Visual Detection of Novel Terrain via Two-Class Classification

Christopher A. Brooks and Karl Iagnemma
Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139
617-452-3262
cabrooks@mit.edu, kdi@mit.edu

ABSTRACT

Remote sensing of terrain characteristics is an important component for autonomous operation of mobile robots in natural terrain. Often this involves classification of terrain into one of a set of a priori known terrain classes. Situations can frequently arise, however, where an autonomous robot encounters a terrain class that does not belong to one of these known classes. This paper proposes an approach for visual detection of novel terrain based on a two-class support vector machine (SVM) for situations when known terrain classes can be confidently associated with only a subset of the training data. Experimental results from a four-wheeled mobile robot in Mars analog terrain demonstrate the effectiveness of this approach.

Categories and Subject Descriptors

I.5.2 [**Pattern Recognition**]: Design Methodology – *classifier design and analysis*

General Terms

Algorithms, Experimentation, Theory

Keywords

Machine vision, robot sensing systems, terrain mapping, image classification.

1. INTRODUCTION

Mobile robots are increasingly being used to explore planetary surfaces and accomplish tasks on Earth, to reduce the cost or hazard of human intervention [1], [2]. In situations when direct human control over these robots is limited, either by a time lag or bandwidth limitations, autonomous mobility is critical to improving robotic system performance.

For mobile robots operating in natural terrain, the ability to distinguish between distinct terrain classes is key to enabling safe autonomous mobility. Often, this is accomplished via a supervised classification paradigm, where a system is trained to recognize a small number of a priori known classes. However, situations frequently arise where a robot encounters terrain with characteristics that are not known a priori. In such situations, it is important that a robot recognize that the terrain does not belong to the set of terrain classes it has been trained to distinguish. Failure to identify novel terrain can lead to loss of mobility (e.g. if an exploration robot mistakes hazardous terrain for benign terrain) or missed scientific opportunity (e.g. if a planetary exploration rover ignores terrain that is of scientific interest).

While significant research effort has been focused on terrain classification little research has addressed the question of how to recognize previously unobserved terrain. For example, Rasmussen presented an approach for combining remotely sensed features to distinguish known terrain classes in a terrestrial environment [3]. In another example, the authors demonstrated the classification of terrain based on visual features in Mars rover imagery [4]. However these approaches are restricted to distinguishing between terrain classes known a priori. Similarly, the work by Dima et al. and Manduchi et al. addressed detection of obstacles based on visual features, but required that the appearance of both obstacles and non-obstacles be known a priori [5], [6]. Other researchers have attempted to relax the need for a priori knowledge of visual features by using a classifier trained on-line using a self-supervised framework [7]-[9], however these approaches are not intended to detect novel terrain. In the related field of image segmentation, several researchers have presented approaches for autonomously dividing images into autonomously generated "classes" (e.g. [10], [11]) but segmentation differs in aim from novelty detection in that it does not distinguish a priori known terrain from unknown terrain. One example of related work is the research by McGuire et al. on the use of computer vision techniques to identify features of interest to geologists in natural outdoor settings [12]. However, this work focuses on detection of a pre-defined measure of "uncommonness" which is unrelated to the novelty of the terrain.

This paper presents an approach to detection of novel terrain in color stereo images. It employs color, visual texture, and range-derived geometric data as features for classification. The innovation of this approach is the use of a two-class classifier for novelty detection. While machine learning techniques have been developed for novelty detection (such as the one-class SVM [13], and distribution modeling with a minimum threshold), the proposed approach is intended for situations when known classes are being identified autonomously by some other means. For

This work was funded by NASA through the Mars Technology Program.

SAC'09, March 8-12, 2009, Honolulu, Hawaii, U.S.A.
Copyright 2009 ACM 978-1-60558-166-8/09/03…$5.00.

example, this approach can be used to visually distinguish between terrain belonging to a class a robot has already come in physical contact with and terrain it has not yet sensed. As such, this approach takes advantage of the ready availability of unlabeled visual data to improve the accuracy of novel class detection, and is appropriate for on-the-fly creation of a novelty detector without a human in the loop. The use of this unlabeled data was inspired by the work on optimal single-class classification strategies by El-Yaniv and Nisensen [14].

