

Obstacle Detection from Overhead Imagery using Self-Supervised Learning for Autonomous Surface Vehicles

Hordur K. Heidarsson and Gaurav S. Sukhatme

Abstract—We describe a technique for an Autonomous Surface Vehicle (ASV) to learn an obstacle map by classifying overhead imagery. Classification labels are supplied by a front-facing sonar, mounted under the water line on the ASV. We use aerial imagery from two online sources for each of two water bodies (a small lake and a harbor) and train classifiers using features generated from each image source separately, followed by combining their output. Data collected using a sonar mounted on the ASV were used to generate the labels in the experimental study. The results show that we are able to generate accurate obstacle maps well-suited for ASV navigation.

I. INTRODUCTION

Autonomous surface vehicles (ASVs) are useful for a variety of tasks in lake and harbor environments. They can be useful tools for environmental monitoring [1], rescue/recovery missions [2], and environmental cleanup [3].

An important ability of an autonomous vehicle is to be able to plan safe and efficient paths through an environment with obstacles. For dynamic obstacles, the vehicle needs to be able to sense its surroundings and react appropriately. Static obstacles can be avoided using a map of the environment. However, maps do not exist for all locations, they can be time consuming to make, and they have to be kept up to date. Further, the degree to which an object or map feature is an obstacle depends on the vehicle. One way to satisfy the need for a current map is to automatically build one from a recent aerial or satellite image of the operating area using classification techniques to estimate the location of obstacles [4]. A problem with this approach is the need for labeled training data. This requires human labor and the resulting classifier might not generalize well to areas outside of the training set because of different terrain or obstacles in the environment. To eliminate the need for manually labeled training data and for better generalization between environments, Sofman et al [5] combined data from overhead imagery and local sensing from a ground robot to estimate traversal costs in an environment.

In this paper we apply this approach to an ASV and perform obstacle detection from overhead imagery by combining data from an on-board sonar with the overhead imagery, in order to build a obstacle map of the environment to use for trajectory

H. K. Heidarsson is with the Robotic Embedded Systems Lab and the Department of Electrical Engineering, University of Southern California, Los Angeles, CA 90089, USA heidarss@usc.edu

G. S. Sukhatme is with the Robotic Embedded Systems Lab and the Department of Computer Science, University of Southern California, Los Angeles, CA 90089, USA gaurav@usc.edu



Figure 1. The USC ASV. Left: The ASV in operation in Echo Park Lake. Right: The ASV on shore with the sonar mounted.

planning and safe velocity control. Compared to [5] we work with a different type of sensor and a different environment. We also incorporate spatial dependencies into our classification. While the classification process is ultimately intended to run online on an ASV, we have chosen not to focus on that in this work but rather investigate the feasibility of the general approach in this domain.

In our previous work [6] we explored the feasibility of using a forward mounted profiling sonar on an ASV for obstacle detection. We investigated the effective range and ability of this setup and found it to be a useful option for obstacle detection for ASVs. During this previous research, we have collected large amounts of forward facing profiling sonar data (with GPS coordinates and compass heading for each measurement). We collected two sets of data, one in Echo Park Lake, Los Angeles and the other in King Harbor, Redondo Beach, both of which are utilized in this paper.

II. PRIOR WORK

Many ASVs have been designed and developed in recent years, some are equipped with radar or cameras for obstacle sensing or mapping [7], [8], but none use sonar for this purpose until recently [6], [9]. ASV literature that specifically deals with the obstacle avoidance problem focuses on obstacle detection for above-surface obstacles. Larson et al. [10] performed ASV navigation using vision and radar, Bandyopadhyay et al. [11] used a laser range finder to perform obstacle avoidance in the Port of Singapore and Snyder et al. [12] performed autonomous navigation in rivers using a camera system. Benjamin et al. broadcasted positions of ASVs to avoid collisions between them [13] without any specific sensing. Most recently, Huntsberger et al. [14] developed stereo-vision based navigation.

Also in related ASV perception/navigation work, Wolf et al. [15] have investigated vision-based methods for target tracking

on the water surface. Martins et al. use vision to dock with an AUV [16]. Leedekerken et al. [9] perform mapping on an ASV by combining data from LADAR above surface and sonar below surface.

