[8] J. M. Paros and L. Weisbord, "How to design flexure hinges," *Mach. Des.*, pp. 151–156, 1965.
[9] T. Yoshikawa, "Manipulability of robotic mechanisms," *Int. J. Robot. Res.*, vol. 4, no. 2, pp. 3–9, 1985.
[10] J. T. Wen and L. S. Wilfinger, "Kinematic manipulability of general constrained rigid multibody systems," *IEEE Trans. Robot. Autom.*, vol. 15, no. 3, pp. 558–567, Jun. 1999.
[11] A. Bicchi and D. Prattichizzo, "Manipulability of cooperation robots with passive joints," in *Proc. IEEE Int. Conf. Robot. Autom.*, 1998, pp. 1038–1044.
[12] F. Park and J. Kim, "Manipulability and singularity analysis of multiple robotic systems: A geometric approach," in *Proc. Int. Conf. Robot. Autom.*, Leuven, Belgium, May 1998, pp. 1032–1037.
[13] A. Ismail-Yahaya and A. Messac, "Required relationship between objective function and Pareto frontier orders: Practical implications," *Am. Inst. Aeronaut. Astronaut. J.*, vol. 39, no. 11, pp. 2168–2174, 2001.
[14] N. Lobontiu, J. S. N. Paine, E. O'Malley, and M. Samuelson, "Parabolic and hyperbolic flexure hinges, motion precision and stress characterization based on compliance closed-form equations," *J. Precision Eng.*, vol. 26, pp. 183–192, 2002.
[15] A. Messac, S. Gupta, and B. Akbulut, "Linear physical programming: A new approach to multiobjective optimization," *Trans. Oper. Res.*, vol. 8, pp. 39–59, Oct. 1996.
[16] A. Ismail-Yahaya and A. Messac, "Effective generation of the Pareto frontier using the normal constraint method," in *Proc. 40th Aerosp. Sci. Meet. Exhib.*, Reno, NV, Jan. 2002, [CD-ROM].
[17] J. Duffy, "The fallacy of modern hybrid control theory that is based on orthogonal complements of twist and wrench spaces," *J. Robot. Syst.*, vol. 7, no. 2, pp. 139–144, Feb. 1990.
[18] J. T. Wen and S. H. Murphy, "Position and force control of robot arms," *IEEE Trans. Autom. Control*, vol. 36, no. 3, pp. 365–374, Mar. 1991.
[19] G. Alici and B. Shirinzadeh, "Kinematics and stiffness analyses of a flexure-jointed planar micromanipulation system for a decoupled compliant motion," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Las Vegas, NV, Oct. 2003, pp. 3282–3287.
[20] E. Amatucci, N. G. Dagalakis, J. A. Kramar, and F. E. Scire, "Performance evaluation of a parallel cantilever biaxial micropositioning stage," in *Proc. Am. Soc. Precision Eng., 15th Annu. Meet.*, Scottsdale, AZ, 2000, [CD-ROM].
[21] N. G. Dagalakis, J. A. Kramar, E. Amatucci, and R. Bunch, "Kinematic modeling and analysis of a planar micropositioner," in *Proc. Am. Soc. Precision Eng., 16th Annu. Meet.*, Crystal City, VA, 2001, pp. 135–138.
[22] J. J. Gorman, N. G. Dagalakis, and B. G. Boone, "Multi-loop control of a nanopositioning mechanism for ultra-precision beam steering," in *Proc. SPIE Conf. Free-Space Laser Commun. Active Laser Illumin. III*, vol. 5160, San Diego, CA, 2003, pp. 170–181.
[23] B. G. Boone, R. S. Bokulic, G. B. Andrewsa, R. L. McNutt, Jr., and N. Dagalakis, "Optical and microwave communications system conceptual design for a realistic interstellar explorer," in *Proc. SPIE Conf. Free-Space Laser Commun. Active Laser Illumin. II*, vol. 4821, Seattle, WA, Jul. 2002, pp. 225–236.

# Vibration-Based Terrain Classification for Planetary Exploration Rovers

Christopher A. Brooks and Karl Iagnemma

***Abstract*—Safe, autonomous mobility in rough terrain is an important requirement for planetary exploration rovers. Knowledge of local terrain properties is critical to ensure a rover's safety on slopes and uneven surfaces. Visual features are often used to classify terrain; however, vision can be sensitive to lighting variations and other effects. This paper presents a method to classify terrain based on vibrations induced in the rover structure by wheel–terrain interaction during driving. This sensing mode is robust to lighting variations. Vibrations are measured using an accelerometer mounted on the rover structure. The classifier is trained using labeled vibration data during an offline learning phase. Linear discriminant analysis is used for online identification of terrain classes, such as sand, gravel, or clay. This approach has been experimentally validated on a laboratory testbed and on a four-wheeled rover in outdoor conditions.**

***Index Terms*—Mobile robots, pattern classification, robot sensing systems, rough terrain.**

## I. Introduction

Planetary exploration rovers are being proposed for missions to increasingly challenging locations [1]. These may include craters, hills, and ravines, where rocky outcrops might yield glimpses of a planet's history. To safely traverse slopes and highly uneven terrain, knowledge of local terrain properties is critical, since terrain conditions can strongly influence rover mobility. For example, a rover might climb a rocky slope with ease, but slide down a sandy slope of the same grade. Similarly, a rover traversing loose sand could become entrenched, where a rover crossing packed soil would face no such danger. Terrain-class knowledge would allow a rover to adapt its control and/or planning strategy to safely and efficiently traverse terrain of varying compositions.

Previous research on terrain classification has focused mainly on remote classification using vision or range data. Navigation systems employing these data have been demonstrated by researchers at the Jet Propulsion Laboratory, the National Institute of Standards and Technology, and Carnegie Mellon University [2]–[4]. Classification algorithms that rely on visual features (such as color or texture) are often sensitive to variations in illumination. In addition, vision-based classifiers usually detect features associated with the topmost terrain surface, which may not be the load-bearing surface of interest. An example of this is vegetation-covered terrain, or crusty Mars terrain covered by a thin layer of drift material. Recently, researchers have developed methods for estimating the location of this load-bearing surface [5]. Note, however, that this is not a classification method.

