Learning Long-Range Terrain Classification for Autonomous Navigation

Max Bajracharya, Benyang Tang, Andrew Howard, Michael Turmon, Larry Matthies

Abstract— This paper describes a method for learning the terrain classification of long-range appearance data from short-range, stereo-based geometry, along with a map representation for utilizing this data to improve autonomous off-road navigation. The continuous, online learning method allows the system to constantly adapt to changing terrain and environmental conditions, while the polar-perspective map representation allows the system to effectively plan with stereo data at long ranges. Various evaluations of the long-range classification and improvements in system performance are described, including results from an independent third-party testing team.

I. INTRODUCTION

The ability to evaluate and exploit long-range perception data is critical for autonomous unmanned ground vehicles (UGVs) to avoid myopic behavior and increase driving speed. The viewing angle and sensor characteristics of both active and passive range sensors typically used by UGVs limit the range at which the terrain can be classified based on geometry. At long ranges, range data becomes too sparse or noisy to evaluate terrain traversability from geometry. For a typical UGV, this range is generally 30-50m; for the vehicle used in this paper, this range is around 10m due to the sensor's lower height and shorter baseline. When autonomously navigating off-road outdoor terrain with only short-range sensors and without a global map, UGVs exhibit problems such as driving into long cul-de-sacs or waiting to avoid large, distant obstacles. Furthermore, because a vehicle must be able to safely stop before hitting an obstacle, the vehicle's speed is limited by the maximum safe distance that can be detected.

Utilizing the long range data available from cameras that can see the horizon enables planning distances to be extended beyond where geometry can be used to classify the terrain, and can consequently be used to both avoid myopic behavior and safely increase the vehicle speed over the terrain. However, long-range perception data presents challenges in both classifying the terrain traversability and then planning on the classified data. Image-based classification of long-range data is difficult because the color and texture of traversable terrain is dependent on the locale (including the time of day and viewing angle) and is not easily generalized. Planning is difficult because range information obtained from stereo has an error that grows quadratically with range.

To address these problems, we combine self-supervised, online learning of long-range terrain traversability with a polar-perspective map representation for long-range planning. Continuous on-line learning of long-range terrain classification using color and texture from geometry-based, short-range terrain classification allows the system to constantly adapt to its local terrain and avoids the need to learn a universal long-range terrain classifier. Then projecting the classified terrain into a polar-perspective map whose cell sizes are a function of the range enables a standard graph planner to make effective use of data out to the limit of stereo matching (zero disparity). A visualization of a polar-perspective map with classified long-range terrain is shown in Figure 1.

The terrain classification and path planning algorithms developed run in real-time as part of a larger navigation system, which includes visual pose estimation and path following. Various feature sets and classification methods have been studied off-line and the final system has been tested extensively onboard an autonomous vehicle in outdoor terrain. The system has been fielded and evaluated by a third-party test team as part of the DARPA Learning Applied to Ground Robots (LAGR) program and has demonstrated the ability to detect and avoid obstacles at long range.


Fig. 1. A 3D view of a polar-perspective map showing traversable (blue), long-range non-traversable (yellow), and unknown (green) terrain, with the best path (white) and stereo point cloud superimposed; along with a corresponding image and its long-range classification (traversable in blue and non-traversable in red)

A. Related Work

The LAGR program's goals include addressing both the need for an adaptable terrain classifier to avoid hand-tuned

systems, as well as long-range perception to avoid myopic behavior [6]. In order to avoid using hand-labeled data, as done in [9] among others, researchers in the program are making use of the concept of generalizing the near-field traversability to the far-field in a self-supervised approach introduced in [10] (extended in [5]). Among others, [12] automatically selects from a set of several binary classifiers, and [3] uses distance normalized images and a convolutional network. The DARPA Grand Challenge has motivated the need for faster driving and the similar approach of extending the appearance of the road out to long ranges has been instrumental in achieving good results [2], [8]. Furthermore, learning traversability of overhead imagery from local data is another method of extending the range of perception and has been effective on terrain where the overhead imagery is available [14]. Overhead imagery makes planning on long-range data easier; however, we consider the case when it is not available. An alternative to projecting long-range data into a map is to use image-based planning, where the plan is generated directly in the image [11].