Section II of this paper describes the two-class classification approach, as well as two baseline approaches which will be used for performance comparison. Section III gives details about experiments conducted to analyze the performance of the proposed approach, and Section IV presents the results of those experiments. Section V presents conclusions and suggests directions for future research.

2. PROPOSED APPROACH

The goal of a novel terrain detection algorithm is to distinguish between terrain belonging to a set of a priori known terrain classes and terrain belonging to an unknown "novel" class. The problem is distinct from traditional two-class classification because labeled examples of the novel class do not exist. Previously proposed approaches to novelty detection, such as those based on a one-class SVM or distribution modeling, are trained using only the labeled examples of known terrain classes [13], [15]. In the scenario considered here, however, it is assumed that information is available about both the a priori known classes and a "world" class, which is a mixture of both known and unknown classes. This situation arises when only a subset of the terrain observed by the robot is associated with the known terrain classes, for example when terrain is labeled as "known" only after the robot has come into physical contact with it, or when only a small fraction of the observed terrain is sensed using some other instrument. The two-class classification approach presented here uses a SVM classifier trained to distinguish the known classes from the world class, thereby implicitly identifying the novel terrain. For comparison, baseline approaches based on a one-class SVM and a distribution model are also presented in Section II.C.

2.1 Theoretical Justification

Here an analysis of a two-class classification framework for novelty detection is presented. Justification for use of this approach is founded on a comparison of an optimal novelty detector and an optimal two-class classifier when the probability density functions of all classes are known. In this section, it is shown that the novelty detector and a two-class classifier can identify the same regions associated with novel and known terrain.

In the following analysis, it is assumed that the world class, W, represents all of the terrains in the robot's environment. W is composed of two disjoint classes: A (known terrain) and B (novel terrain). Any patch of terrain will belong to A or B, but not both. It is also assumed that the probability density functions associated with A and B are known. Here $\mathrm{p}(\mathbf{x}|A)$ and $\mathrm{p}(\mathbf{x}|B)$ represent the probability density of a feature vector $\mathbf{x}$ occurring given that it is associated with terrain class A or B, respectively. In this scenario, the goal of novelty detection is to find the region of the feature space containing the largest fraction of B, while containing only a specified fraction (r) of A (ideally zero). Thus, the novelty detection problem is posed as

$$\mathbf{S}_{ND}^{\ *} = \arg\max_{\mathbf{S}} \int_{\mathbf{S}} \mathrm{p}(\mathbf{x}|B)d\mathbf{x} \ \text{ s.t. } \int_{\mathbf{S}} \mathrm{p}(\mathbf{x}|A)d\mathbf{x} = r \qquad (1)$$

where $\mathbf{S}$ is a region of the feature space, and $\mathbf{S}_{ND}^{\ *}$ is the region of the feature space identified as novel. Note that r is the rate of false positives—instances when known terrain is incorrectly identified as novel. Ideally r can be set to zero, but if there exists some $\mathbf{x}$ where both $\mathrm{p}(\mathbf{x}|A)>0$ and $\mathrm{p}(\mathbf{x}|B)>0$, increasing r will have the benefit of increasing the detection rate of novel terrain, which may be desirable. The solution to the novelty detection problem is

$$\mathbf{S}_{ND}^{\ *} = \{\mathbf{x} : \mathrm{p}(\mathbf{x}|A) \le t\,\mathrm{p}(\mathbf{x}|B)\} \qquad (2)$$

where t is found such that

$$\int_{\mathbf{S}_{ND}^{\ *}} \mathrm{p}(\mathbf{x}|A)d\mathbf{x} = r \quad . \qquad (3)$$

Here t is a monotonic function of r. This solution can be understood intuitively by contradiction. If there is some region outside $\mathbf{S}_{ND}^{\ *}$ where $\mathrm{p}(\mathbf{x}|A) < t\,\mathrm{p}(\mathbf{x}|B)$, then the fraction of B in $\mathbf{S}_{ND}^{\ *}$ could be increased by changing $\mathbf{S}_{ND}^{\ *}$. $\mathbf{S}_{ND}^{\ *}$ is unique as long as there is no region of measure greater than zero where $\mathrm{p}(\mathbf{x}|A) = t\,\mathrm{p}(\mathbf{x}|B)$.