Work on enabling long range sensing using aerial imagery includes [17], [18], [4], but the most closely related work to ours is the work by Sofman et al. [5], where a Bayesian approach is used to perform self-supervised online learning of path costs in off-road environment using aerial imagery in conjunction with LADAR on a ground vehicle. While our work involves a similar idea, we work in a different domain, with a different sensor using different tools.

III. APPROACH

The approach consists of four parts. First the collection and processing of the sonar data, next the processing and feature extraction of the overhead imagery, then combining the two in a binary classification framework and predicting the remainder of the map and finally smoothing and combining the results from the earlier classifications and generating a map with three classes: *obstacle*, *transient/suspect obstacle*, *non-obstacle*. A path planner working with the resulting map would plan paths through the non-obstacle space while avoiding obstacles. The planner would treat the transient obstacles as something it might actually be able to traverse through, but with caution and have some reactive sensing/avoidance handle the obstacles if needed.

We discretize the map to a grid with 1 x 1 m sites. The features for each cell are generated from the collection of pixels within each site and the labeling and classification is done in this 1 x 1 m grid.

A. Sonar Data Processing

The sonar data processing is a modified version of our previous work [6]. A sonar measurement consists of a vector where each element represents echo strength measured from a certain distance. The vector elements are spread equally over the set range of the sonar, which can be set from 1-100 m. In our experience, the effective range when mounted forward is 30-40 m. The sonar data processing extracts potential obstacles from the measurements. The sonar processing pipeline is capable of extracting multiple obstacles from a single measurement, but now we discard all data past the first obstacle found in each measurement, if any. This prevents mis-labeling of anything that is beyond the first obstacle due to reflections. To account for errors in obstacle location, caused by ASV localization error as well as compass error, we draw a 1 m radius circle around each obstacle we find and label each site within that circle as an obstacle.

It is important to note that the sonar measurements will in great likelihood produce inconsistent labels since each site is usually scanned multiple times by the sonar; the first scan could be from a distance, over which a particular obstacle is not visible, and then, once closer, another scan is performed in which the obstacle appears. We use a simple, and perhaps naive, strategy for resolving these labeling inconsistencies: detected obstacle always takes precedence over non-obstacle,

i.e. once we've detected an obstacle in a site, we mark that site having an obstacle, no matter if we get a non-obstacle measurement later.

B. Overhead Imagery Processing and Feature Extraction

As previously stated, we have split the area of interest into 1 x 1 m sites. From overhead imagery, we extract features for each of these sites. We want the features to capture some statistics about the color, texture and structure of the site. Based on literature in texture classification and classification of aerial imagery [19], [20], we have chosen the following set of features: image pixel values in two color spaces (HSV and CIELab), gabor energy features and entropy within sites.

The RGB color model, in which digital images are usually represented, is unsuitable here because of its vulnerability to changes in lighting. Instead we use the HSV color representation which is better suited for perception of the images. For further perceptual uniformity, we also use the CIELab color space too. For each channel in these two color models we take the average intensity value as well as the standard deviation of the intensity values for each site and use that as a feature.

Gabor filters are a family of filters frequently used for edge detection [21] and texture classification [22]. The real component of the gabor filter is formulated as follows:

$$g_{\lambda,\theta,\psi,\sigma,\gamma}(x,y) = \exp\left(-\frac{x'^2 + \gamma^2 y'^2}{2\sigma^2}\right) \cdot \cos\left(2\pi\frac{x'}{\lambda} + \psi\right)$$

where $x' = x \cos \theta + y \sin \theta$ and $y' = -x \sin \theta + y \cos \theta$. For our purposes we use a filter of bandwidth 1, resulting in $\sigma \approx 0.56\lambda$, and aspect ratio $\gamma = 0.5$. We use the filter at 6 different equi-spaced orientations and at 3 different spatial frequencies (defined by $\lambda = \{1, 9, 14\}$), resulting in a filter bank of 18 filters. The filter bank is used to generate so called Gabor energy image by taking the L_2 -norm of the responses from convolution of the L channel of the whole CIELab image with the symmetric and anti-symmetric Gabor filters ($\psi = 0$ and $\psi = \pi/2$, respectively) at each pixel. The gabor energy features we use are then generated by calculating the average value of this response of all the pixels within each site.