Other algorithms have employed range data to emphasize detection of geometric obstacles, such as rocks or steep slopes [6], [7]. These papers do not explicitly classify terrain, or address the issue of the soil itself being a hazard. Such a hazard can be termed a *nongeometric*

Manuscript received February 28, 2005. This paper was recommended for publication by Associate Editor P. Dupont and Editor S. Hutchinson upon evaluation of the reviewers' comments. This work was supported by the NASA Jet Propulsion Laboratory under the Mars Technology Program. This paper was presented in part at the IEEE International Conference on Robotics and Automation, Barcelona, Spain, April 2005.

The authors are with the Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: cabrooks@mit.edu; kdi@mit.edu).

Digital Object Identifier 10.1109/TRO.2005.855994

*hazard* [8]. Previous research at the Massachusetts Institute of Technology (MIT) has attempted to address this issue by developing algorithms to measure wheel sinkage into deformable terrain, identify soil cohesion and internal friction angles, and explicitly estimate terrain traversability [9]–[14]. Again, however, this work does not attempt to sort terrain regions into physically distinct classes.

Vibration-based terrain classification was suggested in 2002 by Iagnemma and Dubowsky as a novel sensing mode for identifying terrain class for hazard detection [15]. This method attempts to classify terrain based on vibrations in the rover structure induced by wheel–terrain interaction during driving. Sadhukhan *et al.* demonstrated vibration-based terrain classification for a high-speed vehicle, but the accuracy deteriorated at low speeds where vibration amplitudes were reduced [16]. Thus, it would not be applicable to planetary rovers, where speeds are expected to be under 15 cm/s.

This paper describes a vibration-based method for terrain classification for planetary rovers. This approach is insensitive to lighting variation. It can also successfully classify load-bearing surfaces that lie below a thin terrain layer of different composition. Thus, the method can complement visual classification approaches, as well as being a stand-alone classifier. In the proposed approach, an accelerometer is mounted to the rover structure. Vibrations sensed online during a traverse are classified based on their similarity to vibrations observed during an offline supervised training phase. This algorithm employs standard signal-processing techniques, including principal component analysis and linear discriminant analysis, to efficiently measure similarity. It is shown here that good experimental classification results can be obtained using even a simple classification algorithm.

## II. Terrain Classification Algorithm Overview

The algorithm presented here takes a signal-recognition approach to classifying terrain based on vibration signals. This is in contrast to an approach that might use a solid mechanics or finite-element model to analytically predict how the rover structure will vibrate in response to interaction with terrain. The proposed algorithm learns to recognize distinct terrain types based on labeled vibration data provided during an offline training phase. During training, the algorithm analyzes these data sets to form a low-dimensional representation of the signals corresponding to each labeled terrain. This information is stored in memory for use by an online classifier. During online classification, measured vibration signals are quickly classified as one of the labeled terrain classes. An overview schematic of the algorithm is shown in Fig. 1.

In this approach, vibration signals are first divided into short segments. These are then converted from time-domain voltage signals into power spectral densities (PSDs). Further analysis is performed in the Fourier domain. Log scaling of the power spectral magnitude is used to reduce the dominating effect of high-magnitude frequency components.

With the signals represented as a time series of Fourier spectra, training is a matter of dividing a high-dimensional space (i.e., the Fourier coefficients) into regions associated with an individual terrain class. To reduce the dimensionality of the comparison, principal component analysis is used [17]. Here only the first $k$ components are retained. The value of $k$ is chosen empirically. Note that principal components are computed during the training phase. These same principal components are used during the classification phase.

To define class boundaries in this principal component space, Fisher linear discriminant analysis is used as a pairwise classifier [18], [19]. In this approach, subclassifiers are created to classify a sample signal as being associated with one of two possible terrains. Separate subclassifiers are used for each possible pair of terrains. For example, for the three-terrain case of sand, gravel, and clay, one classifier would distinguish gravel from sand, another would distinguish gravel from clay,

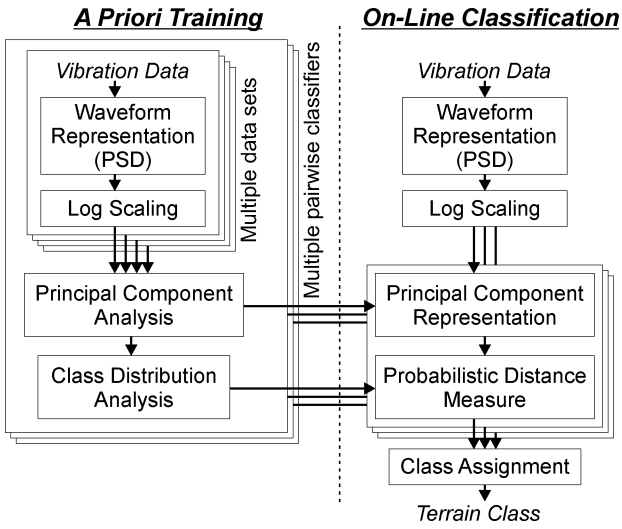


Fig. 1. Overview flowchart for vibration-based terrain-classification algorithm.

and a third would distinguish sand from clay. The linear discriminant is computed as the optimal vector along which to discriminate between the classes in the training data sets. Classification of a new vibration can be done by projecting its principal component representation onto this vector. A number of simple classifiers are available to address the resulting one-dimensional (1-D) classification problem [20], [21].

To accommodate classification of more than two terrains, a voting scheme is used. Each pairwise classifier can cast a "vote" for one of the two terrains it distinguishes, or remain "undecided." The winning terrain class is returned.

## III. Terrain-Classification Algorithm Detailed Description

The terrain-classification algorithm may be divided into two separate phases, *a priori* training and online classification. *A priori* training is computationally intensive, and is performed offline. Online classification is computationally efficient, and is performed during a rover traverse. These phases are described below.