II. LEARNING LONG RANGE TERRAIN CLASSIFICATION

Our basic approach to long-range perception is to use a self-supervised learning method to generalize short-range terrain classification from stereo-based geometry to long-range terrain classification from imagery. This includes classifying terrain locally with stereo data, selecting features from this terrain, learning a two-class classifier of traversable and non-traversable terrain, and then applying this classifier to an image to classify long-range data. We have found that color-based features and a linear support vector machine (SVM) are an effective and practical combination to achieve good performance in real-time.

A. Local Terrain Classification

The traversability of terrain in the near field is classified using the stereo range data produced by stereo camera images. Each pixel with valid stereo data is projected into a cell of an instantaneous 2D Cartesian map and geometry statistics are accumulated for each map cell. Then, each cell is classified as traversable or non-traversable based on a previously learned classifier. Finally, the original image is labeled by back-projecting the map into the image.

The Cartesian map uses a 16cm x 16cm grid size and accumulates the zero, first, and second order statistics (number of points, mean, and variance) of points from a fitted ground plane, as well as a maximum step height, which is the difference between the mean height of adjacent cells. The classifier is a simple histogram-based naive-Bayes classifier. The classifier is learned by adding traversable (terrain the vehicle drives over) and non-traversable (terrain in front of the vehicle when an operator triggers an E-stop) examples to each of the 1D feature histograms. Terrain is then classified by computing the combined probability of traversability for each cell in the map. Although simple and limited to low

dimensionality, this terrain classification approach is fast and robust.

Once the local terrain map has been classified, the map is back-projected into the image using the stereo information for that image. Fixed thresholds on the probability of traversability are used to produce an image labeled with traversable, non-traversable, and unknown terrain types. This image then serves as the training signal for long-range learning. An example of locally classified terrain is shown in Figure 2.


Fig. 2. The original image (top left), false-color (green to blue indicating decreasing disparity) stereo disparity image (top right), a back projection of the local terrain classification (where color represents the probability of traversability, from blue (traversable) to red (non-traversable) (bottom left), and the traversable (blue) and non-traversable (magenta) feature windows selected for the locally classified terrain (bottom right)

B. Feature Selection

To learn a long-range classifier, features are selectively chosen from the back-projected image and accumulated with a sliding window over many frames. Because at long-range, image pixel colors are the vehicle's only sensory input, the problem of long-range learning thus becomes a problem of encoding the robot traversability in appearance. After evaluating several intensity and color features (various color spaces, Gabor filters, and color histograms), a combination of pixel color and normalized color histograms was fielded. Although computationally efficient and simple, the use of a single pixel or window average of raw color, normalized color, or other color space value (such as HSV) is limited in its distinguishing power because it does not account for texture. On the other hand, Gabor filters can capture texture well, but are computationally expensive to use. We reduced the set of filters to a fixed number by using cross-validation to find the set of scales and orientations that maximized classification, but found that color histograms could achieve similar classification rates with significantly lower computation. We use a 2D normalized color histogram, computed using $r = \frac{R}{R+G+B}$ and $g = \frac{G}{R+G+B}$, with 16x16 bins, and generated from a 33x33 pixel patch. The use of the r-g color space over the raw RGB color space offers more lighting independence and a smaller feature length, while the histogram captures a basic notion of texture. A

further improvement to the 2D normalized color histogram was made by using a grid of adaptive, uneven bins. Because the histogram distribution in the 2D r-g color space is very uneven, a constant sampling does not capture the color signal effectively (the resulting histogram will have a zero value in many of the bins all the time). So, for each of the normalized colors r and g, a grid of 16 1D bins are designed so that each bin contains the same number of pixels. Then, these two grids of 1D bins form a grid of 2D bins and are used to compute the histograms during the training and classification phases.

