

Terrain Classification Using Weakly-Structured Vehicle/Terrain Interaction

Amy C. Larson[†], Guleser K. Demir[‡], and Richard Voyles[†]

[†]*Department of Computer Science, University of Minnesota, Minneapolis, MN USA.*

[‡]*Department of Electrical and Electronics Engineering at Dokuz Eylul University, Izmir, Turkey*

Abstract. We present a new terrain classification technique both for effective, autonomous locomotion over rough, unknown terrains and for the qualitative analysis of terrains for exploration and mapping. Our approach requires a single camera with little processing of visual information. Specifically, we derived a *gait bounce* measure from visual servoing errors that results from vehicle-terrain interactions during normal locomotion. Characteristics of the terrain, such as roughness and compliance, manifest themselves in the spatial patterns of this signal and can be extracted using pattern classification techniques. This vision-based approach is particularly beneficial for resource-constrained robots with limited sensor capability. In this paper, we present the gait bounce derivation. We demonstrate the viability of terrain classification for legged vehicles using gait bounce with a rigorous study of more than 700 trials, obtaining an 83% accuracy on a set of laboratory terrains. We describe how terrain classification may be used for gait adaptation, particularly in relation to an efficiency metric. We also demonstrate that our technique may be generally applicable to other locomotion mechanisms such as wheels and treads.

Keywords: terrain classification, mobile robots, legged locomotion

1. Introduction

Creating adept robots for the real world requires adaptation to the environment. Adaptability is especially key when operating in environments where *a priori* information is approximate or unavailable, for example rescue missions inside felled buildings or planetary exploration. When operating in such environments, efficiency can generally be improved by first characterizing the current operational environment, then using this information to optimize some performance metric. Assessment of terrain conditions can be particularly useful because the terrain impacts both path planning and motion control. Operating outdoors highlights this point, where encounters with sand, rocks, grass, or pavement impact efficiency and speed of forward progress. These surfaces do not necessarily obstruct the path, rather require adjustments to motion control for more effective progress.

There are a variety of approaches for terrain assessment, which we review in Section 2, but prior efforts require special equipment or are

computationally expensive. We have developed a terrain classification technique using a measure we call *gait bounce*. This measure can be easily extracted from visual servoing errors obtained from a single camera tracking features in the image plane or from dedicated body-orientation sensors. Intuitively, it is analogous to estimating the roughness of the road, for example, based on the severity of the rocking motion of a vehicle. Our approach provides a rich source of terrain information, yet requires only a single camera, relatively little processing of visual data, and is applicable to legged and to nonlegged platforms. It is particularly beneficial for resource-constrained robots due to its reliance on multi-function sensors.

In Section 3, we present the derivation of *gait bounce* and *gait roll*, which are estimates, based on visual information, of the pitch and roll of the vehicle as it interacts with the terrain. The spatial pattern of the gait bounce encodes terrain conditions, such as roughness and compliance, which we can identify with pattern classification techniques using spatial discriminants, as shown in Section 4.

Our testing platform is the CRAWLER Scout (**C**ylindrical **R**obot for **A**utonomous **W**alking and **L**ifting During **E**mergency **R**esponse, also known as TerminatorBot), as seen in Figure 1. It is a small-scale robot designed for applications of search-and-rescue, planetary exploration, and surveillance. Its versatility in arm motion allows for a variety of gait classes, as well as fine-manipulation. As an active sensor, it can effect a variety of limb-terrain interactions to investigate and define different characteristics of the terrain. As a platform for use in search-and-rescue or planetary exploration, it is sufficiently mobile to traverse the rough terrains of these applications. As a study in adaptation, CRAWLER Scout can provide terrain assessment as an active sensor, then apply that knowledge to adapt its gait for improved locomotion.

2. Related Work

Several techniques have been developed for classification, for traversability assessment, and for characterization of terrain. Motivation of these techniques in the area of mobile robotics splits into two basic areas – for navigational planning and for control. In other words, terrain classification is motivated by the two questions – 1) “Where is the robot going?” (navigation), and 2) How is it getting there?” (control). These questions are closely tied in outdoor environments, where terrain is varied and ease-of-control defines the path, as opposed to indoor environments, where the terrain is assumed constant and obstacles define the path.



Figure 1. Image of CRAWLER Scout emerging from a cinder block, which it can traverse using its narrow-passage gait.

Elevation maps and vision techniques can be used to assess terrain for navigation. In (Langer et al., 1994; Simmons et al., 1995; Gennery, 1999), the authors employed similar approaches using elevation maps, which are derived from stereo vision or laser rangefinders, to assign a traversability measure within a grid map cell based on height, slope, height variability, and/or terrain knowability. A search algorithm generated an optimal path from these measures. Similarly in (Seraji and Howard, 2002; Talukder et al., 2002; Huber et al., 1998), shape, texture, and color analysis aided in generating a path traversability measure. These works *predict* future vehicle-terrain interactions for the purposes of finding a global path and for determining an open-loop locomotion strategy. Our approach *measures* local vehicle-terrain interactions for the purpose of dynamically adapting locomotion.

Other approaches to measuring vehicle-terrain interaction have been developed for wheeled vehicles, some of which are based on the seminal work of Bekker (Bekker, 1969). Yoshida and Hamano (Yoshida and Hamano, 2002) developed a wheel-terrain interaction simulation model and a control law based on estimated wheel-slip to prevent the robot from digging itself into a trench. Iagnemma et al. estimated terrain cohesion and internal friction to calculate drawbar pull (force that pulls a wheel forward) (Iagnemma et al., 2002), and in (Iagnemma et al., 2003), maintained stability using ground contact angle estimates. Talukder et al. (Talukder et al., 2002) used a spring-mass model to estimate terrain compliance both to maintain a safe velocity and to predict vehicle dynamics. In each case, the robot required several sensors to estimate terrain conditions, additionally these techniques are

applicable to wheeled vehicles only. Our technique requires a single forward-looking camera, and it applies to legged vehicles with the potential for use on wheeled vehicles.