### A. A Priori Training

During the *a priori* training phase, the algorithm learns to recognize vibration signatures corresponding to various user-selected terrain types. These are chosen to correspond to terrains of interest that a robot might encounter during field operations.

The first step in *a priori* training is to collect vibration data from representative terrains. Data should be collected for the terrain under a range of conditions spanning those for which the classifier is expected to perform (for example, under varying speeds, wheel-slip conditions, and wheel loads).

The sampled voltage output from the accelerometer is broken into short segments. The duration of these segments should be appropriately scaled to the physical scenario (i.e., a single segment should contain data from a travel distance that is scaled to the wheel diameter, spacing between grousers, spatial variations of terrain, etc.). The PSD of each of these segments is then computed using Welch's method [22], and a log-scaled version of this PSD is stored in a matrix.

For illustration purposes, consider classification of sand and gravel. Data for sand would be stored in a matrix $\mathbf{Y}_{\mathrm{sand}}$ as

$$\mathbf{Y}_{\mathrm{sand}} = \begin{bmatrix} y_{\mathrm{sand},f\min,1} & \cdots & y_{\mathrm{sand},f\min,n} \\ \vdots & \ddots & \vdots \\ y_{\mathrm{sand},f\max,1} & \cdots & y_{\mathrm{sand},f\max,n} \end{bmatrix}. \quad (1)$$

In this representation, each column corresponds to a single time segment, and contains the log PSD components in a frequency range of interest. Each row corresponds to a single frequency and contains the log PSD components for all time segments.

A separate classifier is used to distinguish between terrain pairs. For each classifier, a discrimination vector and terrain-class statistics are produced as follows.

The two matrices describing the training data for each class, $\mathbf{Y}_{\text{sand}} \in \Re^{m \times n_{\text{sand}}}$ and $\mathbf{Y}_{\text{gravel}} \in \Re^{m \times n_{\text{gravel}}}$, are combined to form a complete record of the data $\mathbf{Y} = [\mathbf{Y}_{\text{sand}} \quad \mathbf{Y}_{\text{gravel}}]$, $\mathbf{Y} \in \Re^{m \times n}$, where $m$ is the number of frequency components, and $n = n_{\text{sand}} + n_{\text{gravel}}$ is the total number of time segments. The row mean is then subtracted, to produce the matrix $\hat{\mathbf{Y}}$.

Singular value decomposition [23] is then used to separate $\hat{\mathbf{Y}}$ into three matrices, $\mathbf{U}_a$, $\mathbf{S}_a$, and $\mathbf{V}_a$

$$\hat{\mathbf{Y}} = \mathbf{U}_a \mathbf{S}_a \mathbf{V}_a^{\mathbf{T}}. \tag{2}$$

Only the first $k$ columns of $\mathbf{U}_a$ (i.e., the first $k$ principal components of $\hat{\mathbf{Y}}$) and the upper-left $k \times k$ block of $\mathbf{S}_a$ are retained, in matrices $\mathbf{U}_{\text{signal}}$ and $\mathbf{S}_{\text{signal}}$. The value for $k$ was selected empirically to represent only the subspace of $\mathbf{U}_a$ related to signal (rather than noise) information. Systematic methods for choosing $k$ are available [24]. Using too high a value with a limited amount of training data can be detrimental, as this would overtrain the algorithm (i.e., train it to recognize noise in the training data, reducing its ability to classify new data). In practice, we have used $k = 15$, as it appears to give good signal representation without overfitting. In experiments, the first 15 principal components account for approximately 90% of the variance in the training data.

The principal component representations of $\mathbf{Y}_{\text{sand}}$ and $\mathbf{Y}_{\text{gravel}}$ can then be computed as

$$\mathbf{W}_{\text{sand}} = \mathbf{S}_{\text{signal}}^{-1} \mathbf{U}_{\text{signal}}^{\mathbf{T}} \mathbf{Y}_{\text{sand}} \tag{3}$$

$$\mathbf{W}_{\text{gravel}} = \mathbf{S}_{\text{signal}}^{-1} \mathbf{U}_{\text{signal}}^{\mathbf{T}} \mathbf{Y}_{\text{gravel}}. \tag{4}$$

For all practical cases, if $n > k$, $\mathbf{S}_{\text{signal}}$ will be invertible. If rank $(\mathbf{S}) < k$, a smaller value for $k$ can be used.

Linear discriminant analysis is used to find an optimal vector along which to distinguish the two classes. This vector $\mathbf{d}$ is computed as follows. First, the row means of the classes in the principal component space are stored as $\overline{\mathbf{w}}_{\text{sand}}$ and $\overline{\mathbf{w}}_{\text{gravel}}$. These means are then subtracted from $\mathbf{W}_{\text{sand}}$ and $\mathbf{W}_{\text{gravel}}$ to produce the matrices $\hat{\mathbf{W}}_{\text{sand}}$ and $\hat{\mathbf{W}}_{\text{gravel}}$.

These matrices are scaled by the number of points in the training data and are merged to form the matrix $\hat{\mathbf{W}}$

$$\hat{\mathbf{W}} = \left[ \frac{1}{\sqrt{n_{\text{sand}} - 1}} \hat{\mathbf{W}}_{\text{sand}} \quad \frac{1}{\sqrt{n_{\text{gravel}} - 1}} \hat{\mathbf{W}}_{\text{gravel}} \right]. \tag{5}$$

This matrix is then decomposed into its singular value representation

$$\hat{\mathbf{W}} = \mathbf{U}_b \mathbf{S}_b \mathbf{V}_b^{\mathbf{T}}. \tag{6}$$

Here, $\mathbf{U}_b$ and $\mathbf{S}_b$ are $k \times k$ matrices. The vector $\mathbf{d}$ may then be computed as

$$\mathbf{d} = \mathbf{U}_b \mathbf{S}_b{}^{-1} \mathbf{S}_b{}^{-1} \mathbf{U}_b{}^{\mathbf{T}} \left( \overline{\mathbf{w}}_{\text{sand}} - \overline{\mathbf{w}}_{\text{gravel}} \right). \tag{7}$$

This can be shown to be the Fisher linear discriminant in the $k$-dimensional principal component space.