During the training phase, many training samples can be collected from each image with near-field back-projected labels. However, if all the training data is collected within a single frame, the resulting model may not respond well to the vehicle's changing environment and will lead to an unbalanced data set. Consequently, we collect training data from multiple frames and limit the number of training samples from each frame. Furthermore, because there is generally little frame-to-frame change for small vehicle motions, a frame is only used after the vehicle has moved a specific distance from the previously used one. For our experiments, we limited each frame to have 100 training samples (50 traversable and 50 non-traversable, sampled from all available features in the image), taken every one meter of motion. The training samples are accumulated into a FIFO (first in, first out) queue consisting of 4000 samples; when new training samples are collected, the oldest samples are removed to maintain a constant number of training samples. Figure 2 shows an example of features selected from a single frame.

C. Learning

Using the training samples selected from the near-field labeled images, a linear support vector machine (SVM) with soft margins [13] is learned and used to classify pixels in the far field of the current and future images. The model regularization parameter (C), which penalizes large errors, is chosen to balance computation time and data separation. Using cross-validation experiments, we have fixed C to 500 for our experiments. Furthermore, during evaluation, we leave any samples that are within the SVM margins as unclassified. In practice, we find that these samples account for about 10% to 50% of the data, and in the reported results, we allow 40% of the data to remain unclassified. The results described in this paper require that the training set be balanced, and so the system must be trained with both traversable and non-traversable classes. However, the system has since been extended to learn based on both near and far field data, and so can learn even without non-traversable terrain in the near field.

The linear SVM chosen was initially evaluated alongside a linear discriminant method and non-linear SVM with a radial basis function, and proved to be the best compromise between accuracy and computation speed. For instance, on a typical dataset, the use of a non-linear SVM achieved a correct classification rate of 92%, as compared to 88% for a linear SVM (on near-field data), but required 2-3 times

the computation time. For the linear SVM, we have also considered using a power transform (using $(\text{histogram})^a$ as the input) which is reported to boost accuracy [1], and experimented with several values of a, but only report $a = 0.5$. As a baseline, we also train SVM classifiers with RGB colors and their squares, i.e., the input features being (R, G, B, R^2, G^2, B^2), instead of the color histogram.

III. PLANNING WITH LONG RANGE DATA

Once the terrain in the far field has been classified, it must be effectively utilized by a planner to contribute to improved system performance. Our approach uses a polar-perspective map representation, which implicitly accounts for stereo range error, to make use of data with small stereo disparities. This allows long-range data to be projected into a cost map on which standard graph search techniques can be used to find a best path. Once found, the path is followed by a path following vehicle controller.

A. Map Representation

To incorporate long-range stereo data into a map that can be searched for a path, the semi-persistent, configuration-space grown local Cartesian map and long-range data are fused into a short-term polar-perspective map. The polar-perspective map is a polar map with an origin fixed to the vehicle frame and consists of cells with a fixed angular resolution, but variable range resolution. The range resolution corresponds to stereo disparity, proportional to inverse range, and consequently accounts for stereo range error by accumulating all points that lie within the expected stereo range error. For practical reasons of accumulating enough points in cells close to the origin, the map actually uses a constant range resolution out to a fixed distance, at which point it uses a variable range resolution. For these experiments, the map used consisted of 180 angular cells, covering 180 degrees, and 80 radial cells, covering out to 8m with a constant 20cm cell depth, and out to 64m with a cell size proportional to inverse range.

Because the local pose and near-field stereo-based terrain information is reasonably accurate compared to the local map cell sizes, a local Cartesian map can be maintained over time. However, due to accumulated error in the vehicle pose, this map cannot be kept forever. As a result, a sliding local map is maintained and used to generate a polar-perspective map on each planning cycle. On the other hand, because long-range data has a high uncertainty in range and low uncertainty in angle, it can only be kept for a very short time period. Consequently, only the long-range data from several previous frames is used to populate the polar-perspective map.

