

TERRAIN ESTIMATION METHODS FOR ENHANCED AUTONOMOUS ROVER MOBILITY

Karl Iagnemma, Massachusetts Institute of Technology, U.S.A.

1. INTRODUCTION

Scientific goals for Mars surface exploration during the period 2009-2020 are expected to focus on understanding the planet's climate history, surface geology, and potential for past or present life. To accomplish these goals, rovers will be required to safely access highly rough terrain with a significant degree of autonomy. Terrain areas of interest might include impact craters, rifted basins, and water-carved features such as gullies and outflow channels [1,2]. Such regions are in general highly uneven and sloped, and may be covered with loose drift material that causes rover wheel slippage and sinkage.

Terrain physical properties can strongly influence rover mobility, particularly on uneven and sloped terrain [3]. For example, a rover might navigate up a rocky slope with ease, but slide down a sandy slope of the same grade. Similarly, a rover might easily traverse a region of packed soil, but become entrenched while crossing a patch of loose sand. The effect of terrain properties on rover mobility was exemplified in April–June, 2005, when NASA's Mars Exploration Rover Opportunity became entrenched in loose drift material and was immobilized for several weeks (see Figure 1). Thus, the ability to detect or estimate terrain physical conditions is essential for ensuring rover safety. Such knowledge would allow a rover to intelligently predict its mobility performance and thereby autonomously avoid terrain regions that are potentially non-traversable, such as patches of loose drift material or sloped terrain with low bearing capacity. Knowledge of terrain properties would also allow a system to adapt its control and planning strategies to enhance performance, by maximizing wheel traction or minimizing power consumption.

Terrain property information can be represented in a variety of forms. One representation is a numerical estimate of physical parameters associated with a given wheel-terrain interaction model. This includes, for example, parameters such as cohesion, internal friction angle, shear deformation modulus, etc. A second representation is a numerical assessment of terrain traversability for a given rover system. This might take the form of the value in Newtons of the available wheel thrust on a given terrain patch. Finally, a third way to represent terrain property information is by labeling a given terrain patch with a physical description (e.g. “sandy terrain,” “rocky terrain,” “rigid terrain”). Class labels can be associated with a set of nominal numerical physical parameters, and/or nominal traversability estimates.



Image credit: NASA/JPL

Figure 1. Mosaic of navigation-camera frames from NASA's Mars Exploration Rover Opportunity, presented in a vertical projection, showing the rover entrenched in a dune of highly deformable drift material.

Terrain property information can be employed during two phases of rover navigation: 1) local traversability assessment, and 2) servo-level wheel control. The current state-of-the-art in rover traversability assessment relies on analysis of geometric parameters (i.e. slope and roughness) of a terrain patch, and does not consider “non-geometric” criteria such as soil deformability, allowable thrust, etc. Integration of terrain property information with traversability assessment methods could allow detection of “soil trap” hazards (such as the one described above) and would allow improved prediction of traversability for “geometric” hazards. Integration of terrain property information with servo-level wheel control would allow a rover to enhance its performance by maximizing wheel traction or minimizing power consumption on a given terrain. Methods for accomplishing this will be described later in the chapter. The integration of terrain property information with traversability assessment and control algorithms is an area of active research. In general, terrain estimation methods are excellent candidates for implementation in flight systems since they improve intrinsic rover mobility, require no human supervision, and introduce little additional computational or operational burden.

In this chapter, algorithms for estimating terrain property information for the three representations described above are presented. This work represents research that has been performed over the past five years in the MIT Field and Space Robotics Laboratory. In all algorithms, the case of a rigid wheel traveling through rigid or deformable terrain is considered, since this is the expected condition for planetary exploration vehicles.

First, an algorithm for on-line estimation of terrain cohesion and internal friction angle is presented. The algorithm relies on a simplified form of classical terramechanics equations, and uses a linear-least squares method to estimate terrain parameters in real time. Second, a method

for on-line terrain traversability assessment is presented. The method also relies on a simplified wheel-terrain interaction model, and results in a numerical estimate of the available wheel thrust for a terrain patch. Finally, a method to classify terrain based on vibrations present in the rover structure during driving is presented. The algorithm relies on data from an accelerometer mounted to the rover structure, and employs linear discriminant analysis to compare on-line vibration signatures to those gathered during an off-line training phase. Experimental results from a laboratory wheel-terrain interaction testbed are presented that demonstrate the validity of the three algorithms.

The key requirements of the three algorithms described in this chapter are their accuracy, computational efficiency, and reliance on sensors that are likely to be part of a “standard” sensor suite for future planetary rover systems. Computational efficiency is particularly important, since future rovers can be assumed to have highly constrained computational resources. It should be noted that an algorithm’s reliance on standard rover sensors is highly desirable, since it requires no addition of volume, weight, hardware complexity and cost to the rover system.

2. TERRAIN PARAMETER ESTIMATION

In this section a method for on-line estimation of two key terrain parameters, cohesion (c) and internal friction angle (ϕ), is briefly described. A complete presentation of this material can be found in [4]. Cohesion and internal friction angle are important parameters since they can be used to compute the maximum shear strength, τ_{\max} , of a terrain region, from Coulomb’s equation:

$$\tau_{\max} = c + \sigma_{\max} \tan \phi \quad (1)$$

where σ_{\max} is the maximum normal stress acting on a terrain region. Since soil failure occurs when the maximum shear strength is exceeded, knowledge of c and ϕ can be used to perform model-based prediction of rover traversability on flat and sloped terrain. The parameters c and ϕ are also important scientific parameters for characterizing soils.

To estimate terrain parameters, equations relating the parameters of interest to physically measurable quantities must be developed. A free-body diagram of a driven rigid wheel of radius r and width b traveling through deformable terrain is shown in Figure 2. A vertical load W and a horizontal force DP are applied to the wheel by the vehicle suspension. A torque T is applied at the wheel rotation axis by an actuator. The wheel has angular velocity ω , and the wheel center has linear velocity, V . The angle from the vertical at which the wheel first makes contact with the terrain is denoted θ_1 . The angle from the vertical at which the wheel loses contact with the terrain is denoted θ_2 . Thus, the entire angular wheel-terrain contact region is defined by $\theta_1 + \theta_2$.

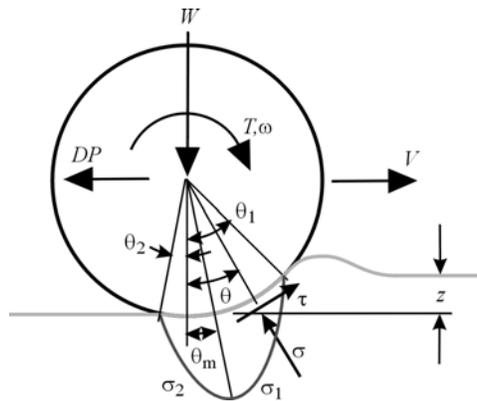


Figure 2. Free-body diagram of rigid wheel on deformable terrain.

A stress region is created at the wheel-terrain interface and indicated by the regions σ_1 and σ_2 . Here σ_1 is the interface section from initial terrain contact (i.e. θ_1) to the point of maximum stress (i.e. θ_m), and σ_2 is the region from the point of maximum stress to final terrain contact (i.e. θ_2). At a given point on the interface, the stress can be decomposed into a component acting normal to the wheel at the wheel-terrain contact point, σ , and a component acting tangential to the wheel at the wheel-terrain contact point, τ . The angle from the vertical at which the maximum stress occurs is denoted θ_m .

