Active Perception of Material and Shape by a Walking Robot

Eric Krotkov
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract—**A legged robot can use its legs for more than locomotion. It can use them to identify the material properties and shape of the terrain it traverses. In this paper, we formulate active perception techniques based on force/torque sensing to identify terrain stiffness and surface friction, and present first results in applying the techniques to characterize terrain traversability. We also describe purposeful contact sensing to update geometric terrain maps in a process that might be called "shape from kicking."**

1 Introduction

A legged robot can use its legs for more than locomotion. It can employ them as probes to test ground compliance before stepping, and to determine surface material properties in order to travel more efficiently. The robot can also use them like a blind person's cane in order to update geometric terrain maps. Active and purposeful use of the legs makes every step an experiment, whose results inform the robot about the environment.

This paper describes how a walking robot can actively conduct perceptual experiments by stepping on and kicking things, and presents first results from initial trials. It also specifies how the results of the experiments may benefit a robot walking over rugged terrain.

One context of this research is the Planetary Rover project at CMU, whose goal is to prototype an autonomous mobile robot for planetary exploration [12]. For this project, researchers have designed and constructed a six-legged walker, called the Ambler, that features orthogonal legs and a circulating gait (Figure 1). Currently, the Ambler walks indoors through artificial obstacle courses. As the Ambler confronts terrain that is more natural, rugged, and hazardous, procedures for active perception may play an important role in autonomous operation.

We distinguish this work from related research in active perception along two dimensions: the sensing modality, and the type of information perceived.

Most active perception research concentrates on the modalities of touch and vision. Much active touch research follows from Gibson's [9] observation that when we touch something with our fingers, the skin stimulation depends on the movements that we make. Similarly, recent active vision research (e.g., [1, 2, 3, 11]) reflects Gibson's [10] view of fixation, accommodation, convergence, and pursuit as dynamic visual activities. We concentrate on a different modality: contact sensing of forces and torques. To the best of our knowledge, Gibson did not consider this modality, although it would not surprise us to learn that he did.

Figure 1: Ambler

The types of information revealed by experiments with legs and terrain include both shape and material. Most related work aims to answer the geometric question, "what does it look like?" Some related work, discussed in Section 2, attempts to answer the question, "what is it made of?" Only a few research efforts, e.g., Pentland and Williams [15], try to answer both questions. We consider material in the next section, and shape in Section 3.

2 Material Properties

Unlike most laboratory robots, a planetary rover must confront harsh and rugged terrain. Without prepared surfaces such as floors and roads, the rover risks traveling over unstable or hazardous regions where it can trigger terrain transitions that may have catastrophic consequences. For example, while climbing an inclined surface, the rover may cause a landslide. Or, like an ice skater, the rover may break through a thin surface covering a deep trough. Or, like a hiker crossing a stream, the rover may cause a supporting rock to capsize.

No known form of mobility can negotiate all such hazardous terrain. Thus, it is advantageous to determine in advance what Bekker [4] calls the ***trafficability*** of the terrain, viz., whether soil will permit the traversal of a given vehicle. This is a difficult problem, not least because it conflates geometric and material effects.

2.1 Related Research

The literature documents a wide variety of techniques to estimate material properties of terrain. Bekker describes why the research field is fertile: "Just as the weather forecast, with which the navigability of air is judged, cannot be expressed in terms of temperature or humidity alone, the trafficability of soils cannot be defined by a single unit," [4], p. 410. We will classify the techniques according to whether or not they require contact sensing.

Non-contact sensing can provide significant information about the material composition of objects. We have shown that analysis of thermal images leads, under suitable conditions, to estimates of material grain size appropriate for distinguishing between dust, sand, and rock [7]. Pentland and Williams [15] formulate estimators for material properties such as stiffness, strength, and mass from visual data (they consider contact modalities as well). The remote sensing literature cites examples of many other techniques, such as α back-scattering and impulse radar, which we do not review further.

Researchers have reported contact sensing approaches using active seismic sensors. The Soviet lunar lander Luna 13 deployed a measuring rod with a gunpowder assembly that punched the rod into the lunar soil. Lunokhod, the machine that the Soviets teleoperated on the moon in 1970, employed a "stomping device" to determine the density and mechanical strength of the lunar soil.

Other researchers concentrate on force sensing to determine material properties. Bicchi et al. [5] instrument a leg-angle-foot system and use it experimentally to assess the deformation of rubber blocks and to estimate coefficients of static friction. Sinha et al. [17] implement exploratory procedures to recover mechanical properties such as penetrability, hardness, and surface roughness using a compliant wrist mounted on a manipulator.

