Proceedings of the 2007 IEEE/RSJ International
Conference on Intelligent Robots and Systems
San Diego, CA, USA, Oct 29 - Nov 2, 2007

Evolving Gaits for Increased Discriminability in Terrain Classification

Amy C. Larson, Richard M. Voyles, Jaewook Bae, and Roy Godzdanker

***Abstract*— Limbs are an attractive approach to certain niche robotic applications, such as urban search and rescue, that require both small size and the ability to locomote through highly rubbled terrain. Unfortunately, a large number of degrees of freedom implies there is a large space of non-optimal locomotion trajectories (gaits), making gait adaptation critical. On the other hand, these extra degrees of freedom open many possibilities for active sensing of the terrain, which is essential information for adapting the gait. In previous work, we developed a metric for terrain classification that makes use of the loping body motion (i.e. gait bounce) during locomotion. In this work we present a framework for evolving gaits to better differentiate the gait bounce signal across terrains. This framework includes a limb/terrain interaction model that estimates gait bounce based on established models of wheel/terrain interaction, and an objective function that can be optimized for terrain discriminability. Additional objective functions for improved locomotion are presented, as well as culling agents that help guide the evolution process away from real-world impossibilities.**

I. INTRODUCTION

Limbs are an attractive approach to certain niche robotic applications, such as urban search and rescue, that require both small size and the ability to locomote through highly rubbled terrain. High degrees of freedom allow limbs to fold and stow to minimize envelope and can also utilize gearing traction to climb and drag over obstacles that are large relative to envelope size. Unfortunately, a large number of degrees of freedom implies there is a large space of non-optimal locomotion trajectories – in fact, many gait trajectories lead to no motion at all, making self-adaptation critical. On the other hand, these extra degrees of freedom open many possibilities for active sensing during locomotion that are not possible with lower degree of freedom actuation mechanisms.

Actively sensing the terrain allows self-adaptation of a robot's gait to improve locomotion efficiency. In previous work [1], [2], we developed a *gait bounce* metric for terrain classification that makes use of the loping body motion of a crawling robot to determine the type of terrain underfoot. The bounce of the gait can be measured visually or by tiltmeter and classified by the characteristic signature. The signature can be qualified and enhanced by the locomotion gait used, hence the robot system acts as an active sensor. Since the gait impacts the ability to distiguish different terrains, we are investigating the optimization of the gait to improve discriminability of gait bounce signatures used for terrain classification.

A. Larson is Adjunct Faculty of Computer Science, University of Minnesota, Minneapolis USA larson@cs.umn.edu
J. Bae is with the Electrical Engineering Department, University of Minnesota, Minneapolis USA baeja@cs.umn.edu
R. Voyles is Faculty of Computer Engineering, University of Denver, CO USA rvoyles@du.edu
R. Godzdamker is with Computer Engineering, University of Denver, CO USA roy.godzdanker@du.edu

Fig. 1. *Image of two TerminatorBots traversing rocks and woodchips.*

We have framed this problem for use in genetic algorithms (GAs). The most critical aspect of GAs is the objective function (or fitness function) that guides selection for the evolutionary process. In this paper we present an objective function that optimizes differences in gait bounce signatures. All work is done in simulation, therefore we also developed a limb/terrain interaction model based on soil characteristics and our specific robot platform to derive an estimation of gait bounce.

Our testing platform is TerminatorBot (Figure 1), which 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, TerminatorBot 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 natural, rough terrains of these applications. As a study in adaptation, TerminatorBot can provide terrain assessment as an active sensor, then apply that knowledge to adapt its gait for improved locomotion.

II. RELATED WORK

Limb-terrain interactions can be measured and analyzed using a variety of sensors to inform motion control. In the simplest case, touch sensors detect the presence or absence of terrain, such as in Hirose's seminal work on control of a quadruped walker [3], or in [4], in which a robot traversed a slatted surface. Force sensors can be used to estimate compliance or slope of the terrain, as shown in [5], [6], [7], to name a few. These approaches do not classify terrain.