It is assumed that the following quantities are known: the vertical load W , torque T , sinkage z , wheel angular speed ω , and wheel linear speed V . Issues related to sensing are discussed later in this section. Force balance equations for the system in Figure 2 can be written as:

$$W = rb \left(\int_{\theta_1}^{\theta_2} \sigma(\theta) \cos \theta \cdot d\theta + \int_{\theta_1}^{\theta_2} \tau(\theta) \sin \theta \cdot d\theta \right), \quad (2)$$

$$DP = rb \left(\int_{\theta_1}^{\theta_2} \tau(\theta) \cos \theta \cdot d\theta - \int_{\theta_1}^{\theta_2} \sigma(\theta) \sin \theta \cdot d\theta \right), \quad (3)$$

$$T = r^2 b \int_{\theta_1}^{\theta_2} \tau(\theta) \cdot d\theta. \quad (4)$$

The shear stress at the wheel-terrain interface can be computed as:

$$\tau(\theta) = (c + \sigma(\theta) \tan \phi) \left(1 - e^{-\frac{r}{k} [\theta_1 - \theta - (1-i)(\sin \theta_1 - \sin \theta)]} \right) \quad (5)$$

where k is the shear deformation modulus, r is the wheel radius, and i is the wheel slip, defined as $i = 1 - (V/r\omega)$ [5]. The normal stress at the wheel-terrain interface is given by:

$$\sigma(z) = \left(\frac{k_c}{b} + k_\phi \right) z^n \quad (6)$$

where b is the wheel width, k_c and k_ϕ are pressure-sinkage moduli, and n is the sinkage exponent [6]. This equation can be expressed as a function of the wheel angular location θ by noting that sinkage is related to θ as $z(\theta) = r(\cos \theta - \cos \theta_1)$. Substituting this relation into Eq. (6) yields expressions for the normal stress distribution along the wheel-terrain interface, as:

$$\sigma_1(\theta) = \left(\frac{k_c}{b} + k_\phi \right) (r(\cos \theta - \cos \theta_1))^n, \quad (7)$$

$$\sigma_2(\theta) = \left(\frac{k_c}{b} + k_\phi \right) \left(r \left(\cos \left(\theta_1 - \theta \frac{(\theta_1 - \theta_m)}{\theta_m} \right) - \cos \theta_1 \right) \right)^n. \quad (8)$$

To develop a parameter estimation algorithm, analytical expressions for the force balance equations (Eqs. (2-4)) are required, since these equations relate physically measurable quantities (W , T , z , ω , V) to the parameters of interest (c , ϕ). However, Eqs. (2-4) are not amenable to closed-form integration, due to their complexity. This motivates the use of an approximate form of the stress equations (Eqs. (5, 7, 8)).

2.1 Equation Simplification

Figure 3 shows typical simulated plots of shear and normal stress distributions (as defined by Eqs. (5) and (7-8), respectively) around the rim of a driven rigid wheel on various terrains at moderate wheel slip. Parameters used in these plots represent a diverse range of terrain types: dry sand, sandy loam, clayey soil, and snow [7,8]. Figure 3 shows that the shear and normal stress distribution curves are approximately linear for a diverse range of terrains.

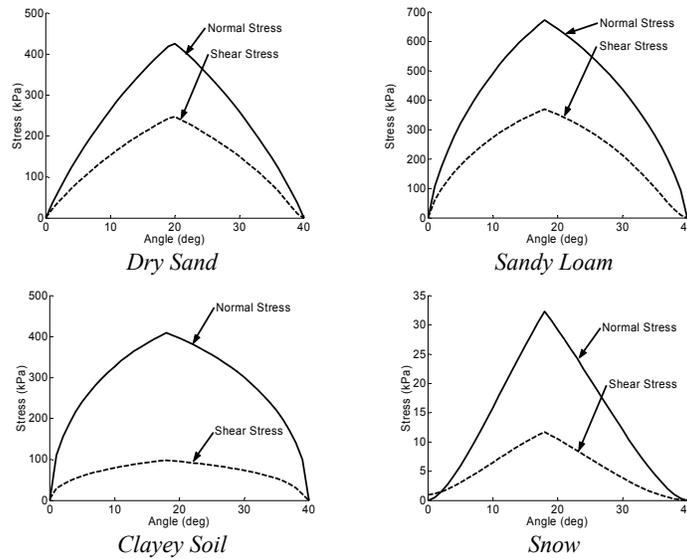


Figure 3. Normal and shear stress distribution for various terrain types at moderate wheel slip ($i = 0.5$)

Based on this observation, linear approximations of the shear and normal stress equations can be written as:

$$\sigma_1^L(\theta) = \frac{\theta_1 - \theta}{\theta_1 - \theta_m} \sigma_m, \quad \sigma_2^L(\theta) = \frac{\theta}{\theta_m} \sigma_m, \quad \tau_1^L(\theta) = \frac{\theta_1 - \theta}{\theta_1 - \theta_m} \tau_m, \quad \tau_2^L(\theta) = c + \frac{\theta}{\theta_m} (\tau_m - c) \quad (9)$$

where σ_m and τ_m are the maximum values of the normal and shear stress, respectively.

Simplified forms of the force balance equations can be written by combining Eqs. (2-4) with Eq. set (9) (with $\theta_2 = 0$, since θ_2 is generally small in practice). Evaluation of the linearized forms of (2-4) yields the following expressions for the normal load and torque:

$$W = \frac{rb}{\theta_m(\theta_1 - \theta_m)} \begin{bmatrix} \sigma_m(-\theta_m \cos \theta_1 + \theta_1 \cos \theta_m - \theta_1) \\ -\tau_m(\theta_m \sin \theta_1 - \theta_1 \sin \theta_m) \\ -c(\theta_1 \sin \theta_m - \theta_m \sin \theta_m - \theta_m \theta_1 + \theta_m^2) \end{bmatrix}, \quad (10)$$

$$T = \frac{r^2 b}{2} (\tau_m \theta_1 + c \theta_m). \quad (11)$$

Two assumptions are made in solving Eqs. (10) and (11) for c and ϕ . The first is that the location of the maximum shear and normal stress occurs at the same location, θ_m . Analysis and simulation have shown that this assumption is reasonable for a wide range of soil types [9]. With this assumption, an additional relation can be written, based on Eq. (5):

$$\tau_m = (c + \sigma_m \tan \phi) \left(1 - e^{-\frac{r}{k} [\theta_1 - \theta_m - (1-i)(\sin \theta_1 - \sin \theta_m)]} \right). \quad (12)$$

The second assumption is that the angular location of maximum stress, θ_m , occurs midway between θ_1 and θ_2 , i.e. $\theta_m = (\theta_1 + \theta_2)/2$. This assumption is reasonable for a wide range of soils at moderate slip ratios [9]. This can be justified by noting that θ_m can be estimated from the relation $\theta_m = (c_1 + ic_2)\theta_1$, where c_1 and c_2 are terrain parameters. The range of c_1 and c_2 is generally $c_1 \approx 0.4$ and $0 \leq c_2 \leq 0.3$ [10].