In the realm of legged vehicles, there are a variety of sensors applied to measuring limb-terrain interaction, including tactile, force, inclinometer, and gyroscopes. In the simplest case, tactile sensors detect the presence or absence of terrain, for example in Hirose's seminal work on control of a quadruped walker (Hirose, 1984), or in (Espenschied et al., 1996), in which a robot traversed a slatted surface. Other sensors can be used to estimate compliance or slope of the terrain for adaptive locomotion, as shown in (Wettergreen et al., 1995; Lewis and Bekey, 2002; Kurazume et al., 2002; Cham et al., 2004), to name a few. Our approach is similarly motivated. Its primary benefit over these other techniques is the sensor requirement. Cameras are typically in use for other purposes, thus no additional hardware is required. This is particularly beneficial for planetary exploration and applications for small-scale robotics, in which minimal weight is crucial. However, our technique does not preclude the use of other sensors, rather it is complementary.

3. Gait Bounce and Gait Roll

Our test platform, CRAWLER Scout, can locomote with a variety of gait classes, including swimming, narrow-passage, and bumpy-wheel for forward motion, and differential and body-shift for turning, as outlined in (Voyles et al., 2001; Voyles and Larson, 2005). The swimming gait is its most often-used gait. Figure 2 depicts a single gait cycle, wherein A-B) the arms are positioned forward to start the gait, B-C) the body is lifted to its pitch angle (also shown in Figure 3), C-D) the shoulders are rotated back to propel the body forward, and E-F) the arms are retracted and repositioned for the next cycle. User-defined parameters, such as pitch angle, elbow span, and length of stroke, define the arm trajectories at each phase. In a variation, the shoulders are rotated at different speeds to turn the body, analogous to turning a wheeled vehicle.

CRAWLER Scout navigates using visual servoing, which can be used to home in on an object by keeping it central in the image during forward locomotion, for example. Any displacement of the object (feature) off-center that is attributable to the robot's movement and not the feature's, is a visual servoing error. We employ sum of squared differences (SSD) template matching to track a feature (*see* (Nelson et al., 1993) for more details). The horizontal displacement, Δx , of the

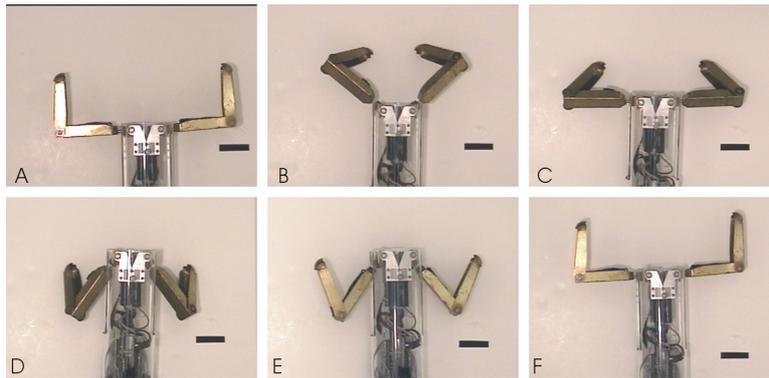


Figure 2. Depicts the various stages of the swimming gait. A-B) the arms are positioned, B-C) the body is lifted, C-D) the shoulders are rotated back for forward movement, and E-F) the arms are repositioned for the next cycle. The black line provides a reference for forward progress.



Figure 3. Lift phase of the swimming gait.

tracked feature from the center of the image is used to provide a corrective turning angle, $\theta = \arcsin(\Delta x/f)$ (where f is the focal length), to orient the body to the homed object. With ground-based robots, we are typically only interested in the horizontal displacement while the vertical displacement is thrown away, but in this work, we retain the latter information to derive gait bounce for terrain classification.

3.1. GAIT BOUNCE DERIVATION

During locomotion, the vertical displacement of a feature in the image plane across time provides a relative measure of the upward tilting motion of the camera, thus the *gait bounce* of the robot. In turn, gait bounce provides an estimate of terrain conditions, analogous to estimating the roughness of a road being traveled based on the intensity of your bounce as a passenger, or to estimating the stiffness of a terrain based on sinkage as you step. To use gait bounce as a terrain classifier,

we derived a formal, quantitative measure of the vertical motion of the robot, which is essentially an estimate of the pitch angle. The gait bounce “sensor” is active during limb-terrain contact portions of the gait cycle, when the motion of the limb against the terrain effects variation in the vertical tilt of the body.

In the simplest case, the robot is moving across a smooth, hard surface and the vision system is tracking a single, fixed feature with a camera. Assuming orthographic projection, the visual servoing error in the vertical plane of the tracked feature (Δy) can be used to derive body angle relative to the ground plane by $\theta = -\arctan(\Delta y/f)$, where f is the focal length. The top, dark dashed line of Figure 4 shows the raw gait bounce signal derived from tracking a single feature as CRAWLER Scout moves across carpet. Before time 90 (time is measured in number of images processed at 30Hz), the arms are positioning themselves for ground contact. From time 100 to 350, CRAWLER Scout is in ground contact – lifting, dragging, then dropping the body to effect forward motion. After time 350, the arms are repositioning for the next cycle and are not in ground contact. The signal generated after time 350 is not due to body tilt, rather body roll, which is problematic.

Body roll is attributable to CRAWLER Scout’s cylindrical shape, which can roll when neither arm is in ground contact, but more generally, it is attributable to variations in elevation between limb-ground (or wheel-ground) contact points commonly seen in rough terrain. A feature in the image plane is displaced vertically when the body rolls, thus it generates an erroneous gait bounce signal. In other words, roll increases the signal strength with no corresponding change of pitch angle. To compensate for this error, we estimate body roll using two fixed features in the image plane.