In the novelty detection scenario, the distribution of the novel class is unknown, so $\mathrm{p}(\mathbf{x}|B)$ cannot be used directly. Instead, we assume it is possible to estimate $\mathrm{p}(\mathbf{x}|W)$, the probability density of the world class, which is given as

$$\mathrm{p}(\mathbf{x}|W) = \gamma\,\mathrm{p}(\mathbf{x}|A) + (1-\gamma)\,\mathrm{p}(\mathbf{x}|B). \qquad (4)$$

Here γ is unknown, indicating the fraction of A in W. Using $\mathrm{p}(\mathbf{x}|W)$, the A vs. W two-class classification problem is

$$\mathbf{S}_{W}^{\ *} = \arg\min_{\mathbf{S}} \left[c_A \int_{\mathbf{S}} \mathrm{p}(\mathbf{x}|A)d\mathbf{x} + c_W \left(1 - \int_{\mathbf{S}} \mathrm{p}(\mathbf{x}|W)d\mathbf{x} \right) \right] \ , \qquad (5)$$

where the region $\mathbf{S}_W^{\ *}$ is the region of the feature space associated with the world class. Here c_A is the cost of misclassifying A as W, and c_W is the cost of misclassifying W as A. The solution to this two-class classification problem in terms of $\mathrm{p}(\mathbf{x}|A)$ and $\mathrm{p}(\mathbf{x}|W)$ is

$$\mathbf{S}_{W}^{\ *} = \left\{ \mathbf{x} : \mathrm{p}(\mathbf{x}|A) \le \frac{c_W}{c_A}\mathrm{p}(\mathbf{x}|W) \right\}. \qquad (6)$$

Written in terms of $\mathrm{p}(\mathbf{x}|A)$ and $\mathrm{p}(\mathbf{x}|B)$, this becomes

$$\mathbf{S}_{W}^{\ *} = \left\{ \mathbf{x} : \mathrm{p}(\mathbf{x}|A) \le \frac{c_W(1-\gamma)}{c_A - \gamma c_w}\mathrm{p}(\mathbf{x}|B) \right\}. \qquad (7)$$

Thus, it can be seen that $\mathbf{S}_{ND}^{\ *}$ and $\mathbf{S}_W^{\ *}$ are both defined as regions where $\mathrm{p}(\mathbf{x}|A)$ is less than some factor multiplied by $\mathrm{p}(\mathbf{x}|B)$. Without knowledge of γ or $\mathrm{p}(\mathbf{x}|B)$, c_A and c_W can be found from (6) such that

$$\int_{\mathbf{S}_W^{\ *}} \mathrm{p}(\mathbf{x}|A)d\mathbf{x} = r \quad . \qquad (8)$$

With these values of c_A and c_W, $\mathbf{S}_{ND}^{\ *}$ and $\mathbf{S}_W^{\ *}$ are identical. Thus, it follows that a two-class classifier distinguishing the known class from the world class can be used for detection of the novel class.

2.2 Detailed Approach

The two-class classification approach proposed here employs a support vector machine to classify terrain based on visual features. This section discusses details about the visual features, the classifier, and the training process.

2.2.1 Description of Visual Features

The feature vector used to represent the terrain was composed of color, visual texture, and range-derived geometric information extracted from images collected by a forward-looking color stereo camera pair.

Color data is available for every pixel as red, green, and blue (RGB) intensities. Other researchers have observed that classification in the RGB space can be significantly affected by changes in illumination, so the color feature used here is the modified hue, saturation, and value (HSV) representation used in [7]. Here hue (an angle) is represented by two elements—sin(hue) and cos(hue)—to eliminate the artificial discontinuity at 2π. Thus, color is represented as a 4-element vector: [sin(hue), cos(hue), saturation, value].