The system of Eqs. (10-12) can be combined into a single equation relating cohesion and internal friction angle, as:

$$c = \frac{\kappa_1 \tan \phi + \kappa_2}{\kappa_3 \tan \phi + \kappa_4} \quad (13)$$

where

$$\begin{aligned} \kappa_1 &= A(\theta_1^2 W r + 4T \sin \theta_1 - 8T \sin(\theta_1/2)), \\ \kappa_2 &= 4T(\cos \theta_1 - 2 \cos(\theta_1/2) + 1), \\ \kappa_3 &= A \theta_1 r^2 b (\sin \theta_1 - 4 \sin(\theta_1/2) + \theta_1) \\ \kappa_4 &= \theta_1 r^2 b (\cos \theta_1 - 2 \cos(\theta_1/2) + 2A \cos \theta_1 - 4A \cos(\theta_1/2) + 2A + 1) \end{aligned}$$

and $A = 1 - e^{-\frac{r}{k} \left[\frac{\theta_1}{2} + (1-i) \left(-\sin \theta_1 + \sin \left(\frac{\theta_1}{2} \right) \right) \right]}$.

Eq. (13) can be rearranged to the following form:

$$\frac{\kappa_2}{\kappa_4} = \frac{\kappa_3}{\kappa_4} c \tan \phi + c - \frac{\kappa_1}{\kappa_4} \tan \phi. \quad (14)$$

The relative contribution of each term in the right-hand side of Eq. (14) was studied numerically over a wide range of parameters that encompasses a broad variety of terrain types [4]. It was found that the maximum contribution of the $(\kappa_3/\kappa_4)c \tan \phi$ term was small compared to the other terms. Thus this term is negligible, and Eq. (14) can be reduced to:

$$\frac{\kappa_2}{\kappa_4} = c - \frac{\kappa_1}{\kappa_4} \tan \phi. \quad (15)$$

Eq. (15) is a single equation in two unknowns. At least two unique instances of Eq. (15) are required to compute c and ϕ . During the parameter estimation process, it is expected that sensor data would be sampled at a frequency of several hertz. For each unique sample j , a unique

instance of Eq. (15) can be written, and the system of equations can be represented as:

$$\mathbf{K}_1 = \mathbf{K}_2 \begin{bmatrix} c \\ \tan \phi \end{bmatrix} \quad (16)$$

with $\mathbf{K}_1 = [\kappa_2^1/\kappa_4^1 \quad \dots \quad \kappa_2^j/\kappa_4^j]^T$, $\mathbf{K}_2 = \begin{bmatrix} 1 & \dots & 1 \\ -\kappa_1^1/\kappa_4^1 & \dots & -\kappa_1^j/\kappa_4^j \end{bmatrix}^T$.

In practice, more than two equations are used to form an estimate of c and ϕ to decrease sensitivity to sensor noise. In this case \mathbf{K}_2 is non-square and Eq. (16) can be solved in a least-squares sense, using the pseudoinverse of \mathbf{K}_2 :

$$\begin{bmatrix} c \\ \tan \phi \end{bmatrix} = (\mathbf{K}_2^T \mathbf{K}_2)^{-1} \mathbf{K}_2^T \mathbf{K}_1. \quad (17)$$

Note that singularity of $(\mathbf{K}_2^T \mathbf{K}_2)^{-1}$ only occurs in the degenerate case where non-unique sensor data is sampled (such as on perfectly flat terrain). All quantities in Eq. (17) can be sensed except the shear deformation modulus k (in the matrices \mathbf{K}_1 and \mathbf{K}_2). In practice, the estimation algorithm exhibits low sensitivity to k , particularly for large wheel radii and high slip ratios. Thus k is chosen as a representative value for deformable terrain.

2.2 Sensing and Implementation Issues

In the preceding analysis it was assumed that the vertical load W , torque T , sinkage z , wheel angular speed ω , and wheel linear speed V could be measured or estimated. Here, methods for measuring or estimating these inputs are discussed, along with other implementation issues.

The vertical load W can be computed from a quasi-static force analysis of the rover, with knowledge of the rover configuration and mass distribution. Quasi-static analysis is valid since dynamic effects are negligible at the low speeds of these vehicles (on the order of several cm/sec). The torque T can be estimated from the current input to the motor and an empirically-determined mapping from current to torque. In applications where large thermal variation is expected (such as Martian surface exploration), motor temperature can be included in this mapping [11]. Note that torque and vertical load could be directly measured if the wheel were instrumented with a multi-axis force sensor. However, this adds cost and complexity to the rover.

The sinkage z can be computed with vision-based techniques or (in some cases) by kinematic analysis of the rover suspension [12,13]. The wheel angular speed ω can be measured with a tachometer. The wheel linear speed V can be computed using inertial measurement unit (IMU) measurements. However, at low speeds IMU velocity measurements can be highly degraded by noise. In this case, visual odometry can yield more accurate results [14].

The sensors described above (i.e. rover configuration sensors, motor current sensor, wheel tachometer, IMU, and vision system) would likely be part of a planetary exploration rover. Thus, all required inputs can be measured or estimated with on-board rover sensors.

An important implementation issue is minimizing sensor noise. Most rover sensors can be modeled as a “true” signal corrupted by white noise. In this case, increasing the number of data points in Eq. (17) acts as an averaging filter and improves estimation accuracy. Other filtering techniques (such as the Kalman filter and its extensions) could also be applied to this problem.

Note that all data points used in a parameter estimate are assumed to be sampled from homogeneous terrain. For example, consider a rover moving at 5 cm/sec with a sensor sampling rate of 5 Hz. If 10 data points are used to compute a parameter estimate, it must be assumed that the terrain is homogeneous within a 10 cm distance. If data is sampled from mixed or

inhomogeneous terrain, resulting parameter estimates will be effective estimates of the combined terrain types. In general, assumptions regarding terrain homogeneity can be formed from *a priori* knowledge of local terrain characteristics. Terrain classification methods could also be used to detect changes in terrain type.

A final implementation issue arises for a rover traveling at constant velocity on flat terrain, where the matrix inverse in Eq. (17) may be poorly conditioned. This occurs because the rover is collecting an identical set of sensor readings at each sampling instance. The ridge regression technique can be used to solve Eq. (17) in these cases:

$$\begin{bmatrix} c \\ \tan \phi \end{bmatrix} = (\mathbf{K}_2^T \mathbf{K}_2 + \delta \mathbf{I})^{-1} \mathbf{K}_2^T \mathbf{K}_1 \quad (18)$$

where δ is a small positive constant that can be optimized by techniques such as cross-validation [15]. In practice (and in the results presented below), natural terrain variation will usually lead to acceptable equation conditioning. Deliberately inducing variable wheel slip also improves parameter estimates on flat terrain.

3. TERRAIN TRAVERSABILITY ASSESSMENT

In this section a method for terrain traversability assessment is briefly described. This work differs from the parameter estimation approach presented in Section 2, in that it does not attempt to estimate parameter values associated with a particular wheel-terrain interaction model. Rather, it derives an explicit estimate of the traversability of a terrain patch by a given rover. This estimate takes the form of the value in Newtons of the available wheel thrust on a given terrain patch.

In this approach, the robot “feels” the terrain by driving a single wheel and measuring the torque and resulting sinkage. The robot is stopped or moving slowly during this process. Measured wheel torque and sinkage are analyzed to predict the available drawbar pull (i.e. net forward thrust) based on a simplified version of the terramechanics model employed in Section 2. A large positive drawbar pull implies that the terrain is easy to traverse. A small or negative drawbar pull implies that the vehicle cannot easily traverse the terrain and may become entrapped. Thus, knowledge of drawbar pull is an important part of traversability analysis.