To estimate body roll, we first calculate the angle of the line segment formed by two tracked features relative to the horizontal plane,

$$\phi_i = \arctan\left(\frac{y2_i - y1_i}{x2_i - x1_i}\right), \quad (1)$$

where $\{x1_i, y1_i\}$ is the $\{x, y\}$ position in the image frame of feature 1 at time i . This angle is attributable both to body roll and to the horizontal misalignment of the two features, but it is only the former that affects gait bounce. An estimate of the latter (or nominal roll) is subtracted from ϕ_i , resulting in a zero-mean sequence of deviations from nominal:

$$\hat{\phi}_i = \phi_i - \frac{\sum_{j=1}^N \phi_j}{N}, \quad (2)$$

where N is the total number of image frames.

To derive roll-compensated gait bounce at each time step, we calculate the position of the center of the line $\{x_i, y_i\}$ relative to the center of the image plane $\{XC, YC\}$ as:

$$\{x_i, y_i\} = \left\{ \frac{x1_i + x2_i}{2} - XC, \frac{y1_i + y2_i}{2} - YC \right\}, \quad (3)$$

then we virtually roll the robot back to an upright position by rotating the point $\{x_i, y_i\}$ about the center of the image plane by $-\hat{\phi}_i$, which is roll at time i . The roll-compensated y-value, \hat{y} , is calculated as:

$$\hat{y}_i = x_i \sin(-\hat{\phi}_i) + y_i \cos(-\hat{\phi}_i). \quad (4)$$

Then we calculate the roll-compensated bounce angle θ using \hat{y} and the focal length f :

$$\theta_i = -\arctan\left(\frac{\Delta \hat{y}_i}{f}\right). \quad (5)$$

In compensating for roll, there is an underlying assumption that the centimeter-scale surface patch traversed during a gait cycle is locally planar. While this is clearly invalid for the rough terrains in which we are interested, it is a reasonable assumption for this calculation. The arm and body contact points form a triangle of support over the local terrain patch. Since the arm contact points typically stay fixed during forward motion (neglecting slip), they do not contribute to the roll during the forward stroke of the gait. The rear (body) contact point is dragged across the nonplanar surface patch, but the resulting erratic motion only manifests itself in gait bounce, not in body roll. Therefore, the planar assumption is reasonable for extracting body roll.

Roll-compensation results in the extraction of an additional characteristic function, similar to gait bounce, that contains terrain-specific information. In the current formulation we ignore gait roll, but we may be discounting some potentially useful information. We surmise that roll may be a complementary discriminator of rocky and level surfaces, as well as of compliant and noncompliant surfaces. In future work, we will attempt to incorporate this additional distinctive signature into our terrain characterization, but in this paper, terrain classification is based solely on gait bounce.

The dotted line of Figure 4 graphs the roll-compensated gait bounce. This new bounce measure reveals another concern, namely the upward data trend due to perspective distortion from what is more accurately modeled as a pinhole camera (not orthographic projection). As the robot moves toward tracked features, they move from the focus of expansion, which is at the center of the image for our fixed camera mounted parallel to the major body axis. Since we do not know depth

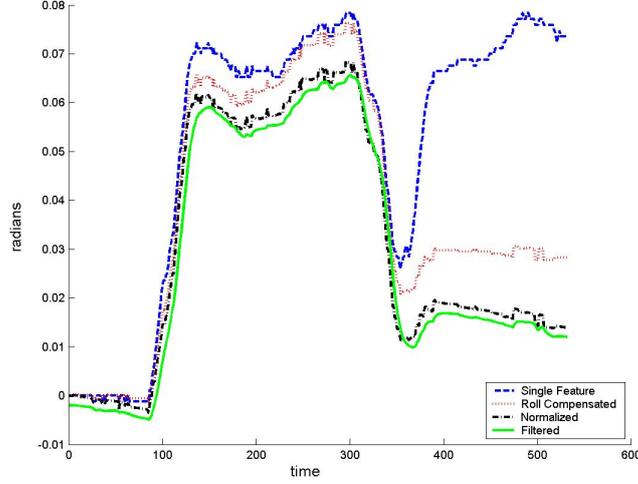


Figure 4. Gait bounce from tracking single object (dash), roll-compensated gait bounce (dot), normalized (dash-dot) for perspective distortion, and filtered (solid) (purposely offset for easier viewing).

to the features, we can't explicitly compensate for perspective. Instead, we crudely approximate the perspective distortion as a linear function of progress through the gait cycle. We estimate the slope of this function by fitting a line to a complete cycle of gait bounce data in the least squares sense. We then normalize for this distortion by subtracting the approximation from the original gait bounce. This results in a new, normalized gait bounce $\hat{\theta}$:

$$\hat{\theta}_i = \theta_i - \theta_{\approx i}, \quad (6)$$

where $\theta_{\approx i}$ is the linear least squares approximation of θ_i . In the case that features get closer than 1 meter, the nonlinearity of the distortion becomes apparent, but the system continues to perform quite well. The black dash-dot line of Figure 4 graphs the normalized gait bounce.

The final issue that we address is local disturbances to the gait bounce. This is a high frequency component relative to the gait cycle frequency and it is filtered out. The filtered gait bounce measure is the solid line of Figure 4 (which is artificially offset from the normalized gait bounce for better viewing). This final gait bounce measure is the result of tracking two features to compensate for roll, of then normalizing the affects of perspective distortion, and finally of filtering out the high frequency component.

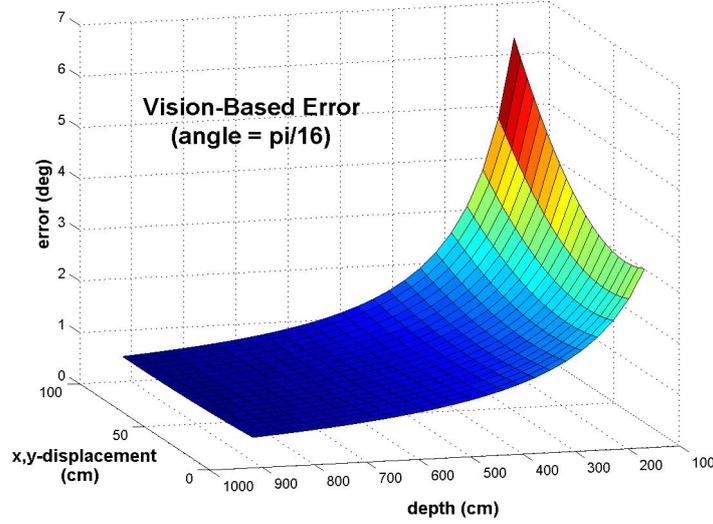


Figure 5. Error of vision-based angle estimation relative to feature placement. x,y -displacement refers to equal feature translation in the horizontal and vertical directions with respect to the image center.

