Enhancing Fuzzy Robot Navigation Systems by Mimicking Human Visual Perception of Natural Terrain Traversability

Ayanna Howard, Edward Tunstel

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California, USA
howard/tunstel@robotics.jpl.nasa.gov

Dean Edwards, Alan Carlson

Mechanical Engineering Department
University of Idaho
Moscow, Idaho, USA
dedwards@uidaho.edu

Abstract

This paper presents a technique for learning to assess terrain traversability for outdoor mobile robot navigation using human-embedded logic and real-time perception of terrain features extracted from image data. The methodology utilizes a fuzzy logic framework and vision algorithms for analysis of the terrain. The terrain assessment and learning methodology is tested and validated with a set of real-world image data acquired by an onboard vision system.

1. Introduction

Autonomous exploration of remote planetary surfaces by outdoor mobile robots, or rovers, is an active pursuit by NASA and other national space agencies. Similar capabilities are pursued by defense agencies for unmanned autonomous vehicles. Rovers must have the ability to operate autonomously and intelligently on challenging terrain with minimal interaction with remote human operators. Rover navigation systems must provide a level of "onboard intelligence" sufficient for long-range traverses in rough, rocky, and poorly-modeled natural terrain, without exposing the rover to undue physical risk or situations that may lead to mission failure.

To ensure mission success, onboard intelligence must be able to assess a rover's ability to traverse terrain regions of varying difficulty en route to designated locations. To enable robots to make autonomous navigation decisions that guide them through the most traversable regions of the terrain, fuzzy logic techniques have been developed for classifying traversability using computer vision-based perception of attributes such as surface roughness, slope, and discontinuity [1, 2]. This paper presents a fuzzy logic system designed to automatically infer terrain traversability from images captured by a robot's onboard vision system. Based on the physical properties of the terrain extracted from the images (and intrinsic knowledge of rover mechanical constraints), the suitability of the terrain for traversal is inferred using a fuzzy logic framework and vision-based algorithms. The fuzzy classification of traversability provides essential perceptual knowledge, which is utilized by a navigation system for robust mobility through rough natural terrain. This paper focuses, in particular, on the enhancement of this capability by applying a nonlinear optimization technique that aims to adjust fuzzy system parameters to achieve perceptual performance closely resembling that of a human expert [3]. In this way, the human expert acts as a supervisor to facilitate the process of teaching a nominal fuzzy terrain classifier to mimic human perception. Sections 2 and 3 describe the terrain assessment and fuzzy classification algorithm. Section 4 describes the methodology for enhancing terrain classification and Section 5 presents the optimization results obtained by using real-world image data.

2. Linguistic Representation of Terrain Features

The first step in classifying the local terrain surrounding the robot involves extracting the terrain features directly accounting for navigation difficulty. To accomplish this task, we have developed a set of vision algorithms used to determine slope and roughness, two important attributes that characterize the difficulty of terrain for traversal by a mobile robot. In this section, we describe the vision algorithms and how the slope and roughness values produced by these algorithms are used to reason about terrain traversability. More detail on the vision algorithms can be found in [2].

2.1 Terrain Roughness Extraction

Terrain *roughness* refers to the coarseness and surface irregularity of the ground to be traversed. Visual perception of rock distribution in a viewable scene is used to determine a measure of terrain surface roughness. First, an algorithm for determining the size and concentration of rocks in a viewable scene is

applied to a pair of stereo camera images taken from the robot's vantage point. A horizon-line extraction program is run that identifies the peripheral boundary of the ground plane. This, in effect, perceives the line at which the ground and the landscaped backdrop intersect. The algorithm then identifies target objects located on the ground plane using a region-growing method [4], which consists of ground signature extraction, edge detection, and obstacle identification (large rocks/boulders and large groups of rocks). In effect, targets that differ from the ground surface are identified and counted as rocks for inclusion in the roughness assessment. The effect of this procedure is illustrated by the sequence of raw and processed images of natural terrain in Fig. 1.

Original Image   Horizon Line Extraction   Rock Detection

Figure 1. Terrain roughness determination

To determine the number of small- and large-sized rocks contained within the image, the number of pixels that comprise a target object are first counted. Those targets with a pixel count less than a user-defined threshold are labeled as belonging to the class of small rocks and those with a count above the threshold are classified as large rocks. All such labeled target objects are then grouped according to their sizes in order to determine the small and large rock concentration parameters. Next, the average separation distance between rocks in the image is calculated. This calculation provides information about the relative amount of free space available for traversal between rocks. The parameters for small and large rock concentration and average separation distance are respectively characterized by fuzzy sets with linguistic labels {FEW, MANY} and {CLOSE, FAR} (see [2] for details). The terrain roughness (normalized) is then represented by three fuzzy sets with linguistic labels {SMOOTH, ROUGH, ROCKY}, defined by membership functions shown in Fig. 2. Thus, roughness is derived from the rock size/concentration and separation parameters of the associated image using the fuzzy logic rules summarized in Table 1 (empty cells ⇒ no effect on the rule).