B. Path Planning

The polar-perspective map represents a graph, like a Cartesian map, and can similarly be searched for a shortest path with standard methods. However, the non-constant distance between cells must be accounted for by the search algorithm. Furthermore, because of the non-penetrating nature of visible imaging and the low range resolution of cells at long range,

the relative cost of lateral and longitudinal motions must be addressed. For instance, if lateral and longitudinal motions were equally weighted, the optimal plan to a waypoint behind a wall in the far-field would travel beyond the wall, which, even though it only occupies the length of one cell, would be well past the waypoint due to the cell size at that range. However, if the cost of lateral motions are kept small, then the optimal plan will initially travel to the side of the wall, but can then travel into the cells being occupied by it. To improve search performance, a look-up-table of costs between four-way connected cells is maintained. For cells beyond 10m, the lateral cost is reduced by a constant factor of 0.5. The standard Dijkstra shortest-path algorithm is used to obtain the best path, and can be computed in approximately 20ms on the onboard 2GHz Intel Core 2 Duo CPU. Once a path has been found, it is followed with a model-based controller, which controls the vehicle's left and right wheel velocities.

IV. EXPERIMENTAL TESTING AND RESULTS

A. System Evaluation Methodology

Quantitatively measuring long-range terrain classification rates and improved system performance based on continuous learning is a challenging task. Ground truth data is difficult and tedious to create, particularly in natural unstructured terrain, and running the entire system enough times over enough courses to statistically show improved performance is extremely time consuming and plagued with the problem of a useful performance metric (the time to achieve the goal is generally not sufficient; one must account for the risk taken by the vehicle).

So, to address these issues, we approach the problem of evaluating the system in several ways. First, we evaluate the long-range terrain classifier both on near-field stereo-classified terrain, as well as split the near field into a near and mid-field, and train the classifier only on the near data, and evaluate it on the mid-field. While this does not strongly show the generalization of the near-field to the far-field, it shows the effectiveness of the learning mechanisms used. Second, we qualitatively evaluate the stability of the classification in the far-field during an approach to obstacles. If the classification is stable during the approach, to the point when it enters the near field, and the obstacles are then classified correctly according to the near-field classification, this indicates that the classifier generalizes reasonably well. Third, we evaluate the distance at which the classifier can correctly label known obstacles. And lastly, we compare the plans generated and executed when populating a map with learned long-range obstacles with both a map with no long-range data and a map populated with fixed heuristic on long-range stereo data. The heuristic considers any tall objects in the distance as an obstacle; a threshold that is a function of range is applied to the variance of height data beyond a specific range to label long-range obstacles.

B. Sub-System Evaluation

To verify the learning machinery implemented, we start by evaluating it only on data where we already have terrain

TABLE I

CLASSIFICATION RATE OF CORRECTLY CLASSIFIED PIXELS WHEN USING THE NEAR SET TO TRAIN AND THE MID SET TO PREDICT (DATA FROM ARROYO SECO PARK SHOWN IN FIGURE 2

	Run 1		Run 2	
	Near	Mid	Near	Mid
SVM ($RGBR^2G^2B^2$)	0.74	0.73	0.83	0.91
SVM (histogram)	0.81	0.86	0.88	0.93
SVM (histogram, $a = 0.5$)	0.85	0.75	0.90	0.93

TABLE II

CLASSIFICATION RATE OF CORRECTLY CLASSIFIED PIXELS FOR THE 4 RUNS IN LAGR TEST 21 (TERRAIN SHOWN IN FIGURE 4), WITH COLOR HISTOGRAM FEATURES

	Run 1	Run 2	Run 3	Run 4
Trained with Run 1	0.95	0.92	0.90	0.91
Trained with Run 2	0.93	0.95	0.85	0.90
Trained with Run 3	0.91	0.88	0.96	0.82
Trained with Run 4	0.88	0.90	0.89	0.96

classification from the near-field classifier. The classified data is split into a set of pixels with a range of less than 5 meters (near-field), and another set with ranges of greater than 5 meters (mid-field). The long-range classifier is then trained only on the near set and evaluated on both the near and mid set separately. The near-field classification rate indicates how well the SVM model fits the training data, while the mid-field classification provides an indication of the range-extension performance. The evaluation results for a run in the Arroyo Seco park (seen in Figure 2) are shown in Table I. This data set covers a distance of 120 meters in 140 seconds. It is split into two sets: Run 1 from 0-50 meters and Run 2 from 50-120 meters. The color histograms are compared to the pixel color and squared color as a reference. For the Run 2 data, all three SVMs perform the range extension well, with rates even better than rates with training data. This may indicate that the mid-range pixels are easier to classify than the near-range pixels because the colors and textures in that region are more homogeneous, but the exact reason is still under investigation.