2.2 Approach and Apparatus

A walker can use its legs as probes to determine material properties of the surface, just as some mountain climbers investigate their surroundings with an ice axe before committing their weight, and safety, to the terrain. Our approach to identifying material properties capitalizes on the mode of locomotion rather than a distinct "material property sensor."

We aim to develop automatic procedures to identify a variety of mechanical and material properties. For the purpose of this paper, we restrict our attention to two specific properties—terrain stiffness and surface friction—and present identification procedures that are neither automatic nor complete.

Figure 2 illustrates the experimental setup. Steel beams approximately 5m above the floor support a prototype leg.[1] Three cables secured to the concrete floor immobilize the leg, preventing it from moving in the plane parallel to the floor. The leg can reach the ground by extending its prismatic joint.

The foot, located at the tip of the leg, is a circular, steel plate 30cm in diameter. Above the foot is a six-axis force-torque sensor.

2.3 Terrain Stiffness

Land locomotion requires the terrain to support the forces applied by the vehicle's mass and actuators. For legged locomotion, vertical foot-soil interactions play a significant role in analyzing support and stability. An important variable is the soil stiffness, in units of force per unit length.

A walking machine can profit in a number of ways from knowledge of terrain stiffness. For example, it can use stiffness information to avoid walking on unstable regions such as a pit covered by low-density drift material like wind-blown sand or dust (less than 1 g/cm^3 on Mars). In addition, stiffness information can improve attitude control, so that the vehicle uses less energy to keep itself level.

[1] Originally, we designed the R-R-P leg for the Ambler walking robot and built one prototype. Since then we have reconfigured the Ambler to use an R-P-P leg.

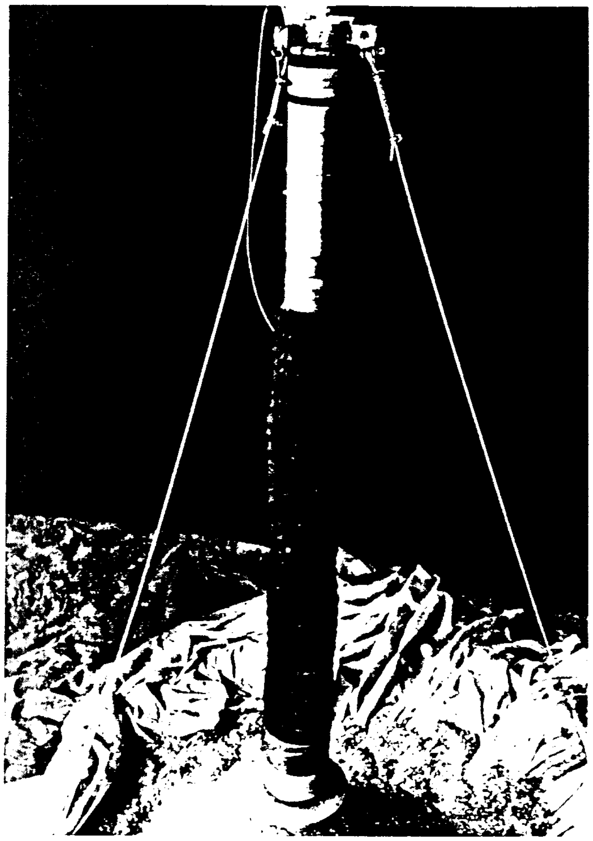

Figure 2: Experimental setup

For a plate (e.g., the foot) contacting a homogeneous, compact soil, we may model the vertical force-displacement response as

$$f = kz^n \ , \tag{1}$$

where f represents the applied normal force, z represents the vertical displacement (sinkage), and k and n are constants. When $n = 1$, (1) reduces to Hooke's law, where k is stiffness. For a plate contacting clays and loose soils, we may model the vertical force-displacement behavior by

$$f = f_0 \left(1 - e^{-z/k_0}\right) \ , \tag{2}$$

where f_0 is the asymptotic load above which the vertical force does not increase with displacement, and k_0 is a constant. Butterfield and Georgiadis [6] advance a compromise between the two:

$$f = f_1 \left[1 + k_2 z - e^{-(k_1 - k_2)z}\right] \ , \tag{3}$$

where f_1 is the intercept of the plastic response curve with the load axis, and k_1 and k_2 are slope parameters.

