

# Surface Identification Using Simple Contact Dynamics for Mobile Robots

Philippe Giguere and Gregory Dudek

**Abstract**—This paper describes an approach to surface identification in the context of mobile robotics, applicable to supervised and unsupervised learning. The identification is based on analyzing the tip acceleration patterns induced in a metallic rod, dragged along a surface that is to be identified. Eight features in time and frequency domains are used for classification. Results show that for ten type of indoor and outdoor surfaces, reliable identification can be achieved (90.0 and 94.6 percent for a 1 and 4 seconds time-window, respectively), using a non-sophisticated classifier (artificial neural network). Demonstration is done on how such a sensor and a simple control strategy can be used to guide a blind robot, using a simulation and a real differential drive robot.

## I. INTRODUCTION

This paper discusses a simple approach to terrain identification, based on the exploitation of contact dynamics. Surface sensing using a feeler can have several advantages for navigation. In the simplest case, the feeler can be used to detect anomalous surface conditions such as holes in the ground. More interestingly, it can be used to probe terrain properties. These properties, for example visco-elasticity or friction, might be challenging to measure accurately using vision or laser range-finders.

Supervised and unsupervised terrain identification can be performed using the information gathered by the feeler over different terrains. In turn, the resulting classification can be exploited in navigation policies. It can also be employed to train higher-dimensionality sensors, such as vision, to remotely recognize terrain types.

### A. Use of Whiskers as Tactile Sensor in Robotics

For a long time, whiskers in animals have been known [1] to be a simple and economical solution for collecting information about surface properties. One of the first examples of touch sensors in robotics is Grey Walter's famous *tortoise* robot [2]. It relied partly on a simple switch to modify the behavior of the robot, in order to avoid obstacles in its quest for reward.

This work was not supported by any organization  
P. Giguere is with the School of Computer Science, McGill University, Montreal, H3A 2A7, Canada philg@cim.mcgill.ca  
G. Dudek is with Faculty of Computer Science, Montreal, H3A 2A7, Canada dudek@cim.mcgill.ca

The task at hand dictates the amount and type of information that needs to be extracted from a tactile sensor. Deflection measurement between the robot and a point of contact have been used in wall-following tasks by Jung *et al.* [4]. A natural extension to this application is to perform profiling or shape measurement using the same deflection information, as in Russell [5], Scholz *et al.* [6], Clements *et al.* [10], or DaeEun *et al.* [11]. Finally, shape and texture extraction has been studied by Schultz *et al.* [12].

### B. Transducer Technology

Various transducer technologies have been employed to measure whisker deformation or vibration. Electrical contacts in Walter [2] or Russel [3], represent the simplest type. More sophisticated schemes include using potentiometers to measure springy whiskers deflection (Jung *et al.* [4], Russel [5]), strain gauges (Pearson *et al.* [7], Schultz *et al.* [12]), load cells (Scholz *et al.* [6], Clements *et al.* [10]), Hall-Effect Sensor (DaeEun *et al.* [11], Hipp *et al.* [13]), vibration or sound picked up by a microphone (Fend *et al.* [14], Roy *et al.* [21]) or changes in optical transmission of fibre optic whiskers (Liu *et al.* [8], Djordjevich *et al.* [9]).

### C. Terrain Identification Using On-board Accelerometers

Techniques to identify terrain relying on acceleration cues for outdoor wheeled vehicles have been explored in the literature (Sadhukan *et al.* [15], Brooks *et al.* [16], Weiss *et al.* [17] [18], DuPont *et al.* [19]). However, their sensitivity is limited by the vehicle itself. A heavy chassis mounted on a suspension system, coupled with air-filled tires, behaves like a damped mass-spring systems. This filters out higher-frequency components of accelerations sensed by accelerometers mounted on the chassis. Weiss *et al.* [18] mention how their results were impacted negatively when using such a vehicle compared to a cart with no suspension and hard plastic wheels.

