2008 IEEE International Conference on
Robotics and Automation
Pasadena, CA, USA, May 19-23, 2008

A Single Camera Terrain Slope Estimation Technique for Natural Arctic Environments

Stephen Williams, Ayanna M. Howard
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332
swilliams8@gatech.edu, ayanna.howard@ece.gatech.edu

Abstract— **Arctic regions present one of the harshest environments on earth for people or mobile robots, yet many important scientific studies, particularly those involving climate change, require measurements from these areas. For the successful deployment of mobile sensors in the arctic, a reliable, fault tolerant, low-cost method of navigating must be developed. One aspect of an autonomous navigation system must be an assessment of the local terrain, including the slope of nearby regions. Presented here is a method of estimating the slope of the terrain in the robot's coordinate frame using only a single camera, which has been applied to both simulated arctic terrain and real images. The slope estimates are then converted into the global coordinate frame using information from a roll sensor, used as an input to a fuzzy logic navigation scheme, and tested in a simulated arctic environment.**

I. INTRODUCTION

Recently, it has been discovered that the giant ice sheets covering Greenland and Antarctica have been shrinking at an accelerated rate [6]. While it is believed that these regions hold important information related to global climate change, there is still insufficient data available to accurately predict the future behavior of these ice sheets. Satellites have been able to map the ice sheet elevations with increasing accuracy, but data about general weather conditions (i.e. wind speed, barometric pressure, etc.) must be measured at the surface. The Greenland Climate Network [14] is a set of fixed weather stations scattered throughout Greenland responsible for collecting this kind of data automatically. However, there are only 18 such stations on an ice sheet measuring over 650,000 sq mi. An analogous network exists in Antarctica with equally sparse coverage [13].

In order to obtain a denser set of measurements, human expeditions must be sent to these remote and dangerous areas. Alternatively, a group of autonomous robotic rovers could be deployed to these same locations, mitigating the cost, effort, and danger of human presence. For this to be a viable solution, a method for navigating arctic terrain must be developed. This paper presents a method of estimating the slope of the upcoming terrain, with an emphasis on the use of simple, inexpensive sensors, which may allow increased autonomy of future arctic rovers. Section II describes various elements that could be encountered in the arctic regions of Greenland or Antarctica, while Section III outlines the current state-of-the-art in autonomous arctic vehicles. Section IV describes a method of estimating the slope of the terrain in a local coordinate frame from a single image, providing examples from an analogous terrain from a glacial region of Colorado. Section V discusses a basic navigational scheme implemented with fuzzy logic control designed to keep the rover on flat ground. Section VI describes the simulation environment used to test the navigation control scheme and discusses the results of those simulations. Finally, conclusions are presented in Section VII.

II. TERRAIN DESCRIPTION

The Greenland and Antarctica land masses are almost entirely covered by huge ice sheets. These ice sheets flow over time, behaving similarly to a river, albeit in slow motion [8]. The central, inland regions are a stable, nearly featureless expanse of compressed snow and ice. A topological survey of Greenland done via satellite [3] shows that in the central region, near a science outpost called Summit Camp, the elevation changes as little as 10 meters over a 20 kilometer expanse. This can also be seen in a recent NASA ICESat survey of Greenland shown in fig. 1 [16]. However, despite its flatness, the snow surface is often sculpted by the high winds present, forming undulating features parallel to the wind direction known as sastrugi. The sastrugi can vary in height from only a few centimeters to over a meter.

As the ice flows toward the edges of Greenland and into the sea, forces build due to differential velocities of different ice sections. These forces can cause nearly vertical fractures in the ice known as crevasses. Crevasses can be as deep as 30 meters and are often covered with snow, making their detection all the more difficult. This snow, referred to as a snow bridge, can be sufficiently strong as to allow passage over the crevasse. Additionally, towards the coast of Greenland, the ice sheet is thin enough that some of the underlying mountain range is visible above the ice surface. The jagged mountain peaks protruding through the ice are called nunataks. Also, towards the outer edge of the ice sheet, the surface is no longer uniform and flat, as seen in the ICESat images in fig. 1(a) and 1(c). These regions can have significant local variability of terrain elevation.