3.2. ERROR ANALYSIS

As with all sensing devices, gait bounce estimation is subject to error. Error is partly attributable to the displacement of the camera relative to the axis of rotation (i.e. the camera is atop the robot but rotation is about the ground contact point). We defined this error as:

$$E = \theta_R - \theta_V \quad (7)$$

where θ_R is the pitch angle of the robot and θ_V is the angle estimation based on the vision data. θ_V is calculated as:

$$\theta_V = -\text{atan}\left(\frac{y' - \hat{y}'}{f}\right) \quad (8)$$

where y' and \hat{y}' are the y -positions of the ideal projection of the feature onto the image plane prior to and after rotation, respectively. This error does not take lens distortion into account.

The error was calculated for a variety of feature positions. Figure 5 graphs the ideal error with respect to feature position when the robot is at a pitch angle of $\pi/16$. Orthographic projection would eliminate this error, however infinite depth features viewed with a pinhole camera equate to orthographic projection. This is reflected in the graph – error decreases as depth increases. Conversely, error increases as the x - and

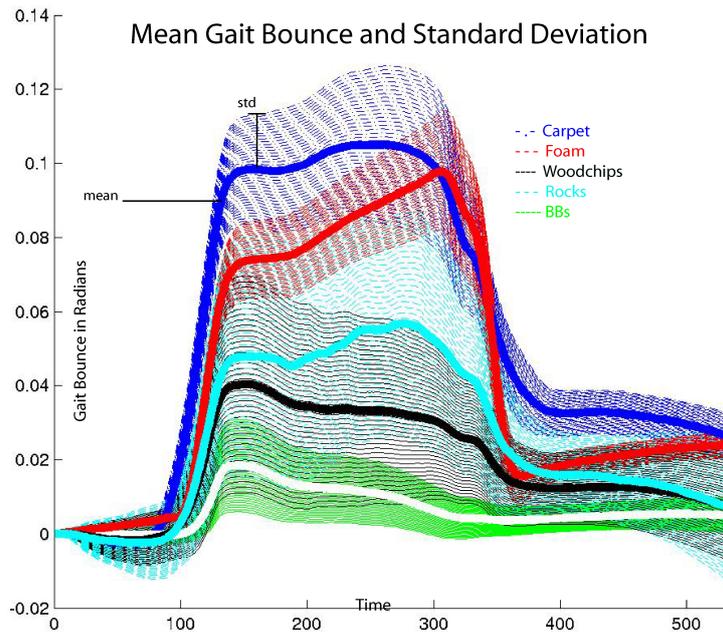


Figure 6. Mean and standard deviation of gait bounce for each tested terrain. The solid line indicates the mean, while the surrounding contours of similar shade signify the standard deviation.

y-displacement of the feature increases (relative to the image center), but depth of the feature is the primary and dominate contributor. Reasonable error is obtained at about 350 cm resulting in an error range from 1° to 3° , depending on the $\{x,y\}$ -position of the feature. Feature depth at 600 cm results in a range of $.5^\circ$ to 1.5° .

Image distortion due both to curvature of the lense and to imperfections in the lense also contributes to error in gait bounce estimation. This error can be managed in two ways. First, distortion is most significant on the edges of the image, thus ideal features are at the center of the image. Second, calibration techniques (see (Willson and Shafer, 1994; Lenz and Tsai, 1988)) exist to mitigate these distortions. Further analysis is needed to assess the range of this error, but a preliminary look indicates that it will not overwhelm the sensor.

4. Terrain Classification

We investigated the application of the gait bounce signal towards terrain classification. The current prototype of CRAWLER Scout is insufficiently hardened for use in some natural terrains, therefore we selected the artificial terrains foam and BBs for our experiments. BB's

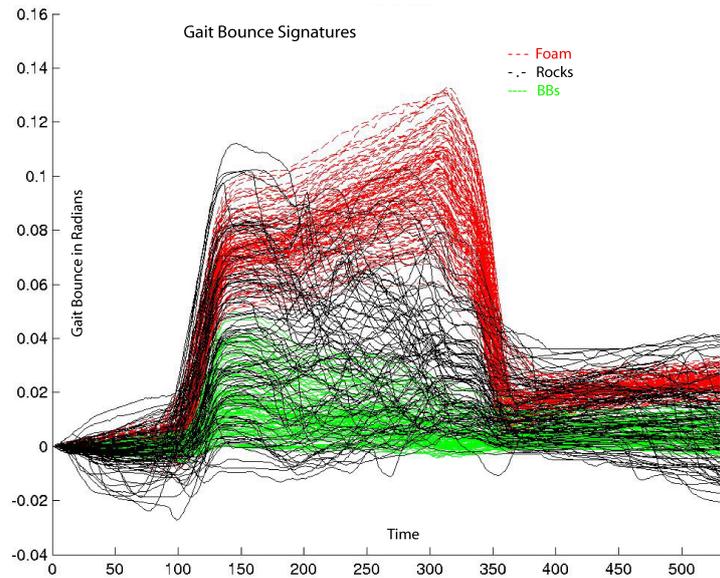


Figure 7. All gait bounce signals generated during experiments from 3 of the 5 tested terrains (the other 2 were not included for clarity). The overlap in signal demonstrates that the classification problem is nontrivial.

simulate highly compliant surfaces, such as sand, and foam simulates semi-compliant surfaces, such as grass. We also used woodchips, rocks, and carpet. These represent a range of terrain characteristics that might be encountered either inside a felled building or in natural terrains.