Wheel-terrain interactions have been measured and analyzed in a method more comparable to our work. Iagnemma et al. have a series of papers related to terrain analysis including a method to estimate terrain cohesion and internal friction [8] and to classifiy terrain by analyzing vibration [9]. Talukder et al. [10] use a spring-mass model to estimate terrain compliance both to maintain a safe velocity and to predict vehicle dynamics. These techniques are applicable only to wheeled vehicles.

Many researches have used learning methods, and genetic algorithms in particular, to evolve gaits [11], [12], [13], [14]. The primary focus of the majority of this work is to optimize the gait for efficient locomotion. While our investigations will also look at optimizing for efficiency, the primary contribution of this work is optimizing for discriminability.

III. Classifying Terrains Using Gait Bounce

Recognition of the terrain type allows self-adaptation of the gait to improve locomotion efficiency (measured either as distance traveled per time or distance per energy). In previous work, we categorized terrain into 5 distinct learned classes using the "gait bounce", which is an estimation of the loping motion of a robot as it locomotes. As the robot moves across an unknown terrain, the bounce of the gait can be measured either visually or with a tiltmeter, providing a characteristic signature for the terrain classification. The signature can be qualified and enhanced by the locomotion gait used, hence the robot system acts as an active sensor of terrain.

For TerminatorBot, a gait bounce signature over a non-compliant, smooth surface using a typical swimming gait (an on-land gait that mimics the breast stroke as shown in Fig. 2) produces almost a flattened bell curve (see Fig. 3). At the beginning and end of the gait, the body rests on the ground and there is no limb/terrain interaction as the robot moves its limbs through the air into position, thus the ends are at zero. The flattened bell shape is produced as the robot lifts itself off the ground, then the inclination levels off as the robot drags its body, until finally it sets its body back on the ground. The robot is cylindrical, thus subject to rolling when both limbs are not in contact with the ground. The final bounce (solid, green line) is derived from 2 visual features and a compensation for the roll.

Figure 4 shows sample gait bounce signatures for three of the five terrain types (rocks, foam (to simulate grass), and bb's (to simulate sand), carpet and woodchips). This figure demonstrates the difficulty of classifying terrains as more types are added, particularly those with similar composition. Consider the class of soils, which are all compliant surfaces but differ by internal frictional forces, thus result in different levels of sinkage when force is applied.

Since the gait impacts the ability to distiguish different terrains, we wanted to explore optimization of the gait to improve the discriminability of gait bounce signatures used for terrain classification.

IV. Evolving Gaits

In our previous work of terrain classification, multiple gait bounce signatures from all terrains were mapped into

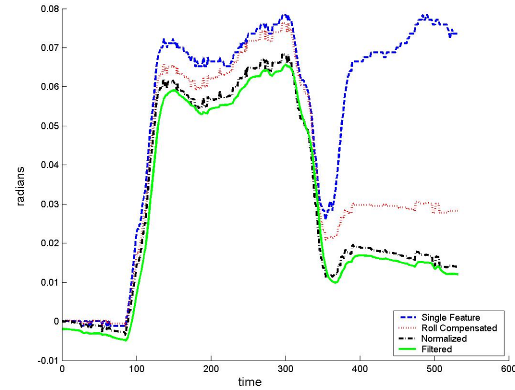

Fig. 3. *Derivation of gait bounce signature using visual information recorded while traversing carpet.*