Another approach to evaluate the classifier is to train it in batch mode on one of several runs on the same course. An example of this on data from the LAGR Test 21 course (shown in Figure 4) is given in Tables II and III. Here, we collect 4000 data samples at all ranges for four runs. Table II shows the performance of the classifier with color histogram features, while Table III shows it for color features. On average, the histogram SVMs have slightly better performance than the color feature SVMs. This is expected because the color histograms accumulate colors over a window and capture local color variation, and are consequently more discriminatory than the single pixel-based color values. There is little variation of performance between runs, indicating that the environments of the four runs are similar.

	Run 1	Run 2	Run 3	Run 4
Trained with Run 1	0.91	0.89	0.92	0.90
Trained with Run 2	0.91	0.89	0.92	0.91
Trained with Run 3	0.90	0.88	0.92	0.90
Trained with Run 4	0.90	0.88	0.92	0.91

Fig. 4. The original image (left column) and its long-range classification (right column), showing traversable (blue) and non-traversable (red) classification of the grove of trees on Test 21 from varying distances (∼30m, ∼15m, ∼7m)

To evaluate the range at which obstacle can be detected by the learned classifier, a course with an easy to identify artificial obstacle of known size was used. The obstacle was made of three 4x8 foot plywood sheets covered in blue material. The vehicle was driven up to the obstacle such that it was in the near-field, and could learn its appearance. Then, the distance at which the obstacle was detected was measured. The obstacle could be detected at all distances closer than 100 meters. Figure 3 shows the obstacle classified at approximately 100 meters and 25 meters. However, the obstacle can only be put into the map when it has some stereo disparity, which was measured during the run at approximately 30 meters.

Fig. 3. The long-range classified image (left column) showing traversable (blue) and non-traversable (red) terrain, a zoomed in view of the classification (center column) and original image (right column) at 100m (top row) and 30m (bottom row)

A similar range of detection can be seen in the results from a real test (LAGR Test 21) through a meadow and approaching a grove of trees. Figure 4 shows the detection of the grove at various distances. More importantly however, these and other tests indicate the stability of classification over distances ranging from 100 meters to 5 meters.

To evaluate the effect of planning with long-range data, we compared the maps and plans generated from learned long-range perception, a fixed height heuristic, and no long-range perception. Figure 5 shows the difference in the map and plan between using a fixed heuristic on height data in the far-field compared to learned terrain classification. The figure is illustrative of several problems with the height heuristic. Because it uses a fixed threshold, not all long-range obstacles are picked up (such as the bushes on the right side of the trees), and because no color information is used, stereo matches on the sky line can cause false obstacles (seen on the left side of the map).

C. System Evaluation

The long-range learning system developed has been fielded on the standard, government provided LAGR vehicle [6]. We have tested it extensively in outdoor terrain consisting mainly of dirt paths and grass fields with traversable and non-traversable brush and grass, along with bushes, trees, logs, and other obstacles. As shown in Figures 1 and 5, the system is capable of learning long-range terrain types and is clearly planning to avoid obstacles in the far field. Figures 6 and 7 show the paths taken by the system with and without long-range learning in the Arroyo Seco and Balboa parks in the Los Angeles, CA area. The terrain in the two parks is very different in appearance (brush with dirt versus trees with grass), but the trained system is able to avoid the distant obstacles.