We walked CRAWLER Scout across these terrains using a periodic swimming gait while its vision system tracked and recorded the position of two stationary features. Currently, we employ manual feature selection, but we are working on an automated process using the KLT tracker (see (Shi and Tomasi, 1994)). Feature locations varied in depth and in $\{x,y\}$ -position throughout the experiments. Data from approximately 200 gait cycles over each of the five terrains was collected and processed off-line to generate a gait bounce signal as described above. Our classification task was to label a gait bounce signal from a single gait cycle with one of the 5 terrains.

Figure 6 shows the mean gait bounce signal for each terrain, as well as the standard deviation. Figure 7 shows the signals generated from all collected samples for three of the five terrains (the other two were omitted for clarity). While these figures show some visually distinct trends for each terrain, there is considerable overlap, making the classification problem nontrivial.

Classification begins with the mapping of an object to a feature vector, in which a feature represents an aspect or character of the ob-

ject. The domains and ranges of these features define the feature space. Pattern classification techniques using spatial discriminants divide this feature space, then assign a class label to each division (i.e. subspace). In the simplest case, a feature vector is given the class label of the subspace to which it is mapped. These pattern classification techniques are distinguished by the functions used to divide the feature space (*see* (Duda et al., 2001) for review).

In our experiments, an object is defined as a gait bounce signal from a single cycle generated from CRAWLER Scout walking across one of the five terrains. Three distinct feature spaces were considered, but each was limited in length due to a discriminant analysis constraint, whereby the length of the feature vector cannot exceed the number of training samples.

The first of the three feature spaces was defined by the subsampling of the gait bounce signal. The sampling rate of image acquisition is 30Hz, but with the subsampling to conform to the discriminant analysis constraint, the sampling rate fell to 6Hz. The resultant gait bounce signal was used as the feature vector (i.e. the first element of the vector is the signal strength at time 1). The second feature space was defined by the application of a fast Fourier Transform (FFT) over the signal. The first 15 magnitude components of the lowest frequencies comprised the feature vector. The remaining frequency components approached 0, having little impact on discrimination. Similarly, the third feature space was defined by applying an FFT over each of 5 manually-selected segments of the signal. By segmenting the data, we could isolate low-magnitude, but key, frequencies for comparison, which might otherwise be overshadowed. The first 10 magnitude components from each segment were concatenated to form the feature vector. Each of the 1000 gait bounce signals generated from the collected data was preprocessed to form each of these 3 feature vectors.

We experimented with a variety of classification techniques to determine the best approach for this particular problem. We used a back-propagation artificial neural network (ANN) with a logarithmic sigmoid transfer activation function for all units, varying the number of hidden units. We looked at a well-known statistical method – discriminant analysis, using linear, quadratic, and logarithmic subspaces to maximize separability. We also looked at support vector machines, using polynomial, radial basis function, and sigmoid kernels.

For each classification method, the data was randomly divided 50 times into train and test subsets, and the success rate of any technique was defined as the average classification rate of the 50 test sets using that technique. Table I summarizes the results. The feature vector resulting from applying the FFT over the entire signal consistently

Table I. Mean and best classification rates over 50 random test sets using artificial neural networks, discriminant analysis, and support vector machines. Feature vectors and user-defined parameters were varied. Total is mean classification rate across all terrains.

| | Total | Carpet | BB | Chip | Foam | Rock |
|------------|-----------|--------|----|------|------|------|
| ANN | | | | | | |
| mean | 62 | 86 | 80 | 26 | 87 | 30 |
| best | 64 | 88 | 88 | 29 | 82 | 30 |
| DA | | | | | | |
| mean | 74 | 93 | 85 | 51 | 88 | 56 |
| best | 83 | 96 | 89 | 70 | 92 | 67 |
| SVM | | | | | | |
| mean | 46 | 57 | 86 | 22 | 53 | 14 |
| best | 64 | 93 | 91 | 31 | 80 | 21 |

Table II. Sample confusion matrix from discriminant analysis classification. 1.0s along the diagonal and 0.0s elsewhere signify perfect classification. Most of the confusion lies with rocks and woodchips, where 19% of woodchips were classified as rocks and 28% of rocks were classified as woodchips.

| | Carpet | BB | Chip | Foam | Rock |
|---------------|-------------|-------------|-------------|-------------|-------------|
| Carpet | 0.91 | 0.05 | 0.02 | 0.01 | 0.0 |
| BB | 0.04 | 0.83 | 0.05 | 0.04 | 0.04 |
| Chip | 0.0 | 0.12 | 0.64 | 0.04 | 0.19 |
| Foam | 0.03 | 0.01 | 0.04 | 0.92 | 0.0 |
| Rock | 0.01 | 0.05 | 0.28 | 0.03 | 0.64 |

produced the best classification. This feature vector and quadratic discriminant analysis (QDA) provided the best classification rate overall at 83%.

A closer look at these results shows that the majority of misclassifications occurs between rocks and woodchips – both rough surfaces that produce stochastic signals. Table II shows a characteristic confusion matrix that highlights this point. 19% of woodchips were classified as rocks and 28% of rocks were classified as woodchips. A cell value indicates the percentage of data that is classified as the column heading but is of the row heading type.

5. Viability of Nonlegged Platforms

Although our terrain classification technique has been investigated here in the context of a legged vehicle, we believe this technique can be extended to other modes of locomotion, including wheels and tracks over nonplanar environments. To support this claim, we performed a brief test on a radio-controlled tank to assess viability. We drove the tank across three of the toughest terrains – rocks, woodchips, and BB’s – and applied the same classification approach.

