available. Consequently, training a perception system under such conditions is significantly more challenging. On the other hand, unsupervised learning approaches allow systems to automatically adapt themselves based on experience, without parameter tuning or data labeling.

A further step towards autonomy is by improving learning flexibility, using infinitely large parameter space. This way, the complexity of the model scales appropriately with the amount of training data. This is exactly what Bayesian nonparametric learning methods have to offer, ensuring a considerable degree of flexibility in statistical modeling when compared to their parametric alternatives. In our previous work [9], we thereby proposed a model based on Dirichlet processes to perform autonomous surface recognition. However, in the context of outdoor robot navigation and for many other applications involving natural phenomena, models yielding power-law behavior, unlike Dirichlet processes, are more appropriate to represent the data.

In this paper, we present a more general framework for autonomous surface recognition based on Pitman–Yor processes, a model generalizing Dirichlet processes and yielding power-law behavior. To this end, we first discuss in Section 2 previous works on tactile sensing systems and Bayesian nonparametrics. Section 3 describes the tactile probe, the data gathering process and the test sets used for the experiments. In Section 4, we demonstrate the surface identification capability of our tactile probe and evaluate the chosen features by ranking them according to their supervised learning performance. Section 5 presents the method used to autonomously learn to differentiate surfaces, without the need to specify data labels or the number of surfaces in the training set. Finally, Section 7 concludes and presents future research directions.

2. Previous work

2.1. Tactile sensing

For human beings, tactile sensory modalities in the fingertips are used to capture multiple object properties such as texture, roughness, spatial features, compliance or friction. These tactile receptors are capable of detecting vibrations as the finger slides [4], making it possible to discriminate between surface textures [10,11], including estimating the spatial frequency of the texture [12]. On the other hand, detecting surface features such as edges or corners [5,8], temperature [13] or compliance [14] does not necessarily require such dragging motion.

Some researchers have focused their attention on artificial skin, particularly on the concept of having flexible and modular components [15–18]. However, the general focus of these skin sensors is more about the identification of contact point locations and pressure forces [17], or the integration of multiple sensing modalities (temperature, acceleration, proximity and vibration) [18], as opposed to our goal of surface identification.

Tactile sensing in robotics is not just confined to skin-covered finger devices. Indeed, tactile sensing technologies not embedded in fingers have been proposed over the years, in the likes of artificial whiskers [19–22], array of whiskers [23,24], or artificial antennas [25]. These devices can estimate surface profile, perform rudimentary object recognition or provide distance estimation.

2.2. Supervised learning with tactile sensing

The targeted application in this paper is surface type identification via tactile sensing. By its direct contact with an object, tactile sensing has the capability to gather information not captured by visual sensor, thereby improving object recognition. One commonly-used physical characteristic employed in surface identification is its texture. By rubbing a tactile probe on a textured surface, vibrations are generated with a temporal periodicity connected to the spatial periodicity of the texture [12].

In the canonical scenario, a tactile probe sensitive to vibrations is rubbed against the investigated surface. Several features are then extracted from the sensor's signal, in the hope of reducing the dimensionality of the problem without significant loss of information. Finally, a classification algorithm is used to recognize the surface, based on the extracted features. Thus, most tactile recognition systems can be categorized based on the sensing technology, features extracted from the sensor signals, or the type of classifier employed.

Many examples of these canonical descriptions are present in the literature, particularly in the context of supervised learning. Hipp et al. [26], for example, presented results regarding texture classification for a system of actuated whiskers. In their case, a magnetometer captured the whisker's vibrations, as it was made of a metal capable of modifying the local magnetic field. The sensor signal was then band-passed between 30 and 150 Hz, and its power spectrum was computed as features. The training consisted in fitting multidimensional Gaussian density estimators on the spectrum. Using a maximum likelihood classification on the testing set, they achieved a success rate of 39% for eight different grades of sandpaper. Fend et al. [27], on the other hand, used a microphone to record the vibrations induced in genuine rat whiskers over 11 surfaces. The features used were the combined and smoothed power spectra of individual sweeps, to generate average power spectra signatures which are more stable. They used an instance-based learning approach to determine how many texture signatures were discernible, using a Euclidean distance between signatures as metric. Overall, Fend and his colleagues concluded that texture identification could be improved by using all whiskers at the same time and by increasing the number of sweeps. However, quantitative results are not readily available from the paper.

Surface classification results can also be improved by focusing on the particular machine learning techniques applied during the classification stage. Jamali and Sammut [28], for example, employed a majority voting scheme to improve the accuracy of surface identification. Their majority voting approach mimicked the strategy used by humans which consists in trying several explorations of the material's surface before reaching a decision. They found that majority voting greatly improves the robustness of classifiers such as naive Bayes, decision trees, naive Bayes tree (NBTrees), boosting on NBTrees and decisions trees. By doing so, they showed that they can distinguish between nine different surfaces, such as carpets, vinyl flooring, tiles, sponge, wood and polyvinyl-chloride (PVC) woven mesh with an accuracy of $95\% \pm 4\%$ on unseen test data, over a test set of 8 surfaces. They also studied the impact of the extracted features on the classification results. In their finger, they used 4 polyvinylidene fluoride (PVDF) sensors to capture vibrations induced during the probing process. Most of their features consisted in the frequencies of the peaks in the amplitude spectrum of the PVDF signals, keeping the n highest peaks of each of the 4 spectra. They showed that larger values of n, thus larger feature vectors, improve classification results. They noted, however, that this increase in performance was limited when $n > 10$.

Fishel et al. [29] adopted an *active classification* approach to the problem of surface identification based on artificial fingers sliding on surfaces. In a way, their work is an extension of the majority voting algorithm used by Jamali and Sammut [28], but with a more informed search strategy. Indeed, when an active

classifier has not reached a sufficient confidence level, then further actions should be taken. These actions are dynamically selected so as to help discriminate between the currently competing hypothesis, as opposed to constantly repeating the same action. To make the problem of identifying an appropriate set of actions tractable, a number of simplifications were made. First, they selected three distinct properties from psychophysics literature that could be used for surface identification (traction, roughness and fineness) and found a feature extraction method to map the measured dragging forces and vibrations of the finger into these 3 properties. Second, they assumed that these 3 features were independent (Naive Bayes). Third, they assumed that each feature was normally distributed, which meant that each surface could be represented internally as a parameter vector of 3 pairs of mean and standard deviation, one for each feature. During a training period, these parameters were evaluated, and the efficiency of all stroking strategies in distinguishing between pairs of known surfaces was established. This resulted in what they call a *Bayesian exploration algorithm*, since the next stroking action (up to 10) of the active classification process was based on Bayesian formulation. Training was performed on a set of 117 surfaces. On the testing sub-set of 16 surfaces, they were able to select the correct match from the complete database of 117 surfaces 95.4% of the time, using the above strategy. Considering the limited number of features (3), this suggests that the BioTac™ tactile probe provides exceptionally good information. Suggestions were made for unsupervised learning strategies, but not implemented.

2.3. Unsupervised learning associated with tactile sensing

Edwards et al. [30] presented one of the few examples of a system that performs some form of autonomous learning of surfaces. Their tactile device included a microphone to record tactile textual features probed by an artificial finger. Signal recordings were mapped to the frequency domain using Fast Fourier transforms. It achieved a 97.6% classification success rate using a *k*-Nearest Neighbors classifier on the data, after Principal Component Analysis (PCA) dimensionality reduction. Clustering of this PCA data showed that similar surfaces are grouped together with few classification errors, thus paving the way for unsupervised learning approaches. One important limitation of this result is that the clustering technique employed, *k*-means, requires that the number of clusters (i.e. surfaces) in the data set be known ahead of time.

Closer to us, Johnsson et al. [31] explored the concept of merging tactile texture information with hardness and performed some unsupervised learning with a derivative of Self-Organizing Map (SOM) neural network. The texture sensing performed by rubbing via a metallic edge against a probed surface while a microphone captured the generated vibrations for 1 s. Hardness was probed directly by measuring material deformation when subjected to a constant pressure. An SOM was trained with unlabeled data to reduce the dimensionality of the signal spectrum (2049 elements) down to a 15×15 matrix. Overall, they showed that combining these two sensing modalities improves the discrimination capabilities of the system, with specific neuron regions in the SOM firing for similar materials. This SOM approach has also been explored by Takamuku et al. [32], for a hand equipped with strain gauges and PVDF (polyvinylidene fluoride) films sensors, sampled at 1.6 KHz. Both squeezing and tapping behaviors were used to explore 7 objects. Again, these experiments showed that an SOM can specialize its neuron to differentiate between categories of objects, based on their physical properties. Specifically, nearby neurons in the SOM would fire for objects with similar physical characteristics. However, no formal classification results were provided and no attempts were made to identify the number of classes in the unlabeled data sets.

The closest experimentations to ours are from Sinapov et al. [33]. They present work about interactive surface recognition and categorization, using a "fingernail" vibrotactile sensor equipped with a triple axis accelerometer to sense vibrations. They used a 2D spectrotemporal histogram as a feature vector, constructed by discretizing the amplitude spectrum in 25 frequency intervals and for 5 different time periods. In the first set of experiments, they classified surfaces in data collected over 20 different common household surfaces using a *k*-Nearest Neighbors and a Support Vector Machine (SVM). They were able to demonstrate that using 5 different exploratory scratching behaviors yield better results (80% classification rate for SVM) than a single one (58% again for SVM), a conclusion similar to Fishel et al. [29]. In a second set of experiments, they performed hierarchical clustering using *spectral clustering* to produce a tree structure capable of differentiating certain categories of surfaces. The greatest number of clusters found was only 7 (out of 20) and produced a recognition rate of 58%.

A related problem concerns the automatic dimensionality reduction in tactile sensing, which has been explored by Kroemer et al. [34]. Their key idea was to explore an object with different sensing modalities (vision and tactile, for example) and use one of the modality (vision) to identify a low dimensionality mapping function for the other (tactile). This identification was done during a training phase, using two proposed variations of the Maximum Covariance Analysis (MCA) technique. Generally speaking, MCA tries to find a mapping which preserves features that are correlated between the sensing modalities for a given sensed object, while discarding those that are not. The underlying assumption for their project was that these correlations between vision and tactile sensing arose because of some physical characteristic of the probed object. The proposed modifications to the original MCA technique allowed them to exploit the inherently weak pairing of data between vision and tactile, since it was difficult to identify the exact trajectory of the tactile sensor in the picture of the object. The original (non-reduced) tactile feature space was in the form of bag-of-words extracted from the vibration signal cepstrum, the latter ensuring phase independence. Testing results over 17 different surface types indicated that the dimensionality reduction mapping allowed a 95.15% success rate with supervised learning, compared to 90.85% with standard PCA. Experiments were also conducted with unsupervised learning, indicating that their technique is also applicable to this paradigm. However, the number of clusters (17) needed to be provided in their experiments, something we want to entirely avoid with our approach.

