

Rule-Based Reasoning and Neural Network Perception for Safe Off-Road Robot Mobility

Edward Tunstel, Ph.D. *

Senior Member of Engineering Staff
Robotic Vehicles Group

Ayanna Howard, Ph.D.

Member of Engineering Staff
Telerobotics Research and Applications Group

Homayoun Seraji, Ph.D.

Senior Research Scientist
Telerobotics Research and Applications Group, Supervisor

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109, USA

* Principal Author Address:

Jet Propulsion Laboratory
4800 Oak Grove Drive, MS 107-102
Pasadena, CA 91109 USA

tunstel@robotics.jpl.nasa.gov

Abstract. *Operational safety and health monitoring are critical matters for autonomous field mobile robots such as planetary rovers operating on challenging terrain. This paper describes relevant rover safety and health issues and presents an approach to maintaining vehicle safety in a mobility and navigation context. The proposed rover safety module is composed of two distinct components: safe attitude (pitch and roll) management and safe traction management. Fuzzy logic approaches to reasoning about safe attitude and traction management are presented, wherein inertial sensing of safety status and vision-based neural network perception of terrain quality are used to infer safe speeds of traversal. Results of initial field tests and laboratory experiments are also described. The approach provides an intrinsic safety cognizance and a capacity for reactive mitigation of robot mobility and navigation risks.*

Keywords: *safe navigation, planetary rovers, neural networks, fuzzy logic, off-road mobility*

1. Introduction

Operational safety and health monitoring for autonomous mobile robots are often treated as secondary research concerns relative to the primary problems of robot navigation and control. The issue is a primary concern, however, in field mobile robot research and applications. Field mobile robots that perform remote missions are expected to take reactive measures to maintain vehicle safety and nominal operation. This is a critical matter for field robots such as planetary rovers operating on challenging terrain. Mobile robots designed to explore remote planetary surfaces are expected to encounter terrain hazards such as extreme slopes, sand or dust-covered pits, ditches, cliffs and otherwise impassable surfaces. In order for such planetary rovers to consistently avoid these mobility hazards and when necessary, negotiate challenging terrain, they must be capable of continuously assessing the safety of their traversal. This is a fundamental requirement for the effective use of rovers in carrying out scientific exploration objectives that require mobility across harsh terrain on planets such as Mars.

In addition to the mobility requirement, significant autonomy is required to enable long-duration missions (several months) involving long-range navigation, and acquisition and analysis of scientific data or surface material samples. All necessary autonomy capabilities must operate within significant constraints on power, computation, weight, and communications bandwidth. To further increase the challenge, many popular and fast state-of-the-art processors that enable advanced capabilities in laboratory research robots are infeasible for planetary rover applications. This is due to the fact that space flight missions require the use of proven, radiation-hardened, or otherwise space flight-qualified electronics that will survive and operate in the harsh temperature and radiation extremes of outer space. The meager

availability of fast and/or powerful space-qualified processors for on-board computation intensifies the need for efficient algorithms, which comply with the practical limitations and constraints while enabling on-board autonomy.

One of the autonomy design considerations for space mobile robot applications is the respective distribution of computational resources among mobility, navigation, and science autonomy. While autonomous mobility is critical, the acquisition of scientific knowledge is the primary objective of the mission. Therefore, mobility-critical autonomy must be designed to consume relatively minimal computational resources. As such, we have concentrated on developing intuitively simple approaches to safe mobility and navigation, which are supported by techniques that enable intelligent control with modest computational overhead.

In this paper, we describe essential components of an overall navigation system that are responsible for maintaining vehicle safety and nominal operation in support of strategic navigation. Fuzzy logic and neural network techniques are used for rule-based reasoning about rover safety when traversing challenging terrain. Visual perception algorithms complement the reasoning approach to realize a practical safety module. The approach is motivated by a desire to emulate human judgement and reasoning as derived from off-road driving heuristics (DeLong 1996). In the following sections, we describe relevant rover safety and health issues, and the fuzzy logic approach to reasoning about them in order to provide an intrinsic safety cognizance and a capacity for reactive mitigation of navigation risks. Field test and experimental results are also described.

2. Safety Reasoning for Mobility

Most existing autonomous rover mobility systems focus on *strategic* navigation goals of the rover and disregard the intrinsic safety and health monitoring of the vehicle itself. In some existing systems (Morrison & Nguyen 1996; Washington et al 1999), the common practice is to consider basic monitoring of individual hardware components for proper operation, without incorporating explicit autonomous reaction or counter-action by the rover. Efficient management of on-board resources, such as power and science data storage capacity, and regulation of energy and internal temperature are common concerns for maintaining vehicle health (Washington et al 1999; Huntsberger & Rose 1998). In addition to vehicle health, operational safety is of primary importance. Navigation systems have also been developed which account for some measure of risk mitigation with respect to accidental damage (as due to tip-over) and/or vehicle entrapment (Rosenblatt 1997; Kelly & Stentz 1997). However, few field mobile robot systems have been reported in the literature that feature efficient implementation of active countermeasures for both vehicle health *and* safety in a comprehensive fashion.