Figure 2. Membership functions for terrain roughness

Small Rocks	Large Rocks	Separation Distance	Roughness
FEW	FEW	—	SMOOTH
MANY	FEW	—	ROUGH
—	MANY	FAR	ROUGH
—	MANY	CLOSE	ROCKY

Table 1. Rule-base for determining roughness

2.2 Terrain Slope Extraction

Terrain *slope* refers to the average incline/decline of the ground surface to be traversed. To obtain the terrain slope from a pair of stereo camera images, we must first calculate the real-world Cartesian (x,y,z) components of the ground plane boundary, or horizon-line. To determine the (x,y,z) components of the horizon-line, Tsai's camera calibration model [5] is used to derive the relationship between the camera image and the real-world object position for a single camera. The images from both cameras are then matched in order to retrieve 3D information.

Determining the position of the largest rocks located along the horizon-line and centered within both images allows the identification and extraction of correlated image points that lie along the horizon-line. These image points are used as input for extraction of the (x,y,z) real-world Cartesian components. Depending on the viewable scene in the pair of stereo images, there may be multiple pairs of such correlated image points. Once all Cartesian points are calculated, they are used to compute the average terrain slope. The terrain slope (normalized) is represented by three fuzzy sets with linguistic labels {FLAT, SLOPED, STEEP}, defined by membership functions shown in Fig. 3.

Figure 3. Membership functions for terrain slope

Figure 4. Membership functions for traversability

3. Fuzzy Classification of Terrain Traversability

As a rover traverses natural terrain, images are periodically acquired for traversability assessment. In each period, terrain in the viewable scene must be classified once the relevant features are extracted. To this end, we have developed a set of fuzzy logic rules that classify the traversability of the terrain based on the roughness and slope measures extracted from the given image data set.

The process of embedding human expert knowledge of terrain traversability in a fuzzy logic system begins when the knowledge base (rule base and membership functions) is developed. The fuzzy sets defined above for roughness and slope allow these terrain characteristics to be represented based on grades of membership, as opposed to a 0 or 1 value. These linguistic variables are used as inputs to a set of fuzzy rules used to classify the terrain. The output from the rule base is a *traversability index*, which represents the relative terrain quality as it relates to the rover's ability to safely traverse terrain in the viewable area. Given imagery of the terrain in the vicinity of the rover, a traversability index τ is defined as a function of the terrain slope and the terrain roughness determined as described above [1]. It is represented by three fuzzy sets with linguistic labels {LOW, MEDIUM, HIGH} defined by the membership functions shown in Fig. 4. Thus, a region of low traversability is unsafe or very difficult for the rover to traverse, while a highly traversable region is safe and relatively easy to traverse. The nominal set of fuzzy logic rules used to infer terrain traversability based on roughness and slope is summarized in Table 2. If the local terrain surrounding the rover is partitioned into adjacent sectors (e.g., front, left and right) the traversability of each sector can be determined using the corresponding imagery.

	SMOOTH	ROUGH	ROCKY
FLAT	HIGH	MEDIUM	LOW
SLOPED	HIGH	MEDIUM	LOW
STEEP	LOW	LOW	LOW

Table 2. Terrain Traversability Rule-base

Initially, the fuzzy membership functions and rules used to determine terrain traversability are defined according to an expert's subjective perception and intuition. Using this subjective representation of the expert's knowledge encoded in the knowledge base, the resulting fuzzy system achieves a terrain assessment behavior that roughly approximates that of the human expert. In order to achieve a better approximation of the expert's behavior, and thereby more closely mimic human expert reasoning, the fuzzy system must be fine-tuned by either modifying the membership functions, the rule-base, or both. In the next section we focus our attention on the input membership functions, for terrain roughness and slope, nominally defined as shown in Figs. 2 and 3.

4. Mimicking Human Classification of Terrain

The ultimate goal for a fuzzy classification system used in practice is to closely mimic the human expert's judgement of the terrain traversability. In this way, mission operators can be reasonably confident that decisions made by a navigation system, that autonomously guides a rover at a remote location, are sound enough to preserve rover safety and ensure mission completion. To achieve such confidence, the system is trained using the expert as a supervisor.

