

Surface Sensing and Classification for Efficient Mobile Robot Navigation

Nicholas Roy¹ Gregory Dudek¹
Paul Freedman²

¹Centre for Intelligent Machines,
School of Computer Science, McGill University,
3480 University St., Montreal, Québec, Canada H3A 2A7

² Centre de recherche informatique de Montral
1801 McGill College Avenue, Suite 800, Montréal Québec H3N 2N4

Abstract

Mobile robot navigation and localization is frequently aided by, or even dependent upon, a good estimate of the rate of dead-reckoning error accumulation. Sensor data can be used for position estimation, but this often involves overhead in acquiring and processing the data. By sensing and then classifying the surface type, an estimate of the rate of error accumulation for dead-reckoning allows us to estimate accurately how often localization, including sensor data acquisition, must be performed.

In the experiments we describe, a boom-mounted microphone is tapped on different floor materials, much as a blind man might tap his cane. The acoustic signature arising from the contact is then used to classify the floor type by comparing a windowed power spectrum of the acoustic signature with one of a family of prototypical signatures generated statistically from the same material. The technique is low-cost, involves limited computational expense, and performs very well.

1 Introduction

Mobile robot navigation and localization is frequently aided by, or even dependent upon, a good estimate of the rate of dead-reckoning error accumulation. Sensor data (in particular sonar or laser range data) can be used for position estimation but often entails delays in order to acquire the measurements or process the data. By observing the material type

Floor type	percentage error
tile	1%
rough floor	8%
carpet	11%

Table 1: Average positional error accumulated over a simple polyhedral path (ten tours of a square of approximately two feet by two feet). In general, the magnitude of the errors depends on various parameters of the trajectory but the relative magnitudes as a function of surface type vary consistently.

over which the robot is moving, an estimate of the rate of error accumulation for dead-reckoning allows us to accurately estimate how often localization, including sensor data acquisition, must be performed. For various floor coverings in our laboratory, for example, the rate of error accumulation varies by a factor of 10 (See Table 1). The system we propose uses dead-reckoning and knowledge of the material over which it is moving to maintain an estimate of its position and uncertainty in position. When this uncertainty exceeds a prescribed limit, re-localization is performed using either sonar or video data [5].

The approach presented in this paper use a simple sensory probing strategy to address the three main problems associated with the sensing-odometry combination. The probe can:

- detect obstacles for collision avoidance,

- detect holes or drop-offs (for example stairways or ramps),
- be used to perform floor surface identification while navigating (the emphasis of this paper).

This third task, in particular, is the most useful in tackling the positional error inherent in odometry. By identifying surface properties such as compliance (e.g. plastic or carpeted surfaces) which would lead to high odometric errors, or hardness (e.g. tiled surfaces) which would lead to lower odometric errors, an appropriate frequency of use for the more resource intensive perception system could be determined.

A boom-mounted microphone is mounted on the front of the robot and is tapped on the ground in front of the mobile robot as it moves, analogous to the way in which a blind person might use their cane. The acoustic/vibrational signature picked up by the microphone and arising from the contact is then used to classify the floor type. In previous studies, we have examined various strategies and control mechanisms and assemblies for performing the tapping [6]. Surface identification is accomplished by preceding the use of the on-line system by a surface learning phase where different samples are classified and prototypes for recognition are generated. A family of several prototypes may be generated for a single material. During the on-line stage, a real signal is acquired by tapping, and identification is performed by comparing a windowed power spectrum of the acoustic signature with the sets of families of prototypical signatures generated during the learning stage. The technique is low-cost and involves only limited computational expense.

The inspiration for the sensory probe presented in this paper is the specialized manner in which certain blind people use a cane easily to detect obstacles and drop-offs and surface height variations. Furthermore, using proprioceptive vibrations that travel up the cane from the impact of the tip against the floor surface, the user can often identify the type of material underfoot [8]. The experienced user can walk a line dividing two different floor materials, using the cane to identify each one and the border between them. It should be noted that this ability does not come simply from the sound of the cane hitting the floor; if the cane is insulated from the user's hand, so that he or she cannot feel the impact but only hear it, the ability to differentiate floor materials is severely impaired. On the other hand, if the sound of the impact is deadened (e.g. with earplugs), the user still has this ability, although accuracy is decreased. By mounting a microphone at the tip of a boom, and using the boom in the same

manner as a cane, a robot can (in principle) detect obstacles, holes, and do surface identification in the same manner as a blind person. This is the subject of the present article.

2 Background

Most work on robotic sensors falls into two categories: tactile sensors and position sensors [1]. Our work also uses what is essentially a tactile sensor, but it is unlike most existing tactile sensors described in the literature. Most work on tactile sensors has been conducted with the goal of identifying shapes by following their contours, just as a human hand can identify a shape using fingers. These sensors are largely feelers, and at best could be used for obstacle detection and drop-off detection, but they do not have the sophistication for surface identification. With the development of slippage detectors, and possibly the implementation of a much finer grain of sensor within a large feeler, surface identification could be similar to that of a human hand, but for this much more development is needed.