Visual texture is a measure of the local spatial variation in intensity in the image. Researchers have proposed many metrics for visual texture, using Gabor filters and local energy methods [16], [17]. This work uses a wavelet-based approach, similar to the one demonstrated in [18]. Here the grayscale image is decomposed with the Haar wavelet. Three scales of wavelets are used, each scale having horizontal, diagonal, and vertical (HDV) wavelets, corresponding to derivative estimates in the horizontal, diagonal, and vertical directions at each length scale. The scales used are 2, 4, and 8 pixels. Because this data can be noisy, the magnitudes of the wavelet coefficients are averaged over windows of 7, 9, and 11 wavelets. Thus, visual texture is represented by a 9-element vector, composed of the window-averaged horizontal, diagonal, and vertical wavelet coefficients at each scale.

Geometry data is available for the terrain through stereo image processing. The raw output of a stereo image processing algorithm is a cloud of range data points. For this work, we divide the points into a grid of 20-cm by 20-cm terrain patches, based on their horizontal position. We then calculate statistics based on the elevation of points within each grid cell. The first element of the geometric feature vector is the average slope of the terrain, defined as the angle the least-squares-fit plane makes with the horizontal plane. The second element is the mean-squared deviation from that plane along its normal. The third element is the variance in height of the pixels, and the fourth is the height difference between the highest and lowest pixels in the patch. Thus, the geometry of each patch is represented at a 4-element vector: [slope, plane fit deviation, height variance, height range]. These features were chosen based on previous work by the authors, which employed an identical feature set for supervised classification of outdoor terrain, with good performance [9].

For the two-class classification approach, the visual features associated with color, visual texture, and geometry were concatenated into a single feature vector with 17 elements. Note that while the color and visual texture features were distinct for each pixel, the geometric features were identical for all pixels in the same terrain patch.

2.2.2 Classifier Description

For this work, the two-class classifier was implemented as a support vector machine using the open-source library LIBSVM [19]. For this work, a Gaussian kernel was used. Parameters of the Gaussian kernel were optimized by cross-validation over a subset of the images used for training.

Classification of the terrain was performed with the terrain represented as a grid of 20-cm by 20-cm terrain patches, with a single class associated with each terrain patch. Each patch contained observations associated with multiple pixels in an image, and the feature vector associated with each pixel was classified separately by the support vector machine. The classification result for the entire patch was calculated by a majority vote of the individual pixel classification results.

2.2.3 Training Process

Two sets of data are necessary to train the two-class classifier: a set of data associated with terrain of a user-selected a priori known class(es), and a set of unlabeled data which represents the distribution of terrain in the world class. Four hundred pixels chosen at random from three hand-labeled stereo images were used to represent the known class. For each pixel, color and texture features were extracted. Geometric features were extracted from the associated 20-cm by 20-cm terrain patch. Two thousand pixels chosen at random from ten unlabeled stereo images were used to represent the world class. The features vectors associated with these pixels were used to train the two-class classifier.

2.3 Baseline Approaches

To assess the performance of the proposed two-class classification approach, it was compared to two existing approaches for novelty detection: one-class SVM and mixture-of-Gaussians distribution modeling.

2.3.1 One-Class SVM

The one-class SVM, described in [13], is a classification framework designed to segment a feature space into "same" and "different" classes. While a traditional SVM operates by finding the hyperplane with the largest margin separating two classes in a Hilbert space, the one-class SVM works by finding the hyperplane with the largest margin separating the "same" class from the origin in that Hilbert space.

As with the two-class SVM, the one-class SVM was implemented using LIBSVM. A Gaussian kernel was used, and parameters of the SVM and its kernel were optimized by cross-validation over a subset of the training images. The one-class SVM was trained using four hundred pixels belonging to the known class from three hand-labeled stereo images, using the same concatenated visual feature vector as the two-class classifier. As above, color and texture data were extracted for each pixel and the geometric features were extracted from the associated terrain patch.