A possible way to mitigate this drawback is to mount the accelerometers on the wheel frame, as in Brooks *et*

al. [16]. However, bypassing the suspension system will not improve detection of surface features that are much smaller than the radius of the wheel, such as small cracks on the ground. Using a sensor system similar to what is proposed here would address this issue, in a temporary basis during a self-supervised vision training phase for example, or on a permanent basis if need be. This type of sensor also makes it possible to sense terrain not only at the contact point of the wheels, but any arbitrary locations under the vehicle.

## II. DESCRIPTION OF THE SENSOR

Surface texture identification requires higher bandwidth transducers than those used typically for deflection measurements. In our system, a solid-state accelerometer is placed at the tip of a dragged aluminum rod, as shown in Fig 1. The hardness and weight of a metallic rod generate significant vibrations while being dragged on a surface. The top end of the rod is joined to the vehicle, maintaining an angle of approximately 40 deg relative to the surface. The accelerometer is oriented to pick up accelerations in the sagittal plane, corresponding to upward and downward motions of the rod.

This system is mechanically passive, low power, robust, and inexpensive. It can be made weather resistant without requiring seals around any rotating or moving parts. Energy used to probe the ground is provided by the vehicle itself as it moves forward. However, no information is collected when the robot is immobile, and backward motion requires the rod to be taken off the ground temporarily.




Fig. 1. Aluminum rod equipped with a single-axis accelerometer on its tip, used as a sensing device. A simple box emulates the chassis of a robot, and contains the data acquisition system.

### A. Selected Features

Important information contained in the phase spectrum is discarded by relying strictly on amplitude spectrum feature. For example, signals in Fig. 3 have identical amplitude spectra but different phase spectra. These phase spectrum differences translate into significant variations in the time-domain. Consequently, features in the time and frequency domain were both selected, within

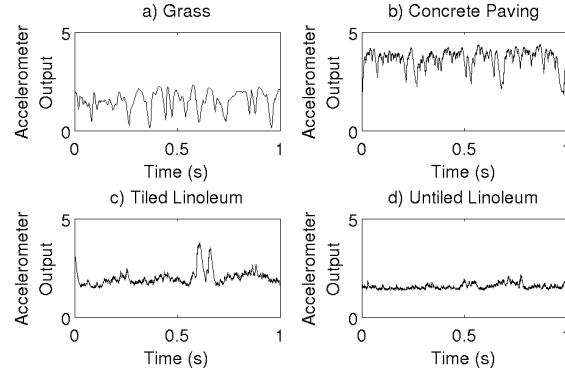


Fig. 2. Accelerometer output for a) grass, b) concrete paving, c) tiled linoleum and d) untiled linoleum surfaces. Grass has few high-frequency components: the soft surface prevents their generation. Paving signal has high-frequency components (due to surface hardness) as well as large jumps (the crack between the bricks). Linoleum surfaces are similar, except at time 0.6 s in tiled linoleum: this correspond to the rod going over a crack between tiles.

non-overlapping time-windows of size  $W$ . Seven of these features are in time-domain:

- mean, variance, skewness, kurtosis and 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 is in frequency-domain:

- sum of higher half of amplitude spectrum.

Some of these features are comparable to the ones used in Weiss *et al.* [17].

The variance of the signal is a good indicator of the amount of vertical motion experienced by the rod, which is large for uneven surfaces such as grass and gravel. The skewness of the histogram helps identify cases with asymmetric distribution of acceleration, typical for surfaces with regular but infrequent asperities such as cracks between tiles. The sign of the skewness plays an important role too, as some terrains tend to have a distribution more skewed to the top than the bottom: compare Fig. 2 b) and c) for example. The presence of high-frequency components helps differentiate between hard and soft surfaces.