2.1 Rover Health and Safety Indicators

The ability of a system to provide substantial safety countermeasures depends upon its capacity for assessing its status with respect to the operating environment. Various observable states, events, and terrain features can be considered for on-line assessment of a rover's operational status. Table I lists a number of possible health and safety indicators associated with rover on-board subsystems, which convey some aspects of rover operational well-being as it relates to safe terrain traversal. Ultimately, a comprehensive autonomous vehicle health and safety system is desired to increase rover survivability. Perhaps consideration of all items in Table I would make this possible, but such complete observability is rare in practice. To this end, we have concentrated on providing some of the basic elements necessary to approach the ultimate goal. To begin with, we will consider the chassis attitude and the terrain type as the prominent factors that affect the rover safety and health. The addition of automated mechanisms for self-regulation of available power and internal operating condition is also planned, guided by examples described in (Huntsberger & Rose 1998; Arkin 1992).

Table I. Rover Health and Safety Indicators.

Health	Safety
Available power	Chassis attitude
Component failure	Terrain type/condition
Component temperature	Wheel slip and sinkage
Drive motor stall	Mechanics compliance

The attitude (pitch and roll) of the vehicle chassis with respect to an inertial reference frame can be monitored in order to avoid instabilities associated with ascent/descent of slopes, traversal of rocky terrain, and turning within the vehicle's curvature constraints. In addition to surface irregularities, the type and condition of the terrain surface provide clues for safety assessment. Human automobile drivers are able to perceive certain road conditions (e.g. oil slicks, potholes, and ice patches) as measures of safety, and react to them in order to reduce the risk of potential accidents. In a similar manner, rover potential safety can be inferred and reacted to based on knowledge of the terrain type or surface condition. Wheel-soil interactions are important mobility considerations in natural terrain. Excessive wheel slippage reduces the effective traction that a rover can achieve and, therefore, its ability to make significant forward progress (not to mention the dramatic effect it can have on the accumulation of errors in estimated position and orientation over distance and time). On soft soils, such as fine-grained sand, excessive wheel slippage can often lead to wheel sinkage and eventual entrapment of the vehicle. Unfortunately, wheel slippage and sinkage are often difficult to measure and

estimate in a simple manner. Some progress has been made, however, in developing statistical estimation approaches for planetary rovers (Wilcox 1994).

At this stage of development, the safety module employs concise fuzzy logic systems that provide autonomous reasoning to facilitate maintenance of stable vehicle attitude and wheel traction on rough terrain based on inertial sensing inputs and neural network perception of terrain texture. The system employs off-road driving heuristics to facilitate avoidance of hazardous vehicle configurations and excessive wheel slippage. In each case, our system is designed to produce safe speed recommendations associated with the current perception of the safety status of the rover. The proposed rover safety and survival system is composed of two distinct subsystems: attitude management and traction management. These subsystems are discussed in the ensuing sections.

3. Safe Attitude Management

For indoor mobile robots, mobility and navigation problems can often be addressed in two dimensions (x and y) since the typical operating environments consist of flat and smooth floors. In sharp contrast, mobility and navigation problems for outdoor rough terrain vehicles are characterized by significantly higher levels of difficulty. This is due to the fact that complex motions in the third dimension (z) occur quite frequently as the vehicle traverses undulated terrain, encountering multi-directional impulsive and resistive forces throughout. The problem is more pronounced for vehicles with more or less rigid suspensions than it is for vehicles with articulated chassis. In any case, sufficient measures must be taken to maintain upright stability of the vehicle in both static and dynamic configurations.

For monitoring chassis attitude, our test vehicle is outfitted with a two-axis inclinometer/tilt sensor, which measures pitch and roll angles relative to a Cartesian reference frame that is aligned with the rover chassis coordinate frame when the vehicle rests on a level surface. With such a sensor, perhaps the simplest approach is to stop rover motion when either axis senses tilt beyond a critical threshold. In a few instances this “wait and see” binary approach may be sufficient. More often than not, however, dynamic effects such as momentum will quickly defeat the simplest approach and cause the rover to reach marginal stability (a point at which the vehicle begins to tip over), or worse yet, to actually tip-over. Even though planetary rovers are typically driven at low speeds (e.g. maximum average speed of ~ 0.3 m/s), more sophistication is required beyond binary threshold reactions. Instead of allowing the vehicle to wait for the roll or pitch to build up to a dangerous threshold before reacting, we have elected to formulate a safety strategy in which the recommended safe speed for the rover is *gradually* modulated in reaction to changes in attitude. When the rover travels on a relatively level surface, a maximum safe speed is recommended. As pitch and/or roll approaches extremes near marginal stability, gradual reductions in safe speed are recommended (culminating at halted motion). At attitudes between these extremes, recommended safe speeds are computed by interpolation via fuzzy sets and logical inference.