3.1 Algorithm Detailed Description

Consider again the free-body diagram of a driven rigid wheel in deformable terrain shown in Figure 2. In traversability assessment we are interested in estimating the drawbar pull force DP acting on the rover suspension by the wheel. Direct computation of Eq. (3) would require knowledge of numerous terrain parameters. However, as shown above the shear and normal stress distributions can be approximated by linear functions with reasonably good accuracy for a wide range of terrains. Based on this observation, an analytical expression for drawbar pull can be written by integrating the linearized forms of the shear and normal stress distributions (Eq. set (9)) over the wheel-terrain contact area, as:

$$DP = \frac{2rb}{\theta_1} \left(\left(2 \cos\left(\frac{\theta_1}{2}\right) - \cos(\theta_1) - 1 \right) \tau_m - \left(2 \sin\left(\frac{\theta_1}{2}\right) - \sin(\theta_1) \right) \sigma_m \right) \quad (19)$$

Again, it is assumed that the angular location of maximum stress, θ_m , occurs at $\theta_1/2$. With Eqs. (10) and (11), expressions for σ_m and τ_m can be derived, as:

$$\tau_m = \frac{2T}{r^2 b \theta_1} \quad (20)$$

$$\sigma_m = \frac{\frac{W\theta_1}{2rb} - \frac{2T}{r^2 b \theta_1} \left(2 \sin\left(\frac{\theta_1}{2}\right) - \sin(\theta_1) \right)}{2 \cos\left(\frac{\theta_1}{2}\right) - \cos(\theta_1) - 1} \quad (21)$$

Eqs. (19-21) can then be combined, as:

$$DP = \frac{4T}{\theta_1^2 r} \left(\frac{\left(2 \cos\left(\frac{\theta_1}{2}\right) - \cos(\theta_1) - 1 \right)^2 + \left(2 \sin\left(\frac{\theta_1}{2}\right) - \sin(\theta_1) \right)^2}{2 \cos\left(\frac{\theta_1}{2}\right) - \cos(\theta_1) - 1} \right) - \frac{W \left(2 \sin\left(\frac{\theta_1}{2}\right) - \sin(\theta_1) \right)}{2 \cos\left(\frac{\theta_1}{2}\right) - \cos(\theta_1) - 1} \quad (22)$$

Eq. (22) describes the drawbar pull as a function of the normal load on the wheel (W), the wheel torque (T) and angular wheel sinkage θ_1 . In real world situations, the sinkage is generally smaller than 60° , which allows accurate small angle approximations to be written for the second-order polynomials in Equation (7), as:

$$2 \sin\left(\frac{\theta_1}{2}\right) - \sin(\theta_1) \cong \frac{\theta_1^3}{8}, \quad 2 \cos\left(\frac{\theta_1}{2}\right) - \cos(\theta_1) - 1 \cong \frac{\theta_1^2}{4}$$

With these simplifications, drawbar pull can be expressed as:

$$DP = \frac{T}{r} \left(1 + \frac{\theta_1^2}{4} \right) - \frac{W\theta_1}{2} \quad (22)$$

and approximated in dimensionless form as:

$$\frac{DP_{est}}{W} \cong \frac{T}{rW} - \frac{\theta_1}{2} \quad (23)$$

This result agrees with empirical observations relating torque, sinkage, and drawbar pull, namely that terrain with low wheel sinkage and high torque (such as hardpack or pavement) generally allows large drawbar pull, and terrain with high sinkage and low torque (such as sand or fine soil) generally allows small or negative drawbar pull. This implies that knowledge of sinkage and torque are important indicators of drawbar pull.

The preceding analyses have been developed for the case of rigid wheels interacting with deformable terrain. For the case of rigid terrain, $\theta = 0$ and Eq. (23) yields $DP = T/r$, which is a reasonable model for the rigid wheel/rigid terrain case. This model's accuracy would be reduced if $T/r > \mu N$, since in this case the measured torque would include contributions from both terrain interaction and wheel angular acceleration. However the rover systems under consideration are expected to travel at speeds of several cm/sec. This suggests that the proposed method can be applied to rigid wheels operating in both rigid and deformable terrain.

3.2 Analysis of Drawbar Pull Prediction

The preceding analysis suggests that DP can be estimated with knowledge of T , W , and z , and without explicit knowledge of terrain parameters (i.e. cohesion, internal friction angle, etc.) This implies that DP is relatively sensitive to perturbations in T , W , and z and relatively insensitive to perturbations in the terrain parameters. Analysis was performed to verify this hypothesis.

A dynamic simulation of a single driven wheel traveling through deformable terrain was developed. The wheel had radius r of 0.1 m and width b of 0.1 m. Simulations were conducted at various slip ratios from 0.1-0.9 and weights of 30 and 80 N. A proportional-derivative control algorithm commanded the wheel. The simulated sampling rate was 20 Hz. Simulations were performed on four distinct terrain types, each defined by a range of parameters. Terrain parameter ranges and further details of the simulation are given in [9]. The estimated drawbar pull of Eq. (23) compared to DP computed from the nonlinear Eqs. (3,5,6).

A set of 2×10^5 simulations was performed. For each simulation the error e was computed as:

$$e = \frac{DP}{W} - \frac{DP_{est}}{W} \quad (24)$$

Here DP has been non-dimensionalized by the wheel normal load W . It was observed that the drawbar pull estimation systematically underestimated the true drawbar pull by a small amount. The mean error in this analysis was -0.102 and the standard deviation was 0.031. The mean value of DP/W over all simulations was 0.481, and thus the mean error of DP_{est} was 21.2%.

It is hypothesized that the systematic error in DP_{est} was due to the fact that the accuracy of the linear approximation employed here decreases as the soil cohesion increases [9]. To reduce the estimation error for high-cohesion soils, a modified version of DP_{est} was proposed as:

$$\frac{DP'_{est}}{W} \cong K_1 \frac{T}{rW} - K_2 \frac{\theta_1}{2} + K_3 \quad (25)$$

The constant ‘‘correction factors’’ K_1 , K_2 , and K_3 were determined via a least-squares fit from data of the previous analysis, as:

$$\mathbf{AK} = \mathbf{B} \quad (26)$$

$$\mathbf{K} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{B} \quad (27)$$

where

$$\mathbf{A} = \begin{bmatrix} \left(\frac{T}{rW}\right)_1 & \left(\frac{z}{r}\right)_1 & 1 \\ \vdots & \vdots & \vdots \\ \left(\frac{T}{rW}\right)_n & \left(\frac{z}{r}\right)_n & 1 \end{bmatrix}, \quad \mathbf{K} = \begin{bmatrix} K_1 \\ K_2 \\ K_3 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \left(\frac{DP}{W}\right)_1 \\ \vdots \\ \left(\frac{DP}{W}\right)_n \end{bmatrix}$$

and $n = 2 \times 10^5$. The correction factors were determined as $K_1 = 1.02$, $K_2 = 0.99$, and $K_3 = -0.09$.

Using these modified values in Eq. (27), a new mean prediction error was computed as:

$$e = \frac{DP}{W} - \frac{DP'_{est}}{W} \quad (28)$$

The mean error in this analysis was -0.003 and the standard deviation was 0.031. Thus the expected error of DP_{est}^i over all simulations was 0.7%. This suggests that the modified drawbar pull prediction equation can yield accurate results in a variety of terrain.

4. CONTACT-BASED TERRAIN CLASSIFICATION

In Sections 2 and 3 methods were presented for representing terrain property information via estimates of terrain parameter values and assessments of terrain traversability, respectively. A third way to represent terrain property information is by defining a class label for a given terrain patch. Class labels generally correspond to generic physical descriptions (e.g. “sandy terrain,” “rocky terrain,” “rigid terrain”), and can be associated with a set of nominal numerical physical parameters, and/or traversability estimates for use in traversability prediction and servo-level wheel control.