In order to further characterize the texture within each site, S_i , we calculate entropy of the pixel intensities within each site from the H channel of the HSV image, $H(S_i) = -\sum_{x \in S_i} p(x) \cdot \log(p(x))$, as one of the features.

C. Classification

The classification process utilizes the Real AdaBoost boosting algorithm [23], a generalization of the basic AdaBoost algorithm by Freund and Schapire [24]. Other classifiers do work as well, but we have chosen AdaBoost for its simplicity, performance and quick run times. There have also been developed online variants of the algorithm [25], in which we are interested for the further development of this work.

Boosting allows for combining multiple weak learners to generate a single strong one. The AdaBoost algorithm works the following way:

Assume we are given a training set $\{(f_1, y_1), \dots, (f_m, y_m)\}$ where $f_1, \dots, f_m \in \mathcal{F}$ is a set of features and y_1, \dots, y_m ,

$y_i \in \{-1, 1\}$, is a set of labels corresponding to the feature vectors. We assign equal weights to all the training samples. Next, for $n = 1, \dots, N$, we train a weak learner based on the weighted training set and get a weak hypothesis, $h_n : \mathcal{F} \rightarrow \{-1, 1\}$, we then calculate a weight, α_n , associated with the hypothesis based on the error rate of the hypothesis. The final step of the loop is to update the feature set weights so that incorrectly classified training examples get more weight when training the next learner.

Finally, the resulting strong classifier, generating labels $c \in \{-1, 1\}$, is as follows:

$$H_N(f_i) = \text{sign} \left(\sum_n \alpha_n h_n(f_i) \right)$$

The training is described in more detail in [23]. The weak learners used in our work are Classification and Regression Trees (CART) [26] with branching factor of 2, from [27].

As a measure of confidence of the classification we use the absolute classification margin, $F(f_i) = \sum_n \alpha_n h_n(f_i)$, to generate a posterior probability using logistic correction [28]: $p(c_i = 1|f_i) = \frac{1}{1 + \exp(-2F(f_i))}$. We will use this measure as an input to our smoothing process.

D. Smoothing

In the classification process we ignored the spatial relationship between neighboring sites. In environments we are interested in, a certain class almost always forms sizable continuous patches on the map, i.e. there are rarely spots of the size of a few sites. When we get such small spots we would like the feature-set to give a strong support for that labeling. Spots with low support that are still misclassified, due to insufficient training data set coverage, artifacts introduced by image stitching and image noise, are something we want to remove. In order to do that we will incorporate the classification results along with the feature set in to a Markov Random Field (MRF) framework. MRFs are one of the tools of the trade in computer vision and allow one to represent spatial dependencies between sites/pixels.

The problem setup is a conventional MRF image denoising problem where the observations come from our classifier but instead of having homogeneous edge weights we generate the weights based on the similarity of the features of the neighboring sites. We set this up as an energy minimization problem [29] where we assign energy to the edges and nodes. The total energy to minimize is:

$$E(c') = \sum_{i \in \mathcal{P}} D_i(c'_i) + \sum_{i,j \in \mathcal{N}} V_{i,j}(c'_i, c'_j)$$

where \mathcal{P} is the set of all sites within the image, \mathcal{N} is the set of all pairs of interacting sites (i.e. all pairs of sites that are in each other's neighborhood), $D_i(\cdot)$ is the data penalty function and $V_{i,j}(\cdot)$ is the interaction potential between neighboring sites and c'_i is the labeling for site i . The site neighborhood size was chosen after testing to be 4 (8-neighborhood did not deliver substantial improvements).

The data penalty function takes into account the confidence of the label assignment from the previous classification step:

$D_i(c'_i) = -\log(p(c_i = c'_i|f_i))$, where $p(c_i = c'_i|f_i)$ is the posterior probability from the previous classification step and, as earlier, f_i is the set of features for site i . Label assignments that agree with the previous classification step will therefore generate lower data penalty than those that disagree. The interaction potential is supposed to favor spatial coherence by penalizing neighboring nodes with different labels. However, when the features for neighboring sites are very different we want the penalty to be smaller than in cases where features are similar, so we incorporate the distance between the feature vectors in to the interaction potential, $V_{i,j}(c_i, c_j) = \beta \frac{|c_i - c_j|}{|f_i - f_j|^2}$, where β is a parameter. The tuning of β was done by hand for this work.