2.4. Bayesian nonparametrics

The rapidly developing field of Bayesian nonparametrics provides new perspectives for developing fully autonomous systems capable of adaptation without external interventions. The Bayesian nonparametric approach consists in constructing probability distributions over infinite dimensional spaces of parameters, where only a finite set of them needs to be explicitly represented, leading to tractable inference. The best known Bayesian nonparametric model is surely the Dirichlet process, which has been extensively studied in the statistical community [35] and has proved to be very useful for density estimation [36].

More recently, Dirichlet process models have been applied to multiple problems related to robotics and autonomous systems. Joseph et al. [37] learned a model of battery time-to-death by representing the different cooling behaviors with an infinite mixture model. A hierarchical Dirichlet process was used by Fox et al. [38] to learn switching linear dynamical models with an unknown number of modes for describing complex dynamical phenomena. The important problem of *imitation learning* has

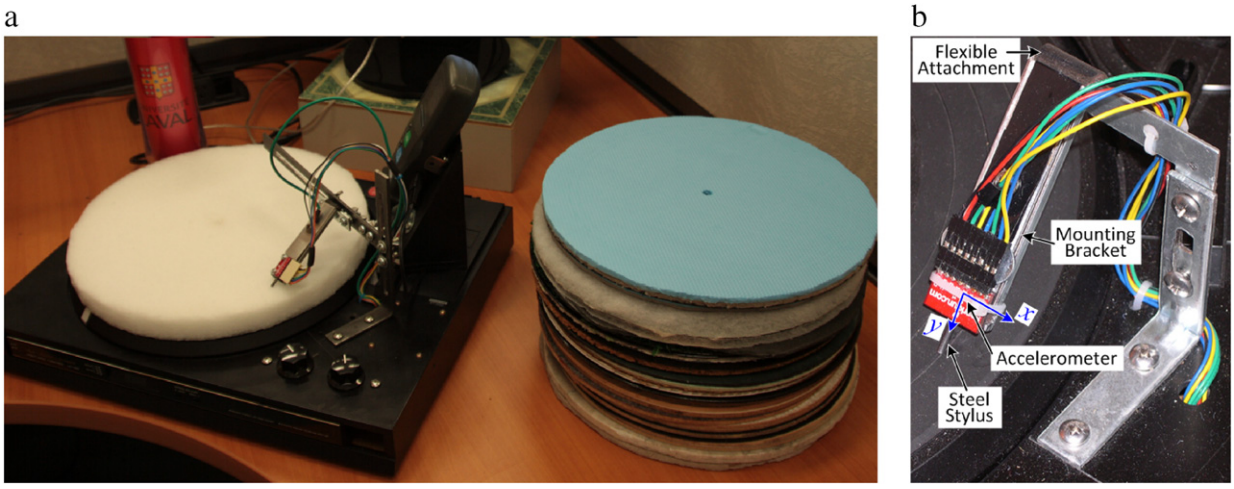

Fig. 1. (a) Turntable used to collect tactile probe data, with the tachometer measuring rotation velocity. The stack of disks, on the right, constitutes the surface test sets. (b) Close-up view of the tactile probe, showing the steel pin and the triple axis accelerometer (in red) attached near the tip. The accelerometer axes orientation are shown in blue, with the unseen z axis pointing away from the reader.

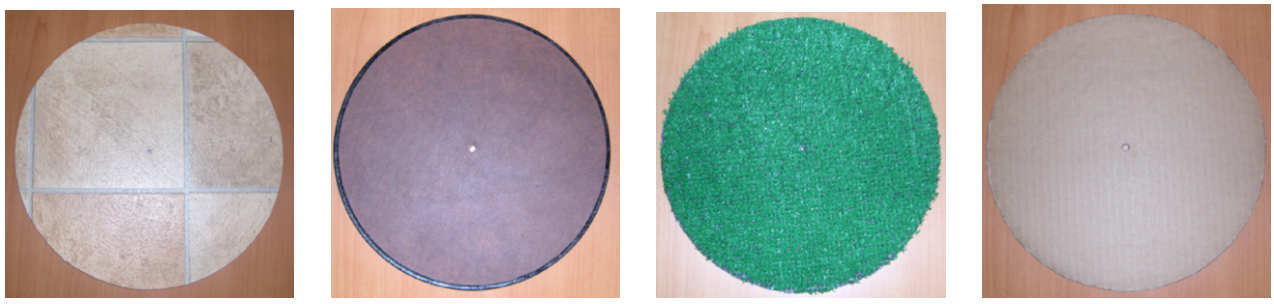

(a) Linoleum. (b) Leatherette. (c) Grass carpet. (d) Cardboard.

Fig. 2. Picture of 4 disks used in the surface identification experiments. The complete test set comprises 28 disks.

also been tackled from a nonparametric Bayesian perspective by some researchers that used hierarchical Dirichlet processes [39] and infinite Gaussian mixture models [40,41] in particular. The problem of automatic classification of chemical sensor data from autonomous underwater vehicles is a task related to the one we consider in this paper. In their work, Jakuba et al. [42] adopted a Bayesian nonparametric learning approach that was able to autonomously identify physically relevant clusters in the data. However, they opted for a variational approximation of a Dirichlet Process Mixture Model based on Gaussian distributions to quickly obtain estimations.

3. Data set gathering with a tactile sensor

3.1. Tactile sensor description

The tactile sensor used to gather the data set, depicted in Fig. 1(b), is an improved version of an earlier prototype [43,44]. This newer version, fully described in Dallaire et al. [9], probes a target surface via a 2.38 mm diameter standard steel stylus. A triple axis digital Micro Electro-Mechanical Systems (MEMS) accelerometer (ADXL345 from Analog Devices) is placed near the tip of the stylus to ensure maximum vibration capture. It is configured to collect data at a rate of 800 Hz. A built-in digital low-pass filter in the accelerometer ensures that there is limited aliasing, as it operates at 3200 Hz. Moreover, the data resolution is set to 3.9 milli-g, with a range of ± 5 g.

Generally speaking, our sensing approach is nearly identical to the one proposed by Sinapov et al. [33], as they also use an ADXL345 sensor mounted on a hard plastic nail to sense surface textures. However, their device is fixed to a robotic finger, which allows us to actively probe surfaces by commanding the robot hand with different exploratory trajectories. In our case, we instead prefer the strategy of moving the probed surface to simplify the design and focus on discrimination capabilities.

Accordingly, our tactile probe is installed on a standard turntable (record player), as shown in Fig. 1(a), at a 45° angle relative to the probed surface. The combined weight of the stylus, accelerometer and mounting bracket is approximately 16 g. The mounting bracket is fixed to the turntable with a flexible attachment made of thick leather. This allows the probe to rotate freely in a plane perpendicular to the X axis, while granting small rotations or motion in other directions. This way, the accelerometer can capture accelerations in all three axes. Using a tunable turntable also has the additional advantage that regulated speeds can be selected (33 or 45 revolutions per minute). However, a unique surface velocity has been used to collect the data for the experiments.

3.2. Data set gathering

The test material comprises various type of surfaces placed onto 12-inch disks, so they can be readily used with the turntable. A small subset of test disks is shown in Fig. 2. During data collection, the angular velocity of the surface is maintained at 45.0 rpm, as measured by a digital tachometer (TC811B, Reliability Direct). The distance between the center of the turntable and the tip is 125.0 mm, giving a constant surface speed of 58.9 cm/s at the stylus' tip.