In this section a method for classifying terrain based on an analysis of a rover’s wheel-terrain interaction vibration signature is presented. A complete presentation of this material can be found in [16]. It should be noted that this approach to terrain classification is unique in that it relies on information generated by physical wheel-terrain interaction. This is in contrast to most classification methods, which rely on non-contact, or remote (i.e. machine vision-based) sensing. Classification algorithms relying on visual features are often sensitive to changes in illumination.

The algorithm presented here learns to recognize distinct terrain types based on labeled vibration data provided during an off-line 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 on-line classifier. During on-line classification, measured vibration signals are quickly classified as one of the labeled terrain classes. An overview schematic of the algorithm is shown in Figure 4.

In this approach, vibration signals are first divided into short segments. These are then converted from time-domain voltage signals into power spectral densities. 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 pair-wise classifier [18,19]. In this approach, sub-classifiers are created to classify a sample signal as being associated with one of two possible terrains. Separate sub-classifiers 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, 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 classification problem [20].

To accommodate classification of more than two terrains, a voting scheme is used. Each pair-wise classifier “votes” for one of the two terrains it distinguishes, or remain “undecided.” The winning terrain class is returned and taken as the best estimate of the terrain class.

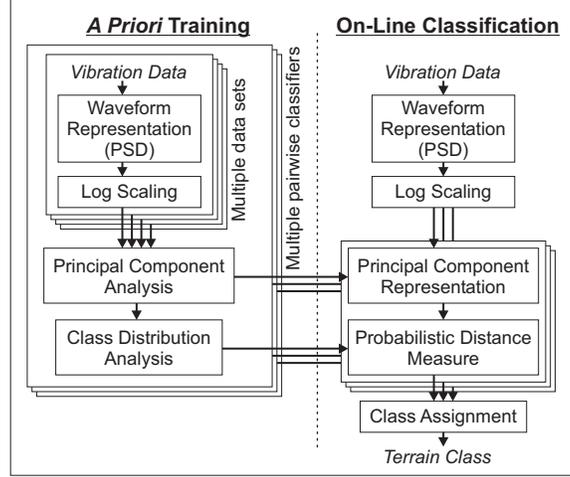


Figure 4. Overview flowchart for vibration-based terrain classification algorithm

4.1 Algorithm Detailed Description

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

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, and load).

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, grouser spacing, spatial variations of terrain, etc.). The power spectral density (PSD) of each of these segments is computed using Welch's method and a log-scaled PSD is stored in a matrix [21].

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

$$\mathbf{Y}_{sand} = \begin{bmatrix} y_{sand, fmin,1} & \cdots & y_{sand, fmin,n} \\ \vdots & \ddots & \vdots \\ y_{sand, fmax,1} & \cdots & y_{sand, fmax,n} \end{bmatrix}. \quad (29)$$

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, $\mathbf{Y}_{sand} \in \mathfrak{R}^{m \times n_{sand}}$ and $\mathbf{Y}_{gravel} \in \mathfrak{R}^{m \times n_{gravel}}$, are combined to form a complete record of the data: $\mathbf{Y} = [\mathbf{Y}_{sand} \quad \mathbf{Y}_{gravel}]$, $\mathbf{Y} \in \mathfrak{R}^{m \times n}$, where m is the number of frequency components, and $n = n_{sand} + n_{gravel}$ is the total number of time segments. The row mean is then subtracted, to produce the matrix \mathbf{Y} .

Singular value decomposition is used to separate $\hat{\mathbf{Y}}$ into the matrices, \mathbf{U}_a , \mathbf{S}_a , and \mathbf{V}_a [22]:

$$\hat{\mathbf{Y}} = \mathbf{U}_a \mathbf{S}_a \mathbf{V}_a^T. \quad (30)$$

Only the first p columns of \mathbf{U}_a (i.e. the first p principal components of $\hat{\mathbf{Y}}$) and the upper-left $p \times p$ block of \mathbf{S}_a are retained, in matrices \mathbf{U}_{signal} and \mathbf{S}_{signal} . The value for p was selected empirically to represent only the subspace of \mathbf{U}_a related to signal (rather than noise) information. Systematic methods for choosing p are available [23]. 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 $p = 15$, as it appears to give good signal representation without overfitting. In experiments, the first 15 principal components account for approximately 90% of the training data variance.

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

$$\mathbf{W}_{sand} = \mathbf{S}_{signal}^{-1} \mathbf{U}_{signal}^T \mathbf{Y}_{sand} \quad (31)$$

$$\mathbf{W}_{gravel} = \mathbf{S}_{signal}^{-1} \mathbf{U}_{signal}^T \mathbf{Y}_{gravel}. \quad (32)$$

For all practical cases, if $n > p$, \mathbf{S}_{signal} will be invertible. If $\text{rank}(\mathbf{S}) < p$, a smaller value for p 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 $\bar{\mathbf{w}}_{sand}$ and $\bar{\mathbf{w}}_{gravel}$. These means are then subtracted from \mathbf{W}_{sand} and \mathbf{W}_{gravel} to produce the matrices $\hat{\mathbf{W}}_{sand}$ and $\hat{\mathbf{W}}_{gravel}$. These matrices are scaled by the number of points in the training data and are merged to form the matrix $\hat{\mathbf{W}}$:

$$\hat{\mathbf{W}} = \left[\frac{1}{\sqrt{n_{sand} - 1}} \hat{\mathbf{W}}_{sand} \quad \frac{1}{\sqrt{n_{gravel} - 1}} \hat{\mathbf{W}}_{gravel} \right] \quad (33)$$

This matrix is then decomposed into its singular value representation:

$$\hat{\mathbf{W}} = \mathbf{U}_b \mathbf{S}_b \mathbf{V}_b^T. \quad (34)$$

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

$$\mathbf{d} = \mathbf{U}_b \mathbf{S}_b^{-1} \mathbf{S}_b^{-1} \mathbf{U}_b^T (\bar{\mathbf{w}}_{sand} - \bar{\mathbf{w}}_{gravel}) \quad (35)$$

This can be shown to be the Fisher linear discriminant in the p -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:

$$d(\mathbf{y}) = \mathbf{d}^T \mathbf{S}_{signal}^{-1} \mathbf{U}_{signal}^T \mathbf{y}, \quad (36)$$

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

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 \bar{d}_{sand} , σ_{sand} , \bar{d}_{gravel} , and σ_{gravel} . The discrimination vector and the terrain class statistics are stored for use in the on-line classification phase of the algorithm.

On-Line 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 power spectral density is computed, and the magnitude is log-scaled and stored in a vector \mathbf{y} .

Pair-wise classifiers then compute the discrimination metric $d(\mathbf{y})$. The Mahalanobis distances from $d(\mathbf{y})$ to the terrain class means, e.g. \bar{d}_{sand} and \bar{d}_{gravel} , are then computed as [24]:

$$md_{sand}(\mathbf{y}) = \frac{|d(\mathbf{y}) - \bar{d}_{sand}|}{\sigma_{sand}} \quad (37)$$

$$md_{gravel}(\mathbf{y}) = \frac{|d(\mathbf{y}) - \bar{d}_{gravel}|}{\sigma_{gravel}} \quad (38)$$

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

A voting scheme merges the results of the various pair-wise classifiers. In this approach, each pair-wise classifier may return a terrain label, or it may return “unknown.” If a pair-wise classifier returns a positive vote for a terrain class, the alternative terrain class receives a negative vote. If the pair-wise 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 assertion that returning “unknown” is preferable to returning the wrong terrain class. Figure 5(a) shows an example of the voting algorithm employed in a three-terrain classifier positively identifying gravel. Figure 5(b) shows an example of the voting algorithm unable to positively identify a terrain.