To find the optimal labeling, c , we minimize the total energy by calculating the min-cut/max-flow using the Boykov-Kolmogorov algorithm [29].

E. Transient Obstacle Detection

We perform the classification process on the aerial images captured at different time instances. The added dimension of time allows us to address non-static obstacles as well as help with addressing noise such as shadows cast from trees and structures that might be predicted to be obstacles by our classifier. Our resulting obstacle map will therefore have three classes for each cell: *obstacle*, *transient/suspect obstacle* and *non-obstacle*.

To combine the results from the two classifications we simply take the intersection of the obstacles from both predicted maps and label those sites as obstacles. The difference between the two maps will be labeled as transient/suspect obstacle and finally the intersection of the non-obstacles will be labeled as non-obstacles. If we assume the labels to be *non-obstacle*, l_n , *obstacle*, l_o , and *transient/suspect*, l_t , the resulting label would be:

$$\text{label}(c_1, c_2) = \mathbf{1}_{c_1=c_2} \cdot c_1 + \mathbf{1}_{c_1 \neq c_2} \cdot l_t$$

where c_1 and c_2 are the predicted class values from each of the classifications.

IV. EXPERIMENTAL RESULTS

We have experimental results from the two aforementioned areas. Echo Park Lake (Fig. 2) is a small man made lake in Los Angeles, CA. Apart from the shoreline, the lakes obstacles include floating islands and a small dock. The second area is King Harbor (Fig. 3), a marina in Redondo Beach, CA. In the marina we have floating docks, docked and moving boats and buoys. The two experiments were done independently and no data was shared between them. The data processing and obstacle map generation for these experiments was done offline after the data gathering.

A. Experimental Platform

The experimental platform used is the Autonomous Surface Vehicle (ASV) designed by the University of Southern California's Robotic Embedded Systems Laboratory (Fig. 1). The ASV is an Ocean Science QBoat-I hull, 2.1 m long and 0.7



Figure 2. Aerial images of Echo Park Lake, Los Angeles. Left: Image from Google Maps, Right: Image from Bing Maps. Apart from the difference in hue and lighting, the main thing to notice is the different direction of shadows between the images.

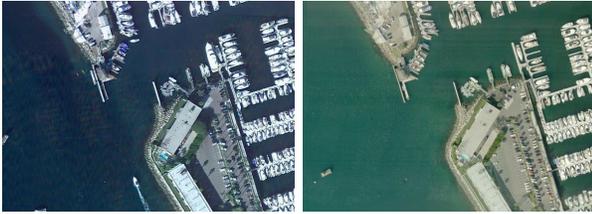


Figure 3. Aerial images of King Harbor, Los Angeles. Left: Image from Google Maps, Right: Image from Bing Maps. Apart from the difference in hue and lighting, the main thing to notice is the difference in transient obstacles present in the images.

m wide at the widest section. The ASV is actuated by two electrical motors and a rudder and is capable of speeds up to 1.6 m/s. The ASV has an on-board computer, GPS, an IMU and a compass. The ASV is controlled by software built using the open-source framework Robot Operating System (ROS).

For the trials described in this paper, the ASV was equipped with an Imagenex 881L Profiling Sonar mounted facing forward scanning in a plane parallel to the water surface. The 881L is a single-beam mechanically scanned multi-frequency sonar with a full scale range from 1 m to 100 m. Along with the sonar measurements, GPS position and compass heading was also recorded.

B. Sonar Data Labels

The ASV was driven manually around the two areas to collect sonar measurements. The sonar data was processed as described above and a small part of it selected for use as labels for the classification process. Figure 4 shows the labels generated by the sonar overlaid on top of aerial images of the areas.