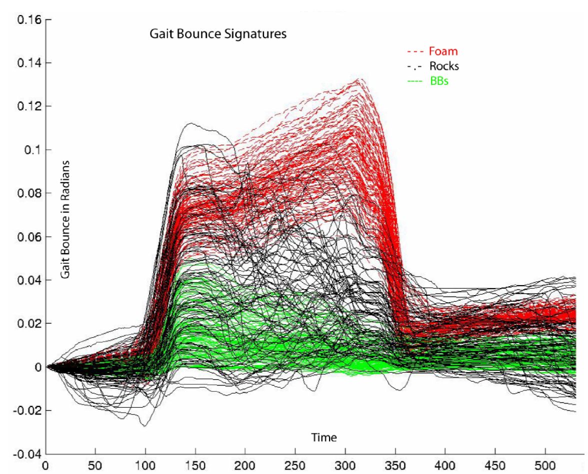

Fig. 4. *3 of the 5 gait bounce signals generated during experiments from all 5 terrains. The overlap in signal demonstrates that the classification problem is nontrivial.*

frequency space using a fast Fourier transform (FFT). Discriminant analysis was then applied to find the subset of frequencies that best distinguished terrains, and a characteristic frequency vector was established for each terrain. Gait bounce from an unknown terrain was similarly processed and then assigned to the terrain to which its frequency vector was the closest.

Success of the above described terrain classification technique relies on the distinguishability of the gait bounce in frequency space for each terrain. Therefore, our goal for evolution is to maximize the distance of the frequency vectors resulting from different terrains.

A. Encoding Gaits

Gaits consist of a series of via points, each signified by a 6-element joint vector $(\vec{\Theta})_i$. Gaits can also be described by a starting position and a series of directional vectors $(\Delta\vec{\Theta})_i$ that specify joint changes. For the GA, the latter approach was used to insure joint differences were limited, especially during mutation and crossover. Additionally, we

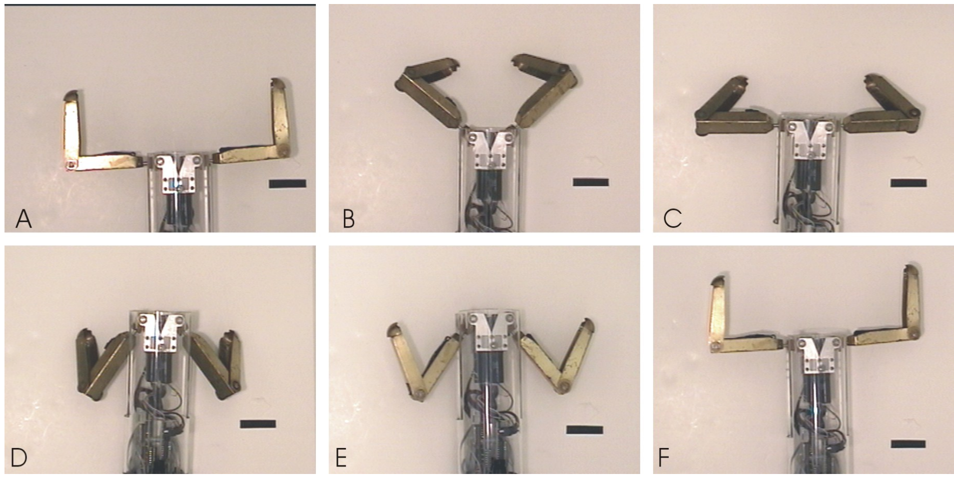

Fig. 2. *Progression of the on-land swimming gait. The arms are brought forward into position (A), rotated to lift the body off the ground (B), then pulled back to propel the body forward (C). The black tape provides a reference for forward progress.*

considered only symmetric gaits, meaning the arms are moving synchronously in mirrored trajectories, thus we need only specify a via as a *3*-element joint vector.

The starting point of the gait was selected randomly within a predefined cartesian space in which the end-effector was in contact with the ground and the arm was collision-free. The inverse kinematic function was applied to the randomly selected (x, y, z) position to derive the joint angles. All gaits contained 8 via points, but only a random number of them were defined, leaving the rest at $(0, 0, 0)$. Any of the 8 steps could be modified during evolution, thus effectively increasing or decreasing the number of vias. To initially define a via point, each joint difference was randomly selected in a range of $[-15, 15]$ degrees. Degrees were selected to ease encoding as 1^o increments are of sufficient precision, thus integers could be used. A gait bitstring consists of the concatenation of the encoded starting position and the subsequent encoded joint angle differences.

