

A Simple Tactile Probe for Surface Identification by Mobile Robots

Philippe Giguere and Gregory Dudek

Abstract—This paper describes a tactile probe designed for surface identification, in a context of all-terrain low-velocity mobile robotics. The proposed tactile probe is made of a small metallic rod with a single-axis accelerometer attached near its tip. Surface identification is based on analyzing acceleration patterns induced at the tip of this mechanically robust tactile probe, while it is passively dragged along a surface. A training data set was collected over ten different indoor and outdoor surfaces. Classification results for an artificial neural network were positive, with 89.9% and 94.6% success rate for 1 and 4 second time-windows of data, respectively. We also demonstrated that the same tactile probe can be used for unsupervised learning of terrains. For 1 second time-windows of data, the classification success rate was only reduced to 74.1%. Finally, a blind mobile robot, performing real-time classification of surfaces, demonstrated the feasibility of this tactile probe as a guidance mechanism.

Index Terms—Tactile Sensing, Terrain Identification, Outdoor Mobile Robots, Accelerometer, Unsupervised Learning

I. INTRODUCTION

TERRAIN identification is a fundamental task for mobile robots operating in unstructured environments. By identifying terrain types, mobile vehicles can better navigate their environment. For example, they can avoid certain terrain types known to be difficult, adjust their gaits based on some prior knowledge about terrains, or use the extra information for mapping and localization purposes. While cameras, range sensors or ground penetrating radar have been considered for terrain identification purposes, tactile sensing can provide an economical and robust solution to this problem, by collecting information directly related to the mechanical properties of terrain surfaces.

Over the years, a number of terrain identification techniques have also been developed, based on vehicle-mounted inertial sensors [1][2][3][4][5][6] or visual sensors tracking targets in the visual plane [7]. These techniques exploit measurable changes in overall system dynamics as the robot is traversing different surfaces. In essence, they represent a form of tactile sensing, with the wheels of the vehicles acting as tactile probes. The quality of the terrain information gathered by such systems is, however, limited by the mechanical response of the wheels and the vehicle assembly itself [5][8]. This mechanical response cannot be modified in a manner that would significantly improve terrain identification, without affecting the mobility of the robot itself.

Philippe Giguere is with the Department of Computer Science and Software Engineering, Laval University, Quebec City, Quebec, Canada. e-mail: philippe.giguere@ift.ulaval.ca.

Gregory Dudek is with the School of Computer Science, McGill University, Montreal, Quebec, Canada. email: dudek@cim.mcgill.ca

Manuscript received xx; revisedxx.

On the other hand, tactile sensors have the potential to be a more discriminative approach to terrain identification. In its simplest expression, a tactile probe can be used to detect anomalous surface conditions, such as holes in the ground or obstacles. More interestingly, a tactile probe can be designed to provide subtle information related to the mechanical properties of terrains, thereby facilitating their identification. This ability to identify terrains can be exploited to enhance robot mobility, by altering behavior based on the terrain type. For example, a blind robot could be made to navigate a bounded area, by relying on surface texture differences between a particular target area and its surroundings.

Tactile sensing can also be used when surface properties, used for identification purposes, cannot be directly observed by optical sensors, such as in the presence of fog or superficial dust. Another potential application employing rich tactile information is the automatic training of monocular vision terrain identification systems for mobile robots. In this paradigm, training examples are collected by a tactile sensing modality, and these examples are used to train a vision system. This topic has been explored in Krebs *et al.* [9], where local information captured by an inertial measurement unit mounted on the robot was mapped to the remote information gathered by a vision system. Angelova *et al.* [10] presented a similar strategy, where the measured slip is used to train a vision system to identify regions of potentially high slip, for a Martian rover. In all these cases, improved local terrain sensing by a tactile probe has the potential to improve the overall performance of the system.

The rest of this paper is divided as follows. In Section II, we first present a number of approaches that have been proposed for surface sensing using tactile probes. In Section III, we present the design of our tactile probe, along with the features extracted for classification purposes. The discrimination capabilities of this tactile probe are evaluated in Section IV with an artificial neural network trained on a data set of ten terrains. Section V shows how the tactile information captured by our tactile probe can also be used in a context of unsupervised learning. Finally, we demonstrate in Section VI how this terrain identification probe can be used to guide a blind robot, based on surface identification.

II. PREVIOUS WORK

Tactile feedback has been considered in several forms by the robotics community, either using indirect measurements from effector performance (such as wheel slip or leg traction) or via direct measurement. Tactile feedback is a fundamental and ubiquitous navigation aid in the animal kingdom; whiskers

in animals have been known to be able to collect some information about surface texture or shape [11], [12]. They compensate for the lack of good vision in rats: the typical visual acuity of a rat is 1 cycle per degree, compared to 30 cycles per degree in humans [13]. In robotics, one of the first examples of touch sensors is Grey Walter's *tortoise* robot [14]. It relied partially on a switch to modify the behavior of the robot, avoiding obstacles in its quest for reward.

One possible application of tactile sensing is obstacle distance or surface shape estimation, using whiskers. Jung *et al.* [15] proposed distance estimation from whisker deflection measured by a potentiometer. Along the same line, Cowan *et al.* [16] used an artificial antenna made of FlexTM sensors to measure distances. In Russell [17], a flexible whisker was mounted on a Puma robotic arm, and the whisker deflections were measured using a potentiometer. Scholz *et al.* [18] and Clements *et al.* [19] employed a whisker attached to the shaft of an electrical motor to perform similar profile measurements. The contact point location on the whisker was estimated by analyzing the force measured at the base of the whisker, by a 6-axis load cell. In Russell *et al.* [20], an array of eight whiskers was used for object recognition by a mobile robot. The deflection of each of the metallic whiskers was measured by a potentiometer. The application of these sensors to terrain identification is limited, as they are not sensitive to surface texture, an important cue in terrain identification.

Closer to what we present in this paper is the use of whiskers for surface texture identification. This topic has been extensively explored in the scientific community, with a number of different designs proposed. For example, Schultz *et al.* [21] used strain gauges, located near the base of an array of whiskers, to pick up friction-induced vibrations. They explored the idea of using power spectral analysis on the strain gauge signals to differentiate between smooth and rough surfaces, although no formal classification results were produced. Their experiments, however, were limited to a small set of artificial surfaces: one smooth surface, and two surfaces with different ridge pattern densities.

Hipp *et al.* [22] present more extensive results regarding surface texture classification, for actuated whiskers. Whisker vibrations were picked up by magnetic sensors, excited by the motions of a small magnet fixed on the base of each whisker. The power spectra of the magnetic sensor signals were computed and a Fisher transform was applied to reduce the dimensionality of the measurements. Classification, using multidimensional Gaussian density estimators, achieved a success rate of 39% for eight different grades of sand paper. Results over natural surfaces or terrains were unspecified.