More telling, however, is the fact that the system was evaluated by an independent test team as part of the LAGR program's monthly field tests. The system was fielded and tested on the LAGR Test 21 course in Maryland during the summer. The course consisted of a start point in the middle of a large meadow and a goal point in an open area beyond a large grove of trees around a ditch. If the system was able to perceive the trees as an obstacle, it would veer right around the trees; if not, it would drive into the trees and need to find a path to the right around the ditch and through tall grass and trees. Images of the approach to the trees can be seen in Figure 4. The path of our system when no long range perception was used is compared to the path when online learning was used in Figure 8. In the learning run, online learning was used and an operator drove the vehicle around a first set of trees so that it could learn their

Fig. 5. The left and right stereo camera views (top row) and a 3D view of the map generated by the height heuristic (middle row) and long-range learning (bottom row); terrain is labeled as traversable (blue), non-traversable long-range (yellow), short-range non-traversable (red), and unknown (green), and the best path is shown in white. With height alone, the bushes to the right of the trees are not detected and incorrect stereo matches cause false obstacles on the left.

appearance, and then to the start point of the previous run. Here it was switched into autonomous mode and drove to the right, circumnavigating the tress. The run with no long-range perception drove up to the trees and was trapped in high grass in the process of finding a route to the goal.

V. CONCLUSIONS AND FUTURE WORK

A. Conclusions

We have implemented and fielded a system capable of learning and using long-range traversability to improve autonomous navigation performance. Comparing various features and classifiers, we have settled on using color-based features and a linear SVM, which has proved to be effective in learning from near-field terrain and practical for a real-time implementation. Using a sliding window of features, learning continuously adapts to the current terrain and en-

Fig. 6. An overhead map view of the paths taken ($\sim$50m long) in the Arroyo Seco park with no long-range perception (red) and with long-range learning (green)

vironmental conditions. To effectively utilize the long-range data, we use its stereo disparity to project it into a polar-perspective map, which can then be used to bias plans around obstacles in the far-field. The map representation enables long-range stereo data to be used by implicitly accounting for its stereo error. The system implemented has been tested extensively in outdoor, unstructured terrain, and evaluated in several ways by an independent test team. The system with learning shows the ability to avoid obstacles at long ranges, which results in a significant improvement in the time and a reduction of the risk taken to achieve goals.

B. Future Work

Despite some success with the fielded system, our experience with long-range learning has revealed several problems that are yet to be addressed. One problem with using a polar-perspective map is that long-range data cannot be accumulated for a very long time, and consequently the vehicle can "forget" about long-range obstacles out of its field of view. This leads the robot to potentially oscillate around two paths if there is one large obstacle in the distance. Furthermore, because the features used do not explicitly take into account variation due to range or intensity, if an obstacle looks very different when viewed from near and far, it can
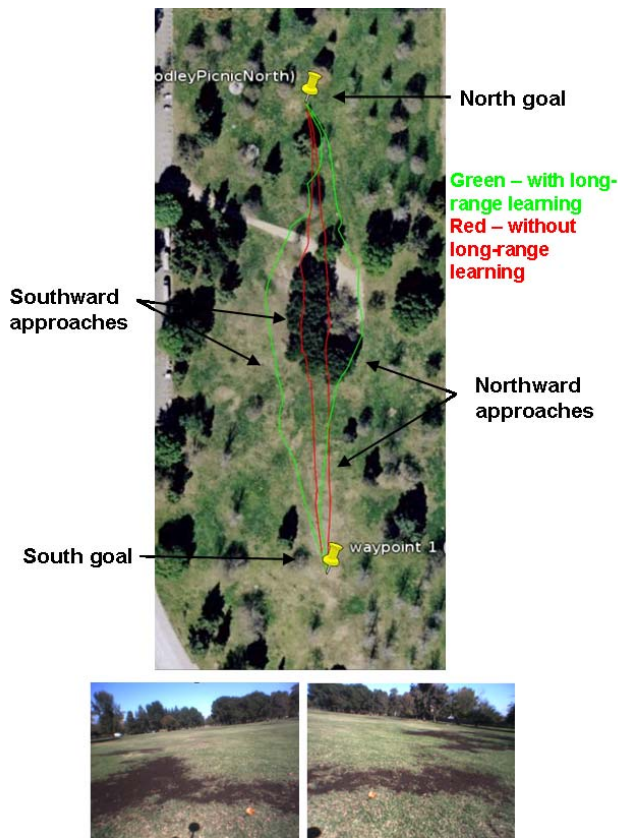
make learning very difficult. This is typical of more complex scenes, particularly in urban settings. This can be addressed by either using features that account for range or by maintaining registered features over many ranges. Investigating other features may also be beneficial to capitalize on scene texture beyond what is used by the current implementation.