## III. SUPERVISED CLASSIFIER TRAINING

The performance of classification in supervised learning, and indirectly the capability of clustering data during unsupervised learning, depends greatly on the quality and the discriminating power of a sensor and selected features on its data set. To evaluate this discriminating

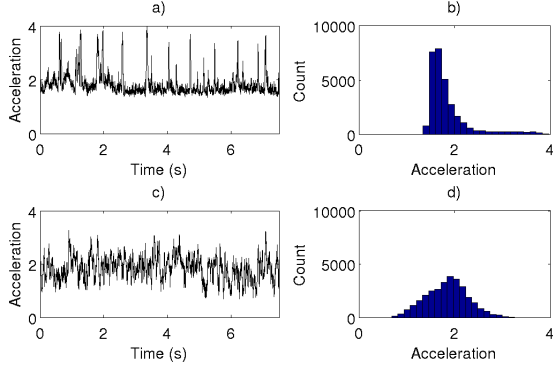


Fig. 3. Two signals a) and c) with identical amplitude spectra but different phase spectra. Signal c) was generated by replacing phase spectrum of a) with random phase values. Important features in time-domain are lost: the histograms b) and d) are significantly different. This points to limitations in pure spectral analysis for terrain identification, at least for this type of sensor.

power, an artificial neural network was selected as a classifier and trained to identify various terrain surfaces with manually collected data.

#### A. Data Collection

The box and rod combination was manually dragged over various surfaces, at an approximate rate of 70 cm/s. The scaled output of the accelerometer was collected at 4 kHz, 8 bit resolution, using an Isaac<sup>TM</sup> Data Acquisition system. After manually labelling the data, between 250 and 500 seconds of valid data was available for each surface. Four outdoor and six indoor horizontal surfaces were sampled: computer room flooring, tiled and untiled linoleum flooring, terazzo, wood bench surface, industrial carpet, grass, packed dirt with small gravel, gravel and paving made of concrete bricks. Fig. 2 shows the accelerometer output for one second, for four different surfaces.

#### B. Classifier Testing Results

The NETLAB [20] library implementation of a neural network was employed. 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 data set was split randomly, with 70 percent of the data going for training and the other 30 percent used for testing.

The success rate over all ten surfaces (Table I) was 90.0 and 94.6 percent for time windows of 1 and 4 seconds, respectively. When only outdoors terrains (grass, paving, large gravel, and small gravel) are considered in the testing phase for a neural networks trained on all ten surfaces, success rate climbs to 98.85 and 99.96 percent for time windows of 1 and 4 seconds, respectively.

TABLE I  
AVERAGE TESTING CLASSIFICATION RATE FOR A NEURAL NETWORK TRAINED TEN TERRAINS.

	Time-Window Size 1 s	Time-Window Size 4 s
Indoor and Outdoor	90.0 %	94.6 %
Outdoor only	98.85 %	99.96 %

This shows that outdoor surfaces are reliably identified, compared to indoor surfaces. One possible reason is that most of the indoor surfaces are smooth with few distinct features. Additionally, the carpet surface was sometimes confused with the packed dirt with small gravel outdoor terrain, as they are both irregular soft surfaces.

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

Fig. 4. Confusion matrix for testing phase of the trained neural network. Time-window size is 1 second. Results shown as percentage, averaged over 25 randomly training sets. Surfaces are sorted in descending order of success. Outdoor terrains are easily identified, while indoor surfaces are more challenging. Entries less than 0.05 percent are rounded down to 0.

#### C. Reducing Dimensionality: Applying PCA

A commonly used method to reduce dimensionality is Principal Component Analysis, or PCA. In PCA, data is projected along principal axes, and  $n$  primary projections are selected. Fig. 6 shows the classification rate as a function of the number of retained dimensions. This graph indicates that over 6 dimensions, no significant improvement is made in classification rate.

### IV. SURFACE CONTOUR FOLLOWING WITH SIMULATED VEHICLE AND A MOBILE ROBOT

Experiments were conducted with both a simulated and a real robot, to validate the idea that a simple contact-based terrain sensor can be used for practical navigation tasks.

Actual \ Predicted	Gr	Pa	La	Di	Co	Ca	Wo	Un	Te	Ti
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	0	100.0	0	0	0	0	0	0	0
Dirt with Small Gravel	0	0	0	100.0	0	0	0	0	0	0
Computer Room	0	0	0	0	99.8	0	0	0	0.2	0
Carpet	0	0	0	1.6	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	0.2	0	5.4	0	86.5	7.9	0
Tiled Linoleum	0	0	0	0.2	3.3	0	4.8	4.8	8.3	78.5

Fig. 5. Confusion matrix for testing phase of the trained neural network. Time-window size is 4 second. Results shown as percentage, averaged over 25 randomly training sets. Surfaces are sorted in descending order of success. Entries less than 0.05 percent are rounded down to 0.