Fend *et al.* [23] used a microphone as transducer to record the vibrations induced in a genuine rat whisker. A servo motor swept an array of eight whiskers, with each whisker being attached on a different microphone. Experiments were conducted over 11 surfaces: seven pieces of sand paper with different grain sizes, foam, mouse pad surface, carpet, and metal. The power spectra of individual sweeps were computed, smoothed, and combined together to generate an average power spectrum. By comparing the Euclidean distances between these power curves, they concluded that texture identi-

fication could be improved by using all whiskers at the same time, and by increasing the number of sweeps. The design of this whisker sensing system might be too fragile in the context of mobile robots in outdoor environments.

Surface texture sensing has also been developed for finger-like devices than can be used with manipulators. Howe *et al.* [24] created an artificial finger, made of a rubber skin mounted on soft foam. A small accelerometer was attached to the inside of the rubber skin, in contact with an object. This device was able to detect slip conditions, as well as provide some form of surface texture sensing. More recently, de Boissieu *et al.* [25] created another finger-like device, by encasing a three-axial MEMS force sensor in a simulated finger of epoxy, covered with a hard rubber skin. Their device exhibited excellent sensitivity to friction forces, which depended on the surface type. This high sensitivity made it possible to differentiate between printed and non-printed areas on a sheet of paper, for dimensions as small as 1 millimeter. Classification experiments were performed on a database of 10 types of paper, with a success rate around 70%, an impressive result considering the small differences in surface textures. The authors noticed, however, significant wear for the rubber tip. This lack of physical robustness reduces the applicability for outdoor surface exploration.

Roy *et al.* [26] presented a case where the tactile probe is not rubbed, but instead tapped against a surface. A small microphone, mounted at the tip of an aluminum boom, recorded the acoustic response of the impact. Classification was achieved by comparing the windowed power spectrum of the recorded impact to a collection of labeled prototype responses. Six surfaces were probed: wood, plastic, cement, metal sheeting (box), solid metal (I-beam), and glass. The classification success rate over this set was 95%, but it is not clear if this success rate would be as high for similar indoor surfaces, such as linoleum and terazzo. Moreover, this approach cannot be used to differentiate surface textures, an important cue in identification.

Seth *et al.* [27] discussed how the behavior of a wheeled robot could be modified by whisker-based sensing. The whiskers were made by gluing two polyamide strips back-to-back, and deformation was measured by sensing the strips' resistivity changes. The robot explored a walled environment, where two strikingly different peg arrangements (vertical or diagonal) were present. Over time, they showed that their simulated neural architecture learned to avoid the peg arrangement associated with a negative reward. The *Whiskerbot* robot (Pearson *et al.* [28]) also served as a demonstration platform for touch-based behavior adaptation. It used whisker arrays, actuated by shape-alloy metal fibers. Strain gauges were used to measure forces, and a simulated Primary Afferent neural network was stimulated by the amplified outputs of the strain gauges. The main goal of their project was to demonstrate a possible neural mechanism by which rodents analyze and interpret whisker vibration patterns; no significant surface classification results were provided.

Some of the work presented in this paper has been introduced previously in Giguere *et al.* [29]. In particular, part of the tactile probe design itself in Section III, the classification

results with an artificial neural network in Section IV and the blind navigation task in Section VI were presented in [29]. The rest of this paper is, however, original contribution.

III. TACTILE PROBE DESIGN

A. Mechanical Design

Our tactile probe was designed to gather as much relevant information as possible, in terms of terrain identification. It also had to be robust, so that it could be readily deployed in the field. The first kind of information that we were interested in collecting was related to the surface profile height $h(x)$. A common approach to modeling surface height $h(x)$ is through a random process, with an associated power spectral density $H(\omega_x)$ [30][31]. This approach has been used with success in civil engineering for the description of road surfaces. Moreover, the majority of terrain identification methods for wheeled vehicles are based on analyzing the power density spectrum of the vibrations of the chassis [1][2][3][4][5], which are directly related to $H(\omega_x)$. The probe had to be sensitive to the mechanical properties of surfaces, such as viscoelasticity. By being sensitive to these properties, our probe had a better chance of being able to distinguish surfaces or terrains of similar profile $H(\omega_x)$, but made of entirely different materials.

The probe design is depicted in Fig 2. The body of the tactile probe was made of solid aluminum, a stiff material, with a Young's modulus of $E_Y = 70$ GPa. This material has been selected so to be harder than most common ground surfaces, except for concrete and asphalt. A simple physical interpretation helps explain this particular choice of material. For instance, it would be difficult to differentiate between a hard and a soft surface if one uses a very flexible probe, since they would both "feel" soft. In the same manner, one can better sense the damping factor of a surface by tapping on it with a hard object, as was done in Roy *et al.* [26].

A similar conclusion regarding the importance of using stiff materials for tactile probes was reached by Lungarella *et al.* [32]. They compared the frequency responses of sensors made of different types of whiskers (polyvinyl fiber, rat whisker and human hair) when rubbed against a moving surface. They noted that the amplitude of vibration of a whisker was monotonically related to its stiffness, with the stiffer whisker (polyvinyl fiber) having vibrations with 17 times more energy than the softest whisker (human hair). This higher level of vibration serves as an indicator of the greater sensitivity of stiff probes.

The top end of the tactile probe was attached to the vehicle, using flexible rubberized fabric. With this fabric attachment, the probe was able to rotate in all three axes, for small angles. Moreover, this attachment reduced the mechanical coupling between the vehicle and the probe. This avoided sending parasitic vibrations generated within the vehicle itself to the rod. A single axis hinge mechanical design, by comparison, would have constrained the rod rotation in pitching motion, and increased unwanted mechanical coupling between the probe and the vehicle. The location of the attachment on the vehicle was selected to maintain the probe at an angle around $\theta_r = 45$ deg, relative to the surface.

Passive tactile probes collect information about surfaces typically by being excited through some sort of forcing action. This action is generally a rubbing or dragging motion over the targeted surface, in a direction parallel to it. This action results in a combination of motion of the contact point perpendicular to the surface, and horizontal motion x . If a surface profile $h(x)$ has large variations, the probe tip motion, as a function of probe location x , will also present large variations. The probe tip motion strongly depends on its shape: a smaller tip is better at tracking fine surface details than a bigger tip. The tip radius effectively acts as a low-pass spatial filter. For this reason, a wheel makes for a poor probing mechanism, compared with a tactile probe with a fine tip. This will naturally limit the terrain identification capacities for wheel-based vehicles relying on inertial sensors mounted on the chassis [1][2][4][5].

B. Choice of Transducer

An important design aspect of any sensing apparatus is the appropriate choice of transducer technology. A transducer is a device that converts one form of energy or physical property into another form, such as voltage, for the purpose of sensing. Ideally, a transducer:

- does not disturb the physical quantity it is measuring,
- has high sensitivity,
- has good resolution,
- has low noise,
- is accurate,
- has a high bandwidth.

