

Extrinsic Calibration of a Monocular Camera for an Autonomous Surface Vessel Using a Forward Facing Sonar

Hordur K. Heidarsson and Gaurav S. Sukhatme

Abstract—We present work towards building a fused sensing system for a small Autonomous Surface Vessel (ASV). We present a technique to estimate the extrinsic calibration of a forward facing monocular camera by using existing range sensors on the vessel along with some improvements on segmenting the water surface in the camera images. We demonstrate our techniques in field experiments and apply the obtained camera calibration to experiment data.

I. INTRODUCTION

Autonomous surface vessels (ASVs) have the potential to become prominent tools in environmental surveying, monitoring [4], [15], [14] and mapping [11]. However, at present, they are not ready to be operated at all times without human interaction in busy environments such as marinas or lakes with boat traffic.

For practical long-term and cost-efficient environmental monitoring purposes, an ASV needs to be able to operate without supervision whenever the need arises and should be adaptable to different situations through reconfiguration of sensors and payload. An ASV that is to be operated at any time, i.e. day or night, and in variety of weather conditions (e.g. sunny, rain, fog) needs to have a robust sensing system that is able to operate in these different conditions. Furthermore, the sensing system needs to be multimodal since relying on a single type of sensor for detecting obstacles in this environment can result in serious errors.

Previously, we have investigated the application of a forward mounted sonar on an ASV for obstacle detection and avoidance [5], [6]. While that solution works, it has some drawbacks. When it detects a vessel, it is not able to detect what type the vessel is. This is a problem if the ASV is to follow the waterways rules of the road, COLREGS (e.g. sailboats have the right of way when the sails are up). That is something that cannot be detected with a sensor below the water surface. Another problem, that is perhaps more immediate, is that in some situations the sonar detects phantom obstacles. These are usually mixing layers in the water that cause the sonar beam to reflect with high strength in some cases causing a phantom obstacle that is very hard to distinguish from a real one based on the sonar data alone. To resolve these problems we examine the utility of adding an additional sensor to complement the sonar.

The available sensors to choose from are lasers, camera and radar. In order to keep cost and size down we have chosen to go with a camera. With a single camera, we need to obtain

scale information to be able to use the data. Fortunately, since we are operating on a plane, we can derive a mapping from the camera image to real-world coordinates on the water surface.

In this paper, we show an approach to compute this mapping from the camera image to the water surface coordinates by using the forward mounted sonar and segmentation of the water surface in the camera images. First we review some relevant related work in the field before introducing the water segmentation technique. Next, we discuss the camera calibration. Finally we show experimental results from field trials with an ASV in a lake.

II. RELATED WORK

A. Sensing for ASVs

The Jet Propulsion Laboratory has made significant advances in sensing for ASVs [7], [16], in particular with stereo and 360° vision. At MIT, Bandyopadhyay et al. [2] have built a laser range finder-based navigation system, and Leedekerken et al., with [11] have used a combination of lidar, radar and sonar to create maps of the environment both the above and under the surface of water. Larson et al. [10] combined vision, radar and nautical charts to model the world and do obstacle avoidance.


However, in contrast to some of the work done by [10], [7], we are dealing with a small ASV which has limited sensor payload capacity and we are also trying to maintain the ASV at a relative low cost. This makes some of the sensing techniques discussed in these papers less viable due to cost, weight or size.

B. Water segmentation and detection

Achar et al. [1] present an approach for self supervised segmentation of river scenes. The approach works very nicely in many situations, but can perform poorly in low-rise environments, i.e. when there is mainly sky above the horizon line. It has some problems with obstacles in the distance. We build on this approach in this paper. Rankin et al. [13] present approaches to water detection based on color variations and sky reflections. The approach is developed for use on surface vehicles and to detect water hazards rather than for an aquatic vehicle with which we are working.

C. Monocular camera calibration

For monocular camera calibration there is a lot of existing work. For general intrinsic calibration the most common



(a) A typical scene from a lake. (b) The ideal water segmentation of a). Blue is water, green is sky, gray is masked out and red are obstacles.



(c) Example of reflections on the water. (d) Example of bright reflections of the surface.

Figure 1: Water segmentation, examples and problems.

method is to use a planar calibration target [17], but there are also methods to do this by observing the structure surrounding the camera while it is moving [3], but the extrinsic or relative pose calibration depends on some external reference coordinate system that you calibrate to. Relative pose calibration can be achieved from data between a camera and an IMU [8].

III. WATER SEGMENTATION

The water segmentation problem is where the water surface in an image captured by an ASV is detected and labeled separately from other things in the image. This can prove difficult: the surface is moving and changing, so tracking methods fail; water has no structural features (corners, etc.); reflections on the surface of the water can be problematic. Examples of a segmented image and these problematic situations can be seen in Figure 1.

For this we extract some features from the image, split it into cells and then use a classifier to classify each cell independently and optionally refine the estimate using a CRF [9].

A. Feature Extraction

Similarly to [1], our feature set consists of some standard features such as RGB color channels, LAB color channels and output from Law's Masks and the pixel height in the image.