The discrimination metric $d(\mathbf{y})$ is defined as the scalar product of $\mathbf{d}$ with the principal component representation of a vibration

$$d(\mathbf{y}) = \mathbf{d}^{\mathbf{T}} \mathbf{S}_{\text{signal}}^{-1} \mathbf{U}_{\text{signal}}^{\mathbf{T}} \mathbf{y} \tag{8}$$

where $\mathbf{y}$ is the log PSD of a data segment.

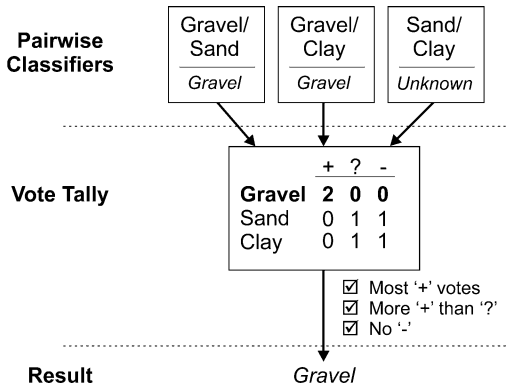


Fig. 2. Schematic of voting scheme positively identifying gravel.

The last step in the *a priori* analysis is to compute the statistics of the discrimination metric for the training data. The means and standard deviations are computed as $\overline{d}_{\text{sand}}$, $\sigma_{\text{sand}}$, $\overline{d}_{\text{gravel}}$, and $\sigma_{\text{gravel}}$.

The discrimination vector and the terrain-class statistics are stored for use in the online classification phase of the algorithm.

### B. Online Classification

During a rover traverse, short segments of vibration sensor data are collected, of the same duration as those use in *a priori* training. For each segment, the PSD is computed, and the magnitude is log-scaled and stored in a vector $\mathbf{y}$.

Pairwise classifiers then compute the discrimination metric $d(\mathbf{y})$. The Mahalanobis distances [25] from $d(\mathbf{y})$ to the terrain-class means, e.g., $\overline{d}_{\text{sand}}$ and $\overline{d}_{\text{gravel}}$, are then computed as

$$md_{\text{sand}}(\mathbf{y}) = \frac{\left| d(\mathbf{y}) - \overline{d}_{\text{sand}} \right|}{\sigma_{\text{sand}}} \tag{9}$$

$$md_{\text{gravel}}(\mathbf{y}) = \frac{\left| d(\mathbf{y}) - \overline{d}_{\text{gravel}} \right|}{\sigma_{\text{gravel}}}. \tag{10}$$

If the difference between the Mahalanobis distances is less than one (i.e., $|md_{\text{sand}}(\mathbf{y}) - md_{\text{gravel}}(\mathbf{y})| < 1$), the pairwise classifier labels the vibration as "unknown." Otherwise, the pairwise classifier labels the vibration as the terrain with the smaller Mahalanobis distance.

A voting scheme merges the results of the various pairwise classifiers. In this approach, each pairwise classifier may return a terrain label, or it may return "unknown." If a pairwise classifier returns a positive vote for a terrain class, the alternative terrain class receives a negative vote. If the pairwise classifier returns "unknown," both classes receive an "unknown" vote.

For a terrain to be positively identified, it must: 1) receive more positive votes than any other terrain class; 2) receive only positive and "unknown" votes; and 3) receive more positive votes than "unknown" votes. These rules were chosen to provide a conservative estimate that would not become drastically more or less conservative with an increased number of terrain classes. This is based on the belief that returning "unknown" is preferable to returning the wrong terrain class. Fig. 2 shows an example of the voting algorithm employed in a three-terrain classifier positively identifying gravel. Fig. 3 shows an example of the voting algorithm unable to positively identify a terrain.

## IV. Experimental Results

The algorithm presented above was experimentally validated on both the Field and Space Robotics Laboratory (FSRL) Wheel–Terrain Interaction Testbed, and the FSRL Technology Testbed Rover, TORTOISE [10].

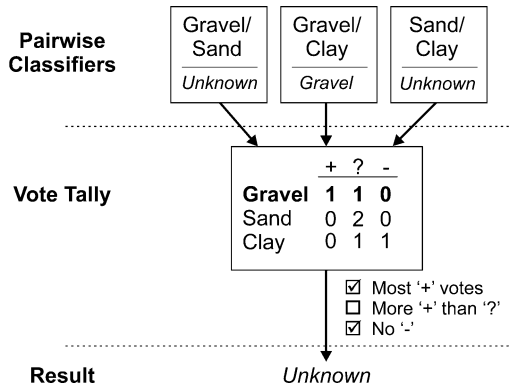


Fig. 3. Schematic of voting scheme resulting in unknown result.

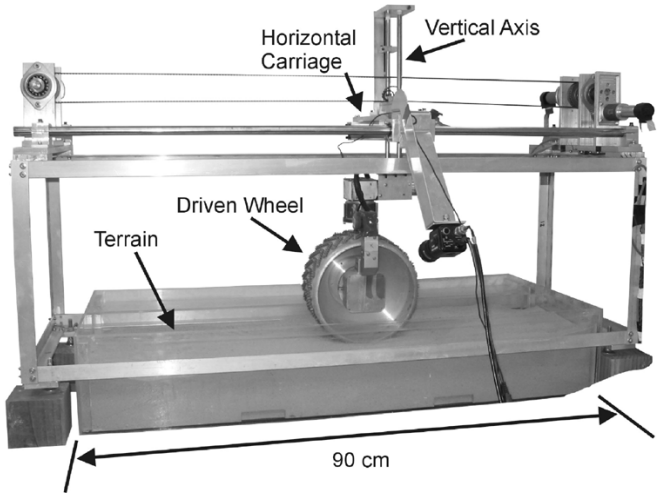


Fig. 4. FSRL wheel–terrain interaction testbed with FIDO wheel.