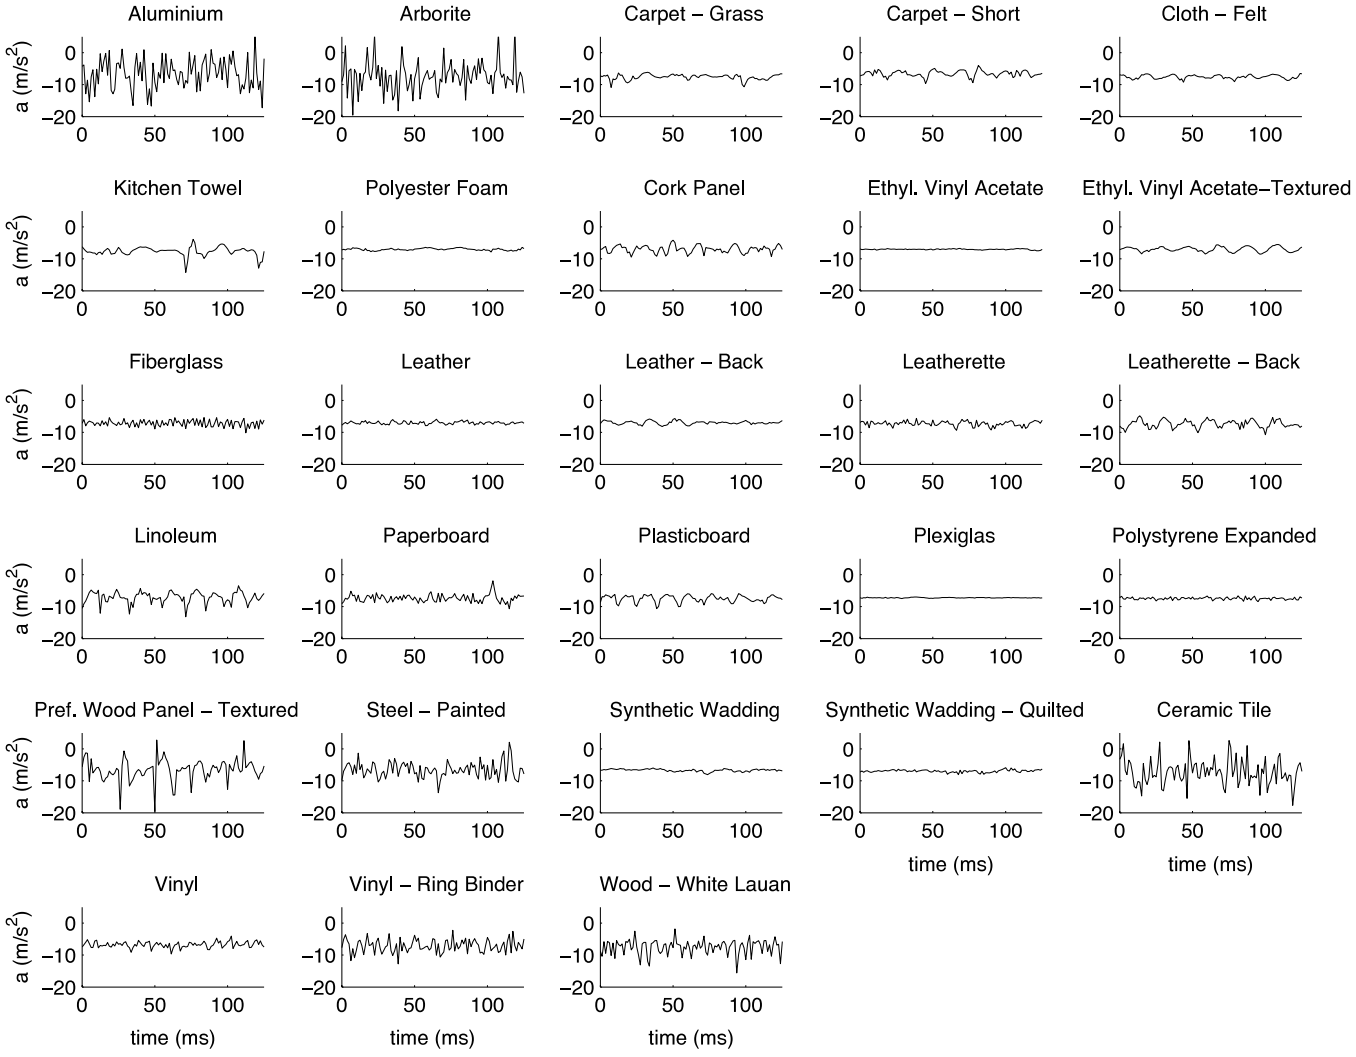

Fig. 3. A single time window of 100 samples of acceleration signals in the Y direction, taken for each of the 28 test surfaces. As would be expected, hard surfaces (aluminum, arborite, tile) produce acceleration signals with rapid variations in time, i.e. with significant energy in the high end of the spectrum, as opposed to soft surfaces (synthetic wadding, polyester foam, carpet). These differences indicate how our tactile probe goes beyond simply sensing surface texture, but rather gathers information about the surface's mechanical properties. Indeed, two smooth surfaces (aluminum and Plexiglas) have completely different signals.

We use 28 surfaces to realistically test the capability of our approach on learning tactile perceptions. In particular, different materials and textures have been selected beforehand in order to cover a wide range of mechanical properties. For example, we use hard surfaces (ceramic tile, aluminum, steel) and soft surfaces (artificial wadding, cotton fabric). Other surfaces are relatively smooth (Plexiglas), some with light texture (wood panel, fiberglass) while others are heavily textured (kitchen towel and grass carpet). As well, we try to cover natural materials (leather, cork panel, paper board) and synthetic materials (leatherette, soft vinyl cover).

The probe's position along the radius of the disk is adjustable via a sliding mechanism. The probe support has an adjustable height to maintain its angle at 45° with respect to the surface, regardless of the thickness of the disks. Combined with the controlled speed of the turntable, this ensures that the data gathering process is similar between disks. Two data sets of raw data, $\mathbf{S}_A$ and $\mathbf{S}_B$, have been collected at the same rotation speed, but at a one-week interval. This was done to ensure that no artificial bias was introduced due to the manipulations when changing disks or changes in the probe's response. Moreover, wear appeared on some surfaces during the first data gathering session, as we collected up to 20 min of recording per surface. This wear was not present at the very beginning of $\mathbf{S}_A$ and appears gradually as time progresses. On the other hand, the surface properties for $\mathbf{S}_B$ are slightly different for the worn surfaces. We believe that this second dataset replicated an exploration strategy where a robot periodically revisited surfaces to ensure that it has a more complete perception of it. A similar approach was taken by Jamali and Sammut [28], where they collected 4 data sets independently. Fig. 3 shows some of the tactile probe's signals for every test surface. In total, 20 min of recording per surface are available in $\mathbf{S}_A$ and 1 min in $\mathbf{S}_B$.

3.3. Selected signal features

We use a set of features that is extracted significantly more from the temporal domain than the spectral domain. This is contrary to the majority of approaches [26–28,30,31,34], that used spectral features extracted from the vibrations. A notable exception was Sinapov et al. [33], who used a spectrotemporal histogram that could capture the temporal evolution of spectrum. As we have shown previously in [44], the amplitude spectrum loses significant information about temporal patterns like spikes, as the phase information is flushed out. We also treat each of the 3 axes

independently, contrary to [33] who combined them together to obtain the magnitude of vibration.

The raw data obtained from the measurements of all three accelerometers is divided into 3-by-W matrices $\mathbf{a}_t$, with W being the width of the time window. The following seven features, adapted from terrain identification work by Weiss et al. [45], are extracted from each row of $\mathbf{a}_t$:

- f_{var}: variance,
- f_{skew}: skewness,
- f_{kurt}: kurtosis,
- f_{m5}: fifth moment,
- f_{sum}: sum of the variation over time:

$$\sum_{t=1}^{W-1}(|a(t) - a(t+1)|)$$

- f_{cros}: number of times 20 uniformly separated thresholds are crossed, and
- f_{hf}: sum of higher half of amplitude spectrum.

Note that these features are the same as those employed in our previous works [43,44], except for one of them. We no longer use the mean values of $\mathbf{a}_t$ to avoid issues with an accelerometer drift or probe angle impacting the classification.

Some explanations can be found for these features f. The variance f_{var} of the signal is a good indicator of the amount of vertical motion experienced by the probe, which is large for uneven surfaces. This kind of feature is correlated to the power of a signal, which was determined in Fishel et al. [29] as a representative of the roughness of textures. Hosoda et al. [46] have also used the variance of the signal from PDVF strips embedded in skins to distinguish between surfaces. The skewness f_{skew} of the signal helps identify surfaces with asymmetric distribution of acceleration, typical for surfaces with regular but infrequent asperities such as cracks or sharp bumps. The sign of the skewness plays an important role too, as some surfaces tend to have a distribution more skewed to the top than the bottom: compare *plasticboard* with *leatherette—back* from Fig. 3 for example. The sum of high-frequency components f_{hf} helps differentiate between hard and soft surfaces, since rubbery surfaces cannot sustain high-frequency vibrations in general. It has some correlation with the spectral centroid feature from Fishel et al. [29] and can also be seen as combining all the upper frequency bins of the spectral features of [26–28,30,31,34].

After extraction from the data set, each feature is normalized with a factor g_i so as to have a standard deviation of 1. When all three axes are used, i.e. using the rows $\mathbf{a}_{1t}$, $\mathbf{a}_{2t}$, and $\mathbf{a}_{3t}$, it produces a 21-dimensional feature vector $\mathbf{x}_i$ per time-window $\mathbf{a}_t$:

$$\mathbf{x}_i = [g_1 f_{\text{var}}(\mathbf{a}_{1t}), g_2 f_{skew}(\mathbf{a}_{1t}), \ldots, g_{21} f_{hifreq}(\mathbf{a}_{3t})]. \quad (1)$$

From observations, we noted that some features f_i had non-Gaussian distributions. When using an appropriate classifiers, such as a Support Vector Machine (see Section 4), non-Gaussian distributions can be easily handled; thus no actions were taken for the supervised learning experiments. However, subsequent feature linearization steps, described in Section 5.3, were required for the unsupervised learning experiments to respect the model assumptions.

4. Supervised classification

4.1. Classifier: support vector machine

We quantified the surface identification performance of our tactile probe by performing supervised learning on combined data from $\mathbf{S}_A$ and $\mathbf{S}_B$. Only the first minute of data was used in each data set, for a total of 2 min where features were extracted using the procedure described in Section 3.3 with time-window size $W = 800$, corresponding to 1 s of data at 800 Hz. The resulting data set of (normalized) features contains 3360 samples and is denoted by $\mathbf{X}$. Classification was performed using Support Vector Machines (SVM) from the LIBSVM library [47]. This is in contrast with our previous results [43,44], where classification was performed with an artificial neural network. We selected a Radial Basis Function (RBF) kernel:

$$K_{kernel}(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2), \quad \gamma > 0,$$

with parameter γ determining the range (radius) of the RBF. The parameter C, describing the weight of the error penalty term in the standard SVM formulation, acts as a regularization term. To avoid overfitting the training data, the values $\gamma = 0.0186$ and $C = 16$ of were found using stratified 10-fold cross-validation. The class attributed to an unknown sample was the one having the highest probability, as evaluated by the SVM. The classification success rate for a given test set was then computed as the number of correct classifications divided by the number of samples in the test set.

4.2. Classification results

The data set $\mathbf{X}$ was divided into two equal sets, with the odd samples placed in $\mathbf{X}_{train}$ and the even samples placed in $\mathbf{X}_{test}$. This ensured that samples in both sets were spread equally over time and balanced between signals from $\mathbf{S}_A$ and $\mathbf{S}_B$: some features f_i extracted from soft surfaces varied over time, as these surfaces experienced some wear during data collection. Using $\mathbf{X}_{train}$ to train the SVM, the 28 surfaces in $\mathbf{X}_{test}$ were almost perfectly identified, giving a success rate of 99.95%. In Fig. 4, we have the confusion matrix corresponding to this test showing that only 1 surface was misclassified.

The previous results indicate that the features extracted using a time-window size $W = 800$ convey enough information to the SVM so that only a few examples per surface are required to achieve near-perfect recognition. However, it is known that some features are more informative than others and have significative impacts on classification performances. As these features consist in our initial step towards dimensionality reduction, we performed a set of experiments, presented in Section 4.3, where each feature is used alone to accomplish the classification task. These experiments aim at measuring the quality of each feature to help determine whether a feature is relevant or not in the context of surface identifications.

4.3. Analyzing our features via feature ranking

Feature extraction is a crucial step in classification systems [48]. In tactile sensing, many different features have been proposed, but little work has been done on clearly identifying which of them carry more information. A possible exception is the work of Fishel et al. [29], where a very small set (3) of features were carefully selected and analyzed. Also, Jamali and Sammut [28] demonstrated that using a greater number of features (peaks in the frequency spectrum) does improve the classification results, although results more or less plateau past 10 of those features. To some extent, the dimensionality reduction approach proposed by Kroemer et al. [34] also constitutes a feature selection.