VI. ACKNOWLEDGMENTS

Fig. 7. An overhead map view of the paths taken (∼100m long) in Balboa park with no long-range perception (red) and with long-range learning (green)

Fig. 8. An overhead map view of the paths taken on Test 21 with no long-range perception (yellow) and with long-range learning (blue); the path in the bottom half of the image shows where the system was trained, the starting point of the runs, and the goal (which is not achieved by the first run); the training to goal distance is ∼50m.

REFERENCES

[1] O. Chapelle, P. Haner, and V.N. Vapnik, Support vector machines for histogram-based image classication. IEEE Transactions on Neural Networks, vol. 10, 5, 1999.

[2] H. Dahlkamp, A. Kaehler, D. Stavens, S. Thrun, and G. Bradski, "Self-supervised monocular road detection in desert terrain". *Robotics: Science and Systems Conference*, 2006.

[3] R. Hadsell, P. Sermanet, J. Ben, A. Erkan, J. Han, U. Muller, and Y. LeCun, "Online Learning for Offroad Robots: Spatial Label Propagation to Learn Long-Range Traversability". *Robotics: Science and Systems Conference*, 2007.

[4] M. Happold, M. Ollis, and N. Johnson, "Enhancing supervised terrain classification with predictive unsupervised learning". *Robotics: Science and Systems Conference*, 2006.

[5] A. Howard, M. Turmon, L. Matthies, B. Tang, A. Angelova, and E. Mjolsness, Towards Learned Traversability for Robot Navigation: From Underfoot to the Far Field. *Journal of Field Robotics*, vol. 23, 11-12, 2007, pp 1005-1017.

[6] L.D. Jackel, E. Krotkov, M. Perschbacher, J. Pippine, and C. Sullivan, The DARPA LAGR Program: Goals, Challenges, Methodology, and Phase I Results. *Journal of Field Robotics*, vol. 23, 11-12, 2007, pp 945-973.

[7] D. Kim, J. Sun, S. Oh, J. Rehg, and A. Bobick, "Traversability classification using unsupervised on-line visual learning for outdoor robot navigation". *International Conference on Robotics and Automation*, 2006.

[8] D. Lieb, A. Lookingbill, and S. Thrun, "Adaptive road following using self-supervised learning and reverse optical flow". *Robotics: Science and Systems Conference*, 2005.

[9] R. Manduchi, A. Castano, A. Talukder, and L. Matthies, Obstacle detection and terrain classification for autonomous off-road navigation. *Autonomous Robots*, vol. 18, 2003, pp 81102.

[10] L. Matthies, M. Turmon, A. Howard, A. Angelova, B. Tang, E. Mjolsness, "Learning for autonomous navigation: extrapolating from underfoot to the far field." *NIPS Workshop on Machine Learning Based Robotics in Unstructured Environments*, 2005.

[11] M.W. Otte, S. Richardson, G. Grudic, J. Mulligan, "Local Path Planning in Image Space for Autonomous Robot Navigation in Unstructured Environments". *International Conference on Intelligent Robots and Systems*, 2007.

[12] M. Procopio, J. Mulligan, G. Grudic, "Long-Term Learning Using Multiple Models for Outdoor Autonomous Robot Navigation". *International Conference on Intelligent Robots and Systems*, 2007.

[13] B. Schölkopf and A.J. Alexander, Learning With Kernels, *The MIT Press*, Cambridge, MA, 2002.

[14] B. Sofman, E. Lin, J. Bagnell, N. Vandapel, and A. Stentz, "Improving robot navigation through self-supervised online learning". *Robotics: Science and Systems Conference*, 2006.

[15] M.J. Swain and D.H. Ballard, "Indexing Via Color Histograms". *International Conference on Computer Vision*, 1990.

[16] B. Tang and D. Mazzoni, "Multiclass Reduced-Set Support Vector Machines". *International Conference on Machine Learning*, 2006.