### A. Experimental Results—FSRL Wheel–Terrain Interaction Testbed

The FSRL Wheel–Terrain Interaction Testbed, shown in Fig. 4, consists of a driven wheel mounted on an undriven vertical axis. The wheel-axis assembly is mounted on a driven carriage, so the wheel forward velocity and angular velocity can be controlled independently. These testbed experiments were conducted using a wheel from the FIDO rover [26] supplied by the Jet Propulsion Laboratory. For these experiments, three terrains were used: landscaping gravel, JSC Mars-1 Soil Simulant [27], and washed beach sand. Landscaping gravel is a mixture of small rounded pebbles ranging in size from 0.5 to 2 cm. JSC Mars-1 Soil Simulant is a dry glassy volcanic ash, developed by Johnson Space Center to represent the Martian soil as observed by Viking Lander 1. It contains fine particles as well as larger solid particles ranging up to 4 cm. Washed beach sand is a homogeneous fine-grained dry sand.

In these experiments, wheel forward velocity ranged from 0.5 to 5 cm/s, with these values chosen to be similar to planned rover missions. Forward velocity was set at a constant value for each trial. The wheel-slip ratio $i$ was varied from 0 to 0.5. (The slip ratio is defined as $i = 1 - V/r\omega$, where $V$ is the wheel forward velocity, $r$ is the wheel radius, and $\omega$ is the wheel angular velocity.) The vertical load on the terrain was varied, as well, from 30 to 50 N, including the weight of the wheel. This variation captures the effect of weight distribution among rover wheels due to travel over uneven terrain.

Vibration signals were sensed using a contact microphone mounted to the wheel frame, as shown in Fig. 5. These signals were collected


Fig. 5. Vibration sensor mounted on the FIDO wheel in the FSRL wheel–terrain interaction testbed.

TABLE I
CLASSIFICATION RESULTS FOR FSRL WHEEL–TERRAIN INTERACTION TESTBED VIBRATION DATA

| | | Classification Result | | | | |
|---|---|---|---|---|---|---|
| | | Gravel | Mars-1 | Sand | Unknown | Total |
| Actual Terrain | Gravel | 302 | 2 | 0 | 8 | 312 |
| | Mars-1 | 5 | 208 | 3 | 61 | 277 |
| | Sand | 0 | 51 | 139 | 86 | 276 |

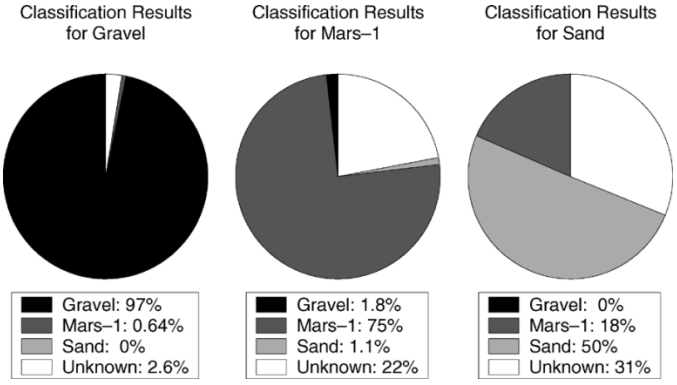


Fig. 6. Classification results for FSRL wheel–terrain interaction testbed vibration data.

using a desktop computer with a sound card. Sixteen-bit samples were collected at a frequency of 44.1 kHz. A single data set consisted of vibration data recorded over a full traverse of the wheel across the 90-cm-long testbed, at a specified load, forward velocity, and angular velocity.

Nearly 2 h of vibration data were collected over the parameter range described above, and the algorithm was tuned using the leave-one-out approach [28]. Tuning consisted of selecting appropriate values for: 1) the range and spacing of frequency components for spectral representation; 2) the number of principal components used to represent the signal space; and 3) the discrimination thresholds for the pairwise classifiers. A single combination of tuned parameters is shared among all pairwise classifiers.

Once the parameters were tuned, the classification accuracy was evaluated. First, the vibration data was randomly divided into training data and test data sets. For each of the three terrains, ten data sets were randomly chosen as test data. This represents approximately 25% of the total data. The remaining data sets were chosen for training.

A three-terrain classifier was trained using the labeled training data sets. Here, a 1 s segment length was used. After the classifier was trained, it was used to classify the test data sets. Classification results are presented in Table I. Values shown are counts of 1-s-long vibration segments. These results are plotted in Fig. 6.

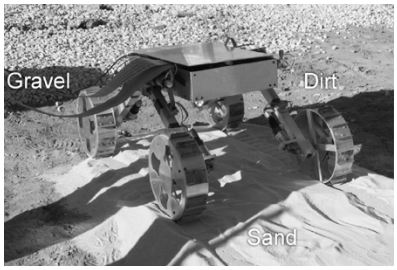


Fig. 7. FSRL technology testbed rover completing three-terrain traverse.

These results show the algorithm's ability to distinguish between multiple terrain types. When attempting to identify gravel-induced vibrations, the algorithm misclassified less than 1% of the test data as Mars-1 or sand. Similarly, when classifying Mars-1 and sand vibration data, less than 1% was misclassified as gravel. This clearly demonstrates the ability of the algorithm to identify terrains which induce obviously distinct vibrations.

The more challenging distinction was between Mars-1 and sand. These two terrains are alike in the fact that they contain small particles which may damp out vibrations in the wheel. Despite this similarity, less than 2% of the Mars-1 vibration data was misidentified as being sand. The difficulty of this distinction reveals itself in the amount of sand vibration data being misidentified as Mars-1. Nevertheless, these misclassifications comprise less than 20% of the sand vibration data, while most of the data is correctly classified.

Considering the inverse problem, of having confidence that the actual terrain matches the classification result, the algorithm performs quite well. Given equal prior likelihoods of the above three terrains, the algorithm is more than 98% confident that terrain identified as gravel is actually gravel. Similarly, the algorithm is more than 97% confident that terrain identified as sand is truly sand. The confidence for Mars-1 is almost 80%.