In order to better understand the importance of each of our features f, we ranked them using a *wrapper* approach [49]. In this approach, a feature is ranked based on the classification success rate it yields when used alone. Fig. 5 shows the ranking when a single feature is extracted from all three axes (yielding a size 3 feature vector), while Fig. 6 shows the ranking when features are extracted from one axis at a time (size 1 feature vector). We can see that features f_{var}, f_{cros} and f_{hf} rank consistently high, while feature

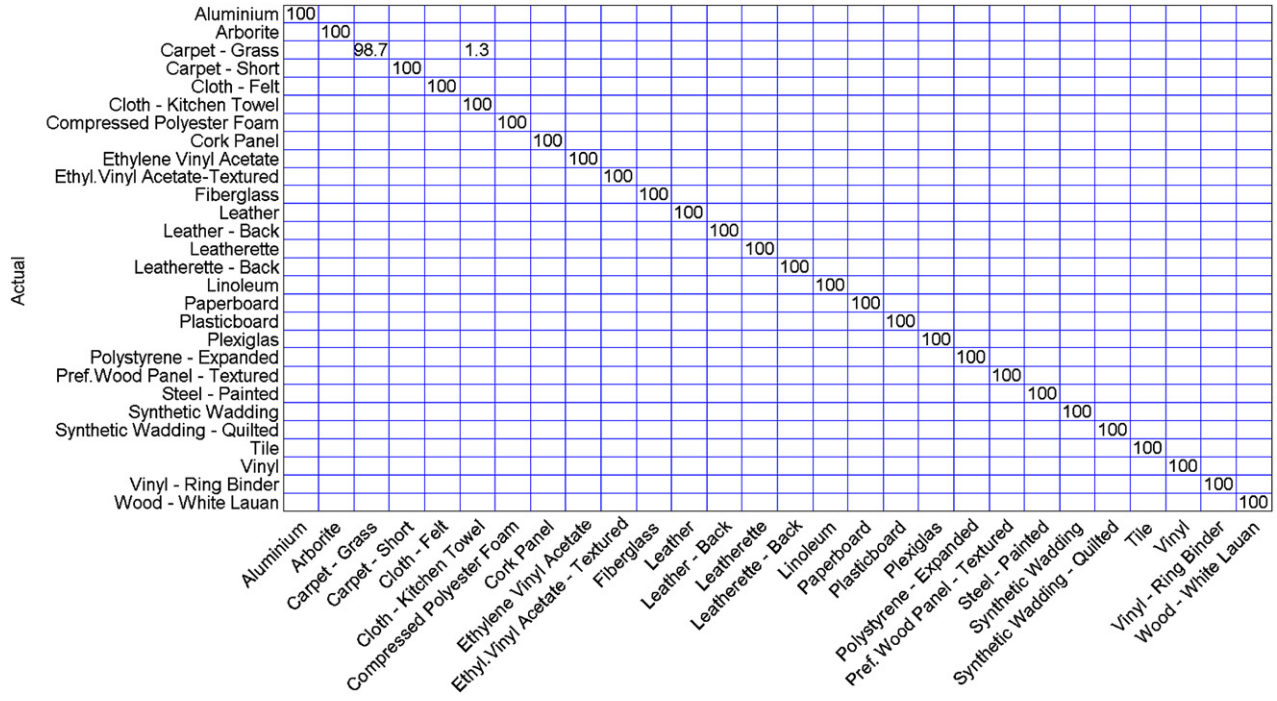

Fig. 4. Confusion matrix for the test set $\mathbf{X}_{test}$ using time-window size of $W = 800$ samples.

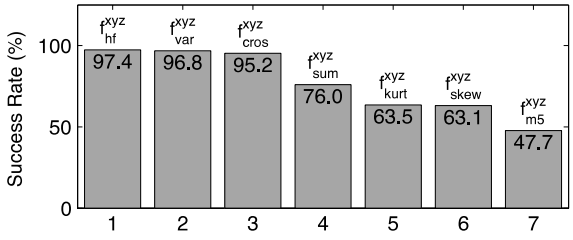

Fig. 5. Ranking of features extracted on all three axes. Values inside the bars are success rate percentages.

f_{sum} ranks in the middle. The richness of information related to f_{var} was also noticed by Fishel et al. [29].

4.4. Link with unsupervised learning

The results in this section show that a supervised learning algorithm was capable of identifying the different surfaces present in a test set. In doing so, we wanted to demonstrate that this data set is, in fact, *learnable*, i.e. that the features f extracted from the tactile probe signals are sufficient for classification purposes. As the unsupervised learning problem is considerably more difficult, this *learnable* condition is necessary but not sufficient to ensure success. In the following section, however, we will show that we were able to perform this unsupervised learning with our approach.

5. Autonomous learning using a Bayesian nonparametric learning approach

One of our long-term objectives is to achieve autonomous learning of low-level perception models for mobile agents by relying on unsupervised learning techniques. This would allow deploying robots in unknown environment to autonomously perform terrain recognition, construct maps and localize itself in the environment. A number of unsupervised learning techniques have been proposed to achieve this task, such as k-means or Normalized-Cut [50]; however a common drawback is the need to provide the number of clusters present in a data set, which is an unknown value in the context we consider. Moreover, these models rely on a single set of parameters to perform classification, thus limiting their flexibility.

A well-known method to estimate Gaussian mixture densities is to use the Expectation–Maximization (EM) algorithm [51], an approach that still requires the number of components to be fixed a priori. To eliminate this constraint, the EM algorithm can be combined with Minimum Description Length (MDL) [52], an agglomerative clustering strategy estimating the number of components which best fit the data. This approach is therefore an interesting alternative to our autonomous learning and is part of our experiments for comparison purpose.

This paper presents an extended and improved version of our previous work [9], which looked at applying Bayesian nonparametric methods to learn a perception model in the absence of labels and the number of clusters in a data set. In that work, we

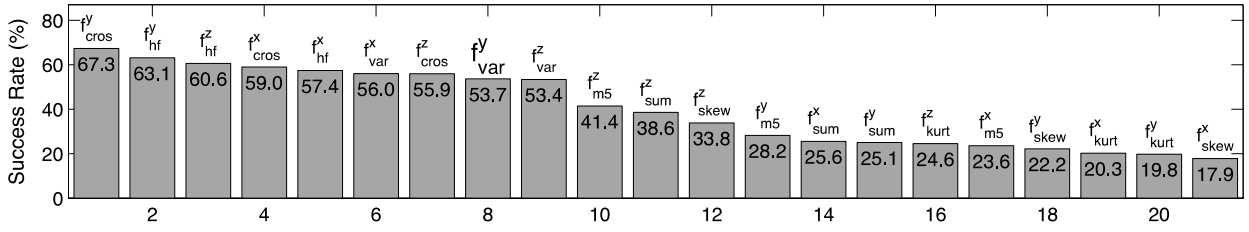

Fig. 6. Ranking of features extracted from one axis at a time. Values inside the bars are success rate percentages.

employed a Dirichlet process. This extended work is based on the Pitman–Yor process, a generalization of the Dirichlet process allowing power-law behavior concerning the population of the clusters. This method estimates the number of active clusters, based on a prior distribution on an infinite dimensional space of parameters. The flexibility provided by such a Bayesian nonparametric approach is highly desirable for autonomous mobile robots trying to learn a real-world environment, since the number of surfaces encountered is unknown and may increase over time.

5.1. Clustering: Pitman–Yor process mixture models

In this section, we introduce a class of random mixture models called Pitman–Yor process mixture models (PYPMMs). They are closely related to species sampling models [53] and have good computational properties. The fundamental aspect of the PYPMM is the Pitman–Yor process (PY) [54]. A Pitman–Yor process $\mathrm{PY}(d, \alpha, G_0)$ is a stochastic process used in Bayesian nonparametrics to define probability distributions over spaces of discrete distributions. It is parameterized by a base distribution G_0, which can be seen as a prior guess of draws from the process. It also has a discount parameter $0 \leq d < 1$ and a strength parameter $\alpha > -d$. Together, these two parameters (α, d) control the variability of the distribution around G_0. Using this distribution on the parameters of mixture models leads to the following specification of PYPMMs:

$$\begin{aligned} G \mid \alpha, G_0 &\sim \mathrm{PY}(d, \alpha, G_0) \\ \theta_i \mid G &\sim G \\ \mathbf{x}_i \mid \theta_i &\sim F(\theta_i) \end{aligned} \tag{2}$$

where F is the data distribution parameterized by θ, such as the mean and covariance of a multivariate Gaussian distribution when modeling Gaussian mixtures. These parameters θ are meant to capture the shape of tactile surface distributions in the feature space of the sensors. The distribution G, corresponding to the unknown mixture distribution, is discrete with probability one and infinite dimensional, providing PYPMMs with a countably infinite number of mixture components. It is precisely this distribution G which is of interest when learning the perception model of our tactile system, where the probabilistic model defined in (2) allows for Bayesian inference. The support of distribution G_0 represents the set of all candidate components to consider in the mixture where $G_0(\theta)$ is the prior probability of a component. Then, G assigns positive probabilities to a subset of these candidate components.

The stick-breaking construction of the Pitman–Yor process, presented by Ishwaran and James [55], is a representation that makes some properties of the distribution explicit. One of them is that a draw G is, with probability one, a weighted sum of an infinite number of point masses such that:

$$G = \sum_{k=1}^{\infty} \pi_k \delta_{\theta_k^*}. \tag{3}$$

In this equation, δ_θ is the unit mass located at θ and parameters are drawn i.i.d. from $\theta_k^* \sim G_0$. These point masses are used to indicate the components comprised in the mixture G. Moreover, the weight sequence $\pi_1, \pi_2, \ldots$, determining the importance of each component (akin to the number of samples per surface in our context), is distributed according to the following process[1]:

$$\begin{aligned} V_k &\sim \mathrm{Beta}(1 - d, \alpha + kd) \\ \pi_k &= V_k \prod_{j=1}^{k-1} (1 - V_j) \quad \text{for } k = 1, \ldots, \infty. \end{aligned} \tag{4}$$

[1] This process is known as a stick-breaking process.

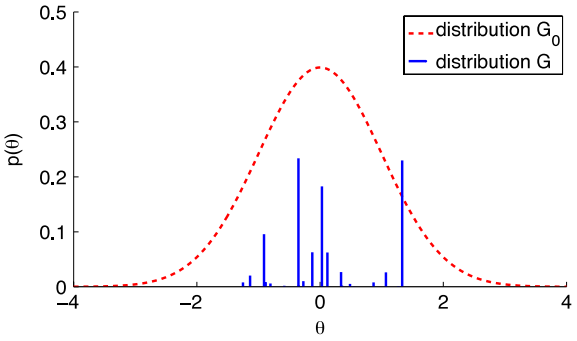

Fig. 7. A random draw G from a Pitman–Yor process is depicted by vertical bars. The Gaussian distribution corresponds to the base distribution G_0.

This process produces an infinite partition of a unit-length stick by recursively breaking it infinitely many times. Each part is interpreted as a probability mass that is randomly positioned according to G_0 in the space of parameters to generate the discrete probability distribution G. When associating a positive mass to a parameter (vector) θ, it indicates the existence of a specific component (a tactile surface) parameterized by θ in the mixture.

As a generic illustrative example, Fig. 7 shows a random sample drawn from a PY with Gaussian base distribution. The vertical bars correspond to the broken parts of the stick, where heights are the respective π values. Note that positions of the bars in the θ space are Gaussianly distributed as G_0 indicates. The resulting random sample G is simply the function depicted by these vertical bars summing to one. Finally, the number of bars is equal to the number of clusters, which is infinite.