978-1-4244-1647-9/08/$25.00 ©2008 IEEE.

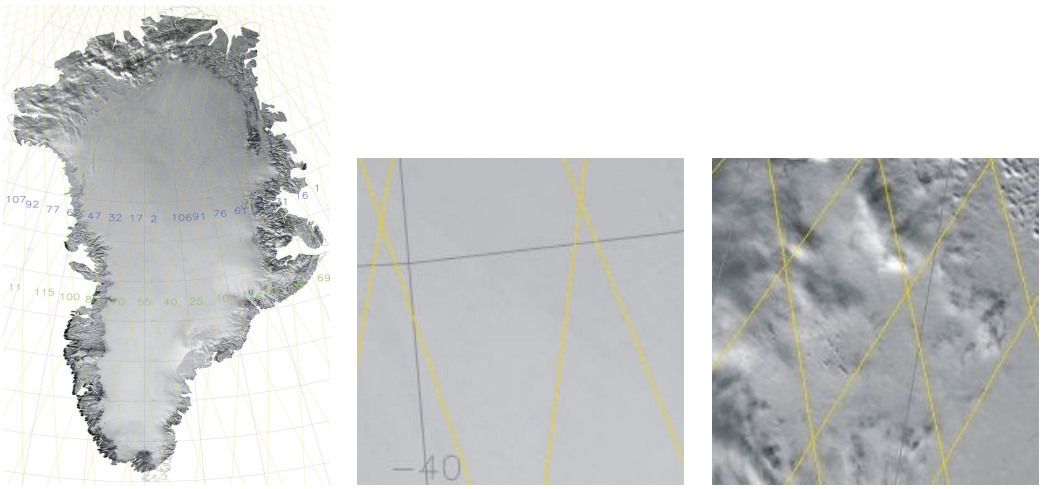

(a) (b) (c)

Fig. 1. (a) A high resolution elevation map of Greenland from ICESat [16], (b) an enlarged view of the nearly featureless area in the center of the ice sheet near Summit Camp, and (c) an enlarged view of a north-western section of the ice sheet near Qaanaaq illustrating an area of more rugged terrain.

III. BACKGROUND ON ARCTIC RESEARCH

While the arctic possesses significant information of scientific value, surprisingly little research is devoted to creating mobile, autonomous robots to collect this data. The CoolRobot [12] developed by Dartmouth College is designed to be a mobile sensor station for use on the Antarctic plateau. This region is essentially featureless, lacking any known obstacles or elevation changes that could impede the robot's motion. Consequently, the navigation system consists solely of a GPS unit and a simple waypoint following algorithm. It has been successfully tested at an analogous site in Greenland, where it operated autonomously for up to eight hours [9].

The Nomad project [2] at Carnegie Melon University was designed for autonomous meteorite location in the Antarctic moraines. The Elephant Moraine, where the Nomad was deployed, is essentially a sheet of blue ice littered with rocks ranging in size from pebbles to large boulders. Due to the search and classify nature of the Nomad's mission, it naturally requires greater sensing capability, and has been outfitted with a high resolution steerable camera, stereo cameras, a scanning laser range finder, and a pitch-roll-yaw sensor, as well as a scientific instrumentation package. Thus, the navigational system of the Nomad is more complex than that of the CoolRobot, consisting of a GPS waypoint system, an obstacle avoidance system based on the laser range finder data, a potential meteorite targeting system based on camera images, and a management system to gracefully switch between the different modes. However, due to the known flatness of the testing and deployment environment, little in way of terrain assessment is required for successful navigation. It has been deployed several times to Antarctica, during which it has autonomously found and identified 5 meteorites.