It should be noted that these results are based solely on 1 s samples of vibration data, and incorporate no memory of prior classifications. An intelligent algorithm on a rover might incorporate an estimate of the likelihood of a transition from one terrain to another to improve overall classification results. Another way to improve terrain-classification accuracy would be to combine the vibration-based classification with visual classification methods.

The use of separate velocity-dependent classifiers was also studied; however, the improvement in classification accuracy was not sufficiently large to warrant added complexity. Other experiments were performed to attempt to detect the presence of a rock covered by a thin (approximately 5 mm) layer of sand. These experiments successfully distinguished between sand-covered rock and deep sand, with classification accuracy rates similar to those presented in Table I.

### B. Experimental Results: FSRL Technology Testbed Rover

The algorithm was also experimentally validated using data collected on the four-wheeled FSRL Technology Testbed Rover, TORTOISE, to study the vibration response of a wheeled rover in outdoor terrain. For these experiments, three terrains were used: sand, gravel, and packed dirt. The sand used was a homogeneous composition of fine-grained dry sand. The gravel was composed of stones ranging from 1 to 3 cm. The packed dirt was thoroughly compacted and solid. Fig. 7 shows the rover during a traverse across all three terrains. For scale, the rover is 80-cm long.

Vibration signals were sensed using a contact microphone mounted to the front right leg of the rover, near the joint with the wheel axle, as seen in Fig. 8. These signals were collected using a laptop computer

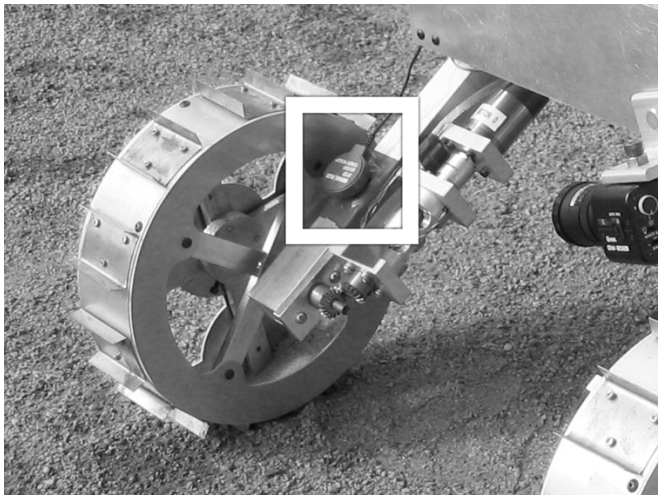

Fig. 8. Vibration sensor mounted on FSRL technology testbed rover.

TABLE II
CLASSIFICATION RESULTS FOR FSRL TECHNOLOGY TESTBED ROVER VIBRATION DATA

| | | Classification Result | | | | |
|---|---|---|---|---|---|---|
| | | Sand | Dirt | Gravel | Unknown | Total |
| Actual Terrain | Sand | 24 | 0 | 0 | 1 | 25 |
| | Dirt | 0 | 21 | 1 | 5 | 27 |
| | Gravel | 0 | 4 | 36 | 4 | 44 |

with a sound card. Sixteen-bit samples were collected at a frequency of 44.1 kHz.

A single data set consisted of vibration data from a 15–30 s traverse of the rover across the terrain at a constant velocity. Velocity was varied from 2 to 5 cm/s. Data from both forward and reverse driving directions were collected. Four single-terrain data sets were collected from each terrain. Once the single-terrain data sets were collected, multiterrain data sets were collected with the rover driving from sand to packed dirt to gravel in a single traverse, in forward and reverse. Three of the four single-terrain data sets from each terrain were used to train a three-terrain classifier, while the fourth data set was used for testing. Training and testing were repeated, such that each data set was used for testing. Due to the increased spacing between the grousers on the rover wheels compared with the wheel on the testbed, the segment length was increased to 3 s.

Table II shows the results by terrain for the classification of single-terrain data sets. The values are counts of the 3-s-long vibration segments. These results are plotted in Fig. 9.

These results show the classification accuracy of the algorithm using vibration data collected on TORTOISE. As with the data from the wheel–terrain testbed, the algorithm demonstrates excellent capability in distinguishing terrains which induce qualitatively different vibrations. Here, none of the data collected on sand was misidentified as gravel or dirt, and vice versa.

The most similar terrain types in these data sets were packed dirt and gravel. Both are solid surfaces unlikely to damp vibrations. It is not surprising that distinguishing between the two terrain types would be challenging. In the results presented above, however, only one 3-s segment of packed dirt was misclassified as gravel. Only four vibration segments of gravel were misclassified as packed dirt.

Given equal prior probabilities of each terrain type, there is a 96% confidence level that terrain labeled as gravel is truly gravel, and a 90% confidence level that terrain labeled as packed dirt is truly dirt. In the

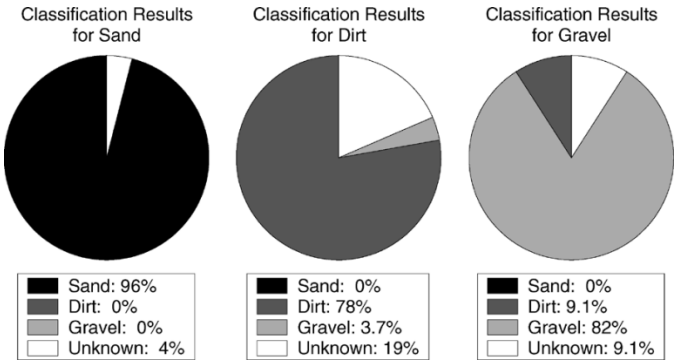


Fig. 9. Classification results for FSRL technology testbed rover vibration data.

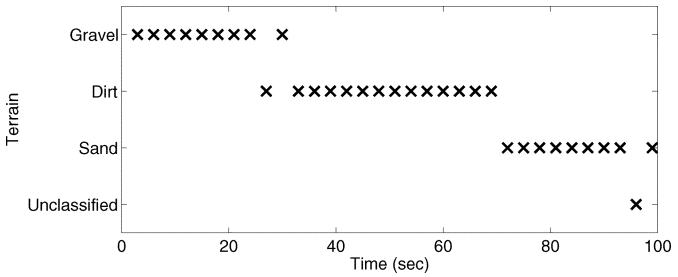


Fig. 10. Classification results for FSRL technology testbed rover traverse of gravel, packed dirt, and sand at 4 cm/s.

test data, no classification of vibration data as sand was wrong, so the estimated confidence level for sand is 100%.