The discreteness and clustering properties of PYs are essential for clustering our tactile data set with PYPMMs. These properties become evident when the random distribution G is marginalized out to obtain the following conditional probability distribution directly in the parameter space:

$$\theta_i \mid \theta_1, \ldots, \theta_{i-1} \sim \sum_{k=1}^{K} \frac{n_k - d}{\alpha + i - 1} \delta_{\theta_k^*} + \frac{\alpha + Kd}{\alpha + i - 1} G_0. \tag{5}$$

In this equation, n_k is the number of occurrences of unique value θ_k^* among the observed values $\theta_1, \ldots, \theta_{i-1}$ and K is the total number of unique values. Eq. (5) comprises a discrete measure (left term) over the K previously observed components, meaning there is a positive probability for an existing value to be resampled. As more observations become available, the discrete measure acquires more mass and the novelty probability (right term) decreases accordingly. Thus, it results in a suitable rich-gets-richer clustering behavior.

From a clustering perspective, the probability for observation $\mathbf{x}_i$ of joining the existing cluster k represented by θ_k^*, implying $\theta_i = \theta_k^*$, or to create its own new cluster is:

$$\begin{aligned} p(c_i = k \mid c_1, \ldots, c_{i-1}) &= \frac{n_{-i,k} - d}{i - 1 + \alpha} \\ p(c_i \neq k \text{ for all } k < K) &= \frac{\alpha + Kd}{i - 1 + \alpha} \end{aligned} \tag{6}$$

where c_i is the cluster indicator of $\mathbf{x}_i$ and $n_{-i,k}$ denotes the size of cluster k, excluding the ith observation. Eq. (6) ensures that large clusters tend to get bigger, since the probability of joining a cluster is roughly proportional to its size. On the other hand, the strength parameter α, which controls the new cluster probability, influences the total number K of non empty clusters. The discount parameter d also acts on K by decreasing the probability of joining small clusters and increasing the probability of creating new ones.

The Pitman–Yor process yields power-law behavior when $0 < d < 1$, which makes it more appropriate for applications involving natural phenomena [56]. When fixing $d = 0$, one recovers the well-known Dirichlet process, a particular case of Pitman–Yor process that does not yield power-law behavior.

To complete the previous prior, we have to specify the data distribution F. Given a set of observations, along with PY parameters G_0, α and d, the PYPMM yields a posterior distribution on the component parameters of a mixture. Exact inference for this posterior is analytically intractable; instead, it is sampled with Markov Chain Monte Carlo (MCMC) methods.

5.2. Pitman–Yor process mixture of Gaussians

The traditional approach to Gaussians mixture model aims to fix the number of components *a priori*. However, from the Bayesian nonparametric perspective, it is not necessary to fix this number, since it can be automatically inferred from the data. Thus, the number of clusters appearing in a data set can be estimated automatically. The PYPMM, as shown in [57] for the special case of Dirichlet processes, can be used to produce infinite Gaussian mixtures, thus avoiding finite representations. This model has been further extended with vague priors on the model hyperparameters in [58].

This section introduces the Pitman–Yor process mixture of Gaussians (PYPMoG) from the general PYPMM. To do so, we define the distribution F in the feature space as multivariate Gaussian:

$$\mathbf{x}_i \mid \theta_i \sim \mathcal{N}(\boldsymbol{\mu}_{c_i}, S_{c_i}^{-1}) \tag{7}$$

where $\boldsymbol{\mu}$ is the mean and S is the inverse covariance matrix. These parameters, which correspond to θ in Section 5.1, should then be assigned a prior distribution G_0 indicating the kind of Gaussians we are likely to use as components. Following the conditionally conjugate model proposed in [58], we define this G_0 prior as:

$$\boldsymbol{\mu}_k \mid \boldsymbol{\xi}, R \sim \mathcal{N}(\boldsymbol{\xi}, R^{-1}) \tag{8}$$

$$S_k \mid \beta, \Sigma \sim \mathcal{W}(\beta, (\beta\Sigma)^{-1}) \tag{9}$$

where $\mathcal{W}$ denotes the Wishart distribution with β degrees of freedom and scale matrix $(\beta\Sigma)^{-1}$. This prior probability distribution on Gaussians determines the expected covariance matrices in the mixture and their expected position in the feature space.

For robustness, we adopt a hierarchical structure and specify vague hyperpriors on hyperparameters. Therefore, the base distribution G_0, corresponding to $p(\boldsymbol{\mu}, S|\boldsymbol{\xi}, R, \beta, \Sigma)$, is considered uncertain and is learned from data. The following prior distribution reflects our uncertainty concerning the unknown distribution (9) over cluster covariances:

$$\Sigma \sim \mathcal{W}\left(D, \frac{1}{D}Q\right), \qquad \left(\frac{1}{\beta - D + 1}\right) \sim \mathcal{G}\left(1, \frac{1}{D}\right) \tag{10}$$

where D is the number of dimensions in the observation space $\mathbf{x}$. Notice that this prior makes it increasingly harder to expect smaller and smaller Gaussians during the inference, thus favoring larger covariance for precision matrices S. Moreover, when assuming properly scaled training data with zero mean and unit covariance, the matrix Q should be the identity matrix. In that case, R^{-1} should also be identity along with $\boldsymbol{\xi}$ being set to zero, fixing the probability distribution (8).

Completing the hyperprior on hyperparameters requires specifying distributions for both α and d. Consequently, an inverse-gamma distribution is defined on α and a uniform distribution is used for d, which can be written as:

$$\alpha^{-1} \sim \mathcal{G}(1/2, 1/2), \qquad d \sim \mathcal{U}(0, 1). \tag{11}$$

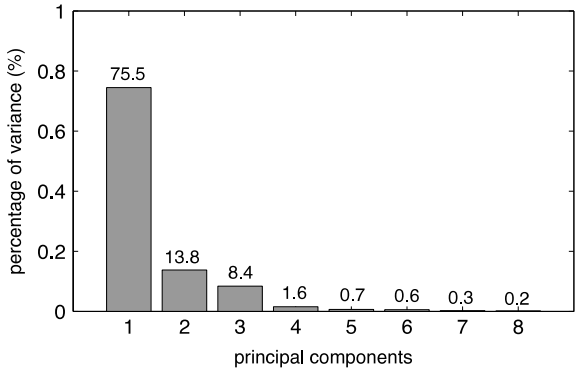

Fig. 8. Percentage of variance of the principal components.

5.3. Experiments with unlabeled data sets

Our final goal is to achieve autonomous learning of surfaces with as little information as possible provided by a human. This task is, by comparison, significantly harder than the supervised learning experiments presented in Section 4. For the unsupervised learning experiments, the information provided was the extracted features; no labels were given. The true number of surfaces was also unspecified, to simulate an agent autonomously collecting a training data set in its environment.

We did not use features f_{skew}, f_{kurt} and f_{m5} in these experiments as their respective distribution had high kurtosis, violating the assumption that features are normally distributed.[2] Fortunately, according to the analysis of features realized in Section 4.3, these 3 features basically coincide with ignoring the less informative features (see Figs. 5 and 6). Consequently, we expect that dropping these features should not significantly decrease the quantity of information contained in the data.

The remaining features were further processed because (i) they were strictly positive and (ii) their variance appeared proportional to the distance from the origin, giving the data sets a conical shape with apex on the origin. As a result, we used $\log(f_{sum})$, $\log(f_{cros})$ and $\log(f_{hf})$ to obtain negative values, equivalent to using log-normal distributions in the mixture. Moreover, we used $\sqrt{f_{\text{var}}}$ on the variance feature since the standard deviations had kurtosis closer to a normal distribution. After these transformations, we performed a Principal Component Analysis (PCA) for dimensionality reduction. We used the 3 principal components having variance 1.18, 0.22 and 0.13, where the fourth and subsequent components had variance 0.02 or less. Fig. 8 shows the corresponding percentage of the total variance held by the first principal components. For the experiments, we constructed multiple training data sets $\mathbf{X}_{train}$ from raw data in $\mathbf{S}_A$ using time-window sizes W of 200, 400, 800 and 1600 samples.

We considered the unsupervised learning task of surface identification as a clustering problem suitable to PYPMoG and thus assumed that the probe's signals coming from the same material were normally distributed in the feature space. The resulting clustering corresponds to the number of Gaussian components found in the data: a single Gaussian per surface according to our hypothesis. To perform posterior inference for this model, we employed a Gibbs sampling approach based on auxiliary parameters to sample posterior mixtures [59]. Markov Chain Monte Carlo methods are often computationally expensive for large data sets; hence, we reduced the size of all training sets to 50 feature vectors $\mathbf{x}_i$ per surface, or in other words, providing 12.5, 25,

[2] Modeling the density of heavy-tail distributions may require multiple Gaussians, leading to an over-clusterization of the data.

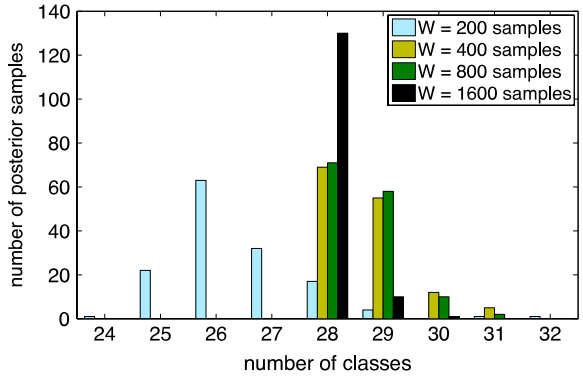

Fig. 9. Estimating the number of surfaces in an unlabeled data set. The most probable numbers of surfaces are 28 ($W = 1600$), 28 ($W = 800$), 28 (for $W = 400$) and 26 ($W = 200$), all close or equal to the true value of 28.

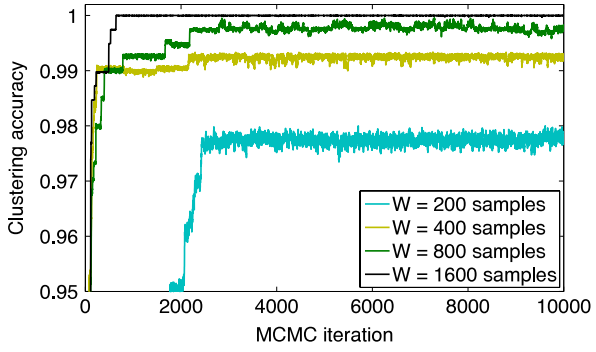

Fig. 10. Clustering accuracy on unlabeled data sets as a function of the MCMC iteration, for different time window sizes W. For these experiments, the true number of clusters was unspecified.

50 and 100 s of data depending on the different window size W. In particular, these training periods were taken at the very beginning of the first gathering procedure, which corresponds to data in $\mathbf{S}_A$.

The parameters used to perform inference are the following. For the Pitman–Yor process prior, the base distribution G_0 is a 3-dimensional Gaussian distribution since $D = 3$. The hyperparameters have prior distributions specified by Eq. (11) where initially $\alpha = 1$ and $d = 0$. As stated earlier, (8) is a unit Gaussian with 0 mean. On the other hand, we have set parameter $\beta = D$ and matrix Σ is diagonal with value 0.01 since we do not expect clusters to have unit variance which would produce a single big cluster due to data normalization. Finally, we used 5 auxiliary parameters during the resampling phase.