The University of Kansas is developing an autonomous snowmobile-based rover for taking ice sheet measurements as part of the PRISM project [5]. Like the Nomad, it is equipped with a camera system as well as a laser range finder for obstacle detection and avoidance. However, unlike the Nomad or CoolRobot, the use of a snowmobile chassis allows for a much wider range of navigable terrain. However, during operation, the PRISM rover will follow, or perform predefined maneuvers, behind a manned unit. Consequently, obstacle avoidance is sufficient in this case, relying on the human operator to set the lead course through acceptable regions. Similarly, the ENEA R.A.S. project [4] has the goal of creating an automated convoy of snowcats, large tracked vehicles, in which the lead vehicle is driven by a human. Again, the snowcats have the capability of traversing significant inclines and can be used in a wide range of arctic terrain, but the determination of acceptable terrain is left solely to the lead human driver.

IV. SLOPE EXTRACTION

Currently, most autonomous driving research is based upon laser range finders or stereo cameras which return quantitative depth information. In contrast, humans navigate through an enormous variety of cluttered environments without explicitly using any such depth data. Instead, we rely on qualitative information, such as "Those mountains are far away", or "That rock is close." Visual cues in our environment, such as relative height and position relative to the horizon, are used to help determine the nature and distance of objects in our field of view. A method of estimating the terrain slope of an arctic scene is outlined below. Similar to human perception, this method relies on visual cues from a single camera to make its estimates. All image processing was done in C++ and utilized the OpenCV libraries [1].

A. Region Extraction

The first step is to extract the region of the image to be analyzed. Such things as distant mountains, clouds, trees,

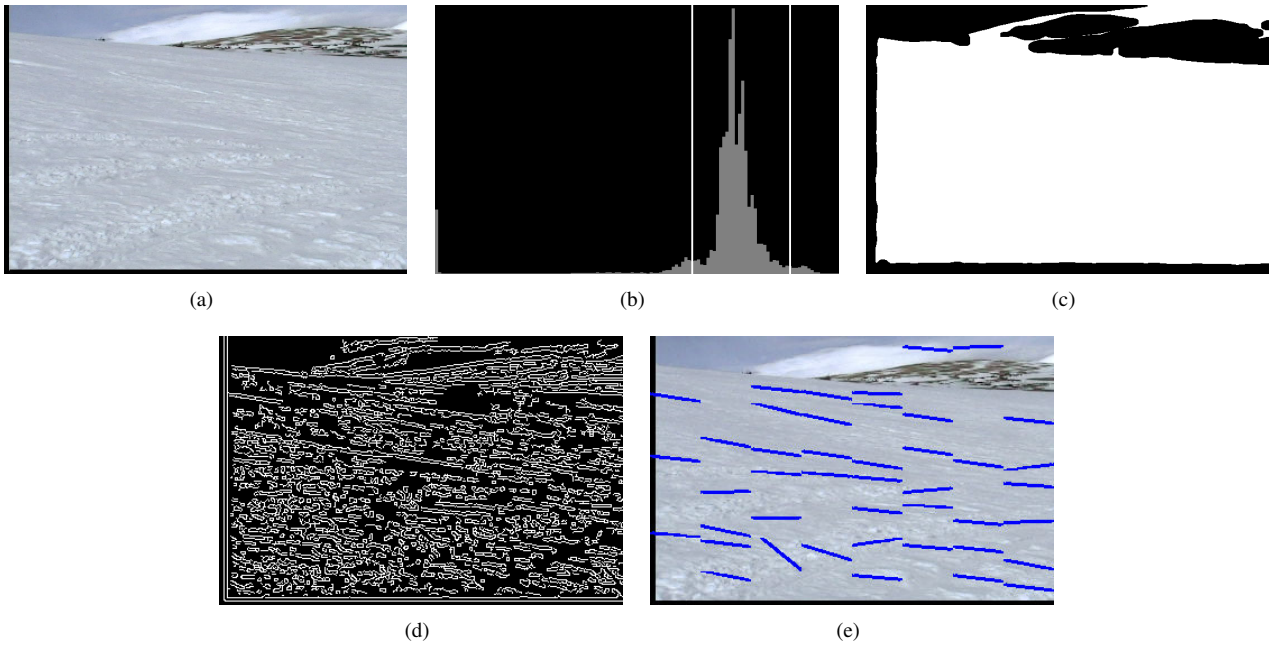

Fig. 2. (a) A sample arctic image from Colorado, (b) the histogram of the sample image with the adaptive thresholds indicated, and (c) the resulting binary mask. (d) The results of a bandpass filter and Canny edge detector on the sample image, and (e) the final slope estimates superimposed on the original image.