Fig. 6. Variation of training and testing classification rates, as a function of the number of principal dimensions kept after performing PCA. Time-window size is 1 second.

### A. Control Algorithm For Surface Contour Following

The control algorithm works as follows. First, it is assumed that the robot is able to execute two simple locomotion commands: forward motion at speed  $V$ , and rotation at angular velocity  $\dot{\theta}$ . With the robot already on the track, it goes forward at a velocity  $V_{nominal}$ . When a sensor detects a transition to a different surface type ( $A$  on Fig. 7), the robot reduces its speed by half, and starts turning at a rate  $\dot{\theta}_{TR}$ . Once the outer sensor is brought back on the track (location  $B$ ), it drives straight for a brief moment in order to move away from the border. It then rotates in the opposite direction at rate  $-\dot{\theta}_{TR}$  for half the duration it took to bring the outer sensor back on the track. This orients the robot parallel to a line passing through  $A$  and  $B$ . If the followed contour is piecewise linear and  $A, B$  correspond to a straight contour of the surface, the robot will be oriented parallel to the track.

1) *Simulation Results:* The simulator operated by allowing a vehicle to move along a 2D terrain. The sensor readings were simulated by playing back the data previously collected in the training phase, based on the




Fig. 7. Control scheme used to follow a surface (light grey). Sensors are shown as darker grey disks, near the corners of the robot (shown as a square). The robot initially drives straight (a). When the left-hand side sensor leaves the track at  $A$ , it triggers a clockwise rotation (b) until the sensor reaches the track again at  $B$ . The robot then moves straight for a brief moment (c), then rotates counterclockwise by angle  $\theta_2 = \theta_1/2$  (d). Finally, it drives straight again (e), in an orientation parallel to a line passing through  $A$  and  $B$ .

terrain under the simulated rod sensor. The simulated environment was a 3-4 m wide track of concrete paving, surrounded by grass (see Fig. 8). The track itself was hand-drawn, resulting in a slightly irregular contour. Curves of various radii were present, in both directions. Simulated sensors were located directly in front of the robot at  $(0.6, 0.6)$  m and  $(0.6, -0.6)$  m. The trained classifier shown in Sec. III-B with a 1 second time-window was used, which strikes a compromise between rapid detection and accuracy. The robot was initially placed on the track at the location marked with a star in Fig. 8, heading in the direction of the side of the track. The nominal speed was  $V_{nominal} = 1$  m/s.

Overall the control strategy and detection results were robust for the presented simulation. The right-hand side sensor spent 94.2 percent of its time on the actual track, while the left-hand side sensor spent 79.0 percent of its time on the track. After more than one thousand loops, the robot was still following the track. The most difficult case was when the robot encounters a sharp 90 deg turn. In those cases, the robot would usually require two exits to negotiate the turn: once to get back on the track, and a second time to orient itself along the straight segment after the curve. This can be seen in the magnified area of Fig. 8.

2) *Experimental Results with a Mobile Robot:* A smaller version of the sensor was built and mounted on the left side of an iCreate™ robot from iRobot™. The robot is propelled using a differential drive, and its motion was remotely controlled via a serial interface. The sensor signals were captured using a standard sound card.




Fig. 8. Clockwise trajectory of the vehicle (thin black curves) as it follows a simulated concrete paving path (grey) surrounded by grass (white). Ten consecutive loops are shown. The starting location is marked with a star on the left-hand side. The robot initially points upward. The magnified view shows the robot trajectory as it negotiates a turn.




Fig. 9. Picture of the iRobot equipped with a rod sensor on the left side. The thick black arrow at the top indicates the direction of motion.