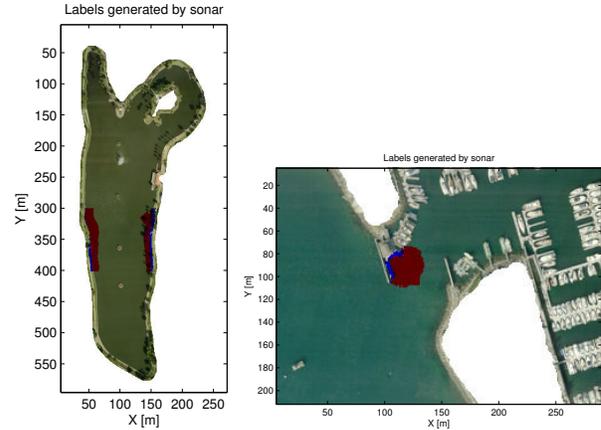


Figure 4. Classification labels for the area of interest, overlaid on aerial images of the area. Obstacles are represented by blue and non-obstacles by red. Top: Echo Park Lake; Bottom: King Harbor. In these images, and all following, the white area has been manually masked out of the image as it is out of the domain of operation.

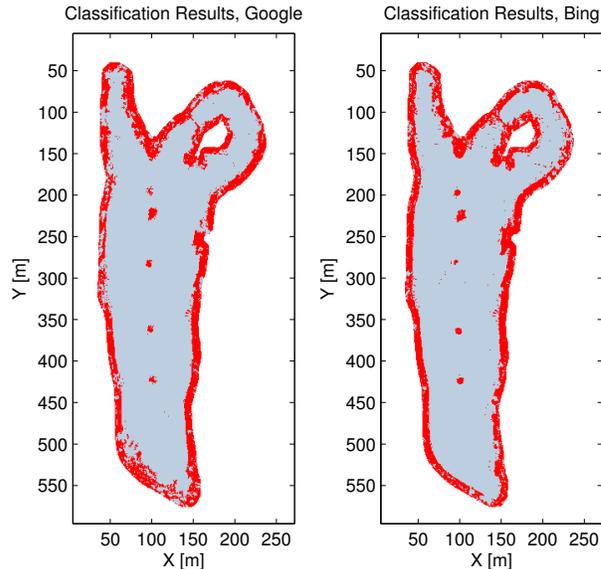


Figure 5. Classification results from Echo Park Lake for each image source. The area outside the lake has been filled white since it is not of any interest. Left: Google, Right: Bing. Obstacles are red, non-obstacles are gray.

C. Classification Results

The classification process was run on images from Google Maps and Bing Maps and the results were then smoothed individually and finally combined as described above. Two image sources are used in order to get the added dimension of time, but because the two image sources are quite different in terms of color and illumination we actually train a classifier independently for each of them instead of using the same one for both as was suggested in the previous section. We have calculated the error rates of the classification by comparing the results to a manually labeled images.

1) *Echo Park Lake*: The resulting obstacle maps are displayed in fig. 5. One can see that the lake shore is correctly classified as obstacle as well as the floating islands and the

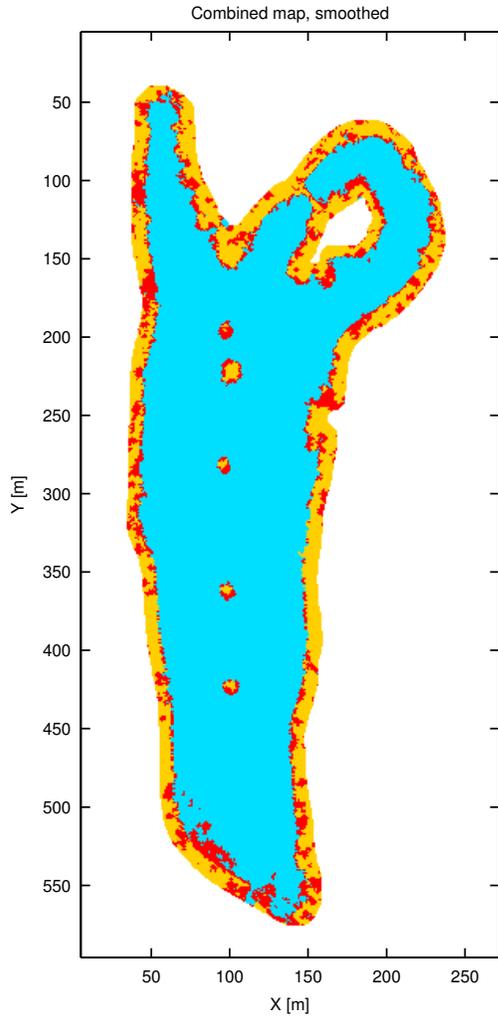


Figure 6. Combined obstacle map for Echo Park Lake. Obstacles are yellow, transient/suspect obstacles are red and non-obstacles are blue.