It is important to point out that tracked and wheeled vehicles do not inject periodicity into the gait bounce the way a legged vehicle does, therefore they are not actively sensing the terrain beyond the action of moving over it. (While it is possible to actively sense the terrain with redundant drive wheels, we did not consider that in these experiments.) For our legged vehicle, we make explicit use of the periodicity of the gait and the unique gait phases. This controlled periodicity dominates any natural periodicity of the terrain, occurring in rolling hills or a terraced landscape, for example. For open-loop, non-periodic locomotion modes, any natural terrain periodicity could skew the results by providing false classification features. Therefore, we deliberately created a BB terrain with periodicity in the contour to demonstrate that our technique can avoid the pitfall of this induced periodicity.

Gait bounce was derived using the equations outlined in Section 3 except for the low-pass filter. We did not apply the low-pass filter to the tank data, because the distinguishing features of the terrain manifest themselves in the high frequency components as compared to the CRAWLER Scout. This is a result of the tank’s higher forward velocity.

Following previous protocol, the gait bounce signal was mapped into frequency space and classified with QDA. The data sets for this experiment were relatively small (11-18), therefore a different testing method was applied. For each feature vector, we determined class membership using a QDA classifier trained on all other feature vectors (i.e. leave-one-out cross-validation). The classification rates for rocks, woodchips, and BB’s were 78, 91, and 94%, respectively. Note that BB’s yielded the highest accuracy despite the mixture of terrain topologies (i.e. some level, some periodic). While this small set is not intended to be conclusive, we think it convincingly demonstrates the potential for general applicability of the technique.

6. Gait Efficiency

A key motivation behind terrain classification is to adapt the gait to best maneuver across the terrain. Gait adaptation refers to gait *class* selection (e.g. swimming or bumpy wheel) or to gait *parameter* selection. Swimming gait parameters include speed of arm motion, height of body lift, length of forward stroke, span of arms, and limb trajectory during the retraction-and-placement phase. To highlight the benefits of adaptation, we developed an efficiency measure and conducted a small experiment to demonstrate its applicability. While other metrics may be used to evaluate performance, efficiency is a dominant feature for resource-constrained robots.

We walked CRAWLER Scout across carpet, woodchips, and foam using two distinct swimming gaits ($S1$ and $S2$), as well as the bumpy wheel gait. The swimming gaits differed only by angle of body lift, whereby $S1$ was pitched 5° and $S2$ was pitched 10° . The bumpy wheel gait simulates a single-spoke wheel by fixing the elbows at 90° , then rotating the arm about the shoulder. Current draw and distance traveled for three consecutive gait cycles was recorded over multiple runs. From this data, we calculated average total energy, J (*Joules*), necessary to execute the three gait cycles. Energy is defined over a continuous space, but we approximated this area as:

$$E = \frac{1}{f} \sum_{t=1}^{n-1} vI_t \quad (9)$$

where f is our sampling rate, n is the number of samples, v is volts (which is constant), and I_t is current draw of sample t . We calculated the amount of energy required to travel a cm as:

$$F = \frac{E}{\Delta d} \quad (10)$$

where Δd is total distance traveled, defining F as our measure of efficiency. Table III outlines the results of this experiment.

First note that the bumpy wheel gait takes significantly more energy to execute per cycle. This is due to its long retraction-and-placement phase necessitated by the current lack of sliprings (the shoulders cannot continuously rotate like a wheel). Power is being consumed during this phase without forward progress of the robot. Bumpy wheel gait does not outperform either of the swimming gaits with respect to efficiency, but when the sliprings are incorporated, bumpy wheel may prove to be the better gait on some terrains.

Second, note from the table that gait $S2$ takes, on average, 61% less J/cm than $S1$ when traversing carpet, yet takes 41% more J/cm

Table III. *Energy and Efficiency calculated from current draw and distance traveled while CRAWLER Scout executed three gait cycles of two distinct swimming gaits (columns 1 and 2) and the bumpy wheel gait (column 3).*

| | Energy (J) | | | Efficiency (J/cm) | | |
|------|------------|------|-------|-------------------|------|------|
| | S1 | S2 | bump | S1 | S2 | bump |
| carp | 63.6 | 71.9 | 106.5 | 3.70 | 1.43 | 2.38 |
| chip | 62.4 | 75.4 | 116.4 | 2.12 | 3.57 | 3.46 |
| foam | 63.7 | 90.3 | 112.7 | 1.76 | 1.59 | 2.51 |

than S1 when traversing woodchips. This effect can be explained by the compliance of the surface. Woodchips are compliant, therefore greater force exerted on the surface results in more sinkage (up to a threshold). The further the arms are buried, the more difficult it is to move. Gait S2's body lift parameter is greater than S1's, thus sinking CRAWLER Scout deeper into the woodchips prior to the forward stroke. This data further supports the idea that efficiency can be improved when the gait is adapted to the known terrain, as has been shown through several other studies (e.g. (Weingarten et al., 2004; Kurazume et al., 2003; Marhefka and Orin, 1997)).

7. Future Directions

As discussed in Section 3, error is sensitive to both camera accuracy and feature selection. With respect to feature selection, depth is key, thus given the choice, distant features should be used. But with a single camera in an unknown environment, depth information is not available. We are developing a collaborative method to extract odometric information using shape from motion that estimates the depth of each feature (see (Demir et al., 2004)). This will allow us to explicitly compute the perspective distortions of each feature, rather than relying on inaccurate linear estimates. In that work, multiple robots function as known landmarks, which are necessary to scale the odometric system. We are working on incorporating odometric and feature depth information.