2.3.2 Distribution Modeling

Another standard approach to novelty detection is distribution modeling, which we implemented using a mixture of Gaussians (MoG) model. This approach is focused on modeling the underlying probability density function of the known class. To classify a new feature vector as known or novel, the probability density function is evaluated and compared to a constant threshold value. If the probability density is above the threshold, the feature is classified as known. Otherwise, the feature is classified as novel. It should be noted that this is equivalent to assuming that the world class is uniformly distributed.

The mixture of Gaussians model was implemented in Matlab using the expectation maximization (EM) procedure for training [15], [20]. The number of Gaussian modes was optimized by cross-validation over a subset of the training images. As with the other classifiers, feature vectors from four hundred pixels belonging to the known class were used for training.

Due to the danger of over-fitting Gaussian modes in a high-dimensional feature space, color, visual texture, and range features were modeled separately in this approach. A separate MoG model was trained for each feature, and the naïve Bayes assumption was made to calculate the combined probability density function as the product of the three separate models.

3. EXPERIMENT DETAILS

The approaches for detection of novel terrain were compared using data from experiments with TORTOISE, a four-wheeled mobile robot developed at MIT, in Mars-analog outdoor terrain.

3.1 Robot Configuration

TORTOISE, shown in Figure 1, is an 80-cm-long, 50-cm-wide, 90-cm-tall robot with 20-cm-diameter rigid aluminum wheels with grousers. The wheels on either side are connected to the main body and mast via a differential.

The robot is outfitted with a forward-looking mast-mounted stereo camera pair, a body-mounted two-axis tilt sensor, and several local-terrain sensors not used for this work. Forward-looking images were captured with the 19-cm-baseline stereo camera pair, a Videre Design “dual DCAM” capable of capturing color images at 640 x 480 resolution. Range data were extracted from the stereo images using Videre Design’s commercial stereo processing software with sufficient accuracy to assign pixels to terrain cells up to a distance of 8 meters. Body pitch and roll were measured with a Crossbow CXTA02 two-axis tilt sensor.

During experiments, the rover traveled at a speed of 6 cm/sec. Stereo images were captured every 1.5 seconds. These data were stored during experiments and processed offline.

3.2 Experimental Environment

Experiments were performed at Wingaersheek Beach, in Gloucester, MA. This is a sandy beach with a mixture of small and large outcrops (relative to the size of the rover) and loose rocks. This site was chosen due to its similarity in appearance to the MER landing sites on Mars. In this environment, sand and rock were considered to be two distinct terrain classes. To demonstrate the ability of the novel terrain detection in a multi-class setting, matted piles of beach grass were used as a third terrain class. These three terrain classes are identified in Figure 1. Visually, the sand appeared as a uniform gray flat surface, rock appeared tan and orange with some steep slopes and fine uniform texture, and beach grass appeared highly-textured with mixed browns and shadows.

Six experimental data sets were collected, each during a rover traverse of at least 10 meters along a straight-line path containing a combination of the terrains. No two paths were identical. In all, 1646 image pairs were collected. During the experiments lighting conditions ranged from overcast conditions with diffuse light to cloudless conditions with low, direct sunlight.

3.3 Data Processing

Data were collected during the experiments and stored for processing offline. In each of the data sets, every fifth image pair was passed to the stereo processing software for range data extraction. Every fourth image pair with range data was hand-labeled to identify a ground-truth terrain class for each pixel. This spacing was chosen to reduce the effect of repeated training and testing on similar images of the same terrain patch. Labeled examples of data corresponding to the known (e.g. sand and beach grass) classes were drawn from the first three hand-labeled images in a data set. Unlabeled examples, corresponding to the “world” class, were drawn from the first 10 (labeled and unlabeled) images with range data in a data set. Performance of the novel terrain detectors were assessed based on the remaining hand-labeled