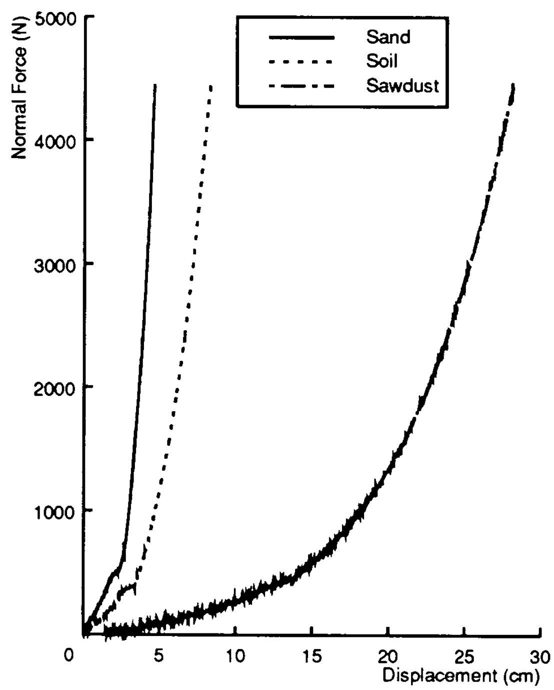

Figure 3: Terrain stiffness results
Force-displacement characteristics for sand (left curve), soil (center), and sawdust (right).

This bilinear equation explicitly represents both elastic and plastic soil responses.

With the experimental setup described above, we can measure f and z as the foot interacts with the terrain. Thus, using this model, the unknowns are f_1, k_1, and k_2. As a first step toward an estimation procedure to identify the unknowns, we conduct experiments in which we step on different materials, including sand (wet and dry), sawdust, and soil. We drive the foot into the terrain at low velocity, typically 1mm/s, in order to minimize dynamic effects such as strain rate hardening. We sample the vertical position and force vectors at a fixed rate, typically 60Hz. We continue stepping until we apply a specified load that is realistic for the Ambler, typically 4450N (1000lb).

Figure 3 plots the results, which reveal a significant qualitative difference between the curves for sand, soil, and sawdust. This difference demonstrates the feasibility of an estimation procedure to distinguish these material types. The quantitative results for sand are consistent with those that Manko [13] and Nagy [14] report.

The left part of the curves corresponds to initial compaction, and the right part of the curves corresponds to an increase in stiffness. The discontinuity between the two parts of the curves, most visible in the sand and soil data, is probably an artifact due to unequal tension in the restraining cables and to the weight of the leg. We observe that the greater is the slope of the terrain, the larger is the vertical displacement required to achieve the initial compaction. Thus, we expect that force-displacement curves for sloped terrain will exhibit more extended initial regions.

2.4 Surface Friction

Land locomotion is not possible without tractive forces, which in turn depend upon the friction characteristics of the terrain. For legged locomotion, knowledge of surface friction properties affords several advantages. For instance, it allows selection of footholds that minimize the risk of slipping while walking. Further, it allows control of propulsion that relies more on secure footholds, thus decreasing tractive effort and increasing energy efficiency.

Coulomb's equation defines the maximum lateral force $T_{\max}$ that the soil can sustain for a given normal load p as

$$T_{\max} = c + p \tan\phi \quad , \tag{4}$$

where c is soil cohesion, and ϕ is the angle of internal friction. Once the lateral force exceeds $T_{\max}$, friction no longer suffices to resist motion.

As above, we can measure p as the foot interacts with the terrain. The unknowns are $T_{\max}$, c, and ϕ. As a first step toward a parameter estimation procedure, we conduct experiments to estimate $T_{\max}$.

We use the same setup as above, but without the immobilizing cables. The material is common highway sand with shear properties of $c = 0$ and $\phi = 36.5$ degrees. We sculpt the sand into hills with slopes that vary up to the angle of repose. We step on the inclined surface with a constant normal force. Then we apply lateral forces directly away from the hill by pulling the leg with a rope, and record the force and position vectors.

Figure 4 plots the results. The figure shows that surface friction and tractive force decrease with increasing slope, and increase with normal force.