and peripheral views of the rover itself can interfere with the estimation process. A method of histogram thresholding, suggested in [4], has been adopted here. It is assumed that the majority of the image is filled with the snowy region. Consequently, in the histogram of the image, the largest peak should be associated with the grayscale values of this region. Thus, an adaptive threshold based on the boundaries of this peak can effectively separate the region of interest from unwanted objects and areas. To locate the peak boundaries, the grayscale image histogram is first computed and the maximum value located. Then, the histogram is searched in both directions for local minima. Criteria on the number of points between consecutive local minima is used to make the process more robust to noise. The two grayscale values found are used to threshold the image, creating a binary mask. The mask is then eroded by several pixels to ensure the edges of the unwanted features are concealed by the mask. Fig. 2 shows a representative image from a glacial region in Colorado, followed by its histogram with the threshold limits, and the resulting mask.

B. Texture Filtering

A human is able to quickly and accurately estimate the slope of the environment, such as that shown in fig. 2(a). From this single image, without any depth information, one can successfully estimate slope and loosely define distances from the camera (such as “close” or “far”). One of the keys in our ability to estimate the slope comes from visual cues in the form of light and dark streaks in the snow that are aligned with the perceived slope. These features are of the greatest interest in the slope estimation process.

The dark and light striations are clearly visible and considerably larger in scale than the snow texture. However, these striations are far from uniform in the region of interest. Clearly then, these features exist in a mid-band of spatial frequency, and can be effectively isolated using a simple bandpass filter. Filters based on Gaussian kernels are generally less likely to produce unwanted ringing effects. Additionally, since Gaussian kernels are linearly separable, the filter can be applied in a computationally efficient manner. Experimentation with the Gaussian standard deviation, σ, of the low pass and high pass filter components has shown that a wide range of values have acceptable performance, but the values are dependent on the image size. Using a σ value for the high pass filter of twice that of the low pass filter generally produces good results. One final step in preparing the image for slope estimation is to convert it into a binary image. This is done by applying a Canny edge detector to the filtered image. The image after these processing steps can be seen in fig. 2(d).

C. Hough Transform

In the final step, the slope of the extracted features is estimated using a Hough transform, which transforms each pixel in the image space into a sinusoidal line in the *rho-theta* parameter space. As each image pixel is transformed, the sinusoids in the parameter domain will tend to intersect if the image pixels are in a straight line. The number of sinusoids that intersect in a particular location is an indication of the strength or confidence of the line. Thus, the *rho-theta* pair corresponding to the maximum confidence values in the parameter space can be selected as being representative of that region. When the pixel is transformed into the parameter space, it loses any sense of its location in the original

image. It is therefore common for many local maxima to occur in a very small neighborhood. This results in having slope data only for a small area in the total region of interest. To overcome this problem, the image is divided into smaller subimages, and the Hough transform is applied separately to each subimage. In this way, the extracted slope can be applied to a specific area of the original image, and slope data will exist for all areas in the image. In each subarea, only one slope is desired, lending itself to the Fast Hough Transform algorithm [10] which employs integer shifts instead of floating point operations, reducing the processing time considerably. The final image with the detected slopes is shown in fig. 2(e).