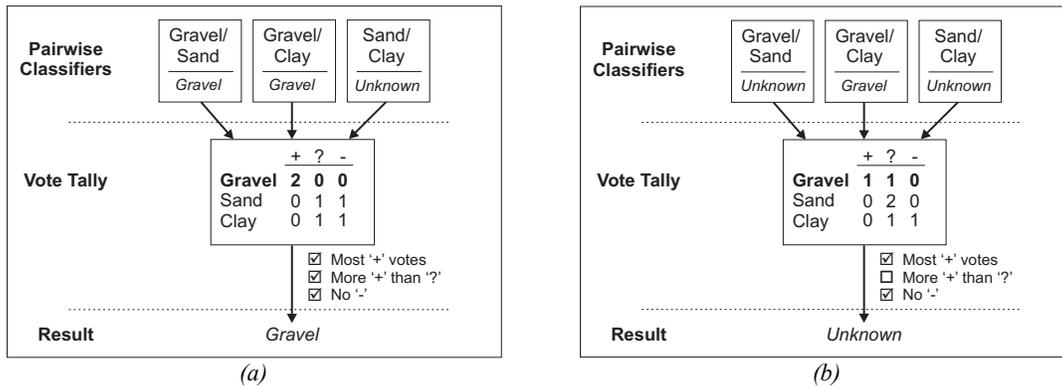


Figure 5. Voting scheme positively identifying gravel (a) and resulting in unknown result (b)

5. EXPERIMENTAL RESULTS

In this section representative experimental results for the three terrain estimation methods are presented. More extensive simulation and experimental results are described in [3,4,9,16].

5.1 Experimental Testbed Description

Experiments were performed on a laboratory terrain characterization testbed, shown in Figure 6. The testbed consists of a driven rigid wheel mounted on an undriven vertical axis. The wheel assembly is mounted to a driven horizontal carriage. By driving the wheel and carriage at

different rates, variable slip ratios can be imposed. The vertical wheel load can be changed by adding weight to the vertical axis.

The testbed is instrumented with encoders to measure angular velocities of both the wheel and the carriage pulley. The carriage linear velocity is computed from the carriage pulley angular velocity. The vertical wheel sinkage is measured with a linear potentiometer. The wheel torque T is measured by a Cooper Instruments torque sensor. The six-component wrench between the wheel and carriage is measured with a JR3 six-axis force/torque sensor. The force sensor allows measurement of the normal load W and drawbar pull DP . Vibration signals arising from wheel-terrain interaction are measured using a contact microphone mounted to the wheel frame. The testbed is controlled by a 133 Mhz PC. The soil bin is 90 cm long, 30 cm wide and 15 cm deep.

For the experiments described below, several terrains were used including landscaping gravel, JSC Mars-1 Soil Simulant, washed beach sand, topsoil, and dried bentonite clay. 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 [26]. It contains fine particles as well as larger solid particles ranging up to 4 cm.

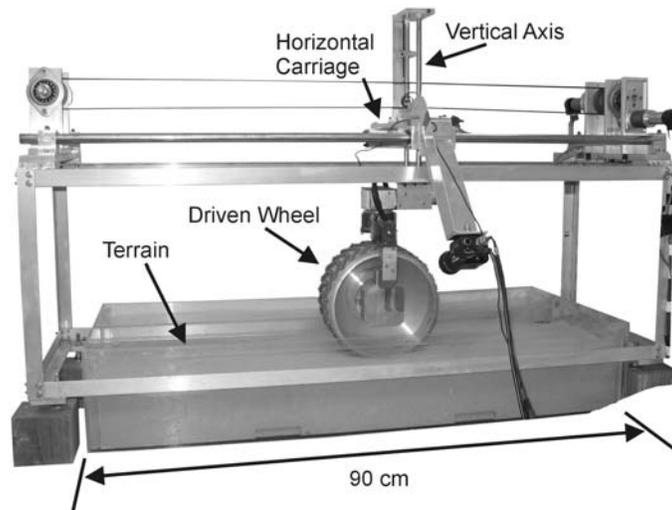


Figure 6. FSRL Wheel-Terrain Interaction Testbed

5.2 Terrain Parameter Estimation Experimental Results

Terrain parameter identification experiments were performed on sand, dried bentonite clay, and compacted moist topsoil. Classical shear failure experiments were first performed to determine c and ϕ for all soils. In these experiments, a vertical load is applied to homogeneous soil through a device called a bevameter, which is translated horizontally until shear failure occurs [6]. By varying the vertical load, a relationship between normal stress and shear stress can be observed. Numerous experiments were run to account for nonuniformity in soil mixing and moisture content. Table 1 summarizes the results of these experiments. These results agree with published values for similar soils [5,10]. A result of shear failure experiments for dried bentonite clay can be seen in Figure 7. This figure shows a linear relationship between normal stress and shear stress, as expected.

	c (kPa)	ϕ (deg)
Washed Sand	0.65 ± 0.24	32.1 ± 2.82
Dried Bentonite Clay	0.48 ± 0.19	33.7 ± 1.99
Compacted Topsoil	0.74 ± 0.24	44.3 ± 2.01

Table 1. Results from shear failure experiments.

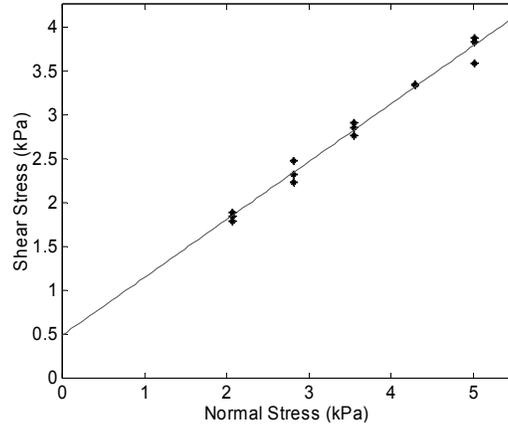


Figure 7. Shear failure experiments for dried bentonite clay. Data points shown with best-fit line

On-line estimation of cohesion and internal friction angle was the performed using the algorithm outlined in Section 2. For these experiments the shear deformation modulus k was assumed to be 0.05. 30 sampled data points were used to compute each parameter estimate.

Figure 8 shows the results of the estimation algorithm for dried bentonite. The estimated cohesion and internal friction angle rapidly converge to values of approximately 0.70 kPa and 32.1° , respectively. These values lie near or within the variance observed in classical shear failure experiments [4].

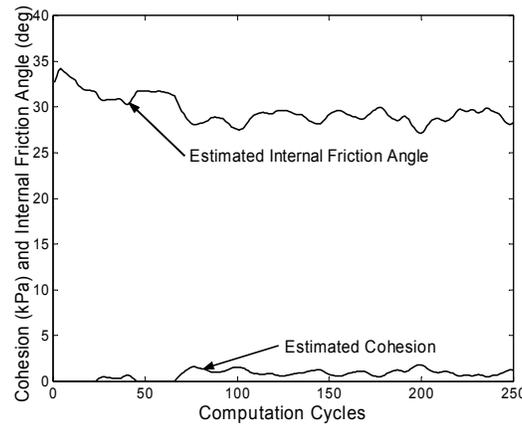


Figure 8. Estimation of cohesion and internal friction angle estimation for dried bentonite

Table 2 summarizes the results of numerous similar experiments. The estimated values are similar to those measured via shear failure experiments. This suggests that the proposed approach can identify c and ϕ of various soils despite noisy sensors. Error and variation in estimated parameters is likely due to nonuniformity in soil mixing and moisture content. Note that this variation is not unique to the proposed method, but is present in any terrain measurement technique. Estimation error is due to sensor noise and error in the assumed value of k . Also, at high slip ratios the testbed wheel exhibited control chatter, which degraded sensor readings.

	c (kPa)	ϕ (deg)
Washed Sand	0.77 ± 0.48	29.6 ± 1.47
Dried Bentonite Clay	0.70 ± 0.35	32.1 ± 2.60
Compacted Topsoil	1.04 ± 0.43	43.7 ± 3.11

Table 2. Results from terrain parameter estimation experiments.