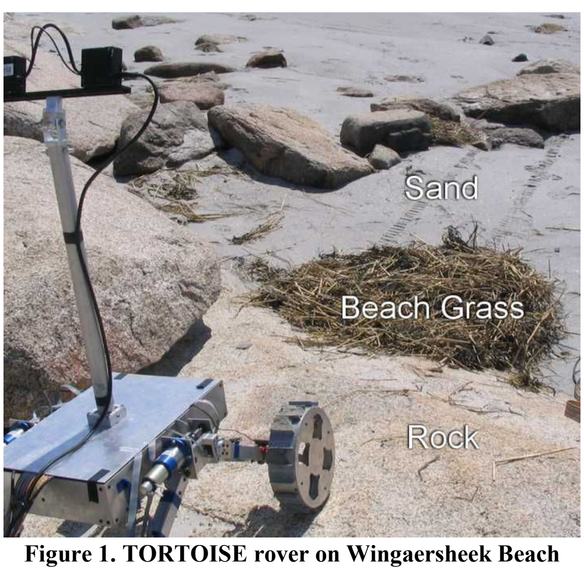

Figure 1. TORTOISE rover on Wingaersheek Beach


Figure 2. Image from stereo camera

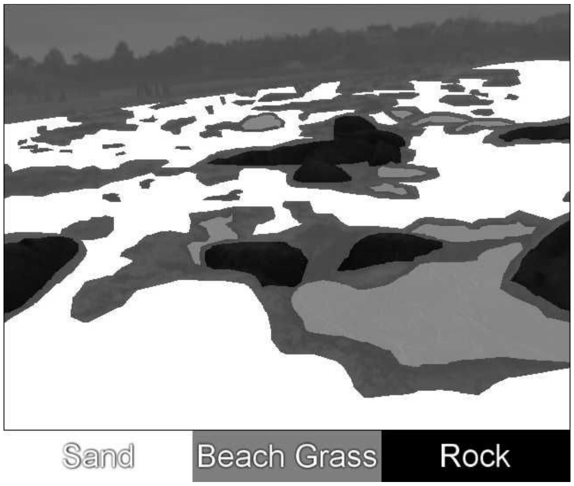

Figure 3. Hand-labeled terrain classes

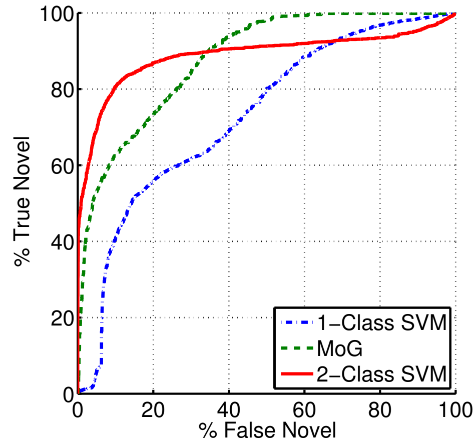

Figure 4. ROC curves for baseline approaches and two-class classification approach for all data sets

images in a data set—75 hand-labeled images in all. Figure 2 shows a sample image from the stereo camera. The corresponding hand-labeled terrain classes are shown in Figure 3. Note that the background in Figure 3 is not labeled, as it is beyond the limit for accurate stereo-based range estimation.

4. RESULTS

The performance of the two-class classification approach was compared to that of the baseline approaches on each test set. Figure 4 shows a receiver operating characteristic (ROC) curve showing the results of each approach across all of the data sets. For this figure, the rock class was used as the novel class, so the labeled training data was drawn only from the sand and beach grass classes. Here the vertical axis indicates the percentage of true novel detection (%TN) (i.e. the fraction of rock terrain patches which were correctly identified as novel) and the horizontal axis indicates the percentage of false novel detection (%FN) (i.e. the fraction of sand or beach grass terrain patches which were incorrectly identified as novel). Each detector forms a curve on the plot, as the threshold for labeling terrain as novel is adjusted. Note that random assignment of terrain as novel would yield a diagonal line from (0,0) to (100,100).