By considering various off-road driving heuristics for traversing rock-fields, ravines, and hills (up-, down-, and side-hill) (DeLong 1996), a set of fuzzy logic rules is formulated to maintain stable rover attitudes for safe navigation. The allowable ranges of pitch and roll are partitioned (based on subjective assessment of the problem and the vehicle characteristics) by fuzzy sets to express the approximate nature of the measurements. Pitch is represented by five fuzzy sets with linguistic labels {NEG-HIGH, NEG-LOW, ZERO, POS-LOW, POS-HIGH}, while roll is partitioned using three fuzzy sets with linguistic labels {NEG, ZERO, POS}. Here, positive and negative are abbreviated by “pos” and “neg”, respectively. Bounds on the allowable ranges for attitude measurements are chosen in accordance with the rover stability constraints and the level of acceptable risk. That is, the actual critical pitch/roll angles that correspond to marginal stability for the rover are scaled down using scalar safety factors (positive constants less than one). The scaled critical angles are used to constrain the universe of discourse such that the attitude behavior responds early to pitch/roll extremes before marginal stability is reached.

Pitch and roll are used to infer rover speed, which is represented by fuzzy sets with linguistic labels {SLOW, MODERATE, FAST}. The fuzzy logic rule set and membership functions for the stable attitude control component are shown in Table II and Figure 1, where the rover speed v is the output and the rover pitch ψ and roll ρ are the inputs. As is common in fuzzy logic control systems, the membership functions used to express uncertainty in the variables of each system component take on triangular and/or trapezoidal shapes (Driankov et al 1993).

		pitch				
roll		NH	NL	ZE	PL	PH
	NEG	SLOW	MOD	MOD	MOD	SLOW
	ZER	MOD	MOD	FAST	MOD	MOD
	POS	SLOW	MOD	MOD	MOD	SLOW

Table II. Rule-Base for Safe Attitude.

Fifteen fuzzy logic rules are employed to map the range of stable attitudes (ψ , ρ) to safe driving speed (v) recommendations. Given specific values for ψ and ρ , each fuzzy logic rule computes a truncated fuzzy set over the range of v partitioned by {SLOW, MODERATE, FAST} (see Figure 1). Let A be the union of all such truncated fuzzy sets. Then the safe speed recommended for maintaining stable attitude is computed using the following relationship according to the Center-of-Gravity defuzzification method (Driankov et al 1993).

$$v = \frac{\sum s \cdot \mu_A(s)}{\sum \mu_A(s)}, \forall s \in [0, v_{\max}] \quad (1)$$

In addition to the fuzzy logic rules in Table II, a crisp rule is applied to set rover speed to zero in the extreme cases when marginal stability is reached and the safest reaction is to stop its motion. However, in contrast to the binary threshold scheme mentioned earlier, as marginal stability is approached the rover speed is smoothly decreased to near zero due to the interpolation provided by the fuzzy logic rules.




Figure 1. Membership functions for safe attitude.

4. Safe Traction Management

In the absence of some measure of control, wheeled vehicles are prone to loss of traction under certain terrain conditions. On dry paved roads, traction performance is maximal for most wheeled vehicles due to the high coefficient of friction/adhesion between the road and the tread. On off-road terrain, however, a variety of surface types are encountered on which rover wheels are susceptible to slippage. Examples include sand, gravel, densely packed soil, ice, mud, and so on. As mentioned above, loss of traction due to excessive wheel slippage can lead to wheel sinkage and ultimately vehicle entrapment. Frequent loss of traction during a traverse from one place to another will also detract significantly from the ability to maintain good position estimates. To improve rover performance, a mechanism for regulating or mitigating wheel slippage is highly desirable.