Our current work extends the above approach in three directions. First, we are replacing the rope with the Ambler's horizontal actuators, and integrating the data acquisition and analysis into the walking control system. Second, we are developing procedures to estimate c and ϕ, in addition to $T_{\max}$. Third, we are developing techniques to distinguish brittle soils from plastic soils. For brittle soils, the evolution of the friction force resembles an underdamped vibration response. For plastic soils, Freitag [8] models the evolution of the friction force as an exponential decay

$$T = T_{\max}\left(1 - e^{-j/k_3}\right) \quad , \tag{5}$$

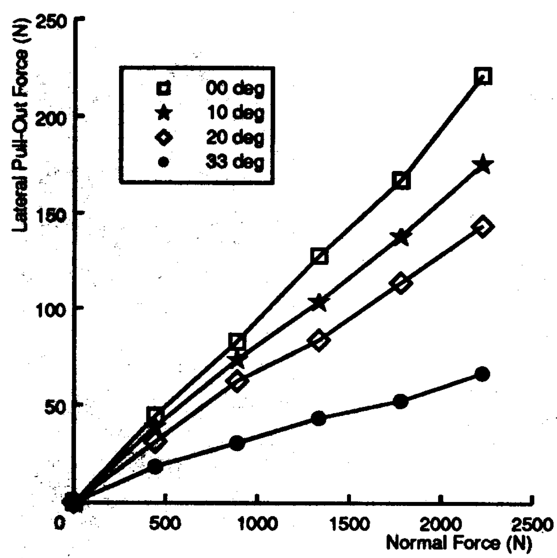

Figure 4: Surface friction results

Unpublished data from P. Nagy, 1989. The data points represent the lateral force required to exceed frictional forces at different slopes, in other words, the maximum tractive force.

where $T_{\max}$ is given by (4), j is lateral displacement, and k_3 is a constant.

3 Terrain Maps

Most rovers use imaging sensors to build terrain maps, or geometric representations of the environment. A legged robot can supplement image information with contact information. It can treat every step as an experiment, much as the blind person uses a cane to learn about the world.

Two recent research efforts define what exactly can be inferred from such a cane. In one, Salisbury analyzes the characteristics of forces and moments through point, line, and surface contacts, and presents an approach to determine geometric features from them. Extending this approach, Salisbury et al. [16] report results in determining contact forces and locations with the Whole Arm Manipulator.

In the other, Tsujimura and Yabuta [18] attach a force/torque sensor to an insensitive probe of arbitrary, but known, shape. When an object strikes the probe and exerts a force on it, the probe transmits the contact force to the sensor. The authors present a method to estimate the contact position from the single sensor, regardless of the magnitude or direction of the contact force.

Let us treat each leg as a cane. It will make one of three possible contacts while taking a step.

1. A foot contacts the terrain. Typically, this collision is desired and controlled.

2. Parts of a leg other than the foot contact the terrain. Typically, this collision is not desired, and may indicate an inadequate map, or poor control, or incorrect trajectory planning.

3. A leg contacts one of the other legs. In this self-destructive case, the machine trips over itself. Here, the problem is not directly related to the terrain representation, and to the extent that blame falls on planning or control, this case is not informative.

In any case, if a leg makes contact, we can recover the contact position $\vec{x}$ using either of the results mentioned earlier. In cases 1 and 2, we can then update the map to show that the space at $\vec{x}$ is occupied. This might be called "shape from kicking."

We do not suggest that a walking robot should rely solely on this kind of information. Walking blindly is an inferior mode of operation: it is dangerous because of the high frequency of collisions with the terrain; it reduces maximum locomotion velocity; and it provides sparse terrain data.

Instead, we suggest that a walking robot capitalize on all contact information by fusing it with other information, e.g., maps constructed by non-contact sensing. Our current research concentrates on one of the problems with this approach: how to update the map.

Let $\vec{x}_c$ represent the estimated contact position, and let Λ_c be the uncertainty in the estimate of $\vec{x}_c$ represented as a covariance matrix. Define $\vec{x}_t$ and Λ_t similarly for the estimated position of a point on the terrain. In general, we expect the uncertainty on the contact position to be substantially smaller than the uncertainty on the terrain position. One formulation of the problem is the following. Given a contact measurement $(\vec{x}_c, \Lambda_c)$ and map $\{(\vec{x}_t, \Lambda_t)\}$, find the smallest rigid motion that registers the local map region with the contact measurement.

Significant work remains to be done both in exploring formulations of the map update problem, and in developing solution techniques.

4 Discussion

Active vision encompasses issues such as where to look next, how to take good pictures, and how to constrain ill-posed visual reconstruction problems. *Active touch* addresses issues of dynamics and temporal integration in tactile sensing. *Active perception*, in our view, subsumes both active vision and active touch, and includes other sensing modalities as well.

In this paper, we have attempted to articulate a view of active perception that includes purposeful force sensing. For the purpose of determining material properties, the key activity is to engage and sustain contact in order to determine force-displacement characteristics. For the purpose of updating geometric terrain maps, the key activity is to sense contact positions (during the kick in "shape from kicking") in order to register them with existing geometric information. We have shown how legged locomotion benefits from both of these examples of active contact perception.