5.3.1. Identifying the number of surfaces

We first evaluated the capacity to recover the exact number of surfaces in the training data sets for various windows size W, by sampling the posterior distribution on the number of clusters. The capability to recover this value is an essential feature of autonomous training of a perception system. Fig. 9 shows the posterior distributions over K for all window sizes. If we consider only the larger window sizes $W \geq 400$, the respective posterior distributions give the highest probability to the true number of surfaces (28). Using the smallest window size $W = 200$ yields slightly inferior estimates but still assigns significant probability to the true value.

Although a sampled posterior mixture might have the correct number of clusters, it is unclear whether all data belonging to a given cluster actually represent a single surface or different surfaces having similar features. This might be particularly true for shorter time-windows, as they generate clusters with larger variances in the feature space.

Table 1
MAP prediction success rates on five data sets $\mathbf{X}^j_{test}$ of unseen samples, with the trained PYPMoG model.

$\mathbf{X}^1_{test}$	$\mathbf{X}^2_{test}$	$\mathbf{X}^3_{test}$	$\mathbf{X}^4_{test}$	$\mathbf{X}^5_{test}$
99.99%	99.99%	99.88%	99.77%	99.70%

5.3.2. Pairwise correct classification

To measure the clustering accuracy, we defined a metric based on pairwise correct classification. It uses the hidden labels to compute a dissimilarity measure with the true clusters. An error is made when two data points are assigned to the same cluster and should not be, or when they are in different clusters and should be in the same one. The resulting clustering accuracy is then 1 minus the error ratio. A score of 1 is achieved only for a perfect clustering. Fig. 10 shows the clustering accuracy obtained as the Markov chain evolved. The initial region corresponds to a burn-in period, and therefore should not be considered until stationarity is (probably) reached.

5.3.3. MAP prediction on unseen data

To predict samples from new data sets $\mathbf{X}^j_{test}$, we used the maximum a posteriori (MAP) estimate learned from the unlabeled data set with $W = 1600$. Data sets $\mathbf{X}^j_{test}$ were comprised of samples $\mathbf{x}_i$ collected immediately after the training set $\mathbf{X}_{train}$, with 2800 samples $\mathbf{x}_i$ in each of them. These data sets $\mathbf{X}^j_{test}$ represent consecutive periods of 100 s of sequential data in $\mathbf{S}_A$, where the overall evaluation is done over 500 s. To determine the class, we computed the likelihood of each Gaussian component, one for each cluster. The label assignment is done by finding the cluster having the greatest likelihood. Classification results with this method are presented in Table 1. As we can see from these results, the Gaussian components found using the PYPMoG method are extremely reliable for classifying unseen data.

It is common to evaluate the accuracy of a model by constructing confusion matrices to easily see where errors are made. Since the labeling is done by the mean of clustering, it is hard to produce such a matrix, especially when the number of classes is incorrect. However, in the case $W = 1600$, the posterior models are almost always perfect, which would often produce diagonal confusion matrices. In particular, the most probable model a posteriori has a perfect clustering and therefore has a 1-to-1 mapping with the true labels. This model was used to produce Table 1 and we reused it to produce the confusion matrix shown in Fig. 11. This confusion matrix has been computed by grouping all 5 test sets $\mathbf{X}^1_{test}$ to $\mathbf{X}^5_{test}$ into a single set $\mathbf{X}^{1:5}_{test}$ and predictions were again done by maximum likelihood.

When looking at results in Table 1, we observe a slight decay of success rate over time. This is due to the alteration of some surfaces, a phenomenon that is becoming increasingly important for the latest test sets. For this reason, identifying surfaces for data in set $\mathbf{S}_B$ is a harder task than the previous one. Recalling that the model was trained with data in $\mathbf{S}_A$ only, the pairwise correct classification result of 99.07% on $\mathbf{S}_B$ is respectable. Fig. 13 shows the confusion matrix corresponding to this result. As a measure of comparison, using Expectation–Maximization (EM) with Minimum Description Length (MDL) [60] to cluster the data and predict unseen data with maximum likelihood produced a pairwise correct classification result of 99.03%.

Another performance metric is the overall classification rate of the model which can be computed by taking the trace of the confusion matrix and divide by the number of surfaces. Using this evaluation method, the PYPMoG achieved a classification success rate of 90.43% while EM with MDL achieved a success rate of 90.29%. This criterion, however, ignores the fact that the model learned some structure among the misclassified data points since errors are not uniformly distributed across surfaces.

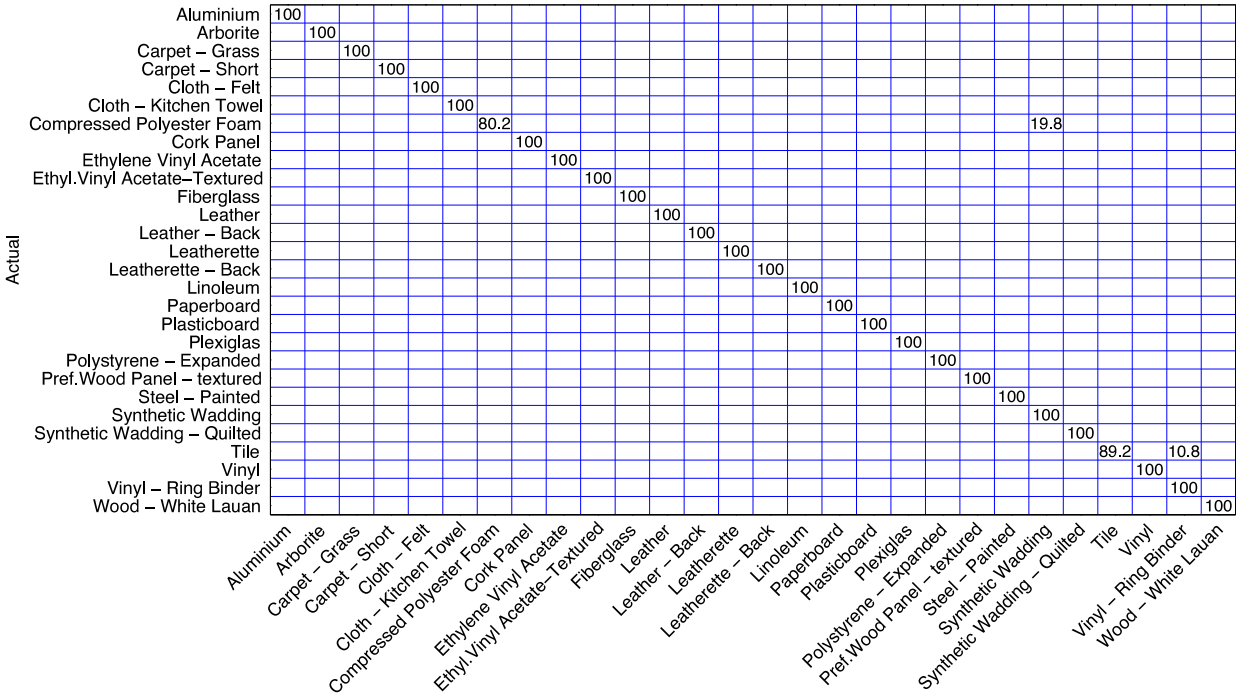

Fig. 11. Confusion matrix for the combined test set $\mathbf{X}_{test}^{1:5}$ using time-window size of $W = 1600$ samples.

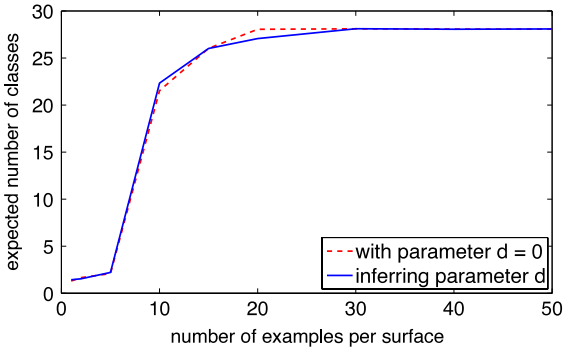

Fig. 12. The expected posterior number of classes as a function of the training set size. All training sets contain the same number of examples for each surface. The blue solid line corresponds to the PYPMoG and the red dashed line is the DPMoG.

5.3.4. The learning curve

A second series of experiments has been conducted to assess the learning curve of the Bayesian nonparametric approach with respect to the number of training examples per surface. Indeed, the fewer the examples provided to the algorithm, the harder it becomes to extract precise structure from the data. Therefore, we trained multiple PYPMoG with data sets containing various numbers of time windows per surface. The number of data points per surface ranged from 1 to 50 examples and was obtained by extracting features from consecutive time windows. Note that in all data sets, the number of examples per surface is exactly the same for all 28 surfaces, thus creating data sets of sizes varying from 28 to 1400 data points. Since we already know from previous experiments that a time-window size $W = 1600$ is enough to achieve a perfect clustering with 50 examples per surface, we only consider this particular time-window size and ignore models with $W \leq 800$.

Fig. 12 shows the learning curve for a PYPMoG with both hyperparameters free and for a PYPMoG having its discount parameter $d = 0$ fixed, namely the former Dirichlet process mixture of Gaussian (DPMoG) used in [9]. The curves represent the expected number of surfaces identified in the training data by averaging over 140 posterior samples for each data set. Fig. 14 shows the expected posterior accuracies for both priors, where we ignored the smallest data sets for their small accuracies.

Table 2
Comparison of clustering accuracies between different unsupervised learning methods, for different time-window size W.

	$W = 200$	$W = 400$	$W = 800$	$W = 1600$
k-means	0.9783	0.9877	0.9947	0.9961
EM	0.9789	0.9807	0.9915	0.9931
EM+MDL	0.9600	0.9891	0.9963	**1.0000**
PY_{post}	0.9775	0.9924	0.9977	0.9999
PY_{MAP}	**0.9800**	**0.9934**	**0.9995**	**1.0000**

5.4. Comparison with other unsupervised learning methods

We compared the PYPMoG results with three other unsupervised learning algorithms: k-means, Expectation–Maximization for Gaussian mixtures and Expectation–Maximization with Minimum Description Length. In contrast to the PYPMoG we employed, the first two algorithms required the true number of clusters to be provided. On the other hand, the last one estimates this number by scoring different models. Each method was run 10,000 times, and only the *best* model was kept for comparison purposes. Table 2 reports the clustering accuracies obtained with the four methods. For the PYPMoG, we used the *averaged posterior* clustering accuracies as well as the maximum model (MAP) from Fig. 10, excluding the burn-in period. The performances obtained by the Bayesian nonparametric method compares favorably with the three other methods, despite the fact that PYPMoG had to estimate the number of clusters in the data set.