In addition we have included an optical flow based motion feature, similar to temporal textures [12]. An optical flow calculation is performed between consecutive frames and the resulting flow is analyzed. The variance over the last few frames is calculated for each cell in the image. Since ripples on water appear to move randomly to the camera (i.e. can't be tracked accurately), this will give a good indication of water in the image. Non-water scenery moves much more consistently in the image and therefore has low variance compared to the water. An example of the output of this feature can be seen in Figure 2.



Figure 2: Optical flow-based variance feature. The sky has fairly high values due to the difficulty of tracking it. The water exhibits higher values than the land and trees in front.

Our final feature set consists of 104 features: 3 RGB, 3 LAB, 8*4*3 Law's, 1 height, 1 motion.

B. Classification

The classification is done using a linear Support Vector Machine (SVM). The SVM is currently trained using manually labeled data; we intend to use self-supervised classification in the future by providing training examples using training examples generated by the onboard sonar. Finally, the resulting output from the SVM is smoothed.

IV. CAMERA CALIBRATION

The camera calibration in this case is the task of finding the homography between the camera image plane to the water surface plane. We assume that our camera has already been calibrated intrinsically or that we have the intrinsic camera calibration matrix, K .

Points on the image plane, p_i , and the euclidean metric water surface plane, P_m , are related by the following equation in homogeneous coordinates:

$$s\tilde{p}_i = K \begin{bmatrix} R & C \end{bmatrix} \tilde{P}_m$$

$$s\tilde{p}_i = K H_{im} \tilde{P}_m$$

where K is the intrinsic camera matrix and H_{ab} is homography between the two planes.

The direct calculation of the homography can be done if accurate measurements for the camera mounting pose is known relative to the ASV and how the ASV sits on the water surface.

When these measurements are not available the transform can be estimated using point correspondences, i.e. if we know the real-world locations of a few points in the the image. The minimum number of correspondences to estimate this is 4. A least squares formulation is used when we have an overdetermined system.

When a large set of correspondences are available, algorithms such as RANSAC, can be used to find a good quality homography while rejecting outliers. OpenCV provides such an implementation which we have used in this work.

We obtain the point correspondences with the help of existing on-board range sensors, in our case, forward mounted sonar. It is worth noting that since all our correspondences lie in the same plane, we can not solve the full camera calibration problem where the intrinsic parameters are also solved, but if



Figure 3: The USC ASV along side USC’s EcoMapper (YSI)




Figure 4: Water segmentation results

we already have the intrinsic parameters, we can solve for the extrinsics with our coplanar set of correspondences.

When the sonar finds a suitable obstacle, we take the camera image for the corresponding time-stamp and do ray-casting from the center bottom of the image outward in the direction where the sonar indicated the obstacle to be. Since we do have the camera intrinsic calibration, we can determine the direction outward that we take in the image. The perspective projection of the camera preserves straight lines, so this ray-casting should intersect with the obstacle in the image. Since we have segmented the water in the image we cast the ray until the first non-water pixel is hit, and its pixel coordinates will be added to our list of point correspondences along with the coordinates measured by the sonar, p_i and P_m .

To determine when the sonar has detected a suitable obstacle we use similar processing as in [5], but some modification is needed for slanted obstacles that have presence both above and under water to return the range to the water-air boundary instead of some closer point under water. Then we apply some heuristics to decide whether to use it.

When the point correspondences have enough coverage of the image plane, they are used for finding the homography as described above. The homography can be improved further when more correspondences have been found.

V. EXPERIMENTAL RESULTS

To verify our approach we conducted field experiments using the USC ASV in San Dimas Lake, CA.

The USC ASV (Figure 3) is a 2 m long, 0.7 m wide boat actuated by two electric thrusters and a rudder. The hull is an OceanScience Q-Boat. The ASV is equipped with a computer, WiFi, RC fail-safe. Its basic sensing consists of a Microstrain IMU & compass and a u-block GPS. Additionally, for these experiments, it was equipped with a forward facing Imagenex profiling sonar, SICK-LMS200 laser range finder and a GoPro Hero camera. The ASV software is based on ROS.

A. Water Segmentation

Since the water segmentation is a crucial step for our approach to work, we will present results on that step separately. For the water segmentation we have manually extracted frames from a few of positions in the video. The frames have been labeled by hand into three classes for training: water, sky,

other. A result from the segmentation can be seen in figure 4.

B. Camera Calibration

The resulting calibration, calculated from 149 automatically chosen point correspondences, was verified by comparing the estimated difference in pose between the camera and the sonar. Figure 5 shows the calibrated transform in action where it was used to project the data from the on-board sonar and laser sensors on to the camera image. Note that in the 3rd image, the sonar projection is somewhat closer (below) than the laser ones, this is due to the sonar getting a hit off the slanted lake bottom, while the laser is hitting further away on land. Figure 6 shows a log-metric distance map for camera image, assuming no obstacles.

VI. CONCLUSIONS

We have presented an approach to obtain the extrinsic calibration between the camera and the plane of the water surface using a sonar.