The characteristics and limitations of a transducer used with a tactile probe determine its range of applications. For the purpose of surface identification, one of the most important factors is the frequency response of the transducer. When a tactile probe is rubbed against a surface, the surface spatial spectrum $H(\omega_x)$ gets partially captured by the tip of the tactile probe. The forward velocity v_c of the probe transforms a spatial spectrum density $H(\omega_x)$ into a temporal spectral density of vibrations. In order to capture high spatial frequencies in $H(\omega_x)$, corresponding to fine surface texture or features such as cracks, the transducer must possess a large bandwidth to capture the high frequency vibrations. The frequency response $E(\omega_t)$ of the transducer is particularly critical if the surface profiles power spectrum $H(\omega_x)$ differ mostly in the region of high spatial frequencies ω_x .

The relationship between transducer frequency response $E(\omega_t)$ and target application of tactile probes can be seen in many whiskers-based sensing projects. For cases where the target application is to only measure distances, a variety of low-frequency transducers (< 10 Hz) have been employed to measure whisker deformation. Electrical contacts, such as in Walter [14] or Russel [33], certainly represent the simplest type. More sophisticated schemes used potentiometers to measure springy whisker deflections (Jung *et al.* [15], Russel [17]).

For cases where tactile probes have been designed for surface identification purposes, a number of large bandwidth transducers have been proposed. Strain gauges (Pearson *et al.* [28], Schultz *et al.* [21]), load cells (Scholz *et al.* [18], Clements *et al.* [19]), Hall-Effect sensors (DaeEun *et al.* [34]),

or miniature magnetic sensors (Hipp *et al.* [22]) have reasonably large bandwidths (100 to 1000 Hz). These transducers can also be employed to perform whisker deflection measurements, since their frequency responses extend down to 0 Hz. Microphones, having a band-pass frequency response with a low cutoff frequency around 10-30 Hz, cannot be used to perform deflection measurements. Their high cutoff frequency, in the tens of kHz, means on the other hand that they can be used for surface texture identification. These microphones have been used, for example, to pick up vibrations generated when a whisker is being rubbed on a surface (Fend *et al.* [23]), or to record the acoustic response of surfaces impacted by a probe (Roy *et al.* [26], Krotkov *et al.* [35].)

In wheeled vehicles, the transducer most frequently used for terrain identification purposes is an on-board accelerometer [1],[2],[5],[4],[8]. This type of transducer has the significant advantage of being more sensitive to rapid changes in motion than any of the transducers listed above in whisker applications. Indeed, the frequency response to probe motion for an ideal accelerometer is proportional to the square of the temporal frequency of motion:

$$E_{accel}(\omega_t) \propto \omega_t^2 \quad (1)$$

The sensitivity of an ideal accelerometer is, however, low for slow variations of position i.e. low frequencies ω_t . At first glance, this sensor would be a poor choice to measure deflection, as the latter corresponds to the lower part of surface spatial spectrum $H(\omega_x)$.

By orienting the transducer and designing the tactile probe properly, the ability to measure slowly changing probe deflections can be recovered. This is accomplished by having a change of the height of the probe tip inducing a change in the orientation of the accelerometer, with respect to the gravity vector \vec{g} . Assuming that the tactile probe is oriented at an angle of $\theta_r = 45^\circ$ deg with respect to the surface, and that the gravity vector \vec{g} is perpendicular to the surface, the frequency response $E_{tilt}(\omega_t)$ of an accelerometer placed on the tip of this probe, as a function of change in tip height is:

$$E_{tilt}(\omega_t) = \frac{g}{l} + \frac{\sqrt{2}}{2} \omega_t^2 \quad (2)$$

With such a design, slow changes in variation of the height of the probe can be measured. This is because the frequency response $E_{tilt}(\omega_t)$, as seen in Fig. 1, extends all the way down to 0 Hz. This leads to the probe configuration shown in Fig. 2. In this design, the upper end of the tactile probe is attached to the vehicle or robot. For terrain sensing purposes, the height of the top-end of the probe is assumed to be fixed. This is due to the support system of the vehicle (wheels or legs), the mass of the robot, and possibly a suspension system, that average out rapid changes in surface height $h(x)$.

The transducer we used on our probe was a solid-state, single axis accelerometer model ADXL05JH from Analog Devices. It had a 5 milli-g resolution and a maximum range of $\pm 5g$. Its -3 dB cutoff frequency was set to 4 kHz using an external capacitor of 0.01 μF , and the attenuation beyond this frequency is -20 dB/decade. The accelerometer was placed near the tip of the probe, as to capture as much vibration




Fig. 1. Frequency response $E_{tilt}(\omega_t)$ described in Eq. 2 of an accelerometer located at the tip of an inclined tactile probe of length $l = 0.1$.




Fig. 2. Design of the tactile probe with an accelerometer. Given its direction of sensing (shown as arrow in the image) and the probe design, the accelerometer is capable of sensing both the angle of the probe θ_r and the acceleration of the tip of the probe. A close-up view of the flexible attachment is shown on the right.

produced at the contact point. The accelerometer sensitivity orientation was in the pitching direction (sagittal plane of the robot), corresponding to upward and downward motions of the probe.

This design was mechanically passive, since the mechanical energy used to probe the surface was provided by the forward motion of the robot. By contrast, active whiskers [17][18][19][22] employed an extra actuator to perform motion. Our tactile probe was also low-power, only requiring a few hundred of mW to power the accelerometer. Given its simplicity and lack of mobile parts, our design is robust, and easy to be made waterproof. Some drawbacks can be associated with this design. The main one is that no information about the surface can be collected when the robot is immobile. Also, the location of the probe is entirely dependent on the position of the robot. Consequently, we cannot perform surface identification independently from the robot position. A lesser problem is that the mechanical design of the rod attachment has to be able to accommodate for backward motions of the vehicle, which is accomplished by allowing the rod to rotate in multiple directions via the flexible attachment.

C. Selected Features for Classification

A significant number of surface identification techniques have relied on features extracted from power spectra, for analyzing vehicle accelerometer data [1][2][3][5], visual servoing errors in a camera [7], or tactile probe vibration patterns [21][22][23][25][26]. Spectral analysis, however, ignores important cues contained in the phase spectrum. For example, signals a) and c) in Fig. 3 have identical amplitude spectra but different phase spectra: the second signal was generated by randomly assigning phase values to the previous one, and

performing the inverse Fourier transform. One can see that these signals are very different in the temporal domain. In particular, the sharp spikes visible in Fig. 3 a) are no longer present in Fig. 3 c). For this reason, we extracted some of the features in time domain, in order to improve the robustness of classification.