Fig. 10 shows the classification results for distinct single traverses of all three terrain types. This plot may be viewed as a strip chart, plotting the result of the terrain classifier against the time the classifier made the identification. It can be clearly observed that the rover is traveling from an area of gravel, to an area of packed dirt, and then onto an area of sand. These results clearly show the effectiveness of the algorithm in classifying multiple terrains during a single traverse. This figure also illustrates the most common error mechanism of the algorithm, which occurs when the rover transitions between homogenous terrain regions.

The algorithm required 37 s of computation on a Pentium III 933 MHz computer for the *a priori* training phase, using 6 min of labeled training data. Online classification was much quicker, taking 475 ms per 3-s vibration segment, using unoptimized Matlab code. An optimized version of the code would be several times faster in both training and classification.

## V. Conclusion

An algorithm has been presented for classifying terrain based on the vibrations in the rover structure induced by wheel–terrain interaction. This algorithm uses linear discriminant analysis to distinguish between each pair of terrain classes, and uses a voting algorithm to arbitrate between pairwise classifiers. The terrain class is returned if the algorithm can uniquely identify one; otherwise, an "unknown" result is returned.

Experimental results demonstrate the classification accuracy using vibration data from two different testing setups. Results from the FSRL Wheel–Terrain Interaction Testbed demonstrated the effectiveness of the algorithm in a controlled laboratory environment. Results from the FSRL Technology Testbed Rover demonstrated the effectiveness of the algorithm in a natural, uncontrolled environment on a wheeled rover. Classification of three distinct terrain types was demonstrated on each platform. Additionally, results from a multiterrain data set from the FSRL Technology Testbed Rover illustrated the use of this algorithm in identifying multiple terrain classes during a rover traverse.

This data shows the potential for vibration-based classification as a standalone terrain classifier, or as a complement to current vision-based terrain-classification approaches. The presented algorithm is a simple, inexpensive, and computationally efficient method for extracting terrain-class information from vibration data. It may be used as a component of a meta-classifier, combining data from multiple sensors along with a memory of past classification results, or a similar approach may be used to define terrain-class probabilities in a Bayesian classifier. This is an area of current research.

## Acknowledgment

The authors would like to acknowledge the support of Drs. R. Volpe and S. Hayati at JPL.

## References

[1] R. Volpe, "Rover functional autonomy development for the Mars mobile science laboratory," in *Proc. IEEE Aerosp. Conf.*, vol. 2, Big Sky, MT, 2003, pp. 643–652.

[2] P. Bellutta, R. Manduchi, L. Matthies, K. Owens, and A. Rankin, "Terrain perception for DEMO III," in *Proc. IEEE Intell. Veh. Symp.*, Dearborn, MI, 2000, pp. 326–332.

[3] C. Rasmussen, "Combining laser range, color, and texture cues for autonomous road following," in *Proc. IEEE Int. Conf. Robot. Autom.*, Washington, DC, 2002, pp. 4320–4325.

[4] N. Vandapel, D. Huber, A. Kapuria, and M. Hebert, "Natural terrain classification using 3-D ladar data," in *Proc. IEEE Int. Conf. Robot. Autom.*, vol. 5, Apr. 2004, pp. 5117–5122.

[5] C. Wellington and A. Stentz, "Learning predictions of the load-bearing surface for autonomous rough terrain navigation in vegetation," in *Proc. Int. Conf. Field, Service Robot.*, Jul. 2003, pp. 49–54.

[6] S. Singh and B. Digney, "Autonomous cross-country navigation using stereo vision," Carnegie Mellon Univ., Robot. Inst., Pittsburgh, PA, Tech. Rep. CMU-RI-TR-99-03, 1999.

[7] R. Manduchi, A. Castaño, A. Talukder, and L. Matthies, "Obstacle detection and terrain classification for autonomous off-road navigation," *Auton. Robots*, vol. 18, pp. 81–102, Jan. 2005.

[8] B. H. Wilcox, "Non-geometric hazard detection for a Mars microrover," in *Proc. AIAA Conf. Intell. Robot. Field, Factory, Service, Space*, vol. 2, Houston, TX, 1994, pp. 675–684.

[9] K. D. Iagnemma, C. A. Brooks, and S. Dubowsky, "Visual, tactile, and vibration-based terrain analysis for planetary rovers," in *Proc. IEEE Aerosp. Conf.*, vol. 2, Big Sky, MT, 2004, pp. 841–848.

[10] C. A. Brooks, "Terrain identification methods for planetary exploration rovers," Master's thesis, Mass. Inst. Technol., Dept. Mech. Eng., Cambridge, MA, 2004.

[11] K. D. Iagnemma, H. Shibly, and S. Dubowsky, "On-line terrain parameter estimation for planetary rovers," in *Proc. IEEE Int. Conf. Robot. Autom.*, Washington, DC, 2002, pp. 3142–3147.

[12] K. D. Iagnemma, S. Kang, H. Shibly, and S. Dubowsky, "Online terrain parameter estimation for wheeled mobile robots with application to planetary rovers," *IEEE Trans. Robot.*, vol. 20, no. 5, pp. 921–927, Oct. 2004.

[13] M. G. Bekker, *Theory of Land Locomotion*. Ann Arbor, MI: Univ. Michigan Press, 1956.

[14] S. Kang, "Terrain parameter estimation and traversability assessment for mobile robots," Master's thesis, Mass. Inst. Technol., Dept. Mech. Eng., Cambridge, MA, 2003.

[15] K. D. Iagnemma and S. Dubowsky, "Terrain estimation for high speed rough terrain autonomous vehicle navigation," in *Proc. SPIE Conf. Unmanned Ground Veh. Technol. IV*, Orlando, FL, 2002, Paper 4715-31.