Krotkov [3] has also recently examined generic surface identification (without motivation from mobile robot navigation) using acoustic signatures in a different framework and with a somewhat different approach than our own. His reported identification rates are slightly lower than those reported here.

Surface identification has been explored to a limited degree within the context of quality control, but the focus is primarily on detecting atypical samples, rather than on classification. These systems rely primarily on defects that change the acoustic or electromagnetic reflectivity of the surface, and as such are able to identify simply everything that is *not* a particular surface. Visual systems have been developed to identify many surfaces with high success rates eg. [9] but there exist floor surfaces which they cannot identify with any notable success rate. Furthermore, these systems are often costly in terms of equipment and processing power. The sensor we describe here has the advantages of being cheap and fast.

3 System Hardware

Our work in this area began with a small condenser-type microphone attached to the tip of a boom (flexible rod) mounted on an RWI B-12 mobile robot equipped with radio antenna controlled from a Sun

workstation, with a separate onboard HC11 micro-controller for controlling the tapping pattern of the boom in front of the robot, much as a blind man would tap his cane [6].

In subsequent work [4], we chose to concentrate on a new mechanical system to support the boom. The results reported here were obtained using a commercial pan and tilt unit connected to a Silicon Graphics (SGI) workstation (there is also a version for the Sun SPARC 1+ workstation) using the 44kHz sampling frequency of the SGI audio port (8000 Hz sampling rate for the SPARCstation).

3.1 Acoustic transduction

The microphone used was selected to be impact resistant, small, and light weight. In some environments, a microphone design that limits sensitivity to ambient sounds may also be desirable. In fact, since a large part of the acoustic signature comes from the vibratory signal of the impact itself, acceptable results may be obtainable by coupling the acoustic transducer to the tapping stick and minimizing the sensitivity to sound (in this case, avoiding the pickup of motor noise from the drive becomes an issue). In practice, we have obtained good results using various low-cost condenser type microphones.

3.2 Pan and Tilt Unit

The experimental methodology employed a two degree-of-freedom stepper-motor based pan-tilt unit (Directed Perception, model PTU-C) to move the boom-mounted microphone.

An appropriate boom must behave with both flexibility and rigidity. Flexibility leads to reduced torque-load on the drive and a reduced risk of damage when the sensor is driven (or overdriven) into a rigid material. On the other hand, rigidity increases the impact noise of the microphone and simplifies modelling of the end-point position and control of the rod (which is complicated by the impacts on different materials that it must sustain).

In previous work [6], a piece of 7.5mm diameter semi-rigid lucite (similar to plexiglass) was used. More recent work [4] demonstrated that a thin hollow aluminum rod proved to be a more stable experimental platform and this was the boom used for the experiments reported here (trading off flexibility against rigidity).

4 Sound Analysis

A simple fixed threshold was used to detect the onset of the impact signal when the boom/microphone makes contact with the floor material. Time domain analysis immediately shows that typical floor materials can have substantially different signal characteristics. For example, linoleum tends to have a bell-shaped decaying waveform with a characteristic double-bell, whereas those waveforms for carpet and tile exhibit much variability.

Although qualitative time domain classification of such signals is used in some contexts, we found existing qualitative classification schemes (attack, decay rate, etc.) insufficiently precise or inapplicable to this class of signals. Instead, we used a continuous frequency domain characterization of the signals [2]. For most materials we have examined, discrimination is possible even with 8 kHz bandwidth although our current implementation can use somewhat higher frequencies. Characteristic spectra including one with high-frequency components are shown in Figure 1.




Figure 1: *Sample characteristic time domain signal signals from metal and wood floor surfaces.*

4.1 Correlation Techniques

Our approach to signal matching is simple and efficient: to perform a normalized cross correlation between an unknown signal $\mathbf{s}(\omega)$ and the signals $\mathbf{k}_i(\omega)$ which are members of a library \mathbf{L} of known prototypes along with their associated typing information:

$$m(\mathbf{s}, \mathbf{k}_i) = \mathbf{s} \circ \mathbf{k}_i. \quad (1)$$

Preliminary classification of the surface type t is thus a matter of selecting the surface type that best matches the candidate signal:

$$t = \text{type}_i \mid \max_i(m(\mathbf{s}, \mathbf{k}_i)) \quad (2)$$

A further improvement was to threshold the cross correlation product. If the product between the unknown signal $\mathbf{s}(\omega)$ and the $\mathbf{k}_i(\omega)$ for type t was above 0.9, then the classification was deemed to be successful; otherwise (values less than 0.9), the signal $\mathbf{s}(\omega)$ was deemed to be “unknown”. This allowed us to determine if a signal $\mathbf{s}(\omega)$ could be correctly identified or not.