D. Results

The slope extraction process was then applied to a set of images taken in a glacial region in Colorado. Fig. 3(a) shows a snowmobile in the foreground of the image. The estimation process ignored the area around the snowmobile and returned correct results for the surrounding regions, illustrating that the masking process can successfully remove unwanted objects from the analysis, particularly man-made objects. Fig, 3(d), however, illustrates the inability of this process to remove unwanted regions, such as gray clouds, which exist in the same color range as the snow, thus creating fictitious slope estimates. The estimation process does seem to be relatively immune to light intensities, as shown in fig. 3(b) and 3(c). These images show the same scene under different lighting conditions, with the second showing significantly more reflected light evidenced by the shiny appearance of the snow. Despite the different lighting conditions, the slope estimates seem stable and accurate. The masking and estimation process are even partially successful in the region of sparse grass shown in fig. 3(d). However, the slope estimates generated in the grassy areas are generally less stable than those of uninterrupted snow. In fig. 3(e) a set of snowmobile tracks is present which exhibit similar properties to the surface striations, and are consequently detected by the slope estimate. A method to detect the presence of tracks and mask them from the slope estimate would be beneficial. Finally, fig. 3(f) shows an image with a lower than normal amount of texture or prominent striations. This process is not able to produce estimates for the full image under these conditions, providing information only where a certain level of confidence is met.

V. NAVIGATION CONTROL SCHEME

Since no measurement of the terrain is being taken, as is done with a laser range finder or a stereo camera system, the output of this process is, at best, an estimate of the terrain. Thus, any navigation or control scheme based on these estimates must be able to handle the inherent noise and uncertainty present in the results. Additionally, attempting to navigate through natural terrain using only vision as an input is something that humans have developed heuristically [7]. A control scheme that could capture this inherently nonlinear heuristic knowledge would be advantageous. Fuzzy logic control provides just such features [11], as well as models an input in terms of a linguistic set, such as “near” and “far”, which parallels how humans think in terms of their environment.

The first step is to convert the slope estimates into fuzzy linguistic sets. For the slope estimates, the following five sets are used to classify each input: *Positive Steep*, *Positive Sloped*, *Flat*, *Negative Sloped*, and *Negative Steep*. Also, when the slope estimates were created, the original image was divided into an 8x8 grid, producing up to 64 slope values. In an attempt to reduce the number of inputs to a