fountain. There is some confusion as to if trees extending over the lake and their shadows are obstacles, but the combined classification in fig. 6 improves on that somewhat. In this figure it can clearly be seen how this allows us to mark the shadows cast by trees at the south side of the lake as transient/suspect obstacles and not as obstacles as they were originally classified by one of the classifiers. One thing to notice is the bridge at the north end of the lake that is classified as an obstacle. Even if there was labeling support at that location, our method would still be hard pressed to classify the bridge as not being an obstacle. The classification error rates are displayed in table I.

2) *King Harbor*: The resulting obstacle maps are displayed in fig. 7. The two maps agree for the most part. We can see that all the docks are correctly classified as obstacles and the boats in the harbor as well. The two maps were then combined into one in fig. 8. Here we can clearly see the benefit of using multiple images as this environment is packed with dynamic obstacles, such as the moving boat coming in to the harbor, two kayaks and of course the different boats parked in the berths. The classification error rates are displayed in table I.

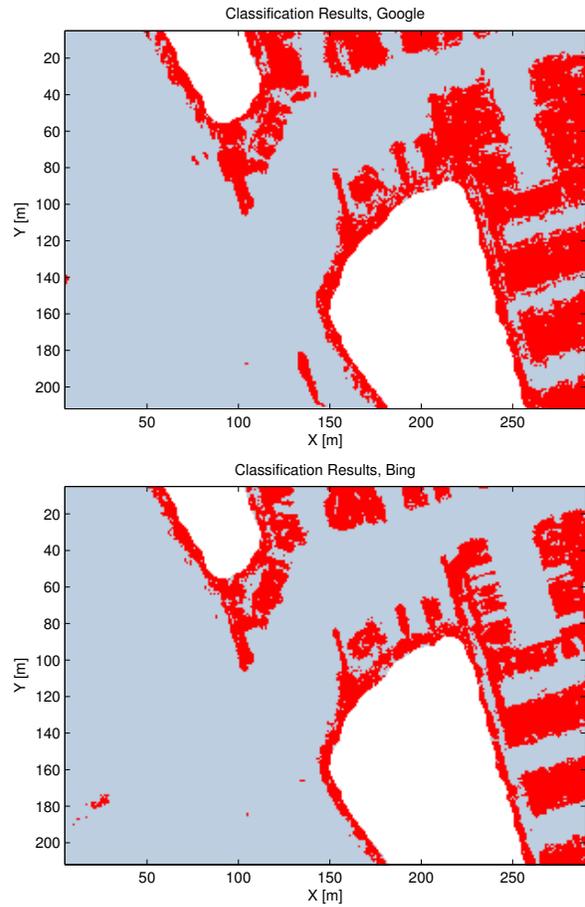


Figure 7. Classification results from King Harbor for each image source. Left: Google, Right: Bing. Obstacles are red, non-obstacles are gray.

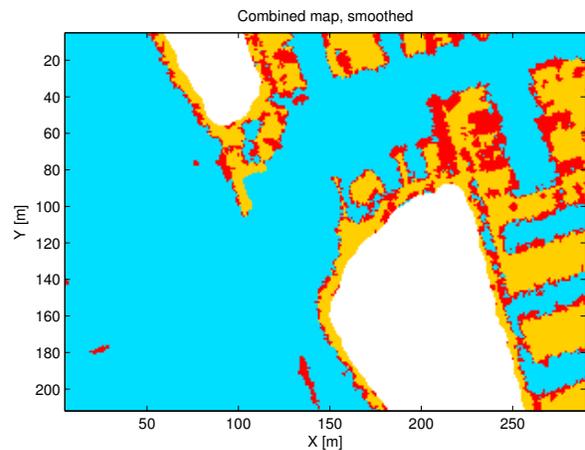


Figure 8. Smoothed and combined obstacle map for King Harbor. Obstacles are yellow, transient/suspect obstacles are red and non-obstacles are blue.