The robot was controlled using the same algorithm as in the simulation. The target surface consisted of a  $3 \times 3$  m carpet placed on top of a tiled floor. A mixture of Gaussians classifier was trained to differentiate these two surfaces. Qualitatively, the robot performed in a similarly fashion to the simulation. Fig. 10 shows the robot executing the re-entry maneuver. A successful test sequence of 5 minutes without failure can be seen in the submitted video. In this sequence, the robot returned on the carpet surface for every exit case, while managing to re-orient itself parallel to the contour of the surface after re-entry.

## V. UNSUPERVISED LEARNING OF TERRAINS WITH A MOBILE ROBOT

This set of experiments investigated the use of the sensor in the context of unsupervised learning. For each experiment run, the robot was manually placed and oriented in such a way so that it encountered two different surfaces along a straight path. A training




Fig. 10. Sequence of pictures taken as the robot executed the control algorithm, after leaving the carpet. Picture sequence is from left-to-right, top-to-bottom. After re-entry, the robot re-oriented itself parallel to the carpet edge.

data set was autonomously collected along this path at a constant speed of  $150$  mm/s, for the predetermined durations displayed in Table II. PCA was applied to the training data, while keeping the two major components. A mixture of two Gaussians classifier was then trained on this modified data set, using an unsupervised learning technique described in Giguere *et al.*[22]. This technique searches for the classifier parameter value  $\vec{\theta}$  that minimizes 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  $c_i \in C$ :

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

Following the training phase, the robot executed a  $180$  deg turn and moved forward at the same constant speed of  $150$  mm/s. The coming sensor information was classified in real time using the trained classifier. When a transition in surface classification was detected, the robot stopped and turned around to move back to the original surface.

This complete unsupervised process was executed for each trial. An experiment was considered successful when the robot drove up to the surface transition and triggered slightly past it. The majority of failures were early false detection of transitions. This could have been due to the non-representativeness of the collected training data, or the failure of the clustering algorithm to properly identify the surfaces in the training data set. Table II presents the results of these tests for three different surface transitions.

One has to bear in mind that a successful trial required

TABLE II

EXPERIMENTAL TESTING RESULTS OF AUTONOMOUS TERRAIN LEARNING AND IDENTIFICATION WITH ROD SENSOR, USING A DIFFERENTIAL DRIVE ROBOT.

Surface Type Combination	Training Data Duration	Success	Failures	Success Rate
Terazzo / Tiled Linoleum	40 s	10	10	50 %
Untiled Linoleum / Tiled Linoleum	30 s	15	5	75 %
Computer Room Flooring/ Carpet	10 s	20	0	100 %

at least ten successful classification in a row. This might explain why the Terazzo/Tiled Linoleum test results were significantly lower than the two others, standing at 50 percent: these two surfaces were not distinct enough to ensure a high success rate of the combined clustering and classification of the data.

## VI. CONCLUSION AND FUTURE WORKS

In this paper a novel sensor used for terrain identification and navigation purposes was presented. A classifier was successfully trained to identify a number of different types of outdoor and indoor surfaces, using a handful of features extracted from the sensor signal in time and frequency domains. It was shown how the dimensionality of the problem can be reduced down to 6 dimensions through PCA, without impacting significantly the classification results. Simulations and tests on a real robot demonstrated how a blind robot can navigate and follow a surface texture. Finally, a few examples of unsupervised learning of terrains were shown.

For future work, we intend on improving the sensor's mechanical design, to allow the robot to move backward without damaging the sensor. These improved sensors would then be deployed on a motorized wheelchair, and used to autonomously navigate in a structured outdoor setting.