In this figure, it can be seen that the best performance was demonstrated by the two-class classification approach, and the worst was demonstrated by the one-class SVM. This difference is particularly significant in the detection of novel terrain with less than 20% false novel detection, where the two-class classifier detected nearly 90% and the one-class SVM detected less than 60%. To summarize these ROC curves, the optimal performance is measured in terms of the maximum difference between %TN and %FN. Using this metric for optimality, the baseline one-class SVM detected 52% of novel terrain while misidentifying 15%, and the baseline MoG distribution modeling detected 89% while misidentifying 33%. The two-class classification approach performed significantly better, detecting 83% of the novel terrain while misclassifying only 12%.

The numerical results summarizing the detection of novel terrain in all data sets is shown in Table 1. Here the two baseline approaches are compared to the proposed two-class classification approach. Mean and standard deviation are compared for each of three metrics: % True Novel, % False Novel, and the area under the curve, P(A) [21]. Note that the numbers in brackets indicate the 95% confidence interval for the mean. For each of these metrics, the two-class classification approach outperformed the baseline approaches. Perhaps the most important feature is that the standard deviations in these metrics are significantly lower for the two-class classification approach than the baseline approaches, indicating that the two-class approach demonstrates more consistent results. Similar results were observed using beach grass as the novel class. Use of sand as the novel class yielded significantly deteriorated performance, because very little training data was available from the known class.

Table 1. Comparison of novel terrain detection approaches

	One-Class SVM	MoG Distribution Modeling	Two-Class SVM
Mean % True Novel	78.4% [57.9% - 98.9%]	85.2% [67.3% - 100.0%]	89.4% [81.9% - 96.9%]
St. Dev. of % True Novel	28.7% [19.7% - 52.4%]	24.8% [17.0% – 45.2%]	10.5% [7.2% – 19.1%]
Mean % False Novel	22.9% [2.8% - 43.1%]	23.2% [9.7% - 36.6%]	14.5% [6.4% - 22.6%]
St. Dev. of % False Novel	28.2% [19.4% - 51.5%]	18.8% [12.9% – 34.3%]	11.4% [7.8% – 20.7%]
Mean P(A)	71.9% [52.1% – 91.6%]	79.7% [66.1% - 93.4%]	88.7% [81.3% – 96.2%]
St. Dev. P(A)	27.6% [19.0% – 50.3%]	19.1% [13.1% – 34.9%]	10.4% [7.1% – 18.9%]

Classification using the two-class approach on a desktop computer ranges from 12-70 s for a 640 x 480 image. Computation time varies based on the complexity of the boundary between the known and novel classes, and the fraction of the image with valid range data.

5. CONCLUSIONS

In this paper, a two-class classification approach was presented to detect novel terrain. This approach uses samples of visual data from unlabeled images to represent the world class for use in training the novel terrain detector. A theoretical analysis showed the common goal of two-class classification and novelty detection. Using data from experiments with a rover in outdoor terrain, this two-class classification approach was compared to two baseline approaches for novelty detection which use only data from known terrain classes during training. The results demonstrate that the two-class classification approach achieves a higher accuracy while reducing the variability between data sets.

This two-class classifier approach for novelty detection can be used in any system where examples of data from a mixture of known and novel classes are available, to improve accuracy beyond what is achievable with traditional novelty detectors.

5.1 Future Work

For a robot in rough terrain, the computational cost of novel terrain detection can be significant. Computation time could likely be reduced by an order of magnitude or more by replacing the Gaussian kernel of the SVM with a polynomial kernel. This is a focus of current work.

6. REFERENCES

[1] NASA. “The Global Exploration Strategy: The Framework for Coordination,” Press Release 07-126, Retrieved February 2, 2008 from http://www.nasa.gov/pdf/178109main_ges_framework.pdf

[2] Committee on Army Unmanned Ground Vehicle Technology, National Research Council. *Technology Development for Army Unmanned Ground Vehicles*, National Academies Press, Washington, DC., 2002.

[3] C. Rasmussen, (2001, December). “Laser Range-, Color-, and Texture-based Classifiers for Segmenting Marginal Roads,” in *Proc. of the Conference on Computer Vision & Pattern Recognition Technical Sketches*, Kauai, HI.