To evaluate the gait, the trajectory was interpolated to simulate the sampling frequency of the gait bounce sensor assuming a constant joint velocity. On the robot, trajectories are similarly interpolated so that the PD-controller maintains a constant velocity.

B. Objective Functions

The objective function is the guiding force of the evolutionary process in GAs. Consider the objective function $f(b)$. The algorithm finds those individuals (represented by the bitstring b) that are *most fit*, meaning those in the population that minimizes $f(b)$. When selecting parents for the next generation, the algorithm is most likely to select those individuals who are more fit.

In general terms for our gait evolution problem of selectivity, the objective is to create a gait that maximizes the differences of the gait bounce signatures (in frequency space) resulting from different terrains. In other words, find a gait that when used to traverse a robot over two different types of terrain, produces two gait bounce signatures that are very far from each other in frequency space.

To evaluate a population using this objective, we created a model of the limb/terrain interaction as the robot traverses a specific terrain using a specific gait. This model estimates the resulting gait bounce signature. At each generation, for each gait (individual) and each terrain, a gait bounce signature is produced with this model and compared in frequency space. The text below describes the model and formally presents the objective function.

1) Limb/Terrain Interaction Model: Bekker[15] did seminal work in modeling wheel/terrain interactions in the 1960's. He derived an equation that determined how much a wheel would sink into the soil given the width (b) and the diameter (D) of the wheel, the load on the wheel (W), and terrain-specific parameters. These parameters are k_c: the coefficient of cohesion, k_φ: the coefficient of friction, and an exponent n. The equation is as follows:

$$s_\tau = \frac{3W}{(3-n)(k_c + bk_\varphi)\sqrt{D}}^{\frac{2}{2n+1}} \tag{1}$$

This specific equation is for a wheel towing a weight across soil. Obviously, TerminatorBot is not a wheeled vehicle, but the fingertip is a cylindrical surface that functions like a wheel and the arms "tow" the weight of the robot, as shown in Fig. 5. For our purposes, we defined wheel width and diameter in terms of the fingertip and calculated load for a static robot. We will continue to refine this model to improve sinkage estimation.

Load (W) on each fingertip was calculated using the following equation:

$$W(\theta) = -\frac{1}{2} - M_b - 2M_a + \frac{M_b D_{cog} + 2M_a D_a}{D_{tip}} \tag{2}$$

where M_b and M_a are the mass of the body and arm, respectively; D_{cog}, D_a, and D_{tip} are the distances along the y-axis from the point of ground contact at the back of the robot to the center of gravity of the body, the center of gravity of the arm, and the fingertip, respectively.

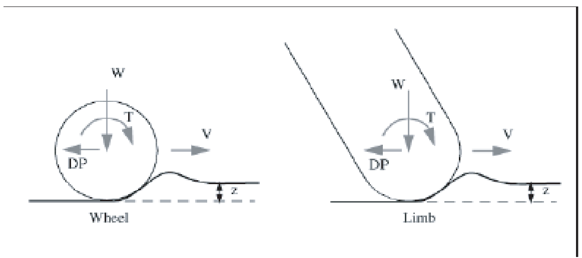

Fig. 5. *Comparison of wheel/terrain and fingertip/terrain interaction.*

To determine the gait bounce for a given terrain, first consider gait bounce along a noncompliant, smooth surface:

$$B_{nc} = atan\frac{-r-z}{l+y} \tag{3}$$

where r and l are the radius and length of the robot body; and z and y are the coordinates of the end-effector in cartesian space. After applying the above interaction model and deriving a sinkage estimate s_τ, the Eq. 3 is modified as follows:

$$B = B_{nc} - \frac{s_\tau}{D_{tip}} \tag{4}$$

The above equations define the process of estimating the gait bounce signature. This is calculated for each gait-terrain pair. These values are used in the objective functions for the genetic algorithm as described below.