4.1 Sensing and Perception Issues

The problem of traction control is not new. It is a common problem in automobile and general transportation vehicle design with a variety of effective solutions. In many cases, solutions are derived from analyses based on the following equation for wheel slip ratio, λ , which is defined non-dimensionally as a percentage of vehicle forward speed, v (Gillespie 1992):

$$\lambda = \left(1 - \frac{v}{r_w \omega_w} \right) \times 100 \quad (2)$$

Here, r_w is the wheel radius and ω_w is the wheel rotational speed. Equation (2) expresses the normalized difference between vehicle and wheel speeds. When this difference is non-zero, wheel slip occurs. The objective of traction control is to regulate λ to maximize traction. This is a relatively straightforward regulation task if v and ω_w are both observable. The wheel rotational speed ω_w is typically available from shaft encoders or tachometers. However, it is often difficult to measure the actual over-the-ground speed v for off-road wheeled vehicles. Nonlinearities and time-varying uncertainties due to wheel-ground interactions further complicate the problem. Despite this, effective solutions have been found for automotive applications of anti-lock (deceleration) and anti-slip (acceleration) control (Mauer 1995; Bauer & Tomizuka 1995; Palm & Storjohann 1997; Van der Burg & Blazevic 1997). In these cases, measurement of v is facilitated by the even surface on which the vehicle travels, or by special sensing arrangements. In (Mauer 1995), an accelerometer is used to measure vehicle speed and the slip ratio is estimated based on deceleration of the four wheels. In (Bauer & Tomizuka 1995), the measurement of vehicle speed is facilitated by the use of magnetic markers alongside the road in an intelligent highway automation system. In this case, the vehicle speed is measured according to travel time between markers. For application to an electrically driven locomotive, the solution in (Palm & Storjohann 1997) makes use of a model of the friction-slip relationship, which is fixed for the wheel-rail interaction. On outdoor terrain, the friction-slip relationship varies with surface type. In large part, the available solutions are not directly transferable to off-road vehicle applications in which the terrain is uneven as opposed to being relatively flat, as is the case for automobiles and locomotives.

The use of an accelerometer to measure off-road vehicle speed is problematic since the gravity effects of traversing longitudinal and lateral slopes will interfere with the measurement. For an accelerometer used to measure horizontal acceleration, any off-horizontal vehicle tilt will be sensed as a change in acceleration; as a result, the integrated velocity will be in error. This is realized in (Van der Burg & Blazevic 1997) where an alternative traction control concept for rovers is considered. In that case, a non-driven “free wheel” is proposed for measuring actual vehicle speed. Another promising solution is proposed for rovers with an articulated chassis, which enables active control of the vehicle center of gravity. For those vehicles, the use of accelerometers in conjunction with rate gyroscopes is suggested (Sreenivasan & Wilcox 1994).

In our work, we have elected to take a simple linguistic approach that does not rely on accurate sensing of vehicle speed. Instead, visual perception of terrain texture is used to infer an appropriate speed of traversal. Results from traction tests performed on the actual rover are used to determine appropriate speeds for a variety of potential surface types. In particular, the rover is tested on different terrain surfaces (e.g., sand, gravel, densely packed soil, etc.) to determine the maximum speeds achieved before the onset of wheel slippage. These tractive speeds are designated by fuzzy linguistic labels {CAUTIOUS, SUBDUED, NOMINAL} to be discussed later. Given this information,

commanded vehicle speed can be modulated during traversal based on visual classification of the terrain surface type in front of the rover. This is analogous to the perception-action process that takes place when a human driver notices an icy road surface ahead and decelerates to maintain traction. For a rover, such speed modulation allows management of traction by mitigating the risk of wheel slippage. The approach is similar in spirit to other fuzzy logic and dynamic feedback control methods (Colyer & Economou 1999; Sarkar & Yun 1998) proposed for appropriately distributing wheel motor torques to improve traction, albeit, after the onset of wheel slippage.

Given the results of actual traction tests, the formulation of fuzzy logic rules to achieve speed modulation is relatively straightforward. The success of the traction management approach depends more heavily on the ability to perceive and classify the various terrain surface types. The problem of off-road surface type identification is formidable for systems equipped with only proximity sensors, range-finders, and/or tactile probes. However, visual image-based classification has been found to be particularly promising (Marra et al 1988). We will now describe an artificial neural network solution to this problem that provides qualitative information about the expected surface traction of terrain immediately in front of the rover. This information is used to infer tractive rover speeds via fuzzy inference.

4.2 Vision-Based Neural Network Terrain Classification

Distinct terrain surfaces reflect different textures in visual imagery. The ability to associate image textures to terrain surface properties such as traction, hardness, or bearing strength has clear benefits for safe autonomous navigation. To provide this capability, we make use of an on-board camera pointed such that its field-of-view (FOV) covers the ground area in front of the rover. That is, the camera is mounted on the robot chassis such that the extracted image will provide a downward looking view of the surface as illustrated in Figure 2. Based on typical surfaces that the rover may encounter, three different texture prototypes are selected: sand, gravel, and compacted-soil (Figure 3).




Figure 2. Camera mounted on rover.



Figure 3. Terrain surface texture images (left to right): sand, gravel, compacted soil.