5.3 Terrain Traversability Assessment Experimental Results

Terrain traversability assessment experiments were performed on dried bentonite clay and moist topsoil. Experiments were first performed to characterize the available DP of each terrain. The testbed wheel was driven under a weight of 30 N at 5.0 cm/sec with constant slip ratio values between 0.1 - 0.8. Drawbar pull measurements were collected at steady state for each trial.

Drawbar pull prediction experiments were then performed. The wheel was driven under the same weight, speed, and slip ratio. Values of T and z were collected at steady state. These values were used to predict drawbar pull with Equation (23). The predicted drawbar pull was then compared to the drawbar pull measured by the force/torque sensor. This comparison for bentonite and topsoil is shown in Figure 9. Note that the predicted DP varies significantly, but generally tends to increase with increasing slip, as expected.

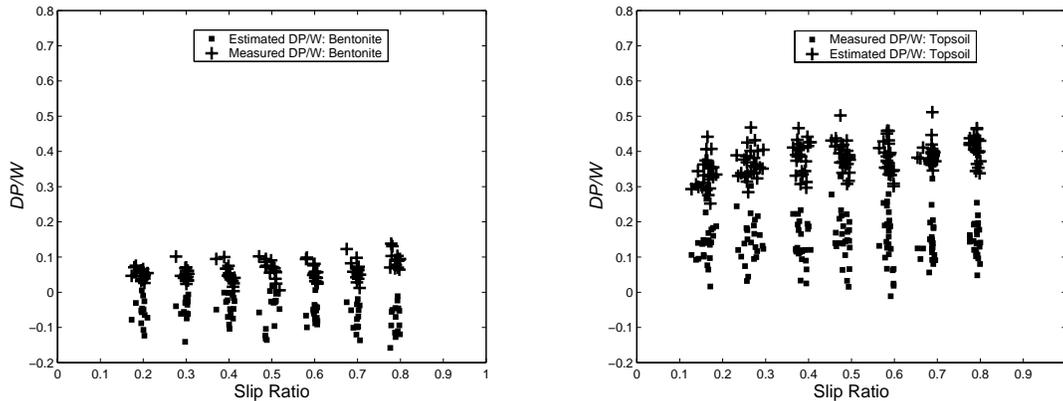


Figure 9. Comparison of predicted and measured DP for bentonite (left) and topsoil (right)

The average difference between the predicted and measured DP/W for bentonite was 0.12. The average difference between the predicted and measured DP/W for topsoil was 0.2. The negative predicted DP for bentonite implies that the wheel-terrain interaction results in a negative thrust, which would tend to decelerate a moving wheel. A wheel at rest would be unable to initiate motion in such a terrain.

These results suggests that while the proposed drawbar pull prediction method accurately captures the general relationship between wheel torque, sinkage, and normal load, correction factors must be employed to improve prediction accuracy. These correction factors compensate for errors introduced by the simplifications described in Section 3.0, and for unmodeled physical effects. One unmodeled effect is significant material transport in granular soils, which creates a wave of soil behind the wheel that contributes to the measured DP . A second unmodeled effect is terrain inhomogeneity, which in this case is caused by nonuniform soil mixing during experiment preparation. In outdoor terrain this effect would likely be due to variation in terrain composition as a function of depth from the surface.

Correction factors K_1 , K_2 , and K_3 were determined via a least-squares fit from bentonite and topsoil experimental data and Eqs. (26,27). The correction factors were determined as $K_1 = 1.5$, $K_2 = 1.0$, and $K_3 = -0.2$. Using these correction factors, the average difference between the predicted and measured DP/W over both data sets was 0.04 and the standard deviation was 0.06. Thus the expected error of DP'_{est} over all experiments was 8.2%. This suggests that the modified drawbar pull prediction equation can yield accurate results in natural terrain.

5.4 Contact-Based Terrain Classification Experimental Results

Contact-based terrain parameter identification experiments were performed on gravel, sand, and the JSC Mars-1 simulant. In these experiments, wheel forward velocity ranged from 0.5 cm/sec to 5 cm/sec, 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 vertical load on the terrain was varied as well, from 30 N 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.

Nearly two hours of vibration data were collected over the parameter range described above, and the algorithm was tuned using the leave-one-out approach. 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 pair-wise classifiers. A single combination of tuned parameters is shared among all pair-wise classifiers.

Next, 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-second segment length was used. After the classifier was trained, it was used to classify the test data sets. Classification results are presented in Table 3. Values shown are counts of 1-second-long vibration segments. These results are plotted in Figure 10.

		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

Table 3. Results from terrain classification experiments.

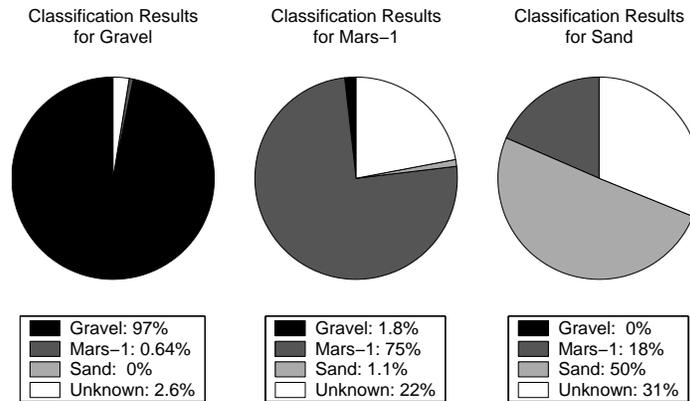


Figure 10. Results from terrain classification experiments

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 half of the data is correctly classified.

Considering the inverse problem—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-second 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.

6. DISCUSSION

The above results have demonstrated that it is possible to estimate terrain physical properties using a variety of methods. The methods are robust and relatively accurate; however it should be noted that all the methods contain error due to various assumptions and real-world implementation issues. The primary sources of error in terrain estimation result from physical phenomena that are difficult to model analytically or semi-empirically, including terrain inhomogeneity and unevenness, and dynamic effects. However it has been shown that these effects do not significantly degrade algorithm performance, suggesting that the above algorithms can be successfully integrated into traversability prediction and servo-level control algorithms.

The mobility enhancement resulting from this integration is algorithm dependent, however in general it can be argued that the introduction of on-line estimates of terrain information would mark an improvement over the current state-of-the-art, which generally rely on simple models of rover-terrain interaction (e.g. static models based on coulomb friction with *a priori* estimates of the friction coefficient). Simple models can be grossly inaccurate predictors of mobility on deformable terrain. They also do not allow estimation of changing terrain conditions.

Recently, researchers have begun developing servo-level control methods for improved rover wheel traction, based on wheel-terrain interaction models and parameter estimates similar to those described in Section 2.0 [27-31]. These works have demonstrated improved mobility performance in sandy, deformable terrain compared to conventional algorithms. Such methods are ideal for implementation in flight systems since they improve intrinsic rover mobility, require no human supervision, and introduce little additional computational or operational burden.