The camera calibration approach used the segmentation output along with sonar data to determine the transform between the camera images and the real-world coordinates relative to the ASV. Comparison with a manually measured calibration shows that our approach is comparably accurate.

VII. FUTURE WORK

In future work we intend to further combine the output of the sonar with the camera to provide robust obstacle avoidance and tracking. This, along with our previous work on generating obstacle maps from overhead imagery, will serve as the base for our navigation system. We also intend to investigate the use of infrared cameras to provide vision at night and to do experiments in various weather conditions.

VIII. ACKNOWLEDGMENTS

We gratefully acknowledge help with experiments and hardware provided by Jon Binney, Megha Gupta, Geofferey Hollinger, Carl Oberg, Ryan Smith and Stephanie Kemna. 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.




Figure 5: Camera images with sonar and laser data overlaid using the estimated transform. Red dots are laser returns and the green dot is the sonar return.




Figure 6: The resulting distance mapping overlaid on the view from the camera.

REFERENCES

- [1] S. Achar, B. Sankaran, S. Nuske, S. Scherer, and S. Singh. Self-supervised segmentation of river scenes. In *2011 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6227–6232. IEEE, May 2011.
- [2] Tirthankar Bandyopadhyay, Lynn Sarcione, and Franz S. Hover. A simple reactive obstacle avoidance algorithm and its application in singapore harbor. In *Field and Service Robotics*, pages 455 – 465, 2009.
- [3] Fenglei Du and Michael Brady. Self-calibration of the intrinsic parameters of cameras for active vision systems. In *Computer Vision and Pattern Recognition, 1993. Proceedings CVPR'93., 1993 IEEE Computer Society Conference on*, pages 477–482, 1993.
- [4] Matthew Dunbabin, James Udy, Alistair Grinham, and Michael Bruenig. Continuous monitoring of reservoir water quality: The wivenhoe project. *Journal of the Australian Water Association*, 36:74–77, 2009.
- [5] Hordur K. Heidarsson and Gaurav S. Sukhatme. Obstacle detection and avoidance for an autonomous surface vehicle using a profiling sonar. In *Robotics and Automation, 2011 IEEE International Conference on*, Shanghai, China, May 2011.
- [6] Hordur K. Heidarsson and Gaurav S. Sukhatme. Obstacle detection from overhead imagery using self-supervised learning for autonomous surface vehicles. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Francisco, California, September 2011.
- [7] Terry Huntsberger, Hrand Aghazarian, Andrew Howard, and David C. Trotz. Stereo vision-based navigation for autonomous surface vessels. *Journal of Field Robotics*, 28(1):3–18, January 2011.
- [8] Jonathan Kelly and Gaurav S. Sukhatme. Visual-inertial sensor fusion: Localization, mapping and sensor-to-sensor self-calibration. *The International Journal of Robotics Research*, 30(1):56–79, 2011.
- [9] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected CRFs with gaussian edge potentials. *arXiv preprint arXiv:1210.5644*, 2012.
- [10] Jacoby Larson, Michael Bruch, Ryan Halterman, John Rogers, and Robert Webster. Advances in autonomous obstacle avoidance for unmanned surface vehicles. Washington, D.C., 2007.
- [11] Jacques C. Leedeckerken, Maurice F. Fallon, and John J. Leonard. Mapping complex marine environments with autonomous surface craft. Delhi, India, December 2010.
- [12] Randal C. Nelson and Ramprasad Polana. Qualitative recognition of motion using temporal texture. *CVGIP: Image Understanding*, 56(1):78–89, July 1992.
- [13] A. L. Rankin, L. H. Matthies, and P. Bellutta. Daytime water detection based on sky reflections. In *2011 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5329–5336. IEEE, May 2011.
- [14] Ryan Smith, Jnaneshwar Das, Hordur K. Heidarsson, Arvind de Menezes Pereira, Filippo Arrichiello, Ivona Cetinic, Lindsay Darjany, M. Garneau, M. Howard, Carl Oberg, M. Ragan, E. Seubert, E. Smith, B. Stauffer, A. Schnetzer, G. Toro-Farmer, David Caron, Burton Jones, and Gaurav S. Sukhatme. USC CINAPS builds bridges: Observing and monitoring the southern california bight. *IEEE Robotics and Automation Magazine*, 17(1):20–30, March 2010.
- [15] Pratap Tokekar, Elliot Branson, Joshua Vander Hook, and Volkan Isler. Coverage and active localization for monitoring invasive fish with an autonomous boat. *IEEE Robotics and Automation Magazine*, 2012.
- [16] Michael T. Wolf, Christopher Assad, Yoshiaki Kuwata, Andrew Howard, Hrand Aghazarian, David Zhu, Thomas Lu, Ashitey Trebi-Ollennu, and Terry Huntsberger. 360-degree visual detection and target tracking on an autonomous surface vehicle. *Journal of Field Robotics*, 27(6):819–833, November 2010.
- [17] Zhengyou Zhang. A flexible new technique for camera calibration. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(11):1330–1334, 2000.