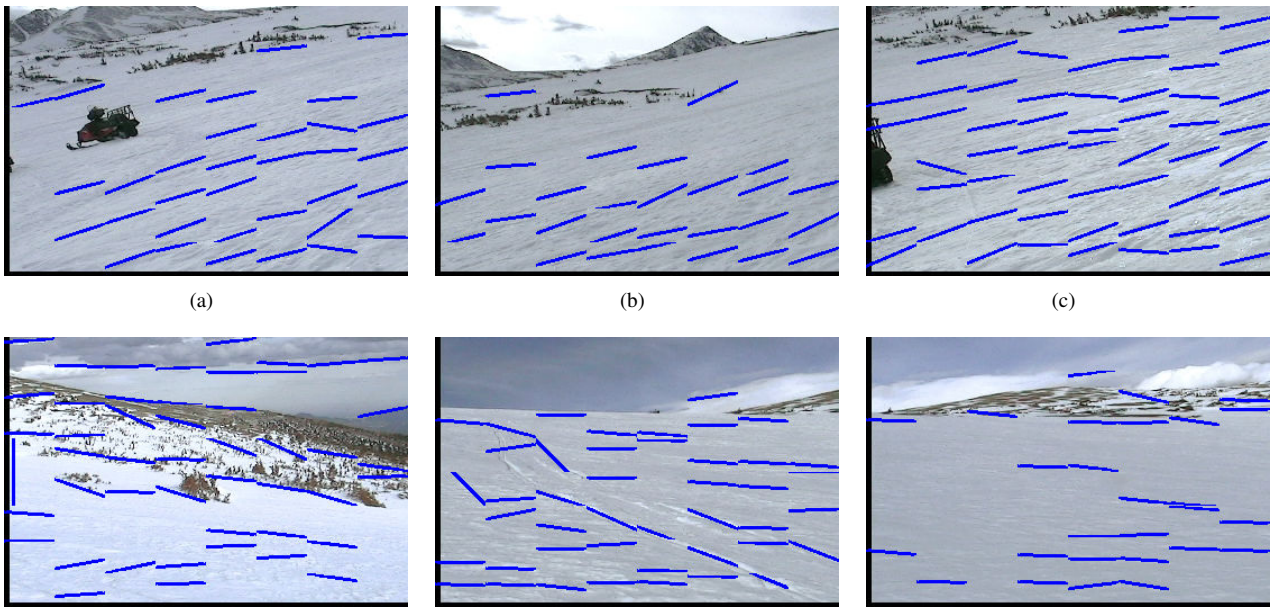

Fig. 3. Slope estimates performed on a variety of images from a glacial region in Colorado.

manageable number, a 2x2 square of available slopes were averaged together reducing the total number to 16. Further, the uppermost row of data generally corresponds to the sky, clouds, or distant mountains. This row of data is ignored by the control system, and, in practice, not calculated at all, bringing the final input count to 12.

A fuzzy rule base was generated in terms of *IF-THEN* statements to control the robot's direction. A human expert was used to generate a set of simple rules with the intention of keeping the robot driving on level ground. Due to symmetry, the rules that are used to turn right mirror the rules to turn left. A total of seven rules, and their corresponding mirrors, were implemented as part of the fuzzy rule base.

VI. SIMULATION

To validate the fuzzy logic navigational control system, a simulated arctic terrain was produced using the Player/Stage/Gazebo project [15], an open source, physics based, 3-D simulation and control package. A series of elliptical hills were placed in the terrain, requiring the robot to make many course corrections in both directions to maintain a level orientation. A topographical map of the simulated terrain is shown in fig. 4(a). A synthetic texture file was then generated, mimicking the dark and light striations present in the actual scenes, and wrapped around the terrain. Fig. 4(b) and 4(c) show the rover in the simulated environment and the augmented view from the rover's camera respectively.

A. Navigational Trials

As long as the rover remains on level ground, the camera's reference frame and global frame will have the same orientation. However, if any body roll is introduced, the values defined as $Flat$ are offset from the global reference frame. As single chip multi-axis inclinometers are inexpensive and readily available, measuring this value seems reasonable. A simulated roll-pitch-yaw sensor was added to the rover, allowing the slope estimates to be transformed back into the global reference frame.

Multiple tests were conducted with the rover's starting position selected at various locations, summarized in Table I. The paths the rover traversed during select trials are superimposed on the elevation map in fig. 4(a). When the rover is initialized on level ground, as was done in trials 1, 2, 3, and 6, the resulting paths remain on level ground and even converge after a short time, demonstrating the robust nature of the control laws. When the rover began on top of a hill, as shown in trial 4, the rover initially turned away from the edge, preferring the flatter section of hill's apex. Once the rover entered the steep, downhill section, the rover turned directly into the slope, reducing the amount of body roll experienced. Immediately after returning to the level area, the rover's heading was sufficiently different from those of the other trials, which were required to navigate around the hill. This was sufficient to cause the path around the next hill to deviated from the standard path of trials 1-3. However, after this turn, the rover rejoins the standard route. Finally, when the rover was initialized in an area surrounded by hills, as was done in trial 5, the rover was able to successfully navigate out and return to the standard route. During the U-turn maneuver required, the rover drives on the shallow outskirts of one of the hills. Due to the minimal slope in this region, the controller lacks sufficient incentive to correct the path immediately. However, as the slope increases, the controller makes an additional path correction, returning it to the standard route.