Obviously, a key assumption of this work is that features are readily available. This is a constraint that cannot always be satisfied, but there are solutions. The gait bounce technique is primarily for small-scale robots that typically function as teams, thus a collaborative approach whereby other robots function as features, such as that used for the odometric system, is reasonable. Alternatively, gait bounce estimation

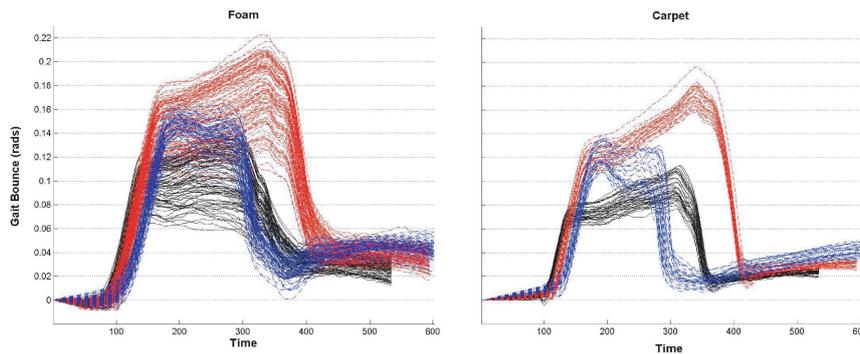


Figure 8. Gait bounce signals for foam (left) and carpet (right). The groupings differ by gait parameters, including pitch angle and span of arms (i.e. angle of elbows).

can be obtained from an inclinometer or accelerometer. The same trained classifiers may potentially be used regardless of sensor. We have included a two-axis, 2-g accelerometer in a new CRAWLER Scout prototype and have begun experimentation.

Using an inclinometer may be viewed as an alternative to the vision technique that we used to estimate gait bounce. We argue that using data and hardware obtained and used for other purposes (vision and visual servoing are required for many search and rescue and planetary exploration applications) ultimately reduces battery usage and weight – a key issue for resource-constrained robots. Furthermore, the use of one does not obviate the other. They may serve as back-up to each other or may be used in tandem to improve estimation.

Tuning gait parameters can improve performance as discussed above, however, it can complicate terrain classification. Traversing the same terrain using different gait parameters will generate distinct gait bounce signals, as shown in Figure 8. This necessitates broad generalization within each terrain classifier, while maintaining separability (distinction) across classifiers.

One solution is to use a handful of gaits (and their corresponding classifiers) to characterize the terrain, then adapt the gait accordingly. However, we are currently working on a more extensible classifier that uses discriminant analysis in combination with a hidden Markov model (see (Larson et al., 2004)). It is based on the segmentation of the signal that separates out high-information content at key points in the gait cycle. For example, in Figure 8, the point of ground-contact appears to be a distinguishing feature between carpet and foam, whereby the gait bounce typically falls below 0 when the robot is traversing the semi-compliant surface foam. Segmentation concentrates classification to persistent features, regardless of gait parameters.

This new approach should be beneficial also in *characterizing* the terrain, which is the ultimate goal of this work. We are not necessarily interested in knowing that the terrain is sand, rather that it has certain compliant characteristics. Again, by segmenting the data, we can isolate portions of the signal (e.g. the ground contact point) that provide specific clues as to the character of the terrain.

8. Conclusions

We presented a method for terrain classification based on vehicle-terrain interaction. It provides valuable feedback for effective locomotion over unknown and rough terrains, and for qualitative analysis of the terrain, both of which are important skills for autonomous operation during search-and-rescue and planetary exploration. Our approach requires only a sensor for gait bounce, which we deduce from a single, forward-looking camera with relatively little processing of the visual data. In fact, no additional image processing is required if feature tracking is implemented for other purposes, such as visual servoing. This reuse of hardware and preprocessed visual information reduces the load and processing time, thus decreasing battery use and operational weight. A rigorous study with more than 700 trials from a real robot demonstrated the viability of the approach for legged vehicles on five laboratory terrains. The average accuracy of the discriminant analysis classifier over a single gait cycle (approximately 3cm) was 83%. Preliminary work with a treaded vehicle demonstrated its potential for application to non-legged platforms. The average accuracy over a smaller benchmark was over 80%.

Acknowledgements

We thank Seth Hulst and Andrew Willouer for their time-consuming efforts in data collection. We thank Berk Yesin for his vision code and Monica LaPoint for her systems work. We also thank the reviewers for their constructive comments.

References

- Bekker, G.: 1969, *Introduction to Terrain-Vehicle Systems*. University of Michigan Press.
- Cham, J., J. Karpick, and M. Cutkosky: 2004, 'Stride period adaptation for a biomimetic running hexapod'. *Int'l Journal of Robotics Research* **23**(2), 141–153.
- Demir, G., R. Voyles, and A. Larson: 2004, 'Motion estimation with cooperatively working multiple robots'. In: *Proc. IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*. p. to appear.
- Duda, R., P. Hart, and D. Stork: 2001, *Pattern Classification*. John Wiley and Sons, Inc., 2nd edition.
- Espenschied, K. S., R. D. Quinn, R. D. Beer, and H. J. Chiel: 1996, 'Biologically based distributed control and local reflexes improve rough terrain locomotion in a hexapod robot'. *Robotics and Autonomous Systems* **18**(1–2), 59–64.
- Gennery, D.: 1999, 'Traversability analysis and path planning for a planetary rover'. *Autonomous Robots* **6**, 131–146.
- Hirose, S.: 1984, 'A study of design and control of a quadruped walking vehicle'. *Int'l Journal of Robotics Research* **3**, 113–133.
- Huber, D., L. Denes, M. Hebert, M. Gottlieb, B. Kaminsky, and P. Metes: 1998, 'A spectro-polarimetric imager for intelligent transportation systems'. In: *Proc. of SPIE - The International Society for Optical Engineering*, Vol. 3207. pp. 94–102.
- Iagnemma, K., A. Rzepniewski, S. Dubowsky, and P. Schenker: 2003, 'Control of robotic vehicles with actively articulated suspensions in rough terrain'. *Autonomous Robots* **14**(1), 5–16.
- Iagnemma, K., H. Shibly, and S. Dubowsky: 2002, 'On-line terrain parameter estimation for planetary rovers'. In: *Proc. of the IEEE Int'l Conference on Robotics and Automation*, Vol. 3. pp. 3142–3147.
- Kurazume, R., A. Byong-Won, K. Ohta, and T. Hasegawa: 2003, 'Experimental study on energy efficiency for quadruped walking vehicles'. In: *Proc. IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*. pp. 613–618.
- Kurazume, R., K. Yoneda, and S. Hirose: 2002, 'Feedforward and feedback dynamic trot gait control for quadruped walking vehicle'. *Autonomous Robots* **12**, 157–172.
- Langer, D., J. Rosenblatt, and M. Hebert: 1994, 'A behavior-based system for off-road navigation'. *IEEE Transactions on Robotics and Automation* **10**(6), 776–783.
- Larson, A., R. Voyles, and G. Demir: 2004, 'Terrain classification through weakly-structured vehicle/terrain interaction'. In: *Proc. of the IEEE Int'l Conference on Robotics and Automation*. pp. 218–224.
- Lenz, R. and R. Tsai: 1988, 'Techniques for calibration of the scale factor and image center for high accuracy 3-D machine vision metrology'. *IEEE Trans. Pattern Analysis and Machine Intelligence* **10**(5), 713–720.
- Lewis, M. and G. Bekey: 2002, 'Gait adaptation in a quadruped robot'. *Autonomous Robots* **12**, 301–312.
- Marhefka, D. W. and D. E. Orin: 1997, 'Gait planning for energy efficiency in walking machines'. In: *Proc. of the IEEE Int'l Conference on Robotics and Automation*. pp. 474–480.
- Nelson, B., N. Papanikolopoulos, and P. Khosla: 1993, *Visual servoing – real-time control of robot manipulators based on visual sensory feedback*, Chapt. Visual servoing for robotic assembly, pp. 129–164. World Scientific Publishing Co. Pte. Ltd.