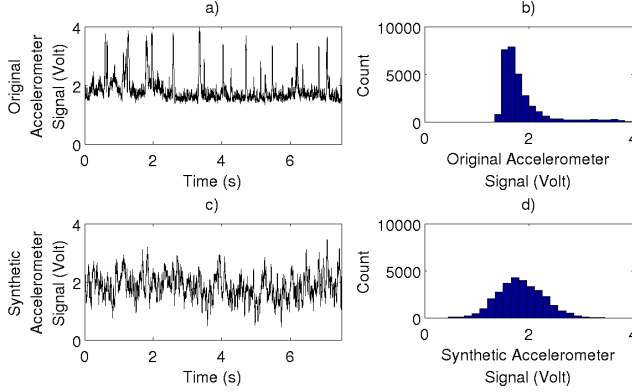


Fig. 3. Two signals a) and c) with identical amplitude spectra but different phase spectra. Signal in a) was captured by dragging the tactile probe over tiled linoleum, while signal in c) was generated by replacing phase spectrum of a) with random phase values. Important features in time-domain that are contained in the phase spectrum are lost: the histograms b) and d) are significantly different. This points to limitations for terrain identification methods based solely on frequency amplitude spectrum.

We employed a set $f_{tactile}$ of eight features extracted from the accelerometer signal $a(t)$ in the time and frequency domains. These features were extracted from time-windows of size W , as was done similarly in Weiss *et al.* [4]. The seven features in the time domain were:

- mean,
- variance,
- skewness,
- kurtosis,
- fifth moment,
- number of times 20 uniformly separated thresholds are crossed,
- sum of the variation over time:

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

and one was in frequency domain:

- sum of higher half of amplitude spectrum.

Overall, the variance of the signal was a good indicator of the amount of vertical motion experienced by the rod, which was typically large for uneven surfaces, such as grass and gravel. The skewness of the histogram helped identify cases with asymmetric distribution of accelerations $a(t)$, seen frequently for surfaces with regular but infrequent asperities, e.g. cracks between tiles. The sign of the skewness played an important role too, as some terrains tended to have a distribution of accelerations $a(t)$ more skewed to the top than the bottom. This asymmetry is visible in Fig. 4 between tiled linoleum a) and *Interblock* e). The presence of high-frequency components helped differentiate hard surfaces from others,

since a stiff surface is required to sustain high-frequency vibrations in the tactile probe.

IV. SUPERVISED LEARNING ON A DATA SET OF TEN COMMON SURFACES

We evaluated the performance of our tactile probe and feature extraction for surface identification. An artificial neural network was selected as a classifier. It was trained to identify various terrain surfaces in a data set $\vec{X}_{tactile}$, described below. The terrain identification performance was equated to the testing error rate of this classifier. More advanced classifier techniques could have been employed, such as Support Vector Machine [36], but the immediate goal of this evaluation was to determine the quality of the information returned by our sensor, and not the classification power of the classifier itself.

A. Data Collection

We attached the tactile probe to a cardboard box of dimension 12 cm x 12 cm x 30 cm. We manually dragged this box using a string, at an approximate forward velocity of 70 cm/s, in order to simulate a low-velocity land rover. Ten different surfaces were sampled during the experiments, over the course of a few days. Six of them were indoor surfaces, and four were outdoor terrains:

- wood bench,
- tiled linoleum flooring,
- untiled linoleum flooring,
- terazzo,
- short hair carpet,
- large tiles for raised computer floors,
- grass,
- paving made of concrete bricks similar to *Interblock*,
- large gravel, and
- small gravel on packed dirt.

Fig. 4 shows these ten surfaces, along with a corresponding sample of the tactile probe acceleration signal. The accelerometer signal was sampled at 4 kHz, with 11-bit resolution, using an IsaacTM stand alone data acquisition system. The fact that the sampling frequency was not twice the cut-off frequency of the accelerometer might have induced some aliasing. However, the good classification results presented later in this Section indicate that the actual measured signals were not significantly aliased. The captured signals were manually spliced, to keep samples captured during a steady motion of the device. The samples were manually labelled, and in total between 250 and 500 seconds of actual data was available for each surface, depending on the amount of data originally captured. This constituted our training data set $\vec{X}_{tactile}$.

B. Classifier

We employed the *NETLAB* [37] library implementation of an Artificial Neural Network as a classifier. The neural network had eight inputs (one for each feature), ten outputs (representing each of the ten terrains), and 20 neurons for each of the two layers. The output activation function was chosen to be linear. These classifier parameters were chosen empirically

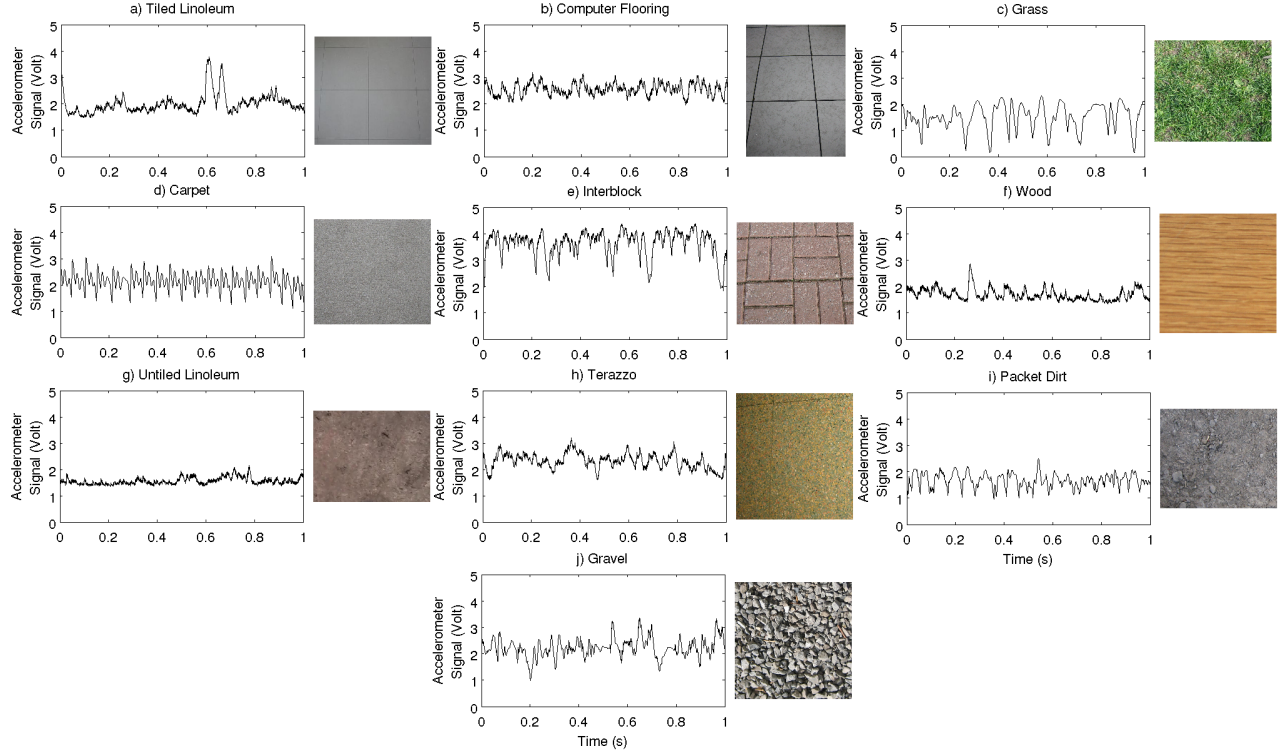


Fig. 4. Picture of the ten surfaces, along with an associated probe signal, used in evaluating the performance of our probe. The probe signals collected over these ten surfaces formed the labelled data set $\tilde{X}_{tactile}$.