Image source	Not smoothed	Smoothed
Echo Park		
Bing	6.6%	5.7%
Google	8.3%	7.8%
King Harbor		
Bing	6.2%	6.1%
Google	8.6%	7.8%

Table I
CLASSIFICATION ERROR RATES

V. CONCLUSIONS AND FUTURE WORK

In this paper we have developed a method for obstacle detection from an overhead image using labels generated from a forward looking sonar attached to an ASV. The results indicate that this is a viable way to generate an obstacle map on the fly for use in path planning or velocity planning.

For future work we plan to address the multiple label problem and incorporate repeated sensor measurements (and a measure of confidence in the measurements) into the estimation process. Furthermore, we plan to frame the problem in a probabilistic framework. We also plan to address problems with occlusions such as bridges and trees. In this work, we have made training data available to the classification process based on location, i.e. we have made the training dataset consist of all labels in a certain region. We plan to make the data sequentially available to the classifier and look at the progression of the classification results. Finally, we plan to formulate this as an active classification problem and plan a path for the ASV with the objective of improving the classifier.

VI. ACKNOWLEDGEMENTS

We gratefully acknowledge help with deployments and hardware provided by Jnaneshwar Das, Carl Oberg, Arvind Pereira, Ryan Smith. This work was supported in part by the ONR Antidote MURI project (grant N00014-09-1-1031), the NOAA MERHAB program (grant NA05NOS4781228), NSF as part of the Center for Embedded Network Sensing (CENS) (grant CCR-0120778), and NSF grant CNS-1035866.