TABLE I
EXPERIMENT DESCRIPTION

Experiment	Description
1	entrance to the course, starting on level ground
2	entrance to the course, starting on level ground
3	entrance to the course, starting on level ground
4	middle of the course, starting on top of hill
5	middle of the course, surrounded by hills
6	end of the course, starting on level ground

B. Numerical Results

The slope estimates of the real arctic images presented in Section IV could only be evaluated qualitatively, as the ground truth information could not be obtained. Within the simulated environment, however, ground truth data was generated, which allowed for quantitative evaluation. Although, it should be noted that as this is a purely visual technique, the slope estimates are sensitive to the level of realism provided by the simulation. Consequently, the following analysis should be viewed only as an indication of potential performance.

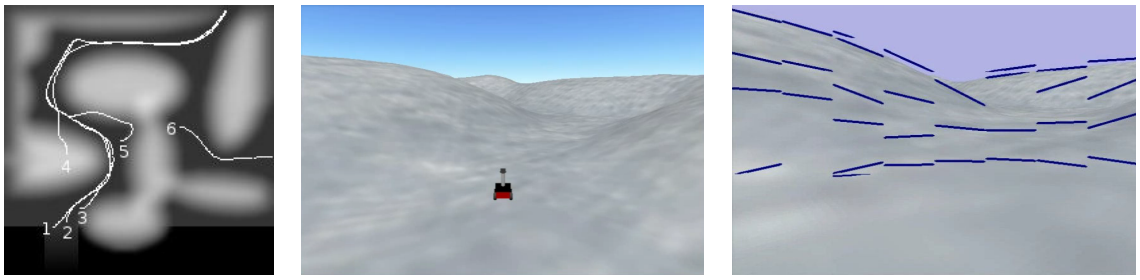

(a) (b) (c)

Fig. 4. (a) A topological map of the simulated terrain with the rover paths superimposed, (b) an illustration of the rover situated in the simulated environment, and (c) the rover's view of the environment with slope estimates.

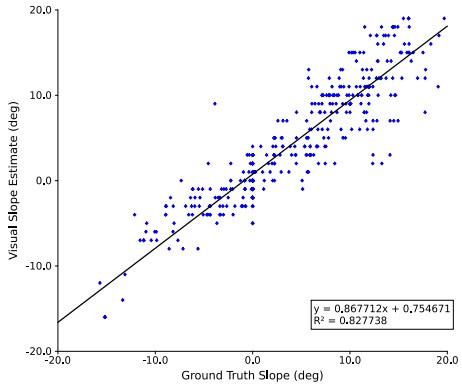

Fig. 5. A plot of the visual slope estimates versus the ground truth data.

The slope estimates generated by this method are two dimensional in nature. However, the actual landscape is a three dimensional entity which can exhibit different slope characteristics in each of those dimensions. When the rover's camera views the landscape, the three dimensional terrain is projected onto a two dimensional plane perpendicular to the camera's line of sight. In order to generate ground truth for the simulation that would be comparable, a similar approach was followed. For each slope estimate, a ray was projected from the camera to a terrain patch at the center of the slope estimate line. Once the intersection of this ray and the terrain is determined, the elevation of the terrain on either side of the intersection point is measured. The ground truth slope is then calculated from the two measured elevations.

With the data collection method in place, the rover was driven manually through a section of the simulated terrain while the generated visual slope estimates and ground truth data were logged approximately once per second. Fig. 5 shows a comparison between the ground truth data and the visual slope estimates obtained during a 60 second traverse. As can be seen, the visual slope estimates are highly correlated with the ground truth data, with a correlation factor above 0.9. The best fit line has a slope of 0.86, indicating that this method tends to mildly underestimate the larger slopes. Over the data set presented, the error between the estimate and ground truth value exhibits a near-zero mean, with a standard deviation of less than 3.0 degrees. The estimates therefore provide a good indication of the terrain slope, which is supported by the success of the navigation trials mentioned above.