TABLE I
AVERAGE TESTING CLASSIFICATION RATE FOR A NEURAL NETWORK
TRAINED ON THE TEN TERRAINS IN THE $\tilde{X}_{tactile}$ DATA SET.

	Time-Window Size 1 s	Time-Window Size 4 s
Indoor and Outdoor	$89.9 \pm 0.4\%$	$94.6 \pm 0.6\%$
Outdoor only	$98.85 \pm 0.08\%$	$99.96 \pm 0.02\%$

by trial and error, and are unlikely to be optimal. Given this sub-optimality, the classification results we achieved here represent a lower bound on the actual discrimination power of the probe itself, for the dataset of surfaces we have sampled.

The performance of this classifier was evaluated through repeated random sub-sampling validation. During this procedure, the data set was split randomly with 70% of the data going for training and 30% used for testing. Results were averaged across all splits. In total, 25 random trials were generated per window size W . Training was accomplished by using a scaled conjugate gradient method, with 2,000 iterations.

C. Classifier Testing Results

The testing success rate over all ten surfaces (Table I) was $89.9 \pm 0.4\%$ and $94.6 \pm 0.6\%$ for time windows of 1 and 4 seconds. These results indicate that our sensor has good discrimination capabilities for this set of surfaces. It compares favourably against Larson *et al.* [7], which reported a success rate of 90.0% for 5 simulated outdoor terrains. Fig. 5 shows the confusion matrix for time windows of 1 second, and Fig. 6 for time windows of 4 seconds, averaged over the 25 random trials. When only outdoor terrains were considered in the testing phase for the neural networks trained on all

ten surfaces, success rate climbed to 98.85% and 99.96% for time windows of 1 and 4 seconds, respectively. This result was expected, as the sampled outdoor surfaces were quite different from one another. By contrast, the sampled indoor surfaces were smooth with few distinct features. For example, the main difference between tiled and untiled linoleum was the presence of small cracks. As a side note, the carpet surface was sometimes confused with the small gravel outdoor terrain. This confusion can probably be explained by the fact that they both were irregular soft surfaces.

We did not directly test the behavior of the trained neural network when samples from unknown terrains were being classified. Instead, we present in Section V a method to train the system with unlabelled data sets. This unsupervised method could then be used to partially alleviate the problems associated with novel terrains.

V. UNSUPERVISED LEARNING FOR TERRAIN IDENTIFICATION WITH TACTILE PROBE DATA

In most terrain identification strategies for ground vehicles, supervised learning methods have been employed to train the classifiers [1][2][3][4][5]. These algorithms require, by definition, a manually labelled training data set. Moreover, these systems are static in terms of knowledge acquisition; they cannot incorporate new knowledge as they explore unknown terrains. Consequently, employing unsupervised learning (clustering) techniques for terrain identification is an attractive option to increase the autonomous capabilities of such vehicles, as the labelling of data samples is no longer required. It is generally more challenging, however, to perform unsupervised learning than supervised learning. In particular, two key aspects can

Actual	Grass	100.0	0	0	0	0	0	0	0	0	0
	Paving	0	99.9	0	0	0	0	0	0	0.1	0
	Large Gravel	0.2	0	99.3	0.2	0	0.2	0	0	0.2	0
	Dirt with Small Gravel	0	0	0.6	94.5	0	4.6	0	0	0.3	0
	Computer Room	0	0.1	0	0	93.6	0	0	0	2.4	3.9
	Carpet	0.5	0	0.4	6.9	0	92.1	0.1	0	0	0
	Wood	0	0	0	0.1	0.1	0.1	84.6	9.6	2.2	3.2
	Untiled Linoleum	0	0	0	0	0	0.1	16.8	80.1	0.6	2.4
	Terazzo	0	0.5	0	0.3	1.2	0.8	12.2	1.1	79.2	4.5
	Tiled Linoleum	0	0.1	0	0.1	7.5	0.1	9.4	2.0	10.7	70.0
		Gr	Pa	La	Di	Co	Ca	Wo	Un	Te	Ti
		Predicted									

Fig. 5. Confusion matrix for testing phase of the trained neural network on the $\vec{X}_{tactile}$ data set. Time-window size was 1 seconds. Results shown as percentage, averaged over 25 randomly training sets. Entries less than 0.05% are rounded down to 0.

Actual	Grass	100.0	0	0	0	0	0	0	0	0	0
	Large Gravel	0	100.0	0	0	0	0	0	0	0	0
	Dirt with Small Gravel	0	0	100.0	0	0	0	0	0	0	0
	Paving	0	0	0	99.9	0	0	0	0	0.1	0
	Computer Room	0	0	0	0	99.8	0	0	0	0.2	0
	Carpet	0	0	1.6	0	0	98.4	0	0	0	0
	Wood	0	0	0	0	0	0	92.4	6.0	0	1.5
	Untiled Linoleum	0	0	0	0	0	0	8.9	87.6	0	3.5
	Terazzo	0	0	0	0	0.2	0	5.4	0	86.5	7.9
	Tiled Linoleum	0	0	0.2	0	3.3	0	4.8	4.8	8.3	78.5
		Gr	La	Di	Pa	Co	Ca	Wo	Un	Te	Ti
		Predicted									

Fig. 6. Confusion matrix for testing phase of the trained neural network on the $\vec{X}_{tactile}$ data set. Time-window size was 4 seconds. Results shown as percentage, averaged over 25 randomly training sets. Entries less than 0.05% are rounded down to 0.

potentially impact the performance of unsupervised learning systems: the quality of the information collected by the sensors and the clustering algorithm itself.

Srebro *et al.* [38] empirically demonstrated that the performance of some clustering techniques depends on how distinct the data clusters are, in feature space. More precisely, they have shown that increasing the relative distance $\frac{|\mu_a - \mu_b|}{\sigma_a}$ between normally distributed classes with means μ_a, μ_b and identical standard deviation σ_a positively impact clustering algorithms. Our tactile sensor increases this relative distance, since the extracted features from the tactile probe measurements vary widely between terrain types, while being relatively stable for the same terrain. This increase in distance is due to the greater sensitivity of our tactile probe to terrain properties, compared to an accelerometer mounted on the chassis of a vehicle. For this reason, our probe is well suited for unsupervised learning of terrains by ground vehicles.