## REFERENCES

- [1] S. B. Vincent. *The function of the vibrissae in the behavior of the white rat*, Behavior Monographs Vol 1, pp.1-81, 1912.
- [2] Walter Grey, *An Electromechanical Animal*, Dialectica (1950) Vol. 4: 4249
- [3] Russell, R.A. *Closing the sensor-computer-robot control loop*. Robotics Age, v 6, n 4, April 1984, p 15-20
- [4] Jung, D.; Zelinsky, A.; *Whisker based mobile robot navigation*. Intelligent Robots and Systems '96, IROS 96, Proceedings of the 1996 IEEE/RSJ International Conference on Volume 2, 4-8 Nov. 1996 Page(s):497 - 504 vol.2
- [5] Russell, R.A., *Using tactile whiskers to measure surface contours*. Robotics and Automation, 1992. Proceedings., 1992 IEEE International Conference on , vol., no., pp.1295-1299 vol.2, 12-14 May 1992
- [6] Scholz, G.R.; Rahn, C.D. *Profile sensing with an actuated whisker*. IEEE Transactions on Robotics and Automation, v 20, n 1, Feb. 2004, p 124-7
- [7] 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, 2007, 15:223-240.
- [8] Liu Renqiang, Wang Shuguo, Liu Pinkuan, Wang Xiaohui. *Curvature optical fiber whiskers for mobile robot guidance*. In High Technology Letters (English Language Edition), v 7, n 3, Sept. 2001, 79-83
- [9] Djordjevic, A. ; Tso, S.K.; Zhu, H.Y.; Pjevalica, V. *AGV guidance by fiber optic tactility*. In Proceedings of the SPIE - The International Society for Optical Engineering, v 3832, 1999, 206-11
- [10] Clements, T.N.; Rahn, C.D. *Three-dimensional contact imaging with an actuated whisker*. IEEE Transactions on Robotics, v 22, n 4, Aug. 2006, p 844-8
- [11] DaeEun Kim and Ralf Moller, *Biomimetic whiskers for shape recognition*, Robotics and Autonomous Systems, Volume 55, Issue 3, 31 March 2007, Pages 229-243.
- [12] Schultz, A.E.; Solomon, J.H.; Peshkin, M.A.; Hartmann, M.J., *Multifunctional Whisker Arrays for Distance Detection, Terrain Mapping, and Object Feature Extraction*. Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on , vol., no., pp. 2588-2593, 18-22 April 2005
- [13] Joerg Hipp, Ehsan Arabzadeh, Erik Zorzin, Jorg Conradt, Christoph Kayser, Mathew E. Diamond, and Peter Knig. *Texture Signals in Whisker Vibrations*, Journal of Neurophysiology, Mar 2006; 95: 1792 - 1799.
- [14] Miriam F., Bovet S., Yokoi H., Pfeifer R. *An active artificial whisker array for texture discrimination*. To appear in Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Las Vegas, October 2003.
- [15] Sadhukan, D., Moore, C. (2003). *Online terrain estimation using internal sensors*. In Proceedings of the Florida conference on recent advances in robotics, Boca Raton, FL, May 2003.
- [16] Brooks C., Iagnemma K., Dubowsky S. *Vibration-based Terrain Analysis for Mobile Robots*. In Proceedings of the 2005 IEEE International Conference on Robotics and Automation, 2005.
- [17] Weiss C., Fröhlich H., Zell A. *Vibration-based Terrain Classification Using Support Vector Machines*. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), October, 2006.
- [18] Weiss C., Fechner N., Stark M., Zell A. *Comparison of Different Approaches to Vibration-based Terrain Classification*. in Proceedings of the 3rd European Conference on Mobile Robots (ECMR 2007), Freiburg, Germany, September 19-21, 2007, pp. 7-12.
- [19] Dupont, E. M., Moore, C. A., Collins, E. G., Jr., Coyle, E. *Frequency response method for terrain classification in autonomous-ground vehicles*. In Autonomous Robots, January, 2008.
- [20] Ian T. Nabney. *NETLAB: algorithms for pattern recognition*. Springer-Verlag, 2002.
- [21] N. Roy, G. Dudek, P. Freedman. *Surface Sensing and Classification for Efficient Mobile Robot Navigation*. Proceedings of the IEEE/RSJ International Conference on Robotics and Automation. Minneapolis, 1996.
- [22] P. Giguere, G. Dudek. *Clustering Sensor Data for Terrain Identification using a Windowless Algorithm*. Proceedings of the Robotics: Science and Systems Conference, The MIT Press, 2008.