The inference procedure we used to learn the PYPMoG model is based on Markov Chain Monte Carlo simulation of the posterior. A single chain of the length of 10 000 takes a few hours to compute and aims for multiple posterior samples. Compared to k-means, EM and EM+MDL, which take only a few minutes of optimization, the PYPMoG is computationally more demanding. However, it is possible to define the inference procedure as an

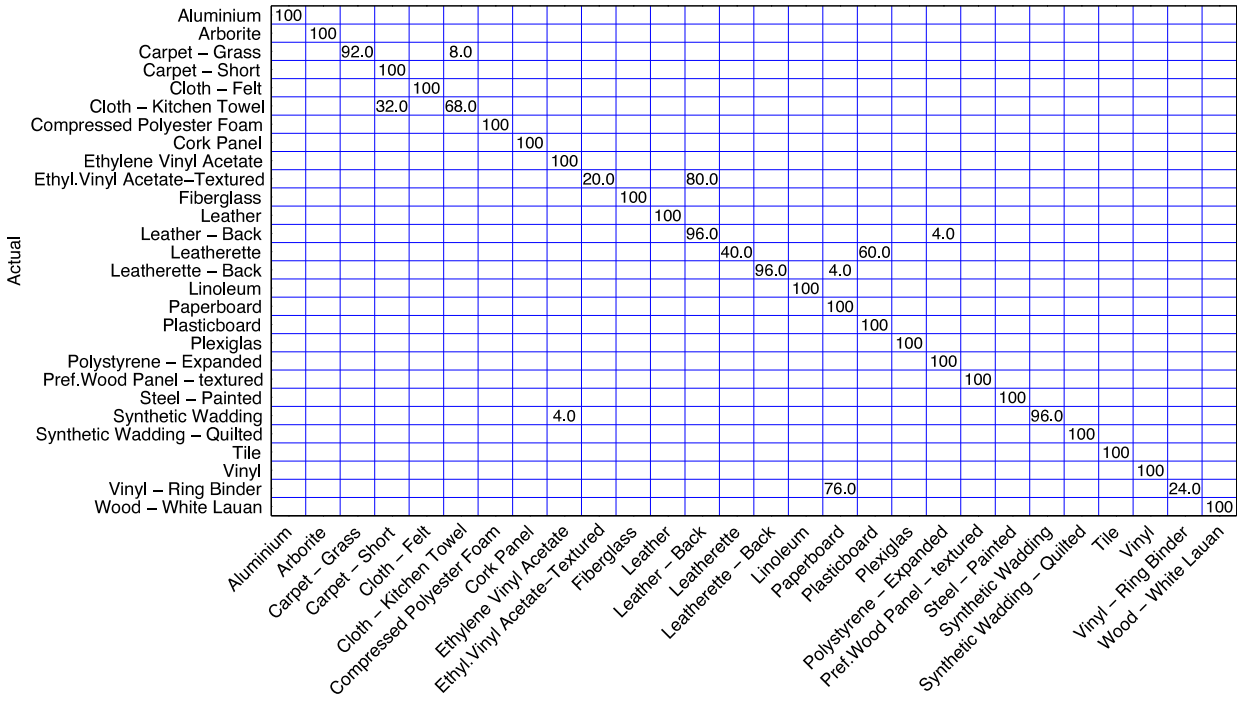

Fig. 13. Confusion matrix for test set $\mathbf{X}^{B}_{test}$ using time-window size of $W = 1600$ samples.

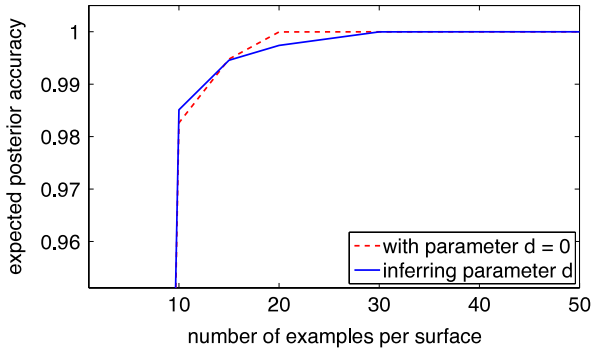

Fig. 14. The expected posterior clustering accuracy as a function of the training set size. All training sets contain the same number of examples for each surface. The blue solid line corresponds to the PYPMoG and the red dashed line is the DPMoG.

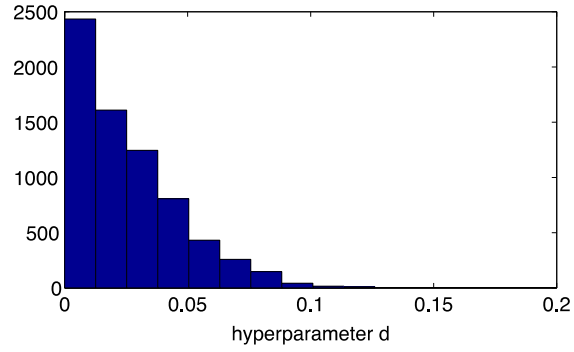

Fig. 15. Histogram of the Markov chain obtained for the discount hyperparameter d. The chain has been obtained for a data set with 50 examples per surface and with $W = 1600$.

optimization of the posterior to directly obtain an MAP estimate without simulating the whole distribution. This modification could significantly reduce the computational cost of learning.

6. Discussion

The effectiveness of our combined tactile sensor and unsupervised learning approach to surface recognition stems from the fact that for every surface, the features were normally distributed and the clusters were well separated. The choice of extracted features and the time-window size from which they are computed have a significant impact on identification performances. Results reported in Table 2 support our observation that smaller time-window size increases the variance of every cluster, which in turn causes more overlapping among the clusters and makes the surface identification harder.

Moreover, to recognize a particular surface, enough data should be available so that the algorithm can distinguish two clusters with smaller variance from one encompassing multiple surfaces and having larger variance. From Fig. 12, we observe that with $W = 1600$, 30 training examples per surface is enough to recognize all surfaces, which represents 60 s of raw data gathered from a single surface. Providing fewer data resulted in the creation of larger clusters representing 2 or more surfaces. What could have influenced the results is that with Bayesian nonparametrics models, the complexity scales with respect to the number of data. In other words, the expected number of observed surfaces gets larger as N grows since the model expects surfaces to be drawn *iid* from the mixture distribution (3). Therefore, providing a single data point per surface ($N = 28$ examples) is biased towards creating one big cluster regrouping every surface, unless we tune hyperparameters.

When looking at the results of Fig. 14, one can argue about the significance of using the PYPMoG over the DPMoG. Moreover, the posterior distribution on hyperparameter d illustrated in Fig. 15 indicates that parameter d contributes slightly to model the distribution on component proportions π. However, in real-world environments, the number of examples per surface will surely not be uniform and one could expect to encounter various surfaces with frequencies distributed according to a power-law distribution, justifying the use of PYPMoG.

When comparing our results with the ones reported in Section 2.3, we first have to consider that most approaches used

different sensors and surfaces, and have often provided the number of surfaces to recognize. Therefore, the success rate of 90.43% that we obtained in our most realistic experiment is essentially only comparable with the work of Sinapov et al. [33] as they used similar sensors and surfaces. In their supervised learning experiment, they obtained a success rate of 80% on surface categorization, and 58% was achieved for unsupervised surface recognition. The inferior success rates can be explained by the fact that their sensor was mounted on a robotic arm, introducing an additional source of uncertainty in the data acquisition process.

7. Conclusion and future work

In this paper, we presented work related to the classification of information coming from a tactile probe based on a triple axis accelerometer. Surface classification relied on features extracted from the accelerometer's measurements in the time and frequency domains. The learnability of a 28-surface data set was first established via the Support Vector Machine classifier.

We then demonstrated that there is a potential for surface identification by autonomous agents, including the difficult case when the number of surfaces in a tactile data set is unknown. This was accomplished through the use of Bayesian nonparametric modeling named the Pitman–Yor process mixture model. Although the clustering performance of the Pitman–Yor process mixture model was similar to those obtained with the Dirichlet process mixture model presented in our previous paper [9], we expect the method presented in this paper to perform significantly better in real environments, where the power-law behavior is desirable.

As part of future work, we are investigating applications for this type of tactile sensing in robotics. For example, a miniature spring-loaded version of this probe could be attached to the finger of a manipulator, possibly imitating a nail [33]. This type of data acquisition has the advantage of avoiding damaging surfaces as multiple areas can be used to perceive the texture. We are also examining how to fuse the rich information obtained from tactile sensing with those shape estimations obtained by manipulation in the context of blind object recognition. Finally, sensor fusion with visual information, possibly borrowing from [34], is also a future area of investigation.

8. Role of the funding source

Financial support for this research was provided by an NSERC Discovery grant from the Government of Canada. The funding agency played no role other than providing financial support.