2) Objective Function for Discriminability: The objective function for discriminability of gait bounce signals takes as input the via points of a gait $\vec{\Theta}$ and produces a distance measure between terrain types using the limb/terrain interaction model. The function is as follows:

$$f_{\Delta B}(\vec{\Theta}) = \frac{\sum_{i=1}^{n-1}(F(f_m(\vec{\Theta},\tau_i)) - F(f_m(\vec{\Theta},\tau_{i-1})))^2}{n-1} \tag{5}$$

where n is the number of terrains, F is the frequency transform function, f_m is the limb/terrain interaction model (i.e. Eq. 4), and the difference is the euclidean distance between vectors. Note that Θ is a matrix of via points that are each a 3-element joint vector, f_m steps through each of the rows of Θ and produces a gait bounce vector. F maps this into a frequency vector. Finally, the average squared euclidian distance is calculated to produce a single value that specifies the fitness of the individual.

There are benefits to simultaneously optimizing multiple objective functions. Two other considerations for this problem are distance travelled and energy consumed during the gait cycle. Both the maximization of distance and minimization of the absolute distance traveled (i.e. optimizing to 0) are useful in different situations. If the terrain may be hazardous to the robot, it would be best to probe the terrain to identify its composition prior to moving. By optimizing to 0, the GA should result in a probing gait. On the other hand, if there is no present danger, it would be more efficient to locomote the robot while classifying terrain. Energy can be particularly critical to field robots, as battery power is often an issue.

To search for a locomoting gait, the evolutionary process maximized the following objective function:

$$f_{\Delta d}(\vec{\Theta}) = \sum_{i=1}^{n-1}(\alpha_i\alpha_{i-1}(-f_k(\vec{\theta}_{i-1})_y + f_k(\vec{\theta}_i)_y) \tag{6}$$

where n is the number of via points, and $f_k(\vec{\theta})_y$ is the y coordinate of the via i after the kinematic function f_k is applied. f_k takes a joint vector and produces an (xyz) coordinate that locates the end-effector of the arm (i.e. the fingertip) in cartesian space. The xyz-origin is in the front center of the robot. Note that distance is measured as a function of the motion of the arms along the axis parallel to the ground (i.e. the y-axis). When the arm is in contact with the ground, effectively the arms stay stationary and the rotation of the shoulder propels the body forward, thus the need for α, which is a "ground contact" boolean variable that equals 1 when the z coordinate falls below the contact point of -37.5. This is necessary because the robot can only locomote when the arms are in contact with the ground.

Similar to Eq. 6, the GA attempted to minimize the following objective function to find a probing gait:

$$f_{d0}(\vec{\Theta}) = \sum_{i=1}^{n-1}(\alpha_i\alpha_{i-1})(-f_k(\vec{\theta}_{i-1})_y + f_k(\vec{\theta}_i)_y)^2 \tag{7}$$

The y-axis runs parallel to the ground, thus changes in this axis (while the robot is in ground contact) signifies forward motion.

Finally, we define the energy function, which is the amount of energy used per radian rotation of the joint (Joule/rad). The robot must apply more torque, thus use more energy, when in contact with the ground. The e_n and e_c signify non-contact and contact energy constants.

$$f_e(\vec{\Theta}) = \sum_{i=1}^{n-1}(e_n(\vec{\theta}_i - \vec{\theta}_{i-1})_n + e_c(\vec{\theta}_i - \vec{\theta}_{i-1})_c) \tag{8}$$

where subscripts n and c signify non-contact and contact for both the energy constants e_n and e_c and rotational changes $\vec{\theta}_i - \vec{\theta}_{i-1}$.