The automated method of classifying terrain surface type is based on a texture analysis approach using a neural network. The method proceeds as follows. Assuming the section of the image just ahead of the front wheels is free of obstacles, a set of 40x40 pixel image blocks (samples) is randomly selected from a camera image of size 320x280 pixels. To reduce the large data dimensionality inherent in typical vision-based applications, a filtering step is performed using a standard technique called Principal Component Analysis (PCA) (Diamantaras & Kung 1996). PCA is a linear optimal method for reducing data dimensionality by identifying the axis about which the desired feature set varies the most. This orthogonal sub-space projection of the image subset permits effective extraction of features embedded in the surface image data set in real time. This technique reduces the dimensionality of the image set while preserving as much of the signal as possible. PCA computes a set of orthonormal eigenvectors (filters) of a data set that captures the greatest correlation between features. The filters associated with a given feature set are derived from the distribution of potential dynamic features embedded in the images. A total of 30 eigenvectors are used to reduce the 40x40 image sample (1600 pixel values) to a pattern set of 30 values by projecting the image data onto the most significant filters. This reduced data set is then used to train a neural network (Figure 4) to associate texture with several surface types. The network output provides the qualitative information needed to make any necessary adjustments to wheel speed in order to maintain traction on the classified surface. After training the network on typical image data representing different surface prototypes, we utilize it to classify the surface types during run-time.




Figure 4. Neural network for surface classification.

4.3 Intuitive Rule-Based Control Reasoning

The neural network classifier is trained to provide texture prototype outputs in the unit interval $[0, 1]$, with 0 corresponding to surfaces of very low traction (e.g., ice) and 1 corresponding to surfaces of very high traction (e.g., dry cement). This is a design decision motivated by a desire to establish some intuitive correlation to actual wheel-terrain coefficients of friction. In this way, we can make a qualitative association between the output of the neural network and expected terrain traction in front of the rover. We will refer to the texture prototype output as the *traction coefficient*, denoted by C_t .

Wheel-terrain friction coefficients for a variety of tread and surface types are widely published in the literature on vehicle mechanics. However, published friction coefficients for identical tread and surface types vary from source to source. This is due to the fact that measured values depend heavily on the variety of tests and test conditions from which they are generated. Nevertheless, common ranges of friction coefficients for given tread and surface types are widely agreed upon. The following are typical estimates of the friction coefficients for rubber tires on various surfaces (Gillespie 1992; Bekker 1969): icy road/snow (0.1), sand (0.3), slippery/wet road (0.4), hard unpaved road (0.65), grass (0.7), dry paved road (0.8-1.0).

Given the uncertainty in associating exact friction coefficients with certain terrain surface types, and the loose correlation provided by the traction coefficient using a neural network, we elect to reason about traction using fuzzy logic. The range of traction coefficients, $[0,1]$, obtained from the neural network classifier is partitioned using three fuzzy sets with linguistic labels {LOW, MEDIUM, HIGH} as shown in Figure 5. Based on these definitions, the following simple fuzzy logic rules are applied to manage rover traction on varied terrain:

IF C_t is LOW, THEN v is CAUTIOUS.

IF C_t is MEDIUM, THEN v is SUBDUE.

IF C_t is HIGH, THEN v is NOMINAL.




Figure 5. Fuzzy sets for traction coefficient.

It is noted that for traction management, the membership functions for the rover speed v are based on results of prior traction tests described in Section 4.1 and can differ from those described in Section 3. However, the shapes of the

membership functions for the rover speed are similar to those shown in Figure 1 and the range of tractive speeds is bound by the interval $[0, v_{max}]$. Given specific values for C_r , each of these fuzzy logic rules computes a truncated fuzzy set over the range of v partitioned by $\{\text{CAUTIOUS}, \text{SUBDUED}, \text{NOMINAL}\}$. If we let T be the union of all such truncated fuzzy sets, then the safe speed recommended for managing traction is computed as follows.

$$v = \frac{\sum s \cdot \mu_T(s)}{\sum \mu_T(s)}, \forall s \in [0, v_{max}] \quad (3)$$

Note that the neural network can be trained to map its inputs directly to the actual range of tractive speeds (rather than the range of C_r). However, in this neuro-fuzzy approach, fuzzy inference serves to account for uncertainties in both the surface classification and the subsequent specification of tractive speed.

The *smaller* of the two rover speeds inferred by the safe attitude and traction management components (Equations 1 and 3, respectively) is issued at each control cycle as the safe speed recommendation v_{safe} . In the overall navigation system, the safety reasoning module focuses on vehicle survivability and health, while the strategic navigation module focuses on mission and goal-directed motion from place to place (Seraji et al 2001). The interface between the safety module and the strategic navigation module is depicted in Figure 6.