In this paper, the demonstration of these benefits has been mainly rhetorical. The goal of our future work is to extend this research beyond the preliminary stages reported here. For active perception of material properties, we plan to distinguish terrains experimentally based on their stiffness and roughness. For active perception of terrain geometry, we intend to explore several formulations of the map update problem and to develop stochastic solution techniques based not on the Kalman filter but on maximum likelihood estimation.

The progress toward active perception reported in this paper derives from the observation that every step is an experiment, that walking is a means both of locomotion and exploration. For active perception in the more distant future, we hope to uncover domains beyond legged locomotion that integrate task performance with learning about the environment.

Acknowledgements

We would like to acknowledge contributions by the entire CMU Planetary Rover research group. We thank J. Martin and G. Roston for their assistance with the experimental facilities, A. John for acquiring data, and D. Manko and P. Nagy for helpful discussions.

This research was sponsored by NASA Grant NAGW 1175. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of NASA or the US Government.

References

[1] J. Aloimonos. Visual Shape Computation. *Proc. IEEE*, 76(8):899–916, August 1988.

[2] R. Bajcsy. Active Perception. *Proc. IEEE*, 76(8):996–1005, August 1988.

[3] D. Ballard. Reference Frames for Animate Vision. In *Proc. Intl. Joint Conf. on Artificial Intelligence*, pages 1635–1641, Detroit, Michigan, 1989.

[4] M. Bekker. *Theory of Land Locomotion*. University of Michigan Press, 1956.

[5] A. Bicchi, L. Bosio, P. Dario, M. Guiggiani, E. Manfredi, and P. Pinotti. Leg-Ankle-Foot System for Investigating Sensor-Based Legged Locomotion. In *Proc. IEEE/RSJ Intl. Workshop Intelligent Robots and Systems*, pages 634–638, Tsukuba, Japan, September 1989.

[6] R. Butterfield and M. Georgiadis. Cyclic Plate Bearing Tests. *Journal of Terramechanics*, 17(3):149–160, 1980.

[7] C. Caillas. Thermal Imaging for Robotic Applications in Outdoor Scenes. Technical Report CMU-RI-TR-90-8, Robotics Institute, Carnegie Mellon University, April 1990.

[8] D. Frietag. Soil Dynamics as Related to Traction and Transport Systems. In *Proc. Intl. Conf. on Soil Dynamics*, pages 605–629, Auburn, Alabama, 1985.

[9] J. Gibson. Observations on Active Touch. *Psychological Review*, 69(6):477–491, November 1982.

[10] J. J. Gibson. *The Perception of the Visual World*. Houghton-Mifflin, Boston, 1950.

[11] E. Krotkov. *Active Computer Vision by Cooperative Focus and Stereo*. Springer-Verlag, New York, 1989.

[12] E. Krotkov, J. Bares, and R. Simmons. Ambler: A Six-Legged Planetary Rover. In *Proc. Intl. Conf. Advanced Robotics*, Pisa, Italy, June 1991.

[13] D. Manko. *A General Model of Legged Locomotion on Natural Terrain*. PhD thesis, Dept. of Civil Engineering, Carnegie Mellon University, April 1990.

[14] P. Nagy and W. Whittaker. Experimental Program for the CMU Mars Rover Single Leg Testbed. In *Proc. 20th Pittsburgh Conf. on Modeling and Simulation*, pages 1825–1829, Pittsburgh, Pennsylvania, May 1989.

[15] A. Pentland and J. Williams. Perception of Non-Rigid Motion: Inference of Shape, Material, and Force. In *Proc. Intl. Joint Conf. Artificial Intelligence*, pages 1565–1570, August 1989.

[16] K. Salisbury, B. Eberman, M. Leven, and W. Townsend. The Design and Control of an Experimental Whole-Arm Manipulator. In *Proc. Intl. Symposium Robotics Research*, pages 261–269, August 1989.

[17] P. Sinha, R. Bajcsy, and R. Paul. How Does a Robot Know Where to Step? Measuring the Hardness and Roughness of Surfaces. In *Proc. IEEE/RSJ Intl. Workshop Intelligent Robots and Systems*, pages 121–126, Tsuchiura, Japan, July 1990.

[18] T. Tsujimura and T. Yabuta. Object Detection by Tactile Sensing Method Employing Force/Torque Information. *IEEE Trans. on Robotics and Automation*, 5(4):444–450, August 1989.