C. Culling Agents

One drawback to genetic algorithms is the time it takes to explore the search space. This process may be sped up by culling those individuals, or parts of the individual, that present an impossibility in the real world. The via points of the gait represent a configuration of the arm, and although the joint value may be in a viable range, the combination of the joints may cause collisions.

Three main collisions were considered in this work, including a left and right arm collision, a point of an arm inside the body, and an arm colliding with itself. Three culling agents were created to test for these conditions. Each agent checked each via of each individual, culling those *vias* that caused a collision prior to analyzing fitness.

The “arm-arm” culling agent checked if either the fingertip or the point of the elbow joint crossed over the center axis of the body (i.e. the x-axis, which runs perpendicular to the ground). Because the arms are symmetric, any crossover of this axis will result in a collision.

The “in-body” culling agent checked for penetration of the body by either the end-effector or the point of the elbow joint. This was established by defining a circle whose radius is equivalent to the orthogonal distance from the center of the robot to the point of interest. If this radius is less than the radius of the robot body, the arm is penetrating. For true collision detection, it must be verified the line segment from the elbow to the fingertip (i.e. the forearm) does not penetrate the body. This will be incorporated in future versions. We need not test for the upperarm, as it is intended to “penetrate” the body, which is appropriately hollowed out around the upperarms.

The “self-collision” agent checked for a joint crossing through $\pi/2$, which is the point at which the elbow is at full flexion. (Imagine folding your elbow to touch your shoulder and then attempting to continue rotation of the forearm through your upper arm.)

V. Experimental Results

We have run some preliminary experiments in simulation to verify that the gait evolution has been properly framed. Some results are shown in Figure 6, where a comparison was made between gaits that were evolved for efficiency (circles), discriminability and distance (triangles), and discriminability only (stars).

Using the terrain model, each of the evolved gaits were run across 6 types of soils ranging from snow to loam. These soil types were distinguished by the coefficients of cohesion and friction as used in Bekker’s equations. From each of these runs, the gait bounce signature was derived then mapped into frequency space and plotted for a visual comparison. The plot is only in 3 dimensions, but the complete vector is $1x128$.

Good selection is based on increased spacing between signatures from similar corresponding soils. The figure shows that the gait evolved for discriminability (stars) shows a greater distinction across the 6 soil types. It should be noted that the gait evolved for discriminability only produced a *probing* gait, meaning it showed no forward progress. Gaits evolved for distance travelled approximately 130 cm per gait.

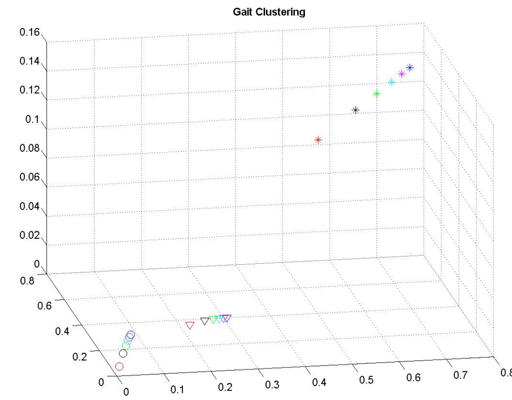

Fig. 6. *Gait bounce signatures plotted in frequency space derived from gaits evolved for efficiency (circles), discriminability and distance (triangles), and discriminability only (stars).*

VI. CONCLUSIONS AND FUTURE WORKS

Actively sensing the terrain allows self-adaptation of a robot’s gait to improve locomotion efficiency. The gait bounce approach to terrain classification has been shown to work for a small class of terrains. Increasing the terrain class will make recognition even more difficult, therefore, in this work we investigated gait evolution for increased discriminability of gait bounce signals used for terrain classification.

We have established a complete framework for evolution including encoding of the gaits, a limb/terrain interaction model that estimates gait bounce, and an objective function that can be optimized for discriminability. We have also presented objective functions for developing both locomoting and probing gaits. Finally, we identified culling agents that help guide the evolution process away from real-world impossibilities.