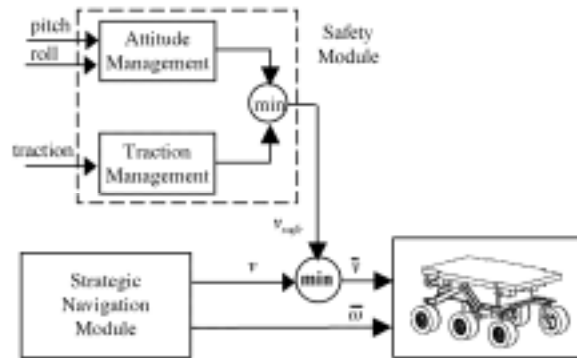


Figure 6. Safety and strategic navigation interface.

As shown in the diagram, the safe speed recommended by the safety module is compared to the strategic speed recommendation. The smaller of the two is taken as the safest speed and is issued as the commanded set-point \bar{v} for translational motion control. Determination of v_{safe} is independent of the behavior fusion process used to compute strategic navigation speeds (Seraji et al 2001). This allows recommended safe speeds to override strategic speeds, if necessary, to ensure vehicle safety. This is also the approach taken in (Murphy & Dawkins 1996) based on the

assertion that distributing speed control across all behaviors makes it difficult to ensure that the interactions will yield a safe speed. Note that the commanded rotational velocity $\bar{\omega}$ of the rover is unaffected by v_{safe} .

5. Results

The rule-based reasoning and neural network perception techniques embedded in the safety module were validated on a mobile robot in off-road terrain. In this section, we discuss the neural network classification capability and describe initial experiments performed to evaluate the effect of the safe attitude and traction components. Physical experiments are performed at an off-road field site, and laboratory computer simulations are also utilized for evaluation.

Two illustrative examples are presented to demonstrate the safe attitude and traction management components as verified in the field and laboratory tests. The first considers reactions to rover pitch and roll during traversal. The second is concerned with mitigation of wheel slippage. As a test vehicle, we use the Pioneer All-Terrain (AT) mobile robot platform, a commercially available robot designed for rough terrain mobility. The rover hardware is enhanced with additional on-board computing (a 333 MHz Pentium II processor), a vision system for real-time terrain assessment, and a tilt sensor (see Figures 7a and 7b). To provide pitch and roll sensing, the Crossbow Technology, Inc. model CXTA02 inclinometer is used, which has a $\pm 75^\circ$ range and 0.05° resolution. The ground-facing camera on the front of the rover is mounted 0.3m above the ground, tilted downward 45° with a 45° FOV. This camera enables surface traction classification (cameras for terrain-based navigation are mounted on the raised platform). We shall now present field test and laboratory experimental results for the safe attitude and traction management separately.




Figure 7a, 7b. Pioneer-AT rover with enhancements.

5.1 Safe Attitude Management Testing

For the safe attitude test, an obstacle-free swath of undulated terrain is chosen. The rover is commanded to traverse the swath with and without the safe attitude component activated. Without active safe attitude management, the rover traverses the terrain at a nominally fast speed recommended by the strategic navigation system based on the fact that no significant obstacles are present. With active attitude management, the rover traverses the terrain at various reduced speeds in response to changes in its pitch and roll according to the fuzzy logic rules in Table II. This reactivity reduces the risk of approaching marginal tilt stability, which leads to tip-over. It also enhances the ability of rigid-suspension vehicles (such as the Pioneer-AT) to maintain wheel contact with the ground. A comparative effect of the safe attitude component is shown in Figure 8. The left picture corresponds to the test without active attitude management; it shows a case where the rover's rear-right wheel loses contact with the ground. The right picture shows the rover at the same approximate location with all wheels in contact with the ground while actively modulating its speed to maintain safe attitude.



Figure 8. Effect of safe attitude control.

To further illustrate the effect of safe attitude management, we exercise the component in a laboratory experiment where the rover traverses a swath of terrain for 10 meters. Synthetic attitude measurements are generated by sinusoidal functions of random amplitude to emulate changes in pitch and roll experienced on a hypothetical undulated and rough terrain. The amplitudes are uniformly distributed random numbers bounded by the maximum stable pitch and roll of the rover. It is assumed that the strategic navigation module recommends a constant normalized speed of 75% (of maximum allowable speed) throughout the traverse. The results of this experiment are shown in Figure 9 in plots of pitch, roll and v_{safe} (normalized) versus distance. The strategic speed is shown in the speed-distance plot as a dashed line.

Observe that v_{safe} is modulated low in response to near-extreme attitudes. This is most apparent when both pitch and roll are simultaneously large in magnitude. Also observe that v_{safe} is consistently lower than the strategic speed, thus exhibiting the caution of the safe attitude component in reaction to cognizance of vehicle safety. The overall

characteristic behavior of the safe attitude component is illustrated by Figure 10, which shows the control decision surface corresponding to the rules of Table I and membership functions of Figure 1 (normalized axes are used).