The other key aspect, the unsupervised learning technique

itself, was previously addressed in complementary work by ourselves [39]. In that work, we presented an unsupervised learning technique that explicitly used temporal coherence in sensor measurements, to improve robustness of clustering. For a mobile robot traversing a number of terrains, a temporal coherence is generally present in sensor measurements affected by the terrain's mechanical properties. This is due to the fact that spatial coherence is generally present in terrains, even in unstructured environments. For example, beaches in littoral settings are narrow patches of sand with a minimum width of a few *meters*. Environments adapted by humans (sidewalks, grassed areas and roads) are patchy by design, occupying areas rarely smaller than a few square *meters*. Briefly speaking, this unsupervised learning technique segments time-series of sensor measurements, using a classifier trained in an unsupervised manner. This training is performed by finding the classifier parameter values $\vec{\theta}$ that minimize an objective function related to the time-variation of the posterior probabilities $p(c_i|\vec{x}, \vec{\theta})$ of samples $x \in X$, for all classes $c_i \in C$:

$$\arg \min_{\vec{\theta}} \sum_{i=1}^{|C|} \frac{\sum_{t=1}^{T-1} (p(c_i|\vec{x}_{t+1}, \vec{\theta}) - p(c_i|\vec{x}_t, \vec{\theta}))^2}{\text{var}(p(c_i|\vec{X}, \vec{\theta}))^2} \quad (3)$$

This unsupervised learning technique is known to perform well for cases where the temporal coherence is weak but still present, e.g. spanning 3 consecutive samples such as \vec{x}_{t-1}, \vec{x}_t and \vec{x}_{t+1} .

Weiss *et al.* [40] also introduced an unsupervised online technique for terrain identification, but based on novelty detection and an implicit assumption of spatial coherence. New terrain types in the sensor measurements were discovered by identifying outliers of a trained Gaussian mixture model on known terrains. Data points considered as novel were collected until a sufficient amount was reached, upon which a new Gaussian in the novelty detector was trained on this accumulated set. The same new data set was used, also, to train a new k -Nearest Neighbors classifier. Data not judged as novel were classified using the already existing k -Nearest Neighbors classifiers. The fact that a novel terrain had to be present for a sufficiently long period of time before another novel terrain could have been admitted indicates an implicit assumption of spatial coherence in their algorithm. For example, 20 novel samples had to be accumulated before training the new Gaussian novelty detector and k -Nearest Neighbors classifier. If data coming from another unseen terrain were captured during this accumulation period, then these novel terrains risk being grouped together.

A. Unsupervised Learning of $\vec{X}_{tactile}$ Data Set

We evaluated the performance of combining the information gathered by our tactile probe with the unsupervised learning algorithm described by Eq. 3, over synthetic time-series \vec{A}_i generated from the tactile probe data set $\vec{X}_{tactile}$. Tests were done for various number $u \in \{2, 3, 4, 5, 7, 10\}$ of terrains. For each value of u , we tested for all possible combinations of terrains in our data set $\vec{X}_{tactile}$:

$$C_u^{10} = \binom{10}{u}$$

Temporal coherence in the time-series \vec{A}_i was simulated by having an average duration within a single terrain of 5 *seconds*. To avoid having the transitions aligned with temporal windows w_t , this duration was normally distributed with a standard deviation of 2 *seconds*. Finally, transitions between the u selected terrains had equal likelihood. The total duration of the time-series \vec{A}_i was variable, depending on the number of terrains u and the number of data samples available for the selected terrains. An example of a time-series \vec{A}_i used for testing is shown in Fig. 7.




Fig. 7. a) Subsequence of one of the time-series \vec{A}_i used in the unsupervised learning of tactile probe data, generated from $u = 3$ different terrains. b) Ground truth displaying the terrain transition.

Each synthetic time-series \vec{A}_i was divided in non-overlapping time-windows w_t of 1 *second*. With $|w_t| = 4000$ samples, a time-series \vec{A}_i contained

$$n_w = \left\lfloor \frac{|\vec{A}_i|}{4000} \right\rfloor$$

time-windows, and the same eight features $f_{tactile}$ presented in Section III-C were extracted from each w_t . This resulted in a 8-by- n_w training set \mathbf{B}_i . After mean centering, Principal Component Analysis was applied on \mathbf{B}_i to find the 8 eigenvectors. The data was projected along these eigenvectors to create \mathbf{Y}_i . All eight components were kept, as this was done mainly to re-align the data along the principal axis.

We employed spherical Gaussian models for the classifiers $p(c_i|\vec{x}_t, \vec{\theta})$ in Eq. 3, trained with five random restarts of simulated annealing, on each data set \mathbf{Y}_i . This classifier was selected, in order to reduce the number of parameters $\vec{\theta}$ of the classifier, compared to the neural network in Section IV. It enabled also a more direct comparison against a Hidden Markov Model (HMM), since the later was using spherical Gaussians as observation models. For reference purposes, three non-temporal clustering methods were also compared on our sensor data set $\vec{X}_{tactile}$: k-means, full covariance Gaussian Mixture Model, and spherical Gaussian Mixture Model, all trained with Expectation-Maximization.

The clustering success rates reported were for time-windows w_t that strictly contained one type of terrain. Clustering results for time-windows containing more than one terrain were not taken into account in the performance measure, but were still

present in the clustering process as they might negatively impact clustering. The optimal mapping between clusters and ground truth was found, and used in computing the average classification success rate of Fig. 8. In these tests, our approach (shown as Cost in legend) ranked first in a statistically significant manner for $u > 2$ terrains. This indicates that a strong emphasis on preserving temporal coherence does improve unsupervised learning of terrains, when employed with a tactile probe with good sensitivity. The same results also indicate that the performance of unsupervised learning over 10 terrains (74.1%) is in the same range as the in case of supervised learning (89.0%), when compared to a random classifier having a $\frac{1}{10} = 10\%$ chance of correctly identifying a terrain. We suspect that employing a more sophisticated classifier in Eq. 3, such as k -Nearest Neighbors, as well as having access to a training data set with more samples, should further close the gap between supervised and unsupervised learning of terrains.

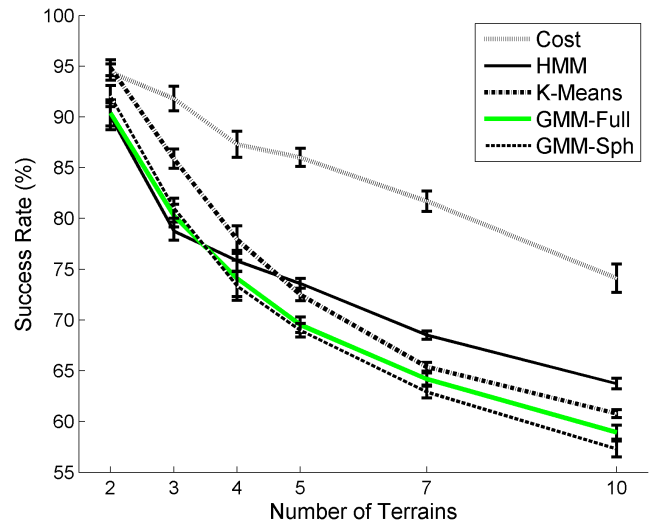


Fig. 8. Results of employing unsupervised learning techniques on the tactile probe data set. Error bars indicate 95% confidence interval.