An important area of current work is integrating the methods presented in this chapter with traversability prediction techniques. The current state-of-the-art in rover traversability prediction relies on analysis of geometric parameters (i.e. slope and roughness) of a terrain patch based on elevation data, and cost assignment based on empirical estimates of rover mobility [32]. This approach does not allow a rover to detect “non-geometric” hazards such as the soil trap described in Section 1.0 that immobilized the Mars Exploration Rover Opportunity. A more sophisticated analysis of rover-terrain interaction would allow detection of such hazards, as well as improved prediction accuracy for “geometric” hazards. This would allow a rover to operate closer to its mobility limits and thereby access otherwise inaccessible terrain regions.

7. CONCLUSIONS AND FUTURE WORK

In this chapter three distinct methods for estimating terrain physical properties have been presented. First, an algorithm for on-line estimation of terrain cohesion and internal friction angle was presented. The algorithm relies on a simplified form of classical terramechanics equations, and uses a linear-least squares method to estimate terrain parameters in real time. Second, a method for on-line terrain traversability assessment was presented. The method also relies on a simplified wheel-terrain interaction model, and results in an explicit estimate of terrain

traversability. Finally, a method to classify terrain based on vibrations present in the rover structure during driving was presented. The algorithm relies on data from an accelerometer mounted to the rover structure, and employs linear discriminant analysis to compare on-line vibration signatures to those gathered during an off-line training phase. Experimental results were presented that demonstrate the validity of the three algorithms.

Current work in this area is focused on fusion of the above methods with terrain analysis techniques based on remote sensing. The goal of this work is to exploit the complementary nature of remote sensing methods with the methods described above. Other current research is focused on accurate traversability prediction based on integration of the models presented above with terrain elevation data and kinematic or dynamic rover models.

9. ACKNOWLEDGEMENTS

This work was supported by the Jet Propulsion Laboratory Mars Technology Program. The author would like to acknowledge the support of Drs. Paul Schenker, Richard Volpe and Samad Hayati at JPL. The author would also like to acknowledge the technical contributions of Steven Dubowsky, Christopher Brooks, and Shinwoo Kang to this work.

REFERENCES

- [1] J. Farmer, "Hydrothermal systems on Mars and their potential analogs on Earth," http://exobiology.nasa.gov/ssx/Selected_Research/exopaleo_ts.html, 1996.
- [2] M. Urquhart and V. Gulick, "Lander detection and identification of hydrothermal deposits," abstract presented at *First Landing Site Workshop for MER*, 2003.
- [3] K. Iagnemma and S. Dubowsky, Mobile Robots in Rough Terrain: Estimation, Motion Planning, and Control with application to Planetary Rovers, Springer, Springer Tracts in Advanced Robotics (STAR) Series, Volume 12, June, 2004.
- [4] K. Iagnemma, S. Kang, H. Shibly, and S. Dubowsky, "On-line terrain parameter estimation for planetary rovers," IEEE Transactions on Robotics, Vol. 20, No. 2, 2004, pp. 921-927.
- [5] J. Wong, Terramechanics and Off-Road Vehicles, Elsevier, 1989.
- [6] G. Bekker, Theory of Land Locomotion, University of Michigan, 1956.
- [7] J. Wong, Theory of Ground Vehicles, John Wiley and Sons, 1976.
- [8] A. Le, D. Rye, and H. Durrant-Whyte, "Estimation of track-soil interactions for autonomous tracked vehicles," Proc. of the IEEE International Conf. on Robotics and Automation, 1997, pp. 1388-1393.
- [9] S. Kang, Terrain Parameter Estimation and Traversability Assessment for Mobile Robots, M.S. Thesis, M.I.T., 2003
- [10] J. Wong, and A. Reece, "Prediction of rigid wheels performance based on analysis of soil-wheel stresses, part I," Journal of Terramechanics, Vol. 4, No. 1, 1967, pp. 81-98.
- [11] J. Matijevic, et al., "Characterization of Martian surface deposits by the Mars pathfinder rover, Sojourner," Science, Vol. 278, No. 5, 1997.
- [12] K. Iagnemma, S. Kang, C. Brooks, and S. Dubowsky, "Multi-sensor terrain estimation for planetary rovers," Proc. of Intl. Symposium on Artificial Intelligence, Robotics and Automation in Space, 2003.
- [13] B. Wilcox, "Non-geometric hazard detection for a Mars microrover," Proc. of the AIAA Conference on Intelligent Robotics in Field, Factory, Service, and Space, Vol. 2, 1994.
- [14] C. Olson, L. Matthies, M. Schoppers, and M. Maimone, "Stereo ego-motion improvements for robust rover navigation," Proc. of IEEE International Conference on Robotics and Automation, 2001.
- [15] G. Golub and C. van Loan, Matrix Computations, JHU Press, 1996.
- [16] C. Brooks and K. Iagnemma, "Vibration-based terrain classification for planetary rovers," IEEE Transactions on Robotics, in press.

- [17] I. T. Jolliffe, Principal Component Analysis, Springer-Verlag, New York, 1986.
- [19] S. Balakrishnama, A. Ganapathiraju, "Linear discriminant analysis—a brief tutorial," http://www.isip.msstate.edu/publications/reports/isip_internal/1998/linear_discrim_analysis/lda_theory.pdf, 1998
- [20] R. Fisher, "The use of multiple measurements in taxonomic problems," Annal of Eugenics, Vol. 7, 1936, pp. 179-188.
- [21] K. Mardia, J. Kent, and J. Bibby, Multivariate Analysis, Academic Press, San Diego, 1979.
- [22] P. 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 and Electroacoustics, Vol. AU-15, pp. 70-73.
- [23] G. Golub and C. Van Loan, "The singular value decomposition and unitary matrices," in Matrix Computations, 3rd ed., Johns Hopkins University Press, Baltimore, 1996.
- [24] E. Malinowski, Factor Analysis in Chemistry, 2nd ed., Wiley & Sons, New York, 1991.
- [25] P. Mahalanobis, "On the generalized distance in statistics," Proc. National Institute of Sciences of India, Vol. 2, No. 1, 1936, pp. 49-55.
- [26] C. Allen, R. Morris, K. Jager, D. Golden, D. Lindstrom, M. Lindstrom, and J. Lockwood, "Martian regolith simulant JSC Mars-1," Proc. 29th Lunar and Planetary Science Conf., 1998.
- [27] K. Yoshida, T. Watanabe, N. Mizuno, G. Ishigami, "Slip, traction control, and navigation of a lunar rover," Proc. of Intl. Symp. on Artificial Intelligence, Robotics, and Automation in Space, 2003.
- [28] K. Yoshida, H. Hamano, "Motion dynamics of a rover with slip-based traction model," Proc. of the IEEE International Conference on Robotics and Automation, 2002
- [29] P. Lamon and R. Siegwart, "Wheel torque control in rough terrain—modeling and simulation," Proc. of the IEEE International Conference on Robotics and Automation, 2005
- [30] P. Lamon, T. Thuer, R. Jordi, and R. Siegwart, "Modeling and optimization of wheeled rovers," Proc. of 8th ESA Workshop on Advanced Space Technologies for Robotics, 2004
- [31] K. Iagnemma and S. Dubowsky, "Traction control of wheeled robotic vehicles with application to planetary rovers," International Journal of Robotics Research, Vol. 23, No. 10, 2004, pp. 1029-1040.
- [32] S. Goldberg, M. Maimone, and L. Matthies, "Stereo vision and rover navigation software for planetary exploration," Proc. of the IEEE Aerospace Conference, 2002