In this work, the objective is to improve the nominal fuzzy system by optimal tuning with respect to a human expert's perceptual classification of terrain, based on images captured by a robot vision system.

There are several parameter search and optimization techniques that are suitable for this problem, including genetic algorithms, artificial neural networks and more conventional optimization algorithms [6]. We employ a conventional optimization algorithm, called the simplex method, to improve the fuzzy system's ability to mimic human expert reasoning by optimizing selected membership function parameters that define terrain roughness and slope. In particular, we apply the Design Optimization Procedure (DOP), which was introduced in [3] and later applied in [7] to design and optimize fuzzy logic controllers. The DOP is based on quantitative performance comparisons between a human expert and a fuzzy system addressing a common problem. It employs simplex optimization to achieve iterative improvement of the fuzzy system toward human performance. The application of the DOP to the traversability assessment problem requires determination of a terrain classification strategy and formulation of a suitable performance index (PI) that indicates the error between the fuzzy system classification and the human expert classification. A suitable terrain classification strategy was presented in section 3; we will now describe the optimization procedure for improving the nominal system and the formulation of a PI to drive that procedure.

4.1 Optimization Procedure

The fuzzy system classification is determined using the nominal (pre-optimized) set of membership functions and fuzzy rules defined above, while the human classification is provided by an expert. Performance improvements in the fuzzy system are obtained in a two-phase process described below. Views of the terrain from the rover's vantage point are presented to the expert, first in the form of raw black-and-white camera images and later as processed camera images (see left and right images of Fig. 1). The fuzzy system always classifies traversability based on quantitative values of roughness and slope that are extracted from each processed camera image (i.e. a reduced data set). Variations between these distinct classifications are minimized by the optimization technique.

In the first phase, the expert classifies the terrain traversability by viewing raw pre-processed images and assessing the ability of the robot to traverse the terrain in each image. A quantitative classification scale from zero (low traversability) to ten (high traversability) is used for the terrain traversability index. A value in this range is furnished by the expert,

and also inferred by the fuzzy system, for each image in a training set containing images of varied terrain. Traversability indices provided by the expert are stored in a database. The absolute error between the fuzzy system's classification (FC) and the human classification (HC) of each image is used to compute the PI as an integral squared error. In the second phase of the process, an attempt is made to further improve the performance of the fuzzy system by allowing the expert to classify traversability using processed terrain images (such as the right-most image of Fig. 1), which represent a reduced data set similar to that actually processed by the fuzzy system. The optimization procedure is applied again for this case to yield a new level of classification performance quantified by the same PI. The basic flow of information is illustrated in the block diagram of Fig. 5.
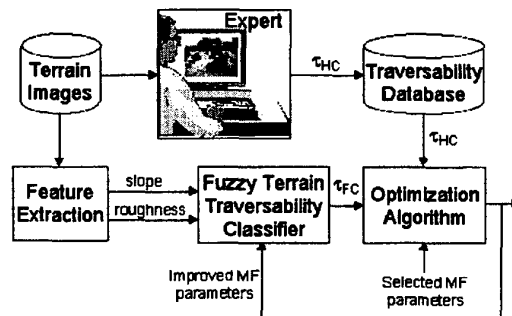

Figure 5. Procedural data flow for optimization.

In either phase, when the fuzzy system's classification is close to the human classification the PI approaches zero; when the error is large, the PI is high. The PI is an important metric in the design approach since it drives the optimization of membership functions such that the inferences made by a fuzzy system closely match inferences made by the human expert. Thus, the problem is similar in structure to a supervised learning problem wherein the optimization technique is used to adjust numerical parameters that define the membership functions for roughness and slope, such that the fuzzy system is encouraged to mimic the expert's classification. Four fuzzy set parameters per input membership function are considered in the optimization. For a trapezoidal fuzzy set, these include the values in the universe of discourse that bound its support and core/nucleus [8]. The parameters continue to be adjusted (thus refining the meaning of the input linguistic variables) until the PI is minimized or reaches a steady and non-decreasing value. This signifies convergence to a tuned set of membership

functions that result in a classification behavior that more closely mimics that of the human expert. After the system is trained, it is used for real-time terrain classification in outdoor rover navigation.

The classification done by the expert in the second phase is later compared with the phase one classification and the classification performed by the fuzzy system. The comparisons between these classifications are valuable in determining if the fuzzy logic system has access to the information required for a correct classification (i.e. an "observability" issue) and if the correct logic has been used. In general, if an expert cannot provide a correct assessment with only the information being presented to the fuzzy logic system, then it becomes very difficult to train the system. Under these circumstances, the information being presented to the fuzzy logic system is not adequate to make a good decision. Additional information from the original image is required. If the expert can determine what information he/she is using to make a correct assessment, then this information must also be processed and presented to the fuzzy logic system. Also, if the expert uses any additional logic with this information, then that logic must also be included in the fuzzy logic system.