While the framework for evolution has been verified and preliminary results look promising, further testing is needed to confirm and quantify the benefits of evolving gaits for discriminability. Additionally, simulation is necessary for the GA approach and can be very useful for general development of the robot. However, simulation results must be verified on the physical robot in the real-world. In extending this work, the primary focus will be in improving the limb/terrain interaction model for soils and developing other non-soil models such as rocks and rubble.

References

[1] A. Larson, R. Voyles, and G. Demir, “Terrain classification through weakly-structured vehicle/terrain interaction,” in *Proc. of the IEEE Int’l Conference on Robotics and Automation*, 2004, pp. 218–224.

[2] A. Larson and R. Voyles, “Terrain classification using weakly-structured vehicle/terrain interaction,” *Autonomous Robots*, vol. 19, pp. 41–52, 2005.

[3] S. Hirose, “A study of design and control of a quadruped walking vehicle,” *Int’l Journal of Robotics Research*, vol. 3, pp. 113–133, 1984.

[4] K. S. Espenschied, R. D. Quinn, R. D. Beer, and H. J. Chiel, “Biologically based distributed control and local reflexes improve rough terrain locomotion in a hexapod robot,” *Robotics and Autonomous Systems*, vol. 18, no. 1–2, pp. 59–64, 1996.

[5] M. Lewis and G. Bekey, “Gait adaptation in a quadruped robot,” *Autonomous Robots*, vol. 12, pp. 301–312, 2002.

[6] R. Kurazume, K. Yoneda, and S. Hirose, “Feedforward and feedback dynamic trot gait control for quadruped walking vehicle,” *Autonomous Robots*, vol. 12, pp. 157–172, 2002.

[7] J. Cham, J. Karpick, and M. Cutkosky, “Stride period adaptation for a biomimetic running hexapod,” *Int’l Journal of Robotics Research*, vol. 23, no. 2, pp. 141–153, 2004.

[8] K. Iagnemma, H. Shibly, and S. Dubowsky, “On-line terrain parameter estimation for planetary rovers,” in *Proc. of the IEEE Int’l Conference on Robotics and Automation*, vol. 3, 2002, pp. 3142–3147.

[9] C. Brooks and K. Iagnemma, “Vibration-based terrain classification for planetary rovers,” *IEEE Transactions on Robotics*, vol. 21, no. 6, pp. 1185–1191, 2005.

[10] A. Talukder, R. Manduchi, R. Castano, K. Owens, L. Matthies, A. Castano, and R. Hogg, “Autonomous terrain characterisation and modelling for dynamic control of unmanned vehicles,” in *Proc.*

IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems, 2002, pp. 708–713.

[11] G. Hornby, S. Takamura, J. Yokono, O. Hanagata, T. Yamamoto, and M. Fujita, "Evolving robust gaits with AIBO," in *Proc. of the IEEE Int'l Conference on Robotics and Automation*, 2000, pp. 3040–3045.

[12] G. HOrnby, S. Takamura, T. Yamamoto, and M. Fujita, "Autonomous evolution of dynamic gaits with two quadruped robots," *IEEE Transactions on Robotics and Automation*, vol. 21, no. 3, pp. 402–410, 2005.

[13] B. Luk, S. Galt, and S. Chen, "Using gaenetic algorithms to establish efficient walking gaits for an eight-legged robot," *Int'l Journal of Systems Science*, vol. 32, no. 6, pp. 703–713, 2003.

[14] R. Jha, B. Singh, and D. Pratihar, "On-line stable gait generation of a two-legged robot using a genetic-fuzzy system," *Robotics and Autonomous Systems*, vol. 53, pp. 15–35, 2005.

[15] G. Bekker, *Introduction to Terrain-Vehicle Systems*. University of Michigan Press, 1969.