- Seraji, H. and A. Howard: 2002, 'Behavior-based robot navigation on challenging terrain: A fuzzy logic approach'. *IEEE Transactions on Robotics and Automation* **18**(3), 308–321.
- Shi, J. and C. Tomasi: 1994, 'Good features to track'. In: *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*. pp. 21–23.
- Simmons, R., E. Krotkov, L. Chrisman, F. Cozman, R. Goodwin, M. Hebert, L. Katragadda, S. Koenig, G. Krishnaswamy, Y. Shinoda, and W. Whitaker: 1995, 'Experience with rover navigation for lunar-like terrain'. In: *Proc. IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, Vol. 1. pp. 441–446.
- Talukder, A., R. Manduchi, R. Castano, K. Owens, L. Matthies, A. Castano, and R. Hogg: 2002, 'Autonomous terrain characterisation and modelling for dynamic control of unmanned vehicles'. In: *Proc. IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*. pp. 708–713.
- Voyles, R. and A. Larson: 2005, 'TerminatorBot: A novel robot with dual-use mechanism for locomotion and manipulation'. *IEEE/ASME Transactions on Mechatronics* **to appear**.
- Voyles, R., A. Larson, K. Yesin, and B. Nelson: 2001, 'Using orthogonal visual servoing errors for classifying terrain'. In: *Proc. IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, Vol. 2. pp. 772–777.
- Weingarten, J., G. Lopes, M. Buehler, R. Groff, and D. Koditschek: 2004, 'Automated gait adaptation for legged robots'. In: *Proc. of the IEEE Int'l Conference on Robotics and Automation*. pp. 2153–2158.
- Wettergreen, D., H. Pangels, and J. Bares: 1995, 'Behavior-based gait execution for the Dante II Walking Robot'. In: *Proc. IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, Vol. 3. pp. 274–279.
- Willson, R. and S. Shafer: 1994, 'What is the center of the image?'. *Journal of the Optical Society of America* **11**(11), 2946–2955.
- Yoshida, K. and H. Hamano: 2002, 'Motion dynamics and control of a planetary rover with slip-based traction model'. In: *Proc. of SPIE - The International Society for Optical Engineering*, Vol. 4715. pp. 275–286.

Contact Author

Richard Voyles
Department of Computer Science and Engineering
University of Minnesota
200 Union Street S.E.
4-192 EE/CS Building
Minneapolis MN 55455
voyles@cs.umn.edu
office: 612 624-8306
fax: 612 625-0572

Figure Legends

Figure 1. Image of CRAWLER Scout emerging from a cinder block, which it can traverse using its narrow-passage gait.

Figure 2. Depicts the various stages of the swimming gait. A-B) the arms are positioned, B-C) the body is lifted, C-D) the shoulders are rotated back for forward movement, and D-E) the arms are repositioned for next cycle.

Figure 3. Lift phase of the swimming gait.

Figure 4. Gait bounce from tracking single object (dash), roll-compensated gait bounce (dot), normalized (dash-dot) for perspective distortion, and filtered (solid) (purposely offset for easier viewing).

Figure 5. Error of vision-based angle estimation relative to feature placement. x,y-displacement refers to equal feature translation in the horizontal and vertical directions with respect to the image center.

Figure 6. Mean and standard deviation of gait bounce for each tested terrain. The solid line indicates the mean, while the surrounding contours of similar shade signify the standard deviation.

Figure 7. All gait bounce signals generated during experiments from 3 of the 5 tested terrains (the other 2 were not included for clarity). The overlap in signal demonstrates that the classification problem is nontrivial.

Figure 8. Gait bounce signals for foam (left) and carpet (right). The groupings differ by gait parameters, including pitch angle and span of arms (i.e. angle of elbows).

Table Legends

Table I. Mean and best classification rates over 50 random test sets using artificial neural networks, discriminant analysis, and support vector machines. Feature vectors and user-defined parameters were varied, as indicated in the text. *Total* is mean classification rate across all terrains.

Table II. Sample confusion matrix from discriminant analysis classification. 1.0s along the diagonal and 0.0s elsewhere signify perfect classification. Most of the confusion lies with rocks and woodchips, where 19% of woodchips were classified as rocks and 28% of rocks were classified as woodchips.

Table III. Energy and Efficiency calculated from current draw and distance traveled while CRAWLER Scout executed three gait cycles of two distinct swimming gaits (columns 1 and 2) and the bumpy wheel gait (column 3).