VI. PRACTICAL APPLICATION: MOBILE ROBOT NAVIGATION BASED ON SURFACE TEXTURE

We explored the concept of blind navigation for mobile robots, using a tactile probe. This navigation was based on the real-time terrain identification of surfaces, using a trained classifier and our tactile probe mounted on a mobile robot. This was analogous to an insect using its antenna to avoid certain areas that have been previously identified as dangerous. The navigation task we selected consisted in having a differential-drive robot navigating a bounded area S_{in} , near its edge. The surface of S_{in} had a distinct surface texture, compared with the surrounding surface S_{out} . Using our tactile probe, the mobile robot was able to detect when it left or re-entered the surface S_{in} . A simple control strategy relying on terrain identification was designed and tested in simulation. Results were validated using a real differential-drive robot, with a medium-sized carpet as the target surface S_{in} , placed on top of a tiled floor representing surface S_{out} .

This control algorithm worked as follows. When the robot was on S_{in} , it went forward at a velocity $V_{nominal}$. When the tactile probe detected a transition to S_{out} , the robot reduced its speed by half, and started turning towards S_{in} . Once the probe located on S_{out} re-entered the target surface S_{in} , it drove straight for a brief moment, in order to move away from the border. It then rotated in the opposite direction, for half the duration it took to bring the outer sensor back on the target area S_{in} . This counter-rotation maneuver was done to orient the robot parallel to the edge of the surface S_{in} .

For this experiment, a smaller version of the tactile probe was built and fitted on a differential-drive wheeled robot. Fig. 9 shows this probe mounted on the left hand side of the iRobotTM CreateTM robot used during this experiment. The transducer signal of the probe was captured using a standard audio input on a 3.0 GHz laptop. The same laptop performed signal classification and executed the control code. Control commands, based on the algorithm described above, were sent to the robot via a serial interface built in the CreateTM robot. The target surface S_{in} to patrol consisted in a 3x3 m carpet placed on top of a tiled floor in a computer lab. We employed the unsupervised learning technique described in Section V, so that training could be achieved on-the-spot without the need to label the data manually. Reusing the train neural network of Section IV was not an option: the body of the tactile probe used in this experiment was different than from the one used to collect $\vec{X}_{tactile}$, and the measurements by a tactile probe depend on its mechanical properties.

Qualitatively the robot performed very well. Five test runs were executed, with an average duration of five minutes per run. Fig. 9 presents a time-lapse series of pictures taken while the robot was executing a re-entry maneuver. In total, the robot exited the carpet area 80 times, and successfully managed to re-enter this surface every time but once. Two main reasons explain this high success rate. First, the tactile probe mounted on the robot was able to reliably differentiate the carpet surface from the computer room flooring. Indeed, the entries in the confusion matrix of Fig. 5 corresponding to carpet and computer room flooring shows no (0) confusion for the $\vec{X}_{tactile}$ data set, thus indicating that these surfaces are reliably identified by the tactile probe. Second, the generally good mobility performance of the CreateTM robot ensured that the sent velocity and rotation commands were well executed.

VII. DISCUSSION AND FUTURE WORK

In this paper, we presented a novel tactile probe for surface and terrain identification, for low-speed mobile robots. The probe design consisted of a metallic rod, with a solid-state accelerometer attached to its tip. Terrain classification was based on eight features extracted from the accelerometer's measurements, in the time and frequency domains. This was in contrast with most prior terrain identification techniques, which usually rely on features extracted in the frequency domain only. The good discrimination capabilities of our tactile probe were established using a non-sophisticated classifier, a neural network, on a data set captured over 10 different indoor and outdoor surfaces. In particular, our probe was




Fig. 9. Series of frames taken during the carpet and computer room flooring experiments with a CreateTM robot. These frames show one of the re-entry maneuvers executed by the robot, after leaving the carpet. Frame sequence is left-to-right, top-to-bottom. After re-entry, the robot oriented itself parallel to the carpet edge, using information gathered during the exit and re-entry.

sensitive enough to be able to differentiate between tiled and untiled linoleum, something that cannot be achieved with an accelerometer mounted on a wheeled vehicle. We, also, demonstrated how the information captured by this tactile probe can be used in the context of unsupervised learning of terrains.

A demonstration for a novel application for tactile probes in mobile robotics was also presented. It consisted in having a mobile robot navigating an area that could be differentiated from its surroundings, using surface texture information captured by our tactile probe. As such, this experiment represented a *weak* form of localization, based on tactile information.

This work might be improved in a number of manners. For example, we are currently assessing the impact of the tactile probe material (wood, plastic, glass and steel) on the sensitivity. The current probe design is using a single-axis accelerometer. Using a 3-axis accelerometer is a natural evolution, and it might be able to capture information about stick and slip behaviors. These have been shown in [25] to be important cues in surface identification. Moreover, the experiments have been conducted at a single velocity (70 cm/s). It is important, therefore, to further investigate how we can perform surface identification at other velocities, possibly employing a technique similar to [2]. In particular, we expect the quality of the information to gradually decrease at lower velocities, in part due to the fact that a lower velocity implies that less surface is probed, per unit of time. Finally, we are looking into how this rich tactile information can be exploited in the context of self-training of a vision system on a mobile robot, similarly to [9]. In particular, we intend to include the concept of curiosity, with the robot navigating towards unknown visual textures, in order to probe them.