References

[1] V. Sujan, S. Dubowsky, Efficient information-based visual robotic mapping in unstructured environments, Int. J. Robot. Res. 24 (4) (2005) 275.
[2] S. Garrido, L. Moreno, D. Blanco, Exploration and mapping using the vfm motion planner, IEEE Trans. Instrum. Meas. 58 (8) (2009) 2880–2892.
[3] T.J. Prescott, B. Mitchinson, R.A. Grant, Vibrissal behavior and function, Scholarpedia 6 (10) (2011) 6642.
[4] R. Howe, M. Cutkosky, Sensing skin acceleration for slip and texture perception, in: IEEE Int'l Conf. on Robotics and Automation, vol. 1, 1989, pp. 145–150.
[5] E. Petriu, W. McMath, S. Yeung, N. Trif, Active tactile perception of object surface geometric profiles, IEEE Trans. Instrum. Meas. 41 (1) (1992) 87–92.
[6] A. Okamura, M. Cutkosky, Feature detection for haptic exploration with robotic fingers, Int. J. Robot. Res. 20 (12) (2001) 925.
[7] E. Petriu, S. Yeung, S. Das, A. Cretu, H. Spoelder, Robotic tactile recognition of pseudorandom encoded objects, IEEE Trans. Instrum. Meas. 53 (5) (2004) 1425–1432.
[8] P. Payeur, C. Pasca, A. Cretu, E. Petriu, Intelligent haptic sensor system for robotic manipulation, IEEE Trans. Instrum. Meas. 54 (4) (2005) 1583–1592.
[9] P. Dallaire, D. Emond, P. Giguere, B. Chaib-draa, Artificial tactile perception for surface identification using a triple axis accelerometer probe, in: Int. Symp. on Robotic and Sensors Environments, ROSE, 2011.
[10] M. Hollins, S. Bensmaïa, The coding of roughness, Can. J. Exp. Psychol. 61 (3) (2007) 184.
[11] R. Johansson, J. Flanagan, Coding and use of tactile signals from the fingertips in object manipulation tasks, Nat. Rev. Neurosci. 10 (5) (2009) 345–359.
[12] H. Muhammad, C. Recchiuto, C. Oddo, L. Beccai, C. Anthony, M. Adams, M. Carrozza, M. Ward, A capacitive tactile sensor array for surface texture discrimination, Microelectron. Eng. 88 (8) (2011) 1811–1813.
[13] S. Lederman, R. Klatzky, Hand movements: a window into haptic object recognition* 1, Cogn. psychol. 19 (3) (1987) 342–368.
[14] S. Chitta, J. Sturm, M. Piccoli, W. Burgard, Tactile sensing for mobile manipulation, IEEE Trans. Robot. 27 (3) (2011) 558–568.
[15] G. Cannata, M. Maggiali, G. Metta, G. Sandini, An embedded artificial skin for humanoid robots, in: IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems, 2008, MFI 2008, IEEE, 2008, pp. 434–438.
[16] G. Cannata, R. Dahiya, M. Maggiali, F. Mastrogiovanni, G. Metta, M. Valle, Modular skin for humanoid robot systems, in: 4th Int. Conf. on Cognitive Systems, 2010.
[17] A. Schmitz, P. Maiolino, M. Maggiali, L. Natale, G. Cannata, G. Metta, Methods and technologies for the implementation of large-scale robot tactile sensors, IEEE Trans. Robot. 27 (3) (2011) 389–400.
[18] P. Mittendorfer, G. Cheng, Humanoid multimodal tactile-sensing modules, IEEE Trans. Robot. 27 (3) (2011) 401–410.
[19] R. Russell, Using tactile whiskers to measure surface contours, IEEE Trans. Robot. Autom. (1992).
[20] G. Scholz, C. Rahn, Profile sensing with an actuated whisker, IEEE Trans. Robot. Autom. 20 (1) (2004) 124–127.
[21] D. Jung, A. Zelinsky, Whisker based mobile robot navigation, in: IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems, IROS, vol. 2, 1996, pp. 497–504.
[22] T. Clements, C. Rahn, Three-dimensional contact imaging with an actuated whisker, IEEE Trans. Robot. 22 (4) (2006) 844–848.
[23] R.A. Russell, J.A. Wijaya, Recognising and manipulating objects using data from a whisker sensor array, Robotica 23 (5) (2005) 653–664.
[24] M.J. Pearson, A.G. Pipe, C. Melhuish, B. Mitchinson, T.J. Prescott, Whiskerbot: a robotic active touch system modeled on the rat whisker sensory system, Adapt. Behav. 15 (3) (2007) 223–240. Animals, Animats, Software Agents, Robots, Adaptive Systems.
[25] N.J. Cowan, E.J. Ma, M. Cutkosky, R.J. Full, A biologically inspired passive antenna for steering control of a running robot, in: International Symposium on Robotics Research, Springer, 2003, pp. 541–550.
[26] J. Hipp, E. Arabzadeh, J. Conradt, E. Zorzin, C. Kayser, M. Diamond, P. König, Texture signals in whisker vibrations, J. Neurophysiology 95 (3) (2006) 1792–1799.
[27] M. Fend, S. Bovet, H. Yokoi, R. Pfeifer, An active artificial whisker array for texture discrimination, in: IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems, IROS, 2003, pp. 1044–1049.
[28] N. Jamali, C. Sammut, Majority voting: material classification by tactile sensing using surface texture, IEEE Trans. Robot. 27 (3) (2011) 508–521.
[29] J.A. Fishel, G.E. Loeb, Bayesian exploration for intelligent identification of textures, Front. Neurorobot. 6 (4) (2012).
[30] J. Edwards, J. Lawry, J. Rossiter, C. Melhuish, Extracting textural features from tactile sensors, Bioinspiration Biomimetics 3 (2008) 1–12.
[31] M. Johnsson, C. Balkenius, Sense of touch in robots with self-organizing maps, IEEE Trans. Robot. 27 (3) (2011) 498–507.
[32] S. Takamuku, G. Gomez, K. Hosoda, R. Pfeifer, Haptic discrimination of material properties by a robotic hand, in: IEEE 6th International Conference on Development and Learning, 2007, ICDL 2007, 2007, pp. 1–6.
[33] J. Sinapov, V. Sukhoy, R. Sahai, A. Stoytchev, Vibrotactile recognition and categorization of surfaces by a humanoid robot, IEEE Trans. Robot. 27 (3) (2011) 488–497.
[34] O. Kroemer, C.H. Lampert, J. Peters, Learning dynamic tactile sensing with robust vision-based training, IEEE Trans. Robot. 27 (3) (2011) 545–557.
[35] S. Ghosal, A. Van der Vaart, Fundamentals of Nonparametric Bayesian Inference, 2011.
[36] T. Ferguson, Bayesian density estimation by mixtures of normal distributions, Recent Adv. Stat. (1983) 287–303.
[37] J. Joseph, F. Doshi-Velez, N. Roy, A Bayesian nonparametric approach to modeling battery health, in: Proceedings of the IEEE International Conference on Robotics and Automation, ICRA 2012, St Paul, MN, 2012.
[38] E. Fox, E. Sudderth, M. Jordan, A. Willsky, Bayesian nonparametric inference of switching dynamic linear models, IEEE Trans. Signal Process. 59 (4) (2011) 1569–1585.
[39] T. Taniguchi, K. Hamahata, N. Iwahashi, Unsupervised segmentation of human motion data using a sticky hierarchical Dirichlet process-hidden Markov model and minimal description length-based chunking method for imitation learning, Adv. Robot. 25 (17) (2011) 2143–2172.
[40] V. Kruger, V. Tikhanoff, L. Natale, G. Sandini, Imitation learning of non-linear point-to-point robot motions using Dirichlet processes, in: IEEE International Conference on Robotics and Automation, ICRA, 2012, 2012, pp. 2029–2034.

[41] S.P. Chatzis, D. Korkinof, Y. Demiris, A nonparametric Bayesian approach toward robot learning by demonstration, Robot. Auton. Syst. 60 (6) (2012) 789–802.
[42] M.V. Jakuba, D. Steinberg, J.C. Kinsey, D.R. Yoerger, R. Camilli, O. Pizarro, S.B. Williams, Toward automatic classification of chemical sensor data from autonomous underwater vehicles, in: 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, 2011, pp. 4722–4727.
[43] P. Giguere, G. Dudek, Surface identification using simple contact dynamics for mobile robots, in: IEEE Int'l Conf. on Robotics and Automation, ICRA, Kobe, Japan, 2009.
[44] P. Giguere, G. Dudek, A simple tactile probe for surface identification by mobile robots, IEEE Trans. Robot. 27 (3) (2013).
[45] C. Weiss, H. Frohlich, A. Zell, Vibration-based terrain classification using support vector machines, in: IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems, IROS, 2006, pp. 4429–4434.
[46] K. Hosoda, Y. Tada, M. Asada, Anthropomorphic robotic soft fingertip with randomly distributed receptors, Robot. Auton. Syst. 54 (2) (2006) 104–109.
[47] C.-C. Chang, C.-J. Lin, LIBSVM: a library for support vector machines, software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm, 2001.
[48] A.K. Jain, R.P.W. Duin, J. Mao, Statistical pattern recognition: a review, IEEE Trans. Pattern Anal. Mach. Intell. 22 (1) (2000) 4–37.
[49] R. Kohavi, G.H. John, Wrappers for feature subset selection, Artif. Intell. 97 (1) (1997) 273–324.
[50] J. Shi, J. Malik, Normalized cuts and image segmentation, IEEE Trans. Pattern Anal. Mach. Intell. 22 (8) (2000) 888–905.
[51] A.P. Dempster, N.M. Laird, D.B. Rubin, Maximum likelihood from incomplete data via the em algorithm, J. R. Stat. Soc. Ser. B 39 (1) (1977) 1–38.
[52] J. Rissanen, A universal prior for integers and estimation by minimum description length, Ann. Statist. (1983) 416–431.
[53] J. Pitman, Some developments of the blackwell-macqueen urn scheme, Lect. Notes Monogr. Ser. (1996) 245–267.
[54] J. Pitman, M. Yor, The two-parameter Poisson–Dirichlet distribution derived from a stable subordinator, Ann. Probab. (1997) 855–900.
[55] H. Ishwaran, L. James, Gibbs sampling methods for stick-breaking priors, J. Amer. Statist. Assoc. 96 (453) (2001) 161–173.
[56] Y. Teh, M. Jordan, Hierarchical Bayesian nonparametric models with applications, Bayesian Nonparametrics 28 (2009) 158.
[57] C. Rasmussen, The infinite Gaussian mixture model, Adv. Neural Inf. Process. Syst. 12 (2000) 554–560.
[58] D. Görür, Nonparametric Bayesian Discrete Latent Variable Models for Unsupervised Learning, Ph.D. Thesis, Max Planck Inst. for Biological Cybernetics, 2007.
[59] R. Neal, Markov chain sampling methods for Dirichlet process mixture models, J. Comput. Graph. Statist. 9 (2) (2000) 249–265.
[60] C.A. Bouman, M. Shapiro, G. Cook, C.B. Atkins, H. Cheng, Cluster: An Unsupervised Algorithm for Modeling Gaussian Mixtures, 1997.

Patrick Dallaire was born in 1982. He received the B.Sc. degree in Computer Science in 2008 and his M.Sc. degree in Artificial Intelligence in 2010 from Laval University, Canada. He is currently a Ph.D. student under the supervision of Prof. Brahim Chaib-draa and Prof. Philippe Giguère, and joined the DAMAS laboratory research group in 2007. His current research interests include reinforcement learning, Bayesian nonparametric learning and structure learning.

Philippe Giguere received a B.A.Sc. in engineering physics from Laval University, Quebec City, in 1996, a M.S. in computer science from Northeastern University, Boston, in 2003 and a Ph.D. degree from McGill University, Montreal, in 2010. He is currently assistant professor at the Department of Computer Science at Laval University in Quebec City. His current research interests include mobile robotics in harsh environments (arctic, underwater), machine learning and tactile sensing.

Daniel Émond was born is 1976. He studied computer engineering at the Royal Military College of Canada then proceed to study Computer Science at Laval University where he received his B.Sc. degree in 2001. He is currently pursuing an M.Sc. degree in Computer Science under the supervision of Prof. Philippe Giguère. His current research interests include machine vision, automation, machine learning mechatronics and bionics.

Brahim Chaib-draa received a Diploma in Computer Engineering from the École Supérieure d'Électricité (SUPELEC) de Paris, Paris, France, in 1978 and a Ph.D. degree in Computer Science from the Université du Hainaut-Cambrésis, Valenciennes, France, in 1990. In 1990, he joined the Department of Computer Science and Software Engineering at Laval University, Quebec, QC, Canada, where he is a Professor and Group Leader of the Decision for Agents and Multi-Agent Systems (DAMAS) Group. His research interests include agent and multiagent computing, machine learning and complex decision making. He is the author of several technical publications. Dr. Chaib-draa is a member of ACM and AAAI and senior member of the IEEE Computer Society.