Figure 9. Speed modulation for attitude management.

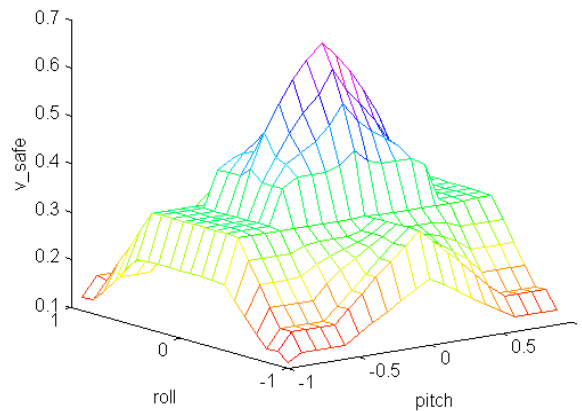


Figure 10. Safe attitude control decision surface.

5.2 Traction Classification and Management Testing

For traction classification, the neural network, described in Section 4.2, uses thirty input nodes corresponding to the set of orthonormal eigenvectors (filters) of the image data set computed using PCA. One output node is used to represent the surface type classification. Twenty processing elements are used in the hidden layer of the neural network. The network is trained using backpropagation to find a set of weights that produce the desired surface type classification for a given training data input. A set of 400 image samples from various images of terrain surface textures was used to train the neural network. After learning, a distinct set of 400 new image samples was presented to

the neural network to test its classification performance. A success rate of 96% was achieved for classifying representative terrain surfaces. The selected network structure proved successful for classifying terrain textures typically encountered by the rover in field experiments run to validate the approach. Figure 11 shows several images of real terrain data properly classified by the trained neural network; these images were not included in the data set used to train the network. The most successful terrain texture classifications can be made in the absence of significant shadows cast by nearby terrain features such as mountains or hills. In the presence of such terrain features, and at different times of day, natural illumination from the sun causes significant shadowing on terrain in front of the rover, which results in misclassification of surface textures (Howard & Seraji 2001). Illumination conditions are known to present difficulties for computer vision-based methods of terrain classification. However, it is expected that the robustness of the neural network classifier can be improved following training using images taken under varied illumination conditions.




Figure 11. Terrain surface images classified by the neural network.

In order to test safe traction management, a benign portion of terrain comprising two distinct surface types (hard compact soil and gravel) is chosen on which the rover will be susceptible to wheel slippage when traversing the surface transition at nominally fast speeds. The scenario is depicted in Figure 12 where the rover is about to transition from a hard compact soil to gravel surface. A comparative field test and laboratory experiment is performed wherein the rover is commanded to traverse the transition with and without the safe traction management component activated. Again, without active traction management, the rover traverses the terrain at a nominally fast speed. With active traction management, the rover reduces its speed upon encountering a surface of lower perceived traction as classified by the vision-based neural network classifier. Safe speed is modulated according to the fuzzy logic rules presented in Section 4.3. This reactivity mitigates the risk of excessive wheel slippage during transitions between, and traversal on, surfaces of different traction characteristics.



Figure 12. Rover approaching surface type transition.

To further illustrate the effect of the safe traction management, we exercise the component in a laboratory experiment where the rover traverses a 12 meter swath of terrain consisting of different surface types for which the traction coefficient C_t is 0.5 for 5m, 0.2 for 3m, and 0.9 for 4m. We assume, for the sake of discussion, that these values correspond to sand, gravel, and concrete, and that the surface texture camera has a ground surface view horizon out to 0.3m in front of the rover wheels. In this experiment, the strategic navigation module recommends a constant normalized speed of 80% throughout the 12m traverse. The results are shown in Figure 13 where the traction coefficient and recommended rover speeds are plotted versus distance. The images of the three terrain surface types corresponding to distance are inset in the figure as well.




Figure 13. Speed modulation for traction management.

As expected, changes in perceived traction result in reactive management of the safe speed recommended by the safe traction component to avoid the risk of excessive wheel slippage. Note that our laboratory experiment accounts for a reaction delay between classification of the surface type and the actual change in set-points for v_{safe} . As in the previous example, v_{safe} is consistently lower than the strategic speed, thus exhibiting the caution of the safety module in reaction to cognizance of changing “road” condition.

6. Conclusions

Safe and autonomous long-range navigation of a rover on hazardous terrain offers significant technical challenges. For rover operation over extended time and distance, some capacity for built-in safety and health cognizance is essential. This paper describes how a nominal level of safety assurance can be achieved with intuitive fuzzy logic rules for intelligent control. Details of a rover safety module designed to support strategic navigation in challenging terrain are presented. The module employs fuzzy logic and a neural network augmented by visual perception algorithms to provide a safety module that is realizable in practical rover computing hardware. The safe attitude and traction management components of the safety module combine to provide active countermeasures to potential vehicle tip-over and excessive wheel slippage. This capability leads to survivable rover systems that are of practical use for performing long-duration missions involving long-range traversal over challenging and high-risk terrain.