REFERENCES

- [1] G. S. Sukhatme, A. Dhariwal, B. Zhang, C. Oberg, B. Stauffer, and D. A. Caron, "Design and development of a wireless robotic networked aquatic microbial observing system," *Environmental Engineering Science*, vol. 24, no. 2, pp. 205–215, 2007.
- [2] F. Arrichiello, H. K. Heidarsson, S. Chiaverini, and G. S. Sukhatme, "Cooperative caging using autonomous aquatic surface vehicles," in *IEEE International Conference on Robotics and Automation*, May 2010, p. 4763–4769.
- [3] S. Bhattacharya, H. K. Heidarsson, G. S. Sukhatme, and V. Kumar, "Cooperative control of autonomous surface vehicles for oil skimming and cleanup," in *IEEE International Conference on Robotics and Automation*, May 2011.
- [4] R. Hudjakov and M. Tamre, "Aerial imagery terrain classification for long-range autonomous navigation," in *International Symposium on Optomechatronic Technologies (ISOT)*, 2009, p. 88–91.
- [5] B. Sofman, E. Lin, J. A. Bagnell, J. Cole, N. Vandapel, and A. Stentz, "Improving robot navigation through self-supervised online learning," *Journal of Field Robotics*, vol. 23, no. 11-12, pp. 1059–1075, 2006.
- [6] H. K. Heidarsson and G. S. Sukhatme, "Obstacle detection and avoidance for an autonomous surface vehicle using a profiling sonar," in *IEEE International Conference on Robotics and Automation*, Shanghai, China, May 2011, pp. 731–736.
- [7] M. Caccia, "Autonomous surface craft: prototypes and basic research issues," in *Mediterranean Conference on Control and Automation*, 2006, pp. 1–6.
- [8] J. Curcio, J. Leonard, and A. Patrikalakis, "SCOUT - a low cost autonomous surface platform for research in cooperative autonomy," in *Proceedings of MTS/IEEE OCEANS*, 2005, pp. 725–729 Vol. 1.
- [9] J. C. Leedeckerken, M. F. Fallon, and J. J. Leonard, "Mapping complex marine environments with autonomous surface craft," in *International Symposium of Experimental Robotics (ISER)*, Delhi, India, Dec. 2010.
- [10] J. Larson, M. Bruch, and J. Ebken, "Autonomous navigation and obstacle avoidance for unmanned surface vehicles," *SPIE PROC. 6230: UNMANNED SYSTEMS TECHNOLOGY VIII*, pp. 17–20, 2006.
- [11] T. Bandyopadhyay, L. Sarcione, and F. S. Hover, "A simple reactive obstacle avoidance algorithm and its application in singapore harbor," in *The 7th International Conference on Field and Service Robots*, 2009, p. 455–465.
- [12] F. D. Snyder, D. D. Morris, P. H. Haley, R. Collins, and A. M. Okerholm, "Autonomous river navigation," in *Proceedings of SPIE, Mobile Robots XVII*, D. W. Gage, Ed., vol. 5609, Dec. 2004, pp. 221 – 232.
- [13] M. Benjamin, J. Curcio, and P. Newman, "Navigation of unmanned marine vehicles in accordance with the rules of the road," in *Robotics and Automation, 2006. ICRA 2006. Proceedings 2006 IEEE International Conference on*, 2006, pp. 3581–3587.
- [14] T. Huntsberger, H. Aghazarian, A. Howard, and D. C. Trotz, "Stereo vision-based navigation for autonomous surface vessels," *Journal of Field Robotics*, vol. 28, no. 1, pp. 3–18, 2011.
- [15] M. T. Wolf, C. Assad, Y. Kuwata, A. Howard, H. Aghazarian, D. Zhu, T. Lu, A. Trebi-Ollennu, and T. Huntsberger, "360-degree visual detection and target tracking on an autonomous surface vehicle," *Journal of Field Robotics*, vol. 27, no. 6, pp. 819–833, 2010.
- [16] A. Martins, J. Almeida, H. Ferreira, H. Silva, N. Dias, A. Dias, C. Almeida, and E. Silva, "Autonomous surface vehicle docking manoeuvre with visual information," in *Robotics and Automation, IEEE International Conference on*, 2007, pp. 4994–4999.
- [17] D. Silver, J. A. D. Bagnell, and A. T. Stentz, "High performance outdoor navigation from overhead data using imitation learning," in *Robotics Science and Systems*, Jun. 2008.
- [18] D. Silver, B. Sofman, N. Vandapel, J. Bagnell, and A. Stentz, "Experimental analysis of overhead data processing to support long range navigation," in *Intelligent Robots and Systems, IEEE/RSJ International Conference on*, 2006, pp. 2443–2450.
- [19] S. Kluckner, T. Mauthner, P. M. Roth, and H. Bischof, "Semantic classification in aerial imagery by integrating appearance and height information," in *Proc. Asian Conference on Computer Vision (ACCV)*, 2009.
- [20] L. Yang, X. Wu, E. Praun, and X. Ma, "Tree detection from aerial imagery," in *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 2009, pp. 131–137.
- [21] R. Mehrotra, K. Namuduri, and N. Ranganathan, "Gabor filter-based edge detection," *Pattern Recognition*, vol. 25, no. 12, pp. 1479–1494, Dec. 1992.
- [22] S. Grigorescu, N. Petkov, and P. Kruizinga, "Comparison of texture features based on gabor filters," in *IEEE Transactions on Image Processing*, vol. 11, 2002, pp. 1160–1167.
- [23] R. E. Schapire and Y. Singer, "Improved boosting algorithms using confidence-rated predictions," *Machine Learning*, vol. 37, no. 3, pp. 297–336, Dec. 1999.
- [24] Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," in *Proceedings of the Second European Conference on Computational Learning Theory*, 1995, p. 23–37.
- [25] C. Leistner, A. Saffari, P. M. Roth, and H. Bischof, "On robustness of on-line boosting - a competitive study," in *IEEE International Conference on Computer Vision (ICCV), Workshops*, 2009, pp. 1362–1369.
- [26] L. Breiman, J. Friedman, R. Olshen, and C. Stone, *Classification and regression trees*. Wadsworth, 1984.
- [27] A. Vezhnevets. GML AdaBoost matlab toolbox, 0.3. [Online]. Available: <http://graphics.cs.msu.ru/science/research/machinelearning/adaboosttoolbox>
- [28] J. Friedman, T. Hastie, and R. Tibshirani, "Additive logistic regression: a statistical view of boosting," *The Annals of Statistics*, vol. 28, no. 2, pp. 337–407, Apr. 2000.
- [29] Y. Boykov and V. Kolmogorov, "An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 9, pp. 1124 – 1137, Sep. 2004.