[16] D. Sadhukhan, C. Moore, and E. Collins, "Terrain estimation using internal sensors," in *Proc. IASTED Int. Conf. Robot. Appl.*, Honolulu, HI, 2004, Paper 447-800.

[17] I. T. Jolliffe, *Principal Component Analysis*. New York: Springer-Verlag, 1986.
[18] S. Balakrishnama and A. Ganapathiraju. (1998) Linear discriminant analysis—A brief tutorial. [Online]. Available: http://www.isip.msstate.edu/publications/reports/isip_internal/1998/linear_discrim_analysis/lda_theory.pdf
[19] R. A. Fisher, "The use of multiple measurements in taxonomic problems," *Ann. Eugen.*, vol. 7, pp. 179–188, 1936.
[20] K. V. Mardia, J. T. Kent, and J. M. Bibby, *Multivariate Analysis*. San Diego, CA: Academic, 1979.
[21] G. J. McLachlan, *Discriminant Analysis and Statistical Pattern Recognition*. New York: Wiley, 1992.
[22] P. D. Welch, "The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms," *IEEE Trans. Audio Electroacoust.*, vol. AU-15, pp. 70–73, Jun. 1967.
[23] G. H. Golub and C. F. Van Loan, "The singular value decomposition and unitary matrices," in *Matrix Computations*, 3rd ed. Baltimore, MD: Johns Hopkins Univ. Press, 1996, pp. 70–71 and 73.
[24] E. R. Malinowski, *Factor Analysis in Chemistry*, 2nd ed. New York: Wiley, 1991.
[25] P. C. Mahalanobis, "On the generalized distance in statistics," *Proc. Nat. Inst. Sci. India*, vol. 2, no. 1, pp. 49–55, 1936.
[26] P. S. Schenker, E. Baumgartner, L. Dorsky, P. Backes, H. Aghazarian, J. Norris, T. Huntsberger, Y. Cheng, A. Trebi-Ollennu, M. Garrett, B. Kennedy, A. Ganino, R. Arvidson, and S. Squyres, "FIDO: A field integrated design and operations rover for Mars surface exploration," in *Proc. 6th Int. Symp. Artif. Intell., Robot., Autom. Space*, Montreal, QC, Canada, 2001, Paper AM012.
[27] C. Allen, R. Morris, K. Jager, D. Golden, D. Lindstrom, M. Lindstrom, and J. Lockwood, "Martian regolith simulant JSC Mars-1," in *Proc. 29th Lunar Planet. Sci. Conf.*, Houston, TX, 1998, Abstract 1690.
[28] T. Jaakolla and D. Haussler, "Probabilistic kernel regression models," in *Proc. 7th Int. Workshop Artif. Intell. Statist.*, San Francisco, CA, 1999, [CD-ROM].

# Distributed Route Planning for Multiple Mobile Robots Using an Augmented Lagrangian Decomposition and Coordination Technique

Tatsushi Nishi, Masakazu Ando, and Masami Konishi

***Abstract*—To enable efficient transportation in semiconductor fabrication bays, it is necessary to generate route planning of multiple automated guided vehicles (AGVs) efficiently to minimize the total transportation time without collision among AGVs. In this paper, we propose a distributed route-planning method for multiple mobile robots using an augmented Lagrangian decomposition and coordination technique. The proposed method features a characteristic that each AGV individually creates a near- optimal routing plan through repetitive data exchange among the AGVs and local optimization for each AGV. Dijkstra's algorithm is used for local optimization. The optimality of the solution generated by the proposed method is evaluated by comparing the solution with an optimal solution derived by solving integer linear programming problems. A near-optimal solution, within 3% of the average gap from the optimal solution for an example transportation system consisting of 143 nodes and 14 AGVs, can be derived in less than 5 s of computation time for 100 types of requests. The proposed method is implemented in an experimental system with three AGVs, and the routing plan is derived in the configuration space, taking the motion of the robot into account. It is experimentally demonstrated that the proposed method is effective for various problems, despite the fact that each route for an AGV is created without minimizing the entire objective function.**

***Index Terms*—Distributed autonomous robotics systems, Lagrangian decomposition and coordination, multiple mobile robots, route planning, semiconductor fabrication bays.**

## I. Introduction

Multiple automated guided vehicles (AGVs) are widely used for transportation in semiconductor fabrication bays. The requests for transportation are given in every few seconds, and therefore, it is required to derive a collision-free routing plan that minimizes the total transportation time in a few seconds. Motion planning of multiple mobile robots has been proven to be a difficult problem, not to be solved in polynomial time with the number of degrees of freedom possessed by the robots [16]. In recent years, the layout for transportation in semiconductor fabrication is growing rapidly with the expansion of the size of silicon wafers. The route-planning problem for multiple mobile robots becomes increasingly difficult with respect to the increase of the number of AGVs. In the field of motion planning for robotics, there have been a wide variety of studies focusing on methods such as potential fields or probabilistic roadmaps [14]. For multiple robot motion-planning problems, a number of algorithms have been proposed that take different approaches, such as graph theory [22] or iterative neighborhood search [10]. The route-planning methods for semiconductor fabrication bays have also been studied. The Petri net [8] approach or genetic algorithm [9] have been reported in the literature.

These algorithms are often characterized by centralized or decentralized approaches. The centralized approach is configured to determine

Manuscript received August 10, 2004; revised February 22, 2005. This paper was recommended for publication by Associate Editor G. Sukhatme and Editor S. Hutchinson upon evaluation of the reviewers' comments. This work was supported by the JSPS under Grant-in-Aid for Scientific Research (B)15360211. This paper was presented in part at the IEEE International Conference on Robotics and Automation, Taipei, Taiwan, R.O.C., September 2003.

The authors are with the Department of Electrical and Electronic Engineering, Graduate School of Natural Science and Technology, Okayama University, Okayama, Japan (e-mail: nishi@cntr.elec.okayama-u.ac.jp; konishi@cntr.elec.okayama-u.ac.jp).

Digital Object Identifier 10.1109/TRO.2005.853489