7. Acknowledgments

The research described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The authors acknowledge experimental assistance of Justin Kao of Caltech and Christopher Sheesley of Univ. of Arkansas, Little Rock.

REFERENCES

- ARKIN, R.C. (1992) Homeostatic control for a mobile robot: dynamic replanning in hazardous environments, *Journal of Robotic Systems*, 9(2), 197-214.
- BAUER, M. and M. TOMIZUKA (1995) Fuzzy logic traction controllers and their effect on longitudinal vehicle platoon systems, Technical Report, UCB-ITS-PRR-95-14, University of California, Berkeley, CA USA.
- BEKKER, M.G. (1969) Introduction to terrain-vehicle systems, University of Michigan Press, Ann Arbor, MI, USA.

COLYER, R.E. and J.T. ECONOMOU (1999) Soft modelling and fuzzy logic control of wheeled skid-steer electric vehicles with steering prioritisation, *Journal of Approximate Reasoning*, 22, 31-52.

DELONG, B. (1996) 4-Wheel freedom: the art of off-road driving, Paladin Press, Boulder, CO USA.

DIAMANTARAS, K. and S.Y. KUNG (1996) Principal component neural networks: theory and applications, Wiley & Sons, New York.

DRIANKOV, D., H. HELLENDORF and M. REINFRANK (1993) An introduction to fuzzy control, Springer-Verlag, Berlin.

GILLESPIE, T.D. (1992) Fundamentals of vehicle dynamics, Society of Automotive Engineers, Inc., Warrendale, PA, USA

HOWARD, A. and H. SERAJI (2001) Vision-based terrain characterization and traversability assessment, *Journal of Robotic Systems*, 18(10), 577-587.

HUNTSBERGER, T.L. and J. ROSE (1998) BISMARC: a biologically inspired system for map-based autonomous rover control, *IEEE Transactions on Neural Networks*, 11(7/8), 1497-1510.

KELLY, A. and A. STENTZ (1997) An approach to rough terrain autonomous mobility, *Proceedings of the International Conference on Mobile Planetary Robots*.

MARRA, M., R.T. DUNLAY and D. MATHIS (1988) Terrain classification using texture for the ALV, *Proceedings of Mobile Robots III*, SPIE 1007, 64-70.

MAUER, G.F. (1995) A fuzzy logic controller for an ABS braking system, *IEEE Transactions on Fuzzy Systems*, 3(4), 381-388.

MORRISON, J.C. and T.T. NGUYEN (1996) On-board software for the Mars Pathfinder microrover, *Proceedings of the 2nd IAA International Conference on Low-Cost Planetary Missions*, IAA-L-0504:1-7.

MURPHY, R. R. and D.K. DAWKINS (1996) Behavioral speed control based on tactical information, *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*, 1715-1721.

PALM, R. and K. STORJOHANN (1997) Slip control and torque optimization using fuzzy logic, in M. Jamshidi, et al. (Eds.), *Applications of Fuzzy Logic: Towards high machine intelligent quotient systems*, 8, Prentice-Hall PTR, NJ, USA.

ROSENBLATT, J.K. (1997) DAMN: A distributed architecture for mobile navigation, *Journal of Experimental and Theoretical AI*, 9(2/3), 339-360.

SARKAR, N. and X. YUN (1998) Traction control of wheeled vehicles using dynamic feedback approach, *Proceedings of the IEEE International Conference on Robotics and Automation*, 413-418.

SERAJI, H., A. HOWARD and E. TUNSTEL (2001) Safe navigation on hazardous terrain, *Proceedings of the IEEE International Conference on Robotics and Automation*, Seoul Korea, 3084-3091.

SREENIVASAN, S.V. and B.H. WILCOX (1994) Stability and traction control of an actively actuated micro-rover, *Journal of Robotic Systems*, 11(6), 487-502.

VAN DER BURG, J. and P. BLAZEVIC (1997) Anti-lock braking and traction control concept for all-terrain robotic vehicles, *Proceedings of the IEEE International Conference on Robotics and Automation*, 1400-1405.

WASHINGTON, R. et al. (1999) Autonomous rovers for Mars exploration, *Proceedings of the 1999 IEEE Aerospace Conference*.

WILCOX, B.H. (1994) Non-geometric hazard detection for a Mars microrover, *Proceedings of the NASA/AIAA Conference on Intelligent Robotics in Field, Factory, Service, and Space*, 675-684.