REFERENCES

- [1] D. Sadhukhan. Autonomous ground vehicle terrain classification using internal sensors. Master's thesis, Florida State University, 2004.
- [2] E.M. DuPont, R.G. Roberts, and C.A. Moore. Speed independent terrain classification. In *Proceeding of the Thirty-Eighth Southeastern Symposium on System Theory*, pages 240–244, March 2006.
- [3] Chris A. Brooks and Karl Iagnemma. Vibration-based terrain classification for planetary exploration rovers. *IEEE Transactions on Robotics*, 21(6):1185–1191, December 2005.
- [4] Christian Weiss, Holger Frohlich, and Andreas Zell. Vibration-based terrain classification using support vector machines. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4429–4434, Oct. 2006.
- [5] Edmond M. Dupont, Carl A. Moore, Jr. Emmanuel G. Collins, and Eric Coyle. Frequency response method for terrain classification in autonomous ground vehicles. *Autonomous Robots*, 24(4):337–347, 2008.
- [6] P. Giguere, G. Dudek, C. Prahacs, and S. Saunderson. Environment identification for a running robot using inertial and actuator cues. In *Robotics: Science and Systems*, Philadelphia, USA, August 2006.
- [7] A.C. Larson, R.M. Voyles, and G.K. Demir. Terrain classification through weakly-structured vehicle/terrain interaction. In *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, 2004.
- [8] Christian Weiss, N. Fechner, M. Stark, and A. Zell. Comparison of different approaches to vibration-based terrain classification. In *Proceedings of the 3rd European Conference on Mobile Robots (ECMR 2007)*, pages 7–12, September 2007.
- [9] Ambroise Krebs, Cédric Pradalier, and Roland Siegwart. Adaptive rover behavior based on online empirical evaluation: Rover–terrain interaction and near-to-far learning. *J. Field Robot.*, 27(2):158–180, 2010.
- [10] A. Angelova, L. Matthies, D. Helmick, and P. Perona. Slip prediction using visual information. In *Robotics: Science and Systems*, August 2006.
- [11] S. B. Vincent. The function of the vibrissae in the behavior of the white rat. *Behavior Monographs*, 1:1–81, 1912.
- [12] P. Leyhausen. *Cat Behavior*. Garland STMP Press, New York, 1979.
- [13] D. Birch and G. H. Jacobs. Spatial contrast sensitivity in albino and pigmented rats. *Vision Research*, 19:933–937, 1979.
- [14] W. Grey. An electromechanical animal. *Dialectica*, 4:42–49, 1950.
- [15] D. Jung and A. Zelinsky. Whisker based mobile robot navigation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 2, pages 497–504 vol.2, Nov 1996.
- [16] Noah J. Cowan, Emily J. Ma, Mark Cutkosky, and Robert J. Full. A biologically inspired passive antenna for steering control of a running robot. In *In International Symposium on Robotics Research*, pages 541–550. Springer, 2003.
- [17] R.A. Russell. Using tactile whiskers to measure surface contours. In *IEEE Transactions on Robotics and Automation*, 1992.
- [18] G.R. Scholz and C.D. Rahn. Profile sensing with an actuated whisker. *IEEE Transactions on Robotics and Automation*, 20(1):124–127, Feb. 2004.
- [19] T.N. Clements and C.D. Rahn. Three-dimensional contact imaging with an actuated whisker. *IEEE Transactions on Robotics*, 22(4):844–848, Aug. 2006.
- [20] R. Andrew Russell and Jaury Adi Wijaya. Recognising and manipulating objects using data from a whisker sensor array. *Robotica*, 23(5):653–664, 2005.
- [21] A.E. Schultz, J.H. Solomon, M.A. Peshkin, and M.J. Hartmann. Multifunctional whisker arrays for distance detection, terrain mapping, and object feature extraction. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2588–2593, April 2005.
- [22] J. Hipp, E. Arabzadeh, J. Conradt, E. Zorzin, C. Kayser, M. Diamond, and P. Knig. Texture signals in whisker vibrations. *Journal of Neurophysiology*, 95(3):1792–1799, 02 2006.
- [23] M. Fend, S. Bovet, H. Yokoi, and R. Pfeifer. An active artificial whisker array for texture discrimination. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 2, pages 1044–1049 vol.2, Oct. 2003.
- [24] R.D. Howe and M.R. Cutkosky. Sensing skin acceleration for slip and texture perception. In *Robotics and Automation, 1989. Proceedings., 1989 IEEE International Conference on*, pages 145–150 vol.1, May 1989.
- [25] F. de Boissieu, C. Godin, B. Guilhamat, D. David, C. Serviere, and D. Baudois. Tactile texture recognition with a 3-axial force MEMS integrated artificial finger. In *Proceedings of Robotics: Science and Systems*, Seattle, USA, June 2009.
- [26] Nicholas Roy, Gregory Dudek, and Paul Freedman. Surface sensing and classification for efficient mobile robot navigation. In *IEEE Transactions on Robotics and Automation*, 1996.
- [27] Anil K. Seth, Jeffrey L. McKinstry, Gerald M. Edelman, and Jeffrey L. Krichmar. Texture discrimination by an autonomous mobile brain-based device with whiskers. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4925–4930, 2004.
- [28] Martin J. Pearson, Anthony G. Pipe, Chris Melhuish, Ben Mitchinson, and Tony J. Prescott. Whiskerbot: A robotic active touch system modeled on the rat whisker sensory system. *Adaptive Behavior - Animals, Animats, Software Agents, Robots, Adaptive Systems*, 15(3):223–240, 2007.
- [29] P. Giguere and G. Dudek. Surface identification using simple contact dynamics for mobile robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, Kobe, Japan, May 2009.
- [30] Louis J. Cote, Frank Kozin, and John L. Bogdanoff. Introduction to a statistical theory of land locomotion–i. *Journal of Terramechanics*, 2(2):17–23, 1965.
- [31] C.J. Dodds and J.D. Robson. The description of road surface roughness. *Journal of Sound and Vibration*, (31):175–183, 1973.
- [32] M. Lungarella, V. V. Hafner, R. Pfeifer, and H. Yokoi. An artificial whisker sensors in robotics. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2931–2936, 2002.
- [33] R.A. Russell. Closing the sensor-computer-robot control loop. *Robotics Age*, 6(4):15–20, April 1984.
- [34] DaeEun Kim and Ralf Möller. Biomimetic whiskers for shape recognition. *Robotics and Autonomous Systems*, 55(3):229–243, 2007.
- [35] Eric Krotkov, Roberta Klatzky, and Nina Zumel. Robotic perception of material: Experiments with shape-invariant acoustic measures of material type. In *Experimental Robotics IV*, pages 204–211. Springer-Verlag, 1995.
- [36] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, September 1995.
- [37] I. T. Nabney. *NETLAB: algorithms for pattern recognition*. Springer-Verlag, 2002.
- [38] N. Srebro, G. Shakhnarovich, and S. Roweis. When is clustering hard? In *PASCAL Workshop on Statistics and Optimization of Clustering*, July 2005.
- [39] P. Giguere and G. Dudek. Clustering sensor data for autonomous terrain identification using time-dependency. *Autonomous Robots*, pages 171–186, March 2009.
- [40] C. Weiss and A. Zell. Novelty detection and online learning for vibration-based terrain classification. *Proceedings of the 10th International Conference on Intelligent Autonomous Systems (IAS)*, pages 16–25, July 2008.



Philippe Giguere received a BSc 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.



Gregory Dudek received the Ph.D. degree from the University of Toronto, Toronto, ON, Canada. He is currently a Professor in the School of Computer Science, McGill University, Montreal, QC, Canada, and a member of the McGill Research Centre for Intelligent Machines. He directs the McGill Mobile Robotics Laboratory. His current research interests include perception for mobile robotics, vision-based robotics, computational vision, and collaborative filtering.