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