VII. CONCLUSIONS

The method outlined in this paper for terrain slope extraction has been shown to be effective both on images of real terrain, and as the input to a real-time fuzzy logic navigational control scheme. The fuzzy logic framework provides adequate performance even with the inherent estimation error, as well as provides a convenient method for translating heuristic human knowledge into control laws. The feature masking approach, originally presented in [4], is particularly effective at removing man-made objects from the analysis area, but sections of unwanted natural objects, such as clouds, often remain. Also, tracks left in the snow, which is an inevitable condition for arctic rovers, are spuriously reported as part of the terrain slope. A method to locate snow tracks and remove them from the processing region could greatly improve the effectiveness of the algorithm as a tool for field robotics.

ACKNOWLEDGMENTS

This work was supported by the National Aeronautics and Space Administration under the Earth Science and Technology Office, Applied Information Systems Technology Program. The authors would also like to express their gratitude to our collaborators Dr. Derrick Lampkin, Pennsylvania State University, for providing the scientific motivation for this research and Dr. Magnus Egerstedt, Georgia Institute of Technology, for providing his experience in multi-agent formations.

REFERENCES

[1] D. Abrosimov *et al.* (2006, Feb) Opencv. intel open source computer vision library. [Online]. Available: http://www.intel.com/technology/computing/opencv/index.htm

[2] D. Apostolopoulos, M. D. Wagner, B. Shamah, L. Pedersen, K. Shillcutt, and W. R. L. Whittaker, "Technology and field demonstration of robotic search for antarctic meteorites," *International Journal of Robotics Research*, Dec 2000.

[3] R. A. Bindschadler *et al.*, "Surface topography of the greenland ice sheet from satellite radar altimetry," National Aeronautics and Space Administration, Tech. Rep. NASA SP-503, 1989.

[4] A. Broggi and A. Fascioli, "Artificial vision in extreme environments for snowcat tracks detection," *IEEE Transactions on Intelligent Transportation Systems*, vol. 3, pp. 162–172, Sep 2002.

[5] S. Brumwell and A. Agah, "Mobility, autonomy, and sensing for mobile radars in polar environments," Information Telecommunication and Technology Center, University of Kansas, Lawrence, KS, Tech. Rep. ITTC-FY2003-TR-20272-02, Nov 2002.

[6] J. L. Chen, C. R. Wilson, and B. D. Tapley, "Satellite gravity measurements confirm accelerated melting of greenland ice sheet," *Science*, Sept 2006.

[7] A. Howard and H. Seraji, "Vision-based terrain characterization and traversability assessment," *Journal of Robotic Systems*, vol. 18, pp. 577–587, Oct 2001.

[8] P. G. Knight, Ed., *Glacier Science and Environment Change*. Blackwell Publishing, 2006.

[9] J. Lever, A. Streeter, and L. Ray, "Performance of a solar-powered robot for polar instrument networks," in *Proceedings of the 2006 IEEE International Conference on Robotics and Automation*, Orlando, Florida, May 2006, pp. 4252–4257.

[10] H. Li, M. A. Lavin, and R. J. L. Master, "Fast hough transform: A hierarchical approach," *Comput. Vision Graph. Image Process.*, vol. 36, no. 2-3, pp. 139–161, 1986.

[11] K. M. Passino and S. Yurkovich, *Fuzzy Control*. Addison-Wesley, 1998.

[12] L. Ray, A. Price, A. Streeter, D. Denton, and J. Lever, "The design of a mobile robot for instrument network deployment in antarctica," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, Barcelona, Spain, April 2005, pp. 2111–2116.

[13] C. Stearns *et al.* (2004, March) Automatic weather stations project and antarctic meteorological research center. [Online]. Available: http://amrc.ssec.wisc.edu/index.html

[14] K. Steffen, J. E. Box, and W. Abdalati. (2005, Feb) Greenland climate network (gc-net). [Online]. Available: http://cires.colorado.edu/science/groups/steffen/gcnet/

[15] R. T. Vaughan, B. Gerkey, and A. Howard. (2007, Aug) The player/stage project. [Online]. Available: http://playerstage.sourceforge.net/

[16] H. J. Zwally and J. DiMarzio. (2005, June) Icesat home page. [Online]. Available: http://icesat.gsfc.nasa.gov/