[4] I. Halatci, C. A. Brooks, and K. Iagnemma, "A Study of Visual and Tactile Terrain Classification and Classifier Fusion for Planetary Exploration Rovers," *Robotica*, Published online by Cambridge University Press 24 Apr 2008, doi:10.1017/S0263574708004360.

[5] C. S. Dima, N. Vandapel, and M. Hebert. “Classifier fusion for outdoor obstacle detection,” in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA) 2004*, 1, pp. 665-671, doi: 10.1109/ROBOT.2004.1307225.

[6] R. Manduchi, A. Castano, A. Talukder, and L. Matthies, L. “Obstacle detection and terrain classification for autonomous off-road navigation,” *Autonomous Robots*, 18, 81-102, May 2005.

[7] B. Sofman, E. Lin, J. A. Bagnell, N. Vandapel, and A. Stentz. “Improving Robot Navigation Through Self-Supervised Online Learning” in *Proc. of Robotics: Science and Systems*, August 2006.

[8] D. Kim, J. Sun, S. M. Oh, J. M. Rehg, and A. F. Bobick. “Traversability Classification using Unsupervised On-line Visual Learning for Outdoor Robot Navigation” in *Proc. of the 2006 IEEE International Conference on Robotics and Automation*, Orlando, Florida. May, 2006.

[9] C. A. Brooks and K. Iagnemma. "Self-Supervised Classification for Planetary Rover Terrain Sensing." in *Proc. of the 2007 IEEE Aerospace Conference*, Big Sky, Montana, March 2007.

[10] R. Manduchi. “Bayesian Fusion of Color and Texture Segmentations” Presented at ICCV’99, Kerkyra, 1999

[11] T. P. Weldon, W. E. Higgins, and D. F. Dunn. “Efficient Gabor Filter Design for Texture Segmentation” *Pattern Recognition* 1996.

[12] P. C. McGuire, E. Diaz Martinez, J. O. Ormö, J. Gomez Elvira, J. A. Rodriguez Manfredi, E. Sebastian Martinez, H. Ritter, R. Haschke, M. Oesker, J. Ontrup, "The Cyborg Astrobiologist: Scouting Red Beds for Uncommon Features with Geological Significance", *International Journal of Astrobiology*, Vol. 4, No. 2, pp. 101-113, 2005.

[13] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson. “Estimating the support of a high-dimensional distribution.” *Neural Computation*, 13(7):1443-1471, 2001.

[14] R. El-Yaniv and M. Nisenson. “Optimal Single-Class Classification Strategies.” in *Advances in Neural Information Processing Systems (NIPS)* 20, MIT Press, Cambridge, MA, 2006.

[15] C. M. Bishop. *Neural networks for pattern recognition*. New York: Oxford University Press. 1995.

[16] C. Bouman and B. Liu. “Multiple Resolution Segmentation of Textured Images,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 13, No. 2, 1991, pp. 99-113.

[17] T. Reed and J. Hans du Buf. “A review of recent texture segmentation and feature extraction techniques,” *CVGIP: Image Understanding*, Vol. 57(3), May 1993, pp. 359-372.

[18] F. Espinal, T. Huntsberger, B. Jawerth, and T. Kubota. “Wavelet-Based Fractal Signature Analysis for Automatic Target Recognition,” *Optical Engineering*, Vol. 37, No. 1, pp. 166-174, January 1998.

[19] Chih-Chung C., and Chih-Jen L. LIBSVM: a library for support vector machines. 2001. Software retrieved January, 2006 available at http://www.csie.ntu.edu.tw/~cjlin/libsvm.

[20] J. Bilmes. “A Gentle Tutorial on the EM Algorithm and its Application to Parameter Estimation for Gaussian Mixture and Hidden Markov Models,” Technical Report, University of Berkeley, 1997.

[21] A. J. Simpson and M. J. Fitter. “What is the best index of detectability?” *Psychological Bulletin*, Vol. 80, No. 6, 1973, pp. 481-488.