Clearly, the difficulty with this method is establishing the characteristic spectra of known surfaces (floor materials). In fact, simple methods such as averaging are inappropriate. A major reason for this is that we have found that the acoustic signature of simple surfaces (even at a fixed location) can take on different prototypical forms, as shown in Figure 2. Note that although the two signals are grossly similar, they exhibit as variety of different peaks and amplitude variations. In addition, we have found that different harmonics can be excited in different proportions due to variations in operating conditions such as temperature and humidity.

By using a small set of alternative characteristic spectra for each material, this problem can be resolved. In the section below, we outline this methodology for generating a *family* of characteristic spectra. This approach has the advantage of being generic, i.e. it makes few assumptions about the specific class of surfaces being dealt with, it can be performed automatically, and is much less case-specific than methods that attempt to infer “higher-level” features for detection.

4.2 Our experiments : generating characteristic spectra

Six categories of floor material were chosen for study: wood (table top), cement (cylinder), plastic




Figure 2: *Two different acoustic signatures for the same glass surface.*

(empty box), glass (empty bottle), and two kinds of metal (I-beam, empty box).

Using these samples, a total of 105 training samples (acoustic signatures) were obtained over a one day period, to generate the characteristic spectra required for our classification purposes, using the mean statistic for each surface: wood (20 samples), cement (10), plastic (20), glass (40), I-beam (10), metal box (5). (Median, maximum and minimum statistics were also used to generate characteristic spectra but they resulted in poorer recognition accuracy.)

As shown in Table 2, different numbers of characteristic spectra were required for the different floor materials, based on qualitative observations of the variability in their acoustic signatures, i.e. the number of characteristic spectra for each floor type was increased until 100% of the training samples could be classified correctly.

Once these characteristic spectra were defined, a total of 60 more samples were obtained for testing purposes, ten for each floor type. The results of the classification are presented in Table 3. In some cases, certain samples were identified as “unknown”; no sam-

Floor type	no. characteristic spectra
wood	3
cement	8
plastic	1
glass	3
I-beam	3
metal box	3

Table 2: *Number of characteristic spectra, by floor type, required for reliable surface classification.*

Floor type	Pct. correctly identified	Pct. unknown
wood	100	0
cement	90	10
plastic	100	0
glass	100	0
I-beam	90	10
metal box	90	10

Table 3: *Results of the floor type classification using characteristic spectra.*

ple was mis-classified.

In summary, over the 165 samples (105 for training + 60 for testing), we obtained a recognition accuracy of $162/165 = 98\%$.

5 Conclusion

The ‘blind person’s cane’ paradigm is an excellent method of secondary obstacle detection and surface identification as an aid to sonar and odometry efficiency. It is fast, being computationally- and time-inexpensive. While our experimental system is not perfect, operating at 98% accuracy, further work shows promise of eliminating errors.

For example, we have already explored the possibility of fully automating the data acquisition and spectra refinement so that the system may learn new surface materials by itself by tapping the new material and attempting to recognize it.

On another front, we have considered using strain gauges at the tip of the boom as an alternative to a microphone, but preliminary considerations suggest that the mounting and signal processing would entail substantial additional complexity.

6 Acknowledgments

The authors gratefully acknowledge the financial support of the National Sciences and Engineering Research Council of Canada, the Centre de recherche informatique de Montréal, and the technical assistance of Paul MacKenzie.

7 References

1. Barth, P.. “Sensor applications in robotics.” *Robotics Engineering* Vol. 8, no. 9, September 1986. pp. 17-20.
2. Brigham, E. The fast Fourier transform and its applications. Englewood Cliffs: Prentice Hall, 1988.
3. Durst, R., Krotkov, E. “Object Classification from Analysis of Impact Acoustics.” In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems* Pittsburgh, Pennsylvania, August 5-9, 1995, pp. 90-95.
4. Gervase, P.. “Final report for the tactile sensor project.” McGill University Final Year Design Project Report, 1994.
5. MacKenzie, P., Dudek, G. “Precise positioning using model-based maps.” In *Proceedings of the IEEE International Conference on Robotics and Automation*, San Diego, California, May 8-13, 1994.
6. Panagakos, A. “A tactile sensor for a mobile robot.” McGill University Summer Project Report, 1992.
7. Russell, R. “Using tactile whiskers to measure surface contours.” In *IEEE Int. Conf. on Robotics and Automation*, Nice, France, May, 1992, pp. 1295–1299.
8. Schenkman, B., Jansson, G. “The detection and localization of objects by the blind with the aid of long-cane tapping sounds.” *Human Factors* Vol. 28, October 1986. pp. 607-18.
9. Tominga, S. “Surface identification using the dichromatic reflection model.” *IEEE Transactions on Pattern Analysis & Machine Intelligence*. Vol. 13, July 1991. pp. 658-70.