5. Results

In this section, we present results of applying the DOP to the nominal fuzzy terrain classifier presented in section 3. The test vehicle is a commercially available mobile robot called the Pioneer-AT (All Terrain). It is outfitted with a custom onboard vision system comprised of three stereo pairs of commercially available CCD cameras. Cameras are mounted on a raised platform (Fig. 6), and oriented for a combined field of view of 180° in front of the rover. A set of 17 images of natural terrain scenes was used as a training set. Two representative images and their processed versions are shown in Fig. 7 — traversability is classified by the expert as HIGH for the top image and LOW for the bottom image.

Optimized membership functions for roughness and slope for the first phase (based on classification of raw images presented to the expert) are shown in Fig. 8. The associated history of the PI during the optimization/learning process is shown in Fig. 9. Results for the second DOP phase (based on classification of processed images by the expert) are shown in the Figs. 10 and 11.


Figure 6. Outdoor mobile robot with vision system


Figure 7. Raw and processed terrain images

Figure 8. Optimized membership functions: phase 1

Figure 9. Optimization performance history: phase 1

Figure 10. Optimized membership functions: phase 2

Figure 11. Optimization performance history: phase 2

These results reveal that the terrain assessment rules are most sensitive to roughness. This is apparent after observing that the optimized membership functions for slope (Figs. 8 and 10) are nearly identical to the nominal set in Fig. 3. Optimized roughness membership functions have been tuned significantly relative to the nominal set in Fig. 2, yielding relative performance improvements of 15% and 16% in phases 1 and 2, respectively. Consideration of a nearly common set of input information in phase 2 by the expert and fuzzy system yielded the better performance enhancement as expected. As can be seen in Figs. 9 and 11, phase 2 begins with a better PI (i.e., closer agreement between expert and fuzzy system) and converges in lesser iterations than phase 1. The richer the set of input data, the more apparent the disparity is between human and robot perception. To overcome the disparity and enhance robot perception through supervised learning it may be necessary to reduce the observability of the human supervisor.

6. Conclusions

This paper describes a methodology for enhancing a fuzzy system's capacity for mimicking human perception. Results of its application to the terrain classification problem for computer vision-based robot navigation are presented. The paper presents the context in which fuzzy logic techniques are applied by explaining the approach to extracting terrain attributes from camera images, using these attributes as fuzzy inputs to infer a traversability index, and acting on this perceptual knowledge to facilitate intelligent autonomous navigation. Results revealing the extent of the classification performance improvements achieved using the approach are discussed. The underlying optimization approach can be generally applied to fuzzy systems of both the Mamdani and Takagi-Sugeno-Kang type [8]. The general description of the methodology and its practical application permits one to assess its utility for enhancing fuzzy system performance in other domains of interest.

7. Acknowledgment

8. References

[1] H. Seraji, "Fuzzy Traversability Index: A new concept for terrain-based navigation," *Journal of Robotic Syst.*, Vol. 17, No. 2, 2000, pp. 75-91.

[2] A. Howard, H. Seraji and E. Tunstel, "An Enhanced Traversability Index for Mobile Robot Navigation," *IEEE Intl. Conf. on Robotics and Automation*, Seoul, Korea, May, 2001.

[3] A. Athalye, D. Edwards, V.S. Manoranjan and A. de Sam Lazaro, "On Designing a Fuzzy Control System Using an Optimization Algorithm," *Fuzzy Sets & Syst.*, Vol. 56, 1993, pp. 281-290.

[4] B. Horn, *Robot Vision*, MIT Press, MA, 1986.

[5] R.Y. Tsai, "A Versatile Camera Calibration technique for high-accuracy 3D machine vision Metrology using Off-the-shelf TV Cameras and Lenses", *IEEE Journal of Robotics and Automation*, Vol. 3, No. 4, 1987, pp. 323-344.

[6] Z. Michalewicz and D.B. Fogel, *How to Solve It: Modern Heuristics*, Springer, New York, 1999.

[7] J.R. Canning and D.B. Edwards, "A Method for Embedding Human Expert Knowledge into a Fuzzy Logic Controller," *ASME 15th Intl. Computers in Engineering Conf.*, Boston, MA, USA, September 1995, pp. 1019-1023.

[8] D. Driankov, H. Hellendoorn and M. Reinfrank, *An Introduction to Fuzzy Control*, Springer-Verlag, Berlin, 1993.